

Optimality of the generalized $c\mu$ rule in the moderate deviation regime

Rami Atar¹ · Subhamay Saha²

Received: 1 May 2016 / Revised: 14 November 2016 / Published online: 19 April 2017
© Springer Science+Business Media New York 2017

Abstract This paper studies a multiclass queueing system with an associated risk-sensitive cost observed in heavy traffic at the moderate deviation scale, accounting for convex queue length penalties. The main result is the asymptotic optimality of a dynamic index policy known from the diffusion-scale heavy traffic literature as the generalized $c\mu$ rule.

Keywords Moderate deviations · Heavy traffic · Risk-sensitive cost · Differential games

Mathematics Subject Classification 60F10 · 60K25 · 93E20 · 49N70

1 Introduction

One of the most appealing results on queueing control problems (QCPs) in heavy traffic, obtained by van Mieghem [9], is the asymptotic optimality (AO) of a dynamic index policy for the multiclass queue with nonlinear holding (as well as queue length) costs. This policy, referred to as the *generalized $c\mu$ rule*, acts as a localized version of the classical $c\mu$ rule, where the parameter c is variable and obtained by feedback from

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

✉ Subhamay Saha
saha.subhamay@iitg.ernet.in

Rami Atar
atar@ee.technion.ac.il

¹ Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa, Israel

² Department of Mathematics, Indian Institute of Technology Guwahati, Assam, India

the system's state. Specifically, a class- i customer that experiences a delay of δ units of time incurs a cost $C_i(\delta)$, where C_i are given smooth, convex functions. The rule is to prioritize the classes according to the index $\mu_i C_i'(\hat{Q}_i^n(t))$, where n is the scaling parameter, μ_i and $\hat{Q}_i^n(t)$ are the corresponding service rate and *diffusion-scaled* queue length at time t , and C_i' denotes the derivative of C_i . The index is thus explicit and simple to compute, and in particular does not require solving a dynamic programming equation, while at the same time it is valid for a relatively rich family of costs.

This work addresses the multiclass scheduling problem in the *moderate deviation heavy traffic* (MDHT) regime, with a *risk-sensitive* (RS) version of the above cost. More precisely, assuming renewal structure for the arrival and potential service time processes, the cost considered is

$$\frac{1}{b_n^2} \log \mathbb{E} \exp \left\{ b_n^2 \int_0^T \sum_i C_i(\tilde{Q}_i^n(t)) dt \right\},$$

for a sequence $b_n \rightarrow \infty$ with $n^{-1/2} b_n \rightarrow 0$, where $\tilde{Q}_i^n := b_n^{-1} n^{-1/2} Q_i^n$ denotes the MD-scaled queue length for class i , and arrival and service rates both scale like n . The main result is the MD-scale AO of the generalized $c\mu$ rule, with the obvious adaptation of determining the index via \tilde{Q}_i^n rather than \hat{Q}_i^n .

In that, this paper contributes to a line of research started in [1], and continued in [2, 3], that addresses QCPs at the MDHT regime. The reader is referred to these papers for background and motivation, as well as discussions of similarities and dissimilarities between diffusion-scale heavy traffic approximations and MDHT, as well as between the latter and QCPs in the large deviation (LD) regime. It has been suggested in [1] that RS control in the MDHT regime is likely to enjoy explicitly computable AO policies much more often than in the LD regime. The papers [2, 3] support this expectation by establishing AO of explicitly computable policies for a complex QCP that involves a free boundary problem. Put in this context, the contribution of the present paper further strengthens the assertion from [1] alluded to above. In relation to this, it should be mentioned that an AO result of the generalized $c\mu$ rule is not to be expected in the LD regime. In fact, when C_i are linear functions, it is known under some conditions (including Markovity) that a RS control problem in the LD regime possesses an AO index rule distinct from the classical $c\mu$ rule (see [4], where a different index rule is shown to be AO in this regime).

The aforementioned paper [1] studies a general RS cost structure, and attains AO of policies described in terms of an underlying differential game. However, the main focus in [1] is to deal with the *existence* of such policies, not with the construction of ones that are in any reasonable sense described explicitly. Indeed, the sole example of a cost function C for which an explicit AO policy is known for this setting is that where C is linear [1]. It is desired to extend the family of cost functions possessing explicit AO policies, and the contribution of this paper can indeed be viewed as such an extension to the collection of convex functions C that adhere to a linear growth condition and certain smoothness assumptions.

Our proof technique borrows from [1] results regarding this game, as well as a general lower bound on the RS performance of the control system. The main body of

work needed to prove our AO result consists of establishing a matching upper bound on the performance under the policy considered.

Let us finally mention that the diffusion-scale heavy traffic analysis of the generalized $c\mu$ rule has reached far beyond [9], most notably in [7], where the rule’s AO was proved in a multi-server setting that allows for servers that differ from each other in terms of the service capabilities offered to the various classes. It is of interest to study this type of extension at the MD scale in future work.

The setting and main results appear in Sect. 2. Specifically, Theorems 2.4 and 2.5 assert the AO of a preemptive and non-preemptive version of the rule, respectively. Section 3 provides some elementary facts about the differential game, while Sects. 4.1 and 4.2 give the proofs of Theorems 2.4 and 2.5, respectively.

Notation For a positive integer k and $\alpha, \beta \in \mathbb{R}^k$, $\alpha \cdot \beta$ will denote the usual inner product, while $\|\cdot\|$ will denote the Euclidean norm. For $T > 0$ and a function $f : [0, T] \rightarrow \mathbb{R}^k$, let $\|f\|_t = \sup_{0 \leq s \leq t} \|f(s)\|$, $t \in [0, T]$. For $\delta > 0$, the δ -oscillation of f is defined as

$$\text{osc}_\delta(f) = \sup\{\|f(s) - f(t)\| : |s - t| \leq \delta, s, t \in [0, T]\}.$$

For a function f defined on \mathbb{R}_+ and $y > 0$, $\text{osc}_\delta(f, y)$ denotes the δ -oscillation of $f|_{[0, y]}$. Denote by $\mathcal{C}([0, T], \mathbb{R}^k)$ the space of continuous functions. The space of functions that are right-continuous with finite left limits (RCLL) will be denoted by $\mathcal{D}([0, T], \mathbb{R}^k)$, and endowed with the Skorohod J_1 metric [5]. It is a well-known fact (see [5]) that this metric space is Polish. Let $\mathcal{AC}([0, T], \mathbb{R}^k)$ denote the class of absolutely continuous mappings from $[0, T]$ to \mathbb{R}^k . Let $\mathcal{AC}_0([0, T], \mathbb{R}^k)$ denote the set of $\psi \in \mathcal{AC}([0, T], \mathbb{R}^k)$ with $\psi(0) = 0$.

2 Model description and results

The model consists of d customer classes and a single server. An infinite capacity buffer is devoted to each customer class, where customers are queued if the server is not available to serve them upon arrival. Within each class, the customers are served on first come first served basis. The server may serve only one customer at a time within a class; however, processor sharing is allowed, in the sense that the server can divide its effort among up to d customers of distinct classes.

