

MANY-SERVER ASYMPTOTICS FOR JOIN-THE-SHORTEST-QUEUE: LARGE DEVIATIONS AND RARE EVENTS

BY AMARJIT BUDHIRAJA¹, ERIC FRIEDLANDER² AND RUOYU WU³

¹*Department of Statistics and Operations Research, University of North Carolina, budhiraj@email.unc.edu*

²*Department of Ecology and Evolution, University of Chicago, efriedlander@uchicago.edu*

³*Department of Mathematics, Iowa State University, ruoyu@iastate.edu*

The join-the-shortest-queue routing policy is studied in an asymptotic regime where the number of processors n scales with the arrival rate. A large deviation principle (LDP) for the occupancy process is established, as $n \rightarrow \infty$, in a suitable infinite-dimensional path space. Model features that present technical challenges include, Markovian dynamics with discontinuous statistics, a diminishing rate property of the transition probability rates, and an infinite-dimensional state space. The difficulty is in the proof of the Laplace lower bound which requires establishing the uniqueness of solutions of certain infinite-dimensional systems of controlled ordinary differential equations. The LDP gives information on the rate of decay of probabilities of various types of rare events associated with the system. We illustrate this by establishing explicit exponential decay rates for probabilities of long queues. In particular, denoting by $E_j^n(T)$ the event that there is at least one queue with j or more jobs at some time instant over $[0, T]$, we show that, in the critical case, for large n and T , $\mathbb{P}(E_j^n(T)) \approx \exp[-\frac{n(j-2)^2}{4T}]$.

1. Introduction. Consider a system of n parallel processors, each processing jobs in its queue at rate 1. Jobs enter the system at rate $n\lambda_n$ with $\lambda_n \rightarrow \lambda \in (0, \infty)$ as $n \rightarrow \infty$. Service times and inter-arrival times are exponentially distributed and are mutually independent. Upon arriving, each job is routed to the shortest available queue by a central dispatcher. This is known as the join-the-shortest-queue (JSQ) routing policy and is a popular model for load balancing among distributed resources in parallel-processing systems that arise in applications of cloud computing, file transfers, database look-ups, etc. (see the survey article [38] and references therein). Denote by $X_i^n(t)$ the proportion of queues at time instant t with i or more jobs. This occupancy process $X^n(t) \doteq (X_1^n(t), X_2^n(t), \dots)$ is a convenient state descriptor for this system. In this work we establish a large deviation principle (LDP) for X^n in the path space $\mathbb{D}_{\mathbb{R}^\infty}$, where for a Polish space S , \mathbb{D}_S denotes the space of all maps from $[0, T]$ to S that are right continuous and have left limits, equipped with the usual Skorokhod topology. This result gives a characterization of exponential decay rates for events of the form $\mathbb{P}(X^n \in H)$, where H is a suitable Borel set in $\mathbb{D}_{\mathbb{R}^\infty}$, in terms of the associated rate function (see Theorem 2.4 for a precise statement). The rate function takes a variational form and is given as the value function of an infinite-dimensional deterministic optimal control problem (see (2.8)).

In general this control problem is intractable and in order to obtain useful information, beyond the fact that certain probabilities of interest converge to 0 at an exponential rate, one needs approximations, for example, by computing costs for sub-optimal control actions. Nevertheless, for some events of interest, one can say more. We illustrate this by studying

Received April 2019; revised March 2020.

MSC2020 subject classifications. 60F10, 90B15, 91B70, 60J75, 34H05.

Key words and phrases. Large deviations, load balancing, discontinuous statistics, diminishing rates, JSQ, jump-Markov processes in infinite dimensions, calculus of variations, infinite-dimensional Skorokhod problem, golden ratio.

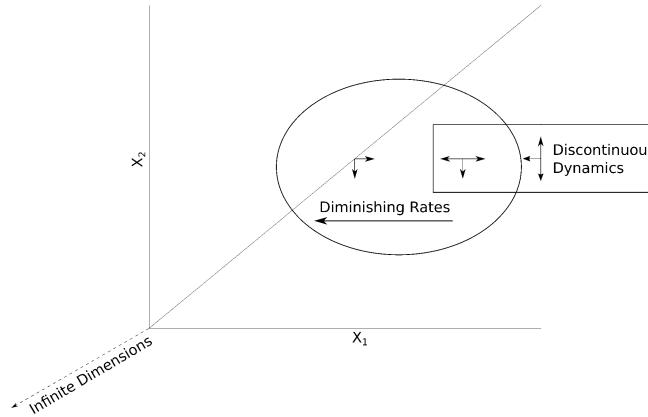


FIG. 1. Three key features of the model: infinite dimensions, discontinuous dynamics, diminishing rates.

the decay rate of probabilities of long queues. For this, we restrict attention to the critical case $\lambda_n \rightarrow 1$ and initial configuration $X_j^n(0) = \mathbf{1}_{\{j=1\}}$ (i.e., all queues are length 1 at time 0). Consider the set $E_j^n(T)$ that represents the event that there is at least one queue with j or more jobs at some time instant over $[0, T]$. In Theorem 2.5 we give an explicit characterization of the exponential decay rate of the probability of such events for $j \geq 3$. In particular, when $j = 3$ and $T = 1$ (or, more generally, when $j - 2 = T$), we obtain the following formula (see (2.13)) in terms of the golden ratio:

$$\mathbb{P}(E_j^n(T)) \approx \exp\left[-nT\left(\ell\left(\frac{1+\sqrt{5}}{2}\right) + \ell\left(\frac{-1+\sqrt{5}}{2}\right)\right)\right] \quad \text{for large } n,$$

where $\ell(x) \doteq x \log x - x + 1$ for $x \geq 0$. For long time horizons, the decay rates take an even more simple form, namely we show that for large n and large T ,

$$(1.1) \quad \mathbb{P}(E_j^n(T)) \approx \exp\left[-\frac{n(j-2)^2}{4T}\right].$$

Although not pursued in this work, techniques used to establish the above explicit asymptotic rates are expected to be useful for other types of events as well, see Conjecture 2.6.

There are several technical challenges in establishing the LDP on the path space (namely Theorem 2.4). These stem from three key features of the model that are illustrated in Figure 1.

Infinite-dimensional dynamics. The state process X^n is an infinite-dimensional Markov process and there is a nontrivial coupling between the different coordinates of the process. As a consequence, the associated rate function in the large deviation analysis is characterized through an infinite-dimensional control problem. Some recent works that have studied large deviation properties of jump-Markov processes in infinite dimensions include [5, 9, 14].

Discontinuous dynamics. The model considered here falls in the class of Markov processes with discontinuous statistics. Roughly speaking, this means that the transition rates change discontinuously at the interface of different regions of the state space. Large deviation analysis of such systems is technically challenging and has been the focus of several works [2, 3, 15, 16, 18, 26, 27, 34]. In the current work an additional challenge is that there are infinitely many regions across which transition behavior changes discontinuously.

Diminishing rates. In the study of large deviation properties of Markov processes for which the transition probability rates decrease (continuously) to zero along some directions, one is led to local rate functions that have poor regularity properties. Some works that have treated large deviation properties of such systems include [1, 5, 20, 30, 37]. The model considered in

the current work has similar features, in fact the setting here is more challenging in that the process X^n may switch among infinitely many regions in which rates diminish along different directions.

We note that all the papers referenced above, except [5], include only one of the three features noted above, while the paper [5] has the first and third feature but not the second. The combination of the three features described above is the main technical challenge in this work and most of the arguments in Section 5, which is the heart of this work, are devoted to overcoming these challenges in establishing a certain uniqueness property.

A LDP for a JSQ system has been obtained in [34] (see also [35]). However, the scaling regime considered in these works is very different from the one of interest here. Specifically, they consider a setting with a *fixed* number of queues for which the arrival and service rates are scaled up by a factor of n , and the LDP is established for the finite-dimensional queue length process scaled down by a factor of $\frac{1}{n}$. In this regime, [34] is in fact able to allow general arrival time distributions, different service time distribution parameters for different queues, and a weighted version of the JSQ policy. For the scaling regime considered here, in which the arrival rate and number of queues n approach infinity, we establish a LDP for the occupancy process (equivalent to the empirical distribution) of n queue length processes. While sacrificing some of the generality of [34], this approach provides insights into questions and situations not addressed in [34]. Specifically, the results here shed light on how probabilities of rare events decay as the size of the system is increased and the load on the system remains constant. Suppose, for example, one is interested in deciding the size of a JSQ system (i.e., the number of servers n), in which each queue has a finite buffer of size k , so that the system experiences an overflow event over the time horizon $[0, T]$ (namely, over the interval $[0, T]$, an arrival occurs to at least one of the queues with a full buffer) with probability at most e^{-6} . Then, in the critical case $\lambda_n \rightarrow 1$ and for an initial configuration where all queues are length 1 at time 0, a calculation based on the asymptotic formula (1.1) tells us that, roughly, one should take $n \approx 24T/(k-1)^2$. Such questions cannot be readily analyzed from the LDP for a fixed size JSQ system established in [34]. Indeed, since the queue lengths in the analysis of [34] are scaled down by a factor of n , the results there will say that, for any fixed size system, the asymptotic (under their scaling) probability that at least one queue will attain a length of m is the same for all $m > 0$. We also note that the proof techniques here are very different than those employed in [34]. In particular, as noted previously, unlike [34], the state descriptor here is infinite-dimensional and the transition rates approach zero in certain directions, which requires a careful analysis of an infinite-dimensional Skorokhod problem and a delicate analysis of an important uniqueness property.

Queuing systems with many parallel servers in the regime where the arrival rate scales with the number of servers have been studied extensively. A significant portion of this body of work concerns the setting in which, upon arrival, jobs join a global queue and a large pool of servers processes jobs from this queue in a FIFO fashion. One of the first works on such queuing systems is by Halfin and Whitt [24] in which the number of servers and arrival rate increase to infinity while the load ρ_n approaches one from below such that $(1 - \rho_n)\sqrt{n} \rightarrow \beta \in (0, \infty)$. Thereafter, this type of scaling has been often referred to as the Halfin–Whitt regime. The setting considered in the current work is different from Halfin–Whitt type queuing systems with a global queue and is motivated by applications in which load balancing is an important concern. In terms of analysis, the setting considered here requires tracking an infinite-dimensional state instead of the size of a single queue. Some of the works that have considered the asymptotics of a JSQ system under a scaling of the form considered in the current work are [4, 8, 21, 33]. In particular, [21] proves a central limit theorem under the heavy traffic scaling $(1 - \lambda_n)\sqrt{n} \rightarrow \beta \in (0, \infty)$ while [33] gives a law of large numbers (LLN) result under the more general condition $\lambda_n \rightarrow \lambda \in (0, \infty)$.

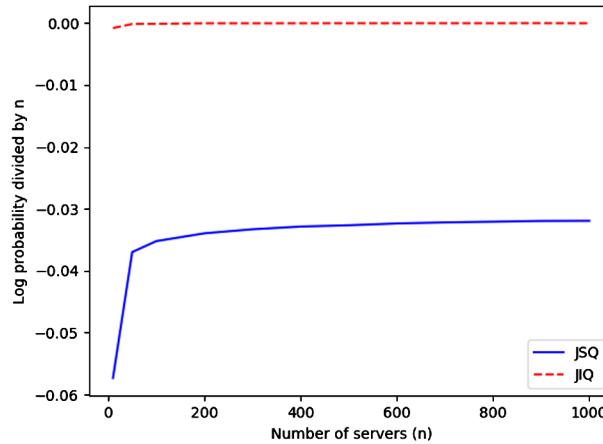


FIG. 2. Monte Carlo estimates for the rate of decay of the probability of $E_3^n(10)$ with $\lambda_n = 0.99$ for JSQ and JIQ starting with all queues of length 1.

Many other load balancing policies have been studied in the literature (see [7, 13, 22, 23, 31, 32, 36, 39] and references therein) and a survey of recent advances can be found in [38]. In particular, [32] considers the JIQ routing policy in which an incoming job is routed to an idle server, if available, and according to another routing policy (e.g., uniformly at random) if there are no idle servers in the system. The authors use a coupling argument to show that, under the heavy traffic condition of [21], JIQ and JSQ behave the same under the diffusion scaling. While JSQ and JIQ look similar in the LLN limit and under the diffusion scaling, the statistical tail behavior of the two systems is expected to be quite different. For example, Figure 2 gives Monte Carlo estimates of $\frac{1}{n} \log \mathbb{P}(E_3^n(10))$ for JSQ and JIQ with $\lambda_n = 0.99$, starting with all queues of length 1. As is clear from this figure, the performance differences between the two policies are more clearly revealed when one considers extreme tail events viewed under a large deviation scaling (see [23] for an alternative point of view for differentiating the performance of JSQ and JIQ systems). The LDP established in this work characterizes the tail statistical behavior of the JSQ system. Establishing a similar result for the JIQ system presents significant new challenges that arise from the analysis of events that include time instants where no idle servers are present. This study will be taken up in a future work. Although not considered here, the large deviation principle given in this work also provides a starting point for developing efficient importance sampling schemes for estimating probabilities of rare events in a fixed size JSQ system (see also Remark 2.4).

We now comment on the proof idea of our main result, Theorem 2.4. The starting point of our analysis is to introduce a convenient representation for the evolution of the state process X^n . This is given in (2.1)–(2.2) through an infinite collection of i.i.d. Poisson random measures (PRM) where each of the PRM corresponds to a stream of events (i.e., arrivals or job completions). The use of PRM allows us to introduce controls on the rates of these events. Roughly speaking, one can view these controls as pushing the state process away from its LLN limit to a “rare-event trajectory” while the cost incurred by the control for perturbing the state process will determine the exponential decay rate of probabilities for such rare paths. The state process has an equivalent and simpler description given in (2.6) through which it can be viewed as a solution of an infinite-dimensional Skorokhod problem for a *free process* Y^n with sample paths in $\mathbb{D}_{\mathbb{R}^\infty}$. The existence and uniqueness of solutions of this Skorokhod problem is established in Section 2.2. The main results of this work, Theorem 2.4, will not only give a LDP for X^n but in fact for the pair (X^n, Y^n) in $\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$. For the proof of the LDP we consider its equivalent formulation in terms of a Laplace principle (see, e.g., [17]). Specifically, to prove the LDP in Theorem 2.4, it suffices to show that every continuous and

bounded function G on the path space $\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$ satisfies the upper and lower bounds associated with this Laplace principle given in (2.11) and (2.12), respectively, and the function I_T in these equations has compact sub-level sets making it a rate function.

