informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 72, No. 4, July–August 2024, pp. 1739–1748
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Technical Note—Asymptotically Optimal Control of Omnichannel Service Systems with Pick-up Guarantees

Xuefeng Gao,[a] Junfei Huang,[b,*] Jiheng Zhang[c]

[a] Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR, China; [b] Department of Decision Sciences and Managerial Economics, CUHK Business School, The Chinese University of Hong Kong, Hong Kong SAR, China; [c] Department of Industrial Engineering and Decision Analytics, The Hong Kong University of Science and Technology, Hong Kong SAR, China

*Corresponding author

Contact: xfgao@se.cuhk.edu.hk, https://orcid.org/0000-0003-2424-8257 (XG); junfeih@cuhk.edu.hk, https://orcid.org/0000-0002-3764-354X (JH); jiheng@ust.hk, https://orcid.org/0000-0003-3025-1495 (JZ)

Received: January 14, 2022
Revised: August 18, 2022
Accepted: November 23, 2022
Published Online in Articles in Advance: April 28, 2023

Area of Review: Stochastic Models

https://doi.org/10.1287/opre.2022.2416

Abstract. Motivated by the recent popularity of omnichannel service systems, we analyze the joint admission and scheduling control of a queueing system with two classes of customers: online and walk-in. Unlike walk-in customers, online customers are given a target time for pick up upon placing an order. Thus, in addition to minimizing the waiting costs of walk-in customers and the rejection cost of both classes, we need to minimize the earliness and tardiness costs of online customers. Such a distinctive objective makes the control problem difficult to analyze. We develop a novel analysis by adopting the idea of proving $H = \lambda G$ to establish an asymptotic relationship between the waiting time cost and the queue length cost under general control policies in the heavy-traffic regime. We also design a policy and show that it is asymptotically optimal.

Funding: X. Gao's research is supported in part by the Hong Kong Research Grants Council [Grants 14201520, 14201421, and 14212522]. J. Huang's research is supported in part by the Hong Kong Research Grants Council [Grants 14500819, 14505820, and 14501621] and the National Natural Science Foundation of China [Grant 72222023]. J. Zhang's research is supported in part by the Hong Kong Research Grants Council [Grants 16208120 and 16214121].

Keywords: omnichannel service • heavy traffic • earliness and tardiness costs • asymptotically optimal control

1. Introduction

The rise of the mobile Internet and the recent COVID-19 pandemic have profoundly changed the way we do business. One such change is that restaurants are now catering to both walk-in and online customers (Chen et al. 2022). Such service systems that use both physical and online channels are referred to as omnichannel service systems (Baron et al. 2023). From an operational perspective, although minimizing the waiting time is still important for walk-in customers, online customers who will be given a promised pick-up time upon ordering have a completely different objective. In the restaurant industry, the quality of the food/beverage can degrade if it is ready before the arrival of online customers. Hence, online customers incur an earliness cost if the preparation of the food/beverage is completed before the promised pick-up time. On the other hand, such customers incur a tardiness or waiting cost if they arrive at the promised times but the food/beverage is not ready (Farahani et al. 2022).

Although many control policies have been designed and analyzed for queueing systems in the literature, few studies have focused on a queueing model with both earliness and tardiness costs (except, for example, Farahani et al. (2022) and Markowitz and Wein (2001)). In this note, we consider the dynamic control of a single-server queue with two classes of customers: online and walk-in. Customers may be rejected by the system to avoid long waiting times (Baron et al. 2023). A control policy decides whether to reject or accept an arriving order from each class, when the server should idle, and which class of orders to serve when not idling. The goal is to minimize the long-run expected average cost, which consists of the earliness and tardiness costs of online orders, the waiting costs of walk-in orders, and the rejection cost of both classes.

As the objective involves the earliness/tardiness/waiting costs of customers and because the waiting time dynamics in multiclass queueing systems are complicated, the joint admission and scheduling control problem appears to be analytically intractable. We resort to heavy-traffic analysis, which, although assuming the system operates near full capacity, has been demonstrated in the queueing literature (Wein 1992) to provide accurate approximations to the underlying stochastic processes. Even in the heavy-traffic framework, the feature that online customers may incur either earliness or tardiness costs, which are functions of the waiting times (from order placement to order completion), distinguishes our problem from existing ones in the literature. To the best of our knowledge, most formulation and analysis on the optimal control of multiclass queueing systems are based on queue length (or workload) except, for example, van Mieghem (1995), which relies on the work-conserving property for nonidling policies. To tackle this challenge, we developed a novel analysis by adopting the idea of proving an extended version of Little's law, known as $H = \lambda G$ (Heyman and Stidham 1980), in the heavy traffic setting. This allows us to convert the optimal control problem involving waiting time costs into a problem involving queue length costs under general control policies, without relying on the work-conserving property. Based on such a conversion, we propose a dynamic control policy and establish its asymptotic optimality when the traffic intensity approaches one.

Indeed, our proposed policy will keep the server idle when the queue of online orders drops below a threshold and there are no walk-in customers. This is primarily due to the earliness cost of the online customers. It is different from the strategic delay proposed by Afeche (2013) in designing incentive-compatible and revenue-maximizing pricing and scheduling policies in queueing systems. The rationale of the strategic delay is to manipulate customers' strategic service class choices. It is also different from the strategic idleness proposed by Baron et al. (2014, 2017) for service networks, where an upstream station is intentionally idled when its downstream station has a long queue to reduce the probability of long waits at the downstream station. Our proposed policy will also reject orders from one class when the workload is too heavy (beyond a threshold). Relying on the heavy-traffic asymptotic analysis, we are able to characterize these thresholds for any given set of model parameters. Finally, for the control part, regarding which class to serve, our policy resembles the $c\mu$ type of policies developed in the literature (Cox and Smith 1961, Atar et al. 2010).

In terms of modeling, this note is most closely related to Baron et al. (2023) and Farahani et al. (2022). Baron et al. (2023) study the two-class omnichannel service systems under the first-come-first-served (FCFS) policy. They use a game-theoretical queueing framework to investigate customers' channel choices. Farahani et al. (2022) study a single-class queue, focusing on online orders with delivery guarantees. They propose a static threshold policy that depends on the remaining time to the due date of the first customer in the queue. They illustrate the efficacy of this heuristic policy by developing lower bounds and numerical experiments. Our model is also related to Armony and Maglaras (2004), in which some customers in a call center choose a call-back option with a guarantee on the maximum delay. In their model, serving a customer earlier than the promised guarantee does not incur a cost and the offered delay guarantee is modeled as a hard constraint.

1.1. Contributions

Our contributions are summarized as follows:

- Formulating the dynamic control problem of an omnichannel service system with two classes of orders, with a fundamentally different cost structure than classical ones.
- Establishing asymptotic relationships between the waiting time cost and queue length cost under general control policies, by adopting the idea of proving $H = \lambda G$.
- Proposing and establishing the asymptotic optimality of a policy with both admission control and capacity allocation.

2. Model

Consider a single-server queue with two classes of customers. Class 1 customers correspond to online orders, and class 2 customers correspond to walk-in orders. We will use customers and orders interchangeably. For $k \in \{1, 2\}$, denote by $E_k(t)$ the number of class k orders having arrived by time t, and assume the process $E_k(\cdot) = \{E_k(t); t \geq 0\}$ is a homogeneous Poisson process with rate λ_k. Both classes of orders may be rejected by the system manager to avoid long waiting times. Denote by $O(\cdot) = (O_k(\cdot); k \in \{1, 2\})$ the two-dimensional rejecting process, with $O_k(t)$ being the number of class k orders rejected up to time t. The actual number of class k orders accepted by time t is then $A_k(t) = E_k(t) - O_k(t)$.

For $k \in \{1, 2\}$, the service times of class k orders are assumed to form an independent and identically distributed sequence of random variables, following an exponential distribution with mean $1/\mu_k$. Equivalently, denote by $S_k(t)$ the number of class k service completions up to time t if the server was continuously serving class k orders in $[0, t]$, and then $S_k(\cdot) = \{S_k(t) : t \geq 0\}$ is a Poisson process with rate μ_k. We assume that $E_k(\cdot)$ and $S_k(\cdot), k \in \{1, 2\}$, are independent. Let $T_k(t)$ denote the cumulative amount of time the server devotes to serving class k orders in $[0, t]$, and $T(\cdot) = (T_k(\cdot); k \in \{1, 2\})$. Then, $S_k(T_k(t))$ is the number of class k orders served by time t.

