Probability in the Engineering and Informational Sciences, **32**, 2018, 353–395.
doi:10.1017/S0269964817000213

DYNAMIC CONTROL OF A SINGLE-SERVER SYSTEM WHEN JOBS CHANGE STATUS

GABRIEL ZAYAS-CABÁN
Center for Healthcare Engineering and Patient Safety,
University of Michigan,
Ann Arbor, MI, USA
E-mail: gzayasca@umich.edu

HYUN-SOO AHN
Ross School of Business,
University of Michigan,
Ann Arbor, MI, USA
E-mail: hsahn@umich.edu

From health care to maintenance shops, many systems must contend with allocating resources to customers or jobs whose initial service requirements or costs change when they wait too long. We present a new queueing model for this scenario and use a Markov decision process formulation to analyze assignment policies that minimize holding costs. We show that the classic $c\mu$ rule is generally not optimal when service or cost requirements can change. Even for a two-class customer model where a class 1 task becomes a class 2 task upon waiting, we show that additional orderings of the service rates are needed to ensure the optimality of simple priority rules. We then show that seemingly-intuitive switching curve structures are also not optimal in general. We study these scenarios and provide conditions under which they do hold. Lastly, we show that results from the two-class model do not extend to when there are $n \geq 3$ customer classes. More broadly, we find that simple priority rules are not optimal. We provide sufficient conditions under which a simple priority rule holds. In short, allowing service and/or cost requirements to change fundamentally changes the structure of the optimal policy for resource allocation in queueing systems.

Keywords: operations research, queues, stochastic dynamic programming, stochastic modeling

1. INTRODUCTION

Scheduling is an important topic in Operations Research and Management Science, for which a central question is how to allocate resources in queueing systems. This question is typically studied by assuming there are one or more customer classes that require service at one or more service stations. Servers decide where to allocate efforts to optimize some performance criterion related to quality of service and/or congestion. Customers are usually idealized as having processing requirements or costs that are known *a priori* and do not change over time, but there are many practical and important situations where this is not the case.

© Cambridge University Press 2017 0269-9648/17 $25.00

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

In emergency departments (EDs), patients are classified or "triaged" based on urgency, required activities, and processing times. Patients within the same triage category have similar urgency and processing requirements, but upon waiting, some patients can see changes in their health. For example, EDs grapple with cohorts of patients who have symptoms (e.g., chest pain, psychosis, abdominal pain) for which a physician cannot make a definitive diagnosis. A physician can run a battery of tests and prepare for the most critical outcome. However, for the large majority of these patients, their condition ameliorates after some time at which point it becomes quicker and less costly to treat the patient.

On the other hand, patients may also deteriorate in the ED, leading to higher hospitalization rate, processing time, and costs. Similarly, patients can be admitted to the inpatient unit, but wait or "board" in a hospital's ED due to insufficient space in the inpatient unit. Excessive boarding times can also lead to adverse outcomes for both boarding patients and patients waiting to be treated, including a further deterioration of the patient's health. In both cases, health providers choose between caring for recently arrived patients or already-deteriorated patients.

Maintenance shops are also faced with jobs that change status. To illustrate, consider the following example by Armstrong [4] on aircrafts maintenance shops. In this setting, aircrafts undergo periodic review in a maintenance shop after being used for a predetermined amount of time. The next aircraft to be inspected may have a considerable number of hours left in its rated life, so there is a question of whether to bring it in for maintenance or to leave the shop empty until its life is exhausted. Conversely, an airplane can reach its rated life before the shop is finished with the previous one; should the aircraft continue to fly or sit idle until the shop is ready?

Here, we analyze this less studied, yet important question at the heart of all these settings: how to allocate servers in a queueing system in which customers or jobs can change status. A server must decide whether to continue to serve a customer whose service requirements and/or costs have changed or switching to serve recently-arrived customers. Their decision determines, considering costs, processing times, and rates at which jobs change status, whether they serve recently-arrived customers before jobs change status or serve customers after they change status.

We focus on a single-server system in which arriving customers with initial service requirements or costs can change while they are waiting. Initial requirements for service can change possibly to become less or more time consuming or less or more costly. There is a Poisson arrival stream to a system with infinite waiting capacity and a single server that serves arriving customers having initial service times that are exponentially distributed. To this setup, we add that initial processing requirements or costs may change after an exponentially distributed period of time and that once changed, the new service for this deteriorated job is also exponentially distributed. Our goal is to provide an optimal server assignment policy where each customer class has (linear) holding costs. We consider the problem of minimizing expected discounted or average costs over an infinite horizon.

Allowing for jobs to change status turns out to be important in determining not only what type of policy is optimal, but also what techniques can be used to show structural properties of optimal policies. When changes in processing requirements or costs are not allowed, the classic $c\mu$ rule is optimal (c.f. Buyukkoc, Varaiya, and Walrand [12]). This is not always the case when jobs can change status, as we show in this paper. The reason is that changing status leads to trade-offs between myopically reducing costs (via the $c\mu$ rule) and reducing costs as quickly as possible (as a consequence of changes in status). For appropriate orderings of the service rates, this trade-off disappears, in which case the $c\mu$ rule is optimal. We identify this ordering. Moreover, since customers may change status between services interchange arguments are difficult to apply except in very simple scenarios. Additionally,

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

we must contend technical challenges associated with unbounded transition rates that are a consequence of jobs changing status.

To analyze this decision-making scenario, the methodology that we use is that of Markov decision processes (MDPs). To our knowledge, we are the first to use MDPs to determine dynamic scheduling for a single, flexible server in a queueing system where customer attributes (e.g., processing times and/or costs) change. We also see our work as having the following contributions:

- When there are two customer classes (i.e., when a class 1 job waits for too long, and then becomes a class 2 job), we show that the classic cμ ordering does not guarantee the optimality of static priority rules. We identify sufficient conditions under which the cμ policy indeed becomes optimal. We then analyze the structure of optimal controls when these conditions are violated. In particular, we provide counterexamples showing that seemingly-intuitive switching curve policies are no longer optimal, and also provide conditions under which they do hold. We also develop and analyze (numerically) heuristics for the case where we were unable to characterize the optimal policy.
- We then consider the general case when there are $n \geq 3$ customer classes. Surprisingly, sufficient conditions in the two-class model do not extend to this more general setting in order to guarantee the optimality of a simple priority rule. We explain why the results do not extend in general, and provide special cases of problem instances under which a simple priority rule holds.
- Lastly, we contend with handling MDPs with unbounded transition rates. In our model, transition rates are unbounded due to customer changing status so that uniformization (c.f. Lippman [24]) is not possible. We use the continuous-time MDP framework to discern relationships of the value functions depending on the starting state. We then use sample path arguments to analyze structural properties of the optimal policy. We also truncate the multi-dimensional state space and linearly decrease the rates as the state reaches the boundary of the truncation. This allows us to obtain additional structural properties for the original unbounded model, by first showing them for each truncated, and hence, uniformizable model, and then taking the limit as the truncation goes to infinity.

The rest of the paper is organized as follows: Section 2 contains a summary of the literature related to our work. Section 3 describes in detail the model we consider, the MDP formulation of the server allocation problem, and the relevant MDP theory that will be used throughout the paper. Section 4 contains our analysis for the two-class model while Section 5 contains an analysis for the general, n-class model. We propose several allocation policies based on our analytical and numerical examples in Sections 4 and 5, respectively, and then study them numerically in Section 6. We conclude with a discussion in Section 7.

2. LITERATURE REVIEW

Work in the area of scheduling is too plentiful to provide a complete review here. Instead, we refer the reader to Pinedo's classic text [28] and the handbook by Leung [23] for systematic and comprehensive treatments and review only the work that closely relates to the current study.

Early work on scheduling with jobs that can change status is motivated by inventory issuing policies (c.f. Brown and Ross [9], Brown and Solomon [10]). In the classic scenario, there are a number of items, each having a unique "value". An item with an initial value is

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

kept in a stockpile and then released, with its field life depending on its release time. Items have released one-by-one upon the death of the previously issued item. The objective is to determine how to order the release of items that maximizes the total field life. A policy which releases items in increasing (decreasing) order of the item rating is called last in, first out (LIFO) [first in, first out (FIFO)]. Brown and Ross [9] consider deterministic field lives and give sufficient conditions for the optimality of LIFO and FIFO policies. Brown and Solomon [10] extend this work to the stochastic setting: they assume a stochastic field life with known distribution that changes in time in a known deterministic fashion. Here also, they provide sufficient conditions for the optimality of LIFO and FIFO policies in the stochastic order sense.

The previous two papers motivated the work by Browne and Yechiali [11]. Browne and Yechiali [11] suppose that there are a finite number of jobs available for processing at time 0, with some initial (random) requirements. If a job's processing time is delayed, the initial processing requirements deteriorates grow linearly in the delay time (with a rate that depends on the job). They assume that a job stops decaying as soon as it starts service. The authors considered the problem minimizing makespan (or variance of makespan) and provide optimal policies (within the class of non-idling, non-preemptive policies) that minimize the expected makespan, and in some cases, provide policies that minimize the variance of makespan. Similar to our problem (but without arrivals), they also consider policies that minimize the sum of weighted completion times, and provide sufficient conditions for when simple priority rules hold.

Browne and Yechiali [11] also point to several open problems. One is a conjecture that the weighted completion time problem is NP-hard, a proof of which was provided by Bachman, Janiak, and Kovalyov [5], and others include extensions from linear to non-linear type of deteriorations, such as processing times that increase exponentially with delays. This has served as baseline for several extensions in this area. For example, Mosheiov [26,27] propose and analyze intuitive classes of policies for scheduling deteriorating jobs with the objective of minimizing flowtime and weighted completion times, respectively. For comprehensive surveys of more recent work, we refer the reader to Alidaee and Womer [2] and Cheng, Ding, and Lin [15], and the more recent monograph by Cai, Wu, and Zhou [13].

Also related is scheduling of jobs that change status motivated by disaster response scenarios (c.f. Argon, Ziya, and Righter [3], Master, Chan, and Bambos [25], and the references therein) where medical resources (e.g., medical care providers, ambulances) may be overwhelmed by an unexpected increase in demand and need to be allocated to minimize treatment delays so as to reduce deaths and other adverse events. To provide insights into the patient triage mass casualty incidents, Argon et al. in [3] consider a single-server clearing system where changes in status correspond to jobs abandoning the system when they wait for too long. In their model, jobs are determined by their lifetime and service time distributions and the objective is to determine optimal or near-optimal dynamic allocation policies that minimize abandonments. More recently, Master et al. in [25] consider a discrete time, multi-server clearing system where changes in status are modeled by a reward function associated with completing each job that decreases over time; that is, they account for each job having an internal value which may decay over time. The function used to model decay is deterministic, non-increasing, but otherwise arbitrary, generalizing the idea considered in Argon et al. [3] of having jobs with deadlines to an arbitrary decay in reward. Because the problem is intractable, they analyze the performance of three heuristics: one which maximizes the expected immediate reward, one which maximizes the expected immediate reward rate, and one that prioritizes jobs with the most imminent deadlines.

We remark that although all of these papers do not consider arrivals (i.e., queueing systems) or random deteriorations time with known distributions, they do provide useful

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

insights into the types of structural properties that an optimal policy might have (e.g., the optimality of FIFO or LIFO policies).

Also related are maintenance and replacement models (c.f. Stengos and Thomas [36], Kawai [20], Kawai [21], Tijms and van der Duyn Schouten [37], and Srinivasan and Parlikad [35]). We refer the reader to work by Haque and Armstrong [19] for a recent, comprehensive survey and briefly mention work whose models somewhat relate to ours. Albright [1] is the first to allow for preventive service to be provided without taking the machine out of operation. The second is the work by Armstrong [4], where machines needing service may delay entering the repair queue under a *double age* preventive repair policy. Under this policy, a working machine that reaches its normal age for overhaul will attempt preventive repair. If all workers are busy at that time, it remains in operation and enters a repair until a worker becomes available, or when the machine reaches its second (higher) age limit, or when the machine breaks down and requires corrective repair. The authors here explore the use of age repair policies and seek an optimal age t at which to shut down machines for preventive repair.

While all of the previous work focuses on optimization, there are also a number of papers which focus on analyzing or approximating queueing systems with "deteriorating" service requirements using fixed policies. In particular, Whitt [38] and Boxma and Vlasiou [8] analyze a G/G/1 queueing system whose inter-arrival and service times depend linearly and randomly on the waiting times via a generalization of Lindley's classic recursion. Whitt [38] assumes that the workload decreases over time, whereas Boxma and Vlasiou [8] assumed that inter-arrival times increase as service rates decrease. Under these assumptions, both papers provide stability conditions for the system and approximations to the waiting time. More recently, and motivated by how delays to medical care can lead to potentially longer lengths of stay, Chan, Farias, and Escobar [14] analyze queueing systems with delay-dependent service times. In contrast to the previous papers, they allow for inter-arrival times that are not a function of service rates, so that the workload never decays. While each case above focuses on systems for a fixed policy, they can also lead to insights for other control models. For example, the author Whitt [38] uses insights derived in this paper to consider the problem of scheduling inter-arrival times, motivated by applications in production systems.

Lastly, applications of MDPs with unbounded rates are becoming widespread (c.f. Guo and Hernández-Lerma [18], Saghafian et al. [34], Legros, Jouini, and Koole [22]). For example, the authors Down, Koole, and Lewis [16] consider the dynamic control of a single server for a parallel system with abandonments. Although their scenario is distinctly different problem from ours, their approach (e.g., doing analysis in continuous time and truncating the state space), serve as useful models for how to analyze optimal controls for problems with unbounded rates. For instance, to obtain the structure of the optimal control, we also truncate the state space. In contrast to Down et al. [16], we linearly decrease the arrival rate until it reaches 0 at the boundary of the truncation. This approach can be justified by using recent results from Prieto-Rumeau and Hernández-Lerma [30], Bhulai, Brooms, and Spieksma [6], and Blok and Spieksma [7]. In particular, this framework allows us to analyze the structure of the optimal policy for the original unbounded model by analyzing the structure of the optimal policy for the simpler, uniformizable truncated model.

3. DYNAMICS AND MDP FORMULATION

In what follows, we use the term "deterioration" to describe events when jobs change status while waiting for the sake of simplicity. We consider the following single-server system with unlimited waiting space. Inter-arrival times and initial service requirements are assumed to

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

be exponential with rates λ and μ_1, respectively. When a customer joins the system, their (random and hidden) "deterioration" time is generated. We assume deterioration times are exponential with rates τ_i for a customer with service rates μ_i and that once the customer deteriorates, his/her service requirements change to an exponential random variable with (finite) rates μ_{i+1} immediately after deterioration ($i = 1, \ldots, n-1$). We also assume that customers that are in service are allowed to deteriorate. Henceforth, we will refer to customers whose service rates are μ_i as class i customers ($i = 1, \ldots, n$). Inter-arrival times, inter-departure times, and times spent before deteriorating constitute mutually independent sequences. Lastly, we assume that there is a non-negative holding cost h_i per job per unit time incurred for holding class i jobs ($i = 1, \ldots, n$). As alluded to, a single-service provider needs to decide who to service: customers of class 1, 2, …, or n. Within each class, the service discipline is first-come-first-served, and service can be preempted. Once a customer completes service, she/he leaves the system forever.

We analyze this decision-making scenario using an MDP formulation. The state space is given by $\mathbb{X} = \{x = (x_1, \ldots, x_n) \,|x_i \in \mathbb{Z}_{\geq 0}, i = 1, \ldots, n\}$, where $\mathbb{Z}_{\geq 0}$ is the set of non-negative integers. That is, the state $x = (x_1, \ldots, x_n)$ is a vector with x_i denoting the number of class i customers in the system. We seek a policy that after each event (i.e., arrival, service completion, or deterioration) time describes where to allocate the server based on the current state and history of states and actions taken (i.e., a non-anticipating policy). The finite horizon, discounted expected cost for a non-anticipating policy π is $v^{\pi}_{t,\alpha}(x) \equiv \mathbb{E}_x \int_0^t e^{-\alpha s} \left[\sum_{i=1}^n h_i Q_i^{\pi}(s)\right] ds$, where $Q_i^{\pi}(s)$ represents the queue length process for class i customers at time $s \geq 0$. For $\alpha > 0$, the infinite horizon discounted expected cost under policy π is $v^{\pi}_{\alpha}(x) \equiv \lim_{t\to\infty} v^{\pi}_{t,\alpha}(x)$. The long-run average cost rate is $\rho^{\pi}(x) \equiv \limsup_{t\to\infty} v^{\pi}_{t,0}/t$. Under either optimality criterion we seek a policy π^* such that $w^{\pi^*}(x) = \inf_{\pi\in\Pi} w^{\pi}(x)$ where Π is the set of all non-anticipating policies and $w = v_\alpha$ or ρ. Since the transition rates are unbounded, we are not guaranteed that each policy yields a regular Markov process. Regularity is guaranteed by showing that Assumptions B.3 and B.5 in the Appendix hold for the current model.

We begin our analysis with two intuitive results that are used throughout. The first is a statement on the monotonicity of the value functions, and says that it is better to have less customers in the system. The second says that there is an optimal policy that does not idle the server whenever there are customers waiting. The latter is used to simplify the optimality equations. In the interest of brevity, and because the proofs are simple but lengthy, we omit the proofs.

PROPOSITION 3.1: *The following hold:*

1. *For all* $x \in \mathbb{X}$

$$v_\alpha(x + e_i) \geq v_\alpha(x), \quad i = 1, \ldots, n,$$

where e_i *denotes the* i*th standard basis vector in* $\mathbb{R}^n$ $(i = 1, \ldots, n)$*. Similarly, if* (g, w) *is a solution to the average cost optimality equations (defined below), the above statements hold with* v_α *replaced with* w*.*

2. *Under the* α*-discounted reward (finite or infinite horizon) or the average reward criterion, there exists a (Markovian) non-idling policy that is optimal.*

To obtain additional properties, we rely on the optimality equations. Proposition 3.1 allows us to restrict attention to non-idling policies. Before we can use the optimality equations, we must first provide conditions that guarantee that the optimality equations under

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

each criterion has a solution. The proofs can be found in the Appendix. To simplify notation, let $c(x) = \sum_{i=1}^{n} h_i x_i$, and for any non-negative function v on $\mathbb{X}$ define the mapping T such that

$$Tv(x) = \lambda[v(x+e_1) - v(x)] + \sum_{i=1}^{n-1} x_i \tau_i [v(x - e_i + e_{i+1}) - v(x)]$$
$$+ \min_{i \in \{1,\ldots,n\}} \{\mathbb{1}_{\{x_i > 0\}} \mu_i [v(x - e_i) - v(x)]\}.$$

for $x \in \mathbb{X}$.

THEOREM 3.2: *Suppose $\alpha > 0$. The following hold:*

1. *The function v_α satisfies the discounted-cost optimality equation (DCOE), i.e.,*

$$\alpha v_\alpha = c + T v_\alpha.$$

2. *There exists a stationary, deterministic policy f_α (depending on the discount factor α) that attains the minimum in the right-hand side of the optimality equations, and hence, f_α is discounted-cost optimal.*

PROOF: See Appendix B. ■

THEOREM 3.3: *Suppose $\lambda < \min_{\{1 \le i \le n\}}\{\mu_i\}$. The following hold:*

1. *There exists a constant g and function w on the state space such that (g, w) satisfies the average-cost optimality equations (ACOEs),*

$$g\vec{1} = c + Tw,$$

where $\vec{1}$ is the vector of ones, g is the optimal average cost, and hence, is unique, and w, known as the relative value function, is unique up to additive constants.