The model is defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Expectation with respect to \mathbb{P} is denoted by \mathbb{E} . The parameters and the processes will be indexed by $n \in \mathbb{N}$, which will act as the scaling parameter. The arrivals occur according to independent renewal processes. Let $\mathcal{I} = \{1, \dots, d\}$. Let $\lambda_i^n, n \in \mathbb{N}, i \in \mathcal{I}$ be given parameters, representing the reciprocal of the mean arrival times for class- i customers. Given are d independent sequences $\{IA_i(l) : l \in \mathbb{N}\}_{i \in \mathcal{I}}$, of i.i.d. positive random variables with mean 1 and variance $\sigma_{IA,i}^2$. Then, the number of arrivals of class- i customers up to time t is given by

$$A_i^n(t) = \sup \left\{ l \geq 0 : \sum_{k=1}^l \frac{IA_i(k)}{\lambda_i^n} \leq t \right\}, \quad t \geq 0.$$

Similarly, we consider another set of parameters $\mu_i^n, n \in \mathbb{N}, i \in \mathcal{I}$, representing the reciprocal mean service times. Also given are d independent sequences $\{ST_i(l) : l \in \mathbb{N}\}_{i \in \mathcal{I}}$ of positive i.i.d. random variables with mean 1 and variance $\sigma_{ST,i}^2$. The potential service process for class- i customers is given by

$$S_i^n(t) = \sup \left\{ l \geq 0 : \sum_{k=1}^l \frac{ST_i(k)}{\mu_i^n} \leq t \right\}, \quad t \geq 0.$$

Namely, $S_i^n(t)$ gives the number of class- i customers served by the time the server has devoted t units of time to that class. Now let $\{b_n\}$ be a sequence satisfying $\lim b_n = \infty$ and $\lim \frac{b_n}{\sqrt{n}} = 0$. We will refer to b_n^{-2} as the *moderate deviation speed*. The arrival and service parameters are assumed to satisfy the following limits as $n \rightarrow \infty$:

$$\begin{aligned} \lambda_i^n/n &\rightarrow \lambda_i \in (0, \infty) \text{ and } \mu_i^n/n \rightarrow \mu_i \in (0, \infty), \\ \tilde{\lambda}_i^n &:= \frac{1}{b_n\sqrt{n}}(\lambda_i^n - n\lambda_i) \rightarrow \tilde{\lambda}_i \in (-\infty, \infty), \\ \tilde{\mu}_i^n &:= \frac{1}{b_n\sqrt{n}}(\mu_i^n - n\mu_i) \rightarrow \tilde{\mu}_i \in (-\infty, \infty). \end{aligned} \tag{2.1}$$

Also, the system is assumed to be critically loaded, that is, $\sum_{i=1}^d \rho_i = 1$, where $\rho_i = \lambda_i/\mu_i, i \in \mathcal{I}$. For $i \in \mathcal{I}$, let Q_i^n denote the number of class- i customers in the system. With $S = \{x = (x_1, \dots, x_d) \in [0, 1]^d : \sum x_i \leq 1\}$, let B^n be a process taking values in S , with the i th component representing the fraction of effort given to the class- i customer. Then, the number of class- i jobs completed by time t is given by

$$D_i^n(t) := S_i^n(T_i^n(t)), \tag{2.2}$$

where

$$T_i^n(t) = \int_0^t B_i^n(s) ds \tag{2.3}$$

is the cumulative amount of time devoted to class- i customers by time t . Assuming, for simplicity, that the system starts empty, we have

$$Q_i^n(t) = A_i^n(t) - D_i^n(t). \tag{2.4}$$

We regard B^n as the *control process*. Given n , we say that B^n is an *admissible* control if it has RCLL sample paths, and

- it is adapted to the filtration

$$\sigma\{A_i^n(s), S_i^n(T_i^n(s)), i \in \mathcal{I}, s \leq t\},$$

- for every $i \in \mathcal{I}$ and $t \geq 0$, one has

$$Q_i^n(t) = 0 \text{ implies } B_i^n(t) = 0. \tag{2.5}$$

Denote the set of admissible controls by \mathcal{B}^n . An admissible control B^n is said to be *non-preemptive* if (i) processor sharing is not allowed, namely B^n takes values in $\{x \in \{0, 1\}^d : \sum x_i \leq 1\}$; and (ii) service is non-interruptible, namely if for some t and i one has $B_i^n(t) = 1$, then $B_i^n(u) = 1$ for $u \in [t, \tau)$, where τ is the time of the next departure, $\tau := \inf\{s \geq t : \Delta^n(s) > \Delta^n(t)\}$, $\Delta^n := \sum_j D_j^n$. Denote the set of non-preemptive admissible controls by $\mathcal{B}^{\#,n}$. Clearly, $\mathcal{B}^{\#,n} \subset \mathcal{B}^n$.

We now introduce the centered and scaled versions of the processes:

$$\begin{aligned} \tilde{A}_i^n(t) &= \frac{1}{b_n \sqrt{n}} (A_i^n(t) - \lambda_i^n t), & \tilde{S}_i^n(t) &= \frac{1}{b_n \sqrt{n}} (S_i^n(t) - \mu_i^n t), \\ \tilde{Q}_i^n(t) &= \frac{1}{b_n \sqrt{n}} Q_i^n(t). \end{aligned} \tag{2.6}$$

Then, it follows from (2.4) that

$$\tilde{Q}_i^n(t) = y_i^n t + \tilde{A}_i^n(t) - \tilde{S}_i^n(T_i^n(t)) + Z_i^n(t), \tag{2.7}$$

where

$$Z_i^n(t) = \frac{\mu_i^n}{n} \frac{\sqrt{n}}{b_n} (\rho_i t - T_i^n(t)), \quad y_i^n = \tilde{\lambda}_i^n - \rho_i \tilde{\mu}_i^n. \tag{2.8}$$

The scaled processes $(\tilde{A}^n, \tilde{S}^n)$ are assumed to satisfy a moderate deviation principle. For that let $I_k, k = 1, 2$, be functions defined on $\mathcal{D}([0, T], \mathbb{R}^d)$ as follows. Denote $\mathcal{AC}_0 = \mathcal{AC}_0([0, T], \mathbb{R}^d)$. For $\psi = (\psi_1, \dots, \psi_d) \in \mathcal{D}([0, T], \mathbb{R}^d)$,

$$\begin{aligned} \mathbb{I}_1(\psi) &= \begin{cases} \sum_{i=1}^d \frac{1}{2\hat{\sigma}_{1,i}^2} \int_0^T \dot{\psi}_i^2 ds, & \text{if } \psi \in \mathcal{AC}_0, \\ \infty, & \text{otherwise,} \end{cases} \\ \mathbb{I}_2(\psi) &= \begin{cases} \sum_{i=1}^d \frac{1}{2\hat{\sigma}_{2,i}^2} \int_0^T \dot{\psi}_i^2 ds, & \text{if } \psi \in \mathcal{AC}_0, \\ \infty, & \text{otherwise,} \end{cases} \end{aligned}$$

where $\hat{\sigma}_{1,i}^2 = \lambda_i \sigma_{IA,i}^2$, $\hat{\sigma}_{2,i}^2 = \mu_i \sigma_{ST,i}^2$. Let $\mathbb{I}(\psi) = \mathbb{I}_1(\psi^{(1)}) + \mathbb{I}_2(\psi^{(2)})$ for $\psi = (\psi^{(1)}, \psi^{(2)}) \in \mathcal{D}([0, T], \mathbb{R}^{2d})$. Note that \mathbb{I} is lower semi-continuous and has compact level sets, and hence is a good rate function.