The system manager controls the system by deciding whether and when an arriving order is rejected and making scheduling decisions about when and whether the server should idle or not and which class of orders to serve when not idling. Mathematically, a dynamic control policy is denoted by $\psi = (O, T)$.

We now formulate the optimal control problem. Following Farahani et al. (2022), each online order (if accepted) is given a static pick-up guarantee δ, say 30 minutes. Denote by $t_{1,i}$ the arrival time, and $D_{1,i}(\psi)$ the departure time, of the ith accepted online order under policy ψ. The due date of the ith accepted online order is then $d_{1,i} = t_{1,i} + \delta$. Let c_e and c_d be the earliness and tardiness costs incurred by an online order per unit of time, respectively. Then, the earliness/tardiness cost incurred by the ith accepted online order is $g(D_{1,i}(\psi) - d_{1,i})$, where

$$g(x) = c_e \cdot x^- + c_d \cdot x^+.$$

Here, $x^+ = \max(0, x)$ and $x^- = \max(0, -x)$. Note that $D_{1,i}(\psi) = t_{1,i} + w_{1,i}(\psi)$, where $w_{1,i}(\psi)$ is the sojourn time of the ith accepted online order under policy ψ, and it includes both the time spent waiting in the queue and the time in service. Then, we have $D_{1,i}(\psi) - d_{1,i} = w_{1,i}(\psi) - \delta$. Because the number of accepted online orders during $[0, t]$ is $A_1(t)$, mathematically, the earliness and tardiness costs incurred by the accepted online orders is given by

$$C_1^{\psi}(t) = \sum_{i=1}^{A_1(t)} g(w_{1,i}(\psi) - \delta).$$

On the other hand, (accepted) walk-in orders may need to wait, and the waiting cost of each walk-in order is denoted by c_w per time unit. Denote by $w_{2,i}(\psi)$ the sojourn time of the ith (accepted) walk-in order under policy ψ. Then, the total waiting costs incurred by the accepted walk-in orders is given by

$$C_2^{\psi}(t) = \sum_{i=1}^{A_2(t)} c_w \cdot w_{2,i}(\psi).$$

Finally, we assume each rejected class k order incurs a cost $\theta_k > 0$; hence, the total cost of rejecting both classes of orders by time t is $\sum_{k=1}^{2} \theta_k O_k(t)$.

Then, for a given policy ψ, the total expected cost incurred by time t is

$$\begin{aligned} V(t, \nu, \psi) &:= \mathbb{E}_{\nu}\left(C_1^{\psi}(t) + C_2^{\psi}(t) + \sum_{k=1}^{2} \theta_k O_k(t)\right) \\ &= \mathbb{E}_{\nu}\left(\sum_{i=1}^{A_1(t)} g(w_{1,i}(\psi) - \delta) + \sum_{i=1}^{A_2(t)} c_w \cdot w_{2,i}(\psi) \right. \\ &\quad \left. + \sum_{k=1}^{2} \theta_k O_k(t)\right), \end{aligned} \tag{1}$$

where the expectation is taken with respect to ν, the distribution of the initial system state that will be discussed in the following. The goal of the system manager is to find the optimal policy to minimize the long run average cost. Formally, we aim to solve the following problem:

$$\inf_{\psi \in \Pi} \left\{ \mathbf{C}^{\psi} := \limsup_{t \to \infty} \frac{1}{t} V(t, \nu, \psi) \right\}, \tag{2}$$

where Π denotes the set of all feasible Markov control policies to be defined later. The cost $V(t, \nu, \psi)$ involves the nonlinear functions of the waiting times. This is different from most of the existing literature on the dynamic control of multiclass queueing systems, in which the control formulation is based on the queue length (or workload). Generally, in multiclass queueing systems, the waiting time dynamics, which depend on future actions, are much more complicated than those of the queue lengths.

We now describe the system state process and define feasible Markov control policies. Under a control policy ψ, denote by $Q_k^{\psi}(t)$ the queue length of class k orders (including the one in service, if any) at time t, for $k \in \{1, 2\}$. For brevity, when the policy is clear from the context, we will drop the notation ψ. Denote by $C(t)$ the order class that is currently in service. That is, $C(t) = k$ means the system is serving a class k order at time t for $k = 1, 2$, and $C(t) = 0$ means the system is idle. The system state process is $\mathfrak{X}(\cdot) = \{\mathfrak{X}(t); t \geq 0\}$ with $\mathfrak{X}(t) = (Q_1(t), Q_2(t), C(t))$.

Denote the process $I(\cdot) = \{I(t); t \geq 0\}$, with $I(t) = t - \sum_{k=1}^{2} T_k(t)$ for $t \geq 0$, which is the cumulative idle time of the system by time t. We also introduce a process $\xi(\cdot) = (\xi_k(\cdot), k \in \{1, 2\})$ with $\xi_k(\cdot) = \{\xi_k(t); t \geq 0\}$, in which $\xi_k(s-)$ indicates the rejecting decision for a virtual new class k order arriving at time s: If $\xi_k(s-) = 1$ and a class k order arrives at s, then that order is rejected; if $\xi_k(s-) = 0$, then that order is accepted. Here, we use $\xi_k(s-)$ to indicate that the decision is made with the information before time s. Then, we have $O_k(t) = \int_0^t \xi_k(s-) dE_k(s)$. Hence, the control of the rejecting process $O(\cdot)$ is via $\xi(\cdot)$.

Definition 1. A control policy $\psi = (O, T)$ is said to be a *feasible Markov control policy* if it is nonanticipating with respect to $\mathfrak{X}(\cdot)$ and

1. $I(\cdot), T(\cdot), O(\cdot)$ are nondecreasing with $I(0) = T(0) = O(0) = 0$,
2. $I(\cdot), T(\cdot)$ are continuous, and
3. $\xi(t)$ can be represented as a measurable function of $\mathfrak{X}(t)$.

For notational brevity, we use $\xi(\cdot)$ to denote the corresponding measurable function, that is, $\xi(t) = \xi(\mathfrak{X}(t))$. One can readily see that under a feasible Markov control policy, the system state process $\mathfrak{X}(\cdot)$ is a continuous-time Markov chain with a countable state space $S = \mathbb{Z}_+^2 \times \{0, 1, 2\}$.

3. Heavy-Traffic Framework

We consider a sequence of queueing systems as before, indexed by $n \in \mathbb{N}$, in the *conventional heavy traffic* regime. The relevant parameters and processes in the nth system will be appended with a superscript n. We assume the arrival rate of class k orders in the nth system is

$$\lambda_k^n = n\lambda_k + \sqrt{n}\beta_k,$$

for some $\beta_k \in \mathbb{R}$, and the service rate is $\mu_k^n = n\mu_k$, where $\lambda_1/\mu_1 + \lambda_2/\mu_2 = 1$. The system state process in the nth system is denoted by $\mathfrak{X}^n(\cdot) = \{\mathfrak{X}^n(t) := (Q_1^n(t), Q_2^n(t), C^n(t)); t \geq 0\}$, whose initial state $\mathfrak{X}^n(0)$ follows distribution ν^n. The control in the nth system is denoted by $\psi^n = (O^n, T^n)$, in which T_k^n records the time allocated to serving class k orders, and O_k^n denotes the number of class k orders rejected. We will focus on feasible Markov control policies.