2. *A deterministic stationary policy f is average cost optimal if and only if f satisfies the minimum in the ACOE.*

PROOF: See Appendix B. ■

A few remarks regarding Theorems 3.2 and 3.3 are in order. Note that Proposition 3.1 implies that we need only consider optimal policies that never idle so long as there's work to be done. It follows that any stationary non-idling policy yields a Markov chain that is irreducible (c.f. Prieto-Rumeau and Hernández-Lerma [30]). In words, this means that starting from any state $x \in \mathbb{X}$, it is possible to reach any other state $x' \in \mathbb{X}$ in finite time, with positive probability. Although the result holds for any n, we illustrate why this is the case for our model when $n = 2$. In this case, any arbitrary starting state (x_1, x_2) can reach state $(0, 0)$ by allowing all of the x_1 class 1 customers to consecutively become class 2 customers, which is then immediately followed by $x_1 + x_2$ consecutive class 2 customer service completions. Similarly, starting from state $(0, 0)$, we can reach any arbitrary state (x_1, x_2) by waiting for $x_1 + x_2$ consecutive arrivals to occur immediately followed by x_2 consecutive class 1 customers becoming class 2 customers.

In addition to irreducibility, we also require an ergodicity condition in order for Theorem 3.3 to hold. This is easily satisfied if $\lambda < \min_{\{1 \le i \le n\}}\{\mu_i\}$. Note that the condition is standard in that the arrival (birth) rate is smaller than the service (death) rate.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

The reason for this is that the number of customers can decrease along any of the n queues ($k = 1, \ldots, n$) and, hence, in order for the system to be stable under the optimal policy, the total arrival rate needs to be smaller than each of these n service rates. In particular, $\lambda < \min_{\{1 \le i \le n\}}\{\mu_i\}$ implies that our control model is uniformly h_A-exponentially ergodic on the class of stationary deterministic policies, where h_A is a Lyapunov function given in Assumption B.3 in the Appendix (c.f. Assumptions B.3 and B.9 in the Appendix). As a result, the vanishing discount approach, a standard technique for dealing with expected average cost by letting the discount rate decreases to zero in the α-discounted cost problem, can be used to obtain Theorem 3.3.

Lastly, note that from these theorems, we can determine when it is optimal to serve a particular class. Theorem 3.2 implies, for example, that when $x_i \ge 1$ in the discounted cost model, it is optimal to serve a class i customer provided $\mu_i[v_\alpha(x - e_i) - v_\alpha(x)] + \mu_j[v_\alpha(x) - v_\alpha(x - e_j)] \le 0$ for all $j \ne i$ such that $x_j \ge 1$. Conditions can similarly be found in the average case. In what follows, any result stated for the average cost criteria assumes that a solution (g, w) to the ACOE exists and that the result holds with v_α replaced with w.

4. OPTIMAL CONTROL FOR TWO JOB CLASSES ($n = 2$)

We first focus on when there are two customer classes. That is, if a class 1 job waits for too long, it becomes a class 2 job. To simplify notation for the two-customer-class case, for any non-negative function v on $\mathbb{X}$ we re-define the mapping T such that

$$\begin{aligned} Tv(i,j) = {} & \lambda[v(i+1,j) - v(i,j)] + i\tau[v(i-1,1) - v(i,0)] \\ & + \mathbb{1}_{\{i>0\}}\mu_1[v(i-1,j) - v(i,j)] + \mathbb{1}_{\{j>0\}}\mu_2[v(i,j-1) - v(i,j)] \\ & + \mathbb{1}_{\{i,j>0\}} \min\{\mu_1[v(i-1,j) - v(i,j)], \mu_2[v(i,j-1) - v(i,j)]\}. \end{aligned}$$

for $i, j \ge 0$.

As before, Theorem 3.2 implies, for example, that in the discounted cost model, for $i, j \ge 1$, it is optimal to serve a class 1 (2) customer if $\mu_1[v_\alpha(i-1,j) - v_\alpha(i,j)] \le (\ge)\mu_2[v_\alpha(i,j-1) - v_\alpha(i,j)]$, with a direct analog in the average cost case.

4.1. The Optimality of Static Priority Policies

In this section, we provide conditions under which one particular customer class should be prioritized. Note that when customers cannot change status (i.e., when $\tau = 0$), the model is simply a special case of the single-server system with parallel queues where the arrival rate for class 2 customers is zero. In this case, it is well known that the classic cμ rule is optimal: if $\mu_1 h_1 \ge \mu_2 h_2$, it is optimal to prioritize class 1 customers except to avoid unforced idling and if $\mu_1 h_1 \le \mu_2 h_2$, it is optimal to prioritize class 2 customers except to avoid unforced idling. One may thus conjecture that when $\tau > 0$, $\mu_1 h_1 \ge \mu_2 h_2$ ($\mu_1 h_1 \le \mu_2 h_2$) also implies it is optimal to prioritize class 1 (2) customers except to avoid unforced idling. The following examples illustrate that this is not always the case.

Example 4.1: Suppose we have the following inputs for the model: $\lambda = 0.01$; $\mu_1 = 1$; $\mu_2 = 5$; $h_1 = 10$; $h_2 = 1.9$; and $\tau = 0.006$. Figure 1 (left panel) displays the average cost optimal scheduling policy when the number of class 1 and class 2 customers in the system is bounded above by 100. For this example, $\mu_1 h_1 > \mu_2 h_2$ but there exist states (i, j) where it is optimal to serve a class 2 customers so that the classic cμ rule does not hold.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

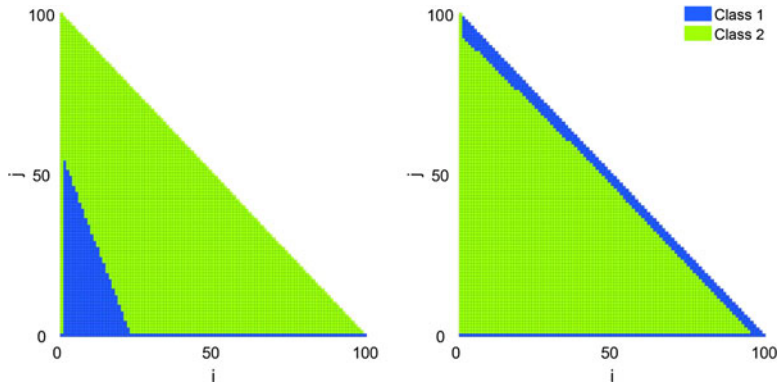


FIGURE **1.** Optimal policy for truncated systems described in Example 4.1 (left) and Example 4.2 (right). Each optimal policy does not satisfy the classic cμ rule, but does has a threshold structure. *Note:* Each box represents a particular state (i, j) of the system $(i, j$ with $i \leq 100$ and $j \leq 100)$ and the colors represent the policy at that state and time (blue = serve class 1; green = serve class 2). States are organized such that the number of class 1 customers (i) varies along the horizontal axis, the number of class 2 customers (j) varies along the vertical axis.

Example 4.2: Suppose we have the following inputs for the model: $\lambda = 0.01$; $\mu_1 = 2$; $\mu_2 = 1$; $h_1 = 1$; $h_2 = 2.1$; and $\tau = 0.001$. Figure 1 (right panel) displays the average cost optimal scheduling policy when the number of class 1 and class 2 customers in the system is bounded above by 100. For this example, $\mu_2 h_2 > \mu_1 h_1$ but there exist states (i, j) where it is optimal to serve a class 1 customers so that the class cμ rules does not hold.

In both examples, we see that the optimality of the cμ rule fails in states where the congestion level is "high". That is, it is optimal to serve class 2 customers even though $\mu_1 h_1 > \mu_2 h_2$ in Example 4.1 when there are a sufficient number of class 1 or class 2 customers. Likewise, it is optimal to serve class 1 even though $\mu_2 h_2 > \mu_1 h_1$ in Example 4.2. In particular, observe that in both examples the optimal policy when there are a sufficient number of class 1 or class 2 customers is to serve whichever customer class has the faster service rate, which are class 2 customers in Example 4.2 and class 1 customers in Example 4.1. Put differently, when the congestion is high, the optimal policy reduces congestion as quickly as possible.

The previous examples illustrate that customer's change in status can make the classic cμ rule fail. A natural question then becomes, are there conditions under which simple priority rules hold? To answer this question, consider the following result:

THEOREM 4.3: *The following hold:*

1. *Under both the α-discounted cost and average cost criteria, if $\mu_1 \geq \mu_2$ and $\mu_1 h_1 \geq \mu_2 h_2$, then it is optimal to serve class* 1 *customers except to avoid unforced idling.*
2. *Under both the α-discounted cost and average cost criteria, if $\mu_2 \geq \mu_1$ and $\mu_2 h_2 \geq \mu_1 h_1$, then it is optimal to serve class* 2 *customers except to avoid unforced idling.*

PROOF: The proof of the first statement uses a sample-path argument and is given in Appendix A. The proof of the second statement is similar to that of the first statement and is therefore omitted. ■

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

If $\mu_1 h_1 \geq \mu_2 h_2$ but $\mu_2 > \mu_1$, it must be the case that $h_1 > h_2$. This means that class 2 customers have the faster service rate *and* lower holding cost rate per customer. In this case, it makes sense to allow customers to change status (from class 1 to class 2) by serving class 2 customers before class 1 customers (c.f. Example 4.1). As a result, the condition $\mu_1 \geq \mu_2$ is required to guarantee that prioritizing class 1 customers is optimal when $\mu_1 h_1 \geq \mu_2 h_2$. Likewise, if $\mu_2 h_2 \geq \mu_1 h_1$ but $\mu_1 > \mu_2$, then $h_2 > h_1$. In this case, class 1 customers have the faster service rate *and* lower holding cost rate per customer. As a result, the optimal policy has an incentive to prevent customers from changing status by prioritizing class 1 customers (c.f. Example 4.2). As a result, the condition $\mu_2 \geq \mu_1$ is required to guarantee that prioritizing class 2 customers is optimal when $\mu_2 h_2 \geq \mu_1 h_1$.

4.2. Monotone Threshold Policies

In the previous section, we have characterized the optimal policy both when $\mu_1 \geq \mu_2$ and $\mu_1 h_1 \geq \mu_2 h_2$ and when $\mu_2 \geq \mu_1$ and $\mu_2 h_2 \geq \mu_1 h_1$. Our next step is to investigate what happens in remaining scenarios, specifically, when

$$\mu_2 > \mu_1 \text{ and } \mu_1 h_1 > \mu_2 h_2$$

and when

$$\mu_1 > \mu_2 \text{ and } \mu_2 h_2 > \mu_1 h_1.$$

For both scenarios, we need the following structural properties of the value function for the discounted cost case and of the relative value function for the long-run average cost case.

PROPOSITION 4.4: *$v_\alpha(\cdot, j)$ is convex in i, $v_\alpha(i, \cdot)$ is convex in j, and $v_\alpha(\cdot, \cdot)$ is supermodular. That is, the following inequalities hold for all $i, j \geq 0$:*

1. $v_\alpha(i+2, j) - 2v_\alpha(i+1, j) + v_\alpha(i, j) \geq 0$;
2. $v_\alpha(i, j+2) - 2v_\alpha(i, j+1) + v_\alpha(i, j) \geq 0$;
3. $v_\alpha(i+1, j+1) + v_\alpha(i, j) - v_\alpha(i+1, j) - v_\alpha(i, j+1) \geq 0$.

Similarly, if (g, w) is a solution to the average reward optimality equations, the above statements hold with v_α replaced with w.

A proof of Proposition 4.4 with successive approximations or sample-path arguments was not apparent to us because the expected sojourn times are unequal at different states. So rather than using these techniques, we slightly adapt methods proposed by Bhulai et al. in [6] and Blok and Spieksma in [7]. The idea is to restrict the number of class 1 customers allowed in the system at any time to be a finite number N and show that Proposition 4.4 holds for an appropriate uniformized model indexed by N, and then let $N \to \infty$ to obtain the result for the original unbounded model. The uniformized model for each N is obtained by decreasing the arrival rate linearly in i from λ when $i = 0$ to 0 when $i = N$. The rate of status change for class 1 customers is then bounded above by $N\tau$ and the MDP is uniformizable. This approach and the proof of these structural results are detailed in Appendix C. We then justify that this approach can be applied to our discounted and average cost model in Appendix D by using recent results from Prieto-Rumeau and Hernández-Lerma [30]. These conditions are Assumptions D.21–D.25 in the Appendix (c.f. Assumptions 4.5 and 4.13, and Theorem 4.18 by Prieto-Rumeau and Hernández-Lerma [30]).

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

4.2.1. $\boldsymbol{\mu_2 > \mu_1}$ *and* $\boldsymbol{\mu_1 h_1 > \mu_2 h_2}$. Let's assume that $\mu_2 > \mu_1$ and $\mu_1 h_1 > \mu_2 h_2$. Taken together, these conditions imply that $\mu_2 > \mu_1$ and $h_1 > h_2$. An example where these conditions may hold was described in the introduction: a cohort of patients whose condition improves over time, at which point they are quicker and less costly to serve. Hence, one might conjecture that if it is optimal to serve a class 1 customer in state (i, j), then it is also optimal to serve a class 1 customer in state $(i+1, j)$ or in state $(i, j+1)$. Example 4.1 in the previous section shows that this is not the case. In particular, as we noted in the previous section, when the congestion level is "high" (i.e., when either i or j is sufficiently high), the optimal policy is to serve class 2 customers instead of class 1 customers.

It turns out that the optimal policy has a surprising structure: if it is optimal to serve class 2 customers in state (i, j), then it is also optimal to serve class 2 customers in state $(i+1, j)$ and also in state $(i, j+1)$. This is indeed the case, as we now show.

THEOREM 4.5: *Suppose* $\mu_2 > \mu_1$ *and* $\mu_1 h_1 > \mu_2 h_2$. *In this case, the following hold:*

1. *Under both the* α*-discounted cost and average cost criteria, if it is optimal to serve a class 2 customer in state* (i, j), *then it is also optimal to serve a class 2 customer in state* $(i, j+1)$.
2. *Under both the* α*-discounted cost and average cost criteria, if it is optimal to serve a class 2 customer in state* (i, j), *then it is also optimal to serve a class 2 customer in state* $(i+1, j)$

For the discounted expected cost problem, observe that for $i, j \geq 1$, the inequality

$$\begin{aligned} &\mu_2[v_\alpha(i,j) - v_\alpha(i,j+1)] - \mu_1[v_\alpha(i-1,j+1) - v_\alpha(i,j+1)] \\ &\quad - \mu_2[v_\alpha(i,j-1) - v_\alpha(i,j)] + \mu_1[v_\alpha(i-1,j) - v_\alpha(i,j)] \leq 0 \end{aligned} \tag{1}$$

implies that if it is optimal to serve a class 2 customer in state (i, j), then it is optimal to serve a class 2 customer in state $(i, j+1)$. Note that the inequalities for the average cost case are analogous. Similarly, for $i, j \geq 1$,

$$\begin{aligned} &\mu_2[v_\alpha(i+1,j-1) - v_\alpha(i+1,j)] - \mu_1[v_\alpha(i,j) - v_\alpha(i+1,j)] \\ &\quad - \mu_2[v_\alpha(i,j-1) - v_\alpha(i,j)] + \mu_1[v_\alpha(i-1,j) - v_\alpha(i,j)] \leq 0 \end{aligned} \tag{2}$$

implies that for the discounted expected cost problem, if it is optimal to serve a class 2 customer in state (i, j), then it is optimal to serve a class 2 customer in state $(i+1, j)$. To prove Proposition 4.5, we show that the two inequalities 1 and 2 hold. We remark that the proofs are given for the discounted cost optimality equations since the proof for the average cost optimality equations are similar.

We begin with the following lemma, which says that it is better to have an extra class 2 customer in the system than an extra class 1 customer in the system. The proof uses a sample-path argument and is delegated to Appendix E.

LEMMA 4.6: *Suppose* $\mu_2 > \mu_1$ *and* $\mu_1 h_1 > \mu_2 h_2$. *In this case, the inequality*

$$v_\alpha(i, j-1) - v_\alpha(i-1, j) \geq 0 \tag{3}$$

holds for all $i, j \geq 1$. *Similarly, in the average reward case with* v_α *replaced with* w *from the ACOE.*

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

Proof of Inequality 1. The proof is by induction and uses the same approach that was used to prove Proposition 4.4 in Appendix C. As a result, we delegate the proof to Appendix F.

Proof of Inequality 2. Observe that Propositions 4.4 together with Lemma 4.6 imply that the terms with coefficients μ_1 in inequality 2 are bounded above by

$$\mu_1[v_\alpha(i+1,j) - v_\alpha(i,j) - v_\alpha(i-1,j+1) + v_\alpha(i-1,j)].$$

The same lemma implies that this last expression is bounded above by

$$\mu_1[v_\alpha(i+1,j) - v_\alpha(i,j) - v_\alpha(i-1,j+1) + v_\alpha(i,j-1)],$$

which is bounded above by

$$\mu_2[v_\alpha(i+1,j) - v_\alpha(i,j) - v_\alpha(i-1,j+1) + v_\alpha(i,j-1)]$$

as a consequence of $\mu_2 > \mu_1$. This completes the proof.

4.2.2. $\boldsymbol{\mu_1 > \mu_2}$ *and* $\boldsymbol{\mu_2 h_2 > \mu_1 h_1}$. We now consider the case $\mu_1 > \mu_2$ and $\mu_2 h_2 > \mu_1 h_1$. Taken together, these conditions imply that $\mu_1 > \mu_2$ and $h_2 > h_1$. These conditions model the examples from the introduction in which a patient/job deteriorates over time at which point they are slower and more costly to service. Hence, one might conjecture that the optimal policy has the following switching curve structure: if it is optimal to serve class 2 customers in state (i,j), then it is also optimal to serve class 2 customers in state $(i,j+1)$ or in state $(i+1,j)$. Example 4.2 shows that this is not the case.

Example 4.2 suggests that the optimal policy has the following switching curve structure: if it is optimal to have the server work on class 1 customers in state (i,j), then it is also optimal to have the server work on class 1 customers in state $(i+1,j)$ and in state $(i,j+1)$. Let's suppose that we could proceed analogously to the previous scenario (i.e., when $\mu_2 > \mu_1$ and $\mu_1 h_1 > \mu_2 h_2$). Then, we would prove that the optimal policy has this structure if we could show the following two inequalities:

$$\begin{aligned}&\mu_1[v_\alpha(i,j) - v_\alpha(i+1,j)] + \mu_2[v_\alpha(i+1,j) - v_\alpha(i+1,j-1)]\\ &\quad - \mu_1[v_\alpha(i-1,j) - v_\alpha(i,j)] - \mu_2[v_\alpha(i,j) - v_\alpha(i,j-1)] \le 0 \end{aligned} \tag{4}$$

and

$$\begin{aligned}&\mu_1[v_\alpha(i-1,j+1) - v_\alpha(i,j+1)] + \mu_2[v_\alpha(i,j+1) - v_\alpha(i,j)]\\ &\quad - \mu_1[v_\alpha(i-1,j) - v_\alpha(i,j)] - \mu_2[v_\alpha(i,j) - v_\alpha(i,j-1)] \le 0. \end{aligned} \tag{5}$$

We would need similar inequalities in the average reward case with v_α replaced with w from the ACOE. However, the proofs of these inequalities cannot proceed in the same way due to customer deteriorations. A sample path or induction argument for the first inequality (4) leads to the second inequality (5) after a deterioration. This implies that to show the first inequality, we must show the second. The second inequality (2), however, does *not* hold for all $i,j \ge 1$ in clearing systems or truncated versions of our model, as shown in Example 4.2. Thus, the usual methods for proving the structure of an optimal policy cannot be used. One possibility is to use interchange arguments (i.e., interchange the order of services). This is also not straightforward since customers may change status between services. Despite this,

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

all our numerical examples (and our own intuition) suggest that the structural result holds. Therefore, we leave the structural properties as a conjecture.

