Workload-dependent dynamic priority for the multiclass queue with reneging*

Rami Atar Anat Lev-Ari

Department of Electrical Engineering
Technion–Israel Institute of Technology
Haifa 32000, Israel

March 6, 2017

Abstract

Scheduling control for a single-server queue with I customer classes and reneging is considered, with linear holding or reneging cost. An asymptotically optimal policy in heavy traffic is identified where classes are prioritized according to a workload-dependent dynamic index rule. Denote by c_i, μ_i and θ_i, $i \in \mathcal{I} := \{1, \dots, I\}$ the queue length cost, service rate and reneging rate, for class-i customers. Then a relabeling of the classes and a partition $0 = w_0 < w_1 < \cdots < w_K = \infty$, $K \le I$, are identified such that the policy acts to always assign least priority to the class i when the rescaled workload is in the interval $[w_{i-1}, w_i)$. The relabeling is such that when workload is withing the lowest [resp., highest] interval $[w_{i-1}, w_i)$, the least priority class is the one with smallest $c\mu$ [resp., greatest θ] value. This result stands in sharp contrast to known fluid scale results where it is asymptotically optimal to prioritize by the fixed $c\mu/\theta$ index. One of the technical challenges is the discontinuity of the limiting queue length process under optimality. Specifically, the limit process is of the form $\psi(\tilde{X}_t)$, $t \ge 0$, where $\tilde{X}_t$ is a one-dimensional diffusion, $\psi : \mathbb{R}_+ \to \mathbb{R}_+^I$ being a piecewise continuous map with set of discontinuities $\{w_i : i = 1, \dots, K-1\}$.

AMS subject classifications: 60F17, 60J60, 60K25, 93E20

Keywords: Multi-class single-server queue, queues with abandonment, diffusion control problem, Bellman equation, dynamic priorities, state space collapse

1 Introduction

In this paper we address asymptotically optimal (AO) scheduling control for the multiclass single-server queue with abandonment in heavy traffic. We show that, in sharp contrast to the behavior of the model in the many-server regime [13], [6], the model is governed by a one-dimensional Brownian control problem (BCP). Whereas in the model without abandonment the well-known $c\mu$ index rule is AO, the index obtained in this paper depends dynamically on the total workload in the system. We give a complete characterization of this dynamic index in terms of the underlying Bellman equation, and prove that it is AO. Under this policy, the queue length process is asymptotic to a process that lies on a piecewise continuous curve, where discontinuities correspond to the workload levels where the index changes.

*Research supported in part by the ISF (Grant 1315/12)

The model considered has a fixed number, I, of classes, where each customer arrives into an infinite buffer dedicated to its class, and may abandon the queue while waiting to be served. The arrival and potential service processes form renewal processes, while abandonment occurs according to an exponential clock. We seek to minimize an infinite horizon discounted cost function associated with abandonment count, that can also be rephrased in terms of queue length. Let b_i, μ_i and θ_i, $i \in \mathcal{I} := \{1, \dots, I\}$ denote cost per reneging, service rate and reneging rate, for class-i customers. The corresponding queue length cost of class-i is $c_i = \theta_i b_i$. Then the policy that is found by solving the BCP, and is shown to be AO in heavy traffic, can be described as follows. For $i \in \mathcal{I}$, let $\varphi_i(y) = \theta_i y - c_i\mu_i$, $y \in \mathbb{R}_+$, and let $\mathcal{K}$ be a minimal subset of $\mathcal{I}$ for which $\max_{\mathcal{K}} \varphi_i(y) = \max_{\mathcal{I}} \varphi_i(y)$ for all $y \in \mathbb{R}_+$. Assume without loss of generality that the classes are relabeled so that $\mathcal{K} = \{1, \dots, K\}$, $K = |\mathcal{K}|$, and $c_i\mu_i$ are increasing in $i \in \mathcal{K}$. Then automatically, also θ_i, $i \in \mathcal{K}$ are increasing. The policy assigns least priority to the class $i \in \mathcal{K}$ when the rescaled workload is in $[w_{i-1}, w_i)$, for a fixed partition $0 = w_0 < w_1 < \cdots < w_K = \infty$ that is determined in terms of a suitable Bellman equation. Thus, in particular, when workload is low, the class that is assigned least priority is the one for which the $c\mu$ value is the smallest. This may be explained by noting that the policy mimics the well-known $c\mu$ rule, that is known to be optimal in the absence of abandonment. Moreover, when workload is high, the class with greatest θ value is assigned the least priority. Then one can say that the policy chooses to reduce workload in the system by activating the abandonment at its greatest possible rate.

The problem has been considered under other parametric regimes. In the many-server heavy traffic asymptotics, known as the Halfin-Whitt regime, an analogous problem has been considered and AO policies have been obtained in [13] and [6] via the analysis of the Hamilton-Jacobi-Bellman equation. However, in contrast to the results found in this paper, the AO policy does not seem to have an explicitly described structure. At the other extreme are the results regarding the many-server fluid limit regime, [3] and [4] that address an explicit, fixed index, namely the $c\mu/\theta$ index, and show that it is AO to prioritize according to it. Notice that the parameters $c_i\mu_i$ and θ_i dictate both this index and the dynamic index of this paper, but in different ways.

Control of the multiclass queue in presence of abandonment has also been studied in the single-server fluid regime, in [18], for the case of two classes. The policy that is found there to be optimal varies dynamically between the $c\mu/\theta$ and the $c\mu$ rules, the transition being determined by a certain switching curve in the queue length space.

The works that are most closely related to this paper are [17], [1], [19] and [12]. The first three do not address AO, but solve the limiting BCP associated with their model in heavy traffic. The first two consider multiclass single-server queue for broader sets of models than that of this paper. Specifically, [17] studies general (as opposed to exponential) abandonment clock, and obtains a dynamic index rule. The structure of this rule is explicit but more complicated than the one obtained with exponential clock, and, in particular, depends on the full abandonment clock distribution. The paper [1] assumes exponential abandonment, but addresses nonlinear cost. It provides a solution in terms of a dynamic index that is expressed by a suitable Bellman equation and the underlying cost function. The third paper, [19], addresses a make-to-order parallel service system with exponential abandonment, long run average cost, and a system manager that may outsource jobs. The optimal allocation rule and the decision to outsource jobs are dynamic, determined by the Bellman equation and depend on the workload level.

A treatment that includes both solving a BCP and establishing AO of the resulting policy is provided in [12], for a related model. The model, often referred to as an *N-system*, has two classes of customers, and two servers that have different capabilities. One of the servers can serve only customers from one of the classes, while the other can serve both. The AO policy acts with two threshold levels. First, the dual-skill server is allowed to 'help' the single-skill server when the number of customers of the class that can be served by both exceeds a threshold. An additional threshold is set on the workload of the dual-skill server, that determines how to prioritize the two classes depending on whether workload exceeds this threshold. The latter aspect of the threshold policy is precisely the one that is generalized in this paper to a multiple number of classes, whereas the additional server aspect is not present in our model.

The model under consideration deserves, so we feel, a treatment separate from the one provided for the N-system in [12] for two main reasons. First, our treatment aims at understanding how the structure of the policy and its AO differ from and extend the case of two classes. Indeed, the structure of the AO described above is not transparent when only two classes are present. Second, we provide an argument for AO which relies on different ideas than those presented in [12], and results in a considerably shorter proof. Moreover, the approach taken in this paper allows us to assume the minimal possible moment assumptions regarding the service time distributions, namely the existence of finite second moments, instead of exponential moments as in [12].

Our proof consists of analyzing the BCP (in Section 3) and showing that its value serves both as an asymptotic lower and upper bound on that of the QCP. The proof of the lower bound (provided in Section 4) relies to some extent on results of [8]. The central part of our contribution is that of the upper bound (given in Section 5). Its idea relies on two crucial ingredients. One is a state-space collapse (SSC) result (Lemma 5.1), which states that within each workload interval $[w_{k-1}, w_k)$, where the index policy is fixed, the multidimensional queue length process lies on a curve determined by the index. The second (Lemma 5.3) is an argument showing that the time spent by the workload process near the points of discontinuities, w_k, is small, by which it follows that the cost incurred at times when the queue length is not in a neighborhood of one of the continuous parts of the curve, can be neglected. The combination of the two yields Theorem 5.1. Along with the results of Section 4, this establishes AO.

We use the following notation. For $x, y \in \mathbb{R}^k$ (k a positive integer), $x \cdot y$ and $\|x\|$ denote the usual scalar product and ℓ_2 norm, respectively. Vectors are regarded column vectors. The transpose of a vector or matrix x is denoted by x'. The standard basis of $\mathbb{R}^k$ is denoted by $\{e_i, 1 \le i \le k\}$. With $\mathbb{R}_+ = [0, \infty)$, for $f : \mathbb{R}_+ \to \mathbb{R}^k$, we let $\|f\|_T = \sup_{t\in[0,T]} \|f(t)\|$, and denote by $|f|_T$ the total variation of f over $[0, T]$. For a Polish space $\mathcal{S}$, we let $\mathbb{C}_{\mathcal{S}}([0, T])$ and $\mathbb{D}_{\mathcal{S}}([0, T])$ denote the set of continuous and, respectively, RCLL functions $[0, T] \to \mathcal{S}$. We endow $\mathbb{D}_{\mathcal{S}} = \mathbb{D}_{\mathcal{S}}(\mathbb{R}_+)$ with the Skorohod J_1 topology. We write $X_n \Rightarrow X$ for convergence in distribution. We use notation such as $X(t)$ and X_t interchangeably, for stochastic processes X and $t \in \mathbb{R}_+$.

For a positive integer k, $m \in \mathbb{R}^k$ and a symmetric, positive matrix $A \in \mathbb{R}^{k\times k}$, *an* (m, A)*-Brownian motion* (BM) is a k-dimensional BM starting from zero, having infinitesimal drift and covariance coefficients m and A, respectively.

The organization of the paper is as follows. The next section introduces the model and main results. In Section 3 is devoted to the study of the BCP, where most of the proofs are deferred to the appendix. Sections 4 and 5 address the lower bound and upper bound, respectively.

2 Model and main results

2.1 The queueing control problem and its scaling

In the queueing control problem (QCP) under consideration, customers of I different classes arrive at the system to receive service from a single server. Customers that cannot be served immediately upon arrival are queued in buffers with infinite room, one dedicated to each class. The server may serve only the customers that wait at the head of the line, and is capable of sharing its effort among them, i.e., work simultaneously on (at most) I jobs of different classes. Since our goal is to study asymptotics, we consider a sequence of systems indexed by $n \in \mathbb{N}$. All random variables (RVs) and stochastic processes introduced below will be defined on a probability space $(\Omega, \mathcal{F}, P)$. Expectation w.r.t. P is denoted by E.

The arrivals follow renewal processes, and the service times are i.i.d. More precisely, let I sequences of i.i.d. positive RVs be given, representing inter-arrival times of the I classes, denoted by $\{IA_i(l), l \in \mathbb{N}\}_{i\in\mathcal{I}}$, with $E[IA_i(1)] = 1$ and $\sigma^2_{i,IA} = \mathrm{var}(IA_i(1)) < \infty$. Set $\mathcal{I} = \{1, 2, \ldots, I\}$. For $i \in \mathcal{I}$, $n \in \mathbb{N}$ let

$\lambda_i^n > 0$ be the reciprocal mean inter-arrival times of class-i customers in the n-th system. Then the number of arrivals of class-i customers in the n-th system up to time t is given by

$$A_i^n(t) = A_i(\lambda_i^n t), \qquad \text{where} \qquad A_i(t) = \sup\Big\{l \geq 0 : \sum_{k=1}^{l} IA_i(k) \leq t\Big\}, \quad t \geq 0.$$

To set up the model for service, let I sequences of i.i.d. positive RVs $\{ST_i(l), l \in \mathbb{N}\}_{i\in\mathcal{I}}$ be given, with $E[ST_i(1)] = 1$ and $\sigma_{i,ST}^2 = \text{var}(ST_i(1)) < \infty$. For $i \in \mathcal{I}$, $n \in \mathbb{N}$, let $\mu_i^n > 0$ be the reciprocal mean service time of class-i customers in the n-th system. Then the number of customers of class i that complete their service by the time the server has devoted to this class t units of time is given by

$$S_i^n(t) = S_i(\mu_i^n t), \qquad \text{where} \qquad S_i(t) = \sup\Big\{l \geq 0 : \sum_{k=1}^{l} ST_i(k) \leq t\Big\}, \quad t \geq 0.$$

Note that the above is different from the actual number of customers to complete service by time t. To introduce the latter, define

$$\mathcal{S} = \Big\{\beta \in \mathbb{R}_+^I : \sum_{i\in\mathcal{I}} \beta_i \leq 1\Big\}$$

and let $\{B_t^n\}$ be a process taking values in $\mathcal{S}$, where its i-th component $B_i^n(t)$ represents the fraction of effort dedicated to class i (recall that processor sharing is allowed). Then

$$T_i^n(t) = \int_0^t B_i^n(s)ds \tag{1}$$

gives the time the server dedicates to class-i customers by time t, and the number of class-i job completions by t is then

$$D_i^n(t) = S_i^n(T_i^n(t)). \tag{2}$$

Let $X_i^n(t)$ denote the number of class-i customers present in the system at time t and $Q_i^n(t)$ the corresponding number in the queue (i.e., not being served). Then

$$Q_i^n(t) = X_i^n(t) - 1_{\{B_i^n(t)>0\}}. \tag{3}$$

Next, to model exponential abandonment, let $\{R_i^0, i \in \mathcal{I}\}$ be I standard Poisson processes, and assume that the number of class-i abandonments up to time t is given by

$$R_i^n(t) = R_i^0\Big(\theta_i^n \int_0^t Q_i^n(s)ds\Big), \tag{4}$$

where $\theta_i^n > 0$ is the abandonment rate for class-i customers. We can now write the balance equation as

$$\begin{aligned} X_i^n(t) &= X_i^n(0) + A_i^n(t) - D_i^n(t) - R_i^n(t) \\ &= X_i^n(0) + A_i^n(t) - S_i^n(T_i^n(t)) - R_i^0\Big(\theta_i^n \int_0^t Q_i^n(s)ds\Big). \end{aligned} \tag{5}$$

We will be interested only in service policies that are non-idling. Therefore we shall require that, for each t,

$$\sum_i X_i^n(t) > 0 \text{ implies } \sum_i B_i^n(t) = 1. \tag{6}$$

Since a class-i job can only be processed when there are such jobs in the system, we also require that, for each t,

$$X_i^n(t) = 0 \text{ implies } B_i^n(t) = 0. \tag{7}$$

For each n, the $1+3I$ objects consisting of the initial condition $X_0^n = (X_i^n(0))$, the I sequences $\{IA_i\}$, the I sequences $\{ST_i\}$ and the I processes $\{R_i^0\}$ are assumed mutually independent. We refer to this tuple as the *stochastic primitives*. The stochastic processes A_i have, by construction, RCLL paths; it is assumed that so do S_i and R_i^0.