Assume that the rejection cost rate θ_k^n, the waiting cost of walk-in orders c_w^n, the earliness/tardiness cost of online orders $g^n(x)$, and the pick-up guarantee δ^n vary with n as follows:

$$\theta_k^n = \frac{\theta_k}{\sqrt{n}}, \quad c_w^n = \frac{c_w}{\sqrt{n}}, \quad g^n(x) = \frac{1}{n} g(\sqrt{n}x) = \frac{g(x)}{\sqrt{n}}, \quad \delta^n = \frac{\delta}{\sqrt{n}}.$$

Denote by $A_k^n(t)$ the number of class k orders accepted by time t. Then, from (1), the expected cumulative cost associated with policy ψ^n by time t is

$$\begin{aligned}\widetilde{V}^n(t, \nu^n, \psi^n) &= \mathbb{E}_{\nu^n}\left(\sum_{i=1}^{A_1^n(t)} g^n(w_{1,i}^n(\psi^n) - \delta^n) + \sum_{i=1}^{A_2^n(t)} c_w^n \cdot w_{2,i}^n(\psi^n) \right.\\ &\quad \left. + \sum_{k=1}^{2} \theta_k^n O_k^n(t)\right)\\ &= \mathbb{E}_{\nu^n}\left(\sum_{i=1}^{A_1^n(t)} \frac{1}{n} g\left(\widetilde{w}_{1,i}^n(\psi^n) - \delta\right)\right.\\ &\quad \left. + \sum_{i=1}^{A_2^n(t)} \frac{1}{n} c_w \cdot \widetilde{w}_{2,i}^n(\psi^n) + \sum_{k=1}^{2} \theta_k \widetilde{O}_k^n(t)\right), \end{aligned} \tag{3}$$

where $\widetilde{w}_{k,i}^n(\psi^n) = \sqrt{n} \cdot w_{k,i}^n(\psi^n)$ and $\widetilde{O}_k^n(t) = \frac{O_k^n(t)}{\sqrt{n}}$ for $k = 1, 2$. The long-run average cost for a given control policy ψ^n is then given by

$$\widetilde{V}^n(\nu^n, \psi^n) := \limsup_{t\to\infty} \frac{1}{t} \widetilde{V}^n(t, \nu^n, \psi^n).$$

We call a sequence of polices $\{\psi^n\}$ *admissible* if, for each n, $\psi^n \in \Pi^n$, where Π^n denotes the set of feasible Markov control policies for the nth system.

Definition 2. (Asymptotic Optimality). A sequence of polices $\{\psi_*^n\}$ is *asymptotically optimal* if it is admissible, and for any other admissible sequence of policies $\{\psi^n\}$,

$$\liminf_{n\to\infty} \widetilde{V}^n(\nu^n, \psi^n) \geq \limsup_{n\to\infty} \widetilde{V}^n(\nu^n, \psi_*^n),$$

for any sequence of initial distributions $\{\nu^n\}$.

4. Analysis and Main Results

In this section, we propose a sequence of control policies and establish its asymptotic optimality.

4.1. Proposed Policies

We first introduce an ordinary differential equation (ODE) that identifies several parameters of the proposed policies. To this end, we first define for $x \in \mathbb{R}_+$,

$$f_1(x) = c_e(x - \lambda_1\delta)^- + c_d(x - \lambda_1\delta)^+ = g(x - \lambda_1\delta), \quad \text{and}$$
$$f_2(x) = c_w x. \tag{4}$$

For each $y \geq -\frac{\lambda_1\delta}{\mu_1}$, consider the following minimization problem:

$$h(y) := \min_{(x_1,x_2)\in\mathbb{R}_+^2} \quad \sum_{k=1}^{2} f_k(x_k) \quad \text{s.t.} \quad \frac{x_1 - \lambda_1\delta}{\mu_1} + \frac{x_2}{\mu_2} = y. \tag{5}$$

One can readily verify that $h(y) = \min\{c_d\mu_1, c_w\mu_2\} \cdot y 1_{\{y>0\}} - c_e\mu_1 y 1_{\left\{-\frac{\lambda_1\delta}{\mu_1} \leq y \leq 0\right\}}$.

We also define the effective rejecting cost κ as follows (if $\theta_1\mu_1 = \theta_2\mu_2$, let $i^* = 1$.):

$$i^* = \operatorname{argmin}\{\theta_k\mu_k : k = 1, 2\} \quad \text{and} \quad \kappa = \theta_{i^*} \cdot \mu_{i^*}. \tag{6}$$

We can show the following result, the proof of which is deferred to Section 5.

Lemma 1. *Let $\sigma^2 = \sum_{k=1}^{2} \frac{2\lambda_k}{\mu_k^2}$. Then, there exist unique constants $l_* \in \left[-\frac{\lambda_1\delta}{\mu_1}, 0\right)$, $u_* > 0$, $\gamma^* > 0$ and a unique (up to an additive constant) real-valued function $\Phi \in C^2\left(\left[-\frac{\lambda_1\delta}{\mu_1}, \infty\right)\right)$ satisfying the differential equation:*

$$\frac{\sigma^2}{2}\Phi''(x) + \sum_{k=1}^{2} \frac{\beta_k}{\mu_k}\Phi'(x) + h(x) = \gamma^*, \quad x \in (l_*, u_*), \tag{7}$$

and $\Phi'(x) \in [0, \kappa]$ for $x \geq -\frac{\lambda_1\delta}{\mu_1}$, $\Phi'(x) = 0$ for $x \in \left[-\frac{\lambda_1\delta}{\mu_1}, l_*\right]$, $\Phi'(x) = \kappa$ for $x \geq u_*$, and $\Phi''(x) = 0$ for $x \notin (l_*, u_*)$. As a consequence, there exists a positive constant C such that for any $x \geq -\frac{\lambda_1\delta}{\mu_1}$,

$$\frac{\sigma^2}{2}\Phi''(x) + \sum_{k=1}^{2} \frac{\beta_k}{\mu_k}\Phi'(x) + h(x) \geq \gamma^*, \tag{8}$$

$$0 \leq \Phi'(x) \leq \kappa, \quad |\Phi''(x)| \leq C. \tag{9}$$

Additionally, $\Phi'''(x)$ exists almost everywhere and $|\Phi'''(x)| \leq C$ whenever it exists.

Equation (7) is a first-order linear ODE for the unknown function $\Phi'(\cdot)$. Note that γ^* and the boundary points l_* and u_* of ODE (7) are unknown and need to be determined; hence, this type of ODE is called a free boundary ODE in the literature (Dai and Yao 2013). Based on the proof of Lemma 1, this free-boundary ODE can be solved numerically by searching the two parameters (γ, w_0) in the associated initial value problem in (35) so that its solution has a maximal value κ on $[0, \infty)$ and a minimal value of zero on $\left[-\frac{\lambda_1\delta}{\mu_1}, 0\right]$.

Remark 1. We summarize how the ODE in Lemma 1 will be used. We use Inequalities (8) and (9) to obtain a lower bound result (for general admissible policies) in Proposition 1. Specifically, Equation (8) is used to bound Ψ_1 in (30), and Equation (9) is used to bound Ψ_2 in (31) and Ψ_3 in (32). On the other hand, we apply

Equation (7) together with the boundary conditions for Φ to analyze the proposed policies in the proof of Proposition 2. In particular, Equation (7) is used to bound Ψ_1, and the boundary conditions are used to bound Ψ_2 and Ψ_3.

To describe the proposed policy, define the diffusion-scaled queue-length processes $\widetilde{Q}^n(\cdot) = \{(\widetilde{Q}_1^n(t), \widetilde{Q}_2^n(t)); t \geq 0\}$ and the nominal workload process $\widetilde{W}^n(\cdot) = \{\widetilde{W}^n(t); t \geq 0\}$ with

$$\widetilde{Q}_k^n(t) = \frac{Q_k^n(t)}{\sqrt{n}}, \quad k = 1, 2, \quad \text{and}$$
$$\widetilde{W}^n(t) := \frac{\widetilde{Q}_1^n(t) - \lambda_1\delta}{\mu_1} + \frac{\widetilde{Q}_2^n(t)}{\mu_2}, \quad \text{for t} \geq 0. \tag{10}$$

We now propose a sequence of feasible Markov control policies $\{\psi_*^n\}$ with ψ_*^n applied to the nth system. The details of the policy ψ_*^n are as follows, with $l_* < 0 < u_*$ given in Lemma 1:

- *Order rejection policy*: Reject a new arriving class i^* order at time t if $\widetilde{W}^n(t) \geq u_*$, and do not reject any new order otherwise.
- *Server scheduling policy*: The service discipline within each class is FCFS, and

 1. If $\frac{\widetilde{Q}_1^n(t) - \lambda_1\delta}{\mu_1} < l_*$: serve class 2 orders if $\widetilde{Q}_2^n(t) > 0$, and keep the server idle if $\widetilde{Q}_2^n(t) = 0$.
 2. If $\frac{\widetilde{Q}_1^n(t) - \lambda_1\delta}{\mu_1} \geq l_*$:
    (a) If $\widetilde{Q}_2^n(t) = 0$, then serve class 1 orders.
    (b) If $\frac{\widetilde{Q}_1^n(t) - \lambda_1\delta}{\mu_1} \leq 0$ and $\widetilde{Q}_2^n(t) > 0$, then serve class 2 orders.
    (c) Otherwise, serve a class 1 order if $c_d\mu_1 \geq c_w\mu_2$, and serve a class 2 order if $c_d\mu_1 < c_w\mu_2$.