Conjecture 4.7: Suppose $\mu_1 > \mu_2$ and that $\mu_2 h_2 > \mu_1 h_1$. In this case, the following hold:

1. If it is optimal to serve a class 1 customer in state (i, j), then it is also optimal to serve a class 1 customer in state $(i + 1, j)$.
2. If it is optimal to serve a class 1 customer in state (i, j), then it is also optimal to serve a class 1 customer in state $(i, j + 1)$

One result that may help establish the aforementioned conjecture is the following lemma, which is the direct analog of Lemma 4.6 given in the previous subsection above, and says that it is better to have an extra class 1 customer in the system than an extra class 2 customer. The proof is similar to the proof of Lemma 4.6 and is omitted.

LEMMA 4.8: *Suppose* $\mu_1 > \mu_2$ *and* $\mu_2 h_2 > \mu_1 h_1$. *In this case, the inequality*

$$v_\alpha(i-1, j) - v_\alpha(i, j-1) \geq 0 \tag{6}$$

holds for all $i, j \geq 1$. *Similarly, in the average reward case with* v_α *replaced with* w *from the ACOE.*

5. OPTIMAL CONTROL FOR N JOB CLASSES ($n \geq 3$)

In Proposition 4.3, we showed that if $\mu_1 \geq \mu_2$ and $\mu_1 h_1 \geq \mu_2 h_2$, then it is optimal to prioritize class 1 customers except to avoid unforced idling. A similar argument can be used to show that, for $n \geq 3$, if $\mu_1 \geq \cdots \geq \mu_n$ and $\mu_1 h_1 \geq \cdots \geq \mu_n h_n$, then it is optimal to prioritize customer class $n - 1$ over customer class n. The reason is that the proof of the first statement in Proposition 4.3 only requires that inter-arrival times are exponential (and not state independent), which is satisfied for the deterioration of class $n - 2$ customers, i.e., the "arrival" process of class $n - 1$ customers. Hence, the system involving only the last two customers is identical to a two-class system to which we can apply the proposition and get the stated result. One might thus conjecture that this result would extend to $n \geq 3$; i.e., if $\mu_1 \geq \cdots \geq \mu_n$ and $\mu_1 h_1 \geq \cdots \geq \mu_n h_n$, then it is optimal to prioritize the most *upstream* customer class except to avoid unforced idling. This turns out not to be the case in general, as the following example (when $n = 3$) illustrates.

Example 5.1: Suppose we have the following inputs for the model: $\lambda = 0.01$; $\mu_1 = 40$; $\mu_2 = 10$; $\mu_3 = 1$; $h_1 = 1$; $h_2 = 2$; $h_3 = 19$; $\tau_1 = 0.1$; and $\tau_2 = 0.1$. Figure 2 displays the average cost optimal scheduling policy when the total number of customers in the system is bounded above by 15. There are states where customer class 3 is prioritized when $j = 0$.

Example 5.1 illustrates that when deciding whether to prioritize an upstream customer class (i.e., customer class 1), the decision maker must now take into account potential changes in status from the most adjacent, downstream customer class (i.e., class 2 customers). In particular, class 2 customers have a relatively high parameter value of $\tau_2 = 0.1$, so they rapidly change status to class 3 customers which have a higher cost and slower service rate. As a result, the server has an incentive to prevent further changes in status by prioritizing class 2 customers.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

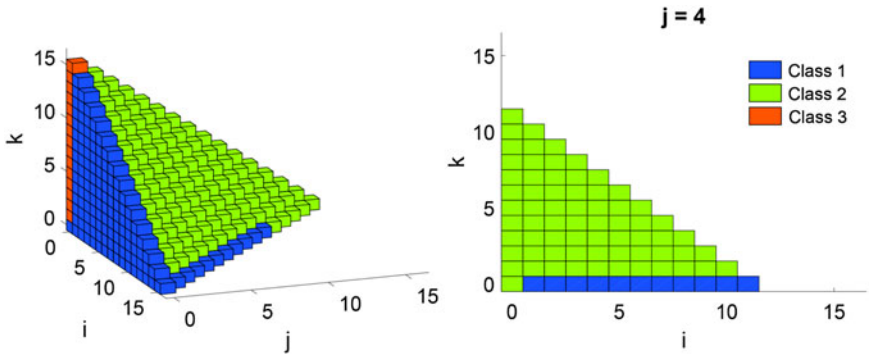


FIGURE **2.** Optimal policy for a truncated system. *Note:* Each box represents a particular state (i, j, k) of the system (i, j, k with $i + j + k \leq 15$) and the colors represent the policy at that state and time (blue = serve class 1; green = serve class 2; red = serve class 3). States are organized such that the number of class 1 customers (i) varies along the horizontal axis, the number of class 2 customers (j) varies along the vertical axis, the number of class 3 customers (k) varies along the vertical axis.

We also showed that for $n = 2$, when $\mu_2 \geq \mu_1$ and $\mu_2 h_2 \geq \mu_1 h_1$, it is optimal to prioritize class 2 customers. However, a similar condition $\mu_n \geq \cdots \geq \mu_1$ and $\mu_n h_n \geq \cdots \geq \mu_1 h_1$ does not imply that it is optimal to prioritize the most *downstream* customer class except to avoid unforced idling (reversing the roles of μ_1 and μ_3 in Example 5.1 provides an example).

Taken together, these examples suggest that when $n \geq 3$, there are no simple orderings of the parameters that guarantee that simple priority rules are optimal. In particular, note from Example 5.1 that Theorem 4.3 may not extend from $n = 2$ to $n = 3$ even in the simplified setting when $\tau_1 = \cdots = \tau_{n-1}$.

However, we can extend our earlier result when all service rates are the same.

PROPOSITION 5.2: *Suppose $\mu_1 = \mu_2 = \cdots = \mu_n$. In this case, under both the α-discounted cost and average cost criteria, $h_1 \geq h_2 \geq \cdots \geq h_n$ ($h_1 \leq h_2 \leq \cdots \leq h_n$) implies that it is optimal to prioritize the most upstream (downstream) customer class except to avoid unforced idling.*

PROOF: See Appendix G. ■

Based on Proposition 5.2 one may think that if $\mu_1 = \mu_2 = \cdots = \mu_n$, then it is always optimal to prioritize the customer class with the highest cost per unit time per customer except to avoid unforced idling. We end this section with the following example (c.f. Figure 2), that illustrates that this also is not the case.

Example 5.3: Suppose we have the following inputs for the model: $\lambda = 0.01$; $\mu_1 = 10$; $\mu_2 = 10$; $\mu_3 = 10$; $h_1 = 19.5$; $h_2 = 20$; $h_3 = 19$; $\tau_1 = 1$; and $\tau_2 = 2$. Figure 3 displays the average cost optimal scheduling policy when the total number of customers in the system is bounded above by 15. Customer class 3 is prioritized over all states except to avoid unforced idling.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

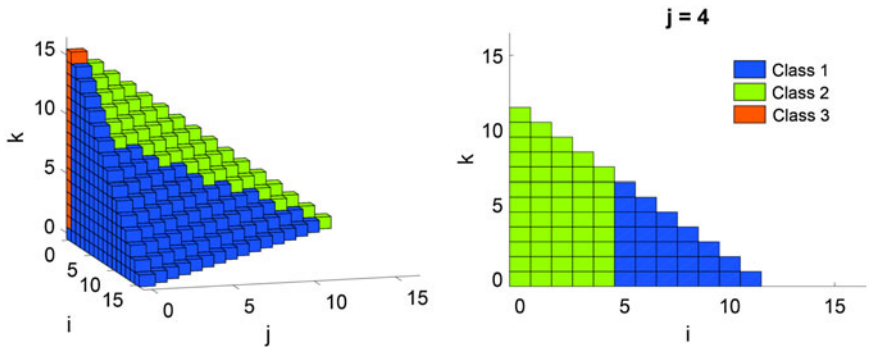


FIGURE **3.** Optimal policy for a clearing system. The optimal policy has a threshold structure. *Note:* Each box represents a particular state (i, j, k) of the system $(i, j, k$ with $i + j + k \leq 15)$ and the colors represent the policy at that state and time (blue = serve class 1; green = serve class 2; red = serve class 3). States are organized such that the number of class 1 customers (i) varies along the horizontal axis, the number of class 2 customers (j) varies along the vertical axis, the number of class 3 customers (k) varies along the vertical axis.

6. NUMERICAL STUDY

While we were unable to characterize the optimal policy for all cases, our numerical study suggests that a switching curve or a threshold-based policy can perform well. For instance, when $\mu_1 > \mu_2$ and $\mu_2 h_2 > \mu_1 h_1$, we observe from our numerical examples that a threshold or switching curve structure result holds: if it is optimal to serve a class 1 customer in state (i, j), then it is also optimal to serve a class 1 customer in states $(i + 1, j)$ and $(i, j + 1)$. In words, the optimal policy eventually prioritizes the customer class with the highest service rate. This section develops and analyzes several heuristics based on this insight.

6.1. Description of Numerical Study

For our numerical study, we truncate the state space so that the resulting state space is finite, we calculate optimal policies for this truncated model using modified policy iteration (c.f. Puterman and Shin [33], Puterman [32]), and compare the average costs of the optimal policies with those of our proposed heuristics. We consider two scenarios:

1. $\mu_n h_n \geq \cdots \geq \mu_1 h_1$;
2. $\mu_1 h_1 \geq \cdots \geq \mu_n h_n$;

to capture the examples from the introduction (e.g., when a patient/job deteriorates over time and when a patient/job improves over time). For each setting, we consider several different examples by changing the arrival rate and the rate at which jobs change status.

For instance, when $n = 2$, we fix $\mu_1 = 4, h_1 = 100$, $\mu_2 = 3, h_2 = 150$, and then reverse μ_1 and μ_2 and of h_1 and h_2 in the second scenario. When $n = 3$, we fix $\mu_1 = 4, h_1 = 100$, $\mu_2 = 3, h_2 = 150$, $\mu_3 = 2, h_3 = 300$ in the first scenario, and then reverse μ_1 and μ_3 and of h_1 and h_3 in the second scenario. In particular, values of λ are chosen so that $\rho := \lambda / \min_i\{\mu_i\} = 0.5, 0.7, 0.9$. Similarly, in the case when $n = 3$, we consider about 70 problem instances with different parameters.

For our study, we truncate the state space to be a rectangular state space. Namely, customers are blocked from arriving or changing status, if the number of customers in a

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

particular class equals the truncation value. A truncation value of 100 is used for $n = 2$ and 20 is used for $n = 3$.

Four types of heuristic policies are compared

h	Prioritize classes in decreasing order of h_i.
Cμ	Prioritize classes in decreasing order of $h_i\mu_i$.
μ	Prioritize classes in decreasing order of μ_i.
Mod-C$\mu(N)$	Prioritize classes in decreasing order of $h_i\mu_i$ if the total number of customers in the system is strictly less than N; otherwise prioritize classes in decreasing order of μ_i.

Heuristic Cμ is simply the well-known cμ rule whereas μ prioritizes the customer class with the fastest service rate. Heuristic Mod-C$\mu(N)$ combines the two: the cμ rule is followed by the total number of customers in the system reaches a fixed and finite level N, after which the server prioritizes the customer class with the fastest service rate. Note policies Mod-C$\mu(1)$ and μ are identical, and Mod-C$\mu(\infty)$ and Cμ are as well. We examine Mod-C$\mu(N)$ for $N = 5, 10, 15$.

Note that in some cases, two priority policies (e.g., h and Cμ or μ and Cμ) can be identical, which we report in the corresponding tables. Additional cases are also considered in which the h priority rule yields different orderings from the μ and Cμ rules. In all cases, however, we observe similar qualitative results regarding the performance of the Mod-C$\mu(N)$ rule proposed in this paper. In this paper, we present the two specific cases (24 problem instances). The results of remaining cases are reported in the Appendix (Tables H.1–H.4 in Appendix H).

6.2. Two Customer Classes ($n = 2$)

We find that heuristics Mod-C$\mu(N)$ for $N = 5$, 10, 15 perform well. They are the best out of the considered heuristics in over half of the instances considered. For those instances in which they are not the best (i.e., when either one of the two static priority rules $C\mu$ or μ perform the best) they are, at worst, no more than 3% away from the optimal policy, and in general, no more than 6.5% away from the optimal. Our results also suggest that it may beneficial to prioritize customer classes in decreasing order of $h_i\mu_i$ by following Cμ when the total number of customers is low and when customers arrive slowly and change status slowly, but otherwise, it is better to stay "close" to heuristic μ.

Consider the first scenario where customers become slower and more expensive to serve (Table 1). When customers arrive and change status slowly (i.e., $\lambda = 1, 1.4$, $\tau_1 = 0.2$), heuristic Mod-C$\mu(5)$ performs the best out of the considered heuristics. When customers arrive or change status more quickly (i.e., $\lambda = 1.8$ or $\tau_1 = 0.4$), heuristic μ performs the best. Also note that when customers arrive or change status quickly, heuristic Cμ performs the worst.

In the second scenario, customers become faster and less expensive to serve (Table 2). In this scenario, heuristics Cμ, Mod-C$\mu(15)$, and Mod-C$\mu(10)$ are close to optimal when $\lambda = 1$ and $\tau_1 = 0.2$, whereas heuristic μ performs the worst when $\lambda = 1, 1.4$ and $\tau_1 = 0.2$. When customers arrive slowly but change status more quickly (i.e., $\lambda = 1$, $\tau_1 = 0.4$), heuristic Mod-C$\mu(5)$ again performs the best, but heuristic μ becomes the best as the load increases. Lastly, heuristic Cμ performs the worst when customers can change status more quickly (i.e., $\tau_1 = 0.4$).

In general, increasing the load (i.e., λ) increases the cost of running the system, since a higher load leads to more customers in the system per unit time. Increasing the rate at which jobs change status τ_1 also makes the system more expensive in the first scenario (Table 1), where more customers change status on average to the slower, higher cost queue.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

TABLE 1. Long-run average holding costs when $n = 2$ and when customers become slower and more expensive to serve. Here, $\mu_1 = 4, h_1 = 100$, $\mu_2 = 3, h_2 = 150$.

τ_1	λ	Optimal	Cμ and h	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	μ
0.2	1	35.92	35.93	35.93	35.93	35.92	36.00
	1.4	59.26	59.39	59.39	59.38	59.29	59.34
	1.8	92.76	93.97	93.93	93.74	93.04	92.79
0.4	1	38.21	38.46	38.46	38.46	38.38	38.21
	1.4	63.65	65.10	65.10	65.05	64.43	63.65
	1.8	100.81	107.69	107.54	106.68	103.24	100.81

TABLE 2. Long-run average holding costs when $n = 2$ and when customers become faster and less expensive to serve. Here, $\mu_1 = 4, h_1 = 150$, $\mu_2 = 4, h_2 = 100$.

τ_1	λ	Optimal	Cμ and h	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	μ
0.2	1	70.02	70.02	70.02	70.02	70.04	70.24
	1.4	118.82	118.90	118.90	118.88	118.85	119.15
	1.8	192.40	193.92	193.83	193.34	192.40	192.61
0.4	1	66.45	66.53	66.53	66.52	66.47	66.54
	1.4	110.64	111.47	111.47	111.42	110.86	110.69
	1.8	173.62	178.74	178.66	177.93	174.59	173.62

Conversely, increasing τ_1 makes the system less expensive in the second scenario (Table 2), where more customers change status on average to the faster, lower cost queue.

6.3. Three Customer Classes ($n = 3$)

Our analysis of the optimal policy in Section 5 suggests that the optimal policy can be extremely complicated when $n \geq 3$. Our numerical results further confirm this. We find that the static priority rules Cμ and μ perform poorly in certain instances. However, similar to when $n = 2$, we find that heuristics Mod-C$\mu(N)$ for $N = 5$, 10, 15 perform very well. They perform the best out of the considered heuristics in 19 out of the 24 instances reported in Tables 3 and 4, and are, at worst, no more than 6% away from the optimal policy. The tables in the Appendix show similar results in that the performance of the Mod-C$\mu(N)$ rule is quite close to that of the optimal policy.

Now we explore the examples in detail. Consider the first scenario where customers may become slower and more expensive to serve (Table 3). When class 1 and class 2 customers change status slowly (i.e., $\tau_1 = \tau_2 = 0.2$), Mod-Cμ(15) and Mod-Cμ(10) perform the best while μ performs the worst. When customers arrive slowly and class 1 customers change status more quickly (i.e., $\lambda = 1, 1.4$, $\tau_1 = 0.4$, $\tau_2 = 0.2$), Mod-Cμ(5) performs the best. When class 2 customers change status more quickly but class 1 customers change status slowly (i.e., $\tau_1 = 0.2$, $\tau_2 = 0.4$), μ performs the best when the load is low, but Cμ performs the best otherwise. When $\lambda = 1$, $\lambda = 1.4$, and $\lambda = 1.8$, heuristics μ, Mod-Cμ(15), and Cμ, perform the best, respectively. When customers arrive slowly and class 1 customers change status more quickly (i.e., $\lambda = 1$, $\tau_1 = \tau_2 = 0.4$), μ performs the best, but when the load is increased, Cμ performs the best.

Consider the second scenario, where customers become faster and less expensive to serve (Table 4). When class 1 and class 2 customers change status slowly (i.e., $\tau_1 = \tau_2 = 0.2$), heuristics Cμ (tied with Mod-Cμ(15)), Mod-Cμ(15), and Mod-Cμ(10) perform the best

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

TABLE 3. Long-run average holding costs when $n = 3$ and when customers become slower and more expensive to serve. Here, $\mu_1 = 4, h_1 = 100$, $\mu_2 = 3, h_2 = 150$, and $\mu_3 = 2, h_3 = 300$.

τ_2	τ_1	λ	Optimal	Cμ and h	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	μ
0.2	0.2	1	36.61	36.62	36.62	36.62	36.67	37.21
		1.4	60.96	61.02	61.01	61.01	61.31	62.63
		1.8	97.85	99.91	98.57	98.36	99.13	101.53
	0.4	1	39.98	40.01	40.01	40.01	39.99	40.43
		1.4	68.84	69.69	69.58	69.34	68.87	69.50
		1.8	115.47	143.30	125.48	120.93	116.19	115.83
0.4	0.2	1	37.21	38.15	38.15	38.15	38.04	37.25
		1.4	62.31	64.94	64.94	64.97	64.49	62.68
		1.8	100.96	106.96	107.18	109.92	113.50	110.70
	0.4	1	41.30	42.14	42.14	42.14	42.20	41.55
		1.4	71.98	73.61	73.61	73.74	75.75	75.36
		1.8	123.32	125.20	125.60	136.55	204.06	219.47

TABLE 4. Long-run average holding costs when $n = 3$ and when customers become faster and less expensive to serve. Here, $\mu_1 = 2, h_1 = 300$, $\mu_2 = 3, h_2 = 150$, and $\mu_3 = 4, h_3 = 100$.