The key ingredient in establishing these results is a variational representation for expected values of exponential functionals of PRM established in [12]. This result can be applied to give variational formulas for the expected values on the left sides of the Laplace principle bounds, (2.11) and (2.12). These formulas are given in terms of controlled analogues of the PRM and state process. Recalling that the uncontrolled PRM correspond to streams of job arrivals or completions, each controlled PRM is constructed using a control process that suitably modulates the rate at which the corresponding stream of events occurs. Using controls $\varphi = (\varphi_i)_{i \in \mathbb{N}_0}$, where φ_i is the control associated with the i th PRM, the controlled state process (\bar{X}^n, \bar{Y}^n) is defined through (3.1). The variational representation from [12] allows us to express the expected values of interest in (2.11) and (2.12) in terms of an infimum of costs associated with these controls and the associated controlled state process (see Lemma 3.1). The utility of this representation is twofold. First, it reduces the majority of the proof of the LDP to arguing tightness and characterization of limits of sequences of carefully chosen controls and controlled state processes. In particular, the limits points (ζ, ψ) of the controlled state processes can be characterized as solutions to a system of controlled ordinary differential equations (ODE), expressed in (2.9)–(2.10) driven by limit control processes $\{\varphi_i\}$. Second, the representation hints at the form of the rate function I_T for the LDP. In view of the above characterization of limiting controlled state processes and in comparing the infimum in Lemma 3.1 to the desired right side of Laplace asymptotics in (2.11) and (2.12), namely,

$$\inf_{(\zeta, \psi) \in \mathcal{C}_T} \{I_T(\zeta, \psi) + G(\zeta, \psi)\},$$

where the space \mathcal{C}_T of trajectories in $\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$ is described in Section 2.3, it is natural to conjecture the following form for the rate function evaluated on a given pair of trajectories $(\zeta, \psi) \in \mathcal{C}_T$. Consider the class $\mathcal{S}_T(\zeta, \psi)$ of all controls $\{\varphi_i\}$ for which the given pair (ζ, ψ) solves the controlled ODE in (2.9)–(2.10). Then the rate function, $I_T(\zeta, \psi)$ suggested by the above considerations is the infimum of the cost,

$$\sum_{i=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_i \ell(\varphi_i(s, y)) ds dy,$$

over all controls $\{\varphi_i\} \in \mathcal{S}_T(\zeta, \psi)$, where ϑ_i is defined above (2.8).

In order to prove the upper bound we consider, for each n , a near-optimal control and controlled process for the infimum in Lemma 3.1, establish the tightness of the sequence of such processes in a suitable space, and characterize the weak limit points of the sequence. This is done in Sections 3.2–3.3. From these properties the Laplace upper bound in (2.11) follows by standard arguments that use Fatou's lemma and certain lower semicontinuity properties, as shown in Section 4. The proof that the function I_T has compact sublevel sets and, thus, is a rate function has many similarities to the proof of the Laplace upper bound and is provided in Section 6.

The main technical challenge in this work is in the proof of the Laplace lower bound (2.12). For this, one starts with the variational expression on the right side of the inequality, namely,

$$\inf_{(\zeta, \psi) \in \mathcal{C}_T} \{I_T(\zeta, \psi) + G(\zeta, \psi)\}.$$

The basic idea is to select a trajectory (ζ^*, ψ^*) in \mathcal{C}_T that is ε -optimal for the above infimum and then select a control $\varphi^* \in \mathcal{S}_T(\zeta^*, \psi^*)$ driving this trajectory which is ε -optimal for the rate function evaluated at (ζ^*, ψ^*) , that is, $I_T(\zeta^*, \psi^*)$. In view of the variational representation of the Laplace functional of interest, given in Lemma 3.1, one would like to construct

a sequence of controlled state processes of the form in (3.1) that converge to (ζ^*, ψ^*) such that the associated cost converges in an appropriate manner as well. There is a natural choice for a sequence of controlled processes that one can attempt to implement for this purpose and it is relatively easy to show that this sequence of controlled processes has the needed tightness properties and that the weak limit points $(\bar{\varphi}, \bar{\zeta}, \bar{\psi})$ of the controls and controlled state processes satisfy $\bar{\varphi} = \varphi^* \in \mathcal{S}_T(\bar{\zeta}, \bar{\psi})$.

However this is where one faces the main obstacle. From the above characterizations of the limit points it is not possible to deduce, in general, that $(\bar{\zeta}, \bar{\psi}) = (\zeta^*, \psi^*)$. The issue is regarding the uniqueness of solutions of the infinite system of controlled ODE described by (2.9)–(2.10) (considered with (ζ, ψ, φ) replaced with $(\zeta^*, \psi^*, \varphi^*)$). Namely, if one is able to say that for a given φ^* there is a unique pair (ζ^*, ψ^*) satisfying this system of equations then one obtains the desired result $(\bar{\zeta}, \bar{\psi}) = (\zeta^*, \psi^*)$. Showing uniqueness of such a system of equations is hard in general and in fact may not hold. Most of Section 5 is devoted to overcoming this challenge. The main result is Lemma 5.1 which says that one can replace (ζ^*, ψ^*) by a nearby pair of trajectories (ζ, ψ) for which the desired uniqueness property does hold with an appropriate near optimal control φ . Key ingredients in the proof of this lemma are Lemmas 5.2–5.4 and together these three lemmas, which perform several delicate surgeries on the original infinite dimensional trajectory (ζ^*, ψ^*) and the associated control φ^* , are at the technical heart of this work. Additional comments on the proofs of these lemmas are given in Section 5.

The paper is organized as follows. Section 2.1 introduces the state dynamics in terms of an infinite collection of PRM and also gives an equivalent representation through an infinite-dimensional Skorokhod map. The properties of this map are studied in Section 2.2. In preparation for the main result, in Section 2.3, we introduce the rate function that governs the LDP. The main result, Theorem 2.4, is then given in Section 2.4. This section also presents Theorem 2.5 which gives our main result on exponential decay rates for probabilities of long queues as an illustration of applications of Theorem 2.4. Other possible applications of this result are discussed briefly in Conjecture 2.6. Section 3 introduces the main variational representation that is the starting point of our analysis and establishes preliminary tightness and limit characterization results that are used in both the Laplace upper bound and lower bound proofs. Proof of the Laplace upper bound (i.e., (2.11)) is completed in Section 4 while the lower bound (i.e., (2.12)) is taken up in Section 5. Section 6 shows that the function I_T introduced in Section 2.3 is indeed a rate function. The results of Sections 4, 5, and 6 together complete the proof of Theorem 2.4. Finally Section 7 gives the proof of Theorem 2.5.

1.1. Notation. The following notation will be used. Fix $T < \infty$. All stochastic processes will be considered over the time horizon $[0, T]$. We denote the Lebesgue measure on a Euclidean space as Leb . Let S be a Polish space. For a set $B \in S$ we denote the closure of B as \bar{B} . The Borel σ -field on S will be denoted as $\mathcal{B}(S)$. Denote by \mathbb{D}_S the collection of all maps from $[0, T]$ to S that are right continuous and have left limits. This space is equipped with the usual Skorokhod topology. Similarly \mathbb{C}_S is the space of all continuous maps from $[0, T]$ to S equipped with the uniform topology. A sequence of \mathbb{D}_S valued random variables is said to be \mathbb{C} -tight if it is tight in \mathbb{D}_S and any weak limit point takes values in \mathbb{C}_S a.s. The space of all continuous and bounded real valued functions on S will be denoted as $\mathbb{C}_b(S)$. For a bounded map $x : S \rightarrow \mathbb{R}$, let $\|x\|_\infty \doteq \sup_{s \in S} |x(s)|$.

Let $(\mathbb{X}, \|\cdot\|)$ be a metric space. We denote by \mathbb{X}^∞ the set of all sequence $x = \{x_i\}_{i \in \mathbb{N}}$ such that $x_i \in \mathbb{X}$ for all $i \in \mathbb{N}$. \mathbb{X}^∞ is equipped with the product topology, which is metrized with

$$d(x, y) \doteq \sum_{i=1}^{\infty} \frac{\|x_i - y_i\| \wedge 1}{2^i}.$$

Let $\ell(z) \doteq z \log(z) - z + 1$ for $z \geq 0$.

2. Model and results. In this section we will describe the model of interest and present our main results. We begin by giving a precise mathematical formulation of the JSQ system in Section 2.1. The state process can equivalently be described through a certain infinite-dimensional Skorokhod map. This map is introduced and studied in Section 2.2. In Section 2.3 we introduce the large deviation problem of interest and present the rate function for the associated LDP. Section 2.4 presents the main result of this work (Theorem 2.4) which, in particular, gives a large deviation principle for the queue occupancy process X^n as the number of servers (and arrival rate) approaches infinity. This LDP can be used to extract information about probabilities of various types of rare events and in Theorem 2.5 we present one such application that provides estimates for probabilities of occurrence of “large queues”.

2.1. Model description. Consider a system of n parallel servers each maintaining its own queue. Jobs arrive in the system according to a Poisson process with rate $n\lambda_n$ where $\lambda_n \rightarrow \lambda$ for some $\lambda \in (0, \infty)$. When a job enters the system, a central dispatcher queries each server and routes the job to the server with the shortest queue. If there are multiple shortest queues, then the tie is broken uniformly at random. Each server processes jobs in its queue using the FIFO protocol and the service times are exponential with mean 1. We assume that the inter-arrival times and service times are mutually independent. The state of the system at time t can be represented as $X^n(t) = (X_0^n(t), X_1^n(t), \dots)$ where $X_i^n(t)$ corresponds to the proportion of queues which are of length i or longer at time t . Note that $X_i^n(t) \in [0, 1]$ and $1 = X_0^n(t) \geq X_1^n(t) \geq X_2^n(t) \geq \dots$ for all $t \in [0, T]$.

We will now give a convenient evolution equation for the state process in terms of a collection of Poisson random measures. For a locally compact metric space \mathbb{S} , let $\mathcal{M}_{\text{FC}}(\mathbb{S})$ represent the space of measures ν on $(\mathbb{S}, \mathcal{B}(\mathbb{S}))$ such that $\nu(K) < \infty$ for every compact $K \in \mathcal{B}(\mathbb{S})$, equipped with the usual vague topology. This topology can be metrized such that $\mathcal{M}_{\text{FC}}(\mathbb{S})$ is a Polish space (see [11] for one convenient metric). A PRM D on \mathbb{S} with mean measure (or intensity measure) $\nu \in \mathcal{M}_{\text{FC}}(\mathbb{S})$ is an \mathcal{M}_{FC} -valued random variable such that for each $H \in \mathcal{B}(\mathbb{S})$ with $\nu(H) < \infty$, $D(H)$ is a Poisson random variable with mean $\nu(H)$ and for disjoint $H_1, \dots, H_k \in \mathcal{B}(\mathbb{S})$, the random variables $D(H_1), \dots, D(H_k)$ are mutually independent random variables (cf. [28]).

Fix $T \in (0, \infty)$ and let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space on which we are given a collection of i.i.d. Poisson random measures $\{D_k(ds dy dz)\}_{k \in \mathbb{N}_0}$ on $[0, T] \times [0, 1] \times \mathbb{R}_+$ with intensity given by the Lebesgue measure. Define the filtration $\{\hat{\mathcal{F}}_t\}_{0 \leq t \leq T}$ as

$$\hat{\mathcal{F}}_t \doteq \sigma\{D_k((0, s] \times H \times B), 0 \leq s \leq t, H \in \mathcal{B}([0, 1]), B \in \mathcal{B}(\mathbb{R}_+)\}$$

and let $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ be the \mathbb{P} -augmentation of this filtration. Using the above collection of PRM we now construct certain point processes with points in $[0, T] \times [0, 1]$ as follows.

Let $\bar{\mathcal{F}}$ be the $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ -predictable σ -field on $\Omega \times [0, T]$. Denote by $\bar{\mathcal{A}}_+$ the class of all $(\bar{\mathcal{F}} \otimes \mathcal{B}([0, 1]))/\mathcal{B}(\mathbb{R}_+)$ -measurable maps from $\Omega \times [0, T] \times [0, 1]$ to \mathbb{R}_+ . For $\varphi \in \bar{\mathcal{A}}_+$ and each $k \in \mathbb{N}_0$, define the counting process D_k^φ on $[0, T] \times [0, 1]$ by

$$D_k^\varphi([0, t] \times H) \doteq \int_{[0, t] \times H} \mathbf{1}_{[0, \varphi(s, y)]}(z) D_k(ds dy dz) \quad \text{for } t \in [0, T], H \in \mathcal{B}([0, 1]).$$

We regard D_k^φ as a controlled random measure, where φ is the control process that can be used to produce a desired intensity. We will write D_k^φ as D_k^θ if $\varphi = \theta$ for some constant $\theta \in \mathbb{R}_+$. In particular we will frequently take $\theta = n$. Note that D_k^θ is a PRM on $[0, T] \times [0, 1]$ with intensity $\theta ds \times dy$.

By using D_0 to represent the arrival process and D_i to represent the departure process from queues with i customers, $i \in \mathbb{N}$, we can now give the state evolution of X^n as follows:

$$(2.1) \quad \begin{aligned} X_1^n(t) &= X_1^n(0) + \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{\{X_1^n(s-) < 1\}} D_0^{n\lambda_n}(ds dy) \\ &\quad - \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{[0, X_1^n(s-) - X_2^n(s-)]}(y) D_1^n(ds dy), \end{aligned}$$

$$(2.2) \quad \begin{aligned} X_i^n(t) &= X_i^n(0) + \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{\{X_{i-1}^n(s-) = 1, X_i^n(s-) < 1\}} D_0^{n\lambda_n}(ds dy) \\ &\quad - \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{[0, X_i^n(s-) - X_{i+1}^n(s-)]}(y) D_i^n(ds dy), \quad i \geq 2. \end{aligned}$$

The first integral on the right side of (2.1) corresponds to incoming jobs that join an empty queue. The indicator in the integral captures the fact that this can happen only when an empty queue is available. The second term on the right side of (2.1) corresponds to completion of jobs by a server with only one job in the queue. The terms in equation (2.2) are interpreted in an analogous manner. By introducing *reflection terms* η_i^n defined by

$$(2.3) \quad \eta_i^n(t) \doteq \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{\{X_i^n(s-) = 1\}} D_0^{n\lambda_n}(ds dy), \quad i \geq 1$$

one can rewrite the state equation as follows. Define a *free process* Y^n as

$$(2.4) \quad \begin{aligned} Y_1^n(t) &= X_1^n(0) + \frac{1}{n} \int_{[0,t] \times [0,1]} D_0^{n\lambda_n}(ds dy) \\ &\quad - \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{[0, X_1^n(s-) - X_2^n(s-)]}(y) D_1^n(ds dy), \end{aligned}$$

$$(2.5) \quad Y_i^n(t) = X_i^n(0) - \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{[0, X_i^n(s-) - X_{i+1}^n(s-)]}(y) D_i^n(ds dy), \quad i \geq 2.$$

Then

$$(2.6) \quad \begin{aligned} X_1^n(t) &= Y_1^n(t) - \eta_1^n(t), \\ X_i^n(t) &= Y_i^n(t) + \eta_{i-1}^n(t) - \eta_i^n(t), \quad i \geq 2. \end{aligned}$$

Written in this manner, X^n can be viewed as a solution of an infinite-dimensional Skorokhod problem as discussed in the next section.

2.2. Skorokhod problem. We now introduce the Skorokhod problem that is associated with the system of equations in the last section. Consider an infinite matrix (namely a map from $\mathbb{N} \times \mathbb{N}$ to \mathbb{R}), R_∞ , defined as

$$R_\infty(i, j) = -\mathbf{1}_{\{j=i\}} + \mathbf{1}_{\{j=i-1, i>1\}} \quad \text{for } (i, j) \in \mathbb{N} \times \mathbb{N}.$$

Define $\mathbb{V} \doteq (-\infty, 1]$ and consider a $M \in \mathbb{N}$. Let \mathbb{V}^∞ (resp. \mathbb{V}^M) denote the space of maps from \mathbb{N} (resp. $\{1, \dots, M\}$) to \mathbb{V} which is equipped with the product topology. The spaces \mathbb{R}^∞ and \mathbb{R}^M are similarly defined with \mathbb{V} replaced by \mathbb{R} . Let $\mathbb{D}_{\mathbb{R}^\infty}^o$ be the subset of $\mathbb{D}_{\mathbb{R}^\infty}$ consisting of paths ψ such that $\psi(0) \in \mathbb{V}^\infty$.