Remark 2. Class i^* in (6) depends on the rejection cost rate θ_k for $k = 1, 2$. In reality, it may be more expensive to reject walk-in customers. Hence, it is likely that i^* would be class 1; that is, it is preferable to reject online customers. Therefore, the proposed order rejection policy is consistent with the practice. See Baron et al. (2023) for an example where a tea shop only rejects online orders when the local store is overwhelmed by waiting orders.

We will append a subscript $*$ to processes under the proposed policy ψ_*^n.

Theorem 1. *The sequence of policies $\{\psi_*^n\}$ is asymptotically optimal.*

The key step in the proof of Theorem 1 is to convert the optimal control problem involving waiting times to the one involving queue lengths (Theorem 2). We then adopt the framework in Gao and Huang (2023) to analyze the problem with queue lengths (see Section 4.3). Our proposed policy is only asymptotically optimal. If the heavy-traffic assumptions are violated, then it may not work well.

4.2. From Waiting Times to Queue Lengths

Given a policy ψ^n and an initial distribution ν_n, we introduce

$$\widetilde{G}^n(t, \nu^n, \psi^n) := \mathbb{E}_{\nu^n}\left[\sum_{k=1}^{2}\int_0^t f_k(\widetilde{Q}_k^n(s))ds + \sum_{k=1}^{2}\theta_k\widetilde{O}_k^n(t)\right], \tag{11}$$

where f_1 and f_2 are given in (4). The following result shows that the cost $\widetilde{V}^n(t, \nu^n, \psi^n)$ in (3) is close to $\widetilde{G}^n(t, \nu^n, \psi^n)$; hence, the original control problem can be converted to one involving queue lengths. The key idea in its proof is based on adapting $H = \lambda G$ (Heyman and Stidham 1980) in the heavy-traffic setting. There are two major technical challenges: (1) constructing appropriate functions to connect the earliness/tardiness costs with certain queue statistics for accepted class 1 orders (see, e.g., (16)–(17)), and (2) analyzing and bounding the expected number of accepted and rejected class 1 orders.

Theorem 2. *There exist constants $C_1, C_2 > 0$, such that*
(a) *For any admissible sequence of policies $\{\psi^n\}$, we have*

$$\liminf_{n\to\infty}\liminf_{t\to\infty}\frac{1}{t}\left[\left(1+\frac{C_1}{\sqrt{n}}\right)\widetilde{V}^n(t, \nu^n, \psi^n) - \widetilde{G}^n(t, \nu^n, \psi^n)\right] \geq 0. \tag{12}$$

(b) *For any admissible sequence of policies $\{\phi^n\}$ such that there exists a unique stationary distribution of $\mathfrak{X}^n(\cdot)$ under the policy ϕ^n for large enough n, we have*

$$\limsup_{n\to\infty}\limsup_{t\to\infty}\frac{1}{t}\left[\widetilde{V}^n(t, \nu^n, \phi^n) - \left(1+\frac{C_2}{\sqrt{n}}\right)\widetilde{G}^n(t, \nu^n, \phi^n)\right] \leq 0. \tag{13}$$

Proof of Theorem 2. We first note that $x^+ = x + x^-$, $A_1^n(t) = E_1^n(t) - O_1^n(t)$ and $\mathbb{E}\left[\frac{E_1^n(t)}{n}\right] = \lambda_1 t + \frac{\beta_1}{\sqrt{n}}t$. Using (3) and (11), we can reorganize the terms and get

$$\begin{aligned}\widetilde{V}^n(t, \nu^n, \psi^n) - \widetilde{G}^n(t, \nu^n, \psi^n) &= (c_e + c_d)\widetilde{\Delta}_0^n(t, \nu^n, \psi^n) \\ &\quad + c_d\delta^n\left(-\beta_1 t + \mathbb{E}_{\nu^n}\left[\widetilde{O}_1^n(t)\right]\right) \\ &\quad + c_d\widetilde{\Delta}_1^n(t, \nu^n, \psi^n) + c_w\widetilde{\Delta}_2^n(t, \nu^n, \psi^n),\end{aligned} \tag{14}$$

where $\delta^n = \delta/\sqrt{n}$ and

$$\widetilde{\Delta}_0^n(t, \nu^n, \psi^n) = \mathbb{E}_{\nu^n}\left[\sum_{i=1}^{A_1^n(t)}\frac{1}{n}\left(\widetilde{w}_{1,i}^n(\psi^n) - \delta\right)^- - \int_0^t(\widetilde{Q}_1^n(s) - \lambda_1\delta)^- ds\right],$$

$$\widetilde{\Delta}_k^n(t, \nu^n, \psi^n) = \mathbb{E}_{\nu^n}\left[\sum_{i=1}^{A_k^n(t)}\frac{1}{n}\widetilde{w}_{k,i}^n(\psi^n) - \int_0^t\widetilde{Q}_k^n(s)ds\right], \text{ for } k = 1, 2.$$

For any admissible policy ψ^n for the nth system, we define for each sample path ω and $s \geq 0$ with $\mathbf{f}_i^n(\omega, s) = 1$ if $t_{1,i}^n + \delta^n > D_{1,i}(\psi^n)$ and $s \in (D_{1,i}(\psi^n), t_{1,i}^n + \delta^n)$, and

$\mathbf{f}_i^n(\omega,s)=0$ otherwise. One can then readily verify that

$$\mathbf{g}_i^n(\omega):=\int_0^\infty \mathbf{f}_i^n(\omega,s)ds=[D_{1,i}(\psi^n)-(t_{1,i}^n+\delta^n)]^- \\ =[w_{1,i}^n(\psi^n)-\delta^n]^-. \tag{15}$$

One can also verify that

$$\mathbf{h}^n(\omega,s):=\sum_{i=1}^\infty \mathbf{f}_i^n(\omega,s)\geq [A_1^n(s)-A_1^n(s-\delta^n)-Q_1^n(s)]^+ \\ =[Q_1^n(s)-(A_1^n(s)-A_1^n(s-\delta^n))]^-, \tag{16}$$

and the inequality becomes an equality if the service discipline is FCFS. Note that (16) holds because $\mathbf{h}^n(\omega,s)$ is the number of class 1 orders who arrive during $(s-\delta^n,s]$ but are not in the system at time s. Equation (16) plays an important role in proving Theorem 2.

We first prove (12). Using (15), (16), and equation (3) in Heyman and Stidham (1980),

$$\sum_{i=1}^{A_1^n(t)}(w_{1,i}^n(\psi^n)-\delta^n)^-\geq\int_0^t (Q_1^n(s)-(A_1^n(s)-A_1^n(s-\delta^n)))^-ds. \tag{17}$$

As a result,

$$\begin{aligned}&\widetilde{\Delta}_0^n(t,\nu^n,\psi^n)\\ &\geq -\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(A_1^n(s)-A_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\ &\geq -\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\ &\quad -\mathbb{E}_{\nu^n}\left[\frac{1}{\sqrt{n}}\int_0^t|O_1^n(s)-O_1^n(s-\delta^n)|ds\right]\\ &\geq -\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\ &\quad -\delta^n\mathbb{E}_{\nu^n}\left[\widetilde{O}_1^n(t)\right],\end{aligned} \tag{18}$$

where $E_1^n(u):=0$ and $O_1^n(u):=0$ if $u<0$, and we used the definition $A_1^n(t)=E_1^n(t)-O_1^n(t)$ for $t\geq 0$, and the fact that

$$\begin{aligned}\int_0^t|(O_1^n(s)-O_1^n(s-\delta^n))|ds&=\int_0^t\int_{s-\delta^n}^s dO_1^n(u)ds\\ &\leq\int_0^t\int_u^{u+\delta^n}dsdO_1^n(u)\\ &=\delta^nO_1^n(t).\end{aligned}$$

Note that (18) holds regardless of the service discipline.