τ_2	τ_1	λ	Optimal	Cμ and h	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	μ
0.2	0.2	1	239.80	239.81	239.81	239.85	242.15	246.05
		1.4	453.58	455.63	455.04	454.78	463.33	469.75
		1.8	848.07	929.90	875.19	860.29	871.57	877.99
	0.4	1	211.91	212.00	211.99	211.95	212.46	215.62
		1.4	381.40	386.29	385.73	383.75	381.76	385.54
		1.8	651.91	736.17	701.62	675.44	652.40	653.29
0.4	0.2	1	237.53	237.61	237.60	237.64	240.30	245.03
		1.4	444.25	447.33	446.69	446.17	457.30	465.98
		1.8	809.05	900.24	845.29	830.18	853.36	863.99
	0.4	1	208.87	209.12	209.11	209.06	209.88	214.16
		1.4	370.33	376.57	376.00	373.90	374.12	380.94
		1.8	623.64	704.53	671.34	645.19	632.44	639.31

when $\lambda = 1$, $\lambda = 1.4$, and $\lambda = 1.8$, respectively. When customers arrive slowly but class 1 customers change status more quickly (i.e., $\lambda = 1, \tau_1 = 0.4, \tau_2 = 0.2$) heuristic Mod-C$\mu$(10) performs the best, but as the load is increased to $\lambda = 1.4, 1.8$, it is Mod-Cμ(5) that performs the best. When class 2 customers change status more quickly (i.e., $\tau_2 = 0.4$), at least one of the Mod-C$\mu(N)$ for $N = 15, 10, 5$ performs the best, depending on the parameter values.

7. CONCLUSION

This paper is concerned with an important, yet less studied scheduling dilemma: how to allocate resources to customers (or jobs) whose initial service or cost requirements can change when they wait for too long. Our work is motivated by applications in service systems, such as those in health care and maintenance shops, where long waits lead to either improved, or deteriorating health or (machine) functioning. To our knowledge, we

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

are the first to consider this scenario using a single-server Markovian system with random changes in status under a holding cost structure.

We use a continuous-time MDP formulation to determine qualitative properties of optimal policies with surprising consequences. For the two-class model, the classic cμ rule no longer holds when jobs are allowed to change status and so, we provide additional conditions for when simple priority rules are optimal. Furthermore, numerical examples illustrate that seemingly intuitive switching curve results also do not hold in general. Perhaps more surprising is that conditions for when priority rules in the two-customer-class model do not extend to the general n-class model (with $n \geq 3$), even with a simplified model (e.g., one with equal "deterioration" rates). We study these scenarios and provide conditions under which some switching curve results hold and under which simple priority rules hold. We remark that our numerical work points to trade-offs between reducing costs myopically via the cμ rule versus reducing congestion as quickly as possible as the reason why intuitive structural properties no longer hold when we allow jobs to change status.

In addition to surprising qualitative results, we must contend with several technical challenges as a result of customers or jobs changing status. For one, we can no longer apply standard adjacent pairwise interchange arguments to determine when priority rules hold. They also lead to unbounded transitions rates, so that the model cannot be analyzed by first uniformizing the process.

There are several avenues for further research. Further characterizing the optimal policy in general when priority rules do not hold is of clear interest. One result in this direction would be a proof of Conjecture 4.7. Another would be determining additional conditions for when priority rules hold (or showing they never hold) when the conditions in Propositions 4.3 and 5.2 do not hold. As observed in this paper, allowing customers (or jobs) to change status make the classic cμ rule and other seemingly structural properties of optimal controls fail, and more broadly, makes the analysis and computation of optimal controls more complicated. Because changing status when customers or jobs wait for too long are well-observed in practice in a broad range of problems, they provide significant technical challenges for optimal control and are a bright research direction.

References

1. Albright, S.C. (1980). Optimal maintenance-repair policies for the machine repair problem. *Naval Research Logistics Quarterly* 27(1): 17–27.
2. Alidaee, B. & Womer, N.K. (1999). Scheduling with time dependent processing times: review and extensions. *Journal of the Operational Research Society* 50(7): 711–720.
3. Argon, N.T., Ziya, S., & Righter, R. (2008). Scheduling impatient jobs in a clearing system with insights on patient triage in mass casualty incidents. *Probability in the Engineering and Informational Sciences* 22(3): 301–332.
4. Armstrong, M.J. (2002). Age repair policies for the machine repair problem. *European Journal of Operational Research* 138(1): 127–141.
5. Bachman, A., Janiak, A., & Kovalyov, M.Y. (2002). Minimizing the total weighted completion time of deteriorating jobs. *Information Processing Letters* 81(2): 81–84.
6. Bhulai, S., Brooms, A.C., & Spieksma, F.M. (2014). On structural properties of the value function for an unbounded jump Markov process with an application to a processor sharing retrial queue. *Queueing Systems* 76(4): 425–446.
7. Blok, H. & Spieksma, F.M. (2015). Countable state Markov decision processes with unbounded jump rates and discounted cost: optimality equation and approximations. *Advances in Applied Probability* 47(4): 1088–1107.
8. Boxma, O.J. & Vlasiou, M. (2007). On queues with service and interarrival times depending on waiting times. *Queueing Systems* 56(3-4): 121–132.
9. Brown, M. & Ross, S.M. (1973). Optimal issuing policies. *Management Science* 19(11): 1292–1294.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

10. Brown, M. & Solomon, H. (1973). Optimal issuing policies under stochastic field lives. *Journal of Applied Probability* 10(4): 761–768.
11. Browne, S. & Yechiali, U. (1990). Scheduling deteriorating jobs on a single processor. *Operations Research* 38(3): 495–498.
12. Buyukkoc, C., Varaiya, P., & Walrand, J. (1985). The $c\mu$-rule revisited. *Advances in Applied Probability* 17(1): 237–238.
13. Cai, X., Wu, X., & Zhou, X. (2014). *Optimal stochastic scheduling*. Springer.
14. Chan, C.W., Farias, V.F., & Escobar, G.J. (2016). The impact of delays on service times in the intensive care unit. *Management Science*. doi: 10.1287/mnsc.2016.2441.
15. Cheng, T.C.E., Ding, Q., & Lin, B.M.T. (2004). A concise survey of scheduling with time-dependent processing times. *European Journal of Operational Research* 152(1): 1–13.
16. Down, D.G., Koole, G., & Lewis, M.E. (2011). Dynamic control of a single-server system with abandonments. *Queueing Systems* 67(1): 63–90.
17. Guo, X. & Hernández-Lerma, O. (2003). Continuous-time controlled Markov chains with discounted rewards. *Acta Applicandae Mathematica* 79(3): 195–216.
18. Guo, X. & Hernández-Lerma, O. (2009). *Continuous-time Markov decision processes*. Berlin, Heidelberg: Springer.
19. Haque, L. & Armstrong, M.J. (2007). A survey of the machine interference problem. *European Journal of Operational Research* 179(2): 469–482.
20. Kawai, H. (1983). An optimal ordering and replacement policy of a Markovian degradation system under complete observation. *Journal of Operations Research Society of Japan* 26: 279–290.
21. Kawai, H. (1984). An optimal inspection and replacement policy of a Markovian deterioration system. In S. Osaki & Y. Hatoyama (eds.), *Stochastic models in reliability theory*. Lecture Notes in Economics and Mathematical Systems, vol 235. Berlin, Heidelberg: Springer, pp. 177–186.
22. Legros, B, Jouini, O., & Koole, G. (2014). A uniformization approach for the dynamic control of multi-server queueing systems with abandonments. Technical report, Working paper.
23. Leung, J.Y.T. (2004). *Handbook of scheduling: algorithms, models, and performance analysis*. New York, NY: Chapman and Hall/CRC Press.
24. Lippman, S.A. (1975). Applying a new device in the optimization of exponential queuing systems. *Operations Research* 23(4): 687–710.
25. Master, N., Chan, C.W., & Bambos, N. (2016). Myopic policies for non-preemptive scheduling of jobs with decaying value. *Probability in the Engineering and Informational Sciences* 1–36. doi: 10.1017/S0269964816000474.
26. Mosheiov, G. (1991). V-shaped policies for scheduling deteriorating jobs. *Operations Research* 39(6): 979–991.
27. Mosheiov, G. (1996). λ-shaped policies to schedule deteriorating jobs. *Journal of the Operational Research Society* 47(9): 1184–1191.
28. Pinedo, M.L. (2012). *Scheduling: theory, algorithms, and systems*. Boston, MA: Springer US.
29. Prieto-Rumeau, T. & Hernández-Lerma, O. (2012). Discounted continuous-time controlled Markov chains: convergence of control models. *Journal of Applied Probability* 49(4): 1072–1090.
30. Prieto-Rumeau, T. & Hernández-Lerma, O. (2012). *Selected topics on continuous-time controlled Markov chains and Markov games*, vol. 5. London, UK: Imperial College Press.
31. Prieto-Rumeau, T. & Lorenzo, J.M. (2010). Approximating ergodic average reward continuous-time controlled Markov chains. *IEEE Transactions on Automatic Control* 55(1):201–207.
32. Puterman, M.L. (1994). *Markov decision processes: discrete stochastic dynamic programming*. Wiley Series in Probability and Mathematical Statistics. New York: John Wiley and Sons, Inc.
33. Puterman, M.L. & Shin, M.C. (1978) Modified policy iteration algorithms for discounted markov decision problems. *Management Science* 24(11): 1127–1137.
34. Saghafian, S., Hopp, W.J., Van Oyen, M.P., Desmond, J.S., & Kronick, S.L. (2014). Complexity-augmented triage: a tool for improving patient safety and operational efficiency. *Manufacturing & Service Operations Management* 16(3): 329–345.
35. Srinivasan, R. & Parlikad, A.K. (2014). Semi-Markov decision process with partial information for maintenance decisions. *IEEE Transactions on Reliability* 63(4): 891–898.
36. Stengos, D. & Thomas, L.C. (1980). The blast furnaces problem. *European Journal of Operational Research* 4(5): 330–336.
37. Tijms, H.C. & van der Duyn Schouten, F.A. (1985). A Markov decision algorithm for optimal inspections and revisions in a maintenance system with partial information. *European Journal of Operational Research* 21(2): 245–253.
38. Whitt, W. (1990). Queues with service times and interarrival times depending linearly and randomly upon waiting times. *Queueing Systems* 6(1): 335–351.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

APPENDIX

Appendix A: Proof of Theorem 4.3

Proof of Statement 1 in Theorem 4.3. The proof is given for the discounted expected cost model. The proof of the long-run average cost case is similar. Note that the optimality equations imply that it is optimal to prioritize a class 1 customer in state (i, j) with $i, j \geq 1$ when

$$\mu_1[v_\alpha(i-1,j) - v_\alpha(i,j)] + \mu_2[v_\alpha(i,j) - v_\alpha(i,j-1)] \leq 0. \tag{A.1}$$

We thus show (A.1) using a sample path argument. Fix $i, j \geq 1$ and start four processes on the same probability space. Processes 1–4 begin in states $(i-1,j), (i,j), (i,j)$, and $(i,j-1)$, respectively. Processes 2 and 4 use stationary optimal policies, which we denote by π_2 and π_4, respectively. In what follows, we show how to construct (potentially sub-optimal) policies for Processes 1 and 3, which we denote by π_1 and π_3, so that

$$\mu_1[v_\alpha^{\pi_1}(i-1,j) - v_\alpha^{\pi_2}(i,j)] + \mu_2[v_\alpha^{\pi_3}(i,j) - v_\alpha^{\pi_4}(i,j-1)] \leq 0. \tag{A.2}$$

Since π_1 and π_3 are potentially sub-optimal, (A.1) follows from (A.2). In what follows, discounting is suppressed without any loss of generality.

Observe that starting from (A.2), the immediate costs incurred at the next event are $(\mu_2 h_2 - \mu_1 h_1)t_1 \leq 0$, where t_1 is the time of the next event and the inequality is due to the assumption that $\mu_2 h_2 \leq \mu_1 h_1$. Moreover, if the relative position (as measured by the current states) of the four processes at the next event remains the same, then we may relabel the initial states and continue from the beginning of the argument. This occurs when any of the uncontrolled events occur that are seen by all four processes. It also occurs when π_2 and π_4 both serve the same customer class $k \in \{1, 2\}$ by letting π_1 and π_3 also serve the same customer class k customer provided there is one or more customer class k customer in all four processes, and the next event service completion. Consider now the other cases.

Case A.1: Customer status changes.

Suppose the first event is a change in status in Processes 2–4 only (with probability $\tau/\lambda + \mu_1 + \mu_2 + i\tau$), after which all processes follow an optimal control. Suppressing the probability of this event in the expression, it follows that the remaining costs on the left-hand side of the inequality in (A.2) for this case are

$$\begin{aligned}&(\mu_2 h_2 - \mu_1 h_1)\, t_1 + \mu_1[v_\alpha(i-1,j) - v_\alpha(i-1,j+1)] + \mu_2[v_\alpha(i-1,j+1) - v_\alpha(i-1,j)]\\ &= (\mu_2 h_2 - \mu_1 h_1)\, t_1 + (\mu_1 - \mu_2)\,[v_\alpha(i-1,j) - v_\alpha(i-1,j+1)].\end{aligned}$$

Our assumption $\mu_1 \geq \mu_2$ and Proposition 3.1 imply this last expression is bounded above by 0.

Case A.2: Customer service completions.

Suppose that policies $\pi_i (i = 2, 3, 4)$ all serve a class 1 customer whereas π_1 serves a class 2 customer. If the first event is a class 2 service completion in Process 1 (with probability $\mu_2/\lambda + \mu_1 + \mu_2 + i\tau$), after which all processes follow optimal controls, then the remaining costs on the left-hand side of the inequality in (A.2) (with the denominator of the probability

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

suppressed) are

$$\mu_2(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_2\Big[\mu_1\big[v_\alpha(i-1,j-1) - v_\alpha(i,j)\big] + \mu_2\big[v_\alpha(i,j) - v_\alpha(i,j-1)\big]\Big]. \quad \textbf{(A.3)}$$

If the first event is a class 1 service completion in Processes 2–4 (with probability $\mu_1/\lambda + \mu_1 + \mu_2 + i\tau$), again after which optimal controls are used, then the remaining costs are

$$\mu_1(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_1\Big[\mu_2\big[v_\alpha(i-1,j) - v_\alpha(i-1,j-1)\big]\Big]. \quad \textbf{(A.4)}$$

Adding expressions (A.3) and (A.4) yields

$$(\mu_1 + \mu_2)(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_2\Big[\mu_1\big[v_\alpha(i-1,j) - v_\alpha(i,j)\big] + \mu_2\big[v_\alpha(i,j) - v_\alpha(i,j-1)\big]\Big].$$

The terms inside the brackets in this last expression are implied by the left-hand side of the inequality in (A.2). That is, we may restart the argument from here.

Suppose that policy π_2 serves a class 1 customer, whereas π_4 serves a class 2 customer. In this case, let π_1 serve a class 2 customer and π_3 a class 1 customer. If the first event is a class 2 service completion in Processes 1 and 4 (with probability $\mu_2/\lambda + \mu_1 + \mu_2 + i\tau$), after which all processes follow optimal controls, then the remaining costs on the left-hand side of the inequality in (A.2) are

$$\mu_2(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_2\Big[\mu_1\big[v_\alpha(i-1,j-1) - v_\alpha(i,j)\big] + \mu_2\big[v_\alpha(i,j) - v_\alpha(i,j-2)\big]\Big]. \quad \textbf{(A.5)}$$

If the first event is a class 1 (with probability $\mu_1/\lambda + \mu_1 + \mu_2 + i\tau$) service completion in Processes 2 and 3, again after which optimal controls are used, then the remaining costs are

$$\mu_1(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_1\Big[\mu_2[v_\alpha(i-1,j) - v_\alpha(i,j-1)]\Big]. \quad \textbf{(A.6)}$$

Adding (A.5) and (A.6) we obtain

$$(\mu_1 + \mu_2)(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_2\Big[\mu_1\big[v_\alpha(i-1,j-1) - v_\alpha(i,j-1)\big] + \mu_2\big[v_\alpha(i,j-1) - v_\alpha(i,j-2)\big]\Big] + \mu_2\Big[\mu_1\big[v_\alpha(i-1,j) - v_\alpha(i,j)\big] + \mu_2\big[v_\alpha(i,j) - v_\alpha(i,j-1)\big]\Big].$$

The expression inside the first pair of brackets above is the same expression that is on the left-hand side of the inequality in (A.1) but with the role of state (i,j) replaced with state $(i,j-1)$. For this expression, we may relabel the states and continue as though we had started in these states. The expression inside the second pair of brackets above equals the left-hand side of the inequality in (A.1). For this expression, we may restart the argument from here.

Suppose that policy π_2 serves a class 2 customer whereas π_4 serves a class 1 customer. In this case, let π_1 serve a class 2 customer and π_3 a class 1 customer. If the first event is a class 2 service completion in Processes 1 and 2 (with probability $\mu_2/\lambda + \mu_1 + \mu_2 + i\tau$), after which all processes follow optimal controls, then the remaining costs on the left-hand side of the inequality in (A.2) are

$$\mu_2(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_2\Big[\mu_1\big[v_\alpha(i-1,j-1) - v_\alpha(i,j-1)\big] + \mu_2\big[v_\alpha(i,j) - v_\alpha(i,j-1)\big]\Big]. \quad \textbf{(A.7)}$$

If the first event is a class 1 service completion in Processes 3 and 4 (with probability $\mu_1/\lambda + \mu_1 + \mu_2 + i\tau$), again after which optimal controls are used, then the remaining costs are

$$\mu_1(\mu_2 h_2 - \mu_1 h_1)t_1 + \mu_1\Big[\mu_1[v_\alpha(i-1,j) - v_\alpha(i,j)] + \mu_2\big[v_\alpha(i-1,j) - v_\alpha(i-1,j-1)\big]\Big]. \quad \textbf{(A.8)}$$

Adding (A.7) and (A.8) we obtain

$$(\mu_1 + \mu_2)(\mu_2 h_2 - \mu_1 h_1)t_1 + (\mu_1 + \mu_2)\Big[\mu_1\big[v_\alpha(i-1,j) - v_\alpha(i,j)\big] + \mu_2\big[v_\alpha(i,j) - v_\alpha(i,j-1)\big]\Big].$$

The expression inside the pair of brackets above is the expression on left-hand side of the inequality in (A.1). We may thus restart the argument from here.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

In every case save one we may (possibly) relabel the states and restart the argument. By doing so we can wait until Case A.1 occurs. In particular, Case A.1 yields the result.