Assumption 2.1 *The sequence $(\tilde{A}^n, \tilde{S}^n)$ satisfies the LDP with speed b_n^{-2} and rate function \mathbb{I} in $\mathcal{D}([0, T], \mathbb{R}^{2d})$, that is:*

- for any open set $G \subset \mathcal{D}([0, T], \mathbb{R}^{2d})$

$$\liminf \frac{1}{b_n^2} \log \mathbb{P}((\tilde{A}^n, \tilde{S}^n) \in G) \geq - \inf_{\psi \in G} \mathbb{I}(\psi);$$

- for any closed set $F \subset \mathcal{D}([0, T], \mathbb{R}^{2d})$

$$\limsup \frac{1}{b_n^2} \log \mathbb{P}((\tilde{A}^n, \tilde{S}^n) \in F) \leq - \inf_{\psi \in F} \mathbb{I}(\psi).$$

A sufficient condition for this assumption is the existence of finite exponential moments for the random variables $IA_i(1)$ and $ST_i(1)$ (a precise statement is provided below, in Assumption 2.3). For a proof, as well as considerably weaker sufficient conditions, see [8].

The cost to be considered is defined in terms of functions $C_i, i \in \mathcal{I}$, that are strictly increasing, strictly convex, nonnegative, continuously differentiable functions from \mathbb{R}_+ to \mathbb{R}_+ , satisfying $C_i(0) = C'_i(0) = 0$. A linear growth condition is assumed, namely $C_i(x) \leq u_1 x$ for all $x \in \mathbb{R}_+, i \in \mathcal{I}$, where u_1 is a constant. Denote $C(x) = \sum_i C_i(x_i)$, where $x = (x_1, \dots, x_d)$. For more about the assumptions on C , see Remark 2.6 below. Given n , the cost, that is of RS type, associated with the control $B^n \in \mathcal{B}^n$ is given by

$$J^n(B^n) = \frac{1}{b_n^2} \log \mathbb{E} \left[\exp \left(b_n^2 \int_0^T C(\tilde{Q}^n(t)) dt \right) \right]. \tag{2.9}$$

The two value functions of interest are given by

$$V^n = \inf_{B^n \in \mathcal{B}^n} J^n(B^n), \quad \hat{V}^n = \inf_{B^n \in \mathcal{B}^{\#,n}} J^n(B^n).$$

Now, we describe the differential game which will govern the asymptotic behavior of the above control problem. Let $\theta = (\mu_1^{-1}, \dots, \mu_d^{-1})$ and let $y = (y_1, \dots, y_d)$ where $y_i = \tilde{\lambda}_i - \rho_i \tilde{\mu}_i$. Denote $\mathcal{P} = \mathcal{C}_0([0, T], \mathbb{R}^{2d})$, the set of continuous functions starting from 0, and

$$\mathcal{E} = \{ \zeta \in \mathcal{C}([0, T], \mathbb{R}^d) : \theta \cdot \zeta \text{ starts from zero and is non-decreasing} \}.$$

Endow both the spaces with the uniform topology. Let ρ be the mapping from $\mathcal{D}([0, T], \mathbb{R}^d)$ into itself defined by

$$\rho[\psi]_i(t) = \psi_i(\rho_i t), \quad t \in [0, T], \quad i \in \mathcal{I}.$$

Given $\psi = (\psi^{(1)}, \psi^{(2)}) \in \mathcal{P}$ and $\zeta \in \mathcal{E}$, the dynamics for the data (ψ, ζ) are defined by

$$\varphi = y\iota + \psi^{(1)} - \rho[\psi^{(2)}] + \zeta, \tag{2.10}$$

where $\iota : [0, T] \rightarrow [0, T]$ is the identity map. The dynamics are considered with the constraint

$$\varphi_i(t) \geq 0, \quad t \geq 0, i \in \mathcal{I}. \tag{2.11}$$

The game is defined in the sense of Elliot and Kalton [6], for which we need the notion of strategies. A measurable mapping $\alpha : \mathcal{P} \rightarrow \mathcal{E}$ is called a *strategy* if it satisfies a causality property. Namely, for every $\psi = (\psi^{(1)}, \psi^{(2)})$, $\tilde{\psi} = (\tilde{\psi}^{(1)}, \tilde{\psi}^{(2)}) \in \mathcal{P}$ and $t \in [0, T]$,

$$\begin{aligned} (\psi^{(1)}, \rho[\psi^{(2)}])(s) = (\tilde{\psi}^{(1)}, \rho[\tilde{\psi}^{(2)}])(s) \quad \forall s \in [0, t] \text{ implies} \\ \alpha[\psi](s) = \alpha[\tilde{\psi}](s) \quad \forall s \in [0, t]. \end{aligned}$$

A strategy α is said to be *admissible* if, whenever $\psi \in \mathcal{P}$ and $\zeta = \alpha[\psi]$, the corresponding dynamics (2.10) satisfy the non-negativity constraint (2.11). We denote the set of all admissible strategies by \mathcal{A} . Then, the cost associated with $(\psi, \zeta) \in \mathcal{P} \times \mathcal{E}$ is given by

$$c(\psi, \zeta) = \int_0^T C(\varphi(t))dt - \mathbb{I}(\psi),$$

where φ gives the dynamics for the data (ψ, ζ) . The value of the game is defined as

$$V = \inf_{\alpha \in \mathcal{A}} \sup_{\psi \in \mathcal{P}} c(\psi, \alpha[\psi]).$$

Before stating the main result, we require one more assumption.

Assumption 2.2 For any constant u ,

$$\limsup_{n \rightarrow \infty} \frac{1}{b_n^2} \log \mathbb{E} \left[e^{b_n^2 u (\|\tilde{A}^n\|_T + \|\tilde{S}^n\|_T)} \right] < \infty.$$

Again, this assumption holds when $IA_i(1)$ and $ST_i(1)$ have finite exponential moments. More precisely, consider

Assumption 2.3 There exists $u_0 > 0$ such that $\mathbb{E}[e^{u_0 IA_i(1)}]$ and $\mathbb{E}[e^{u_0 ST_i(1)}]$, $i \in \mathcal{I}$, are finite.