We now introduce the diffusion scaling and the heavy traffic condition. First, we assume that there are constants $\lambda_i, \mu_i, \theta_i \in (0,\infty)$ and $\hat{\lambda}_i, \hat{\mu}_i \in \mathbb{R}$ such that, as $n \to \infty$,

$$\frac{\lambda_i^n}{n} \to \lambda_i, \qquad \frac{\mu_i^n}{n} \to \mu_i, \qquad \theta_i^n \to \theta_i,$$

$$\hat{\lambda}_i^n = \frac{\lambda_i^n - n\lambda_i}{\sqrt{n}} \to \hat{\lambda}_i, \qquad \hat{\mu}_i^n = \frac{\mu_i^n - n\mu_i}{\sqrt{n}} \to \hat{\mu}_i.$$

We assume that the system is in heavy traffic in the sense that $\sum_{i=1}^I \rho_i = 1$ where $\rho_i = \lambda_i/\mu_i$. We denote the scaled headcount process and the scaled reneging count by

$$\hat{X}_i^n(t) = n^{-1/2} X_i^n(t), \qquad \hat{R}_i^n(t) = n^{-1/2} R_i^n(t). \tag{8}$$

It is assumed throughout that the scaled initial condition $\hat{X}_0^n = n^{-1/2} X_0^n$ converges in L^2, as $n \to \infty$, to a deterministic vector $x_0 \in \mathbb{R}_+^I$. We regard B_t^n as the *control* of the n-th system. We are interested only in control processes that are determined via observations from the past events of the system. To this end we formulate the following.

Definition 2.1. (Admissible control, QCP) *Fix n. A process B^n taking values in $\mathcal{S}$, having RCLL sample paths, is called an* **admissible control** *for the n-th system if the following hold. Let the processes D^n, Q^n, R^n, X^n be defined by the primitives (X_0^n, A^n, S^n, R_0^n) and the control B^n via equations* (2), (3), (4), (5). *Then*

i. B^n is adapted to the filtration $\sigma\{A_i^n(s), D_i^n(s), R_i^n(s), i \in \mathcal{I}, s \le t\}$;

ii. P-a.s., for all t and i, $X_i^n(t) \ge 0$, and (6) *and* (7) *hold.*

Denote the class of all admissible controls for the n-th system by $\mathcal{B}^n$. Given the primitives and a control, call the corresponding tuple (D^n, Q^n, R^n, X^n) the **controlled processes**.

We are interested in a cost that accounts for the abandonment count in diffusion scale. Fix $\alpha > 0$ and $b \in (0,\infty)^I$ and let the cost function for the n-th system be given by

$$\hat{J}^n(B^n) = E\Big(\int_0^\infty e^{-\alpha t} b' d\hat{R}_t^n\Big), \tag{9}$$

where $\hat{R}^n$ is a scaled version of R^n via (8), and R^n is a component (i.e., the third) of the controlled process corresponding to B^n. The value of the QCP is defined as

$$\hat{V}^n = \inf_{B^n \in \mathcal{B}^n} \hat{J}^n(B^n). \tag{10}$$

Clearly, $\hat{J}^n$ and $\hat{V}^n$ depend on the initial data X_0^n, but we consider the initial data as part of the model's primitives and therefore do not specify this dependence explicitly in the notation.

Remark 2.1. *As can be seen in equation* (25) *on Section 3, the cost can be translated into a linear cost of the queue length. Thus, the analysis of this paper can be applied to a more general cost structure, which includes linear queue length costs in addition to linear abandonment costs.*

2.2 Structure of the AO policy

Our goal is to study the asymptotic behavior of the QCP as $n \to \infty$. We do so by (i) linking it to a BCP that can be fully analyzed, and (ii) constructing an AO control for the QCP. While the details of the BCP and its solution are lengthy (and provided in Section 3), the structure of the proposed AO policy is simple, and can be presented with little additional notation. We thus turn to the description of the policy.

The policy is based on dynamic prioritization. It allows **no processor sharing**, in the sense that only one customer is served at each given time; service is **non-interruptible**, namely once a job starts being served it is served until completion; and it is **non-idling**, in the sense that the server operates (at full capacity) whenever there is at least one customer in the system. (Note that in contrast to the restrictions above on the proposed policy, Definition 2.1 considers policies that allow for processor sharing and service interruptions, and so the QCP is concerned with this broader class of policies.) In order to identify the policy one only needs to specify which class the server admits into service each time it becomes free (note that any new arrival into an empty system is immediately served).

To this end we will need a notion of priority. Let $(i_1, i_2, \ldots, i_I)$ be a permutation of the set of classes $\mathcal{I}$. Suppose that at time t the server becomes free and there are customers in the system. By saying that the next customer to be served is picked **according to the ordering** $(i_1, \ldots, i_I)$ we mean that one selects the minimal k such that at least one i_k-class customer is present (that is, $X_t^{i_k,n} > 0$) and admits into service the first customer in line from that class.

We denote $m_i = 1/\mu_i$ and call $m = (m_i)_{i\in\mathcal{I}}$ the workload vector. We call

$$\tilde{X}^n = m'\hat{X}^n \tag{11}$$

the **scaled workload process** (as a rule we use '$\hat{\ }$' and '$\tilde{\ }$' to denote I-dimensional and, respectively, 1-dimensional diffusion scaled processes). The policy we are interested in prioritizes the classes based on the current value of $\tilde{X}^n$.

Definition 2.2. (cyclic dynamic priority) *Let a non-empty subset $\mathcal{K}$ of $\mathcal{I}$ be given along with a partition $\{L_k\}_{k\in\mathcal{K}}$ of $[0,\infty)$ consisting of intervals L_k of the form $[a,b)$. Here $a \in [0,\infty)$ and $b \in (0,\infty]$. The* **cyclic dynamic priority policy** *associated with this partition, denoted by $\mathbf{P}(\{L_k\}_{k\in\mathcal{K}})$, is a non-interruptible, non-idling policy that allows no processor sharing, and acts as follows.*

When, at time t, the server becomes available and some of the buffers are non-empty, the current value of the scaled workload $\tilde{X}_t^n$ is computed via (11) and the **least priority class** *k is set to be the unique i for which $\tilde{X}_t^n \in L_i$. The customer to be served is then picked according to the ordering*

$$(k+1, k+2, \ldots, I, 1, 2, \ldots, k).$$

Example 2.1. *In this example we show how priorities are determined given the data $\mathcal{K}$ and $\{L_k\}$. Suppose $\mathcal{I} = \{1,2,3,4\}$, $\mathcal{K} = \{1,2,3\}$ and $L_1 = [0,1)$, $L_2 = [1,5)$ and $L_3 = [5,\infty)$. Then when the scaled workload is below level 1, priority is assigned according to the ordering $(2,3,4,1)$ (with 2 being highest and 1 lowest priority). Similarly, when the workload is between 1 and 5 [resp., exceeds 5] priority is assigned according to $(3,4,1,2)$ [resp., $(4,1,2,3)$].*

To describe the specific choice of partition to be considered, let $\varphi_i : \mathbb{R}_+ \to \mathbb{R}$ denote the affine map

$$\varphi_i(y) = \theta_i y - c_i\mu_i, \qquad y \in \mathbb{R}_+, \tag{12}$$

for each $i \in \mathcal{I}$, where, throughout, $c_i = \theta_i b_i$. Let $F(y) = \max_i \varphi_i(y)$.

To specify the parameters of the policy, we introduce the following equation that serves as the *Bellman equation* corresponding to a stochastic control problem to be introduced in Section 3, namely

$$-\frac{\tilde{\sigma}^2}{2}\frac{d^2v}{dx^2} - \tilde{y}\frac{dv}{dx} + xF\Big(\frac{dv}{dx}\Big) + \alpha v = 0, \qquad 0 < x < \infty, \tag{13}$$

where

$$\tilde{y} = \sum_i m_i(\hat{\lambda}_i - \rho_i \hat{\mu}_i), \qquad \tilde{\sigma} = \Big(\sum_i m_i^2 \lambda_i (\sigma_{i,IA}^2 + \sigma_{i,ST}^2)\Big)^{1/2}. \tag{14}$$

When considered with the boundary condition $\frac{dv}{dx}(0) = 0$ and the growth condition $|v(x)| \le C(1+x)^C$, $x \in [0,\infty)$ (some constant C), equation (13) has a unique classical solution, denoted throughout by v (see Proposition 3.1). Moreover, v is a convex function (ibid.). (The precise form of the Bellman equation is linked to the underlying dynamics of a certain 1-dimensional controlled process, to be presented below in Section 3.2. Specifically, the Neumann boundary condition at zero owes to the fact that it is constrained to lie in $\mathbb{R}_+$).

Let $\mathcal{K}$ be a minimal subset of $\mathcal{I}$ such that $\max_{k\in\mathcal{K}} \varphi_k(y) = \max_{i\in\mathcal{I}} \varphi_i(y)$ for all y, where the term 'minimal' means that every strict subset of $\mathcal{K}$ does not have this property. For concreteness, a construction of such a set appears below in Remark 2.2(a).

Let $K = |\mathcal{K}|$. Assume, without loss of generality, that the classes $i \in \mathcal{I}$ are labeled in such a way that $\mathcal{K} = \{1,\dots,K\}$, and

$$c_1\mu_1 \le c_2\mu_2 \le \cdots \le c_K\mu_K. \tag{15}$$

Next, for $k \in \mathcal{K}$, let L'_k be the unique interval $[a_k, b_k)$ where φ_k maximizes φ_i over all i. It is easy to see that $\{L'_k\}$ form a partition of $\mathbb{R}_+$. Since v is convex and satisfies the boundary condition at zero, it follows that $\frac{dv}{dx}$ is a nondecreasing function from $\mathbb{R}_+$ to itself, starting at zero. As a result, there exists a partition $\{L_k\}$ of $\mathbb{R}_+$ such that L_k is of the form $[w_{k-1}, w_k)$, with $0 = w_0 < w_1 < \cdots < w_K = \infty$, and is the inverse image of L'_k under $\frac{dv}{dx}$

$$\frac{dv}{dx}(y) \in L'_k \text{ iff } y \in L_k.$$

Let the resulting cyclic dynamic priority policy $\mathbf{P}(\{L_k\}_{k\in\mathcal{K}})$ be denoted by $\mathbf{P}^*$, and let the control process B^n obtained under this policy be denoted by $B^{n,*}$. The main result of this paper is that the policy thus defined is AO.

Example 2.2. *Here we show how the functions $\{\varphi_i\}$, $i \in \mathcal{I}$ determine the intervals $\{L'_k\}$, $k \in \mathcal{K}$. Consider the functions φ_i, $i = 1,2,3,4$ defined in Figure 1. Then $\mathcal{I} = \{1,2,3,4\}$, and we can see that $\mathcal{K} = \{1,2,3\}$. Moreover, $L'_1 = [0,1)$, $L'_2 = [1,3)$ and $L'_3 = [3,\infty)$. Note that $c_1\mu_1 = 1 < c_2\mu_2 = 3 < c_3\mu_3 = 9$ and $\theta_1 = 1 < \theta_2 = 3 < \theta_3 = 5$. Class 1 [resp., 2, 3] is assigned least priority when $\frac{dv}{dx}(\tilde{X}^n)$ lies in L'_1 [resp., L'_2, L'_3]. Class 4 is never assigned the least priority.*

Remark 2.2. **(a)** *A construction of a set $\mathcal{K}$ with the aforementioned property is as follows. Fix a subset $\mathcal{I}' \subset \mathcal{I}$ such that*

- *$\mathcal{I}'$ has no two distinct members i and j for which the functions φ_i and φ_j are identical, and*
- *the collection of functions $\{\varphi_i, i \in \mathcal{I}'\}$ equals $\{\varphi_i, i \in \mathcal{I}\}$.*

Let $\mathcal{K}$ consist of all $k \in \mathcal{I}'$ for which $\varphi_k(y) > \max\{\varphi_i(y) : i \in \mathcal{I}', i \ne k\}$ for some y.
(b) *Although no constraints have been put on the parameters c_i, μ_i and θ_i, $i \in \mathcal{I}$, it is interesting to note that not only the parameters $\{c_k\mu_k\}$, $k \in \mathcal{K}$ are ordered in increasing order (by the requirement (15)), but so are the parameters $\{\theta_k\}$, $k \in \mathcal{K}$. (This is not difficult to see, by (12) and the specific choice of the index set $\mathcal{K}$. Indeed, if $c_i\mu_i < c_j\mu_j$ for some $i, j \in \mathcal{K}$, then the definition of $\mathcal{K}$ implies that $\varphi_j(y) > \varphi_i(y)$ for all large y, hence $\theta_j > \theta_i$). In addition, the definition of the index set $\mathcal{K}$ implies that the index $c_1\mu_1$ is the smallest among all $\{c_i\mu_i\}_{i\in\mathcal{I}}$, and θ_K is the largest among all $\{\theta_i\}_{i\in\mathcal{I}}$ (as can be seen by taking $y = 0$ and $y \to \infty$ in the definition of $F(y)$). This observation has the following interpretation. When workload is low (specifically, lies in the lowest interval among the L_k's), the policy assigns least priority to the class that has least $c\mu$ value. When it is high, it assigns least priority to the class that has greatest θ value. A possible intuitive explanation of this is as follows. When workload is low, so is the total*

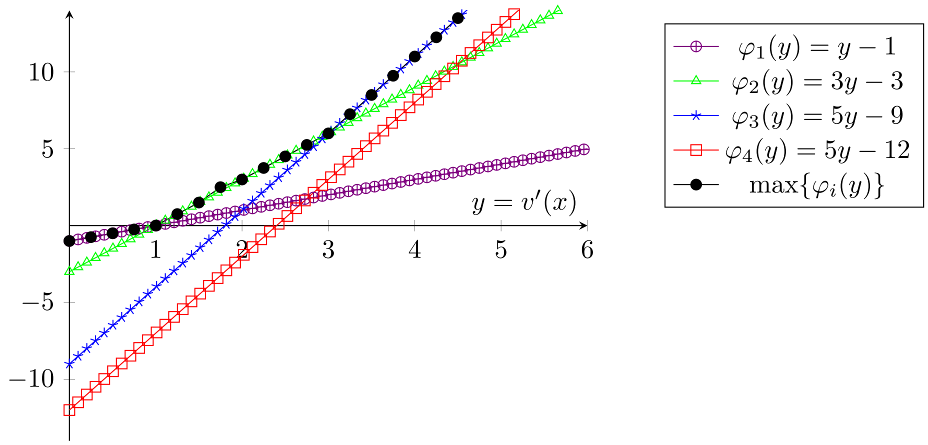

Figure 1: *The functions φ_i and their maximum. The maximizing classes are 1, 2 and 3.*

abandonment rate, and the policy mimics the one that is optimal in absence of abandonment, namely the usual $c\mu$ rule, that leaves the least priority to the class with the least $c\mu$ value. When workload is very large, the abandonment rate is significant, and the policy may then utilize abandonment to reduce the workload by activating the largest possible rate of abandonment.