Appendix B: Existence of a Solution to the Optimality Equations

In this section, we show that the optimality equations have a solution. Let $\mathbb{A} = \{1, 2, \ldots, n\}$ denote the action space with action $a = i \in \mathbb{A}$ denoting action serve a class i customer. The set of actions available at state $x \in \mathbb{X}$ is $A(x) \subset \mathbb{A}$. Note that $\mathbb{A}$, and hence $A(x)$ for all $x \in \mathbb{X}$, are Polish spaces with the discrete metric on the set of natural numbers. For states $y \neq x$, let $q(y|x, a)$ be the rate at which a process leaves x and goes to y given that action a is chosen. Moreover, let $-q(x|x, a)$ be the rate at which a Markov process leaves state x under action a. Note for this study, the transition rate kernel is *conservative* (i.e., $\sum_{y\in\mathbb{X}} q(y|x, a) = 0$ for all $x \in \mathbb{X}, a \in A(x)$) and *stable* (i.e., $q(x) := \sup\{-q(x|x, a) : a \in A(x)\} < \infty$ for all $x \in \mathbb{X}$). The following notation is adopted from [18,30] and will be used throughout this section.

- For any measurable function $h \geq 1$ on $\mathbb{X}$, we define the h-weighted supremum norm $\|\cdot\|_h$ of a real-valued measurable function f on $\mathbb{X}$ by

$$\|f\|_h := \sup_{x\in\mathbb{X}}\{h(x)^{-1}|f(x)|\}$$

  and the Banach space $\mathcal{B}_h(\mathbb{X}) := \{f : \|f\|_h < \infty\}$.
- Let $K := \{(x, a) \in \mathbb{X} \times \mathbb{A} : a \in A(x)\}$ denote the family of state-action pairs.
- Let Π be the set of all randomized Markov policies, Π_s the set of all randomized Markov stationary policies, and F the set of all stationary, non-idling policies.

Customer changing status imply that the transition rates are unbounded. As a result, we verify that each process generated by each *Markov* policy yields a transition kernel that (in the one-dimensional case) has row sums equal to one for all time. That is to say, we do not have an infinite number of transitions in finite time. To do so, we verify the following assumption from [17] (see Assumption B.3 in [17]).

Assumption B.3: *There exists a sequence of subsets of* $\mathbb{X}_m \subset \mathbb{X}$*, a non-decreasing function* $h_A \geq 1$ *on* $\mathbb{X}$ *and constants* $b_A \geq 0$ *and* $c_A \in \mathbb{R}$ *such that*

1. $\mathbb{X}_m \uparrow \mathbb{X}$ *and for each* $m \geq 1$, $\sup\{q(x)|x \in \mathbb{X}_m\} < \infty$.
2. $\inf\{h_A(x)|x \notin \mathbb{X}_m\} \to \infty$ *as* $m \to \infty$.
3. $\sum_{y\in\mathbb{X}} q(y|x, a)h_A(y) \leq c_A h_A(x) + b_A$ *for all* $(x, a) \in K$.

Lemma B.4: *Fix* $\gamma \in (0, \alpha)$*, where* $\alpha > 0$ *is a fixed, positive discount rate, and take* $h_A(x) = h_A((x_1, \ldots, x_n)) = \exp\{\epsilon(\sum_{i=1}^n x_i)\}$ *with* $\epsilon = \log(\gamma/\lambda + 1) > 0$*. It follows that Assumption B.3 holds with* $\mathbb{X}_m = \{x = (x_1, \ldots, x_n)|\sum_{i=1}^n x_i \leq m\}$*;* $h_A(x) = h_A((x_1, \ldots, x_n)) = \exp\{\epsilon(\sum_{i=1}^n x_i)\}$*;* $c_A = \gamma$*; and* $b_A = 0$.

Proof: The first two statements of Assumption B.3 are clear. Also note that $h(\cdot)$ is a non-decreasing function that is bounded below by 1. To show the third statement, fix $x \in \mathbb{X}$ and expand expression $\sum_{y\in\mathbb{X}} h_A(y)q(y|x, a)$ to obtain

$$\sum_{y\in\mathbb{X}} q(y|x, a)h_A(y) = \exp\left\{\epsilon\left(\sum_{i=1}^n x_i\right)\right\}\left[\lambda\big[\exp\{\epsilon\} - 1\big] + \sum_{i=1}^n \mu_i \mathbb{1}(a = i)\big[\exp\{-\epsilon\} - 1\big]\right.$$

$$\leq \lambda\big[\exp\{\epsilon\} - 1\big]\exp\left\{\epsilon\left(\sum_{i=1}^n x_i\right)\right\}$$

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

$$
\begin{aligned}
&= \gamma \exp\left\{\epsilon\left(\sum_{i=1}^{n} x_i\right)\right\} \\
&= \gamma h_A\left((x_1, \ldots, x_n)\right) \\
&= \gamma h_A(x),
\end{aligned}
$$

as desired. ■

In addition to Assumption B.3, we will need to verify the following assumption from [17] (see Assumption B.5 in [17]).

Assumption B.5: *With c_A and $h_A(\cdot)$ as in Assumption B.3:*

1. *Either $c_A \leq 0$, or $c_A - \alpha < 0$ when $c_A > 0$; and*
2. *There exist non-negative constants M_1 and M_2 such that*

$$\left|\sum_{i=1}^{n} x_i h_i\right| \leq M_1 + M_2 h_A(x)$$

for every $x \in \mathbb{X}$.

Lemma B.6: *Assumption B.5 holds with $M_1 = 0$ and $M_2 = \max_{i\in\{1,\ldots,n\}}\{h_i\}/\epsilon$.*

Proof: Note that we have chosen $c_A = \gamma$ in Assumption B.3 so that $0 < c_A < \alpha$, and hence, the first statement clearly holds.

For the second statement, note that

$$
\begin{aligned}
\frac{\left|\sum_{i=1}^{n} x_i h_i\right|}{h_A(x)} &= \frac{\sum_{i=1}^{n} x_i h_i}{\exp\{\epsilon(\sum_{i=1}^{n} x_i)\}} \\
&\leq \frac{\sum_{i=1}^{n} x_i h_i}{1 + \epsilon(\sum_{i=1}^{n} x_i)} \\
&\leq \frac{\max_{\{i\in\{1,\ldots,n\}\}}\{h_i\}(\sum_{i=1}^{n} x_i)}{1 + \epsilon(\sum_{i=1}^{n} x_i))} \\
&\leq \frac{\max_{\{i\in\{1,\ldots,n\}\}}\{h_i\}}{\epsilon} = M_2,
\end{aligned}
$$

as desired. ■

Assumptions B.3 and B.5 imply that the function v_α is the unique solution within $\mathcal{B}_h(\mathbb{X})$ of the discounted-cost optimality equation (see part (2) of Theorem 3.2 in [17]).

The next assumption we need to check is Assumption B.7 (c.f. Assumption B.7 in [17]). That is, Assumptions B.3–B.7 imply that Theorem 3.2 in [17] holds. In particular, there exists an optimal (deterministic) stationary optimal policy, and this policy attains the minimum in the right-hand side of the discounted cost optimality equations.

Assumption B.7:

1. *For each $x \in \mathbb{X}$, the set $A(x)$ is compact;*
2. *The functions $q(y|x,a)$, $c(x,a) \equiv \sum_{i=1}^{n} h_i x_i$, and $\sum_{y\in\mathbb{X}} q(y|x,a)h_A(y)$ are continuous in $a \in A(x)$ for each fixed $x, y \in \mathbb{X}$; and*

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

3. *Given h_A, the function in Assumption B.3, there exists a non-negative function $h_C : \mathbb{X} \to \mathbb{R}$ and constants $c_C > 0$, $b_C \geq 0$, and $M_C > 0$ such that $q(x)h_A(x) \leq M_C h_C(x)$ for every $x \in \mathbb{X}$ and*

$$\sum_{y \in \mathbb{X}} q(y|x,a)h_C(y) \leq c_C h_C(x) + b_C$$

for all $(x,a) \in K$.

LEMMA B.8: *Assumption B.7 holds with $h_C(x) = \left(\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i x_i\right) h_A(x) = q(x)h_A(x)$ where h_A is the function from Assumption B.3; $c_C = c_A + \sum_{i=1}^{n-1}\tau_i + \left(\sum_{i=1}^{n-1}\tau_i\right) c_A/\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\}$ where the constant c_A is the constant from Assumption B.3; $b_C = 0$; and $M_C = 1$.*

PROOF: The first statement follow as a consequence of $A(x)$ being finite for all $x \in \mathbb{X}$.

The second function in the second statement, $c(x,a)$, does not depend on $a \in A(x)$ (i.e., it is constant in $a \in A(x)$), and hence, is continuous in $a \in A(x)$ for each fixed $x \in \mathbb{X}$. To complete the proof of the second statement, it remains to show that the first and third functions in the second statement are continuous in $a \in A(x)$ for each fixed $x, y \in \mathbb{X}$. To this end, recall that a real valued function f on $A(x)$ is continuous in $a \in A(x)$ if it is continuous at each point $a \in A(x)$. The latter is the case provided that

$$\forall \varepsilon > 0 \quad \exists \delta > 0 \quad \forall a' \in A(x) \quad ((d_{A(x)}(a', a) < \delta) \implies (d_{\mathbb{R}}(f(a'), f(a)) < \epsilon)),$$

where $d_{A(x)}(\cdot,\cdot)$ and $d_{\mathbb{R}}(\cdot,\cdot)$, respectively, are metrics on $A(x)$ (e.g., discrete metric on $\mathbb{N}$) and $\mathbb{R}$ (e.g., distance metric on $\mathbb{R}$). Fix $x, y \in \mathbb{X}$. Suppose $d_{A(x)}(\cdot,\cdot)$ is the discrete metric on $A(x) \subset A \subset \mathbb{N}$ and $d_{\mathbb{R}}(\cdot,\cdot)$ the usual distance metric on $\mathbb{R}$. The result will follow by choosing $0 < \delta \leq 1$. In the latter case, $d_{A(x)}(a', a) < \delta$ implies $a' = a$, and hence, $q(y|x,a') = q(y|x,a)$ and $\sum_{y\in\mathbb{X}} q(y|x,a')h_A(y) = \sum_{y\in\mathbb{X}} q(y|x,a)h_A(y)$, so that the two functions are continuous at $a \in A(x)$, and hence, are continuous in $a \in A(x)$ for each fixed $x, y \in \mathbb{X}$.

Next, observe that $h_C(x)$ is a non-negative function and satisfies $q(x)h_A(x) \leq M_C h_C(x)$ for every $x \in \mathbb{X}$. Lastly, fix $x \in \mathbb{X}$ and note that $q(y) \leq \lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i\,(x_i+1)$, for every state $y \neq x$ that is reachable from x. As a result, we obtain the following upper bound for $\sum_{y\in\mathbb{X}} q(y|x,a)h_C(y) = \sum_{y\in\mathbb{X}} q(y|x,a)q(y)h_A(y)$:

$$\left(\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i\,(x_i+1)\right)\left(\textstyle\sum_{y\in\mathbb{X}} q(y|x,a)h_A(y)\right)$$
$$+ \left(\sum_{i=1}^{n-1}\tau_i\right)\left(\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i x_i\right) h_A(x).$$

Assumption B.3 implies that this last expression is bounded above by

$$\left(\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i\,(x_i+1)\right) c_A h_A(x) + \left(\sum_{i=1}^{n-1}\tau_i\right)\left(\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i x_i\right) h_A(x)$$
$$= \left(\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i x_i\right) c_A h_A(x) + \left(\sum_{i=1}^{n-1}\tau_i\right)\left(\lambda + \max_{i\in\{1,\ldots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i x_i\right) h_A(x)$$
$$+ \left(\sum_{i=1}^{n-1}\tau_i\right) c_A h_A(x).$$

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

This last expression is bounded above by

$$\begin{aligned}
&\left(c_A + \sum_{i=1}^{n-1} \tau_i\right)\left(\lambda + \max_{i\in\{1,\dots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i x_i\right) h_A(x) \\
&\quad + \frac{\left(\sum_{i=1}^{n-1}\tau_i\right)c_A}{\lambda + \max_{i\in\{1,\dots,n\}}\{\mu_i\}}\left(\lambda + \max_{i\in\{1,\dots,n\}}\{\mu_i\} + \sum_{i=1}^{n-1}\tau_i x_i\right) h_A(x) \\
&= \left(c_A + \sum_{i=1}^{n-1}\tau_i + \frac{\left(\sum_{i=1}^{n-1}\tau_i\right)c_A}{\lambda + \max_{i\in\{1,\dots,n\}}\{\mu_i\}}\right) h_C(x),
\end{aligned}$$

as desired. ■

Our next step is show that if $\lambda < \min_{i\in\{1,\dots,n\}}\{\mu_i\}$, then there exist a solution to the average cost optimality equations. To do this, we need to verify Assumption B.9 (c.f. Assumption 3.11 in [30]). That is, Assumptions B.3–B.9 imply that Theorem 3.20 in [30] holds. In particular, there exist $g^* \geq 0$ and $w \in \mathcal{B}_h(\mathbb{X})$ satisfying the average cost optimality equations; g^* is the optimal average cost, and hence is unique, but w is unique up to additive constants; and any deterministic stationary $f \in F$ is average-cost optimal if and only if f attains the minimum in the ACOE.

Assumption B.9:

1. *There exist constants $c_D > 0$ and $b_D \geq 0$, and a finite set $D \subset \mathbb{X}$ such that*

$$\sum_{y\in\mathbb{X}} q(y|x,a)h_A(y) \leq -c_D h_A(x) + b_D \mathbb{1}_D(x)$$

   for all $(x,a) \in K$ and $h_A(\cdot)$ is the function from Assumption B.3.
2. *Each $f \in F$ is irreducible.*
3. *The control model is h_A-exponentially ergodic uniformly on F, i.e., there exist constants $\delta, R > 0$ such that*

$$\sup_{f\in F} |\mathbb{E}_x^f[u(x(t))] - \mu_f(u)| \leq Re^{-\delta t}\|u\|_{h_A} h_A(x)$$

   for every $x \in \mathbb{X}$, $u \in \mathcal{B}_{h_A}(\mathbb{X})$, and $t \geq 0$, where $x(t)$ is the state process and μ_f is the unique invariant probability measure for $f \in F$.

Lemma B.10: *Suppose $\lambda < \min_{i\in\{1,\dots,n\}}\{\mu_i\}$. There exist $c_D > 0$, $b_D \geq 0$, and a finite subset D of the state space $\mathbb{X}$ such that Assumption B.9 holds.*

Proof: Let $D = \{\vec{0} = (0,\dots,0)\}$ be the singleton set consisting of the zero vector. It follows that the first statement in Assumption B.9 holds if $x \in D$ (i.e., if $x = (0,\dots,0)$). As a result, fix $x \notin D$. This implies that $x_i \geq 1$ for some $i \in \{1,\dots,n\}$. As a consequence of Proposition 3.1, assume that $a = 1$ (i.e., the action is serve a class 1 customer) as the argument for $a = i, i \geq 2$ is identical. In this case, note that

$$\frac{\left(\sum_{y\in\mathbb{X}} q(y|x,a)h_A(y)\right)}{h_A(x)} = \lambda\big[\exp\{\epsilon\} - 1\big] + \mu_1\big[\exp\{-\epsilon\} - 1\big]. \tag{B.1}$$

Next, recall that the only restriction we have imposed on γ in Assumption B.3 is that it satisfies $0 < \gamma < \alpha$. As a result, we can choose $\gamma = \lambda/k$ for some large enough natural number $k \geq 1$ given

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

below so that $0 < \gamma < \alpha$ holds and so that a condition we specify below also holds. In particular, note that $\epsilon = \log(1/k+1) > 0$ and expression (B.1) simplifies to

$$\frac{1}{k+1}\left[\frac{k+1}{k}\lambda - \mu_1\right].$$

Letting $k \to \infty$, we have that $(k+1/k) \to 1$. Since $\lambda < \mu_1$, we may choose a large enough k^* so that $(k^*+1/k^*)\lambda - \mu_1 < 0$ (and $0 < \gamma < \alpha$). The result now follows by letting $c_D = \frac{1}{2}[\mu_1 - (k^*+1/k^*)\lambda] > 0$ and $b_D = c_D + \lambda[\exp\{\epsilon\} - 1]$.

The second statement follows by noting that for any non-idling, stationary deterministic policy f, all states in the Markov process generated by f communicate.

To show the third statement, it is enough to show that there exist a Lyapunov function $h' : \mathbb{X} \to \mathbb{R}$ (i.e., a monotone non-decreasing function with $h' \geq 1$) such that $h'(x)/h_A(x)$ is monotone non-decreasing and $\lim_{x\to\infty} h'(x)/h_A(x) = \infty$. Moreover, there exist constants $c'_D \in \mathbb{R}$ and $b'_D \geq 0$ such that

$$\sum_{y\in\mathbb{X}} q(y|x,a)h'(y) \leq c'_D h'(x) + b'_D$$

for all $(x,a) \in K$ (c.f. Assumption 2.10 in [30]). To this end, let $h' = \exp\{2\epsilon(\sum_{i=1}^n x_i)\}$, $c'_D = \lambda\big[\exp\{2\epsilon\} - 1\big]$, and $b'_D = 0$. Then h' is a Lyapunov function, $h'(x)/h_A(x) = h_A(x)$ is monotone non-decreasing and $\lim_{x\to\infty} h'(x)/h_A(x) = \lim_{x\to\infty} h_A(x) = \infty$. Lastly, fix $x \in \mathbb{X}$ and note that

$$\begin{aligned}\frac{\left(\sum_{y\in\mathbb{X}} q(y|x,a)h'(y)\right)}{h'(x)} &= \lambda\big[\exp\{2\epsilon\} - 1\big] + \sum_{i=1}^n \mu_i \mathbb{1}(a=i)\big[\exp\{-2\epsilon\} - 1\big] \\ &\leq \lambda\big[\exp\{2\epsilon\} - 1\big],\end{aligned}$$

as desired ■

Appendix C: Structural Properties

To prove Proposition 4.4, we restrict the number of class 1 customers allowed in the system at any time to be N, where N is finite, show that Proposition 4.4 holds for an appropriate uniformized model indexed by N, and then let $N \to \infty$ to obtain the result for the original unbounded model. The uniformized model for each N is obtained by decreasing the arrival rate linearly in i from λ when $i = 0$ to 0 when $i = N$. The "deterioration" rate for class 1 customers is then bounded above by $N\tau$ and the MDP is uniformizable.

In particular, fix $N \in \bar{\mathbb{N}} = \{1, 2, \ldots, \infty\}$ and a stationary deterministic policy f. Let $q_{N,f}((i',j')|(i,j), a^f_{(i,j)})$ denote the rate at which the process leaves state (i,j) and goes to (i',j') with $(i',j') \neq (i,j)$ given that action $a^f_{(i,j)}$ is chosen. In this section, action $a^f_{(i,j)} = 0(1)$ denotes serve a class 1 (2) customer. Moreover, let $-q_{N,f}((i,j)|(i,j), a^f_{(i,j)})$ be the rate at which the process leaves state (i,j) under action $a^f_{(i,j)}$. Rates for this *perturbed* MDP are

$$q_{N,f}((i',j')|(i,j), a^f_{(i,j)}) = \begin{cases} \lambda(1 - \frac{i}{N}) & \text{if } (i',j') = (i+1,j) \\ \mu_1 & \text{if } (i',j') = (i-1,j) \text{ and } a^f_{(i,j)} = 0 \\ \mu_2 & \text{if } (i',j') = (i,j-1) \text{ and } a^f_{(i,j)} = 1 \\ i\tau & \text{if } (i',j') = (i-1,j+1) \\ \quad -\lambda(1 - \frac{i}{N}) - \cdots & \\ \mu_1 \mathbb{1}\{a^f_{(i,j)} = 0, i \geq 1\} - \cdots & \text{if } (i',j') = (i,j). \\ \mu_2 \mathbb{1}\{a^f_{(i,j)} = 1, j \geq 1\} - i\tau & \end{cases}$$

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

Note that the rate $q_{\infty,f}$ coincides with the rates for the original unbounded model and rates $q_{N,f}$ for each $N < \infty$ are uniformly bounded.