Then, Assumption 2.3 is a sufficient condition for Assumption 2.1 (by Theorem 6.2 of [8]) and for Assumption 2.2 (by Proposition 2.1 of [1]).

Now, we define a particular control that will be referred to as the *preemptive generalized $c\mu$ rule*. This policy gives preemptive priority to the class i for which

$\mu_i C'_i(\tilde{Q}_i^n) \geq \mu_j C'_j(\tilde{Q}_j^n)$ for all j , where ties are broken in some predefined manner. To define it precisely, we need some additional notation. Given a set of d real numbers $A = \{\alpha_i, i \in \mathcal{I}\}$, denote $\arg \max A = \{i : \alpha_i \geq \max_j \alpha_j\}$, and let $\arg \max^* A$ be the smallest member of $\arg \max A$. The control, that we denote by $B^{*,n}$, is defined by setting

$$B_i^{*,n}(t) = 1_{\{\tilde{Q}^n(t) \in Q_i\}}, \quad i \in \mathcal{I}, \tag{2.12}$$

where $Q_i, i \in \mathcal{I}$, partition $\mathbb{R}_+^d \setminus \{0\}$ according to

$$Q_i = \{q \in \mathbb{R}_+^d \setminus \{0\} : \arg \max^* \{\mu_j C'_j(q), j \in \mathcal{I}\} = i\}, \quad i \in \mathcal{I}. \tag{2.13}$$

(Thus, in case of a tie, priority is given to the lowest index.) Note that if, for some $i, q \in Q_i$, we have $q_i > 0$ thanks to the assumption that, for all $i, C'_i(x) = 0$ iff $x = 0$; as a result, the queue selected for service is nonempty.

It is easy to see that Eq. (2.12), along with Eqs. (2.2)–(2.4) and (2.6) uniquely define the processes $B^n = B^{*,n}, D^n, T^n, Q^n$ and \tilde{Q}^n , based on the data (A^n, S^n) . Moreover, the process $B^{*,n}$ thus defined is an admissible control.

It follows from the results of [1] that, under Assumptions 2.1 and 2.2,

$$\lim_{n \rightarrow \infty} V^n = V. \tag{2.14}$$

While this result is crucial in validating the heuristic that the asymptotics are governed by the differential game, it leaves open the important problem of finding explicitly computable AO policies. This problem has not been treated in [1] beyond the linear case. Our goal here is to fill in this gap for the family of costs described above.

In view of (2.14), a sequence of control policies B^n is said to be *asymptotically optimal* if it achieves the limit, namely

$$\lim_{n \rightarrow \infty} J^n(B^n) = V. \tag{2.15}$$

The first main result that we prove is the asymptotic optimality of the preemptive generalized $c\mu$ rule.

Theorem 2.4 *Let Assumptions 2.1 and 2.2 hold. Then $B^{*,n}$ is asymptotically optimal.*

The non-preemptive version of the generalized $c\mu$ rule is a control, denoted $B^{\#,n}$, that upon completion of a job selects a customer from the class i for which $\tilde{Q}_i^n \in Q_i$. Namely, if τ is any time of departure (a jump time of the process Δ^n), then $B^{\#,n}(\tau) = 1_{\{\tilde{Q}^n(\tau) \in Q_i\}}$. Note that the job departing at time τ is not counted in $\tilde{Q}^n(\tau)$, due to right-continuity. Also note that if the system is empty right after a departure, the above definition sets $B^{\#,n}(\tau) = 0$; hence, it is consistent with (2.5). One must also mention the non-idling condition: When a customer is admitted into an empty system, it is immediately served. Again, this set of conditions along with the equations alluded to above uniquely define the processes involved; these details are skipped.

Our second main result requires slightly stronger assumptions.

Theorem 2.5 *Let Assumption 2.3 hold. Then $\lim_{n \rightarrow \infty} \hat{V}^n = V$. Furthermore, $B^{\#,n}$ is asymptotically optimal.*

Remark 2.6 (a) The linear growth condition on C is required to assure finiteness of the cost J^n . The finiteness of J^n is indeed guaranteed by means of Assumption 2.2, by appealing to the linear growth condition on C and the relation (2.7). Furthermore, the linear growth condition is also used in step 5 in the proof of Theorem 4.2 in [1]. The conditions on C' are necessary for the existence of a continuous minimizing curve (see Lemma 3.1) which plays a central role in the solution of the limiting differential game. Such conditions on C' have appeared in [7], which has established the heavy traffic AO of the generalized $c\mu$ rule in the multi-server setup.

(b) Functions C that satisfy our assumptions include

$$C(x) = \sum_i (a_i + b_i x_i^{p_i})^{1/p_i} - \sum_i a_i^{1/p_i},$$

for constants $a_i > 0, b_i > 0, p_i > 1$. For example, one may take $C(x) = \sum_i (1 + b_i x_i^2)^{1/2} - d$.

3 Solution of the game

In this section, we give a minimizing strategy for the game described in the previous section. For that we require the following lemma about the existence of a continuous minimizing curve.

Lemma 3.1 *There exists a continuous function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+^d$ such that*

$$\begin{cases} \theta \cdot f(w) = w, \\ C(f(w)) = \inf\{C(q) : q \in \mathbb{R}_+^d, \theta \cdot q = w\}, \end{cases} \quad w \in \mathbb{R}_+. \quad (3.1)$$

This function satisfies

$$\mu_1 C'_1(f_1(w)) = \dots = \mu_d C'_d(f_d(w)), \quad w \in \mathbb{R}_+. \quad (3.2)$$

Proof First of all note that, by our assumptions, C'_i are continuous strictly increasing bounded functions, null at zero, for each i . Let $M_i = \mu_i \|C'_i\|_\infty$. Without loss of generality let $M_1 = \min M_i$. Then for any $c \in [0, M_1)$ there exists $q^c = (q_1^c, \dots, q_d^c)$ such that

$$(\mu_1 C'_1(q_1^c), \dots, \mu_d C'_d(q_d^c)) = c(1, \dots, 1).$$

Moreover, $q^0 = (0, \dots, 0)$. Now consider the function $F : [0, M_1) \rightarrow \mathbb{R}_+$ defined as

$$F(c) = \theta \cdot q^c.$$

Recall $\theta_i = \mu_i^{-1}$. By our assumptions, F is continuous, $F(0) = 0$ and $\lim_{c \rightarrow M_1} F(c) = \infty$. So by the intermediate value theorem, for any $w \in \mathbb{R}_+$, there exists such that $F(c) = w$. Note that the strict monotonicity of C'_i implies that of $c \mapsto q_i^c$, hence that of F . As a result, the solution c to $F(c) = w$ is unique. Hence, for every $w \in \mathbb{R}_+$, there exists a unique member of \mathbb{R}_+^d , that with an abuse of notation we denote by q^w , satisfying

$$\mu_1 C'_1(q_1^w) = \dots = \mu_d C'_d(q_d^w) \quad \text{and} \quad \theta \cdot q^w = w. \tag{3.3}$$

Let $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+^d$ be defined by $f(w) = q^w$. Note that, because of the facts that $C'_i(0) = 0$ and C'_i is strictly increasing, it follows from (3.3) that $f_i(w) > 0$ whenever $w > 0$. Thus, by a Lagrange multiplier argument and by the strict convexity assumption, we obtain that f satisfies (3.1). Relation (3.2) holds by (3.3).