(c) *An interesting question is whether there should be a connection to the $c\mu/\theta$ rule, that is known to be AO in a different parametric regime [3, 4, 5]. Specifically, a fixed priority rule that prioritizes according to the index $c_i\mu_i/\theta_i$, in a system with a finite number of classes, reneging customers and many servers, is AO at fluid scale. The main differences between the setting of the aforementioned papers and this paper, apart from the different scaling regimes, are that an ergodic cost is considered in [3, 4, 5] while here we treat a discounted cost, and the deterministic control problem one obtains under fluid limits versus the stochasticity governing the limit model (the BCP) that one has in the present setting. This might lead one to conjecture that focusing on a small discount parameter α, and letting the Brownian motion diffusion coefficient in the BCP tend to zero (thus making the problem look more like a deterministic fluid control problem) the optimal solution obtained here converges to the $c\mu/\theta$ rule. We leave this question open.*

We can now state the main result.

Theorem 2.1. *Let v denote the unique classical solution of the Bellman equation (13) with the boundary condition $\frac{dv}{dx}(0) = 0$ and the growth condition $|v(x)| \le C(1+x)^C$, $x \in [0,\infty)$ (for some constant C). Recall that $x_0 \in \mathbb{R}_+^I$ is the limit of the scaled initial conditions $\hat{X}_0^n$. Then*
i. The limit value is determined by the function v. In particular,

$$\lim_{n\to\infty} \hat{V}^n = v(m'x_0).$$

ii. The policy $\mathbf{P}^*$ is AO, that is,

$$\lim_{n\to\infty} \hat{J}^n(B^{n,*}) = v(m'x_0).$$

Remark 2.3. *The only characteristic of the cyclic dynamic priority policy used in the proof is the way it selects the least priority class, specifically, when $\tilde{X}_t^n \in L_k$, class k gets the least priority. We have restricted our attention to this particular policy for concreteness, but, in fact, Theorem 2.1 is valid for any dynamic priority policy that respects the above rule, regardless of how it treats the classes of higher priority. There is therefore some room for flexibility in choosing the policy.*

3 The BCP

In this section we derive the BCP from the model equations by taking formal limits. While this control problem is obtained heuristically, the rigorous justification of the claim that it governs the asymptotics of the QCP will be a consequence of the results, established in later sections. The proofs of this section's results are standard, but we do provide them in the appendix, for completeness. We first obtain an I-dimensional BCP, then transform it into a 1-dimensional one that we call a reduced BCP (RBCP), relate the latter to the Bellman equation (13), and finally show some elementary properties of the solution to (13).

3.1 Derivation of the BCP

In addition to the vectors $m = (m_i) = (\mu_i^{-1})$ introduced above, we will use the notation

$$\Theta = \mathrm{diag}(\theta), \qquad M = \mathrm{diag}(m), \qquad q = M^{-1}\Theta b, \qquad y_i = \hat{\lambda}_i - \rho_i\hat{\mu}_i. \tag{16}$$

Note that, with this notation, $c = \Theta b$. Also, let

$$\Theta^n = \mathrm{diag}(\theta^n), \qquad c^n = \Theta^n b, \qquad y_i^n = \hat{\lambda}_i^n - \rho_i\hat{\mu}_i^n. \tag{17}$$

Recall the scaled versions of the processes X^n and R^n introduced before. Additional scaled processes are as follows:

$$\hat{A}_i^n(t) = \frac{A_i^n(t) - \lambda_i^n t}{\sqrt{n}}, \qquad \hat{S}_i^n(t) = \frac{S_i^n(t) - \mu_i^n t}{\sqrt{n}}, \qquad \hat{D}_i^n(t) = \hat{S}_i^n(T_i^n(t)),$$
$$\hat{Q}_i^n(t) = \frac{Q_i^n(t)}{\sqrt{n}}, \qquad \hat{Y}_i^n(t) = \frac{\mu_i^n}{\sqrt{n}}\big(\rho_i t - T_i^n(t)\big), \tag{18}$$

and

$$\hat{W}_i^n(t) = y_i^n t + \hat{A}_i^n(t) - \hat{S}_i^n(T_i^n(t)). \tag{19}$$

Diving by $\sqrt{n}$ in (5) and using the above definitions we obtain, with

$$e_i^{(1),n}(t) = \frac{R_i^0\Big(\theta_i^n \int_0^t Q_i^n(s)ds\Big) - \theta_i^n \int_0^t Q_i^n(s)ds}{\sqrt{n}}, \tag{20}$$

$$e_i^{(2),n}(t) = \int_0^t (\theta_i^n \hat{X}_i^n(s) - \theta_i \hat{Q}_i^n(s))ds, \tag{21}$$

$e^n = -e^{(1),n} + e^{(2),n}$, the equation

$$\hat{X}_t^n = \hat{X}_0^n + \hat{W}_t^n - \int_0^t \Theta \hat{X}_s^n ds + \hat{Y}_t^n + e_t^n. \tag{22}$$

Note that our assumption $\sum_1^I \rho_i = 1$, along with the property $\sum_{i\in\mathcal{I}} B_i^n \le 1$, imply that

$$Y_t^{\#,n} := \sum_{i\in\mathcal{I}} (\mu_i^n)^{-1}\hat{Y}_i^n(t) \tag{23}$$

is a non-negative, non-decreasing process, for each n.

Taking limits in (22), y^n tends to y. Moreover, by the central limit theorem for renewal processes, $(\hat{A}_i^n(t), \hat{S}_i^n(t))$ converges to a pair of independent BMs with drift zero and diffusion coefficients $\sqrt{\lambda_i}\sigma_{i,IA}$ and $\sqrt{\mu_i}\sigma_{i,ST}$, resp. (see details in the proof of Theorem 4.1). The term e^n vanishes in the limit (ibid.). As far as $\hat{X}_0^n$ is concerned, we denoted its limit by x_0, but in this section we write the initial condition for the dynamics as $x \in \mathbb{R}_+^I$. Recall (16) and denote $y = (y_i)$. Denote also $\sigma = \text{diag}(\sigma_i)$, $\sigma_i = \lambda_i^{1/2}(\sigma_{i,IA}^2 + \sigma_{i,ST}^2)^{1/2}$. We obtain

$$X_t = x + W_t - \int_0^t \Theta X_s ds + Y_t, \tag{24}$$

where $\{W_t\}$ is an I-dimensional (y,σ)-BM. A limiting version of the aforementioned property of $\hat{Y}^n$ is that $m'Y$ is non-negative and non-decreasing.

We regard Y as *control* and X as the corresponding *controlled process*. The precise definition is as follows.

Definition 3.1. (Admissible control, BCP) *An admissible control system with initial condition $x \in \mathbb{R}_+^I$ is a tuple $(\bar{\Sigma}, W_t, Y_t, X_t)$, where $\bar{\Sigma} = (\bar{\Omega}, \bar{\mathcal{F}}, \{\bar{\mathcal{F}}_t\}, \bar{P})$ is a filtered probability space, $\{W_t\}$ is an I-dimensional $\bar{\mathcal{F}}_t$-adapted (y,σ)-BM, the processes $\{X_t\}$ and $\{Y_t\}$ have sample paths in $\mathbb{D}_{\mathbb{R}^I}(\mathbb{R}_+)$ and are $\bar{\mathcal{F}}_t$-adapted, and the following hold:*

i. For all $t, s \geq 0$, $W_{t+s} - W_t$ is independent of $\bar{\mathcal{F}}_t$ under $\bar{P}$,

ii. X_t defined in (24) *satisfies $X_t \in \mathbb{R}_+^I$ for all t $\bar{P}$-a.s.,*

iii. The process $m'Y$ is non-negative and non-decreasing.

We write $\mathcal{A}(x)$ for the class of admissible controls for the initial condition x. When we write $Y \in \mathcal{A}(x)$ it will be understood that this process carries with it a filtered probability space and the processes W and X. Moreover, with a slight abuse of notation, we write E for the expectation corresponding to this probability space.

For a limit version of the cost $\hat{J}^n$ defined in (9), use (4) to replace $\hat{R}_i$ by $\theta_i \int_0^\cdot \hat{X}_i^n(t)dt$ to arrive at

$$J(x,Y) = E\Big(\int_0^\infty e^{-\alpha t} c' X_t dt\Big). \tag{25}$$

The BCP is concerned with minimizing J over all admissible controls. The value function for the BCP is thus given by

$$V(x) = \inf_{Y \in \mathcal{A}(x)} J(x,Y), \qquad x \in \mathbb{R}_+^I. \tag{26}$$

3.2 The 1-dimensional problem and the Bellman equation

The problem can be reduced into a control problem with 1-dimensional dynamics by projecting the controlled process onto the workload vector m, that is a SSC property. Denoting $\tilde{x} = m'x$, $\tilde{W} = m'W$ (by which $\tilde{W}$ is a $(\tilde{y},\tilde{\sigma})$-BM, where we recall the notation $(\tilde{y},\tilde{\sigma})$ from (14) and note that $\tilde{y} = m'y$ and $\tilde{\sigma} = (\sum_i m_i^2 \sigma_i^2)^{1/2}$), $\tilde{X} = m'X$ and $\tilde{Y} = m'Y$, we have from (24),

$$\tilde{X}_t = \tilde{x} + \tilde{W}_t - \int_0^t m'\Theta X_s ds + \tilde{Y}_t.$$

Let

$$\mathcal{S}_1 = \Big\{\beta \in \mathbb{R}_+^I : \sum_{i \in \mathcal{I}} \beta_i = 1\Big\},$$

and let U be an $\mathcal{S}_1$-valued process given by $MX/m'X$ if $m'X > 0$ and an arbitrary, fixed, element of $\mathcal{S}_1$ otherwise. Then we have

$$X = m'XM^{-1}U = \tilde{X}M^{-1}U, \tag{27}$$

and so we can write an equation involving only $\tilde{X}$ and $(U, \tilde{Y})$, namely

$$\tilde{X}_t = \tilde{x} + \tilde{W}_t - \int_0^t \theta' U_s \tilde{X}_s ds + \tilde{Y}_t. \tag{28}$$

The above equation will serve as the dynamics for the reduced problem, with $(U, \tilde{Y})$ and $\tilde{X}$ acting as control and controlled processes, respectively.

Since X_i represents the ith queue length in the limit model, and $\tilde{X} = m'X$ is the total workload, the process $U_i = m_i X_i / m'X$ corresponds to the fraction of class-i workload. Thus U describes how workload is distributed among the various classes. In the prelimit model, the distribution of workload among classes is affected by the policy and the system dynamics, but there is no way to determine it directly, and so it is not natural to regard it as a control process. However, it makes perfect sense to consider this process as control in the limit model under consideration.

Note that the cost can also be written in terms of $\tilde{X}$ and U only, as

$$\tilde{J}(\tilde{x}, (U, \tilde{Y})) = E\Big(\int_0^\infty e^{-\alpha t} q' U_t \tilde{X}_t dt \Big). \tag{29}$$

Definition 3.2. (Admissible control, RBCP) *An admissible control system with initial condition $\tilde{x} \in \mathbb{R}_+$ is a tuple $(\tilde{\Sigma}, \tilde{W}, U, \tilde{Y}, \tilde{X})$, where $\tilde{\Sigma} = (\tilde{\Omega}, \tilde{\mathcal{F}}, \{\tilde{\mathcal{F}}_t\}, \tilde{P})$ is a filtered probability space, $\{\tilde{W}_t\}$ is a 1-dimensional $\tilde{\mathcal{F}}_t$-adapted $(\tilde{y}, \tilde{\sigma})$-BM, and the processes $\{U_t\}$, $\{\tilde{Y}_t\}$, $\{\tilde{X}_t\}$ have sample paths in $\mathbb{D}_{\mathcal{S}_1}(\mathbb{R}_+)$, $\mathbb{D}_{\mathbb{R}^I}(\mathbb{R}_+)$ and $\mathbb{D}_{\mathbb{R}}(\mathbb{R}_+)$, resp., are $\tilde{\mathcal{F}}_t$-adapted and the following hold:*

i. For all $t, s \geq 0$, $\tilde{W}_{t+s} - \tilde{W}_t$ is independent of $\tilde{\mathcal{F}}_t$ under $\tilde{P}$,

ii. $\tilde{X}_t$ defined in (28) satisfies $\tilde{X}_t \geq 0$ for all t $\tilde{P}$-a.s.,

iii. The process $\tilde{Y}$ is non-negative and non-decreasing.

We write $\tilde{\mathcal{A}}(\tilde{x})$ for the class of admissible controls for the initial condition $\tilde{x}$. The value function of the RBCP is defined as

$$\tilde{V}(\tilde{x}) = \inf_{(U,\tilde{Y}) \in \tilde{\mathcal{A}}(\tilde{x})} \tilde{J}(\tilde{x}, (U, \tilde{Y})), \qquad \tilde{x} \in \mathbb{R}_+. \tag{30}$$

Remark 3.1. *Note that any admissible control system for which*

$$\int_{[0,\infty)} 1_{\{\tilde{X}_t > 0\}} d\tilde{Y}_t = 0, \tag{31}$$

satisfies the bound $\tilde{X}_t \leq \tilde{x} + 2\|\tilde{W}\|_t$, a fact that will be used in the sequel. To see this, note that if, for $t > 0$, $\tilde{X}_t > 0$ then, with $\sigma = \sup\{s < t : \tilde{X}_s = 0\}$, we have on $\{\sigma = -\infty\}$, $\tilde{X}_t \leq \tilde{x} + \tilde{W}_t$, and on $\{\sigma \in [0, t]\}$, using (28),

$$\tilde{X}_t = \tilde{X}_t - \tilde{X}_\sigma \leq \tilde{W}_t - \tilde{W}_\sigma + \tilde{Y}_t - \tilde{Y}_\sigma = \tilde{W}_t - \tilde{W}_\sigma \leq 2\|\tilde{W}\|_t.$$

This observation is used in the sequel (specifically, in Lemma 4.2).

It is shown in the appendix (Lemma A.1) that the BCP and RBCP are equivalent in the sense that $V(x) = \tilde{V}(m'x)$, $x \in \mathbb{R}_+^I$. Analyzing the BCP via the RBCP has the advantage that the associated Bellman equation is one-dimensional, namely it is an ODE. The following result links RBCP value function to that equation.

For $x \in [0, \infty)$ let $i(x)$ denote the index $i \in \mathcal{K}$ for which $x \in L_i$. The result below is concerned with the stochastic differential equation (SDE) for a pair of processes $(\tilde{X}, \tilde{Y})$, where both processes take values in $\mathbb{R}_+$ with $\tilde{Y}$ non-decreasing, taking the form

$$\tilde{X}_t = \tilde{x} + \tilde{W}_t - \int_0^t \theta_{i(\tilde{X}_s)} \tilde{X}_s ds + \tilde{Y}_t, \qquad \int_{[0,\infty)} 1_{\{\tilde{X}_t > 0\}} d\tilde{Y}_t = 0. \tag{32}$$

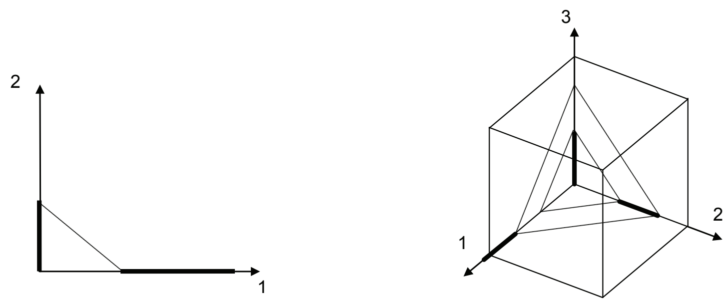

Figure 2: *As workload level increases from zero, the target queue length varies along a discontinuous curve. The axes labeled* 1, 2 *(and* 3*) correspond to queue lengths of the same labels, and the curve is shown in thick line. The discontinuities correspond to workload levels at which the priority index switches.*

Here, $\tilde{W}$ is as in Definition 3.2.