In addition, from lemma 1 in Stidham (1974), we have $\sum_{i=1}^{A_k^n(t)}w_{k,i}^n(\psi^n)\geq\int_0^tQ_k^n(s)ds$ for each class $k=1,2$, which immediately implies that

$$\widetilde{\Delta}_k^n(t,\nu^n,\psi^n)\geq 0,\quad \text{for } k=1,2. \tag{19}$$

Using Inequalities (18) and (19), we then infer from (14) that

$$\begin{aligned}&\widetilde{V}^n(t,\nu^n,\psi^n)-\widetilde{G}^n(t,\nu^n,\psi^n)\\ &\geq -(c_e+c_d)\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\ &\quad -c_d\delta^n\beta_1t-c_e\delta^n\mathbb{E}_{\nu^n}\left[\widetilde{O}_1^n(t)\right].\end{aligned}$$

By Definition (3), we have $\mathbb{E}_{\nu^n}[\delta^n\widetilde{O}_1^n(t)]\leq\frac{\delta}{\sqrt{n}}\frac{1}{\theta_1}\widetilde{V}^n(t,\nu^n,\psi^n)$. It follows that

$$\begin{aligned}&\left(1+\frac{\delta c_e}{\sqrt{n}\theta_1}\right)\widetilde{V}^n(t,\nu^n,\psi^n)-\widetilde{G}^n(t,\nu^n,\psi^n)\\ &\geq -(c_e+c_d)\frac{1}{\sqrt{n}}\mathbb{E}_{\nu^n}\left[\int_0^t|(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\sqrt{n}\delta|ds\right]\\ &\quad -c_d\delta^n\beta_1t.\end{aligned} \tag{20}$$

Because $E_1^n(\cdot)$ is a Poisson process with rate $\lambda_1^n=\lambda_1n+\sqrt{n}\beta_1$ and $\delta^n=\delta/\sqrt{n}$, we have

$$\lim_{n\to\infty}\lim_{t\to\infty}\frac{1}{\sqrt{n}}\frac{1}{t}\mathbb{E}_{\nu^n}\left(\int_0^t|(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\sqrt{n}\delta|ds\right)=0. \tag{21}$$

Thus, we immediately obtain (12) from (20).

We next prove (13). If $t\geq t_{1,i}^n+\delta^n$, then $\mathbf{f}_i^n(\omega,s)=0$ for $s\geq t$, and hence $\int_0^t\mathbf{f}_i^n(\omega,s)ds=\int_0^\infty\mathbf{f}_i^n(\omega,s)ds$. Thus, we have

$$\begin{aligned}\int_0^t\mathbf{h}^n(\omega,s)ds&\geq\sum_{t_{1,i}^n\leq t-\delta^n}\int_0^t\mathbf{f}_i^n(\omega,s)ds\\ &=\sum_{t_{1,i}^n\leq t-\delta^n}\int_0^\infty\mathbf{f}_i^n(\omega,s)ds=\sum_{t_{1,i}^n\leq t-\delta^n}\mathbf{g}_i^n(\omega).\end{aligned}$$

It then follows from (15) and (16) that

$$\begin{aligned}&\widetilde{\Delta}_0^n(t,\nu^n,\phi^n)\\ &=\mathbb{E}_{\nu^n}\left[\sum_{i=1}^{A_1^n(t)}\frac{1}{n}(\widetilde{w}_{1,i}^n(\phi^n)-\delta)^--\frac{1}{\sqrt{n}}\int_0^t(Q_1^n(s)-(A_1^n(s)\right.\\ &\quad\left.-A_1^n(s-\delta^n)))^-ds\right]\\ &\quad+\mathbb{E}_{\nu^n}\left[\frac{1}{\sqrt{n}}\int_0^t(Q_1^n(s)-(A_1^n(s)-A_1^n(s-\delta^n)))^-ds\right.\\ &\quad\left.-\int_0^t(\widetilde{Q}_1^n(s)-\lambda_1\delta)^-ds\right]\\ &\leq\mathbb{E}_{\nu^n}\left[\sum_{i=A_1^n(t-\delta^n)+1}^{A_1^n(t)}\frac{1}{n}\left(\widetilde{w}_{1,i}^n(\phi^n)-\delta\right)^-\right]\\ &\quad+\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\ &\quad+\delta^n\mathbb{E}_{\nu^n}\left[\widetilde{O}_1^n(t)\right]\\ &\leq\mathbb{E}_{\nu^n}\left[\frac{\delta(A_1^n(t)-A_1^n(t-\delta^n))}{n}\right]\\ &\quad+\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\ &\quad+\delta^n\mathbb{E}_{\nu^n}\left[\widetilde{O}_1^n(t)\right],\end{aligned}$$

where in the first inequality we apply a similar argument as in (18). We can then infer that

$$
\begin{aligned}
&\widetilde{V}^n(t,\nu^n,\phi^n)-\widetilde{G}^n(t,\nu^n,\phi^n)\\
&\le (c_e+c_d)\mathbb{E}_{\nu^n}\left[\frac{\delta(A_1^n(t)-A_1^n(t-\delta^n))}{n}\right]-c_d\delta^n\beta_1 t\\
&\quad+(c_e+c_d)\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\
&\quad+(c_e+2c_d)\delta^n\mathbb{E}_{\nu^n}\left[\widetilde{O}_1^n(t)\right]+c_d\widetilde{\Delta}_1^n(t,\nu^n,\psi^n)+c_w\widetilde{\Delta}_2^n(t,\nu^n,\psi^n).
\end{aligned}
$$

By (11), we have $\mathbb{E}_{\nu^n}[\delta^n\widetilde{O}_1^n(t)]\le\frac{\delta}{\sqrt{n}}\frac{1}{\theta_1}\widetilde{G}^n(t,\nu^n,\phi^n)$. Moreover, we can compute that $\mathbb{E}_{\nu^n}\left[A_1^n(t)-A_1^n(t-\delta^n)\right]\le\mathbb{E}_{\nu^n}\left[E_1^n(t)-E_1^n(t-\delta^n)\right]=(\lambda_1 n+\beta\sqrt{n})\cdot\frac{\delta}{\sqrt{n}}$. Hence,

$$
\begin{aligned}
&\widetilde{V}^n(t,\nu^n,\phi^n)-\left(1+\frac{\delta(c_e+2c_d)}{\sqrt{n}\theta_1}\right)\widetilde{G}^n(t,\nu^n,\phi^n)\\
&\le (c_e+c_d)\cdot(\lambda_1+\beta/\sqrt{n})\cdot\frac{\delta^2}{\sqrt{n}}\\
&\quad+(c_e+c_d)\mathbb{E}_{\nu^n}\left[\int_0^t\left|\frac{1}{\sqrt{n}}(E_1^n(s)-E_1^n(s-\delta^n))-\lambda_1\delta\right|ds\right]\\
&\quad-c_d\delta\beta_1 t/\sqrt{n}+c_d\widetilde{\Delta}_1^n(t,\nu^n,\psi^n)+c_w\widetilde{\Delta}_2^n(t,\nu^n,\psi^n).
\end{aligned}\tag{22}
$$

Because $\mathfrak{X}^n(\cdot)$ has a unique stationary distribution under the policy ϕ^n for large enough n, we can obtain from Little's law for stationary systems (Serfozo 2012, chapter 6.4) that

$$\lim_{t\to\infty}\frac{1}{t}\widetilde{\Delta}_k^n(t,\nu^n,\phi^n)=0,\quad \text{for } k = 1,\, 2.$$

By combining (21), we deduce from (22) that (13) holds. This completes the proof. □

4.3. Analysis of the Problem Involving the Cost of Queue Lengths

From Theorem 2, we can focus on the analysis of the expected total cost (11) that involves the cost of queue lengths. For fixed n, the control problem involving (11) has a similar structure to the problem in Benjaafar et al. (2010), where they consider the optimal control of a production-inventory system with both backorders and lost sales. It can also be formulated as a continuous time Markov decision process (MDP) and numerically solved, for example, by policy iteration. However, computing such an MDP policy is computationally expensive in general. In this note, we develop the policy ψ_*^n, which is asymptotically optimal in heavy traffic and easier to compute and implement. We also emphasize that our original control problem has an objective that involves the earliness/tardiness/waiting costs of customers, and this problem cannot be directly studied using MDP because the dynamics of the waiting times depend on future decisions.

Proposition 1 shows that γ^* in Lemma 1 serves as a lower bound for the long-run average cost under any admissible sequence of control policies $\{\psi^n\}$ and initial distributions $\{\nu^n\}$. Proposition 2 shows that the long-run average cost under the proposed policy is asymptotically γ^*. The proofs of these two results are deferred to Section 5.

Proposition 1. *For any admissible sequence of control policies $\{\psi^n\}$ and initial distributions $\{\nu^n\}$, we have*

$$\limsup_{n\to\infty}\limsup_{t\to\infty}\frac{1}{t}\widetilde{G}^n(t,\nu^n,\psi^n)\ge\gamma^*.\tag{23}$$

Proposition 2. *Under the proposed sequence of policies $\{\psi_*^n\}$, there exists a unique stationary distribution of $\mathfrak{X}_*^n(\cdot)$ for large enough n, and*

$$\lim_{n\to\infty}\lim_{t\to\infty}\frac{1}{t}\widetilde{G}^n(t,\nu^n,\psi_*^n)=\gamma^*.\tag{24}$$

Theorem 1 easily follows from Theorem 2 and Propositions 1 and 2.