Next, we prove the continuity of f . Denote $\theta_{\min} = \min_i \theta_i$. For that fix $\varepsilon > 0$, and take $\delta = \varepsilon \theta_{\min} / 2$. We claim that $|w - w'| < \delta$ implies $\sum |q_i^w - q_i^{w'}| < \varepsilon$. Suppose not. Then, there exist w and w' with $0 < w - w' < \delta$ and $\sum |q_i^w - q_i^{w'}| \geq \varepsilon$. By the identity $w = \theta \cdot f(w)$, we have $w - w' = \sum \theta_i q_i^w - \sum \theta_i q_i^{w'} < \delta$. But $\sum |\theta_i (q_i^w - q_i^{w'})| \geq \varepsilon \theta_{\min} = 2\delta$. Note that if a_i are any constants satisfying $0 < \sum a_i < \delta$ and $\sum |a_i| \geq 2\delta$ then there exist i and j for which $a_i < 0 < a_j$. Therefore, there must exist i and j such that $q_i^w > q_i^{w'}$ and $q_j^{w'} > q_j^w$. By the monotonicity of C'_i for each i , we have

$$C'_i(q_i^w) > C'_i(q_i^{w'}) \quad \text{and} \quad C'_j(q_j^w) < C'_j(q_j^{w'}).$$

This contradicts

$$\mu_i C'_i(q_i^w) = \mu_j C'_j(q_j^w), \quad \mu_i C'_i(q_i^{w'}) = \mu_j C'_j(q_j^{w'}),$$

which follows from (3.3). Hence, we have proved the claim. □

Next, we describe the solution of the game. For that, consider the one-dimensional Skorohod map Γ from $\mathcal{D}([0, T], \mathbb{R})$ to itself given by

$$\Gamma[z](t) = z(t) - \inf_{s \in [0, t]} [z(s) \wedge 0], \quad t \in [0, T]. \tag{3.4}$$

From (2.10), for $\psi = (\psi^{(1)}, \psi^{(2)}) \in \mathcal{P}$ and $\zeta \in \mathcal{E}$, the dynamics of the differential game are given by $\varphi = \xi + \zeta$, where

$$\xi = \gamma \iota + \psi^{(1)} - \rho[\psi^{(2)}].$$

We associate with each $\psi \in \mathcal{P}$ a 4-tuple $(\varphi[\psi], \xi[\psi], \zeta[\psi], \mathbf{w}[\psi])$ given by

$$\begin{aligned} \xi[\psi] &= \gamma \iota + \psi^{(1)} - \rho[\psi^{(2)}], & \mathbf{w}[\psi] &= \Gamma[\theta \cdot \xi[\psi]], \\ \varphi[\psi] &= f(\mathbf{w}[\psi]), & \zeta[\psi] &= \varphi[\psi] - \xi[\psi]. \end{aligned}$$

The following result has been proved in Proposition 3.1 of [1].

Proposition 3.2 *The map ζ is an admissible strategy. Moreover, it is a minimizing strategy, namely, $V = \sup_{\psi \in \mathcal{P}} c(\psi, \zeta[\psi])$.*

4 Proofs of main results

4.1 Proof of Theorem 2.4

In this section, we keep Assumptions 2.1 and 2.2 in force and shorten the notation $B^{*,n}$ to B^n . Our goal is to argue that

$$\limsup_n J^n(B^n) \leq V. \tag{4.1}$$

As we mentioned above, $\lim_n V^n = V$, and so Theorem 2.4 will follow once (4.1) is established.

To this end, given a constant Δ define the set

$$\mathcal{D}_\Delta = \{\psi \in \mathcal{D}([0, T], \mathbb{R}^{2d}) : \mathbb{I}(\psi) \leq \Delta\}.$$

Then, by the definition of the rate function \mathbb{I} , \mathcal{D}_Δ is a compact set containing absolutely continuous functions starting from zero, with derivative having L^2 norm uniformly bounded. By the compactness, there exists a constant M (depending on Δ) such that $\|\psi^{(1)}\|_T + \|\psi^{(2)}\|_T \leq M$ whenever $\psi \in \mathcal{D}_\Delta$. Also, by the L^2 bound, the members of \mathcal{D}_Δ are equicontinuous. For any $\tilde{\psi} \in \mathcal{D}_\Delta$ and any $r > 0$ define

$$\mathcal{A}_r(\tilde{\psi}) = \{\psi \in \mathcal{D}([0, T], \mathbb{R}^{2d}) : d(\psi, \tilde{\psi}) < r\},$$

where d is the Skorohod metric. We recall from [5] that

$$d(\psi, \tilde{\psi}) = \inf_{g \in \mathcal{Y}} \|g\|^\circ \vee \|\psi - \tilde{\psi} \circ g\|_T,$$

where \mathcal{Y} is the set of strictly increasing, continuous functions from $[0, T]$ onto itself, and

$$\|g\|^\circ = \sup_{0 \leq s < t \leq T} \left| \log \frac{g(t) - g(s)}{t - s} \right|.$$

Now, for any $g \in \mathcal{Y}$,

$$\begin{aligned} \|\psi(t) - \tilde{\psi}(t)\| &\leq \|\psi(t) - \tilde{\psi}(g(t))\| + \|\tilde{\psi}(g(t)) - \tilde{\psi}(t)\|, \quad t \in [0, T], \\ \|g(\cdot) - \cdot\|_T &\leq T(e^{\|g\|^\circ} - 1). \end{aligned}$$

Thus, by equicontinuity, for any $\eta > 0$ there exists $r > 0$ such that, for any $\tilde{\psi} \in \mathcal{D}_\Delta$,

$$\psi \in \mathcal{A}_r(\tilde{\psi}) \text{ implies } \|\psi - \tilde{\psi}\|_T < \eta. \tag{4.2}$$

Since \mathcal{D}_Δ is compact and \mathbb{I} is lower semi-continuous, it is possible to choose a finite collection of members, $\psi^1, \psi^2, \dots, \psi^N$, of \mathcal{D}_Δ , and positive constants r_1, r_2, \dots, r_N , such that, denoting $\mathcal{A}^k = \mathcal{A}_{r_k}(\psi^k)$, one has $\mathcal{D}_\Delta \subset \cup_k \mathcal{A}^k$, and

$$\inf\{\mathbb{I}(\psi) : \psi \in \overline{\mathcal{A}^k}\} \geq \mathbb{I}(\psi^k) - \frac{\varepsilon}{2}, \quad k = 1, 2, \dots, N. \tag{4.3}$$