The form of the Bellman equation (13) is well known; see Section IV.5 of [11]. The following proposition states the existence of a unique C^2 solution and its properties.

Proposition 3.1. *Consider equation* (13) *for* $v \in C^2(\mathbb{R}_+)$ *with the boundary condition* $\frac{dv}{dx}(0) = 0$ *and the growth condition* $|v(x)| \le C(1+x)^C$*, for some constant* C*. Then*
i. There exists a unique solution to this equation, denoted by v*.*
ii. One has $\tilde{V} = v$*, where* $\tilde{V}$ *is defined in* (30)*.*
iii. Weak existence and uniqueness hold for solutions $(\tilde{X}, \tilde{Y})$ *to equation* (32)*. Setting* $U_t = e_{i(\tilde{X}_t)}$*,* $t \ge 0$*, (where we recall that* $\{e_i\}$ *denote the standard basis) gives rise to a solution of the RBCP, namely*

$$\tilde{J}(\tilde{x}, (U, \tilde{Y})) = \tilde{V}(\tilde{x}).$$

iv. v *is a convex function.*

This result is proved in the appendix. A solution to the BCP can be obtained from that of the RBCP, provided above, by appealing to Lemma A.1(ii) in the appendix.

Remark 3.2. *The combination of the relation* $U = e_{i(\tilde{X})}$ *stated in Proposition 3.1(iii) and* (27) *shows how the (multidimensional) state process is reconstructed from the (one-dimensional) workload process, namely* $X = \tilde{X}\mu^{-1}_{i(\tilde{X})}e_{i(\tilde{X})}$*. Thus* X *evolves along the piecewise continuous curve* $(y, \psi(y))$*,* $y \in \mathbb{R}_+$*, where* $\psi(y) = y\mu^{-1}_{i(y)}e_{i(y)}$*,* $y \in \mathbb{R}_+$*. The points of discontinuities of* ψ *correspond to levels of workload* y *where the priority index* $i(y)$ *changes. Examples of possible* $(y, \psi(y))$ *curves are depicted in Figure 2, for dimensions 2 and 3.*

4 Lower bound

In this section we prove that the function v serves as a lower bound on the limit inferior of the cost functions, providing a first step towards the proof of Theorem 2.1.

Theorem 4.1. *For any sequence of admissible controls* $B^n \in \mathcal{B}^n$*,* $\liminf_{n\to\infty} \hat{J}^n(B^n) \ge V(x_0)$*.*

The lemmas below will be used in the proof. For any locally integrable function φ denote

$$\mathfrak{I}\varphi = \int_0^{\cdot} \varphi(t)dt.$$

Recall the definition (9) of $\hat{J}^n$. For $B^n \in \mathcal{B}^n$ let

$$\hat{J}^{(1),n}(B^n) = E\Big(\int_0^\infty e^{-\alpha t}(c^n)'\hat{Q}^n_t dt\Big), \qquad \hat{J}^{(2),n} = \alpha E\Big(\int_0^\infty e^{-\alpha t}(c^n)'\Im\hat{Q}^n_t dt\Big). \tag{33}$$

Lemma 4.1. *For $n \in \mathbb{N}$ and $B^n \in \mathcal{B}^n$ one has $\hat{J}^n(B^n) = \hat{J}^{(1),n}(B^n) = \hat{J}^{(2),n}(B^n)$.*

Proof. The first identity is a consequence of the martingale property of

$$R^0_i\Big(\theta^n_i \int_0^t Q^n_i(s)ds\Big) - \theta^n_i \int_0^t Q^n_i(s)ds,$$

that follows by an argument provided in the appendix of [7]. The second identity is achieved by integration by parts, and the fact that, for n fixed, $e^{-\alpha t}\Im\hat{X}^n_i(t)$ converge to 0 a.s. as $t \to \infty$. From equation (5),

$$\hat{Q}^n_t \le \hat{X}^n_t \le \hat{X}^n_0 + \frac{1}{\sqrt{n}}A^n_t,$$

where $\hat{X}^n_0$ converge to a deterministic vector x and $A^n_i(t)$ is a renewal process with finite expectation λ^n_i and therefore $A^n_i(t)/t \to 1/\lambda^n_i$ as $t \to \infty$. This implies the convergence of $e^{-\alpha t}\Im\hat{X}^n_t$ to 0 a.s. as required. □

The next lemma is concerned with a variation of the definition of the cost for the BCP. It is based on a definition of admissible control systems that is broader than that given by Definition 3.1. Namely, define $\bar{\mathcal{A}}(x)$ as in Definition 3.1 with the exception that the requirement that $\{X_t\}$ and $\{Y_t\}$ have sample paths in $\mathbb{D}_{\mathbb{R}^I}(\mathbb{R}_+)$ is replaced by the requirement that these $\mathbb{R}^I$-valued processes are progressively measurable. For brevity we write this as $Y \in \bar{\mathcal{A}}(x)$. Note that $\mathcal{A}(x) \subseteq \bar{\mathcal{A}}(x)$.

This extended class of controls is introduced for the following reason. Our technique is based on tightness of the processes $\Im\hat{X}^n$ and $\Im\hat{Y}^n$ rather than $\hat{X}^n$ and $\hat{Y}^n$. Limits of these processes are to be proved to have Lipschitz continuous, hence a.e. differentiable, sample paths. Now, the derivatives of these processes, that are candidates for the processes X and Y from the BCP, need not be RCLL, and so the class of controls $\mathcal{A}(x)$ is too small for this purpose. Using instead the class $\bar{\mathcal{A}}(x)$ is possible thanks to an argument used in [8] (as mentioned in the proof of Theorem 4.1 below) by which progressively measurable a.e. derivatives always exist.

Lemma 4.2. *Given $Y \in \bar{\mathcal{A}}(x)$ let*

$$\bar{J}(x, Y) = \alpha E\Big(\int_0^\infty e^{-\alpha t}c'\Im X(t)dt\Big),$$

and $\bar{V}(x) = \inf_{Y\in\bar{\mathcal{A}}(x)} \bar{J}(x, Y)$. Then $V(x) = \bar{V}(x)$.

Proof. For any $Y \in \mathcal{A}(x)$, integration by parts and positivity of $c'X$ give $\bar{J}(x, Y) \le J(x, Y)$. Along with the inclusion $\mathcal{A}(x) \subseteq \bar{\mathcal{A}}(x)$ this gives $\bar{V}(x) \le V(x)$.

To prove the reverse inequality let $\varepsilon > 0$. Consider an admissible control Y_ε and the corresponding processes $(X_\varepsilon, W_\varepsilon)$ in $\bar{\mathcal{A}}(x)$, such that $\bar{J}(x, Y_\varepsilon) \le \bar{V}(x) + \varepsilon$. Let $T > 0$ and define the tuple $(\bar{X}, \bar{Y}, \bar{W})$ to equal $(X_\varepsilon, Y_\varepsilon, W_\varepsilon)$ on $[0, T)$. Note by (24) (that is part of Definition 3.1) that the limit $m'\bar{X}(T-)$ is well-defined. Let $(U, \tilde{Y})$ be an arbitrary admissible control for the RBCP, for which (31) holds, starting from $\tilde{X}(T) = m'\bar{X}(T-)$, and set (X, Y, W) to be the corresponding admissible control for the BCP given by the transformation that appears in the second part of Lemma A.1. Define $(\bar{X}, \bar{Y}, \bar{W})$ on $[T, \infty)$ to equal (X, Y, W). Then, $(\bar{X}, \bar{Y}, \bar{W}) \in \bar{\mathcal{A}}(x)$ and by Remark 3.1, $m'\bar{X}_t \le m'x + 2\|m'W\|_t$, for $t \ge T$. Thus, for some constant c_1,

$$\bar{J}(x, \bar{Y}) - \bar{J}(x, Y_\varepsilon) = \alpha E\Big(\int_T^\infty e^{-\alpha t}c'(\Im\bar{X}(t) - \Im X_\varepsilon(t))dt\Big) \le c_1 E\int_T^\infty e^{-\alpha t}(1 + \|W\|_t)dt = \kappa(T),$$

and $\kappa(T) \to 0$ as $T \to \infty$. Thus,

$$\bar{J}(x, \bar{Y}) \le \bar{J}(x, Y_\varepsilon) + \kappa(T) \le \bar{V}(x) + \varepsilon + \kappa(T).$$

To complete the proof it suffices to show that $V(x) \le \bar{J}(x, \bar{Y})$. Note, however, that $\bar{Y} \in \bar{\mathcal{A}}(x)$ may not be an element of $\mathcal{A}(x)$. We use the reduced BCP to overcome this problem in the following way: $Y_\varepsilon \in \bar{\mathcal{A}}(x)$ implies that $m'Y_\varepsilon$ is non-decreasing. Thus, $\tilde{Y} = \lim_{s\downarrow t} m'Y_\varepsilon(s)$ has RCLL sample paths and so the same is true for $m'\tilde{X}_t = m'x + m'W_t - \int_0^t m'\Theta\tilde{X}_s ds + \tilde{Y}_t$. Denote $\tilde{x} = m'x$ and $\tilde{\mathcal{U}}_i(s) = (m'\tilde{X}_s)^{-1} m_i \tilde{X}_i(s)$. Then, $(U, \tilde{Y}) \in \tilde{\mathcal{A}}(\tilde{x})$ and Lemma A.1 gives $V(x) = \tilde{V}(\tilde{x}) \le \tilde{J}(\tilde{x}, (U, \tilde{Y})) = \bar{J}(x, Y_\varepsilon)$. Taking $\varepsilon \to 0$ and $T \to \infty$ shows the claim. This completes the proof. □

Recall $\tilde{X}^n$ defined in (11) and define $\tilde{e}^n_t = m'e^n_t$, $\tilde{W}^n_t = m'\hat{W}^n_t$, $H^n_t = m'\Theta\Im\hat{X}^n$ and $\tilde{Y}^n_t = m'\hat{Y}^n_t$. Then by (22),

$$\tilde{X}^n_t = \tilde{X}^n_0 + \tilde{W}^n_t - H^n_t + \tilde{Y}^n_t + \tilde{e}^n_t. \tag{34}$$

A sequence of processes with sample paths in $\mathbb{D}_{\mathbb{R}^k}$ is said to be *C-tight* if it is tight and, in addition, any subsequential limit has, with probability 1, continuous sample paths. A useful characterization of C-tightness is as follows (see Proposition VI.3.26 of [15]): *C-tightness of X^n is equivalent to*

C1. The sequence of RVs $\|X^n\|_T$ is tight for every fixed $T < \infty$, and

C2. For every $T < \infty$, $\varepsilon > 0$ and $\eta > 0$ there exist N and $\delta > 0$ such that

$$n \ge N \text{ implies } P(w_T(X^n, \delta) > \eta) < \varepsilon,$$

where

$$w_T(f, \delta) = \sup_{0\le s<u\le s+\delta\le T} \|f_u - f_s\|.$$

Lemma 4.3. *i. For each T, the sequence $\|\hat{W}^n\|_T \vee \|\hat{X}^n\|_T \vee \|\hat{Y}^n\|_T$ is tight, and e^n converge to zero in probability.*
ii. The sequence $(\tilde{X}^n, \tilde{W}^n, H^n, \tilde{Y}^n)$ is C-tight. Moreover, let (X^1, W^1, H^1, Y^1) be a subsequential limit. Then, with $Z_t := \tilde{x} + W^1_t - H^1_t$,

$$X^1_t = Z_t + \sup_{s\le t}(-Z_s)^+, \quad Y^1_t = \sup_{s\le t}(-Z_s)^+. \tag{35}$$

Proof. By the central limit theorem for renewal processes (see [10], Section 17), the processes $(\hat{A}^n_t, \hat{S}^n_t)$ converge in law to a pair of mutually independent I-dimensional BMs with zero drift and diffusion coefficients $\mathrm{diag}(\sqrt{\lambda_i}\sigma_{i,IA})$ and $\mathrm{diag}(\sqrt{\mu_i}\sigma_{i,ST})$, respectively. Recalling the definition (19) of $\hat{W}^n$ and the fact that $0 \le \frac{d}{dt}T^n_i(t) \le 1$, it follows that $\{\hat{W}^n\}$ is a C-tight sequence of processes and thus so is $\tilde{W}^n$.

Let T be fixed. Then, immediately we have tightness of the RVs $\|\hat{W}^n\|_T$. The proof that $\|\hat{X}^n\|_T \vee \|\hat{Y}^n\|_T$ are tight is attained is several steps. The first step relies on (34) and the property that, for every n, $Y^{\#,n}$ (defined in (23)) satisfies

$$\int 1_{\{\tilde{X}^n>0\}} dY^{\#,n} = 0 \quad \text{a.s.} \tag{36}$$

Indeed, by (1), (18), (23) and the fact $\sum_i \rho_i = 1$, we have $\frac{d}{dt}Y^{\#,n} = n^{-1/2}(1 - \sum_i B^n_i)$, and therefore (36) follows from (6) (that is valid due to Definition 2.1(2)).