As alluded to, each *perturbed* and truncated process is uniformizable for $N < \infty$. We can determine structural properties of the value function for $N < \infty$ via successive approximations. Let $\bar{\mu} = \max\{\mu_1, \mu_2\}$, and without loss of generality, let $\Psi_N := \lambda + \bar{\mu} + N\tau = 1$ for $N < \infty$. We fix $N < \infty$ in this section and subsequently suppress the dependence on N. For example, the uniformized discount factor $\gamma_N = \Psi_N/\Psi_N + \alpha$ will simply be denoted by γ. The discrete-time optimality criteria are defined for a fixed policy f by

$$v^f_{M,\gamma}(x) := \sum_{m=0}^{M-1} \gamma^m c\left(X^f_m, d_m(X^f_m)\right), \tag{C.1}$$

$$v^f_{\gamma}(x) := \lim_{M\to\infty} v^f_{M,\gamma}(x), \tag{C.2}$$

where $\{X^f_m, m \geq 0\}$ denotes the stochastic process representing the state at decision epoch m. Equations (C.1) and (C.2) define the M-stage expected discounted cost and the infinite horizon expected discounted cost, respectively. Again, in each case we define optimal values $y(i,j) := \inf_{f\in F} y^f(i,j)$, where $y = v_{M,\gamma}$ or v_γ depending on the optimality criterion.

If we let $v_{m,\gamma}$ be zero outside its state space for all $m \geq 0$, then the (discrete time) finite horizon optimality equations for $i < N$ (with $v^f_{0,\gamma} = 0$) are

$$v_{m+1,\gamma}(i,j) = ih_1 + jh_2 + \gamma\left[\lambda\left(1 - \frac{i}{N}\right)v_{m,\gamma}(i+1,j) + i\tau v_{m,\gamma}(i-1,j+1)\right]$$
$$+\gamma\begin{cases} (N\tau + \bar{\mu})v_{m,\gamma}(0,0) & \text{if } i = j = 0 \\ \mu_1 v_{m,\gamma}(i-1,0) + \left(\frac{i\lambda}{N} + (N-i)\tau + \bar{\mu} - \mu_1\right)v_{m,\gamma}(i,0) & \text{if } 0 < i \leq N, j = 0 \\ \mu_2 v_{m,\gamma}(i,j-1) + \left(N\tau + \bar{\mu} - \mu_2\right)v_{m,\gamma}(0,j) & \text{if } i = 0, 0 < j \\ \min_{s\in\{0,1\}}\left[s\mu_1 v_{m,\gamma}(i-1,j) + (1-s)\mu_2 v_{m,\gamma}(i,j-1)\right. & \\ \quad \left. +\left(\frac{i\lambda}{N} + (N-i)\tau + \bar{\mu} - s\mu_1 - (1-s)\mu_2\right)v_{m,\gamma}(i,j)\right] & \text{if } 0 < i \leq N, 0 < j. \end{cases}$$

The proof is by induction (where $v_{0,\gamma} \equiv 0$). Note that inequalities 1–3 trivially hold for $m = 0$. Assume inequalities 1–3 hold for m and consider them for $m+1$.

Proof of Inequality 1 in Proposition 4.4. We start by showing $v_{m+1,\gamma}(i+2,j) - 2v_{m+1,\gamma}(i+1,j) + v_{m+1,\gamma}(i,j) \geq 0$. There are a couple of cases to consider.

Case C.11: $i + 2 < N$.

We simplify the expression by collecting similar terms. Terms that involve λ are given by

$$\lambda\left[\left(1 - \frac{i+2}{N}\right)v_{m,\gamma}(i+3,j) - 2\left(1 - \frac{i+1}{N}\right)v_{m,\gamma}(i+2,j) + \left(1 - \frac{i}{N}\right)v_{m,\gamma}(i+1,j)\right]$$
$$+\lambda\left[\frac{i+2}{N}v_{m,\gamma}(i+2,j) - 2\frac{i+1}{N}v_{m,\gamma}(i+1,j) + \frac{i}{N}v_{m,\gamma}(i,j)\right].$$

After a little algebra, this last expression can be rewritten as

$$\lambda\left(1 - \frac{i+2}{N}\right)\left[v_{m,\gamma}(i+3,j) - 2v_{m,\gamma}(i+2,j) + v_{m,\gamma}(i+1,j)\right]$$
$$+\frac{2\lambda}{N}\left[v_{m,\gamma}(i+1,j) - v_{m,\gamma}(i+2,j)\right] + \frac{\lambda i}{N}\left[v_{m,\gamma}(i+2,j) - 2v_{m,\gamma}(i+2,j) + v_{m,\gamma}(i,j)\right]$$
$$+\frac{2\lambda}{N}\left[v_{m,\gamma}(i+2,j) - v_{m,\gamma}(i+1,j)\right],$$

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

which simplifies to

$$\lambda\left(1-\frac{i+2}{N}\right)\Big[v_{m,\gamma}(i+3,j)-2v_{m,\gamma}(i+2,j)+v_{m,\gamma}(i+1,j)\Big]$$
$$+\frac{\lambda i}{N}\Big[v_{m,\gamma}(i+2,j)-2v_{m,\gamma}(i+1,j)+v_{m,\gamma}(i,j)\Big].$$

The inductive hypothesis implies that this last expression is non-negative. The terms with coefficients τ are

$$\tau\Big[(i+2)v_{m,\gamma}(i+1,j+1)-2(i+1)v_{m,\gamma}(i,j+1)+iv_{m,\gamma}(i-1,j+1)\Big]$$
$$+\tau\Big[(N-i-2)v_{m,\gamma}(i+2,j)-2(N-i-1)v_{m,\gamma}(i+1,j)+(N-i)v_{m,\gamma}(i,j)\Big],$$

which simplifies to

$$i\tau\Big[v_{m,\gamma}(i+1,j+1)-2v_{m,\gamma}(i,j+1)+v_{m,\gamma}(i-1,j+1)\Big]$$
$$+(N-i-2)\tau\Big[v_{m,\gamma}(i+2,j)-2v_{m,\gamma}(i+1,j)+v_{m,\gamma}(i,j)\Big]$$
$$+2\tau\Big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\Big].$$

The inductive step (i.e., the convexity of $v_{m,\gamma}(\cdot,j)$ and supermodularity of $v_{m,\gamma}(\cdot,\cdot)$) implies that this last expression is non-negative. The final terms involve $\bar{\mu}$, μ_1, and μ_2. For these terms, there are several cases to consider. First observe that if $j=0$, then the result trivially follow since, as consequence of Proposition 3.1, it is optimal to serve a class 1 customer except to avoid unforced idling. As a result, suppose $j\geq 1$ and consider the other cases.

Subcase C.12: Suppose that

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i+1,j)+(1-s)\mu_2 v_{m,\gamma}(i+2,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i+2,j)\big]$$
$$=\mu_1 v_{m,\gamma}(i+1,j)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i+2,j)$$

and

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i-1,j)+(1-s)\mu_2 v_{m,\gamma}(i,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i,j)\big]$$
$$=\mu_1 v_{m,\gamma}(i-1,j)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i,j).$$

In this case, the terms involving $\bar{\mu}$, μ_1, and μ_2 are bounded below by

$$\mu_1 v_{m,\gamma}(i+1,j)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i+2,j)-2\mu_1 v_{m,\gamma}(i,j)-2(\bar{\mu}-\mu_1)v_{m,\gamma}(i+1,j)$$
$$+\mu_1 v_{m,\gamma}(i-1,j)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i,j)$$
$$=\mu_1\Big[v_{m,\gamma}(i+1,j)-2v_{m,\gamma}(i,j)+v_{m,\gamma}(i-1,j)\Big]$$
$$+(\bar{\mu}-\mu_1)\Big[v_{m,\gamma}(i+2,j)-2v_{m,\gamma}(i+1,j)+v_{m,\gamma}(i,j)\Big],$$

and the inductive hypothesis implies that this last expression is bounded below by 0.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

Subcase C.13: Suppose that

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i+1,j)+(1-s)\mu_2 v_{m,\gamma}(i+2,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i+2,j)\big]$$
$$=\mu_1 v_{m,\gamma}(i+1,j)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i+2,j)$$

and

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i-1,j)+(1-s)\mu_2 v_{m,\gamma}(i,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i,j)\big]$$
$$=\mu_2 v_{m,\gamma}(i,j-1)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i,j).$$

In this case, the terms involving $\bar{\mu}$, μ_1, and μ_2 are bounded below by

$$\begin{aligned}&\mu_1 v_{m,\gamma}(i+1,j)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i+2,j)-2\mu_1 v_{m,\gamma}(i,j)-2(\bar{\mu}-\mu_1)v_{m,\gamma}(i+1,j)\\&\quad+\mu_1 v_{m,\gamma}(i-1,j)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i,j)\\&=\mu_1\Big[v_{m,\gamma}(i+1,j)-2v_{m,\gamma}(i,j)+v_{m,\gamma}(i-1,j)\Big]\\&\quad+(\bar{\mu}-\mu_1)\Big[v_{m,\gamma}(i+2,j)-2v_{m,\gamma}(i+1,j)+v_{m,\gamma}(i,j)\Big],\end{aligned}$$

and the inductive hypothesis implies that this last expression is bounded below by 0.

Subcase C.14: Suppose that

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i+1,j)+(1-s)\mu_2 v_{m,\gamma}(i+2,j-1)+\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i+2,j)\big]$$
$$=\mu_2 v_{m,\gamma}(i+2,j-1)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i+2,j).$$

and that

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i-1,j)+(1-s)\mu_2 v_{m,\gamma}(i,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i,j)\big]$$
$$=\mu_2 v_{m,\gamma}(i,j-1)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i,j).$$

In this case, the terms involving $\bar{\mu}$, μ_1, and μ_2 are bounded below by

$$\begin{aligned}&\mu_2 v_{m,\gamma}(i+2,j-1)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i+2,j)-2\mu_2 v_{m,\gamma}(i+1,j-1)-2(\bar{\mu}-\mu_2)v_{m,\gamma}(i+1,j)\\&\quad+\mu_2 v_{m,\gamma}(i,j-1)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i,j)\\&=\mu_2\Big[v_{m,\gamma}(i+1,j-1)-2v_{m,\gamma}(i+1,j-1)+v_{m,\gamma}(i,j-1)\Big]\\&\quad+(\bar{\mu}-\mu_2)\Big[v_{m,\gamma}(i+2,j)-2v_{m,\gamma}(i+1,j)+v_{m,\gamma}(i,j)\Big],\end{aligned}$$

and this last expression is bounded below by 0 by the inductive hypothesis.

Subcase C.15: Suppose that

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i+1,j)+(1-s)\mu_2 v_{m,\gamma}(i+2,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i+2,j)\big]$$
$$=\mu_2 v_{m,\gamma}(i+2,j-1)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i+2,j)$$

and that

$$\min_{s\in\{0,1\}}\big[s\mu_1 v_{m,\gamma}(i-1,j)+(1-s)\mu_2 v_{m,\gamma}(i,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i,j)\big]$$
$$=\mu_1\mathbb{1}_{\{i-1\geq 1\}}v_{m,\gamma}(i-1,j)+(\bar{\mu}-\mu_1\mathbb{1}_{\{i-1\geq 1\}})v_{m,\gamma}(i,j).$$

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

Once again, the terms involving $\bar{\mu}$, μ_1, and μ_2 are bounded below by

$$\begin{aligned}
&\mu_2 v_{m,\gamma}(i+2,j-1) + (\bar{\mu}-\mu_2)v_{m,\gamma}(i+2,j) - 2\mu_2 v_{m,\gamma}(i+1,j-1) - 2(\bar{\mu}-\mu_2)v_{m,\gamma}(i+1,j)\\
&\quad + \mu_2 v_{m,\gamma}(i,j-1) + (\bar{\mu}-\mu_2)v_{m,\gamma}(i,j)\\
&= \mu_2\Big[v_{m,\gamma}(i+2,j-1) - 2v_{m,\gamma}(i+1,j-1) + v_{m,\gamma}(i,j-1)\Big]\\
&\quad + (\bar{\mu}-\mu_2)\Big[v_{m,\gamma}(i+1,j) - 2v_{m,\gamma}(i+1,j) + v_{m,\gamma}(i,j)\Big],
\end{aligned}$$

and this last expression is bounded below by 0 by the inductive hypothesis. This concludes the proof of convexity for the case $i+2<N$.

Case C.16: $i+2=N$

The terms with coefficients λ are now given by

$$\frac{2\lambda}{N}\Big[v_{m,\gamma}(i+1,j) - v_{m,\gamma}(i+2,j)\Big] + \lambda\left[v_{m,\gamma}(i+2,j) - 2\frac{N-1}{N}v_{m,\gamma}(i+1,j) + \frac{N-2}{N}v_{m,\gamma}(i,j)\right].$$

After a little algebra, this last expression can be rewritten as

$$\begin{aligned}
&\frac{2\lambda}{N}\Big[v_{m,\gamma}(i+1,j) - v_{m,\gamma}(i+2,j)\Big] + \frac{2\lambda}{N}\Big[v_{m,\gamma}(i+2,j) - v_{m,\gamma}(i+1,j)\Big]\\
&\quad + \frac{\lambda(N-2)}{N}\Big[v_{m,\gamma}(i+2,j) - 2v_{m,\gamma}(i+1,j) + v_{m,\gamma}(i,j)\Big],
\end{aligned}$$

which simplifies to

$$\frac{\lambda(N-2)}{N}\Big[v_{m,\gamma}(i+2,j) - 2v_{m,\gamma}(i+1,j) + v_{m,\gamma}(i,j)\Big].$$

The inductive hypothesis implies that this last expression is non-negative. The terms with coefficients τ and the terms with coefficients $\bar{\mu}$, μ_1, and μ_2 for this case are the same as in the previous case (when $i+1<N$), and hence, are non-negative.

Proof of Inequality 2 in Proposition 4.4. Next, we show $v_{m+1,\gamma}(i,j+2) - 2v_{m+1,\gamma}(i,j+1) + v_{m+1,\gamma}(i,j) \geq 0$. Collecting similar terms, we have that the terms involving λ are

$$\begin{aligned}
&\lambda\Big(1-\frac{i}{N}\Big)\Big[v_{m,\gamma}(i+1,j+2) - 2v_{m,\gamma}(i+1,j+1) + v_{m,\gamma}(i+1,j)\Big]\\
&+ \frac{\lambda i}{N}\Big[v_{m,\gamma}(i,j+2) - 2v_{m,\gamma}(i,j+1) + v_{m,\gamma}(i,j)\Big].
\end{aligned}$$

The inductive hypothesis (i.e., the convexity of $v_{m,\gamma}(i,\cdot)$) implies that this last expression is non-negative. The terms with coefficients τ are

$$\begin{aligned}
&i\tau\Big[v_{m,\gamma}(i-1,j+3) - 2v_{m,\gamma}(i-1,j+2) + v_{m,\gamma}(i-1,j+1)\Big]\\
&\quad + (N-i)\tau\Big[v_{m,\gamma}(i,j+2) - 2v_{m,\gamma}(i,j+1) + v_{m,\gamma}(i,j)\Big],
\end{aligned}$$

which (by the inductive step) are also non-negative. Lastly, note that by replacing the roles of i and j and of μ_1 and μ_2 in the proof of inequality 1 above, we can show that the terms involving $\bar{\mu}$, μ_1, and μ_2 for 2 are also non-negative. This completes the proof for 2.

Proof of Inequality 3 in Proposition 4.4. We complete the proof of the proposition by showing $v_{m+1,\gamma}(i+1,j+1) + v_{m+1,\gamma}(i,j) - v_{m+1,\gamma}(i+1,j) - v_{m+1,\alpha}(i,j+1) \geq 0$. There are a couple of cases to consider.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

Case C.17: $i+1 < N$.

Terms involving λ are

$$\lambda\left[\left(1-\frac{i+1}{N}\right)v_{m,\gamma}(i+2,j+1)+\left(1-\frac{i}{N}\right)v_{m,\gamma}(i+1,j)-\left(1-\frac{i+1}{N}\right)v_{m,\gamma}(i+2,j)\right.$$
$$\left.-\left(1-\frac{i}{N}\right)v_{m,\gamma}(i+1,j+1)\right]+\lambda\left[\frac{i+1}{N}v_{m,\gamma}(i+1,j+1)\right.$$
$$\left.+\frac{i}{N}v_{m,\gamma}(i,j)-\frac{i+1}{N}v_{m,\gamma}(i+1,j)-\frac{i}{N}v_{m,\gamma}(i,j+1)\right],$$

which simplifies to

$$\lambda\left(1-\frac{i+1}{N}\right)\Big[v_{m,\gamma}(i+2,j+1)+v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i+2,j)-v_{m,\gamma}(i+1,j+1)\Big]$$
$$+\frac{\lambda i}{N}\Big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\Big]$$
$$+\frac{\lambda}{N}\Big[v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i+1,j+1)\Big]+\frac{\lambda}{N}\Big[v_{m,\gamma}(i+1,j+1)-v_{m,\gamma}(i+1,j)\Big].$$

This last expression equals

$$\lambda\left(1-\frac{i+1}{N}\right)\Big[v_{m,\gamma}(i+2,j+1)+v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i+2,j)-v_{m,\gamma}(i+1,j+1)\Big]$$
$$+\frac{\lambda i}{N}\Big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\Big].$$

The inductive hypothesis implies that this last expression is non-negative. The terms with coefficients τ are

$$\tau\Big[(i+1)v_{m,\gamma}(i,j+2)+iv_{m,\gamma}(i-1,j+1)-(i+1)v_{m,\gamma}(i,j+1)-iv_{m,\gamma}(i-1,j+2)\Big]$$
$$+\tau\Big[(N-i-1)v_{m,\gamma}(i+1,j+1)+(N-i)v_{m,\gamma}(i,j)-(N-i-1)v_{m,\gamma}(i+1,j)$$
$$-(N-i)v_{m,\gamma}(i,j+1)\Big],$$

which simplifies to

$$i\tau\Big[v_{m,\gamma}(i,j+2)+v_{m,\gamma}(i-1,j+1)-v_{m,\gamma}(i,j+1)-v_{m,\gamma}(i-1,j+2)\Big]$$
$$+\tau\Big[v_{m,\gamma}(i,j+2)-v_{m,\gamma}(i,j+1)\Big]$$
$$+(N-i-1)\tau\Big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\Big]$$
$$+\tau\Big[v_{m,\gamma}(i,j)-v_{m,\gamma}(i,j+1)\Big].$$

This last expression equals

$$i\tau\Big[v_{m,\gamma}(i,j+2)+v_{m,\gamma}(i-1,j+1)-v_{m,\gamma}(i,j+1)-v_{m,\gamma}(i-1,j+2)\Big]$$
$$+(N-i-1)\tau\Big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\Big]$$
$$+\tau\Big[v_{m,\gamma}(i,j+2)-2v_{m,\gamma}(i,j+1)-v_{m,\gamma}(i,j)\Big],$$

which is non-negative by the inductive step. We next consider terms which involve $\bar{\mu}$, μ_1, and μ_2. Note that if it is optimal to serve a class 1 (2) customer in states $(i+1,j+1)$ and (i,j), then, by

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

also serving a class 1 (2) customer in states $(i+1,j)$ and $(i,j+1)$, the terms with coefficients $\bar{\mu}$, μ_1, and μ_2 will be bounded below by 0. Consider now the other cases.