We fix such $\{\psi^k\}$ and $\{r_k\}$. By (4.2) we may, and will, assume without loss of generality that $\max_k r_k$ is so small that

$$\psi \in \mathcal{A}^k \text{ implies } \sum_{j=1}^d \sum_{m=1}^2 \|\psi_j^{(m)} - \psi_j^{k,(m)}\|_T \leq \frac{\delta}{8\theta_{\max}}, \quad k = 1, 2, \dots, N, \tag{4.4}$$

where $\theta_{\max} = \max_j \theta_j$, and $\delta > 0$ is chosen so that for f as in Lemma 3.1,

$$\max_j \text{osc}_\delta(f_j, L) < \frac{\varepsilon}{4}, \quad L := 4\theta_{\max}(\|y\|_1 T + d(M + 1)). \tag{4.5}$$

For $k = 1, 2, \dots, N$, let

$$(\varphi^k, \xi^k, \zeta^k, w^k) = (\varphi[\psi^k], \xi[\psi^k], \zeta[\psi^k], \mathbf{w}[\psi^k]).$$

Thus, φ^k gives the dynamics corresponding to ψ^k and ζ^k . Let $\Lambda^n = \|\tilde{A}^n\|_T + \|\tilde{S}^n\|_T$, and define

$$\Omega_k^n = \{(\tilde{A}^n, \tilde{S}^n) \in \mathcal{A}^k\}, \quad k = 1, 2, \dots, N. \tag{4.6}$$

Fix $\varepsilon > 0$. Then, as in the proof of Theorem 4.2 in [1], it is enough to show that there exists a constant c_1 such that for all n sufficiently large

$$\|\tilde{Q}^n\|_T \leq c_1(1 + \Lambda^n) \tag{4.7}$$

and

$$\|\tilde{Q}^n - \varphi^k\|_T \leq c_1\varepsilon \quad \text{on } \Omega_k^n, \quad k = 1, 2, \dots, N. \tag{4.8}$$

Once the two estimates are established, the rest of the proof proceeds as in Step 5 of Theorem 4.2 in [1]. To this end, let $\theta^n = ((\mu_1^n)^{-1}n, \dots, (\mu_d^n)^{-1}n)$. Then, it follows from (2.7) and (2.12) that

$$\theta^n \cdot \tilde{Q}^n(t) = \theta^n \cdot Y^n(t) + \frac{\sqrt{n}}{b_n} \int_0^t 1_{\{\theta^n \cdot \tilde{Q}^n(s)=0\}} ds, \tag{4.9}$$

where $Y_i^n(t) = y_i^n t + \tilde{A}_i^n(t) - \tilde{S}_i^n(T_i^n(t))$. Since $\theta^n \cdot \tilde{Q}^n$ is nonnegative and the last term on the right-hand side increases only on the set of times $\{t : \theta^n \cdot \tilde{Q}^n(t) = 0\}$, it

follows that $\theta^n \cdot \tilde{Q}^n$ is a solution to the one-dimensional Skorohod problem for $\theta^n \cdot Y^n$. In particular,

$$\theta^n \cdot \tilde{Q}^n = \Gamma[\theta^n \cdot Y^n].$$

Thus, by (3.4),

$$\|\theta^n \cdot \tilde{Q}^n\|_T + \|\theta^n \cdot Z^n\|_T \leq 3\|\theta^n \cdot Y^n\|_T. \tag{4.10}$$

By the assumptions (2.1) on the limit parameters, and the definitions of the constants θ^n, θ, y^n and y , one has $\theta^n \rightarrow \theta$ and $y^n \rightarrow y$. Using the non-negativity of \tilde{Q}^n , it follows from (4.10) that (4.7) holds, for a suitable constant c_1 that does not depend on n .

Note that Λ^n is bounded by $M + 1$ on $\cup_k \Omega_k^n$. As a result, we also have the following uniform bound on \tilde{Q}^n , namely for all large n , $\|\tilde{Q}^n\|_T \leq c_2$ on the event $\cup_k \Omega_k^n$, where c_2 does not depend on n or ε (but may depend on Δ). A more concrete bound will be needed for $\theta^n \cdot \tilde{Q}^n$. Namely, on $\cup_k \Omega_k^n$,

$$\begin{aligned} \|\theta^n \cdot \tilde{Q}^n\|_T &\leq 3\|\theta^n \cdot Y^n\|_T \leq 4\theta_{\max} \left(\|y\|_1 T + \sum_{j=1}^d (\|\tilde{A}_j^n\|_T + \|\tilde{S}_j^n\|_T) \right) \\ &\leq 4\theta_{\max} \left(\|y\|_1 T + d(M + 1) \right) = L, \end{aligned} \tag{4.11}$$

by (4.5).

We next prove (4.8). Let

$$G_i^n(t) = \frac{\theta_i^n}{\theta_i} \tilde{Q}_i^n(t) - f_i(\theta^n \cdot \tilde{Q}^n(t)), \quad t \in [0, T], i \in \mathcal{I}.$$

We first show that, for all n large, on Ω_k^n , one has

$$\|G_i^n\|_T < c_3\varepsilon, \quad \text{for all } i, \tag{4.12}$$

for some constant c_3 that does not depend on n or ε . Note that a sufficient condition for (4.12) is that, for each i ,

$$\inf_{0 \leq t \leq T} G_i^n(t) \geq -\varepsilon. \tag{4.13}$$

To see this, recall that from the definition of f we have

$$\sum_i \theta_i G_i^n = 0. \tag{4.14}$$

Hence, if (4.13) holds then

$$\theta_i G_i^n = - \sum_{j:j \neq i} \theta_j G_j^n \leq \varepsilon \sum_{j:j \neq i} \theta_j,$$

by which (4.12) holds.

We thus turn to prove (4.13). Fix i . Arguing by contradiction, assume that (4.13) is false. Recall that we assume the initial condition $\tilde{Q}^n(0) = 0$. Moreover, the jump sizes of the process \tilde{Q}^n , hence those of G^n , are uniformly small when n is large, by appealing to the continuity of f and the uniform bound on \tilde{Q}^n alluded to above. By these considerations, it follows that, provided n is sufficiently large, there must exist times σ^n and τ^n such that $0 \leq \sigma^n < \tau^n \leq T$, and

$$G_i^n(\tau^n) < -\varepsilon, \quad G_i^n(\sigma^n) \geq -\frac{\varepsilon}{2}, \quad G_i^n(t) < -\frac{\varepsilon}{4} \text{ for all } t \in [\sigma^n, \tau^n]. \quad (4.15)$$