DEFINITION 2.1. Let $\psi \in \mathbb{D}_{\mathbb{R}^\infty}^o$. Then $(\phi, \eta) \in \mathbb{D}_{\mathbb{V}^\infty \times \mathbb{R}^\infty}$ solves the Skorokhod problem (SP) for ψ associated with the reflection matrix R_∞ if the following hold:

(i) $\phi(t) = \psi(t) + R_\infty \eta(t)$ for all $t \in [0, T]$, namely

$$\phi_1(t) = \psi_1(t) - \eta_1(t), \quad \phi_i(t) = \psi_i(t) + \eta_{i-1}(t) - \eta_i(t) \quad \text{for all } i \geq 2 \text{ and } t \in [0, T].$$

(ii) For each $i \in \mathbb{N}$, (a) $\eta_i(0) = 0$, (b) η_i is nondecreasing, and (c) $\int_0^T \mathbf{1}_{\{\phi_i(s) < 1\}} d\eta_i(s) = 0$.

On the domain $E \subset \mathbb{D}_{\mathbb{R}^\infty}^o$ on which there is a unique solution to the SP we define the Skorokhod map (SM) $\Gamma : E \rightarrow \mathbb{D}_{\mathbb{V}^\infty}$ as $\Gamma(\psi) = \phi$ if (ϕ, η) solves the SP posed by ψ . Also, define the map $\bar{\Gamma} : E \rightarrow \mathbb{D}_{\mathbb{R}^\infty}$ by $\bar{\Gamma}(\psi) = \eta$. For $\psi, \tilde{\psi} \in \mathbb{D}_{\mathbb{R}^\infty}$, let

$$\|\psi - \tilde{\psi}\|_\infty \doteq \sum_{i=1}^{\infty} \frac{\|\psi_i - \tilde{\psi}_i\|_\infty}{2^i}.$$

The following result gives the wellposedness and regularity of the above infinite-dimensional Skorokhod problem.

LEMMA 2.2. *The SP is well defined on all of $\mathbb{D}_{\mathbb{R}^\infty}^o$ (namely $E = \mathbb{D}_{\mathbb{R}^\infty}^o$) and the SM is Lipschitz continuous in the following sense: For all $\psi, \tilde{\psi} \in \mathbb{D}_{\mathbb{R}^\infty}^o$,*

$$\|\Gamma(\psi) - \Gamma(\tilde{\psi})\|_\infty \leq 4\|\psi - \tilde{\psi}\|_\infty, \quad \|\bar{\Gamma}(\psi) - \bar{\Gamma}(\tilde{\psi})\|_\infty \leq 2\|\psi - \tilde{\psi}\|_\infty.$$

PROOF. We first prove uniqueness. Fix $\psi \in \mathbb{D}_{\mathbb{R}^\infty}^o$. Suppose there are two solutions (ϕ, η) and $(\tilde{\phi}, \tilde{\eta})$. For each $M \in \mathbb{N}$, define the matrix $R_M = -I_{M \times M} + P_M$, where $I_{M \times M}$ is the $M \times M$ identity matrix and $P_M(i, j) = \mathbf{1}_{\{j=i-1, i>1\}}$. Let $\psi^M, \phi^M, \tilde{\phi}^M, \eta^M, \tilde{\eta}^M \in \mathbb{D}_{\mathbb{R}^M}$ be defined as

$$(\psi_i^M, \phi_i^M, \tilde{\phi}_i^M, \eta_i^M, \tilde{\eta}_i^M) = (\psi_i, \phi_i, \tilde{\phi}_i, \eta_i, \tilde{\eta}_i), \quad i = 1, \dots, M.$$

Then (ϕ^M, η^M) solves the M -dimensional Skorokhod problem for ψ^M associated with the domain \mathbb{V}^M and reflection matrix R_M , namely

$$\phi^M = \psi^M + R_M \eta^M, \quad \phi^M \in \mathbb{D}_{\mathbb{V}^M},$$

$$\eta(0) = 0, \quad \eta \text{ is right continuous and nondecreasing,} \quad \int_0^T \mathbf{1}_{\{\phi_i(s) < 1\}} d\eta_i(s) = 0.$$

Furthermore, $(\tilde{\phi}^M, \tilde{\eta}^M)$ also solves the same M -dimensional Skorokhod problem for ψ^M . Since P_M has spectral radius less than 1, it is well known from [19, 25] that this M -dimensional SP has a unique solution. Thus we must have $(\phi^M, \eta^M) = (\tilde{\phi}^M, \tilde{\eta}^M)$. Since $M \in \mathbb{N}$ is arbitrary, we have the desired uniqueness property.

Next we prove existence. The solutions of the finite-dimensional SP have the following consistency property: Suppose for some $M \in \mathbb{N}$, (ϕ^M, η^M) is the solution to the M -dimensional SP for ψ^M (associated with (\mathbb{V}^M, R_M)) and let $1 \leq m \leq M$. Define $(\psi^{m,M}, \phi^{m,M}, \eta^{m,M}) \in \mathbb{D}_{\mathbb{R}^{3m}}$ as

$$(\psi_i^{m,M}, \phi_i^{m,M}, \eta_i^{m,M}) = (\psi_i^M, \phi_i^M, \eta_i^M), \quad i = 1, \dots, m.$$

Then $(\phi^{m,M}, \eta^{m,M})$ solves the m -dimensional SP for ψ^m (associated with (\mathbb{V}^m, R_m)). Fix $\psi \in \mathbb{D}_{\mathbb{R}^\infty}^o$ and for $M \in \mathbb{N}$ define ψ^M as before. Let (ϕ^M, η^M) be the solution of the M -dimensional SP for ψ^M . Define $(\phi, \eta) \in \mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$ as $(\phi_n, \eta_n) \doteq (\phi_n^n, \eta_n^n)$ for $n \in \mathbb{N}$. From the consistency property noted above, (ϕ, η) is a solution to the infinite-dimensional SP. This gives existence.

Finally we prove the Lipschitz property. Fix $\psi, \tilde{\psi} \in \mathbb{D}_{\mathbb{R}^\infty}^o$ and let (ϕ, η) and $(\tilde{\phi}, \tilde{\eta})$ be solutions to the infinite-dimensional SP for ψ and $\tilde{\psi}$ respectively. For each $M \in \mathbb{N}$, since

$$\phi_M(t) = \psi_M(t) + \eta_{M-1}(t) - \eta_M(t),$$

$$\tilde{\phi}_M(t) = \tilde{\psi}_M(t) + \tilde{\eta}_{M-1}(t) - \tilde{\eta}_M(t),$$

where we define $\eta_0 = \tilde{\eta}_0 = 0$, it follows that (ϕ_M, η_M) and $(\tilde{\phi}_M, \tilde{\eta}_M)$ are solutions to the one-dimensional SP (associated with the domain $\mathbb{V} = (-\infty, 1]$) for $\psi_M + \eta_{M-1}$ and $\tilde{\psi}_M + \tilde{\eta}_{M-1}$, and hence

$$\begin{aligned} \|\eta_M - \tilde{\eta}_M\|_\infty &\leq \|(\psi_M + \eta_{M-1}) - (\tilde{\psi}_M + \tilde{\eta}_{M-1})\|_\infty \\ &\leq \|\psi_M - \tilde{\psi}_M\|_\infty + \|\eta_{M-1} - \tilde{\eta}_{M-1}\|_\infty, \\ \|\phi_M - \tilde{\phi}_M\|_\infty &\leq 2\|(\psi_M + \eta_{M-1}) - (\tilde{\psi}_M + \tilde{\eta}_{M-1})\|_\infty \\ &\leq 2\|\psi_M - \tilde{\psi}_M\|_\infty + 2\|\eta_{M-1} - \tilde{\eta}_{M-1}\|_\infty. \end{aligned}$$

This means

$$\|\eta_M - \tilde{\eta}_M\|_\infty \leq \sum_{i=1}^M \|\psi_i - \tilde{\psi}_i\|_\infty, \quad \|\phi_M - \tilde{\phi}_M\|_\infty \leq 2 \sum_{i=1}^M \|\psi_i - \tilde{\psi}_i\|_\infty$$

and hence

$$\begin{aligned} \|\eta - \tilde{\eta}\|_\infty &= \sum_{k=1}^{\infty} \frac{\|\eta_k - \tilde{\eta}_k\|_\infty}{2^k} \leq \sum_{k=1}^{\infty} \sum_{i=1}^k \frac{\|\psi_i - \tilde{\psi}_i\|_\infty}{2^k} = \sum_{i=1}^{\infty} \sum_{k=i}^{\infty} \frac{\|\psi_i - \tilde{\psi}_i\|_\infty}{2^k} \\ &= \sum_{i=1}^{\infty} \frac{\|\psi_i - \tilde{\psi}_i\|_\infty}{2^{i-1}} = 2\|\psi - \tilde{\psi}\|_\infty. \end{aligned}$$

Similarly,

$$\|\phi - \tilde{\phi}\| = \sum_{k=1}^{\infty} \frac{\|\phi_k - \tilde{\phi}_k\|_\infty}{2^k} \leq 4\|\psi - \tilde{\psi}\|_\infty.$$

This completes the proof. \square

REMARK 2.1. It is easy to verify that if $\psi \in \mathbb{D}_{\mathbb{R}^\infty}^o$ is such that ψ_i is continuous (resp. absolutely continuous) for each i and $\zeta = \Gamma(\psi)$, $\eta = \bar{\Gamma}(\psi)$, then ζ_i and η_i are continuous (resp. absolutely continuous) for each i .

2.3. Rate function. For each $n \in \mathbb{N}$, (X^n, Y^n) is a $\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$ valued random variable. In this work we show that as $n \rightarrow \infty$, the sequence $\{(X^n, Y^n)\}_{n \in \mathbb{N}}$ satisfies a LDP in the above space. We begin by introducing the rate function that will govern this large deviation principle.

Recall that we assume $\lambda_n \rightarrow \lambda \in (0, \infty)$ as $n \rightarrow \infty$. Fix $x \in \mathbb{V}^\infty$ such that $x_i \geq 0$ for every i and $\sum_{i=1}^{\infty} x_i < \infty$. Let $\mathcal{C}_T(x)$ be the subset of $\mathbb{C}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$ consisting of all functions (ζ, ψ) such that:

- (i) $\zeta_i(0) = \psi_i(0) = x_i$. ζ_i and ψ_i are absolutely continuous on $[0, T]$.
- (ii) For all $t \in [0, T]$, $1 = \zeta_0(t) \geq \zeta_1(t) \geq \zeta_2(t) \geq \dots \geq 0$.
- (iii) $\sup_{t \in [0, T]} \sum_{i=1}^{\infty} \zeta_i(t) < \infty$.

(iv) For some $\eta \in \mathbb{C}_{\mathbb{R}^\infty}$, (ζ, η) solves the Skorokhod problem for ψ associated with the reflection matrix R_∞ . Namely,

$$(2.7) \quad \zeta_i(t) = \psi_i(t) + \eta_{i-1}(t) - \eta_i(t), \quad t \in [0, T], i \in \mathbb{N},$$

where $\eta_0(t) = 0$ and for every $i \geq 1$, $\eta_i(0) = 0$, η_i is nondecreasing, and

$$\int_0^T \mathbf{1}_{\{\zeta_i(s) < 1\}} \eta_i(ds) = 0.$$

Note that property (iii) implies there exists a smallest $M = M(\zeta) \in \mathbb{N}$ such that

$$\sup_{t \in [0, T]} \zeta_M(t) < 1.$$

Thus one only needs to consider a M -dimensional SP for ψ^M (associated with (\mathbb{V}^M, R_M)).

REMARK 2.2. Indeed, an analogous M exists for every realization of X^n . However, this M depends on the random realization and there is no uniform M that works for all realizations of X^n . As a result, it is convenient to work with the infinite-dimensional SP introduced in Section 2.2 when proving tightness and convergence properties in Section 3.

We now define the rate function. Recall $\ell(z) = z \log(z) - z + 1$ for $z \geq 0$, and let $\vartheta_0 \doteq \lambda$ and $\vartheta_i \doteq 1$ for $i \in \mathbb{N}$. For $(\zeta, \psi) \notin \mathcal{C}_T(x)$, define $I_{T,x}(\zeta, \psi) \doteq \infty$. For $(\zeta, \psi) \in \mathcal{C}_T(x)$, define

$$(2.8) \quad I_{T,x}(\zeta, \psi) \doteq \inf_{\varphi \in \mathcal{S}_T(\zeta, \psi)} \left\{ \sum_{i=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_i \ell(\varphi_i(s, y)) ds dy \right\},$$

where the set $\mathcal{S}_T(\zeta, \psi)$ consists of all $\varphi = (\varphi_i)_{i \in \mathbb{N}_0}$, where each $\varphi_i : [0, T] \times [0, 1] \rightarrow \mathbb{R}_+$, such that

$$(2.9) \quad \begin{aligned} \psi_1(t) &= x_1 + \lambda \int_{[0, t] \times [0, 1]} \varphi_0(s, y) ds dy \\ &\quad - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_1(s) - \zeta_2(s)]}(y) \varphi_1(s, y) ds dy, \end{aligned}$$

$$(2.10) \quad \psi_i(t) = x_i - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_i(s) - \zeta_{i+1}(s)]}(y) \varphi_i(s, y) ds dy, \quad i \geq 2.$$

Intuitively (2.9)–(2.10) represents the ODE limit of the free process Y^n with controls φ modulating the rates at which jobs enter and leave the system. Note that when φ_i is taken to be 1 for each i in the above equations, $\{(\zeta_i, \psi_i)\}_{i \in \mathbb{N}_0}$ correspond to the law of large numbers limit of the constrained and free processes $\{(X_i^n, Y_i^n)\}_{i \in \mathbb{N}_0}$. Clearly, with this choice of $\{\varphi_i\}_{i \in \mathbb{N}_0}$, the cost on the right side of (2.8) is zero which verifies that the rate function evaluated at the LLN limit is 0. For a general pair $\{(\zeta_i, \psi_i)\}_{i \in \mathbb{N}_0}$, the rate function is obtained by considering all controls $\{\varphi_i\}_{i \in \mathbb{N}_0}$ that produce the pair $\{(\zeta_i, \psi_i)\}_{i \in \mathbb{N}_0}$ through the system of equations in (2.9)–(2.10) and by then taking infimum over the cost for all such controls as on the right side of (2.8).

2.4. Main result. We now present the main result of this work. First we introduce the following assumption on the initial values $X^n(0)$.

ASSUMPTION 2.3. There exist a sequence of x^n and x in \mathbb{V}^∞ such that a.s.

$$\text{for every } i \in \mathbb{N}, \quad X_i^n(0) = x_i^n \rightarrow x_i \quad \text{as } n \rightarrow \infty,$$

$$1 = x_0^n \geq x_1^n \geq x_2^n \geq \dots \geq 0, \quad \sup_{n \in \mathbb{N}} \sum_{i=1}^{\infty} x_i^n < \infty.$$

REMARK 2.3. Note that Assumption 2.3 along with Fatou's lemma imply that $\sum_{i=1}^{\infty} x_i < \infty$.