5. Additional Proofs

In this section, we collect the proofs of Propositions 1 and 2 and Lemma 1. For the sake of notational simplicity, in the following proofs, we denote by C the generic positive constants that are independent of n, although the value of C may differ from line to line.

5.1. Proofs of Propositions 1 and 2

The proofs of Proposition 1 and 2 are similar to the proofs of theorems 2 and 1 in Gao and Huang (2023); therefore, we have only outlined the key steps.

We first describe the dynamic equations of the diffusion-scaled queue length processes $\widetilde{Q}^n(\cdot)=\{(\widetilde{Q}_1^n(t),\widetilde{Q}_2^n(t));\, t\ge 0\}$ and the nominal workload process $\widetilde{W}^n(\cdot)=\{\widetilde{W}^n(t);\, t\ge 0\}$ defined in (10). Define $\rho_k=\lambda_k/\mu_k$ to be the nominal workload of class k orders for $k=1,2$. Let

$$
\begin{aligned}
&\widetilde{X}_k^n(t)=\frac{S_k^n(T_k^n(t))-\mu_k^n T_k^n(t)}{\sqrt{n}}-\frac{E_k^n(t)-\lambda_k^n t}{\sqrt{n}},\quad\text{and}\\
&\widetilde{Y}_k^n(t)=\sqrt{n}(\rho_k t-T_k^n(t)).
\end{aligned}\tag{25}
$$

Then, the dynamics of $\widetilde{Q}^n$ under a control policy ψ^n is $\left(\text{recall that } \widetilde{O}_k^n(t)=\frac{O_k^n(t)}{\sqrt{n}}\right)$:

$$
\begin{aligned}
\widetilde{Q}_k^n(t)&=\widetilde{Q}_k^n(0)-\frac{S_k^n(T_k^n(t))}{\sqrt{n}}+\frac{E_k^n(t)}{\sqrt{n}}-\widetilde{O}_k^n(t)\\
&=\widetilde{Q}_k^n(0)-\widetilde{X}_k^n(t)+\mu_k\widetilde{Y}_k^n(t)+\beta_k t-\widetilde{O}_k^n(t).
\end{aligned}\tag{26}
$$

Define $\widetilde{I}^n(t):=\sum_{k=1}^2\widetilde{Y}_k^n(t)$. Using $\sum_{k=1}^2\rho_k=1$, one has $\widetilde{I}^n(t)=\sqrt{n}\left(t-\sum_{k=1}^2 T_k^n(t)\right)=\sqrt{n}I^n(t)$, where $I^n(\cdot)$ is the cumulative idle time process. Then, from (10) and (26), for $t \ge 0$,

$$\widetilde{W}^n(t) = \widetilde{W}^n(0) - \widetilde{\mathbb{X}}^n(t) + \widetilde{I}^n(t) + \sum_{k=1}^{2} \frac{\beta_k}{\mu_k} t - \widetilde{\mathbb{O}}^n(t), \tag{27}$$

in which

$$\widetilde{\mathbb{X}}^n(t) := \sum_{k=1}^{2} \frac{\widetilde{X}_k^n(t)}{\mu_k} \quad \text{and} \quad \widetilde{\mathbb{O}}^n(t) := \sum_{k=1}^{2} \frac{\widetilde{O}_k^n(t)}{\mu_k}. \tag{28}$$

We apply Ito's formula, with Φ in Lemma 1, to the semimartingale $\widetilde{W}^n$ in (27), and take Taylor expansion to obtain

$$\begin{aligned}\Phi(\widetilde{W}^n(t)) &= \Phi(\widetilde{W}^n(0)) + \int_0^t \Phi'(\widetilde{W}^n(s-))d\widetilde{W}^n(s)\\ &\quad + \sum_{s\le t:|\triangle\widetilde{\mathbb{X}}^n(s)|>0} \left(\frac{1}{2}\Phi''(\widetilde{W}^n(s-))(\triangle\widetilde{\mathbb{X}}^n(s))^2\right)\\ &\quad - \sum_{s\le t:|\triangle\widetilde{\mathbb{O}}^n(s)|>0} \left(\frac{1}{2}\Phi''(\widetilde{W}^n(s-))(\triangle\widetilde{\mathbb{X}}^n(s))^2\right)\\ &\quad + \sum_{s\le t:|\triangle\widetilde{W}^n(s)|>0} \frac{1}{6}\Phi'''(\Xi^n(s))(\triangle\widetilde{W}^n(s))^3,\end{aligned}$$

where the notation $\triangle\mathbb{Y}(t) = \mathbb{Y}(t) - \mathbb{Y}(t-)$ denotes the jump of a process $\mathbb{Y}$ at time t, and we use $|\triangle\widetilde{W}^n(t)| = |\triangle\widetilde{\mathbb{X}}^n(t)| - |\triangle\widetilde{\mathbb{O}}^n(t)|$. The term $\Xi^n(s)$ lies in between $\min\{\widetilde{W}^n(s), \widetilde{W}^n(s-)\}$ and $\max\{\widetilde{W}^n(s), \widetilde{W}^n(s-)\}$. Under a feasible Markov control policy, if the Markov chain $\mathfrak{X}^n(\cdot)$ has a stationary distribution, denoted by π^n, then we can use (25), (27), and (28) to obtain

$$\begin{aligned}0 &= \mathbb{E}_{\pi^n}[\Phi(\widetilde{W}^n(t))] - \mathbb{E}_{\pi^n}[\Phi(\widetilde{W}^n(0))]\\ &= \Psi_1(\pi^n, t) + \Psi_2(\pi^n, t) + \Psi_3(\pi^n, t),\end{aligned} \tag{29}$$

where

$$\begin{aligned}&\Psi_1(\pi^n, t)\\ &= \mathbb{E}_{\pi^n}\left[\int_0^t \left(\sum_{k=1}^{2} \frac{\beta_k}{\mu_k}\Phi'(\widetilde{W}^n(s-)) + \sum_{k=1}^{2} \frac{\lambda_k}{\mu_k^2}\Phi''(\widetilde{W}^n(s-))\right)ds\right],\end{aligned} \tag{30}$$

$$\begin{aligned}&\Psi_2(\pi^n, t)\\ &= \mathbb{E}_{\pi^n}\left[-\int_0^t \Phi'(\widetilde{W}^n(s-))d\widetilde{\mathbb{O}}^n(s) + \int_0^t \Phi'(\widetilde{W}^n(s-))d\widetilde{I}^n(s)\right],\end{aligned} \tag{31}$$

$$\begin{aligned}\Psi_3(\pi^n, t) &= \sum_{k=1}^{2} \frac{1}{\mu_k^2} \times \mathbb{E}_{\pi^n}\left[\int_0^t \frac{1}{2}\Phi''(\widetilde{W}^n(s-))\left(\frac{\beta_k}{\sqrt{n}}ds - \frac{\mu_k}{\sqrt{n}}d\widetilde{Y}_k^n(s)\right)\right]\\ &\quad - \mathbb{E}_{\pi^n}\left[\sum_{s\le t:|\triangle\widetilde{\mathbb{O}}^n(s)|>0} \left(\frac{1}{2}\Phi''(\widetilde{W}^n(s-))(\triangle\widetilde{\mathbb{X}}^n(s))^2\right)\right]\\ &\quad + \mathbb{E}_{\pi^n}\left[\sum_{s\le t:|\triangle\widetilde{W}^n(s)|>0} \frac{1}{6}\Phi'''(\Xi^n(s))(\triangle\widetilde{W}^n(s))^3\right].\end{aligned} \tag{32}$$

Proof of Proposition 1. As in Gao and Huang (2023) (see their lemma 4), to prove (23), it suffices to consider n with $n \in \left\{n : \inf_{\nu^n, \psi^n \in \Pi^n} \limsup_{t\to\infty} \frac{1}{t}\widetilde{G}^n(t, \nu^n, \psi^n) \le \gamma^* + 1\right\}$. For such n, one can prove the existence of a stationary distribution of $\mathfrak{X}^n(\cdot)$, and it suffices to show

$$\limsup_{n\to\infty} \limsup_{t\to\infty} \frac{1}{t}\widetilde{G}^n(t, \pi^n, \psi^n) \ge \gamma^*, \tag{33}$$

where π^n is the stationary distribution of the Markov chain $\mathfrak{X}^n(\cdot)$ under the policy ψ^n. As a result, we can apply (29) and analyze the three terms (Ψ_1, Ψ_2, and Ψ_3) in the following.