It follows from the last assertion of (4.15) that, for all n sufficiently large, for $t \in [\sigma^n, \tau^n]$,

$$\tilde{Q}_i^n(t) - f_i(\theta^n \cdot \tilde{Q}^n(t)) < 0, \quad (4.16)$$

where we used the uniform bound on \tilde{Q}^n and the convergence $\theta^n \rightarrow \theta$. Also, by (4.14) and (4.15), for every $t \in [\sigma^n, \tau^n]$ there exists a j (depending on t) such that $G_j^n(t) > c'\varepsilon$ for some constant $c' > 0$. Therefore, provided that n is large, for every $t \in [\sigma_n, \tau_n]$ there exists a j such that

$$\tilde{Q}_j^n(t) - f_j(\theta^n \cdot \tilde{Q}^n(t)) > 0. \quad (4.17)$$

By (3.2),

$$\mu_i C'_i(f_i(\theta^n \cdot \tilde{Q}^n)) = \mu_j C'_j(f_j(\theta^n \cdot \tilde{Q}^n)).$$

Using (4.16) and (4.17) and the strict monotonicity of C'_m for all m , we have that

$$\text{for every } t \in [\sigma^n, \tau^n] \text{ there exists a } j \text{ such that } \mu_i C'_i(\tilde{Q}_i^n(t)) < \mu_j C'_j(\tilde{Q}_j^n(t)). \quad (4.18)$$

By the way the policy is defined, specifically (2.12) and (2.13), we obtain that on the interval $[\sigma^n, \tau^n]$ class i does not receive any service, namely $B_i^n = 0$.

The first two statements of (4.15) yield

$$\frac{\theta^n}{\theta_i} (\tilde{Q}_i^n(\tau^n) - \tilde{Q}_i^n(\sigma^n)) + \frac{\varepsilon}{2} \leq f_i(\theta^n \cdot \tilde{Q}^n(\tau^n)) - f_i(\theta^n \cdot \tilde{Q}^n(\sigma^n)). \quad (4.19)$$

Fix a sequence $\{a_n\}$ having the properties that $a_n \rightarrow 0$ but $\sqrt{nb_n}^{-1} a_n \rightarrow \infty$. We distinguish between two cases.

Case I: $\tau^n - \sigma^n \leq a_n$.

Since on the interval $[\sigma^n, \tau^n]$ class i receives no service, the queue length can only increase over the interval, as follows from the balance equation (2.4), (2.2), and $B_i^n = 0$. Thus $\tilde{Q}_i^n(\tau^n) \geq \tilde{Q}_i^n(\sigma^n)$. We obtain from (4.19) that

$$f_i(\theta^n \cdot \tilde{Q}^n(\tau^n)) - f_i(\theta^n \cdot \tilde{Q}^n(\sigma^n)) \geq \frac{\varepsilon}{2} > \frac{\varepsilon}{4}. \tag{4.20}$$

(Working with the bound $\varepsilon/4$ instead of $\varepsilon/2$ will be useful in the proof provided in Sect. 4.2.) Recalling that $\theta^n \cdot \tilde{Q}^n$ is uniformly bounded by L (see (4.11)), we have from (4.5) that

$$\theta^n \cdot \tilde{Q}^n(\tau^n) - \theta^n \cdot \tilde{Q}^n(\sigma^n) \geq \delta,$$

where δ is as in (4.5). From (4.16), it follows that in the interval $[\sigma^n, \tau^n]$, $\theta^n \cdot \tilde{Q}^n > 0$, because the function f vanishes at zero. Thus, by (4.9),

$$\theta^n \cdot (Y^n(\tau^n) - Y^n(\sigma^n)) \geq \delta.$$

Hence, using the definition of Y^n , the convergence $\theta^n \rightarrow \theta$, and the fact that T_i^n is Lipschitz with constant 1, we obtain

$$\delta \leq 2\theta_{\max} \left[2\|y\|_1 a_n + \sum_{j=1}^d \text{osc}_{a_n}(\tilde{A}_j^n) + \sum_{j=1}^d \text{osc}_{a_n}(\tilde{S}_j^n) \right].$$

Now, on Ω_k^n ,

$$\text{osc}_{a_n}(\tilde{A}_j^n) \leq 2\|\tilde{A}_j^n - \psi_j^{k,(1)}\|_T + \text{osc}_{a_n}(\psi_j^{k,(1)}).$$

A similar statement holds for \tilde{S}^n . Hence, on Ω_k^n ,

$$\delta \leq 2\theta_{\max} \left[2\|y\|_1 a_n + \frac{\delta}{4\theta_{\max}} + \sum_{j=1}^d \sum_{m=1}^2 \left\{ \text{osc}_{a_n}(\psi_j^{k,(m)}) \right\} \right],$$

where we used (4.4). By the equicontinuity of the functions in the class \mathcal{D}_Δ , we obtain $\delta \leq \delta/2 + \eta_n$ for some $\eta_n \rightarrow 0$. Hence, for all n large enough we arrive at a contradiction.

Case 2: $\tau^n - \sigma^n \geq a_n$.

On Ω_k^n , the right-hand side of (4.19) is bounded. As for its left-hand side, we have from (2.7) that

$$\begin{aligned} \tilde{Q}_i^n(\tau^n) - \tilde{Q}_i^n(\sigma^n) &= y_i^n(\tau^n - \sigma^n) + \tilde{A}_i^n(\tau^n) - \tilde{A}_i^n(\sigma^n) \\ &\quad - (\tilde{S}_i^n(T_i^n(\tau^n)) - \tilde{S}_i^n(T_i^n(\sigma^n))) \\ &\quad + \frac{\mu_i^n}{n} \frac{\sqrt{n}}{b_n} \rho_i(\tau^n - \sigma^n), \end{aligned} \tag{4.21}$$

where we used the fact that T_i^n does not increase on the interval $[\sigma_n, \tau_n]$. On the right-hand side of (4.21), all terms are bounded, except the last term which tends to infinity as $n \rightarrow \infty$. Thus again for large n , we arrive at a contradiction.

Summarizing the two cases, we conclude that (4.13), and hence (4.12), holds on Ω_k^n provided that n is sufficiently large.

Finally, we use (4.12) to show (4.8). Since $\theta^n \rightarrow \theta$ and \tilde{Q}^n is bounded on Ω_k^n , there exists a constant c_4 such that on Ω_k^n , for all large n ,

$$\|\tilde{Q}^n - f(\theta^n \cdot \tilde{Q}^n)\|_T < c_4\varepsilon. \tag{4.22}$$

Now using the facts that $\theta^n \cdot \tilde{Q}^n = \Gamma[\theta^n \cdot Y^n]$ and $\theta \cdot \varphi^k = \Gamma[\theta \cdot \xi^k]$, where $\xi^k = y_l + \psi^{k,(1)} - \rho[\psi^{k,(2)}]$, and the Lipschitz continuity of the map Γ , we have for all n sufficiently large,

$$\begin{aligned} \|\theta^n \cdot \tilde{Q}^n - \theta \cdot \varphi^k\|_T &\leq 2\|\theta^n \cdot Y^n - \theta^n \cdot \xi^k\|_T + 2\|\theta^n - \theta\|_T \\ &\leq c_5\|Y^n - \xi^k\|_T + \varepsilon \\ &\leq c_5[\|\tilde{A}^n - \psi^{k,(1)}\|_T + \|\tilde{S}^n \circ T^n - \rho[\psi^{k,(2)}]\|_T] + 2\varepsilon. \end{aligned}$$