Assumption 2.3 will be taken to hold throughout this work and $\{x^n\}$ and x as in Assumption 2.3 will be fixed. Thus we will not note this condition explicitly in our results and will suppress x in the notation when writing $\mathcal{C}_T(x)$ or $I_{T,x}$.

THEOREM 2.4. *The function I_T defined in (2.8) is a rate function on $\mathbb{D}_{\mathbb{R}^{\infty} \times \mathbb{R}^{\infty}}$. The sequence (X^n, Y^n) satisfies a large deviation principle on $\mathbb{D}_{\mathbb{R}^{\infty} \times \mathbb{R}^{\infty}}$ with rate function I_T .*

PROOF. From the equivalence between a LDP and a Laplace principle (cf. Section 1.2 of [17]), it suffices to establish the following three statements.

(1) Laplace upper bound: For all $G \in \mathbb{C}_b(\mathbb{D}_{\mathbb{R}^{\infty} \times \mathbb{R}^{\infty}})$,

$$(2.11) \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} e^{-nG(X^n, Y^n)} \leq - \inf_{(\zeta, \psi) \in \mathcal{C}_T} \{I_T(\zeta, \psi) + G(\zeta, \psi)\}.$$

(2) Laplace lower bound: For all $G \in \mathbb{C}_b(\mathbb{D}_{\mathbb{R}^{\infty} \times \mathbb{R}^{\infty}})$,

$$(2.12) \quad \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E} e^{-nG(X^n, Y^n)} \geq - \inf_{(\zeta, \psi) \in \mathcal{C}_T} \{I_T(\zeta, \psi) + G(\zeta, \psi)\}.$$

(3) I_T is a rate function, namely for each $M \in [0, \infty)$, $\{(\zeta, \psi) \in \mathcal{C}_T : I_T(\zeta, \psi) \leq M\}$ is compact.

Statements (1) and (2) are proved in Sections 4 and 5, respectively, while the proof of the third statement is given in Section 6. \square

The LDP given by the above theorem is useful in obtaining estimates for probabilities of various types of rare events in the JSQ system. We now consider one such example. Consider the critical (heavy traffic) regime, namely $\lambda = 1$. Suppose all queues are of length 1 at time 0 (i.e., $X_1^n(0) = x_1 = 1$ and $X_i^n(0) = x_i = 0$ for $i \geq 2$). Consider the problem of estimating the probability that a queue length will be at least $j \geq 3$ at some time instant in the time interval $[0, T]$. This corresponds to estimating the probability of (X^n, Y^n) taking values in the following (relatively) open set:

$$G_j = \left\{ (\zeta, \psi) \in \bar{\mathcal{C}}_T \mid \sup_{t \in [0, T]} \zeta_j(t) > 0 \right\},$$

where

$$\bar{\mathcal{C}}_T = \{(\zeta, \psi) \in \mathcal{C}_T \mid \zeta_1(0) = 1, \zeta_i(0) = 0 \text{ for } i \geq 2\}.$$

We also consider the closed set F_j , obtained by a slight enlargement of G_j , defined as follows:

$$F_j = \left\{ (\zeta, \psi) \in \bar{\mathcal{C}}_T \mid \sup_{t \in [0, T]} \zeta_{j-1}(t) = 1 \right\}.$$

The following result characterizes the decay rate of probabilities of G_j , F_j and shows that for large time intervals the probability of a queue buildup of length j or higher at any station decays exponentially, approximately, with rate $e^{-\frac{n(j-2)^2}{4T}}$.

THEOREM 2.5. *For every $j \geq 3$,*

$$\begin{aligned}
 (2.13) \quad & \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(X^n \in G_j) \\
 &= \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(X^n \in F_j) \\
 &= -T \ell \left(\frac{\frac{j-2}{T} + \sqrt{4 + (\frac{j-2}{T})^2}}{2} \right) - T \ell \left(\frac{-\frac{j-2}{T} + \sqrt{4 + (\frac{j-2}{T})^2}}{2} \right)
 \end{aligned}$$

and

$$\lim_{T \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{T}{n} \log (\mathbb{P}(X^n \in G_j)) = \lim_{T \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{T}{n} \log (\mathbb{P}(X^n \in F_j)) = -\frac{(j-2)^2}{4}.$$

Proof of the above theorem is given in Section 7.

REMARK 2.4. The above result shows that the probability of having a queue of length j or larger at any point in a time interval $[0, T]$ decays exponentially at rate $e^{-nI(j, T)}$, where $I(j, T)$ is given by the expression in (2.13). In addition, the proof gives information on how such an event is likely to occur. The optimal path given by (7.3) corresponds to the behavior that the number of jobs per queue increases linearly at rate $(j-2)/T$ until reaching the queue length of $(j-1)$ at time T . Such information can also be used to design accelerated Monte-Carlo sampling algorithms for approximating related probabilities for any fixed sized system by, for example, drawing samples from a proposal distribution which places more weight on sample paths close to this trajectory (see, e.g., [10], Part IV, and references therein).

CONJECTURE 2.6. *Similar techniques that are used to prove Theorem 2.5 are expected to be useful for studying decay rates of other types of events as well. In particular, we make the following two conjectures. A rigorous analysis of the probability asymptotics of events in the two conjectures is left as an open problem.*

(1) *Consider the event that at some time instant in $[0, T]$, some of the servers are busy with long queues, but the remaining are idle. Such an event signals an undesirable inefficiency or lack of balance in the system. More precisely, the event of interest is*

$$\{X_1^n(t) = X_2^n(t) = \cdots = X_j^n(t) = c, X_{j+1}^n(t) = 0 \text{ for some } t \in [0, T]\}$$

for some $j \geq 2$ and $c \in (0, 1)$. This event corresponds to $n(1-c)$ queues being idle while the remaining nc queues being of length j . We conjecture that the most likely manner in which this event occurs is that first, the system reaches a state with nc queues of length j and the remaining $n(1-c)$ queues of length $j-1$ at some time instant $t \in [0, T]$, and then those nc queues remain in that state while the other $n(1-c)$ queues decrease down to zero. This result will allow the identification of an optimal trajectory and a characterization of the decay rate of the probability of the above event.

(2) *Consider the event that at some time instant in $[0, T]$, some of the servers have queues of length $j \geq 2$ while the remaining are of length 1, namely*

$$\{X_1^n(t) = 1, X_2^n(t) = \cdots = X_j^n(t) = c, X_{j+1}^n(t) = 0 \text{ for some } t \in [0, T]\}$$

for some $j \geq 2$ and $c \in (0, 1)$. This is similar to the first event but simpler to analyze. Once more, there is a natural guess for the most likely manner in which this event occurs (which is similar to the conjecture in (1), except that the other $n(1-c)$ queues decrease down to one instead of zero) using which one can identify the optimal trajectory in the path space whose cost characterizes the decay rate of the probability. We conjecture that in this case the optimal trajectory is piecewise linear.

3. Representation and weak convergence of controlled processes. In this section we give several preparatory results that are needed for the proofs of both the upper and the lower bounds (i.e., (2.11) and (2.12)). Section 3.1 presents a variational representation from [12] that is the starting point of our analysis. In Section 3.2 we prove tightness of certain families of controls and controlled processes which arise from the variational representation of Section 3.1. Finally, Section 3.3 presents a result which characterizes the distributional limit points of this collection of processes.

3.1. Variational representation. Recall that $\bar{\mathcal{A}}_+$ denotes the class of $(\bar{\mathcal{F}} \otimes \mathcal{B}([0, 1]))/\mathcal{B}(\mathbb{R}_+)$ -measurable maps from $\Omega \times [0, T] \times [0, 1]$ to \mathbb{R}_+ . For each $m \in \mathbb{N}$ let

$$\bar{\mathcal{A}}_{b,m} \doteq \left\{ (\varphi_k)_{k \in \mathbb{N}_0} : \varphi_k \in \bar{\mathcal{A}}_+ \text{ for all } k \in \mathbb{N}_0, \text{ for all } (\omega, t, y) \in \Omega \times [0, T] \times [0, 1] \right. \\ \left. \frac{1}{m} \leq \varphi_k(\omega, t, y) \leq m \text{ for } k \leq m \text{ and } \varphi_k(\omega, t, y) = 1 \text{ for } k > m \right\}$$

and let $\bar{\mathcal{A}}_b \doteq \bigcup_{m=1}^{\infty} \bar{\mathcal{A}}_{b,m}$. For any $\varphi \in \bar{\mathcal{A}}_b$ we denote by $(\bar{X}^{n,\varphi}, \bar{Y}^{n,\varphi}, \bar{\eta}^{n,\varphi})$ the controlled analogues of (X^n, Y^n, η^n) obtained by replacing the PRMs in (2.1)–(2.5) with controlled point processes, $D_0^{n\lambda_n\varphi_0}$ and $D_i^{n\varphi_i}$, $i \in \mathbb{N}$. Namely, the state evolution equations for the controlled processes are as follows:

$$(3.1) \quad \begin{aligned} \bar{Y}_1^{n,\varphi}(t) &= x_1^n + \frac{1}{n} \int_{[0,t] \times [0,1]} D_0^{n\lambda_n\varphi_0}(ds dy) \\ &\quad - \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{[0, \bar{X}_1^n(s-) - \bar{X}_2^n(s-)]}(y) D_1^{n\varphi_1}(ds dy), \\ \bar{Y}_i^{n,\varphi}(t) &= x_i^n - \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{[0, \bar{X}_i^n(s-) - \bar{X}_{i+1}^n(s-)]}(y) D_i^{n\varphi_i}(ds dy), \quad i \geq 2, \\ \bar{X}_1^{n,\varphi}(t) &= \bar{Y}_1^{n,\varphi}(t) - \bar{\eta}_1^{n,\varphi}(t), \quad \bar{X}_i^{n,\varphi}(t) = \bar{Y}_i^{n,\varphi}(t) + \bar{\eta}_{i-1}^{n,\varphi}(t) - \bar{\eta}_i^{n,\varphi}(t), \quad i \geq 2, \\ \bar{\eta}_i^{n,\varphi}(t) &= \frac{1}{n} \int_{[0,t] \times [0,1]} \mathbf{1}_{\{\bar{X}_i^{n,\varphi}(s-) = 1\}} D_0^{n\lambda_n\varphi_0}(ds dy), \quad i \geq 1. \end{aligned}$$

When it is clear from context which controls are being used we may simply write $(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n)$ to represent the controlled processes.

Let $\vartheta_0^n \doteq \lambda_n$ and $\vartheta_i^n \doteq 1$ for $i \in \mathbb{N}$. The following variational representation will be instrumental in proving both (2.11) and (2.12). For a proof we refer the reader to [12], Theorem 2.1. We remark that the representation in [12] is given for the setting of a single PRM. However the result given in the lemma below, which is formulated in terms of a countable collection of PRM, follows immediately on considering a single PRM on the augmented space $[0, T] \times [0, 1] \times \mathbb{R}_+ \times \mathbb{N}_0$ with intensity $\text{Leb} \otimes \varrho$ where Leb is the Lebesgue measure on $[0, T] \times [0, 1] \times \mathbb{R}_+$ and ϱ is the counting measure on \mathbb{N}_0 . Proof is omitted.

LEMMA 3.1. *Let $G \in \mathbb{C}_b(\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty})$. Then*

$$(3.2) \quad \begin{aligned} & -\frac{1}{n} \log \mathbb{E} e^{-nG(X^n, Y^n)} \\ &= \inf_{\varphi^n \in \bar{\mathcal{A}}_b} \mathbb{E} \left\{ \sum_{i=0}^{\infty} \int_{[0,T] \times [0,1]} \vartheta_i^n \ell(\varphi_i^n(s, y)) ds dy + G(\bar{X}^n, \bar{Y}^n) \right\}. \end{aligned}$$

3.2. Tightness. In this section we prove a key tightness result which says that if the costs are appropriately bounded then the corresponding collection of controls and controlled processes is tight. We begin by describing the topology on the space of controls. For $M \in (0, \infty)$, denote by S_M the collection of all $h = \{h_k\}_{k \in \mathbb{N}_0}$, where $h_k : [0, T] \times [0, 1] \rightarrow \mathbb{R}_+$ for each $k \in \mathbb{N}_0$ and

$$\sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \ell(h_k(s, y)) \, ds \, dy \leq M.$$

Any h_k as above can be identified with a finite measure ν^{h_k} on $[0, T] \times [0, 1]$ by the following relation:

$$\nu^{h_k}(H) \doteq \int_H h_k(s, y) \, ds \, dy, \quad H \subset \mathcal{B}([0, T] \times [0, 1]).$$

The space \mathbb{M} of finite measures on $[0, T] \times [0, 1]$ is equipped with the weak convergence topology and the space \mathbb{M}^∞ is equipped with the corresponding product topology. Using the above identification, each element in S_M can be mapped to an element of the Polish space \mathbb{M}^∞ and the space S_M with the inherited topology is compact (see [9], Lemma A.1).

We record the following elementary lemma for future use. See [5], Lemma 3.2, for a proof.

LEMMA 3.2. *Let $\ell(x) = x \log(x) - x + 1$. Then the following properties hold for $\ell(x)$:*

- (a) *For each $\beta > 0$, there exists $\gamma(\beta) \in (0, \infty)$ such that $\gamma(\beta) \rightarrow 0$ as $\beta \rightarrow \infty$ and $x \leq \gamma(\beta)\ell(x)$, for $x \geq \beta \geq 1$.*
- (b) *For $x \geq 0$, $x \leq \ell(x) + 2$.*

The following is the main tightness result of this section. As a convention, we will take $\bar{X}_0^n(t) = \bar{Y}_0^n(t) = 1$ for all $t \in [0, T]$.

LEMMA 3.3. *Suppose that $\{\varphi^n\}$ is a sequence in $\bar{\mathcal{A}}_b$ such that for some $M_0 \in (0, \infty)$*

$$(3.3) \quad \sup_{n \in \mathbb{N}} \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \ell(\varphi_k^n(s, y)) \, ds \, dy \leq M_0 \quad a.s.$$

Denote by $(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n)$ the controlled processes associated with φ^n , given by (3.1) (replacing φ with φ^n). Then, regarding φ^n as a S_{M_0} valued random variable, the sequence $\{(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n)\}_{n \in \mathbb{N}_0}$ is tight in $\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty \times \mathbb{R}^\infty} \times S_{M_0}$. Furthermore the collection $\{(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n)\}_{n \in \mathbb{N}_0}$ is \mathbb{C} -tight.

PROOF. Since S_{M_0} is compact the tightness of $\{\varphi^n\}_{n \in \mathbb{N}_0}$ is immediate. From Lemma 2.2 and since $(\bar{X}^n, \bar{\eta}^n) = (\Gamma(\bar{Y}^n), \bar{\Gamma}(\bar{Y}^n))$, it suffices now to show that $\{\bar{Y}^n\}_{n \in \mathbb{N}}$ is \mathbb{C} -tight in $\mathbb{D}_{\mathbb{R}^\infty}$.