Let $e^{(3),n} = \tilde{Y}^n - Y^{\#,n}$ and $e^{\#,n} = \tilde{e}^n + e^{(3),n}$. Then (34) gives

$$\tilde{X}^n_t = \tilde{X}^n_0 + \tilde{W}^n_t - H^n_t + Y^{\#,n}_t + e^{\#,n}_t. \tag{37}$$

Now, equation (37) along with the facts that $\tilde{X}^n$ is non-negative, $Y^{\#,n}$ is non-negative and non-decreasing, and (36) holds, implies that $(\tilde{X}^n, Y^{\#,n})$ solves the 1-dimensional Skorohod problem for the data $Z^n := \tilde{X}_0^n + \tilde{W}^n - H^n + e^{\#,n}$. In particular,

$$\tilde{X}_t^n = Z_t^n + \sup_{s\le t}\big(-Z_s^n\big)^+, \qquad Y_t^{\#,n} = \sup_{s\le t}\big(-Z_s^n\big)^+, \qquad \text{for all } t\ge 0. \tag{38}$$

As a result, $|\tilde{X}_t^n| \le 2\|Z^n\|_t$. Since m_i are positive constants and $\hat{X}_i^n \ge 0$, it follows that, for some constants c_1, c_2 (all constants c_1, c_2, etc. introduced below are independent of n and t),

$$M_t^n := \|\hat{X}^n\|_t \le c_1\|\tilde{X}^n\|_t \le c_2(\|\tilde{X}_0^n\| + \|\tilde{W}^n\|_t + \|e^{\#,n}\|_t) + c_2\int_0^t M_s^n ds. \tag{39}$$

In the next step we analyze the error terms $e^{(i),n}$ for $i = 1,2,3$. First, by (3), $\|\hat{Q}^n\|_t \le M_t^n$. Moreover, by (20),

$$\|e^{(1),n}\|_t \le \|\bar{R}_0^n\|_{c_3 t M_t^n}, \tag{40}$$

where c_3 is an upper bound on $\|\theta_n\|$, and

$$\bar{R}_0^n(t) := n^{-1/2}(R^0(n^{1/2}t) - n^{1/2}t).$$

Recalling that R_i^0 is a standard Poisson process, it follows from the law of iterated logarithm that for every $\varepsilon > 0$ there exists a random, finite constant $\kappa = \kappa(\varepsilon, i, \omega)$, such that $|R_i^0(t) - t| \le \kappa + \varepsilon t$ for all $t \in \mathbb{R}_+$, w.p.1. Hence, for each fixed $\varepsilon > 0$, as $n \to \infty$,

$$P\{\text{there exists } t \in [0,\infty) \text{ such that } |\bar{R}_0^n(t)| > \varepsilon(1+t)\} \to 0.$$

Thus with $\varepsilon > 0$ fixed, there exists a sequence of events Ω^n with $P(\Omega^n) \to 1$, such that, on Ω^n, one has for all t, $\|e^{(1),n}\|_t \le \varepsilon(1 + c_3 t M_t^n)$. We choose ε so small that $\varepsilon < 1$ and $c_2\varepsilon c_3 T < 1/2$. We get by (39), for all $t \le T$, on Ω^n,

$$\frac{1}{2}M_t^n \le c_2(1 + \|\tilde{X}_0^n\| + \|\tilde{W}^n\|_t + \|e^{(2),n}\|_t + \|e^{(3),n}\|_t) + c_2\int_0^t M_s^n ds. \tag{41}$$

Next, by (3) and (21), $|e_t^{(2),n}| \le tM_t^n\|\theta_n - \theta\| + \|\theta\|n^{-1/2}t$. By (23) and the definition of $e^{(3),n}$,

$$|e_t^{(3),n}| \le \max_i \frac{|\mu_i^n - \mu_i|}{\mu_i^n \mu_i}\|\hat{Y}_t^n\| =: \varepsilon_n\|\hat{Y}_t^n\|.$$

Using (22), there exists a constant c_4 (that depends on T) such that for $t \le T$,

$$\|\hat{Y}_t^n\| \le c_4(\|\hat{X}^n\|_t + \|\hat{W}^n\|_t + \|e^n\|_t). \tag{42}$$

Going back to (41), assuming n is so large that the bound we have just obtained for $e^{(2),n}$ gives, for $t \le T$, $|e_t^{(2),n}| \le M_t^n/(4c_2) + 1$, we have

$$\frac{1}{4}M_t^n \le c_2[2 + \|\tilde{X}_0^n\| + \|\tilde{W}^n\|_t + \varepsilon_n c_4(M_t^n + \|\hat{W}^n\|_t + \|e^n\|_t)] + c_2\int_0^t M_s^n ds.$$

Using once again the bounds used before on $\|e^{(1),n}\|_t$ and $\|e^{(2),n}\|_t$ (note that e^n depends on $e^{(1),n}$ and $e^{(2),n}$ but not on $e^{(3),n}$) and noting that $\varepsilon_n \to 0$, we finally obtain for all large n, all $t \le T$, on Ω^n,

$$M_t^n \le c_5(1 + \|\tilde{X}_0^n\| + \|\hat{W}^n\|_t) + c_5\int_0^t M_s^n ds,$$

which, by Gronwall's lemma gives $M^n_T \le c_5(1+\|\tilde X^n_0\|+\|\hat W^n\|_T)e^{c_5 T}$ on the same event. The tightness of the initial conditions and of $\|\hat W^n\|_T$ now gives that of M^n_T.

The last step is to use (42) and the bounds on $\|e^{(1),n}\|_t$ and $\|e^{(2),n}\|_t$ to conclude that for large n, on Ω^n, for $t\le T$ one has

$$\|\hat Y^n\|_T \le c_4(M^n_T+\|\hat W^n\|_T+1+c_6 M^n_T),$$

by which $\|\hat Y^n\|_T$ are tight. The first assertion of part (i) of the lemma is now proved.

Using the bound (40), the tightness on M^n_T and the functional LLN by which $\|\bar R^n_0\|_t \Rightarrow 0$ for any fixed t, gives $\|e^{(1),n}\|_T \Rightarrow 0$. The bound on $e^{(2),n}$ and the tightness of M^n_T also give $\|e^{(2),n}\|_T \Rightarrow 0$. Similarly, the bound on $e^{(3),n}$ and the tightness of $\|\hat Y^n\|_T$ give $\|e^{(3),n}\|_T \Rightarrow 0$. This completes the proof of part (i).

For part (ii), note that, since M^n_T are tight (for T fixed, as before), H^n are automatically C-tight. Having established the convergence of $\|e^{\#,n}\|_T$, it follows that Z^n are also C-tight. The transformation (38) from Z^n to $Y^{\#,n}$ preserves the modulus of continuity, and we conclude that $Y^{\#,n}$ and, in turn, $\tilde X^n$ are also C-tight. The convergence of the error term $e^{(3),n}$ to zero shows that the same is true for $\tilde Y^n$. Finally, using the relation (38), the continuity (in the uniform topology) of the Skorohod map, and again the fact $e^{(3),n}\to 0$, gives (35). □

Proof of Theorem 4.1. With Lemmas 4.1, 4.2 and 4.3 at hand, the proof of the result is almost identical to that of Theorem 3.1 of [8]. In the proof, Lemma 4.3(i) is used to show C-tightness of the sequence $(\mathfrak{I}\hat Q^n, \mathfrak{I}\hat Y^n, \hat W^n)$. Then, given any subsequential limit, which is denoted by $(\mathfrak{I}X, \mathfrak{I}Y, W)$, the components $\mathfrak{I}X, \mathfrak{I}Y$ can be shown to have Lipschitz continuous sample paths with a.e. derivatives that are progressively measurable. Denote the latter by X, Y. As a consequence, $Y\in\bar{\mathcal{A}}(x)$, and Lemmas 4.1, 4.2, together with Fatou's Lemma, gives

$$\liminf_{n\to\infty}\hat J^n(B^n)=\liminf_{n\to\infty}\hat J^{(2),n}(B^n)\ge \bar J(x,Y)\ge \inf_{Y\in\bar{\mathcal{A}}(x)}\bar J(x,Y)=\bar V(x)=V(x).$$

□

5 Upper bound

In this section we prove an upper bound on the limit of the cost functions. Theorem 2.1 is a direct conclusion from Theorem 5.1 below, Theorem 4.1 and Lemma A.1.

Recall that $\mathbf{P}^*$ denotes the dynamic priority $\mathbf{P}(\{L_k\}_{k\in\mathcal{K}})$ and the corresponding control process is denoted by $B^{n,*}$.

Theorem 5.1. *One has*

$$\limsup_{n\to\infty}\hat J^n(B^{n,*})\le V(x_0).$$

A crucial ingredient of the proof is to show that the multi-dimensional normalized queue length process follows closely the curve dictated by the BCP solution (such as in Figure 2), a result often referred to as SSC. Traditionally, this term has been used in the case of a continuous limiting curve. For example, under fixed priority, the curve lies on the axis corresponding to the class of least priority. In this paper the limit curve has discontinuities, but apart from that the phenomenon is similar.

We now explain the main ideas of the proof. It relies on the C-tightness of the normalized workload processes $\tilde X^n$, already established in Section 4, Lemma 4.3. In addition to the scaling parameter n we introduce a parameter $\varepsilon>0$. The latter is used to define intervals in workload space, L^ε_k and $L^{\varepsilon/2}_k$ (k being the index for the class having least priority), and in turn a sequence of random times σ^k_j and τ^k_j, $j=1,2,3,\dots$ constructed in such a way that during each time interval $I^k_j:=[\sigma^k_j,\tau^k_j]$, the workload is within $L^{\varepsilon/2}_k$. Figure 3 below shows how these intervals are constructed in terms of passage times of the levels $w_k\pm\varepsilon$ and $w_k\pm\varepsilon/2$ (the exact details appear in the sequel). This way the priority rule remains fixed within each I_j^k, with k having least priority. Transitions from one such interval to another require the normalized workload process to travel at least $\varepsilon/2$ distance. Significantly, the order of limits is to take $n \to \infty$ first and then $\varepsilon \to 0$. By this one achieves the following. *For each fixed* ε, given any finite time T, the number of time intervals I_j^k that intersect $[0, T]$ (in other words, $\#\{j : \sigma_j^k \le T\}$) does not grow to infinity with $n \to \infty$. Instead, this sequence of random variables, indexed by n, (for fixed ε) is tight, as proved in part (ii) of Lemma 5.1. This is true thanks to the C-tightness alluded to above: a process that is uniformly close to a continuous process cannot cross an $\varepsilon/2$-wide strip too many times in a finite time interval. The significance of having these as tight random variables is that it makes it sufficient to analyze *one excursion*, to $L_k^{\varepsilon/2}$, (such as the jth excursion occurring during the time interval I_j^k) and establish a SSC result as we do in Lemma 5.1(i). Deducing from this the result regarding the complete set of excursions (performed in Lemma 5.1(ii)) is automatic thanks to the tightness of the number of these excursions.

Next, several moment estimates are obtained in Lemma 5.2.

An estimate on the time spent ε-close to the discontinuity set is then established in Lemma 5.3. The convergence of the normalized workload process allows one to reduce the task of estimating that time into a similar question in terms of the limit process. The latter is analyzed by means of Ito's formula applied to a carefully chosen test function. Then this argument is developed, along with an argument of uniqueness of solutions to the underlying SDE, in the remaining part of the proof of Lemma 5.3 to showing the convergence of the prelimit (pre-expectation) cost to that of the RBCP.

These elements are finally combined in the proof of Theorem 5.1.

We now turn to the rigorous construction. Throughout, the superscript '*' is removed from the notation of all processes.

Recall the intervals L_k, $k \in \mathcal{K}$. Denote the boundary points by $w_0, w_1, \dots, w_{K-1}, w_K$ where $w_0 = 0$ and $w_K = \infty$, i.e. $L_k = [w_{k-1}, w_k)$ for $k = 1, \dots, K$. The priority ordering depends on the interval L_k to which $\tilde{X}^n$ belongs. For $0 < \varepsilon < \min_{k=1,\dots,K-1}(w_k - w_{k-1})/3$ we define new intervals

$$L_1^\varepsilon = [0, w_1 - \varepsilon), \quad L_2^\varepsilon = (w_1 + \varepsilon, w_2 - \varepsilon), \dots, \quad L_{K-1}^\varepsilon = (w_{K-2} + \varepsilon, w_{K-1} - \varepsilon), \quad L_K^\varepsilon = (w_{K-1} + \varepsilon, \infty).$$

Denote the union of the intervals L_k^ε by L^ε and its complement in $\mathbb{R}_+$ by M^ε.

For each $k \in \mathcal{K}$, define inductively a sequence of times $\tau_0^k < \sigma_1^k < \tau_1^k < \sigma_2^k < \cdots$ as $\tau_0^k = \varepsilon_1 > 0$ (where ε_1 is a constant) and, for $j = 1, 2, \dots$,

$$\sigma_j^k = \inf\{t > \tau_{j-1}^k : \tilde{X}_t^n \in L_k^\varepsilon\},$$
$$\tau_j^k = \inf\{t > \sigma_j^k : \tilde{X}_t^n \notin L_k^{\varepsilon/2}\},$$

representing the entrance time into L_k^ε and exit time from $L_k^{\varepsilon/2}$ on the j-th visit after time ε_1. This construction is illustrated in Figure 3. Clearly, these random times depend on n, ε_1 and ε, but this is not indicated in the notation, in order to keep it simple. Define also, for $j = 1, 2, \dots$,

$$\nu_j^k = \inf\{t \in [0, \sigma_j^k] : \tilde{X}_s^n \in L_k \text{ and } \tilde{X}_s^n > 0 \text{ for all } s \in [t, \sigma_j^k]\}.$$

Note that on each interval $[\nu_j^k, \tau_j^k]$ the priority is fixed, where class k receives the least priority.

Lemma 5.1. *i. Fix* $k \in \mathcal{K}$, $j \in \mathbb{N}$, *and* $T > 0$. *Then, as* $n \to \infty$,

$$\max_{i : i \ne k} \max_{t \in I_j^k \cap [\varepsilon_1, T]} \hat{X}_i^n(t) \to 0 \quad \textit{in probability.}$$

ii. Fix $k \in \mathcal{K}$ *and* $T > 0$. *Then*

$$\varDelta_k^n := \sup_{t \in [\varepsilon_1, T]} 1_{\{\tilde{X}_t^n \in L_k^\varepsilon\}} \|\hat{X}^n(t) - \mu_k \tilde{X}_t^n e_k\| \to 0, \quad \textit{in probability.}$$

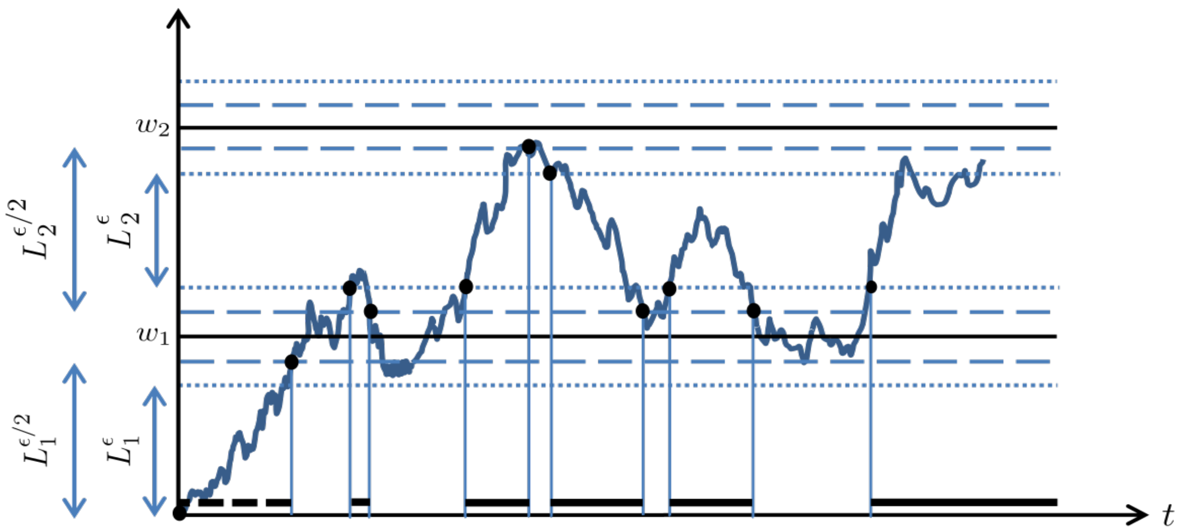

Figure 3: *Workload as a function of time and corresponding random times. Solid [resp., dashed, dotted] lines show the levels w_k [resp., $w_k \pm \varepsilon/2$, $w_k \pm \varepsilon$] for $k = 1, 2$. Time intervals $[\sigma_j^1, \tau_j^1]$ [resp., $[\sigma_j^2, \tau_j^2]$] are shown in thick dashed [resp., solid] lines. These time intervals begin on entering the set L_k^ε and end on exiting $L_k^{\varepsilon/2}$. Throughout such an interval, the priority remains fixed.*