Now by (4.10) and the definition of Z^n , we have

$$\|\rho_i l - T_i^n\|_T \rightarrow 0, \quad \text{as } n \rightarrow \infty.$$

Thus, on Ω_k^n , for all sufficiently large n , there exists a constant c_6 such that

$$\|\theta^n \cdot \tilde{Q}^n - \theta \cdot \varphi^k\|_T < c_6\varepsilon. \tag{4.23}$$

Now, from the definition of φ^k we have $\varphi^k = f(\theta \cdot \varphi^k)$. Since, on Ω_k^n , \tilde{Q}^n and φ^k are bounded, by the continuity of f there exists a constant c_7 such that

$$\|f(\theta^n \cdot \tilde{Q}^n) - \varphi^k\|_T < c_7\varepsilon. \tag{4.24}$$

Combining (4.22) and (4.24), we obtain (4.8). This concludes the proof. □

4.2 Proof of Theorem 2.5

Since the control space $\mathcal{B}^{\#,n}$ is contained in \mathcal{B}^n , we automatically have $\hat{V}^n \geq V^n$. Hence, it suffices to prove the asymptotic optimality of the non-preemptive generalized $c\mu$ rule, $B^{\#,n}$. The proof for that follows the same steps as for the preemptive case with some modifications, which we describe next. The first deviation from the proof of Theorem 2.4 is related to (4.18). While (4.18) is a valid statement, we cannot deduce from it that $B_i^n = 0$ on the interval $[\sigma^n, \tau^n]$ (here and in what follows, we suppress the symbol # in $B^{\#,n}$). Indeed, (4.18) assures that every time the server *becomes available* during this interval customers from classes other than i are selected for service, but it is possible that at the time σ^n a class- i customer is already in service. In this case, one has $B_i^n = 1$ until this service completes. Note that (4.19) remains valid. In what follows, we describe how Cases 1 and 2 of the proof of Theorem 2.4 are modified to address this change. With the modified version of these two cases, the remainder of the proof proceeds as before.

For Case 1, one can no longer say that class i receives no service. However, at most one customer of that class is served during $[\sigma^n, \tau^n]$, by which $\tilde{Q}_i^n(\tau^n) - \tilde{Q}_i^n(\sigma^n) \geq -b_n^{-1}n^{-1/2}$. With (4.19), using $\theta_i^n \rightarrow \theta_i$ and assuming n is large, we obtain

$$f_i(\theta^n \cdot \tilde{Q}^n(\tau^n)) - f_i(\theta^n \cdot \tilde{Q}^n(\sigma^n)) \geq \frac{\varepsilon}{2} - \frac{2}{b_n\sqrt{n}} > \frac{\varepsilon}{4}.$$

As a result, the final conclusion of (4.20) is still valid, and we can proceed exactly as before.

As for Case 2, let

$$ST_j^n = \max_{l \leq A_j^n(T)} \frac{ST_j(l)}{\mu_j^n}$$

denote the maximal class- j service time among all jobs arriving up to time T . Equation (4.21) need not hold because of the possible customer in service at time σ^n . However, the time devoted to this customer is bounded by ST_i^n . As a result, in place of (4.21), we may write

$$\begin{aligned} \tilde{Q}_i^n(\tau^n) - \tilde{Q}_i^n(\sigma^n) &\geq y_i^n(\tau^n - \sigma^n) + \tilde{A}_i^n(\tau^n) - \tilde{A}_i^n(\sigma^n) \\ &\quad - (\tilde{S}_i^n(T_i^n(\tau^n)) - \tilde{S}_i^n(T_i^n(\sigma^n))) \\ &\quad + \frac{\mu_i^n \sqrt{n}}{n} \rho_i(\tau^n - \sigma^n) - \frac{\mu_i^n \sqrt{n}}{n} \rho_i ST_i^n. \end{aligned}$$

The proof will proceed as before if we show that the last term in the above display is bounded off an event of negligible probability. More precisely, it suffices to show that for every constant c_1 , we have for large n ,

$$\mathbb{P}\left(\frac{\sqrt{n}}{b_n} \max_j ST_j^n > 1\right) \leq e^{-c_1 b_n^2}. \tag{4.25}$$

To this end, note first that, since \tilde{A}^n satisfies the MDP, there exists a constant c_3 such that for any constant c_1 , for all n large enough,

$$\mathbb{P}(\max_j A_j^n(T) > c_3 n) \leq e^{-c_1 b_n^2}.$$

Then, with c_3 as above we have

$$\begin{aligned} \mathbb{P}\left(ST_j^n > \frac{b_n}{\sqrt{n}}\right) &\leq c_3 n \mathbb{P}\left(ST_j(1) > \frac{\mu_j^n b_n}{\sqrt{n}}\right) + \mathbb{P}(A_j^n(T)) > c_3 n \\ &\leq c_3 n \mathbb{P}\left(ST_j(1) > \frac{\mu_j b_n \sqrt{n}}{2}\right) + e^{-c_1 b_n^2} \\ &\leq c_4 n e^{-\frac{u_0 \mu_j}{2}(b_n \sqrt{n})} + e^{-c_1 b_n^2} \\ &\leq 2e^{-c_1 b_n^2}, \end{aligned}$$

where c_4 is a constant, and the second last line follows from Assumption 2.3. It follows that for any c_1 , (4.25) holds for all n large enough.

This completes the proof. \square

References

1. Atar, R., Biswas, A.: Control of the multiclass G/G/1 queue in the moderate deviation regime. *Ann. Appl. Probab.* **24**, 2033–2069 (2014)
2. Atar, R., Cohen, A.: A differential game for a multiclass queueing model in the moderate-deviation heavy-traffic regime. *Math. Oper. Res.* **41**(4), 1354–1380 (2016)
3. Atar, R., Cohen, A.: Asymptotically optimal control for a multiclass queueing model in the moderate deviation heavy traffic regime. Preprint (2015)
4. Atar, R., Goswami, A., Shwartz, A.: Asymptotic optimality of a fixed priority rule for a queueing problem in the large deviation regime. *Electron. Commun. Probab.* **19**(11), 1–13 (2014)
5. Billingsley, P.: *Convergence of Probability Measures*, 2nd edn. Wiley, New York (1999)
6. Elliot, R.J., Kalton, N.J.: *The Existence of Value in Differential Games*. American Mathematical Society, Providence (2014)
7. Mandelbaum, A., Stolyar, A.: Scheduling flexible servers with convex delay costs: heavy-traffic optimality of the generalized $c\mu$ rule. *Oper. Res.* **52**, 836–855 (2004)
8. Puhalskii, A.A., Whitt, W.: Functional large deviation principles for first-passage-time processes. *Ann. Appl. Probab.* **7**(2), 362–381 (1997)
9. van Mieghem, J.A.: Dynamic scheduling with convex delay costs: the generalized $c\mu$ rule. *Ann. Appl. Probab.* **5**(3), 809–833 (1995)