In order to verify tightness of $\{\bar{Y}^n\}_{n \in \mathbb{N}}$ we appeal to Aldous' tightness criteria (cf. Theorem 2.2.2 in [29]) for each $\{\bar{Y}_i^n\}_{n \in \mathbb{N}}, i \in \mathbb{N}$. The tightness of $\{\bar{Y}_i^n(t)\}_{n \in \mathbb{N}}$ in \mathbb{R} for each $t \in [0, T]$ follows from the following estimate:

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbb{E}|\bar{Y}_i^n(t)| &\leq \limsup_{n \rightarrow \infty} x_i^n + \limsup_{n \rightarrow \infty} \mathbb{E} \int_{[0, t] \times [0, 1]} [\lambda_n \varphi_0^n(s, y) + \varphi_i^n(s, y)] \, ds \, dy \\ &\leq x_i + (\lambda + 1)M_0 + 2(\lambda + 1)T, \end{aligned}$$

where the last inequality uses (3.3), Lemma 3.2(b), and Assumption 2.3. We next verify the following condition on the fluctuations of $\{\bar{Y}^n\}_{n \in \mathbb{N}}$:

$$(3.4) \quad \limsup_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \sup_{\tau \in \mathcal{T}^\delta} \mathbb{E}[|\bar{Y}_i^n(\tau + \delta) - \bar{Y}_i^n(\tau)|] = 0,$$

where \mathcal{T}^δ is the set of all $[0, T - \delta]$ -valued stopping times. For any $L > 1$, it follows from the definition of \bar{Y}^n , the triangle inequality, Lemma 3.2(a), and (3.3) that

$$\begin{aligned}
& \mathbb{E}[|\bar{Y}_i^n(\tau + \delta) - \bar{Y}_i^n(\tau)|] \\
& \leq \lambda_n \mathbb{E} \int_{[\tau, \tau + \delta] \times [0, 1]} \varphi_0^n(s, y) ds dy + \mathbb{E} \int_{[\tau, \tau + \delta] \times [0, 1]} \varphi_i^n(s, y) ds dy \\
& = \lambda_n \mathbb{E} \int_{[\tau, \tau + \delta] \times [0, 1]} \varphi_0^n(s, y) \mathbf{1}_{\{\varphi_0^n(s, y) \leq L\}} ds dy \\
& \quad + \mathbb{E} \int_{[\tau, \tau + \delta] \times [0, 1]} \varphi_i^n(s, y) \mathbf{1}_{\{\varphi_i^n(s, y) \leq L\}} ds dy \\
& \quad + \lambda_n \mathbb{E} \int_{[\tau, \tau + \delta] \times [0, 1]} \varphi_0^n(s, y) \mathbf{1}_{\{\varphi_0^n(s, y) > L\}} ds dy \\
& \quad + \mathbb{E} \int_{[\tau, \tau + \delta] \times [0, 1]} \varphi_i^n(s, y) \mathbf{1}_{\{\varphi_i^n(s, y) > L\}} ds dy \\
& \leq (\lambda_n + 1) \delta L + (\lambda_n + 1) M_0 \gamma(L)
\end{aligned}$$

and thus

$$\limsup_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \sup_{\tau \in \mathcal{T}^\delta} \mathbb{E}[|\bar{Y}_i^n(\tau + \delta) - \bar{Y}_i^n(\tau)|] \leq (\lambda + 1) M_0 \gamma(L).$$

The property in (3.4) now follows upon sending $L \rightarrow \infty$. This proves the tightness of $\{Y^n\}_{n \in \mathbb{N}}$. Finally, \mathbb{C} -tightness follows upon noting that jump sizes of Y_i^n are bounded by $\frac{2}{n}$. \square

3.3. Characterization of limit points. Suppose that $\{\varphi^n\}$ is a sequence as in Lemma 3.3. Then from the lemma we have the tightness of the vector sequence $\{(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n)\}_{n \in \mathbb{N}_0}$. In this section we characterize the limit points of this sequence. It will be convenient to consider the following compensated point processes:

$$(3.5) \quad \bar{D}_k^{n\vartheta_k^n\varphi_k^n}(ds dy) \doteq D_k^{n\vartheta_k^n\varphi_k^n}(ds dy) - n\vartheta_k^n\varphi_k^n(s, y) ds dy, \quad n \in \mathbb{N}, k \in \mathbb{N}_0.$$

Define compensated processes \tilde{B}^n as

$$\begin{aligned}
(3.6) \quad \tilde{B}_1^n(t) & \doteq \frac{1}{n} \int_{[0, t] \times [0, 1]} \bar{D}_0^{n\lambda_n\varphi_0^n}(ds dy) \\
& \quad - \frac{1}{n} \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \bar{X}_1^n(s-) - \bar{X}_2^n(s-)]}(y) \bar{D}_1^{n\varphi_1^n}(ds dy),
\end{aligned}$$

$$(3.7) \quad \tilde{B}_i^n(t) \doteq \frac{1}{n} \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \bar{X}_i^n(s-) - \bar{X}_{i+1}^n(s-)]}(y) \bar{D}_i^{n\varphi_i^n}(ds dy), \quad i \geq 2$$

which allows us to write

$$\begin{aligned}
(3.8) \quad \bar{Y}_1^n(t) & = X_1^n(0) + \tilde{B}_1^n(t) + \lambda_n \int_{[0, t] \times [0, 1]} \varphi_0^n(s, y) ds dy \\
& \quad - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \bar{X}_1^n(s) - \bar{X}_2^n(s)]}(y) \varphi_1^n(s, y) ds dy,
\end{aligned}$$

$$(3.9) \quad \bar{Y}_i^n(t) = X_i^n(0) - \tilde{B}_i^n(t) - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \bar{X}_i^n(s) - \bar{X}_{i+1}^n(s)]}(y) \varphi_i^n(s, y) ds dy, \quad i \geq 2.$$

The following lemma characterizes the limit points of $\{(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n)\}_{n \in \mathbb{N}_0}$.

LEMMA 3.4. Suppose that $\{\varphi^n\}$ is a sequence as in Lemma 3.3. Suppose also that the associated sequence $\{(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n)\}_{n \in \mathbb{N}_0}$ converges along a subsequence, in distribution, to $(\bar{X}, \bar{Y}, \bar{\eta}, \varphi)$ given on some probability space $(\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$. Then the following hold \mathbb{P}^* -a.s.

(a) Equations (2.9)–(2.10) are satisfied with (ζ, ψ, φ) replaced by $(\bar{X}, \bar{Y}, \bar{\varphi})$.

(b) $(\bar{X}, \bar{Y}) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\bar{X}, \bar{Y})$. In particular, $(\bar{X}, \bar{Y}, \bar{\eta})$ satisfy the following system of equations:

$$(3.10) \quad \bar{X}_1(t) = \bar{Y}_1(t) - \bar{\eta}_1(t),$$

$$(3.11) \quad \bar{X}_i(t) = \bar{Y}_i(t) + \bar{\eta}_{i-1}(t) - \bar{\eta}_i(t), \quad i \geq 2,$$

and for every $i \in \mathbb{N}$, $\bar{\eta}_i(0) = 0$, $\bar{\eta}_i$ is nondecreasing, and $\int_0^t \mathbf{1}_{\{\bar{X}_i(s) < 1\}} \bar{\eta}_i(ds) = 0$.

PROOF. Assume without loss of generality that convergence occurs along the whole sequence. Recall the expression for \bar{Y}^n given in (3.8)–(3.9) and the definition of \tilde{B}^n given in (3.6)–(3.7). It follows from Doob's inequality and Lemma 3.2(b) that for each $k \in \mathbb{N}_0$

$$(3.12) \quad \begin{aligned} \mathbb{E} \left(\sup_{0 \leq t \leq T} |\tilde{B}_k^n(t)|^2 \right) &\leq \frac{1}{n} \mathbb{E} \int_{[0, T] \times [0, 1]} [\lambda_n \varphi_0^n(s, y) + \varphi_k^n(s, y)] ds dy \\ &\leq \frac{1}{n} \mathbb{E} \int_{[0, T] \times [0, 1]} [\lambda_n (\ell(\varphi_0^n(s, y)) + 2) + \ell(\varphi_k^n(s, y)) + 2] ds dy \\ &\leq \frac{1}{n} (\lambda_n + 1) (M_0 + 2T) \rightarrow 0. \end{aligned}$$

By appealing to the Skorokhod representation theorem (cf. [6], Theorem 6.7), we can assume without loss of generality that $(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n, \tilde{B}^n) \rightarrow (\bar{X}, \bar{Y}, \bar{\eta}, \varphi, 0)$ a.s. on $(\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$, and thus the rest of the argument will be made a.s. on $(\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$. From the \mathbb{C} -tightness proved in Lemma 3.3 $(\bar{X}, \bar{Y}, \bar{\eta})$ takes values in $\mathbb{C}_{\mathbb{R}^\infty \times \mathbb{R}^\infty \times \mathbb{R}^\infty}$. Using the triangle inequality

$$(3.13) \quad \begin{aligned} &\left| \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \bar{X}_i^n(s) - \bar{X}_{i+1}^n(s)]} (y) \varphi_i^n(s, y) ds dy \right. \\ &\quad \left. - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \bar{X}_i(s) - \bar{X}_{i+1}(s)]} (y) \varphi_i(s, y) ds dy \right| \\ &\leq \int_{[0, t] \times [0, 1]} |\mathbf{1}_{[0, \bar{X}_i^n(s) - \bar{X}_{i+1}^n(s)]} (y) - \mathbf{1}_{[0, \bar{X}_i(s) - \bar{X}_{i+1}(s)]} (y)| \varphi_i^n(s, y) ds dy \\ &\quad + \left| \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \bar{X}_i(s) - \bar{X}_{i+1}(s)]} (y) (\varphi_i^n(s, y) - \varphi_i(s, y)) ds dy \right| \end{aligned}$$

for each $i \in \mathbb{N}$. Since $\text{Leb}_t\{(s, y) : y = \bar{X}_i(s) - \bar{X}_{i+1}(s)\} = 0$, where Leb_t is the Lebesgue measure on $[0, t] \times [0, 1]$, we have

$$|\mathbf{1}_{[0, \bar{X}_i^n(s) - \bar{X}_{i+1}^n(s)]} (y) - \mathbf{1}_{[0, \bar{X}_i(s) - \bar{X}_{i+1}(s)]} (y)| \rightarrow 0$$

as $n \rightarrow \infty$ for Leb_t -a.e. $(s, y) \in [0, t] \times [0, 1]$. From (3.3) and the super-linearity of ℓ , one has the uniform integrability of $(s, y) \mapsto \varphi_i^n(s, y)$ with respect to the normalized Lebesgue measure on $[0, T] \times [0, 1]$. The above two observations imply that, as $n \rightarrow \infty$,

$$(3.14) \quad \int_{[0, t] \times [0, 1]} |\mathbf{1}_{[0, \bar{X}_i^n(s) - \bar{X}_{i+1}^n(s)]} (y) - \mathbf{1}_{[0, \bar{X}_i(s) - \bar{X}_{i+1}(s)]} (y)| \varphi_i^n(s, y) ds dy \rightarrow 0.$$

Recalling the topology on S_{M_0} , the convergence $\varphi^n \rightarrow \varphi$ and $\lambda_n \rightarrow \lambda$ implies that

$$(3.15) \quad \left| \int_{[0,t] \times [0,1]} \mathbf{1}_{[0, \bar{X}_i(s) - \bar{X}_{i+1}(s))}(y) (\varphi_i^n(s, y) - \varphi_i(s, y)) ds dy \right| \rightarrow 0,$$

$$(3.16) \quad \lambda_n \int_{[0,t] \times [0,1]} \varphi_0^n(s, y) ds dy \rightarrow \lambda \int_{[0,t] \times [0,1]} \varphi_0(s, y) ds dy.$$

Combining (3.8)–(3.9) with (3.12) and (3.13)–(3.16) completes the proof of (a).

We now prove part (b). The validity of (3.10)–(3.11) is immediate from the fact that these equalities hold with $(\bar{X}, \bar{Y}, \bar{\eta})$ replaced with $(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n)$. We now check that $(\bar{X}, \bar{Y}) \in \mathcal{C}_T$ by verify properties (i)–(iv) in the definition of \mathcal{C}_T . The absolute continuity property in property (i) follows from the absolute continuity of \bar{Y}_i , which follows from part (a) of the lemma, and property (iv) (whose proof is given below), together with properties of the Skorokhod map. The remaining statements in properties (i)–(ii) of the definition of \mathcal{C}_T are immediate from Assumption 2.3 and the fact that for every t , $1 = \bar{X}_0^n(t) \geq \bar{X}_1^n(t) \geq \dots$. Property (iii) follows on noting that

$$\begin{aligned} \mathbb{E} \left[\sup_{0 \leq t \leq T} \sum_{i=1}^{\infty} \bar{X}_i(t) \right] &\leq \liminf_{n \rightarrow \infty} \mathbb{E} \left[\sup_{0 \leq t \leq T} \sum_{i=1}^{\infty} \bar{X}_i^n(t) \right] \\ &\leq \liminf_{n \rightarrow \infty} \sum_{i=1}^{\infty} x_i^n + \liminf_{n \rightarrow \infty} \int_{[0,T] \times [0,1]} \lambda_n \varphi_0^n(s, y) ds dy < \infty, \end{aligned}$$

where the first inequality follows from Fatou's lemma, and the last inequality uses Assumption 2.3 and (3.3). Finally we verify part (iv). Fix $k \in \mathbb{N}$. That $\bar{\eta}_k(\cdot)$ is nondecreasing follows from the fact that the property holds for all $\bar{\eta}_k^n(\cdot)$. We now verify that

$$(3.17) \quad \int_0^T (1 - \bar{X}_k(s)) \bar{\eta}_k(ds) = 0.$$

Note that $\bar{\eta}_k^n(t)$ can only increase when $\bar{X}_k^n(t-) = 1$ and thus

$$\int_0^T (1 - \bar{X}_k^n(t-)) \bar{\eta}_k^n(ds) = 0.$$

From this we have

$$\begin{aligned} (3.18) \quad \left| \int_0^T (1 - \bar{X}_k(s)) \bar{\eta}_k(ds) \right| &= \left| \int_0^T (1 - \bar{X}_k(s)) \bar{\eta}_k(ds) - \int_0^T (1 - \bar{X}_k^n(s-)) \bar{\eta}_k^n(ds) \right| \\ &\leq \left| \int_0^T (1 - \bar{X}_k(s)) \bar{\eta}_k(ds) - \int_0^T (1 - \bar{X}_k(s)) \bar{\eta}_k^n(ds) \right| \\ &\quad + \int_0^T |\bar{X}_k^n(t-) - \bar{X}_k(t)| \bar{\eta}_k^n(dt). \end{aligned}$$

It follows from the fact that $\bar{\eta}_k^n \rightarrow \bar{\eta}_k$ as a finite measure on $[0, T]$ and $\bar{X}_k(s) \in \mathbb{C}_b([0, T] : \mathbb{R})$ that

$$(3.19) \quad \left| \int_0^T (1 - \bar{X}_k(s)) \bar{\eta}_k(ds) - \int_0^T (1 - \bar{X}_k(s)) \bar{\eta}_k^n(ds) \right| \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Finally, continuity of \bar{X}_k and uniform convergence of \bar{X}_k^n to \bar{X}_k gives

$$\int_0^T |\bar{X}_k^n(t-) - \bar{X}_k(t)| \bar{\eta}_k^n(dt) \leq \left(\sup_{0 \leq t \leq T} |\bar{X}_k^n(t-) - \bar{X}_k(t)| \right) \bar{\eta}_k^n(T) \rightarrow 0,$$

which, combined with (3.18) and (3.19), gives (3.17) verifying property (iv). Thus we have shown that $(\bar{X}, \bar{Y}) \in \mathcal{C}_T$ a.s. The fact that $\varphi \in \mathcal{S}_T(\bar{X}, \bar{Y})$ is now immediate from part (a) of the lemma. \square