Proof. i. We simplify the notation by writing σ, τ, ν for σ_j^k, τ_j^k and ν_j^k. In the first step we show that, as $n \to \infty$,

$$\max_{i:i\neq k} \hat{X}_i^n(\sigma) 1_{\{\sigma \le T\}} \to 0 \quad \text{in probability.} \tag{43}$$

This is achieved by analyzing the time interval $[\nu, \sigma]$. Denote

$$G^n(t) = \sum_{i:i\neq k} \frac{n}{\mu_i^n} \hat{X}_i^n(t).$$

It follows from the definition of the policy that, during any time interval $[a, b]$ on which (i) $G^n > 0$ (specifically, customers from classes other than k are present throughout the interval), and (ii) $\tilde{X}^n \in L_k$, the server must dedicate its full capacity to processing jobs from classes other than k. Thus

$$\sum_{i:i\neq k} T_i^n[a, b] = b - a. \tag{44}$$

For such an interval, a calculation based on (19)–(23) and (44) gives the following balance equation for G^n:

$$\begin{aligned} G^n[a, b] &= \sum_{i:i\neq k} \frac{n}{\mu_i^n} \Big(\hat{W}_i^n[a, b] - \int_a^b \theta_i \hat{X}_i^n(s) ds + \hat{Y}_i^n[a, b] + e_i^n[a, b] \Big) \\ &= L^n[a, b] - n^{1/2} \rho_k (b - a), \end{aligned} \tag{45}$$

where

$$L^n(t) = \sum_{i:i\neq k} \frac{n}{\mu_i^n} \Big(\hat{W}_i^n(t) - \int_0^t \theta_i \hat{X}_i^n(s) ds + e_i^n(t) \Big).$$

Note that L^n are C-tight, as follows by the C-tightness of $\hat{W}^n$, the uniform (on compacts) convergence of e^n, and the boundedness in probability of $\hat{X}^n$ (all shown in Lemma 4.3).

To prove (43), it suffices to show that, for any $\delta > 0$ (that we keep fixed in what follows),

$$P(G^n(\sigma) > \delta, \sigma \le T) \to 0. \tag{46}$$

Consider the event

$$\Omega_n = \{G^n(t) > \delta/2 \text{ for all } t \in [\nu, \sigma]\}.$$

We address (46) by considering separately the event Ω_n and its complement. First, note that on the event Ω_n, during the time interval $[\nu, \sigma]$ there are always some customers in the system from classes other than k. In particular, (45) is valid with ν for a and σ for b, and so on the event $\{G^n(\sigma) > \delta\} \cap \Omega_n$ one has

$$L^n[\nu, \sigma] - n^{1/2}\rho_k(\sigma - \nu) = G^n(\sigma) - G^n(\nu) \ge -G^n(\nu),$$

hence fixing a sequence $r_n > 0$, $r_n \to 0$, $n^{1/2} r_n \to \infty$, one has

$$P(\{G^n(\sigma) > \delta\} \cap \Omega_n \cap \{\sigma \le T\}) \le P(\sigma - \nu \le r_n, \sigma \le T) + P(n^{1/2} r_n \rho_k \le 2\|L^n\|_T + \|G^n\|_T).$$

Since $\|L^n\|_T$ and $\|G^n\|_T$ are tight RVs, the last term converges to zero.

To show that $P(\sigma - \nu \le r_n, \sigma \le T) \to 0$, note that by construction $\sigma \ge \varepsilon_1 > 0$, and moreover, either $\nu = 0$ or $|\tilde{X}^n[\nu, \sigma]| \ge \varepsilon/2$ must hold. Since ε_1 is fixed and $r_n \to 0$,

$$P(\sigma - \nu \le r_n, \sigma \le T) \le P(w_T(r_n, \tilde{X}^n) \ge \varepsilon/2) \to 0,$$

as follows from the C-tightness of $\tilde{X}^n$ (proved in Lemma 4.3).

Next, on the event Ω_n^c, there exists a time $\eta \in [\nu, \sigma)$ such that $G^n(\eta) \le \delta/2$. We may assume without loss of generality that $G^n(t) > 0$ for all $t \in [\eta, \sigma]$. As a result, (45) is valid with η for a and σ for b. Thus

$$P(\{G^n(\sigma) > \delta\} \cap \Omega_n^c, \sigma \le T) \le P(L^n[\eta, \sigma] - n^{1/2}\rho_k(\sigma - \eta) \ge \delta/2, \sigma \le T).$$

Splitting according to whether $\sigma - \eta > r_n$ or $\le r_n$ shows that the above probability is bounded by

$$P(2\|L^n\|_T \ge n^{1/2} r_n \rho_k) + P(w_T(r_n, L^n) \ge \delta/2).$$

Both probabilities in the above display converge to zero as $n \to \infty$ by the C-tightness of L^n and the properties $r_n \to 0$ and $n^{1/2} r_n \to \infty$. This completes the proof of (46).

On our second step we prove part (i). Due to (43), it suffices to prove that, for any fixed $\delta > 0$,

$$P(\sigma \le T, G^n(\sigma) \le \delta, \sup_{[\sigma, \tau \wedge T]} G^n \ge 2\delta) \to 0. \tag{47}$$

On the event indicated in (47) there must exist a random time $\sigma_1 \le \tau_1 := \tau \wedge T$ such that $\sigma \le \sigma_1 \le \tau_1$, $G^n(\sigma_1) \le \delta$, and $G^n > 0$ on $[\sigma_1, \tau_1]$. Now, since the new interval $[\sigma_1, \tau_1]$ is a subset of $[\sigma, \tau]$, we have $\tilde{X}^n \in L_k$ on this interval, and so (45) is again valid, with σ_1 for a and τ_1 for b. Hence the probability in (47) is bounded by

$$P(L^n[\sigma_1, \tau_1] - n^{1/2}\rho_k(\tau_1 - \sigma_1) \ge \delta).$$

The above is further bounded by

$$P(2\|L^n\|_T \ge n^{1/2}\rho_k r_n) + P(w_T(r_n, L^n) \ge \delta),$$

where both probabilities converge to zero by C-tightness of L^n. This completes the proof of part (i).

ii. If $t \in [\varepsilon_1, T]$ and $\tilde{X}_t^n \in L_k^\varepsilon$ then, by the definition of σ_j^k and τ_j^k, one has $t \in [\sigma_j^k, \tau_j^k \wedge T]$ for some j for which $\sigma_j^k \le T$. Denoting

$$Z_k^n = \#\{j : \sigma_j^k \le T\},$$

we thus have, for $i \neq k$,

$$\sup_{t\in[\varepsilon_1,T]} 1_{\{\tilde X^n_t\in L^\varepsilon_k\}}\hat X^n_i(t) \le \max_{j\le Z^n_k}\max_{i:i\neq k}\max_{t\in I^k_j\cap[\varepsilon_1,T]}\hat X^n_i(t).$$

Moreover, by (11), $\tilde X^n = m'\hat X^n$, and so $|\hat X^n_k(t) - \mu_k \tilde X^n_t| \le c_1 \max_{i:i\neq k}\hat X^n_i(t)$, for some constant c_1. It follows that, for some constant c_2,

$$\Delta^n_k \le c_2 \max_{j\le Z^n_k}\max_{i:i\neq k}\max_{t\in I^k_j\cap[\varepsilon_1,T]}\hat X^n_i(t).$$

Hence to deduce the result from part (i) of the lemma it suffices to show that $\{Z^n_k\}_{n\in\mathbb{N}}$ is a tight sequence of RVs. To this end, note that, by the definition of σ^k_j and τ^k_j, $|\tilde X^n[\sigma^k_j,\tau^k_j]| \ge \varepsilon/2$ for every j. Hence

$$P(Z^n_k \ge K+1) \le P(w_T(K^{-1}T, \tilde X^n) \ge \varepsilon/2).$$

Since $\tilde X^n$ are C-tight, the limit superior of the RHS of the above display can be made arbitrarily small by selecting K large. Hence Z^n_k are tight, and the result follows. □

Lemma 5.2. *i. For every $T \ge 0$ there exists a constant a_1 such that $E\|\hat X^n\|^2_T \le a_1$.*

ii. For every $T \ge 0$ there exists a constant a_2 such that $E|H^n|_T + E|Y^{\#,n}|_T \le a_2$.

iii. For every $\varepsilon > 0$ there exist constants a_3 and n_0 such that, for all $n \ge n_0$,

$$E\|\hat X^n\|_t \le a_3 e^{\varepsilon t}, \qquad t \ge 0.$$

Proof. In this proof, the symbols c, c_1, c_2 denote positive constants that do not depend on n or t. We will use equations (36), (37), as well as the bounds $c_1\|\tilde X^n\|_t \le \|\hat X^n\|_t \le c_2\|\tilde X^n\|_t$. First, let us argue along the lines of Remark 3.1, that

$$|\tilde X^n_t| \le |\tilde X^n(0)| + 2\|\tilde W^n\|_t + 2\|e^{\#,n}\|_t. \tag{48}$$

Recall that $\tilde X^n$ is non-negative and H^n is non-negative and non-decreasing. Given $t > 0$, if $\tilde X^n_s > 0$ for all $s \le t$ then, by (36), $Y^{\#,n}_t = 0$, and thus (48) holds. Otherwise, let $\sigma = \sup\{s \in [0,t] : \tilde X^n_s = 0\}$. Then $\tilde X^n_{\sigma-} = 0$, and by (36), $Y^{\#,n}_t = Y^{\#,n}_{\sigma-}$. Hence by (37),

$$\tilde X^n_t = \tilde X^n_t - \tilde X^n_{\sigma-} \le \tilde W^n_t - \tilde W^n_{\sigma-} + e^{\#,n}_t - e^{\#,n}_{\sigma-},$$

and (48) follows.

The L^2 convergence assumption on $\hat X^n_0$ implies $E\|\tilde X^n_0\|^2 \le c$. Also, from equation (19) and the fact that $T^n_i(t) \le t$ and y^n_i converge, we have $\|\hat W^n\|_t \le c(t + \|\hat A^n\|_t + \|\hat S^n\|_t)$. It is well known that the scaled renewal processes satisfy $E\|\hat A^n\|^2_t + E\|\hat S^n\|^2_t \le c(1+t)$, thanks to the second moment assumption; see eg. equation (172) in [9]. As a result,

$$E\|\tilde W^n\|_t \le cE\|\hat W^n\|^2_t \le c(1+t^2). \tag{49}$$

We thus have by (48)

$$E\|\tilde X^n\|^2_t \le c(1+t^2) + cE\|e^{\#,n}\|^2_t. \tag{50}$$

Let us then analyze the error term $e^{\#,n}$. Recall that

$$e^{\#,n}_t = m'(-e^{(1),n} + e^{(2),n}) + e^{(3),n}_t,$$

where $e^{(1),n}$ and $e^{(2),n}$ are given by (20) and (21), and

$$e_t^{(3),n} = m'\hat{Y}_t^n - (m^n)'\hat{Y}_t^n.$$

Recall from the proof of Lemma 4.1 that $\{e_i^{(1),n}(t)\}_t$ is a martingale, for each n and i. By the Burkholder-Davis-Gundy inequality,

$$E\|e_i^{(1),n}\|_t^2 \leq cE[e_i^{(1),n}, e_i^{(1),n}]_t,$$

where $[M, M]_t$ denotes the quadric variation of a process M. Since $e_i^{(1),n}$ is piecewise smooth and the size of each of its jumps is $n^{-1/2}$, $[e_i^{(1),n}, e_i^{(1),n}]_t$ is given by n^{-1} times the number of its jumps by time t. The latter is given by $R_i^0(n^{1/2}\theta_i^n \int_0^t \hat{Q}_i^n(s)ds)$. Thus

$$E\|e_i^{(1),n}\|_t^2 \leq cn^{-1}ER_i^0(n^{1/2}\theta_i^n \int_0^t \hat{Q}_i^n(s)ds) = cn^{-1}n^{1/2}\theta_i^n E\int_0^t \hat{Q}_i^n(s)ds, \tag{51}$$

where the last identity again follows the martingale property. Hence, using (3),

$$E\|e^{(1),n}\|_t^2 \leq cn^{-1/2}\int_0^t E\|\hat{Q}^n(s)\|ds \leq cn^{-1/2}\int_0^t E\|\hat{X}^n(s)\|ds. \tag{52}$$

Next, by (3),

$$\|e^{(2),n}\|_t \leq cn^{-1/2}t. \tag{53}$$

As for the term $e^{(3),n}$, it is shown in the proof of Lemma 4.3 that $|e_t^{(3),n}| \leq \varepsilon_n\|\hat{Y}^n\|_t$ where $\varepsilon_n \to 0$ is a deterministic sequence. By (22) and the positivity of $\hat{X}_i^n(0)$,

$$\|\hat{Y}_t^n\| \leq \|\hat{X}_t^n\| + \|\hat{W}_t^n\| + \|e_t^{(1),n}\| + \|e_t^{(2),n}\| + c\int_0^t \|\hat{X}_s^n\|ds.$$

As a result,

$$\|\hat{Y}^n\|_t \leq \|\hat{X}^n\|_t + \|\hat{W}^n\|_t + \|e^{(1),n}\|_t + \|e^{(2),n}\|_t + c\int_0^t \|\hat{X}_s^n\|ds, \tag{54}$$

where we used the monotonicity in t of the last term. Using Jensen's inequality,

$$\|e^{(3),n}\|_t^2 \leq c\varepsilon_n^2(\|\hat{X}^n\|_t^2 + \|\hat{W}^n\|_t^2 + \|e^{(1),n}\|_t^2 + \|e^{(2),n}\|_t^2) + c\varepsilon_n^2 t\int_0^t \|\hat{X}^n\|_s^2 ds. \tag{55}$$

Combine (50), (52), (53) and (55) to obtain

$$M_t^n := E\|\hat{X}^n\|_t^2 \leq c(1+t^2) + cn^{-1/2}\int_0^t E\|\hat{X}^n(s)\|ds + cn^{-1/2}t + c\varepsilon_n^2 M_t^n + c\varepsilon_n^2 t\int_0^t M_s^n ds. \tag{56}$$

For the second term on the RHS use Jensen's inequality and $\sqrt{y} \leq 1+y$, $y \geq 0$, to bound the integrand $E\|\hat{X}^n(s)\|$ by $1 + M_s^n$. The fourth term on the RHS is bounded by $\frac{1}{2}M_t^n$ for all sufficiently large n. As a result,

$$M_t^n \leq c(1+t^2) + cn^{-1/2}\int_0^t (1+M_s^n)ds + cn^{-1/2}t + c\varepsilon_n^2 t\int_0^t M_s^n ds.$$

Thus

$$M_t^n \leq c(1+t^2) + c(n^{-1/2} + \varepsilon_n^2 t)\int_0^t M_s^n ds.$$

Using Gronwall's lemma gives

$$M_t^n \leq c(1+t^2)e^{cn^{-1/2}t + c\varepsilon_n^2 t^2} \leq c(1+t^2)e^{c(1+t^2)},$$

establishing part (i) of the lemma.