Subcase C.18: Suppose that

$$\min_{s\in\{0,1\}}\left[s\mu_1 v_{m,\gamma}(i,j+1)+(1-s)\mu_2 v_{m,\gamma}(i+1,j)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i+1,j+1)\right]$$
$$=\mu_1 v_{m,\gamma}(i,j+1)+(\bar{\mu}-\mu_1)v_{m,\gamma}(i+1,j+1).$$

and

$$\min_{s\in\{0,1\}}\left[s\mu_1 v_{m,\gamma}(i-1,j)+(1-s)\mu_2 v_{m,\gamma}(i,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i,j)\right]$$
$$=\mu_2 v_{m,\gamma}(i,j-1)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i,j).$$

In this case, the terms involving $\bar{\mu}$, μ_1, and μ_2 are bounded below by

$$(\bar{\mu}-\mu_1)\big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\big]$$
$$+\mu_2\big[v_{m,\gamma}(i,j+1)-2v_{m,\gamma}(i,j)-v_{m,\gamma}(i,j-1)\big],$$

and the inductive hypothesis implies that this last expression is bounded below by 0.

Subcase C.19: Suppose that

$$\min_{s\in\{0,1\}}\left[s\mu_1 v_{m,\gamma}(i,j+1)+(1-s)\mu_2 v_{m,\gamma}(i+1,j)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i+1,j+1)\right]$$
$$=\mu_2 v_{m,\gamma}(i+1,j)+(\bar{\mu}-\mu_2)v_{m,\gamma}(i+1,j+1)$$

and

$$\min_{s\in\{0,1\}}\left[s\mu_1 v_{m,\gamma}(i-2,j)+(1-s)\mu_2 v_{m,\gamma}(i-1,j-1)+\big(\bar{\mu}-s\mu_1-(1-s)\mu_2\big)v_{m,\gamma}(i-1,j)\right]$$
$$=\mu_1 v_{m,\gamma}(i-1,j)+(\bar{\mu}-\mu_1)v^N_{m,\gamma}(i,j).$$

In this case, the terms involving $\bar{\mu}$, μ_1, and μ_2 are bounded below by

$$(\bar{\mu}-\mu_2)\big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\big]$$
$$+\mu_1\big[v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j)-v_{m,\gamma}(i-1,j)\big],$$

and the inductive hypothesis implies that this last expression is bounded below by 0. This concludes the proof of convexity for the case $i+1<N$.

Case C.20: $i+1=N$.

The terms with coefficients τ, μ_1, and μ_2 are the same as in the previous case (when $i+1<N$) and, hence, are non-negative. The terms involving λ are given by

$$\frac{\lambda}{N}\Big[v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i+1,j+1)\Big]$$
$$+\lambda\left[v_{m,\gamma}(i+1,j+1)+\frac{N-1}{N}v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-\frac{N-1}{N}v_{m,\gamma}(i,j+1)\right],$$

which (after a little algebra) can be rewritten as

$$\frac{\lambda(N-1)}{N}\Big[v_{m,\gamma}(i+1,j+1)+v_{m,\gamma}(i,j)-v_{m,\gamma}(i+1,j)-v_{m,\gamma}(i,j+1)\Big],$$

and the inductive hypothesis implies this expression is non-negative.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

It follows these arguments that inequalities 1–3 hold for each $m \geq 0$, and hence (by letting $m \to \infty$) for v_γ. The result for the discounted cost problem follows by letting $N \to \infty$. By letting $\alpha = 0$, the same proof gives the result for the average cost criterion.

Appendix D: Approximating Controlled Markov Chains

To obtain structural properties of the value and relative value function (c.f. Proposition 4.4 and inequalities 1–2) we restricted the number of class 1 customers allowed in the system at any time to be N, where N is finite. In particular, we showed that these properties hold for an appropriate uniformized model indexed by N, and then let $N \to \infty$ to obtain the result for the original unbounded model. Slightly adapting the Smoothed Rate Truncation method from Bhulai et al. [6], the uniformized model for each N was obtained by decreasing the rate linearly in i from λ when $i = 0$ to 0 when $i = N$ so that the "deterioration" rate for class 1 customers is bounded above by $N\tau$.

This section is dedicated to justifying this approach. To do this, we can use recent results from Prieto-Rumeau and Hernández-Lerma [30] (c.f. Theorems 4.11 and 4.17 by Prieto-Rumeau and Hernández-Lerma [30]), or Bhulai et al. [6] for the average cost problem (c.f. Theorem 3.3 by Bhulai et al. [6]), or Blok and Spieksma [7] (c.f. Theorem 5.1 by Blok and Spieksma [7]) for the discounted cost problem. In this section, we will follow Prieto-Rumeau and Hernández-Lerma [30], based on Prieto-Rumeau and Hernández-Lerma [29] (for the discounted cost case) and on Prieto-Rumeau and Lorenzo [31] (for the average cost case). To this end, let

$$\mathcal{M} := \{\mathbb{X}, A, (A(x)), (q(y|x,a)), (c(x))\}$$

denote our original, unbounded control model whose optimal discounted and average costs and optimal policies we want to approximate. Consider the sequence of control models

$$\mathcal{M}_N := \{\mathbb{X}_N, A, (A(x)), (q_N(y|x,a)), (c(x))\}$$

for $N = 1, 2, \ldots$ with $X_N = \{(i,j) | i \leq N\}$ and where the entries of rate matrix Q_N^f are given by $q_{N,f}\left(\left(i',j'\right) | (i,j), a\right)$ from Appendix B with $a^f_{(i,j)}$ replaced with a and with $a = 1(2)$ denoting the action of serving a class 1 (2) customer. Note that $q_N(y|x,a)$ is continuous and hence measurable in $A(x)$ for each fixed $x, y \in X_N$. Additionally, the transition rates (by definition) are conservative and stable.

In this section, we verify conditions (Assumptions D.21–D.25) that guarantee that the sequence of control models $\mathcal{M}_N$ converges to $\mathcal{M}$ in the sense of Definitions 4.8 (for the discounted expected cost case) and 4.14 (for the long-run average costs case) by Prieto-Rumeau and Hernández-Lerma [30]. To that end, we note that the definitions for the original control $\mathcal{M}$, can be easily modified to account for the control models $\mathcal{M}_N$. For instance, given $u : \mathbb{X}_N \to \mathbb{R}$, its h norm is defined as

$$\|u\|_h := \sup_{x \in \mathbb{X}_N} \left\{ \frac{u(x)}{h(x)} \right\},$$

and $\mathcal{B}_h$ is the Banach space of functions on $\mathbb{X}_N$ with finite h-norm. The class of Markov policies for each control model $\mathcal{M}_N$ is denoted by Φ_N. Also, we denote by F_N the set of stationary deterministic policies for the control model $\mathcal{M}_N$; i.e., $f : \mathbb{X}_N \to A$ is in F_N if $f(x) \in A(x)$ for all $x \in \mathbb{X}$. Similarly, $K_N = \{(x,a) \in \mathbb{X}_N \times A | a \in A(x)\}$.

Assumption D.21: *With c_A, b_A, and $h_A(\cdot)$ as in Assumption B.3 and c'_D, b'_D, and $h'(\cdot)$ as in Assumption B.9, the following hold for each $N \geq 1$ and for all $(x,a) \in K_N$:*

1. $\sum_{y \in \mathbb{X}} q_N(y|x,a) h_A(y) \leq c_A h(x) + b_A$;
2. $\sum_{y \in \mathbb{X}} q_N(y|x,a) h'(y) \leq c'_D h'(x) + b'_D$.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

LEMMA D.22: *Assumption D.21 holds.*

PROOF: For first statement, fix $N \geq 1$ and $(x,a) \in K_N$. A simple calculation yields

$$\begin{aligned}\sum_{y\in\mathbb{X}} q_N(y|x,a)h_A(y) &= \exp\{\epsilon(i+j)\}\Big[\lambda(1-i/N)\big[\exp\{\epsilon\}-1\big] \\ &\qquad + \mu_1\mathbb{1}(a=1)\big[\exp\{-\epsilon\}-1\big] + \mu_2\mathbb{1}(a=2)\big[\exp\{-\epsilon\}-1\big]\Big] \\ &\leq \lambda\big[\exp\{\epsilon\}-1\big]\exp\{\epsilon(i+j)\} \\ &= \gamma\exp\{\epsilon(i+j)\} \\ &= \gamma h_A(x),\end{aligned}$$

so that the first statement follows. For the second statement, a similar calculation yields

$$\begin{aligned}\frac{\left(\sum_{y\in\mathbb{X}} q(y|x,a)h'(y)\right)}{h'(x)} &= \lambda(1-\frac{i}{N})\big[\exp\{2\epsilon\}-1\big] + \mu_1\mathbb{1}(a=1)\big[\exp\{-2\epsilon\}-1\big] \\ &\qquad + \mu_2\mathbb{1}(a=2)\big[\exp\{-2\epsilon\}-1\big] \\ &\leq \lambda\big[\exp\{2\epsilon\}-1\big],\end{aligned}$$

as desired. ■

Assumption D.21 implies for any discount rate α such that $0 < \gamma < \alpha$ the following hold (c.f. Theorem 4.11 by Prieto-Rumeau and Hernández-Lerma [30]):

1. For every $x \in \mathbb{X}$, $\lim_{N\to\infty} v_\gamma = v_\alpha$.
2. If $f_N \in F_N$ for $N \geq 1$, is a discount cost optimal policy for $\mathcal{M}_N$, where F_N is the set of stationary deterministic policies for $\mathcal{M}_N$, then any limit policy $f \in F$ of $\{f_N\}_{N\geq 1}$ is also discount optimal for $\mathcal{M}$.

ASSUMPTION D.23: *With $h_A(\cdot)$ as in Assumption B.3 and c_D, b_D, and D as in Assumption B.9, the following hold for each $N \geq 1$ and for all $(x,a) \in K_N$:*

1. *If $\lambda < \min\{\mu_1,\mu_2\}$, then*

$$\sum_{y\in\mathbb{X}} q_N(y|x,a)h_A(y) \leq -c_D h_A(x) + b_D\mathbb{1}_D(x).$$

2. *For every $x, y \in \mathbb{X}$, $q_N(y|x,\cdot)$ converges to $q(y|x,\cdot)$ uniformly on $A(x)$.*

LEMMA D.24: *Assumption D.21 holds.*

PROOF: For each fixed $N \geq 1$ and $(x,a) \in K_N$, the first statement can be proved in exactly the same way that the first statement in Assumption B.9 was proved. The second statement follows since $A(x)$ is finite for all $x \in \mathbb{X}$. ■

ASSUMPTION D.25: *With $h_A(\cdot)$ as in Assumption B.3, there exist constants $\nu > 0$, $c_G > 0$, and $b_G \geq 0$ such that, for each $N \geq 1$,*

$$\sum_{y\in\mathbb{X}} q_N(y|x,a)h_A^{2+\nu}(y) \leq -c_G h_A^{2+\nu}(x) + b_G.$$

LEMMA D.26: *Suppose $\lambda < \min\{\mu_1,\mu_2\}$. In this case, Assumption D.25 holds with c_G given in the proof below and $b_G = c_G + \lambda[\exp\{(2+\nu)\epsilon\} - 1]$.*

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

PROOF: The case when $x = (0, 0)$ is immediate. Suppose $x \neq 0$, and as a consequence of Proposition 3.1, assume that $a = 1$. The proof when $a \geq 2$ is exactly the same. This yields

$$\begin{aligned}\frac{\left(\sum_{y\in\mathbb{X}} q(y|x,a)h_A^{2+\nu}(y)\right)}{h_A^{2+\nu}(x)} &= \lambda\left(1-\frac{i}{N}\right)\left[\exp\{(2+\nu)\epsilon\}-1\right] \\ &\quad + \mu_1\mathbb{1}(a=1)\left[\exp\{-(2+\nu)\epsilon\}-1\right] \\ &\quad + \mu_2\mathbb{1}(a=2)\left[\exp\{-(2+\nu)\epsilon\}-1\right] \\ &\leq \lambda\left[\exp\{(2+\nu)\epsilon\}-1\right] + \mu_1\left[\exp\{-(2+\nu)\epsilon\}-1\right] \\ &= \lambda\left[\left(1+\frac{1}{k}\right)^{2+\nu}-1\right] + \mu_1\left[\left(\frac{1}{1+1/k}\right)^{2+\nu}-1\right] \\ &= \left[1-\left(\frac{1}{1+1/k}\right)^{2+\nu}\right]\left[\lambda\frac{(k+1)^{2+\nu}}{k^{2+\nu}}-\mu_1\right].\end{aligned}$$

Following the proof of the first statement in Assumption B.9, we can choose $k^* \geq 1$ large enough so that this last expression is negative (as a consequence of $\lambda < \mu_1$). The result follows by choosing

$$c_G = -\left[1-\left(\frac{1}{1+1/k}\right)^{2+\nu}\right]\left[\lambda\frac{(k+1)^{2+\nu}}{k^{2+\nu}}-\mu_1\right] \qquad \blacksquare$$

Assumptions D.21 and D.23 above imply that the following statements hold (c.f. Theorem 4.18 by Prieto-Rumeau and Hernández-Lerma [30]):

1. $\lim_{n\to\infty} g_N^* = g^*$, where g_N^* (g^*) is the optimal gain for the control model $\mathcal{M}_N$ ($\mathcal{M}$).
2. If $f_N \in F_N$ is, for each $N \geq 1$, and average cost optimal policy for $\mathcal{M}_N$, then any limit policy of $\{f_N\}$ is average cost optimal for $\mathcal{M}$.

Appendix E: Proof of Lemma 4.6

We use a sample path argument. Start two processes on the same probability space with Processes 1 and 2 starting in states $(i, j-1)$ and $(i-1, j)$, respectively. Process 1 uses a stationary optimal policy, which we denote by π_1. We construct a potentially sub-optimal policy for Process 2, which we denote by π_2, so that

$$v_\alpha^{\pi_1}(i, j-1) - v_\alpha^{\pi_2}(i-1, j) \geq 0. \tag{E.1}$$

Since π_2 is potentially sub-optimal, Eq. (3) follows from (E.1).

Note that starting from (E.1), the (immediate) costs at the next event are $(h_1 - h_2)\, t_1 \geq 0$, where t_1 is the time of the next event and the inequality is due to the assumption that $\mu_1 h_1 > \mu_2 h_2$ and $\mu_2 > \mu_1$. Moreover, if the relative position (as measured by the current states) of the two processes at the next event remains the same, then we may relabel the initial states and continue from the beginning of the argument. This occurs when any of the uncontrolled events occur that are seen by both processes. It also occurs when π_1 serves a class $k \in \{1, 2\}$ and there is at least one class k customer in Process 2, by letting π_2 also serve a class k customer, and the next event is a service completion. Consider now the other cases.

Case E.27: Customer status changes.

Suppose that the first event is a change in status in Process 1 only (with probability $\tau/\lambda + \mu_1 + \mu_2 + i\tau$), after which both processes follow an optimal control. In this case, it follows that the remaining costs on the left-hand side of the inequality in (E.1) (with the probability

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

of this event suppressed) are

$$(h_1 - h_2)\, t_1 + v_\alpha(i-1, j) - v_\alpha(i-1, j) = (h_1 - h_2)\, t_1 \geq 0.$$

Case E.28: Customer service completions.

Suppose that policy π_1 serves a class 1 customer whereas π_2 serves a class 2 customer. If the first event is a class 1 service completion in Process 1 (with probability $\mu_1/\lambda + \mu_1 + \mu_2 + i\tau$), after which both processes follow optimal controls, then the remaining costs on the left-hand side of the inequality in (E.1) (with the denominator of the probability suppressed) are

$$\mu_1(h_1 - h_2)t_1 + \mu_1\Big[v_\alpha(i-1, j-1) - v_\alpha(i-1, j)\Big]. \tag{E.2}$$

If the first event is a class 2 service completion in Process 2 (with probability $\mu_2/\lambda + \mu_1 + \mu_2 + i\tau$), again after which optimal controls are used, then the remaining costs (with the denominator of the probability suppressed) are

$$\mu_2(h_1 - h_2)t_1 + \mu_2\Big[v_\alpha(i, j-1) - v_\alpha(i-1, j-1)\Big]. \tag{E.3}$$

Adding expressions (E.2) and (E.3) yields

$$\begin{aligned}&(\mu_1 + \mu_2)(h_1 - h_2)t_1 + (\mu_2 - \mu_1)\left[v_\alpha(i, j-1) - v_\alpha(i-1, j-1)\right]\\ &\quad + \mu_1\left[v_\alpha(i, j-1) - v_\alpha(i-1, j)\right].\end{aligned}$$

Proposition 3.1 implies that the expression inside the first pair of brackets above is nonnegative. The expression inside the second pair of brackets is implied by the expression on the left-hand side of the inequality in (3). That is, for this last expression, we may restart the argument from here.

Appendix F: Proof of Inequality 1

The proof is by induction (where $v_{0,\gamma} \equiv 0$). Note that inequality 1 trivially holds for $m = 0$. Assume 1 holds for m (so that, for $i, j \geq 1$, if it is optimal to serve a class 2 customer in state (i, j), then it is also optimal to serve a class 2 customer in state $(i, j+1)$ at epoch $m+1$) and consider them for $m+1$. There are a couple of cases to consider. In what follows, discounting is suppressed without any loss of generality.

Case F.29: $i < N$.

We simplify the expression by collecting similar terms. Terms that involve λ in the expression on the left-hand side of the inequality in 1 are given by

$$\begin{aligned}&\mu_1\left[\lambda\left(1 - \frac{i}{N}\right)\left[v_{m,\gamma}(i+1, j+1) + v_{m,\gamma}(i, j) - v_{m,\gamma}(i+1, j) - v_{m,\gamma}(i, j+1)\right]\right.\\ &\quad \left. + \lambda\frac{i-1}{N}\left[v_{m,\gamma}(i, j+1) + v_{m,\gamma}(i-1, j) - v_{m,\gamma}(i, j) - v_{m,\gamma}(i-1, j+1)\right]\right],\end{aligned}$$

whereas terms that involve λ in the expression on the right-hand side of the inequality in 1 are given by

$$\begin{aligned}&\mu_2\left[\lambda\left(1 - \frac{i}{N}\right)\left[v_{m,\gamma}(i+1, j+1) - 2v_{m,\gamma}(i+1, j) + v_{m,\gamma}(i+1, j-1)\right]\right.\\ &\quad \left. + \frac{\lambda i}{N}\left[v_{m,\gamma}(i, j+1) - 2v_{m,\gamma}(i, j) + v_{m,\gamma}(i, j-1)\right]\right].\end{aligned}$$

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

The inductive hypothesis implies the result for these terms. The terms with coefficients τ in the expression on the left-hand side of the inequality in 1 are given by

$$\begin{aligned}\mu_1\Big[&(i-1)\tau\big[v_{m,\gamma}(i-1,j+2)+v_{m,\gamma}(i-2,j+1)-v_{m,\gamma}(i-1,j+1)-v_{m,\gamma}(i-2,j+2)\big]\\&+(N-i)\tau\big[v_{m,\gamma}(i,j+1)+v_{m,\gamma}(i-1,j)-v_{m,\gamma}(i,j)-v_{m,\gamma}(i-1,j+1)\big]\\&+\tau\big[v_{m,\gamma}(i-1,j+2)-2v_{m,\gamma}(i-1,j+1)+v_{m,\gamma}(i-1,j)\big]\Big],\end{aligned}$$

whereas the terms with coefficients τ on the expression in the right-hand side of the inequality in 1 are given by

$$\begin{aligned}\mu_2\Big[&i\tau\big[v_{m,\gamma}(i-1,j+2)-2v_{m,\gamma}(i-1,j+1)+v_{m,\gamma}(i-1,j)\big]\\&+(N-i)\tau\big[v_{m,\gamma}(i,j+1)-2v_{m,\gamma}(i,j)+v_{m,\gamma}(i,j-1)\big]\Big].\end{aligned}$$

The inductive hypothesis implies the result for these terms.