Note that Φ satisfies Condition (8). Hence, using the definition of h in (5), we obtain $\Psi_1(\pi^n, t) \ge \gamma^* t - \mathbb{E}_{\pi^n}\left[\int_0^t h(\widetilde{W}^n(s-))ds\right] \ge \gamma^* t - \mathbb{E}_{\pi^n}\left[\int_0^t \sum_{k=1}^2 f_k(\widetilde{Q}_k^n(s-))ds\right]$. Using (9), the fact that $\widetilde{I}^n(\cdot)$ is nondecreasing and the definition of κ in (6), we have $\Psi_2(\pi^n, t) \ge -\kappa\mathbb{E}_{\pi^n}\left[\widetilde{\mathbb{O}}^n(t)\right] \ge -\sum_{k=1}^2 \theta_k \mathbb{E}_{\pi^n}\left[\widetilde{O}_k^n(t)\right]$. Using the definitions of $\widetilde{W}^n$, $\widetilde{\mathbb{X}}^n$, and the fact that $|\Phi''(x)|, |\Phi'''(x)| \le C$, we can make a similar argument to the one in the proof of theorem 2 in Gao and Huang (2023) and deduce that

$$\begin{aligned}|\Psi_3(\pi^n, t)| &\le \hat{\Psi}_3(\pi^n, t) + \mathbb{E}_{\pi^n}\left[\frac{C}{2\sqrt{n}}\widetilde{\mathbb{O}}^n(t)\right]\\ &\le \hat{\Psi}_3(\pi^n, t) + \frac{C}{\sqrt{n}}\widetilde{G}^n(t, \pi^n, \psi^n),\end{aligned} \tag{34}$$

where

$$\begin{aligned}\hat{\Psi}_3(\pi^n, t) &:= \left|\mathbb{E}_{\pi^n}\left[\sum_{k=1}^{2} \frac{1}{\mu_k}\int_0^t \Phi''(\widetilde{W}^n(s-))\frac{1}{2\sqrt{n}}d\widetilde{Y}_k^n(s)\right]\right|\\ &\quad + \mathbb{E}_{\pi^n}\left[\int_0^t \sum_{k=1}^{2} \frac{C}{\mu_k^2}\frac{|\beta_k|}{2\sqrt{n}}ds\right]\\ &\quad + \frac{C}{6\sqrt{n}}\frac{\mathbb{E}_{\pi^n}\left[\sum_{k=1}^2 (\lambda_k^n t + \mu_k^n T_k^n(t))\right]}{n}.\end{aligned}$$

It then follows from (29) that $\left(1 + \frac{C}{\sqrt{n}}\right)\widetilde{G}^n(t, \pi^n, \psi^n) - \gamma^* t \ge -\hat{\Psi}_3(\pi^n, t)$. Similar to (46) in Gao and Huang (2023), one can show that $\lim_{n\to\infty}\lim_{t\to\infty}\hat{\Psi}_3(\pi^n, t)/t = 0$. It then follows that (33) holds. Hence, the proof of Proposition 1 is complete. □

Proof of Proposition 2. For large enough n, under ψ_*^n and any initial distribution ν^n, there is a stationary distribution π_*^n of $\mathfrak{X}_*^n(\cdot)$ (see proposition 1 in Gao and Huang (2023)). Moreover, every state of $\mathfrak{X}_*^n(\cdot)$ can reach the state (0, 0, 1), so there is at most one positive recurrent class. Hence the stationary distribution is unique. Therefore, we have $\lim_{n\to\infty}\lim_{t\to\infty}\frac{1}{t}\widetilde{G}^n(t, \nu^n, \psi_*^n) = \lim_{n\to\infty}\lim_{t\to\infty}\frac{1}{t}\left[\widetilde{G}^n(t, \pi_*^n, \psi_*^n)\right]$. As a result, we can assume the initial distribution of $\mathfrak{X}_*^n(\cdot)$ is π_*^n and then apply (29). For notational simplicity, we use $\mathbb{E}$ to denote $\mathbb{E}_{\pi_*^n}$.

Because $\Phi'(x)=0$ for $x\in\left[-\frac{\lambda_1\delta}{\mu_1},l_*\right]$, $\Phi'(x)=\kappa$ for $x\geq u_*$, and $\Phi''(x)=0$ for $x\notin(l_*,u_*)$, one can verify that $\frac{\sigma^2}{2}\Phi''(x)+\sum_{k=1}^2\frac{\beta_k}{\mu_k}\Phi'(x)+h(l_*\vee(x\wedge u_*))=\gamma^*$. From (30), we can infer that

$$\begin{aligned}\Psi_1(\pi_*^n,t)&=\gamma^*t-\mathbb{E}\left[\int_0^t h\left(l_*\vee(\widetilde{W}_*^n(s)\wedge u_*)\right)ds\right]\\&=\gamma^*t-\mathbb{E}\left[\sum_{k=1}^2\int_0^t f_k(\widetilde{Q}_{k*}^n(s))ds\right]\\&\quad+\mathbb{E}\left[\sum_{k=1}^2\int_0^t f_k(\widetilde{Q}_{k*}^n(s))ds-\int_0^t h\left(l_*\vee(\widetilde{W}_*^n(s)\wedge u_*)\right)ds\right].\end{aligned}$$

Under the policy ψ_*^n, when $\widetilde{\mathbb{O}}_*^n$ jumps (i.e., an order is rejected), we have $\widetilde{W}_*^n(t)\geq u_*>0$. On the other hand, when $\widetilde{I}_*^n(s)$ increases (i.e., the system becomes idle), we have $\frac{\tilde{Q}_{1*}^n(t)-\lambda_1\delta}{\mu_1}<l_*$ and $\widetilde{Q}_{2*}^n(t)=0$. Hence, $\widetilde{W}_*^n(t)=\frac{\tilde{Q}_{1*}^n(t)-\lambda_1\delta}{\mu_1}+\frac{\tilde{Q}_{2*}^n(t)}{\mu_2}<l_*$. From Lemma 1, $\Phi'(x)=0$ for $x\in\left[-\frac{\lambda_1\delta}{\mu_1},l_*\right]$, and $\Phi'(x)=\kappa$ for $x\geq u_*$. Hence, by (31), we have $\Psi_2(\pi_*^n,t)=-\kappa\cdot\mathbb{E}[\widetilde{\mathbb{O}}_*^n(t)]$. In addition, under the proposed policy ψ_*^n, only class i^* orders will be rejected. Using the definition of $\kappa=\theta_{i^*}\cdot\mu_{i^*}$, we have $\kappa\cdot\widetilde{\mathbb{O}}_*^n(t)=\theta_{i^*}\widetilde{O}_{i^**}^n(t)=\sum_{k=1}^2\theta_k\widetilde{O}_{k*}^n(t)$. It follows that $\Psi_2(\pi_*^n,t)=-\sum_{k=1}^2\theta_k\mathbb{E}[\widetilde{O}_{k*}^n(t)]$. Finally, $\Psi_3(\pi_*^n,t)$ can be upper bounded, as in (34). Hence, from (29), we obtain

$$\begin{aligned}\left(1-\frac{C}{\sqrt{n}}\right)\widetilde{G}^n(t,\pi_*^n,\psi_*^n)-\gamma^*t&\leq\hat{\Psi}_3(\pi_*^n,t)\\&\quad+\mathbb{E}\left[\left|\sum_{k=1}^2\int_0^t f_k(\widetilde{Q}_{k*}^n(s))ds-\int_0^t h(l_*\vee(\widetilde{W}_*^n(s)\wedge u_*))ds\right|\right].\end{aligned}$$

Similar to the proof of Proposition 1, we can show that $\lim_{n\to\infty}\lim_{t\to\infty}\hat{\Psi}_3(\pi_*^n,t)/t=0$. Moreover, if we divide the last term of the right-hand side of the previous equation by t, send t to infinity first, and then n to infinity, we obtain that the limit is zero by establishing a steady-state state-space collapse result (similar to proposition 3 in Gao and Huang (2023)). Then, $\lim_{n\to\infty}\lim_{t\to\infty}\frac{1}{t}\left[\widetilde{G}^n(t,\pi_*^n,\psi_*^n)\right]\leq\gamma^*$. Together with Proposition 1, we obtain (24). □