4. Laplace upper bound. This section is devoted to the proof of the Laplace upper bound (2.11). Fix $G \in \mathbb{C}_b(\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty})$. From the variational representation in Lemma 3.1, for all $n \in \mathbb{N}$, we can select a control $\tilde{\varphi}^n \in \tilde{\mathcal{A}}_b$ such that

$$(4.1) \quad \begin{aligned} & -\frac{1}{n} \log \mathbb{E} e^{-nG(X^n, Y^n)} \\ & \geq \mathbb{E} \left\{ \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k^n \ell(\tilde{\varphi}_k^n(s, y)) ds dy + G(\bar{X}^{n, \tilde{\varphi}^n}, \bar{Y}^{n, \tilde{\varphi}^n}) \right\} - \frac{1}{n}. \end{aligned}$$

This shows that

$$\sup_{n \in \mathbb{N}} \mathbb{E} \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k^n \ell(\tilde{\varphi}_k^n(s, y)) ds dy \leq 2\|G\|_\infty + 1 \doteq M_G.$$

By a standard localization argument (see, e.g., [12], Proof of Theorem 4.2) it now follows that for any fixed $\sigma > 0$ there is a $M_0 \in (0, \infty)$ and a sequence $\varphi^n \in \tilde{\mathcal{A}}_b$ taking values in S_{M_0} a.s. such that, for every n , the expected value on the right side of (4.1) differs from the same expected value, but with $\tilde{\varphi}^n$ replaced by φ^n throughout, by at most σ . In particular,

$$(4.2) \quad \begin{aligned} & -\frac{1}{n} \log \mathbb{E} e^{-nG(X^n, Y^n)} \\ & \geq \mathbb{E} \left\{ \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k^n \ell(\varphi_k^n(s, y)) ds dy + G(\bar{X}^{n, \varphi^n}, \bar{Y}^{n, \varphi^n}) \right\} - \frac{1}{n} - \sigma. \end{aligned}$$

Now we can complete the proof of the Laplace upper bound. Since φ^n are in S_{M_0} a.s., from Lemma 3.3 we have the tightness of $(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n)$. Assume without loss of generality that $(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n)$ converges along the whole sequence, in distribution, to $(\bar{X}, \bar{Y}, \bar{\eta}, \varphi)$, given on some probability space $(\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$. By Lemma 3.4 we have $(\bar{X}, \bar{Y}) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\bar{X}, \bar{Y})$ a.s. \mathbb{P}^* . Using (4.2), Fatou's lemma, and the definition of I_T in (2.8)

$$\begin{aligned} & \liminf_{n \rightarrow \infty} -\frac{1}{n} \log \mathbb{E} e^{-nG(X^n, Y^n)} \\ & \geq \liminf_{n \rightarrow \infty} \mathbb{E} \left\{ \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k^n \ell(\varphi_k^n(s, y)) ds dy + G(\bar{X}^n, \bar{Y}^n) - \frac{1}{n} - \sigma \right\} \\ & \geq \mathbb{E}^* \left\{ \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy + G(\bar{X}, \bar{Y}) \right\} - \sigma \\ & \geq \inf_{(\zeta, \psi) \in \mathcal{C}_T} \{I_T(\zeta, \psi) + G(\zeta, \psi)\} - \sigma, \end{aligned}$$

where the second inequality is a consequence of a lower semicontinuity property established in Lemma A.1 in [9]. Since $\sigma \in (0, 1)$ is arbitrary, this completes the proof of the Laplace upper bound.

5. Laplace lower bound. This section is devoted to the proof of the Laplace lower bound (2.12). The proof, given in Section 5.1, proceeds by first selecting a trajectory (ζ^*, ψ^*) which is near-optimal in the infimum on the right side of the lower bound and then an associated control φ^* which is near-optimal for the rate function evaluated at (ζ^*, ψ^*) . One then leverages the results of Section 3 to construct a sequence of controls and associated controlled process $(\bar{X}^n, \bar{Y}^n, \varphi^n)$ which are tight and have weak limit points $(\bar{\psi}, \bar{\zeta}, \bar{\varphi})$ that satisfy $\bar{\varphi} = \varphi^* \in \mathcal{S}_T(\bar{\zeta}, \bar{\psi})$. However, it is not guaranteed that $(\zeta^*, \psi^*) = (\bar{\zeta}, \bar{\psi})$ as it is not necessarily true that the control φ^* drives a unique set of solutions to the controlled ODE, (2.9)–(2.10).

In particular, if $\varphi^* \in \mathcal{S}_T(\zeta^*, \psi^*)$ and $\varphi^* \in \mathcal{S}_T(\bar{\zeta}, \bar{\psi})$ then, in general, (ζ^*, ψ^*) need not equal $(\bar{\zeta}, \bar{\psi})$. This obstacle is overcome in Lemma 5.1, which says that one can select a trajectory and control (ζ, ψ, φ) suitability close to $(\zeta^*, \psi^*, \varphi^*)$ for which this uniqueness property does hold. This uniqueness result is at the technical heart of this work and its proof relies on several other intermediate results which proceed by successively modifying the original trajectories and controls so that the modified trajectories remain close to the original trajectories while changing the cost only slightly, and such that the final set of trajectories and controls do have the desired uniqueness property. To do this, first Lemmas 5.3 and 5.4 form an inductive argument showing that one can approximate any trajectory with one in which the shortest queue length π changes only a finite number of times. This relies on taking any candidate trajectory $(\bar{\zeta}, \bar{\psi})$ and “smoothing” out small excursions of the trajectories over which π may change an infinite number of times. Then, in Lemma 5.2, we further modify the trajectory so that it is well-behaved in the neighborhood of the finitely many time instants at which π changes values. With these properties in hand, we then proceed to the proof of the uniqueness result, Lemma 5.1. The main step in the proof follows from an argument by contradiction which hinges on introducing “ ε -gaps” in the spatial thinning used to define the controls for the PRM, see (5.6). These gaps alter each control $\varphi_i(t, y)$ so that they are zero in the ε -neighborhood around $\zeta_k(t) - \zeta_{k-1}(t)$ without altering the resulting state trajectory. The basic idea is that if there were two different solutions to the ODE then there would be some time τ when they diverged. In the time right after τ , the two trajectories must be extremely close which due to the ε -gap property will say that the time derivative of the ODE must remain the same in a small interval beyond τ , triggering the contradiction. A similar “ ε -gap” argument in a very different context has recently also been used in [5]. The technical details of this argument rely heavily on the regularity properties established in Lemmas 5.2–5.4. In order to make clear the motivation behind the various lemmas in this section and improve the overall readability we present these lemmas out of order. An outline of how these various lemmas are used and the organization of this section is as follows:

- Proof of the lower bound (2.12) is a consequence of the main uniqueness result, Lemma 5.1. The statement of Lemma 5.1 and proof of (2.12) based on this lemma are given in Section 5.1.
- Lemma 5.1 relies on Lemma 5.2 which asserts existence of approximating trajectories for which π changes at finitely many time instants and that have suitable regularity in the neighborhood of the boundaries of the finitely many pieces of the approximating trajectories over which π is constant. The statement of Lemma 5.2 and proof of Lemma 5.1 based on this lemma are given in Section 5.2.
- Section 5.3 is the longest subsection and it is devoted to the proof of Lemma 5.2. This lemma is a consequence of Lemmas 5.3 and 5.4 that construct approximating trajectories in which the shortest queue π changes a finite number of times. These lemmas (statements and proofs) are given in Section 5.3 as well.

5.1. *Proof of lower bound.* The following lemma is key to the proof of the lower bound (2.12). It says that, given a trajectory $(\zeta^*, \psi^*) \in \mathcal{C}_T$, one can select a trajectory (ζ, ψ) which is suitably close to (ζ^*, ψ^*) and a control φ such that (ζ, ψ) is the unique trajectory driven by φ . We will equip the space $\mathbb{C}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$ with the distance: for $(\zeta, \psi), (\tilde{\zeta}, \tilde{\psi}) \in \mathbb{C}_{\mathbb{R}^\infty \times \mathbb{R}^\infty}$,

$$\begin{aligned} d((\zeta, \psi), (\tilde{\zeta}, \tilde{\psi})) &\doteq \|(\zeta, \psi) - (\tilde{\zeta}, \tilde{\psi})\|_\infty \doteq \|\zeta - \tilde{\zeta}\|_\infty + \|\psi - \tilde{\psi}\|_\infty \\ &= \sum_{i=1}^{\infty} \frac{\|\zeta_i - \tilde{\zeta}_i\|_\infty \wedge 1}{2^i} + \sum_{i=1}^{\infty} \frac{\|\psi_i - \tilde{\psi}_i\|_\infty \wedge 1}{2^i}, \end{aligned}$$

where recall that for $f \in \mathbb{C}([0, T] : \mathbb{R})$, $\|f\|_\infty = \sup_{0 \leq t \leq T} |f(t)|$.

LEMMA 5.1. Fix $\sigma \in (0, 1)$. Given $(\zeta^*, \psi^*) \in \mathcal{C}_T$ with $I_T(\zeta^*, \psi^*) < \infty$, there exists $(\zeta, \psi) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\zeta, \psi)$ such that:

- (a) $\|(\zeta, \psi) - (\zeta^*, \psi^*)\|_\infty \leq \sigma$.
- (b) $\sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy \leq I_T(\zeta, \psi) + \sigma \leq I_T(\zeta^*, \psi^*) + 2\sigma$.
- (c) If $(\tilde{\zeta}, \tilde{\psi})$ is another pair in \mathcal{C}_T such that $\varphi \in \mathcal{S}_T(\tilde{\zeta}, \tilde{\psi})$, then $(\tilde{\zeta}, \tilde{\psi}) = (\zeta, \psi)$.

Proof of this lemma will be given in Section 5.2. We now complete the proof of the lower bound using this result. Fix $G \in \mathbb{C}_b(\mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty})$ and $\sigma \in (0, 1)$. Select a trajectory (ζ^*, ψ^*) which is σ -optimal for the RHS of (2.12), namely

$$(5.1) \quad I_T(\zeta^*, \psi^*) + G(\zeta^*, \psi^*) \leq \inf_{(\zeta, \psi) \in \mathcal{C}_T} \{I_T(\zeta, \psi) + G(\zeta, \psi)\} + \sigma.$$

By continuity of G and Lemma 5.1, we can find $(\bar{\zeta}, \bar{\psi}) \in \mathcal{C}_T$ and $\bar{\varphi} \in \mathcal{S}_T(\bar{\zeta}, \bar{\psi})$ such that the uniqueness property in Lemma 5.1 holds (with φ replaced by $\bar{\varphi}$) and

$$(5.2) \quad \begin{aligned} \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\bar{\varphi}_k(s, y)) ds dy + G(\bar{\zeta}, \bar{\psi}) &\leq I_T(\bar{\zeta}, \bar{\psi}) + G(\bar{\zeta}, \bar{\psi}) + \sigma \\ &\leq I_T(\zeta^*, \psi^*) + G(\zeta^*, \psi^*) + 2\sigma. \end{aligned}$$

Define the sequence of controls $\varphi^n \in \bar{\mathcal{A}}_b$ as

$$\begin{aligned} \varphi_i^n(s, y) &\doteq \frac{1}{n} \mathbf{1}_{\{\bar{\varphi}_i(s, y) \leq \frac{1}{n}\}} + \bar{\varphi}_i(s, y) \mathbf{1}_{\{\frac{1}{n} < \bar{\varphi}_i(s, y) < n\}} + n \mathbf{1}_{\{\bar{\varphi}_i(s, y) \geq n\}}, \quad i \leq n, \\ \varphi_i^n(s, y) &\doteq 1, \quad i > n. \end{aligned}$$

Then there is a $M_0 \in (0, \infty)$ such that the sequence $\{\varphi^n\}$ satisfies (3.3). Furthermore, it is easily checked that $\varphi^n \rightarrow \bar{\varphi}$ (in S_{M_0}). It then follows from Lemmas 3.3 and 3.4 that $(\bar{X}^n, \bar{Y}^n, \bar{\eta}^n, \varphi^n)$ is tight and any limit point $(\bar{X}, \bar{Y}, \bar{\eta}, \varphi)$, given on some probability space $(\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$, satisfies $(\bar{X}, \bar{Y}) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\bar{X}, \bar{Y})$ a.s. From the fact that $\varphi^n \rightarrow \bar{\varphi}$ we must have $\varphi = \bar{\varphi}$. Thus $\bar{\varphi} \in \mathcal{S}_T(\bar{X}, \bar{Y})$ and since we also have $\bar{\varphi} \in \mathcal{S}_T(\bar{\zeta}, \bar{\psi})$, we must have $(\bar{X}, \bar{Y}) = (\bar{\zeta}, \bar{\psi})$ a.s. \mathbb{P}^* from the uniqueness property noted above. Noting that $\ell(\varphi_k^n(s, y)) \leq \ell(\bar{\varphi}_k(s, y))$ for all $n \in \mathbb{N}_0$ and $(s, y) \in [0, T] \times [0, 1]$, it then follows from the variational representation (3.2) and (5.1)–(5.2) that

$$\begin{aligned} &\limsup_{n \rightarrow \infty} -\frac{1}{n} \log \mathbb{E} e^{-nG(X^n, Y^n)} \\ &\leq \limsup_{n \rightarrow \infty} \mathbb{E} \left\{ \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k^n \ell(\varphi_k^n(s, y)) ds dy + G(\bar{X}^n, \bar{Y}^n) \right\} \\ &\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\bar{\varphi}_k(s, y)) ds dy + G(\bar{\zeta}, \bar{\psi}) \\ &\leq \inf_{(\zeta, \psi) \in \mathcal{C}_T} \{I_T(\zeta, \psi) + G(\zeta, \psi)\} + 3\sigma. \end{aligned}$$

The inequality in (2.12) now follows upon sending $\sigma \rightarrow 0$.

5.2. *Proof of Lemma 5.1.* In this section we prove Lemma 5.1 using an intermediate result, Lemma 5.2. Proof of Lemma 5.2 is given in Section 5.3. Consider $(\zeta, \psi) \in \mathcal{C}_T$. Informally, we will view $\zeta_k(t)$ as the (asymptotic analogue of) fraction of queues with k or more

jobs. Let $\pi(t) \doteq \max\{k : \zeta_k(t) = 1\}$ represent the shortest queue length at time t . It is easily verified that since ζ is continuous,

$$(5.3) \quad \pi(t) \text{ is upper semicontinuous.}$$

The following lemma shows that any $(\zeta, \psi) \in \mathcal{C}_T$ can be suitably approximated by a trajectory for which the associated $\pi(t)$ changes only a finite number of times, in an appropriately regular manner.