To prove part (iii) we use the bounds on the error terms without taking squares. Let $N^n_t = E\|\hat{X}^n\|_t$. Then by (48) and (49),

$$N^n_t \le c(1+t) + 2E\|e^{\#,n}\|_t.$$

Now,

$$E\|e^{(1),n}\|_t \le \sqrt{E\|e^{(1),n}\|_t^2} \le 1 + E\|e^{(1),n}\|_t^2 \ \le 1 + cn^{-1/2}\int_0^t N^n_s ds,$$

by (52). By (54),

$$E\|e^{(3),n}\|_t \le \varepsilon_n\Big(N^n_t + c(1+t) + E\|e^{(1),n}\|_t + E\|e^{(2),n}\|_t + c\int_0^t N^n_s ds\Big).$$

Thus, for n such that $\varepsilon_n < \frac{1}{2}$,

$$N^n_t \le c(1+t) + c(n^{-1/2} + \varepsilon_n)\int_0^t N^n_s ds.$$

Hence by Gronwall's lemma, $N^n_t \le c(1+t)e^{c(n^{-1/2}+\varepsilon_n)t}$. This proves part (iii).

As for part (ii), the claim follows from part (i): First,

$$|H^n|_T \le c\int_0^T \|\hat{X}^n_s\| ds,$$

hence a bound on its expectation follows from the bound we have on $\|\hat{X}^n\|_T$. Next, recall from (23) that $Y^{\#,n}$ is non-negative and non-decreasing. Hence $|Y^{\#,n}|_T = Y^{\#,n}_T$. According to (37), it suffices to show that $E(|\tilde{X}^n_T| + |\tilde{W}^n_T| + |H^n_T| + |e^{(1),n}_T| + |e^{(2),n}_T| + |e^{(3),n}_T|) \le c$, where c does not depend on n. These estimates all follow directly from the proof of part (i). Thus follows (ii). □

Recall the notation $\tilde{X}^n$, $\tilde{W}^n$, H^n, $\tilde{Y}^n$ and $\tilde{e}^n$ and the relation (34) that they satisfy. We will next take limits in this equation and argue that the limits form a solution to the SDE with reflection (32). Recall that weak existence and uniqueness hold for equation (32), and let $(\tilde{X}, \tilde{Y}, \tilde{W})$ be a solution. Denote $H_t = \int_0^t \theta_{i(\tilde{X}_s)}\tilde{X}_s ds$, by which

$$\tilde{X}_t = \tilde{x} + \tilde{W}_t - H_t + \tilde{Y}_t \ge 0, \qquad \int 1_{\{\tilde{X}>0\}} d\tilde{Y} = 0.$$

Denote

$$K^n_t = \int_0^t e^{-\alpha s}(c^n)'\hat{Q}^n_s ds. \tag{57}$$

We are concerned with the convergence of K^n because by Lemma 4.1, $\hat{J}^n(B^{n,*}) = EK^n_\infty$.

Lemma 5.3. *Fix $T < \infty$. Then*

$$\limsup_n E\int_0^T 1_{\{\tilde{X}^n_t \in M^\varepsilon\}} dt \le c\varepsilon, \tag{58}$$

where c does not depend on ε. Moreover,

$$(\tilde{X}^n, \tilde{W}^n, H^n, \tilde{Y}^n) \Rightarrow (\tilde{X}, \tilde{W}, H, \tilde{Y}), \qquad K^n_T \Rightarrow \int_0^T e^{-\alpha t} q_{i(\tilde{X}_t)}\tilde{X}_t dt. \tag{59}$$

Proof of Lemma 5.3. Recall from Lemma 4.3 that we have $\tilde{W}^n \Rightarrow \tilde{W}$, and that the sequence $(\tilde{X}^n, \tilde{W}^n, H^n, \tilde{Y}^n)$ is C-tight. Fix a convergent subsequence and denote by (X^1, W^1, H^1, Y^1) its weak limit. Proving the first part of (59) amounts to showing that (X^1, W^1, H^1, Y^1) and $(\tilde{X}, \tilde{W}, H, \tilde{Y})$ are equal in law; note that W^1 and $\tilde{W}$ are equal in law.

The main estimate towards showing (59) is (58), which we prove first. By Lemma 4.3,

$$X_t^1 = \tilde{x} + W_t^1 - H_t^1 + Y_t^1, \qquad Y_t^1 = \sup_{s\le t}(-\tilde{x} - W_s^1 + H_s^1)^+. \tag{60}$$

Recall from the proof of Lemma 4.3 that $\|e^{(3),n}\|_T \Rightarrow 0$, where $e^{(3),n} = \tilde{Y}^n - Y^{\#,n}$. Hence Y^1 is also a weak limit of $Y^{\#,n}$. Hence, by Fatou's lemma, the uniform bound on $E|H^n|_T$ and $E|Y^{\#,n}|_T$ obtained in Lemma 5.2 implies that the variation of H^1 and Y^1 over $[0,T]$ satisfies

$$E(|H^1|_T + |Y^1|_T) \le c. \tag{61}$$

Let $g = g^\varepsilon : \mathbb{R}_+ \to [0,1]$ be a continuous function, that equals 1 on M^ε and vanishes on the set $\{x : \mathrm{dist}(x, M^\varepsilon) > \varepsilon\}$. Denote $G = \int_0^\cdot g(x)dx$ and $\hat{G} = \int_0^\cdot G(x)dx$. Note that $G(x) \le c\varepsilon$ and $\hat{G}(x) \le c\varepsilon x$ for all x. Applying Ito's formula to $\hat{G}(X_T^1)$,

$$\hat{G}(X_T^1) = \hat{G}(\tilde{x}) + \int_0^T G(X_t^1)(dW_t^1 - dH_t^1 + dY_t^1) + c\int_0^T g(X_t^1)dt.$$

Hence

$$E\int_0^T g(X_t^1)dt \le c\varepsilon + c\varepsilon E|X_T^1| + c\varepsilon(T + E|H^1|_T + |Y^1|_T).$$

Using (61) gives

$$E\int_0^T g(X_t^1)dt \le c\varepsilon. \tag{62}$$

Since $\tilde{X}^n \Rightarrow X^1$ and $g(x) \ge 1_{M^\varepsilon}(x)$ for all $x \in \mathbb{R}_+$, we have (58) valid along the subsequence. Since, moreover, the constants do not depend on the chosen subsequence, (58) follows.

Now we show $(X^1, H^1, Y^1) = (X, Y, H)$. We have

$$X_t = x + W_t + \int_0^t b(X_s)ds + Y_t \ge 0, \qquad \int 1_{\{X>0\}}dY = 0.$$

Also, it has a unique strong solution. Rewrite (34) as

$$\tilde{X}_t^n = m'\hat{X}_t^n = \tilde{X}_0^n + \tilde{W}_t^n - \int_0^t m'\Theta\hat{X}_s^n ds + \tilde{Y}_t^n + \tilde{e}_t^n,$$

and recall by Lemma 4.3 that $\tilde{e}^n \Rightarrow 0$. Write the above as

$$\begin{aligned}
\tilde{X}_t^n &= \tilde{X}_0^n + \tilde{W}_t^n - \int_0^{\varepsilon_1} m'\Theta\hat{X}_s^n ds - \int_{\varepsilon_1}^{t\vee\varepsilon_1}(1 - g^\varepsilon(\tilde{X}_s^n))m'\Theta\hat{X}_s^n ds \\
&\quad - \int_{\varepsilon_1}^{t\vee\varepsilon_1} g^\varepsilon(\tilde{X}_s^n)m'\Theta\hat{X}_s^n ds + \tilde{Y}_t^n + \tilde{e}_t^n \\
&= \tilde{X}_0^n + \tilde{W}_t^n - \int_0^{\varepsilon_1} m'\Theta\hat{X}_s^n ds - \int_{\varepsilon_1}^{t\vee\varepsilon_1}(1 - g^\varepsilon(\tilde{X}_s^n))\theta_{i(\tilde{X}_s^n)}\tilde{X}_s^n ds \\
&\quad - \int_{\varepsilon_1}^{t\vee\varepsilon_1} g^\varepsilon(\tilde{X}_s^n)m'\Theta\hat{X}_s^n ds + \tilde{Y}_t^n + \hat{e}_t^n,
\end{aligned} \tag{63}$$

where, using the convergence $\Delta^n \Rightarrow 0$ from Lemma 5.1, one has $\hat{e}^n \Rightarrow 0$. Thus

$$\tilde{X}_t^n = \tilde{X}_0^n + \tilde{W}_t^n - \int_0^t (1 - g^\varepsilon(\tilde{X}_s^n))\theta_{i(\tilde{X}_s^n)}\tilde{X}_s^n ds + \tilde{Y}_t^n + \hat{e}_t^n + \eta^n(t, \varepsilon, \varepsilon_1), \tag{64}$$

where

$$\sup_{t\in[0,T]} |\eta^n(t, \varepsilon, \varepsilon_1)| \le c\varepsilon_1 \|\hat{X}^n\|_T + c\|\hat{X}^n\|_T \int_0^T g^\varepsilon(\tilde{X}_t^n)dt. \tag{65}$$

We now take weak limits in (64), along the chosen subsequence. Note that the set of discontinuities of $y \mapsto i(y)$ is given by $\{w_k\}$, and that g^ε vanishes on a neighborhood of this set. As a result, we have uniform continuity of the coefficient $y \to (1 - g^\varepsilon(y))\theta_{i(y)}y$, (for each $\varepsilon > 0$). Since all terms in (64) except the last one converge uniformly on $[0, T]$, the last one must also converge. We denote its limit $\eta(t, \varepsilon, \varepsilon_1)$. We have

$$X_t^1 = \tilde{x} + W_t^1 - \int_0^t (1 - g^\varepsilon(X_s^1))\theta_{i(X_s^1)}X_s^1 ds + Y_t^1 + \eta(t, \varepsilon, \varepsilon_1),$$

where, using the fact $g^\varepsilon \le 1_{M^{3\varepsilon}}$ with (58) and (65), we have for $\delta > 0$,

$$\limsup_{\varepsilon\to 0} \limsup_{\varepsilon_1\to 0} P(\|\eta(\cdot, \varepsilon, \varepsilon_1)\|_T > \delta) = 0.$$

Thus

$$X_t^1 = \tilde{x} + W_t^1 - \int_0^t \theta_{i(X_s^1)}X_s^1 ds + Y_t^1 + \eta(t, \varepsilon, \varepsilon_1) + \tilde{\eta}(t, \varepsilon),$$

where $|\tilde{\eta}(t, \varepsilon)| \le c\|X^1\|_T \int_0^T g^\varepsilon(X_s^1)ds$ and so by (62),

$$\limsup_{\varepsilon\to 0} P(\|\tilde{\eta}(\cdot, \varepsilon)\|_T > \delta) = 0.$$

It follows that

$$X_t^1 = \tilde{x} + W_t^1 - \int_0^t \theta_{i(X_s^1)}X_s^1 ds + Y_t^1.$$

We also have $X^1 \ge 0$, and by (60), $\int X^1 dY^1 = 0$. Using weak uniqueness of solutions to the SDE, stated in Proposition 3.1(iii), and the fact that W^1 and $\tilde{W}$ are equal in law gives the claimed equality in law of (X^1, W^1, H^1, Y^1) and $(\tilde{X}, \tilde{W}, H, \tilde{Y})$. This proves the first part of (59).

The convergence of K_T^n stated in (59) is proved by writing, similar to (63),

$$\begin{aligned} K_T^n &= \int_0^T e^{-\alpha t}(c^n)'\hat{Q}_t^n dt \\ &= \int_0^T e^{-\alpha t}c'\hat{X}_t^n dt + e^{(4),n} \\ &= \int_0^T e^{-\alpha t}(1 - g^\varepsilon(\tilde{X}_t^n))c'\hat{X}_t^n dt + \int_0^T e^{-\alpha t}g^\varepsilon(\tilde{X}_t^n)c'\hat{X}_t^n dt, \end{aligned}$$

and repeating the argument provided above, taking $n \to \infty$ then $\varepsilon \to 0$. □

Proof of Theorem 5.1. Recall that, by Lemma 4.1, $\hat{J}^n(B^{n,*}) = EK_\infty^n$. It follows from Lemma 5.2(iii) that $\limsup_n E\int_0^\infty e^{-\alpha t}\|\hat{X}_t^n\|dt < \infty$. Hence, given $\varepsilon > 0$, there exists $T < \infty$ such that

$$\limsup_n E\int_T^\infty e^{-\alpha t}\|\hat{X}_t^n\|dt < \varepsilon.$$

Fix such T. Then $\limsup_{n\to\infty} EK^n_\infty \le \limsup_{n\to\infty} EK^n_T + c\varepsilon$, where c does not depend on ε. Lemma 5.2(i) implies $E\int_0^T \|\hat{X}^n_t\|^2 dt \le c$, where c does not depend on n. Hence follows the uniform integrability of K^n_T, which, along with the weak convergence stated in (59) gives

$$\limsup_{n\to\infty} \hat{J}^n(B^{n,*}) \le E\int_0^T e^{-\alpha t} q_{i(\tilde{X}_t)}\tilde{X}_t dt + c\varepsilon \le E\int_0^\infty e^{-\alpha t} q_{i(\tilde{X}_t)}\tilde{X}_t dt + c\varepsilon.$$

By (29) and Proposition 3.1(iii), the RHS above equals $\tilde{V}(\tilde{x}) + c\varepsilon$. Since ε is arbitrary, it follows that $\limsup_{n\to\infty} \hat{J}^n(B^{n,*}) \le \tilde{V}(\tilde{x}) = V(x)$, where the last identity is provided in Lemma A.1(iii). ☐

A Appendix

A.1 Elementary properties of the BCP and RBCP

In Lemma A.1 we prove the equivalence between the BCP and RBCP. This result is needed because the lower bound is given in terms of the BCP value function V, whereas the upper bound is linked to the RBCP value function, $\tilde{V}$. Then, in Lemma A.2 we prove properties such as monotonicity and convexity of the RBCP value function.