The final terms involve μ_1 and μ_2, and there are several cases to consider. For these terms, we rewrite 1 slightly as follows:

$$\begin{aligned}&(\mu_2-\mu_1)\left[v_{m+1,\gamma}(i,j+1)-2v_{m+1,\gamma}(i,j)+v_{m+1,\gamma}(i,j-1)\right]\\&\quad+\mu_1[v_{m+1,\gamma}(i,j-1)+v_{m+1,\gamma}(i-1,j+1)-v_{m+1,\gamma}(i,j)-v_{m+1,\gamma}(i-1,j)]\geq 0. \qquad \textbf{(F.1)}\end{aligned}$$

Next, observe that if it is optimal to serve a class 1 customer in states $(i,j+1)$, $(i,j-1)$, and $(i-1,j+1)$ at epoch m, then, by serving a class 1 customer in states (i,j) and $(i-1,j)$, the terms with coefficients μ_1 and μ_2 in (F.1) will be bounded below by 0. Similarly, if it is optimal to serve a class 2 customer in states $(i,j+1)$, $(i,j-1)$, and $(i-1,j+1)$ at epoch m, then, by serving a class 2 customer in states (i,j) and $(i-1,j)$, the terms with coefficients μ_1 and μ_2 in (F.1) will be bounded below by 0. Also note that (by the inductive hypothesis) we can assume that if it is optimal to serve a class 1 customer in state $(i,j+1)$ then it is also optimal to serve a class 1 customer in state $(i,j-1)$ in epoch m. Consider now the other cases.

Subcase F.30: Suppose it is optimal to serve a class 2 customer in states $(i,j+1)$ and $(i,j-1)$ and optimal to serve a class 1 customer in state $(i-1,j+1)$ at epoch m.

In this case, a little algebra implies that the terms involving μ_1 and μ_2 are bounded below by

$$\begin{aligned}&+\mu_2\Big[\,(\mu_2-\mu_1)\left[v_{m,\gamma}(i,j)-2v_{m,\gamma}(i,j-1)+v_{m,\gamma}(i,j-2)\right]\\&+\mu_1[v_{m,\gamma}(i,j-2)+v_{m,\gamma}(i-1,j)-v_{m,\gamma}(i,j-1)-v_{m,\gamma}(i-1,j-1)]\Big],\\&+\mu_1\Big[\,(\mu_2-\mu_1)\left[v_{m,\gamma}(i-1,j+1)-2v_{m,\gamma}(i-1,j)+v_{m,\gamma}(i-1,j-1)\right]\\&+\mu_1[v_{m,\gamma}(i-1,j-1)+v_{m,\gamma}(i-2,j+1)-v_{m,\gamma}(i-1,j)-v_{m,\gamma}(i-2,j)]\Big].\end{aligned}$$

The inductive hypothesis implies that this last expression is bounded below by 0.

Subcase F.31: Suppose it is optimal to serve a class 2 customer in states $(i,j+1)$ and $(i-1,j+1)$ and optimal to serve a class 1 customer in state $(i,j-1)$ at epoch m.

Without loss of generality, we can assume that the optimal policy is indifferent between serving a class 1 versus a class 2 customer in state (i,j) and that it is optimal to serve a class 2 customer in state $(i-1,j)$ at epoch m. Under these assumptions, a little algebra shows that for this case, the terms involving μ_1, and μ_2 equal 0.

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

Subcase F.32: Suppose it is optimal to serve a class 1 customer in states $(i, j-1)$ and $(i-1, j+1)$ and optimal to serve a class 2 customer in state $(i, j+1)$ at epoch m.

In this case, a little algebra yields that the terms involving μ_1, and μ_2 are bounded below by

$$\mu_1\Big[\,(\mu_2-\mu_1)\,[v_{m,\gamma}(i-1,j+1)-2v_{m,\gamma}(i-1,j)+v_{m,\gamma}(i-1,j-1)]$$
$$+\,\mu_1[v_{m,\gamma}(i-1,j-1)+v_{m,\gamma}(i-2,j+1)-v_{m,\gamma}(i-1,j)-v_{m,\gamma}(i-2,j)]\Big],$$

and the inductive hypothesis implies that this last expression is bounded below by 0.

Subcase F.33: Suppose it is optimal to serve a class 1 customer in states $(i, j+1)$ and $(i, j-1)$ and optimal to serve a class 2 customer in state $(i-1, j+1)$ at epoch $m+1$.

In this case, a little algebra yields that the terms involving μ_1, and μ_2 are bounded below by

$$(\mu_2-\mu_1)\Big[\,(\mu_2-\mu_1)\,[v_{m,\gamma}(i,j+1)-2v_{m,\gamma}(i,j)+v_{m,\gamma}(i,j-1)]$$
$$+\,\mu_1[v_{m,\gamma}(i,j-1)+v_{m,\gamma}(i-1,j+1)-v_{m,\gamma}(i,j)-v_{m,\gamma}(i-1,j)]\Big],$$

and the inductive hypothesis implies that this last expression is bounded below by 0.

Case F.34: $i = N$.

Terms that involve λ on the expression on the left-hand side of the inequality in 1 are given by

$$\mu_1\frac{\lambda(N-1)}{N}\Big[v_{m,\gamma}(i,j+1)+v_{m,\gamma}(i-1,j)-v_{m,\gamma}(i,j)-v_{m,\gamma}(i-1,j+1)\Big],$$

whereas the terms with coefficients λ on the expression on the right-hand side of the inequality in 1 are given by

$$\mu_2\lambda\Big[v_{m,\gamma}(i,j+1)-2v_{m,\gamma}(i,j)+v_{m,\gamma}(i,j-1)\Big].$$

The inductive hypothesis implies the result for these terms. The terms with coefficients τ and the terms with coefficients μ_1 and μ_2 for this case are the same as in the previous case (when $i < N$), and hence, are non-negative. This completes the proof of 1.

It follows from these arguments that inequality 1 holds for each $m \geq 0$, and hence (by letting $m \to \infty$) for v_γ. The result for the discounted cost problem follows by letting $N \to \infty$. By letting $\alpha = 0$, the same proof above gives the result for the average cost optimality equations.

Appendix G: Proof of Proposition 5.2

The result is shown in the discounted cost case. The average case is similar. Assume $h_1 \geq \cdots \geq h_n$. Fix $x \in \mathbb{X}$ and start two processes on the same probability space. Process 1 assigns the server to class i customers with $x_i > 0$ while Process 2 assigns it to class $j > i$ with $x_j > 0$. After the first event, both processes follow an optimal policy. We show that the expected cost of Process 2, is at least as high as that of Process 1. In what follows, discounting is suppressed without any loss of generality.

If the first event is an arrival, both processes enter state $x + e_1$, we relabel this as the new starting state and restart the argument. Similarly, if there is a change in status seen by the two processes, then they also enter the same state, and we relabel the starting states and continue. Now, with probability $\mu/(\lambda + x_1\tau_1 + \cdots + x_{n-1}\tau_{n-1} + 2\mu)$ a class i customer completes service first. Process 1 sees the service and Process 2 does not. The difference in the (expected) costs is

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

$\mu/(\lambda + x_1\tau_1 + \cdots + x_{n-1}\tau_{n-1} + 2\mu)\big(v_\alpha(x - e_i) - v_\alpha(x)\big)$. If a class j customers completes service first the difference is $\mu/(\lambda + x_1\tau_1 + \cdots + x_{n-1}\tau_{n-1} + 2\mu)\big(v_\alpha(x) - v_\alpha(x - e_j)\big)$. A little algebra yields that the result holds provided

$$v_\alpha(x - e_i) - v_\alpha(x - e_j) \leq 0. \tag{G.1}$$

We continue to follow the sample paths of two processes (on the same probability space) as shown in G.1 via a sample path argument. Processes 1 and 2 begin in states $x - e_i$ and $x - e_j$, respectively. Process 2 uses a stationary optimal policy, which we denote by π_2. In what follows, we construct a (potentially sub-optimal) policy for Process 1, which we denote by π_1, so that

$$v_\alpha^{\pi_1}(x - e_i) - v_\alpha^{\pi_2}(x - e_j) \leq 0. \tag{G.2}$$

Since π_1 is potentially sub-optimal, (G.1) follows from (G.2).

Note that starting from (G.2) the (immediate) costs at the next event are $(h_j - h_i)t_1 \leq 0$, where t_1 is the time of the next event and the inequality is due to the assumption that $h_j \leq h_i$. Moreover, if the relative position (as measured by the current states) of the two processes at the next event remains the same, then we may relabel the initial states and continue from the beginning of the argument. This occurs when any of the uncontrolled events occur that are seen by the two processes (i.e., arrivals and changes of status seen by both processes). It also occurs when π_2 serves a class k customer (by letting π_1 also serve a class k customer provided there is at least one such customer) and the next event is a class k customer service completion. Consider now the other cases.

Case G.35: Customer status change.

Suppose the next event is a change in status in Process 2 only (with probability $\tau_i/(\lambda + x_1\tau_1 + \cdots + x_{n-1}\tau_{n-1} + 2\mu)$), after which both processes follow optimal controls. If $j = i + 1$, then the two processes couple and the remaining costs on the left-hand side of the inequality in (G.2) are $\big(h_j - h_i\big)\, t_1 \leq 0$. Otherwise, the remaining costs (with the probability of the event suppressed) are

$$\big(h_j - h_i\big)\, t_1 + [v_\alpha(x - e_i) - v_\alpha(x - e_i + e_{i+1} - e_j)].$$

Note that if the expression inside the brackets above is non-positive, then it is optimal to prioritize class $i + 1$ customers over class j customers. In other words, we may replace the role of class i customers with that of class $i + 1$ customers and restart the argument.

Suppose the next event is a change in status in Process 1 only (with probability $\tau_j/(\lambda + x_1\tau_1 + \cdots + x_{n-1}\tau_{n-1} + 2\mu)$), after which both processes follow optimal controls. In this case, the remaining costs on the left-hand side of the inequality in (G.2) (with the probability of the event suppressed) are

$$\big(h_j - h_i\big)\, t_1 + [v_\alpha(x - e_i - e_j + e_{j+1}) - v_\alpha(x - e_j)].$$

Note if that the expression inside the brackets above is non-positive, then it is optimal to prioritize class i customers over class $j + 1$ customers. In other words, we may replace the role of class j customers with that of class $j + 1$ customers and restart the argument.

Case G.36: Customer service completions.

Suppose that $x_i = 1$ and that policy π_2 has the server work at station i. In this case, let π_1 have the server work at station j. If the first event is a class i service completion in Process 2

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

(with probability $\mu/(\lambda + x_1\tau_1 + \cdots + x_{n-1}\tau_{n-1} + 2\mu)$), after which both processes follow optimal controls, then the remaining costs on the left-hand side of the inequality in (G.2) are

$$(h_j - h_i)t_1 + v_\alpha(x - e_i) - v_\alpha(x - e_i - e_j). \tag{G.3}$$

If the first event is a class j service completion in Process 1 (with probability $\mu/(\lambda + x_1\tau_1 + \cdots + x_{n-1}\tau_{n-1} + 2\mu)$), after which optimal controls are used, then the remaining costs are

$$(h_j - h_i)t_1 + v_\alpha(x - e_i - e_j) - v_\alpha(x - e_j). \tag{G.4}$$

Since both events occur with the same probability, we can add expressions (G.3) and (G.4) to obtain that the remaining costs for this case are

$$2(h_j - h_i)t_1 + [v_\alpha(x - e_i) - v_\alpha(x - e_j)]. \tag{G.5}$$

Note that expression inside the second pair of brackets in (G.5) above is implied by the left-hand side of inequality (G.2). That is, we may restart the argument from here.

The result now follows by (possibly) relabeling the states or relabeling the classes and continuing until the two processes couple. In the latter case, the remaining rewards on the left-hand side are non-positive. The case with $h_1 \leq \cdots \leq h_n$ can be similarly shown. The proof is omitted for the sake of brevity.

Appendix H: Tables

See Tables H.1–H.4.

TABLE **H.1.** Long-run average holding costs when $n = 3$. Here, $\mu_1 = 20$, $h_1 = 1$, $\mu_2 = 5$, $h_2 = 15$, and $\mu_3 = 15$, $h_3 = 10$ so that C-μ, μ, and h yield distinct priority rules.

τ_1	τ_2	λ	Optimal	C-μ	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	μ	h
0.2	0.2	2.5	0.23	0.23	0.23	0.23	0.23	0.24	0.23
0.2	0.2	3.5	0.35	0.35	0.35	0.35	0.35	0.37	0.35
0.2	0.2	4.5	0.49	0.49	0.49	0.49	0.5	0.52	0.5
0.4	0.2	2.5	0.33	0.33	0.33	0.33	0.33	0.33	0.33
0.4	0.2	3.5	0.51	0.52	0.52	0.52	0.52	0.52	0.52
0.4	0.2	4.5	0.75	0.8	0.79	0.78	0.76	0.76	0.82
0.2	0.4	2.5	0.22	0.22	0.22	0.22	0.23	0.23	0.23
0.2	0.4	3.5	0.34	0.34	0.34	0.34	0.35	0.36	0.34
0.2	0.4	4.5	0.49	0.49	0.49	0.49	0.49	0.51	0.49
0.4	0.4	2.5	0.32	0.32	0.32	0.32	0.32	0.33	0.32
0.4	0.4	3.5	0.5	0.51	0.51	0.5	0.5	0.51	0.51

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

TABLE **H.2.** Long-run average holding costs when $n = 3$. Here, $\mu_1 = 15$, $h_1 = 10$, $\mu_2 = 4$, $h_2 = 15$, and $\mu_3 = 20$, $h_3 = 1$ so that C-μ, μ, and h yield distinct priority rules.

τ_1	τ_2	λ	Optimal	C-μ	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	μ	h
0.2	0.2	2.5	2.14	2.14	2.14	2.14	2.14	2.14	2.22
0.2	0.2	3.5	3.27	3.27	3.27	3.27	3.28	3.28	3.51
0.2	0.2	4.5	4.65	4.65	4.65	4.65	4.65	4.65	5.25
0.4	0.2	2.5	2.27	2.27	2.27	2.27	2.27	2.28	2.5
0.4	0.2	3.5	3.51	3.51	3.51	3.51	3.51	3.51	4.22
0.4	0.2	4.5	5.02	5.02	5.02	5.02	5.02	5.03	30.91
0.2	0.4	2.5	2.13	2.13	2.13	2.13	2.13	2.13	2.21
0.2	0.4	3.5	3.25	3.25	3.25	3.25	3.25	3.26	3.47
0.2	0.4	4.5	4.61	4.61	4.61	4.61	4.61	4.62	5.12
0.4	0.4	2.5	2.25	2.25	2.25	2.25	2.25	2.26	2.44
0.4	0.4	3.5	3.46	3.46	3.46	3.46	3.47	3.47	4.03
0.4	0.4	4.5	4.94	4.94	4.94	4.94	4.94	4.95	24.78

TABLE **H.3.** Long-run average holding costs when $n = 3$. Here, $\mu_1 = 2$, $h_1 = 150$, $\mu_2 = 3$, $h_2 = 120$, and $\mu_3 = 4$, $h_3 = 100$ so that Cμ and μ are identical (and only report Cμ), but h yields a priority rule distinct from Cμ.

τ_1	τ_2	λ	Optimal	C-μ and μ	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	h
0.2	0.2	2.5	1312.77	1312.77	1312.77	1312.77	1312.77	6760.38
0.2	0.2	3.5	3088.8	3088.8	3088.8	3088.8	3088.8	7199.98
0.2	0.2	4.5	4014.34	4014.34	4014.34	4014.34	4014.34	7279.99
0.4	0.2	2.5	849.83	849.83	849.83	849.83	849.83	6733.92
0.4	0.2	3.5	2827.21	2827.21	2827.21	2827.21	2827.21	7199.98
0.4	0.2	4.5	5258.71	5258.71	5258.71	5258.71	5258.71	7279.99
0.2	0.4	2.5	1238.26	1246.85	1246.85	1246.85	1246.85	6759.83
0.2	0.4	3.5	2801.46	2801.63	2801.63	2801.63	2801.63	7199.98
0.2	0.4	4.5	3633.97	3633.97	3633.97	3633.97	3633.97	7279.99
0.4	0.4	2.5	790.46	790.46	790.46	790.46	790.46	6732.08
0.4	0.4	3.5	2290.22	2290.22	2290.22	2290.22	2290.22	7199.98
0.4	0.4	4.5	6086.91	6086.91	6086.91	6086.91	6086.91	7279.99

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press

TABLE **H.4.** Long-run average holding costs when $n = 3$. Here, $\mu_1 = 4$, $h_1 = 100$, $\mu_2 = 3$, $h_2 = 120$, and $\mu_3 = 2$, $h_3 = 150$ so that Cμ and μ are identical (and only report Cμ), but h yields a priority rule distinct from Cμ.

τ_1	τ_2	λ	Optimal	C-μ and μ	Mod-Cμ(15)	Mod-Cμ(10)	Mod-Cμ(5)	h
0.2	0.2	2.5	205.04	205.04	205.04	205.04	205.04	6223.18
0.2	0.2	3.5	1744.77	1745.54	1745.54	1745.54	1745.54	6973.21
0.2	0.2	4.5	6553.33	6685.68	6685.68	6685.68	6685.68	7030.59
0.4	0.2	2.5	231.42	231.42	231.42	231.42	231.42	6792.99
0.4	0.2	3.5	2280.41	2280.43	2280.43	2280.43	2280.43	7071.09
0.4	0.2	4.5	6672.96	6697.99	6697.99	6697.99	6697.99	7124.74
0.2	0.4	2.5	213.2	213.34	213.34	213.34	213.34	6703.39
0.2	0.4	3.5	1877.53	1879.72	1879.72	1879.72	1879.72	7083.38
0.2	0.4	4.5	6566.7	6685.8	6685.8	6685.8	6685.8	7140.74
0.4	0.4	2.5	246.05	246.05	246.05	246.05	246.05	6901.37
0.4	0.4	3.5	2460.69	2460.71	2460.71	2460.71	2460.71	7174.59
0.4	0.4	4.5	6684.83	6698.07	6698.07	6698.07	6698.07	7228.35

https://doi.org/10.1017/S0269964817000213 Published online by Cambridge University Press