LEMMA 5.2. *Fix $\sigma \in (0, 1)$. Given $(\tilde{\zeta}, \tilde{\psi}) \in \mathcal{C}_T$ with $I_T(\tilde{\zeta}, \tilde{\psi}) < \infty$, there exists $(\zeta, \psi) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\zeta, \psi)$ such that:*

- (i) *There exist $N \in \mathbb{N}$, $c_i \in \mathbb{N}_0$ for all $i = 0, 1, \dots, N-1$, and a finite partition $0 = t_0 < t_1 < \dots < t_N = T$, such that $\pi(t) = c_i$ for $t \in (t_i, t_{i+1})$.*
- (ii) *For any i , if $\pi(t_i) > c_i$, there exists some $\Delta t_i \in (0, t_{i+1} - t_i)$ such that $\varphi_0(t, y) = 0$ for $t \in (t_i, t_i + \Delta t_i)$.*
- (iii) $\|(\zeta, \psi) - (\tilde{\zeta}, \tilde{\psi})\|_\infty \leq \sigma$.
- (iv) $I_T(\zeta, \psi) \leq I_T(\tilde{\zeta}, \tilde{\psi}) + \sigma$.
- (v) $\sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy \leq I_T(\zeta, \psi) + \sigma$.
- (vi) *There exists $M \in \mathbb{N}$ such that $\varphi_k = 1$ for $k \geq M$.*

The proof of this lemma is deferred to Section 5.3. We now have all of the ingredients needed to prove Lemma 5.1. To simplify the notation, let for $(\zeta, \psi) \in \mathcal{C}_T$ and $t \in [0, T]$,

$$(5.4) \quad r_k(t) \doteq \zeta_k(t) - \zeta_{k+1}(t).$$

Informally, this represents the fraction of queues with length k at time instant t .

PROOF OF LEMMA 5.1. We first give the basic idea of the proof. We will refer to Lemma 5.2(i)–(vi) simply as properties (i)–(vi), since they will be frequently used. Fix $(\zeta^*, \psi^*) \in \mathcal{C}_T$ with $I_T(\zeta^*, \psi^*) < \infty$. We will begin by approximating this trajectory by a more regular trajectory (ζ, ψ) of the form given in Lemma 5.2. Next, we will make additional modifications to the associated control so that one has the desired uniqueness property in part (c). Ultimately the uniqueness will be argued through a proof by contradiction. The purpose of these further modifications to the controls is to ensure that, should there be two possible trajectories driven by the modified control, then the time derivative of the state process will be the same in a small time interval following the moment at which these trajectories diverge, triggering the contradiction. This is accomplished by introducing “ ε -gaps” in the spatial thinning used to define the control process. The basic idea is to set the controls equal to zero in ε -neighborhoods of each $r_k(t)$ and, in some cases, at the boundary between regions, so that a small divergence from r_k does not cause a change in the time derivative of the state process. The controls are also reweighted so that they do not change the state trajectory (ζ, ψ) . We then show that this modified control increases the cost only slightly and has all the desired properties in Lemma 5.1.

Approximating (ζ, ψ) by a more regular trajectory. Fix $\sigma \in (0, 1)$. Let $(\zeta, \psi) \in \mathcal{C}_T$ be as in Lemma 5.2 with $(\tilde{\zeta}, \tilde{\psi})$ replaced with (ζ^*, ψ^*) and σ replaced with $\sigma/3$. Denote the associated control by $\hat{\varphi}$. Thus $\hat{\varphi} \in \mathcal{S}_T(\zeta, \psi)$. It is now immediate that part (a) and the second inequality in part (b) of Lemma 5.1 are satisfied. Furthermore, since ℓ is a convex function and $\ell(1) = 0$ we can assume without loss of generality (and no change to the state trajectory) that

$$(5.5) \quad \begin{aligned} \hat{\varphi}_0(t, y) &= \hat{\varphi}_0(t) \\ \hat{\varphi}_k(t, y) &= \hat{\varphi}_k(t) \mathbf{1}_{[0, r_k(s)]}(y) + \mathbf{1}_{[r_k(s), 1]}(y), \quad k \in \mathbb{N}. \end{aligned}$$

An ε -gap modification to the thinning function. Fix $\varepsilon > 0$ and define a new control φ^ε with an ε -gap around $r_k(s)$, defined as

$$(5.6) \quad \begin{aligned} \varphi_0^\varepsilon(t, y) &= \hat{\varphi}_0(t, y), \\ \varphi_k^\varepsilon(t, y) &= \frac{\hat{\rho}_k(t)}{1 - \varepsilon} \mathbf{1}_{[0, (1-\varepsilon)r_k(t))}(y) + \mathbf{1}_{[(1+\varepsilon)r_k(t), 1]}(y), \quad k \in \mathbb{N}. \end{aligned}$$

This new control φ^ε results in the same trajectory (ζ, ψ) and the (possible) increase in cost can be estimated as follows:

$$\begin{aligned} &\sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} [\vartheta_k \ell(\varphi_k^\varepsilon(s, y)) - \vartheta_k \ell(\hat{\varphi}_k(s, y))] ds dy \\ &= \sum_{k=1}^{\infty} \int_0^T \left[(1 - \varepsilon)r_k(s) \ell\left(\frac{\hat{\rho}_k(s)}{1 - \varepsilon}\right) + 2\varepsilon r_k(s) \ell(0) - r_k(s) \ell(\hat{\rho}_k(s)) \right] ds \\ &= \sum_{k=1}^{\infty} \int_0^T r_k(s) \left[\hat{\rho}_k(s) \log\left(\frac{1}{1 - \varepsilon}\right) + \varepsilon \right] ds \\ &\leq \sum_{k=1}^{\infty} \int_0^T r_k(s) \left[(\ell(\hat{\rho}_k(s)) + 2) \log\left(\frac{1}{1 - \varepsilon}\right) + \varepsilon \right] ds \\ &\leq \left(I_T(\zeta, \psi) + \frac{1}{3}\sigma \right) \log\left(\frac{1}{1 - \varepsilon}\right) + 2T \log\left(\frac{1}{1 - \varepsilon}\right) + T\varepsilon. \end{aligned}$$

Since $I_T(\zeta, \psi) < \infty$ there exists an $\varepsilon_0 > 0$ such that

$$(5.7) \quad \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} [\vartheta_k \ell(\varphi_k^\varepsilon(s, y)) - \vartheta_k \ell(\hat{\varphi}_k(s, y))] ds dy \leq \frac{1}{3}\sigma$$

for all $\varepsilon \leq \varepsilon_0$.

A final modification of controls. We make one last modification. Fix $\varepsilon < \varepsilon_0 \wedge \frac{1}{3T}\sigma$. Define $\varphi \doteq \varphi^\varepsilon$ except that, for $t \in (t_i, t_{i+1})$, $i \in \{0, 1, \dots, N-1\}$, if $c_i \geq 2$, set $\varphi_{c_i-1}(t, y) = \mathbf{1}_{[\varepsilon, 1]}(y)$. This new control φ still satisfies property (ii) and results in the same trajectory, namely $\varphi \in \mathcal{S}_T(\zeta, \psi)$, since $\pi(t) = c_i$ for $t \in (t_i, t_{i+1})$ by property (i). Furthermore, the control φ only incurs a small additional cost which can be estimated as follows:

$$\sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} [\vartheta_k \ell(\varphi_k(s, y)) - \vartheta_k \ell(\varphi_k^\varepsilon(s, y))] ds dy \leq T\varepsilon \ell(0) \leq \frac{1}{3}\sigma.$$

Combining this, (5.7), and property (v), yields

$$\begin{aligned} \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy &\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k^\varepsilon(s, y)) ds dy + \frac{1}{3}\sigma \\ &\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\hat{\varphi}_k(s, y)) ds dy + \frac{2}{3}\sigma \\ &\leq I_T(\zeta, \psi) + \sigma, \end{aligned}$$

which completes the proof of part (b).

We now show that with the above choice of φ , part (c) holds. Suppose there is another pair $(\tilde{\zeta}, \tilde{\psi}) \in \mathcal{C}_T$ such that $\varphi \in \mathcal{S}_T(\tilde{\zeta}, \tilde{\psi})$. Define time τ by

$$\tau \doteq \inf\{t \in [0, T] : (\tilde{\zeta}(t), \tilde{\psi}(t)) \neq (\zeta(t), \psi(t))\} \wedge T.$$

We argue by contradiction and suppose that $\tau < T$. Then $\tau \in [t_i, t_{i+1})$ for some $i \in \{0, 1, \dots, N-1\}$. Let $\tilde{K} \doteq \pi(t_i) \geq K \doteq c_i$ where the inequality is a consequence of the upper semicontinuity of π (see (5.3)). By continuity, we have $(\tilde{\zeta}(\tau), \tilde{\psi}(\tau), \tilde{\eta}(\tau)) = (\zeta(\tau), \psi(\tau), \eta(\tau))$. We claim that there exists some $\delta \in (0, t_{i+1} - \tau)$ such that

$$(5.8) \quad \tilde{\eta}_k(t) = \eta_k(t), \quad t \in [\tau, \tau + \delta], k \in \mathbb{N}.$$

Assuming for the moment that the claim holds, we now complete the proof of part (c). Define Y and Δ_k by

$$Y(t) \doteq \sum_{k=0}^{\infty} |\tilde{r}_k(t) - r_k(t)|,$$

$$\Delta_k(t) \doteq \tilde{\zeta}_k(t) - \zeta_k(t), \quad k \in \mathbb{N}_0,$$

where r_k and \tilde{r}_k are as defined in (5.4) using ζ and $\tilde{\zeta}$ respectively. We first argue that for $Y(t)$, the differentiation, under the summation over k , with respect to t is valid for a.e. $t \in [\tau, \tau + \delta]$. From the dominated convergence theorem, it suffices to give a summable bound on $\frac{d}{dt} |\tilde{r}_k(t) - r_k(t)|$ for $k \geq M \vee (\tilde{M} + 2)$, where M is as in property (vi) and $\tilde{M} \doteq \min\{j \in \mathbb{N} : \sup_{t \in [0, T]} \zeta_j(t) < 1, \sup_{t \in [0, T]} \tilde{\zeta}_j(t) < 1\}$ (which is finite by property (iii) of \mathcal{C}_T). Note that for $k \geq M \vee (\tilde{M} + 2)$ and a.e. $t \in [0, T]$, $\eta_{k-1}(t) = \eta_k(t) = \eta_{k+1}(t) = \tilde{\eta}_{k-1}(t) = \tilde{\eta}_k(t) = \tilde{\eta}_{k+1}(t) = 0$. Hence,

$$\begin{aligned} \frac{d}{dt} |\tilde{r}_k(t) - r_k(t)| &= \frac{d}{dt} |\tilde{\psi}_k(t) - \tilde{\psi}_{k+1}(t) - \psi_k(t) + \psi_{k+1}(t)| \\ &\leq \tilde{r}_k(t) + \tilde{r}_{k+1}(t) + r_k(t) + r_{k+1}(t) \leq \tilde{\zeta}_k(t) + \zeta_k(t) \leq 2x_k, \end{aligned}$$

where the equality uses (5.4) and (2.7), the first inequality uses (2.10) and property (vi), and the last inequality uses (2.10) again. Since x_k is summable by Remark 2.3, we have the desired property of differentiation under the summation.

Differentiating Y we have that for a.e. $t \in [\tau, \tau + \delta]$,

$$\begin{aligned} Y'(t) &= \sum_{k=0}^{\infty} [(\tilde{r}'_k(t) - r'_k(t)) \mathbf{1}_{\{\tilde{r}_k(t) > r_k(t)\}} + (r'_k(t) - \tilde{r}'_k(t)) \mathbf{1}_{\{\tilde{r}_k(t) < r_k(t)\}}] \\ &= \sum_{k=0}^{\infty} [(\Delta'_k(t) - \Delta'_{k+1}(t)) \mathbf{1}_{\{\tilde{r}_k(t) > r_k(t)\}} + (\Delta'_{k+1}(t) - \Delta'_k(t)) \mathbf{1}_{\{\tilde{r}_k(t) < r_k(t)\}}] \\ &= \sum_{k=1}^{\infty} \Delta'_k(t) [\mathbf{1}_{\{\tilde{r}_k(t) > r_k(t)\}} - \mathbf{1}_{\{\tilde{r}_k(t) < r_k(t)\}} - \mathbf{1}_{\{\tilde{r}_{k-1}(t) > r_{k-1}(t)\}} + \mathbf{1}_{\{\tilde{r}_{k-1}(t) < r_{k-1}(t)\}}], \end{aligned}$$

where the last line follows from rearranging terms and the fact that $\Delta'_0(t) = 0$. For any $k \geq 1$, when $\tilde{r}_k(t) > r_k(t)$, we have

$$\begin{aligned} &\mathbf{1}_{\{\tilde{r}_k(t) > r_k(t)\}} - \mathbf{1}_{\{\tilde{r}_k(t) < r_k(t)\}} - \mathbf{1}_{\{\tilde{r}_{k-1}(t) > r_{k-1}(t)\}} + \mathbf{1}_{\{\tilde{r}_{k-1}(t) < r_{k-1}(t)\}} \\ &= 1 - \mathbf{1}_{\{\tilde{r}_{k-1}(t) > r_{k-1}(t)\}} + \mathbf{1}_{\{\tilde{r}_{k-1}(t) < r_{k-1}(t)\}} \geq 0 \end{aligned}$$

and

$$\Delta'_k(t) = \tilde{\zeta}'_k(t) - \zeta'_k(t) = \tilde{\psi}'_k(t) - \psi'_k(t) = \int_0^1 (\mathbf{1}_{[0, r_k(t))}(y) - \mathbf{1}_{[0, \tilde{r}_k(t))}(y)) \varphi_k(t, y) dy \leq 0$$

for a.e. $t \in [\tau, \tau + \delta]$, by (5.8) and (2.7). Similarly, when $\tilde{r}_k(t) < r_k(t)$, we have

$$\mathbf{1}_{\{\tilde{r}_k(t) > r_k(t)\}} - \mathbf{1}_{\{\tilde{r}_k(t) < r_k(t)\}} - \mathbf{1}_{\{\tilde{r}_{k-1}(t) > r_{k-1}(t)\}} + \mathbf{1}_{\{\tilde{r}_{k-1}(t) < r_{k-1}(t)\}} \leq 0$$

and $\Delta'_k(t) \geq 0$ for a.e. $t \in [\tau, \tau + \delta]$. From these we have $Y'(t) \leq 0$ and hence $Y(t) = Y(\tau) = 0$ for $t \in [\tau, \tau + \delta]$. Therefore $\tilde{r}_k(t) = r_k(t)$ and hence $\tilde{\zeta}_k(t) = \zeta_k(t)$ for all $t \in [\tau, \tau + \delta]$ and $k \geq 1$. Thus, we have shown that $\zeta(t) = \tilde{\zeta}(t)$, and therefore $\tilde{\psi}(t) = \psi(t)$, for $t \in [\tau, \tau + \delta]$. This contradicts the definition of τ and completes the proof of part (c).

Finally we verify the claim that there exist δ such that (5.8) holds. Recall that $\tilde{K} = \pi(t_i)$ and $K = c_i$. Consider the following two possible cases: (1) $\tau = t_i$ and $\tilde{K} > K$ and (2) $\tau \in (t_i, t_{i+1})$ or $\tilde{K} = K$.

Case (1): $\tau = t_i$ and $\tilde{K} > K$. In this case, we simply take $\delta = \Delta t_i$. From property (ii) $\varphi_0(t, y) = 0$ for $t \in (\tau, \tau + \delta)$ and so we have $\psi'_k(t) \leq 0$ and $\tilde{\psi}'_k(t) \leq 0$ for each $k \geq 1$. Therefore, $\eta_k, \tilde{\eta}_k$ stay constant over $[\tau, \tau + \delta]$ and hence (5.8) holds.