Lemma A.1. *i. Given an admissible control Y_t for the I-dimensional BCP (Definition 3.1) defined on a filtered probability space $(\bar{\Omega}, \bar{\mathcal{F}}, \{\bar{\mathcal{F}}_t\}, \bar{P})$ with initial condition x, let $\tilde{x} = m'x$ and define $\tilde{X}_t = m'X_t$, $\tilde{Y}_t = m'Y_t$ and $U_i(s) = (m'X_s)^{-1} m_i X_i(s)$ on the same probability space. Then $(U, \tilde{Y}) \in \tilde{\mathcal{A}}(\tilde{x})$ and $\tilde{J}(\tilde{x}, (U, \tilde{Y})) = J(x, Y)$.*

ii. Given an admissible control $(U, \tilde{Y})$ for the RBCP (Definition 3.2) defined on a filtered probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \{\tilde{\mathcal{F}}_t\}, \tilde{P})$ with initial condition $\tilde{x}$, let W_t be any $\tilde{\mathcal{F}}_t$ adapted I-dimensional BM such that $m'W_t = \tilde{W}_t$, and $x \in \mathbb{R}^I_+$ be any vector such that $m'x = \tilde{x}$. Define $X_t = \tilde{X}_t M^{-1} U_t$ and $Y_i(t) = X_i(t) - x_i - W_i(t) + \theta_i \int_0^t X_i(s) ds$ for all $i \in \mathcal{I}$ on the same probability space. Then, $Y_t = (Y^1_t, \ldots, Y^I_t) \in \mathcal{A}(x)$ and $J(x, Y) = \tilde{J}(\tilde{x}, (U, \tilde{Y}))$.

iii. $V(x) = \tilde{V}(\tilde{x})$.

Proof. i. By its definition, $\{\tilde{W}_t\}$ is a 1-dimensional $\tilde{\mathcal{F}}_t$-BM with the stated parameters. The processes $\{U_t\}$, $\{\tilde{Y}_t\}$, $\{\tilde{X}_t\}$ have sample paths in $\mathbb{D}_{\mathcal{S}_1}(\mathbb{R}_+)$, $\mathbb{D}_{\mathbb{R}^I}(\mathbb{R}_+)$ and $\mathbb{D}_{\mathbb{R}}(\mathbb{R}_+)$, resp., and are $\tilde{\mathcal{F}}_t$-adapted. Also, since Y is admissible one has $X \ge 0$ $\bar{P}$-a.s., and $m'X \ge 0$ $\bar{P}$-a.s. The nonnegativity and monotonicity of $\tilde{Y}$ follows from the fact that Y satisfies Definition 3.1. Thus $(U, \tilde{Y}) \in \tilde{\mathcal{A}}(\tilde{x})$. Regarding the cost functions, notice that $X_t = \tilde{X}_t M^{-1} U_t$. This equation, along with the definition of $\tilde{X}$ implies $\sum c_i \theta_i X_i(t) = \sum c_i \theta_i \mu_i U_i(t) \tilde{X}(t)$ and thus $\tilde{J}(\tilde{x}, (U, \tilde{Y})) = J(x, Y)$.

ii. By definition and the assumption $(U, \tilde{Y}) \in \tilde{\mathcal{A}}(\tilde{x})$, one has $Y_t \in \mathcal{A}(x)$. The relation $X_t = \tilde{X}_t M^{-1} U_t$ also implies $J(x, Y) = \tilde{J}(\tilde{x}, (U, \tilde{Y}))$.

iii. This is an immediate consequence of parts 1 and 2. ☐

Lemma A.2. *The value function $\tilde{V}$ is non-decreasing, convex on $\mathbb{R}_+$, and its right-derivative at zero is well defined and equal to 0.*

Proof. Monotonicity follows easily from the structure of the dynamics and cost (28), (29). Indeed, given $0 \le x_1 < x_2$ and an ε-optimal control system for x_2, denoted $(\tilde{\Sigma}, \tilde{W}, U, \tilde{Y}, \tilde{X})$, one constructs on the probability space $\tilde{\Sigma}$ a new collection of processes by adding to $\tilde{Y}$ the constant $x_2 - x_1$. Namely, let $Y^* = \tilde{Y} + x_2 - x_1$. Then $(\tilde{\Sigma}, \tilde{W}, U, Y^*, \tilde{X})$ is a control system for x_1. Moreover, the two costs agree. Therefore, upon optimizing over control systems for x_1 and taking the limit $\varepsilon \to 0$, one obtains $\tilde{V}(x_1) \le \tilde{V}(x_2)$.

For convexity of $\tilde{V}$, the argument uses two transformations of the RBCP. One is simply to go from the RBCP to the original BCP (using the equivalence stated in Lemma A.1). In particular, since $\tilde{V}$

and V are related via $\tilde{V}(m'x) = V(x)$, $x \in \mathbb{R}_+^I$, is suffices to prove the convexity of V. Second, there is an equivalence between the weak and the strong formulation of the BCP, where the former is the formulation that we work with in this paper, namely Definition 3.1. The latter is concerned with a fixed probability space, endowed with an I-dimensional BM. In the context of the current problem this equivalence is justified by the methods of [2].

Given $x_1, x_2 \in \mathbb{R}_+^I$, consider two corresponding ε-optimal controls Y_1, Y_2, (defined on a common probability space, supporting the BM W). Let X_1 and X_2 denote the corresponding controlled processes defined via (24). Given $p \in [0,1]$ and $q = 1-p$, it is clear that $Y := pY_1 + qY_2$, $X := pX_1 + qX_2$ define a control and a corresponding controlled process for the initial condition $x = px_1 + qx_2$. By (25), $J(x, Y) = pJ(x_1, Y_1) + qJ(x_2, Y_2)$, and so, upon optimizing over controls for x and taking the limit $\varepsilon \to 0$, the convexity of V follows.

As for the final assertion, that is a standard result, we only provide a sketch. Fix $h > 0$. We already have that $\tilde{V}(0) \le \tilde{V}(h)$. Given an ε-optimal control $(U, \tilde{Y})$ for the RBCP with initial condition 0, let $\tilde{X}$ be the corresponding controlled process given by (28), namely

$$\tilde{X}_t = \tilde{W}_t - \int_0^t \theta' U_s \tilde{X}_s ds + \tilde{Y}_t, \qquad t \ge 0.$$

Let X^* denote the solution to (28) with initial condition h and $\tilde{Y} = 0$, namely

$$X_t^* = h + \tilde{W}_t - \int_0^t \theta' U_s X_s^* ds, \qquad t \ge 0.$$

Let $\tau = \inf\{t \ge 0 : \tilde{X}_t \ge X_t^*\}$. Let $(\hat{U}, \hat{Y}, \hat{X}) = (U, 0, X^*)$ on $[0, \tau)$ and $(\hat{U}, \hat{Y}, \hat{X}) = (U, \tilde{Y} - \tilde{Y}_\tau, \tilde{X})$ on $[\tau, \infty)$. This gives rise to a control for the initial condition h. The difference between the costs, by (29), is given by $E \int_0^\tau e^{-\alpha t} q' U(X_t^* - \tilde{X}_t) dt$. Standard estimates on BM, by which $E(\tau)$ is of order h^2, yield that $\tilde{V}(h) - \tilde{V}(0) \le ch^2$ for all small h. The result follows. □

A.2 The Bellman equation

The viscosity-sense solvability of Bellman equations of the form considered in this paper is well known (see Chapter V of [11]). See Section IV.5 of [11], equations (5.8) and (3.2) for the Bellman equation and the form of the Hamiltonian $\mathcal{H}$, respectively. However, classical, (i.e., C^2) solutions do not in general exist for such equations. In this paragraph we prove Proposition 3.1, addressing classical solutions and solution to the RBCP. The reader is referred to [12] for a different approach.

Proof of Proposition 3.1. The argument is based on existence of C^2 solutions satisfying Dirichlet boundary conditions, as well as on a verification theorem, by which the solution of a Bellman equation equals the value of a control problem. We use such a verification three times in the proof, in slightly different settings. However, rather than providing the details of these verification theorems, that are standard, we refer the reader to Chapter IV of [11] for several such results in similar context.

Step 1. Fix $p > 0$. We show that a C^2 solution to the Bellman equation (13) exists on $[0, p]$, with $\frac{dv}{dx}(0) = 0$ and $v(p) = \tilde{V}(p)$. We address this by formulating an exit time control problem on $[0, p]$, for which existence of C^2 solutions for the corresponding Bellman equation is known. For an initial condition $x \in [0, p]$, consider the process

$$Z_t = x + \tilde{W}_t - \int_0^t b(U_s, Z_s) ds, \tag{66}$$

where $\tilde{W}_t$ is as in Definition 3.2 (i.e., a 1-dimensional $(\tilde{y}, \tilde{\sigma})$-BM), and for $u \in \mathcal{S}_1$ and $x \in \mathbb{R}$, $b(u, x) = \theta' ux$, and U is a control process taking values in $\mathcal{S}_1$. Denote $\tau = \inf\{t \ge 0 | Z_t = p \text{ or } Z_t = 0\}$. Consider the value function

$$H(x) = \inf_U J(x,U) = \inf_U E\Big(\int_0^\tau e^{-\alpha t} q' U_t Z_t dt + e^{-\alpha\tau} \tilde{V}(Z_\tau) \Big), \qquad 0 \le x \le p. \tag{67}$$

Note that $H(x) = \tilde{V}(x)$ for $x = 0$ and $x = p$. The Bellman equation associated with the problem is precisely (13) (considered for $0 < x < p$). By Theorem 5.1 of [14], there exists a $C^2(0,p) \cap C[0,p]$ function, that we will denote by h^p, that satisfies (13) with the boundary conditions $h(0) = \tilde{V}(0)$, $h(p) = \tilde{V}(p)$. A verification theorem (for example, as in Theorem IV.5.1 of [11]) shows that $h^p(x) = H(x)$ for $x \in [0,p]$. Next, the dynamic programming principle [11] gives that $H = \tilde{V}$ on $[0,p]$. Hence by the last assertion of Lemma A.2, the right-derivative of h^p at zero is well defined and equal to zero. Thus h^p is a classical solution to (13) on $(0,p)$ with boundary conditions $\frac{dh}{dx}(0) = 0$ and $h(p) = \tilde{V}(p)$.

Step 2. Take $p \to \infty$. Taking into account that p is arbitrary, it follows that $\tilde{V}$ is a classical solution to (13) with the boundary condition $\frac{dv}{dx}(0) = 0$.

For uniqueness of solutions to (13) we need to consider the equation with the additional boundary condition at infinity, as asserted. To this end, one first shows that $\tilde{V}$ satisfies such a condition. Indeed, it follows from Remark 3.1 that $\tilde{V}(x) \le cE\int_0^\infty e^{-\alpha t}(x + \|\tilde{W}\|_t)dt$ for $c = 2\|q\|$. Since $\tilde{W}_t$ is a BM with constant drift, the constant $E\int_0^\infty e^{-\alpha t}\|\tilde{W}\|_t dt$ is finite, and so $\tilde{V}(x) \le c(1 + x)$ for c that does not depend on $x \in \mathbb{R}_+$.

A verification-type argument for the control problem on $\mathbb{R}_+$ will now show that any classical solution, satisfying the two boundary conditions stated, equals the value function $\tilde{V}$. This completes the proof of parts (i) and (ii).

As for part (iii), weak existence and uniqueness of solutions to the SDE (32) without the reflection term $\tilde{Y}$ are well known (see Propositions 5.3.6 and 5.3.10 of [16]). The proof via Girsanov's theorem can be modified in a straightforward way, appealing to the continuity of the 1-dimensional Skorohod map, to obtain an analogous result for the SDE with reflection.

The assertion regarding the optimality of the control system defined via (32) and setting $U_t = e_{i(\tilde{X}_t)}$ is based, once again, on a standard verification-type argument.

Finally, the convexity of v (equivalently, $\tilde{V}$) has been shown in Lemma A.2. □

Acknowledgement. The authors are grateful to the associate editor and three referees for their numerous helpful comments that considerably helped improve the exposition.

References

[1] Barış Ata and Mustafa H. Tongarlak. On scheduling a multiclass queue with abandonments under general delay costs. *Queueing Syst.*, 74(1):65–104, 2013.

[2] Rami Atar and Amarjit Budhiraja. Singular control with state constraints on unbounded domain. *Ann. Probab.*, 34(5):1864–1909, 2006.

[3] Rami Atar, Chanit Giat, and Nahum Shimkin. The $c\mu/\theta$ rule for many-server queues with abandonment. *Oper. Res.*, 58(5):1427–1439, 2010.

[4] Rami Atar, Chanit Giat, and Nahum Shimkin. On the asymptotic optimality of the $c\mu/\theta$ rule under ergodic cost. *Queueing Syst.*, 67(2):127–144, 2011.

[5] Rami Atar, Haya Kaspi, and Nahum Shimkin. Fluid limits for many-server systems with reneging under a priority policy. *Mathematics of Operations Research*, 39(3):672–696, 2013.

[6] Rami Atar, Avi Mandelbaum, and Martin I. Reiman. Scheduling a multi class queue with many exponential servers: asymptotic optimality in heavy traffic. *Ann. Appl. Probab.*, 14(3):1084–1134, 2004.

[7] Rami Atar and Martin I. Reiman. Asymptotically optimal dynamic pricing for network revenue management. *Stochastic Systems*, 2(DOI: 10.1214/12-SSY062.):232–276, 2012.

[8] Rami Atar and Mark Shifrin. An asymptotic optimality result for the multiclass queue with finite buffers in heavy traffic. *Stochastic Systems*, to appear, 2015.

[9] Steven L. Bell and Ruth J. Williams. Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy. *Ann. Appl. Probab.*, 11(3):608–649, 2001.

[10] Patrick Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999. A Wiley-Interscience Publication.

[11] Wendell H. Fleming and H. Mete Soner. *Controlled Markov processes and viscosity solutions*, volume 25 of *Stochastic Modelling and Applied Probability*. Springer, New York, second edition, 2006.

[12] Samim Ghamami and Amy R. Ward. Dynamic scheduling of a two-server parallel server system with complete resource pooling and reneging in heavy traffic: asymptotic optimality of a two-threshold policy. *Math. Oper. Res.*, 38(4):761–824, 2013.

[13] J. Michael Harrison and Assaf Zeevi. Dynamic scheduling of a multiclass queue in the Halfin-Whitt heavy traffic regime. *Oper. Res.*, 52(2):243–257, 2004.

[14] Philip Hartman. *Ordinary differential equations*. Birkhäuser, Boston, Mass., second edition, 1982.

[15] Jean Jacod and Albert N. Shiryaev. *Limit theorems for stochastic processes*, volume 288 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1987.

[16] Ioannis Karatzas and Steven E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991.

[17] Jeunghyun Kim and Amy R. Ward. Dynamic scheduling of a $GI/GI/1+GI$ queue with multiple customer classes. *Queueing Syst.*, 75(2-4):339–384, 2013.

[18] Maialen Larranaga, Urtzi Ayesta, and Ina Maria Verloop. Dynamic fluid-based scheduling in a multi-class abandonment queue. *Performance Evaluation*, 70(10):841–858, 2013.

[19] Melanie Rubino and Baris Ata. Dynamic control of a make-to-order, parallel-server system with cancellations. *Operations Research*, 57(1):94–108, 2009.