5.2. Proof of Lemma 1

We extend the domain of the function h in (5) so that $h(y)=\min\{c_d\mu_1,c_w\mu_2\}\cdot y1_{\{y>0\}}-c_e\mu_1y1_{\{y\leq 0\}}$. By lemma 15 in Gao and Huang (2023), we have the following three results:

(1) For $w_0,\gamma\in\mathbb{R}$, there is a unique continuously differentiable solution $w(x;w_0,\gamma)$ to the ODE:

$$\frac{\sigma^2}{2}w'(x)+\sum_{k=1}^2\frac{\beta_k}{\mu_k}w(x)+h(x)=\gamma,\ \text{for } x\in\mathbb{R},\qquad \text{subject to } w(0)=w_0; \tag{35}$$

(2) For $w_0<\kappa$, there is a unique $\gamma_+(w_0)$, such that $\max_{x\geq 0}w(x;w_0,\gamma_+(w_0))=\kappa$, and the function $\gamma_+(\cdot)$ is continuous and strictly decreasing; and

(3) There is $w_*\in(0,\kappa)$, such that $\max_{x\geq 0}w(x;w_*,\gamma_+(w_*))=\kappa$ and $\min_{x\leq 0}w(x;w_*,\gamma_+(w_*))=0$.

Denote by $\underline{l}=\operatorname{argmin}_{x\leq 0}w(x;w_*,\gamma_+(w_*))$. We consider two cases.

Case 1. If $\underline{l}\geq-\frac{\lambda_1\delta}{\mu_1}$: Set $\gamma_*=\gamma_+(w_*)$, $l_*=\underline{l}$, $v(x)=w(x;w_*,\gamma_+(w_*))$, and $u_*=\operatorname{argmax}_{x\geq 0}v(x)$;

Case 2. If $\underline{l}<-\frac{\lambda_1\delta}{\mu_1}$: From part 1 of lemma 15 in Gao and Huang (2023), $w(\cdot;w_*,\gamma_+(w_*))$ has no local maximizer on $(-\infty,0)$ and cannot be a constant in any interval; hence, it is strictly increasing on $[\underline{l},0]$ and $w\left(-\frac{\lambda_1\delta}{\mu_1};w_*,\gamma_+(w_*)\right)>0$. Next, we show $w\left(-\frac{\lambda_1\delta}{\mu_1};0,\gamma_+(0)\right)<0$. Let w_0 decrease from w_* to zero, then $\gamma_+(w_0)$ increases; hence, from part 1 of lemma 15 in Gao and Huang (2023), $w(\underline{l};w_0,\gamma_+(w_0))$ decreases; thus, $w(\underline{l};0,\gamma_+(0))<0$. As $-\frac{\lambda_1\delta}{\mu_1}\in(\underline{l},0)$, and $w(\cdot;0,\gamma_+(0))$ has no local maximizer on $(-\infty,0)$ and $w(\underline{l};0,\gamma_+(0))<0=w(0;0,\gamma_+(0))$, one has $w\left(-\frac{\lambda_1\delta}{\mu_1};0,\gamma_+(0)\right)<0$. As a result, there exists $\overline{w}_*\in(0,w_*)$, such that $w\left(-\frac{\lambda_1\delta}{\mu_1};\overline{w}_*,\gamma_+(\overline{w}_*)\right)=0$. Set $\gamma_*=\gamma_+(\overline{w}_*)$, $l_*=-\frac{\lambda_1\delta}{\mu_1}$, $v(x)=w(x;\overline{w}_*,\gamma_+(\overline{w}_*))$ and $u_*=\operatorname{argmax}_{x\geq 0}v(x)$.

For both cases, we have the constants l_*,u_*,γ^*, and the function $v(\cdot)\in C^1([l_*,u_*])$ satisfies $\frac{\sigma^2}{2}v'(x)+\sum_{k=1}^2\frac{\beta_k}{\mu_k}v(x)+h(x)=\gamma^*$ for $x\in(l_*,u_*)$, with $v(x)\in[0,\kappa]$ for $x\in[l_*,u_*]$, $v(l_*)=0$, $v(u_*)=\kappa$, and $v'(u_*)=0$. The uniqueness of the constants l_*,u_*,γ^* and the function v can be established similarly to the proof of uniqueness in lemma 2 of Gao and Huang (2023). Let $\Phi(x)=\int_{l_*}^x\overline{v}(y)dy$ with $\overline{v}(x)=0\times 1_{\left\{x\in\left[-\frac{\lambda_1\delta}{\mu_1},l_*\right)\right\}}+v(x)\times 1_{\{x\in[l_*,u_*)\}}+\kappa\times 1_{\{x\in[u_*,\infty)\}}$. Then, it is straightforward to check that $\Phi(\cdot)$ satisfies the conditions in Lemma 1. The uniqueness of the function $\Phi(\cdot)$ follows from that of v. The proof is therefore complete. □

Acknowledgments

The authors thank the editors and referees for constructive comments and suggestions.

References

Afeche P (2013) Incentive-compatible revenue management in queueing systems: Optimal strategic delay. *Manufacturing Service Oper. Management* 15(3):423–443.

Armony M, Maglaras C (2004) On customer contact centers with a call-back option: Customer decisions, routing rules, and system design. *Oper. Res.* 52(2):271–292.

Atar R, Giat C, Shimkin N (2010) The $c\mu/\theta$ rule for many server queues with abandonment. *Oper. Res.* 58(5):1427–1439.

Baron O, Berman O, Krass D, Wang J (2014) Using strategic idleness to improve customer service experience in service networks. *Oper. Res.* 62(1):123–140.

Baron O, Berman O, Krass D, Wang J (2017) Strategic idleness and dynamic scheduling in an open-shop service network: Case study and analysis. *Manufacturing Service Oper. Management* 19(1):52–71.

Baron O, Chen X, Li Y (2023) Omnichannel services: The false premise and operational remedies. *Management Sci.* Forthcoming.

Benjaafar S, ElHafsi M, Huang T (2010) Optimal control of a production-inventory system with both backorders and lost sales. *Naval Res. Logist.* 57(3):252–265.

Chen M, Hu M, Wang J (2022) Food delivery service and restaurant: Friend or foe? *Management Sci.* 68(9):6539–6551.

Cox DR, Smith W (1961) *Queues* (Chapman & Hill/CRC, Boca Raton, FL).

Dai JG, Yao D (2013) Brownian inventory models with convex holding cost, part 1: Average-optimal controls. *Stochastics Systems* 3(2):442–499.

Farahani MH, Dawande M, Janakiraman G (2022) Order now, pickup in 30 minutes: Managing queues with static delivery guarantees. *Oper. Res.* 70(4):2013–2031.

Gao X, Huang J (2023) Asymptotically optimal control of make-to-stock systems. *Math. Oper. Res.* Forthcoming.

Heyman DP, Stidham S Jr (1980) The relation between customer and time averages in queues. *Oper. Res.* 28(4):983–994.

Markowitz DM, Wein LM (2001) Heavy traffic analysis of dynamic cyclic policies: A unified treatment of the single machine scheduling problem. *Oper. Res.* 49(2):246–270.

Serfozo R (2012) *Introduction to Stochastic Networks*, vol. 44 (Springer Science & Business Media, Boston).

Stidham S Jr (1974) A last word on $L = \lambda W$. *Oper. Res.* 22(2):417–421.

van Mieghem JA (1995) Dynamic scheduling with convex delay costs: The generalized $c\mu$ rule. *Ann. Appl. Probability* 5(3):809–833.

Wein LM (1992) Dynamic scheduling of a multiclass make-to-stock queue. *Oper. Res.* 40(4):724–735.

Xuefeng Gao is an associate professor at the Department of Systems Engineering and Engineering Management at The Chinese University of Hong Kong. His research interests include stochastic modeling, analysis, and control.

Junfei Huang is an associate professor in the Department of Decision Sciences and Managerial Economics at the Chinese University of Hong Kong. His research interests are in asymptotic analysis and optimal control of queueing systems and their applications in manufacturing and services.

Jiheng Zhang is a professor in the Department of Industrial Engineering and Decision Analytics at the Hong Kong University of Science and Technology. His research interests are in applied probability, stochastic modeling and optimization, data analysis, numerical methods, and algorithms.