Case (2): $\tau \in (t_i, t_{i+1})$ or $\tilde{K} = K$. In this case $\pi(\tau) = K$ and hence $\tilde{r}_K(\tau) = r_K(\tau) > 0$ and $\tilde{r}_k(\tau) = r_k(\tau) = 0$ for $k \leq K - 1$. Recall the fixed ε introduced below (5.7). By continuity, there exists some $\delta \in (0, t_{i+1} - \tau)$ such that for $t \in [\tau, \tau + \delta]$,

$$(5.9) \quad \tilde{r}_K(t) > 0, \quad r_K(t) > 0,$$

$$(5.10) \quad |\tilde{r}_K(t) - r_K(t)| \leq \varepsilon r_K(t),$$

$$(5.11) \quad \tilde{r}_k(t) < \varepsilon, \quad r_k(t) < \varepsilon, \quad k \leq K - 1.$$

The first inequality in (5.9) implies that $\tilde{\pi}(t) \leq K$ for $t \in [\tau, \tau + \delta]$, where $\tilde{\pi}(t) \doteq \max\{k : \tilde{\zeta}_k(t) = 1\}$. Fix $t \in [\tau, \tau + \delta]$. Since $\pi(t) = K \geq \tilde{\pi}(t)$, we have $\zeta_{K+1}(t) < 1, \tilde{\zeta}_{K+1}(t) < 1$, and hence (5.8) holds for $k \geq K + 1$ using the fact that $\eta_k, \tilde{\eta}_k$ do not increase for these coordinates. It now remains to show (5.8) for $k \leq K$. We consider the following three different cases.

$K = 0$: In this case, (5.8) holds trivially for $k \leq K$.

$K = 1$: We only need to check (5.8) for $k = 1$. Note that, for $t \in [\tau, \tau + \delta]$,

$$\begin{aligned} \tilde{\psi}_1(t) &= \tilde{\psi}_1(\tau) + \lambda \int_{[\tau, t] \times [0, 1]} \varphi_0(s, y) ds dy - \int_{[\tau, t] \times [0, 1]} \mathbf{1}_{[0, \tilde{r}_1(s))}(y) \varphi_1(s, y) ds dy \\ &= \psi_1(\tau) + \lambda \int_{[\tau, t] \times [0, 1]} \varphi_0(s, y) ds dy - \int_{[\tau, t] \times [0, 1]} \mathbf{1}_{[0, r_1(s))}(y) \varphi_1(s, y) ds dy = \psi_1(t), \end{aligned}$$

where the second line uses (5.10) and the ε -gap property of φ_1 in (5.6). By the uniqueness of solutions of the one-dimensional SP, $\tilde{\eta}_1(t) = \eta_1(t)$ for $t \in [\tau, \tau + \delta]$, which gives (5.8) for $k = 1$.

$K \geq 2$: In this case $\varphi_{K-1}(t, y) = \mathbf{1}_{[\varepsilon, 1]}(y)$ (see below (5.7)). This and (5.11) yield,

$$\tilde{\psi}'_{K-1}(t) = \mathbf{1}_{\{K=2\}} \lambda \int_0^1 \varphi_0(t, y) dy - \int_0^1 \mathbf{1}_{[0, \tilde{r}_{K-1}(t))}(y) \varphi_{K-1}(t, y) dy = \psi'_{K-1}(t) \geq 0$$

for a.e. t . Therefore, for $t \in [\tau, \tau + \delta]$, $\tilde{\psi}_{K-1}(t) = \psi_{K-1}(t)$, $1 \geq \tilde{\zeta}_{K-1}(t) \geq \tilde{\zeta}_{K-1}(\tau) = 1$, and thus

$$\tilde{\zeta}_k(t) = \zeta_k(t) = 1, \quad k \leq K - 1,$$

$$\tilde{\psi}_k(t) = \psi_k(t), \quad k \leq K - 1,$$

$$\begin{aligned} \tilde{\psi}_K(t) &= \tilde{\psi}_K(\tau) - \int_{[\tau, t] \times [0, 1]} \mathbf{1}_{[0, \tilde{r}_K(s))}(y) \varphi_K(s, y) ds dy \\ &= \psi_K(\tau) - \int_{[\tau, t] \times [0, 1]} \mathbf{1}_{[0, r_K(s))}(y) \varphi_K(s, y) ds dy = \psi_K(t), \end{aligned}$$

where the first line uses property (i) and the decreasing order of $\zeta_k, \tilde{\zeta}_k$, the second line uses (2.9) and (2.10), and the last line uses (5.10) and the ε -gap property of φ_K in (5.6). These together imply

$$\tilde{\eta}_k(t) = \eta_k(t), \quad k \leq K,$$

by the uniqueness of solutions of the K -dimensional SP (associated with (\mathbb{V}^K, R_K) ; see proof of Lemma 2.2). Hence (5.8) holds, which completes the proof. \square

5.3. Proof of Lemma 5.2. In this section we present the proof of Lemma 5.2. A key ingredient is Lemma 5.4 in which we show that one can suitably approximate a $(\tilde{\zeta}, \tilde{\psi}) \in \mathcal{C}_T$ that has a finite cost by a more regular trajectory for which the length of the shortest queue $\pi(t)$ switches only a finite number of times. This is accomplished by “smoothing” out small excursions during which $\pi(t)$ may change an infinite number of times. Proof of Lemma 5.4 uses an inductive argument and for ease of presentation we first present the key inductive step separately in Lemma 5.3. Using this result we then complete the proof of Lemma 5.4. Finally, we use Lemma 5.4 to complete the proof of Lemma 5.2.

LEMMA 5.3. *Fix $\sigma \in (0, 1)$ and an integer $K \geq 2$. Suppose $(\tilde{\zeta}, \tilde{\psi}) \in \mathcal{C}_T$ and $I_T(\tilde{\zeta}, \tilde{\psi}) < \infty$. Further suppose that there is a $\tilde{N} \in \mathbb{N}$ and a finite partition $0 = \tilde{t}_0 < \tilde{t}_1 < \dots < \tilde{t}_{\tilde{N}} = T$ such that on each $(\tilde{t}_i, \tilde{t}_{i+1})$, $\tilde{\pi}(t)$ is either less than K , or some constant $\tilde{c}_i \geq K$. Then there exists a $(\zeta, \psi) \in \mathcal{C}_T$ such that:*

- (a) *There exists a $N \in \mathbb{N}$ and a finite partition $0 = t_0 < t_1 < \dots < t_N = T$, such that on each (t_i, t_{i+1}) , $\pi(t)$ is either less than $K - 1$, or some constant $c_i \geq K - 1$.*
- (b) $\|(\zeta, \psi) - (\tilde{\zeta}, \tilde{\psi})\|_\infty \leq \sigma$.
- (c) $I_T(\zeta, \psi) \leq I_T(\tilde{\zeta}, \tilde{\psi}) + \sigma$.

PROOF. Let σ, K and $(\tilde{\zeta}, \tilde{\psi}) \in \mathcal{C}_T$ be as in the statement of the lemma. Fix $\tilde{\varphi} \in \mathcal{S}_T(\tilde{\zeta}, \tilde{\psi})$ such that

$$(5.12) \quad \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\tilde{\varphi}_k(s, y)) ds dy \leq I_T(\tilde{\zeta}, \tilde{\psi}) + \frac{\sigma}{2}.$$

Let $\tilde{\eta} = \bar{\Gamma}(\tilde{\psi})$, where $\bar{\Gamma}$ was introduced below Definition 2.1, that is, $(\tilde{\zeta}, \tilde{\eta})$ solves the SP for $\tilde{\psi}$ associated with the reflection matrix R_∞ . Since $(\tilde{\zeta}_k, \tilde{\psi}_k)_{k=1}^{K-1}$ are uniformly continuous on $[0, T]$, there exists some $\varepsilon \leq \sigma/4\tilde{N}K(\lambda + 1)$ such that

$$(5.13) \quad \begin{aligned} & \|(\tilde{\zeta}_k(s_1), \tilde{\psi}_k(s_1)) - (\tilde{\zeta}_k(s_2), \tilde{\psi}_k(s_2))\| \\ & \leq \frac{\sigma}{4\tilde{N}K(\lambda + 1)} \quad \text{for all } 1 \leq k \leq K - 1, |s_1 - s_2| \leq \varepsilon. \end{aligned}$$

From the finiteness of the cost we can assume without loss of generality that ε is such that

$$(5.14) \quad \sum_{k=0}^{\infty} \int_{B \times [0, 1]} \vartheta_k \ell(\tilde{\varphi}_k(s, y)) ds dy \leq \frac{\sigma}{4\tilde{N}K(\lambda + 1)} \quad \text{whenever } \text{Leb}(B) \leq \varepsilon,$$

where Leb is the Lebesgue measure on $[0, T]$. With the above preparation, let $(\zeta(0), \psi(0), \eta(0)) \doteq (\tilde{\zeta}(0), \tilde{\psi}(0), \tilde{\eta}(0))$ and consider the interval $(\tilde{t}_i, \tilde{t}_{i+1})$ for each $i = 0, 1, \dots, \tilde{N} - 1$. In the argument that follows we will inductively construct $(\zeta, \psi, \eta, \varphi)$. The sets A, B, C, U, U_j will be introduced which depend on i , however, for ease of notation we will sometimes suppress this dependence on i when it is clear from context. Consider the following two possible cases for $\tilde{\pi}$.

Case 1: $\tilde{\pi}(t) = \tilde{c}_i \geq K$ for every $t \in (\tilde{t}_i, \tilde{t}_{i+1})$. In this case we define

$$(5.15) \quad \begin{aligned} \zeta(t) & \doteq \tilde{\zeta}(t), & \psi(t) & \doteq \psi(\tilde{t}_i) + \tilde{\psi}(t) - \tilde{\psi}(\tilde{t}_i), \\ \eta(t) & \doteq \eta(\tilde{t}_i) + \tilde{\eta}(t) - \tilde{\eta}(\tilde{t}_i), & \varphi(t, y) & \doteq \tilde{\varphi}(t, y) \end{aligned}$$

for $t \in (\tilde{t}_i, \tilde{t}_{i+1}]$, $y \in [0, 1]$.

Case 2: $\tilde{\pi}(t) < K$ for every $t \in (\tilde{t}_i, \tilde{t}_{i+1})$. Consider the set U of time instants in which the shortest queue is less than $K - 1$, namely

$$U \doteq \{t \in (\tilde{t}_i, \tilde{t}_{i+1}) : \tilde{\pi}(t) < K - 1\}.$$

From the upper semicontinuity of $\tilde{\pi}(t)$ it follows that U is open and hence $U = \bigcup_{j=1}^{\infty} U_j$ for some disjoint open intervals U_j . Since the Lebesgue measure $\text{Leb}(U) < \tilde{t}_{i+1} - \tilde{t}_i < \infty$, we can express $U = A \cup B$, where $A \equiv A_i = \bigcup_{j=1}^m U_j$ is a union of finitely many U_j 's for some $m \in \mathbb{N}$, and $B \equiv B_i = \bigcup_{j=m+1}^{\infty} U_j$ with $\text{Leb}(B) < \varepsilon$. Let $C = (\tilde{t}_i, \tilde{t}_{i+1}) \setminus U$ be the set of time instants in $(\tilde{t}_i, \tilde{t}_{i+1})$ at which the shortest queue length is $K - 1$, that is,

$$(5.16) \quad C \doteq \{t \in (\tilde{t}_i, \tilde{t}_{i+1}) : \tilde{\pi}(t) = K - 1\},$$

and define the new trajectory as follows. When the shortest queue is of length $K - 1$ or on a long excursion from $K - 1$ (i.e., $t \in A \cup C$) the trajectory remains unchanged. Namely, let

$$(5.17) \quad \zeta(t) \doteq \tilde{\zeta}(t), \quad \varphi(t, y) \doteq \tilde{\varphi}(t, y), \quad t \in A \cup C, y \in [0, 1].$$

Over short excursions from $K - 1$, namely when $t \in B$, we “smooth out” the trajectories by setting the shortest queue equal to $K - 1$ as follows. For $t \in B$ and $y \in [0, 1]$, let

$$(5.18) \quad \begin{aligned} \zeta_k(t) &\doteq 1, & k \leq K - 1, & \zeta_k(t) \doteq \tilde{\zeta}_k(t) < 1, & k \geq K, \\ \varphi_0(t, y) &\doteq 0, \\ \varphi_k(t, y) &\doteq 0, & k \leq K - 1, & \varphi_k(t, y) \doteq \tilde{\varphi}_k(t, y), & k \geq K. \end{aligned}$$

Having defined ζ and φ over $(\tilde{t}_i, \tilde{t}_{i+1})$, let for $t \in (\tilde{t}_i, \tilde{t}_{i+1})$,

$$(5.19) \quad \begin{aligned} \psi_1(t) &\doteq \psi_1(\tilde{t}_i) + \int_{[\tilde{t}_i, t] \times [0, 1]} \varphi_0(s, y) ds dy \\ &\quad - \int_{[\tilde{t}_i, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_1(s) - \zeta_2(s))}(y) \varphi_1(s, y) ds dy, \\ \psi_k(t) &\doteq \psi_k(\tilde{t}_i) - \int_{[\tilde{t}_i, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \varphi_k(s, y) ds dy, \quad k \geq 2, \\ \eta(t) &\doteq \eta(\tilde{t}_i) + \int_{[\tilde{t}_i, t] \setminus B} \tilde{\eta}(ds) = \eta(\tilde{t}_i) + \int_{[\tilde{t}_i, t] \cap (A \cup C)} \tilde{\eta}(ds), \end{aligned}$$

and define $(\zeta(\tilde{t}_{i+1}), \psi(\tilde{t}_{i+1}), \eta(\tilde{t}_{i+1}))$ by continuity.

Now we verify that (ζ, ψ, η) is the required trajectory, namely: (1) $(\zeta, \psi) \in \mathcal{C}_T$, (2) $\eta = \bar{\Gamma}(\psi)$, (3) $\varphi \in \mathcal{S}_T(\zeta, \psi)$ and, (4) parts (a)–(c) of the lemma are satisfied.

We will refer to $(\tilde{t}_i, \tilde{t}_{i+1})$ as a type 1 (resp. type 2) interval if it corresponds to Case 1 (resp. Case 2) and begin by making the following observations.

- For $k \geq K$, $(\zeta_k, \psi_k, \eta_k) = (\tilde{\zeta}_k, \tilde{\psi}_k, \tilde{\eta}_k)$. Indeed, the equality of the first coordinate $(\zeta_k = \tilde{\zeta}_k)$ is immediate from (5.15), (5.17), and (5.18). The second coordinate equality $(\psi_k = \tilde{\psi}_k)$ follows from the fact that if $\psi_k(\tilde{t}_i) = \tilde{\psi}_k(\tilde{t}_i)$ then, by (5.15) and the second line of (5.19), $\psi_k(t) = \tilde{\psi}_k(t)$ for all $t \in [\tilde{t}_i, \tilde{t}_{i+1}]$. Similarly, the equality for the third coordinate $(\eta_k = \tilde{\eta}_k)$ follows from the fact that if $\eta_k(\tilde{t}_i) = \tilde{\eta}_k(\tilde{t}_i)$, then (5.15), the third line of (5.19), and the fact that $\tilde{\eta}_k$ stays constant over $[\tilde{t}_i, \tilde{t}_{i+1}]$ for a type 2 interval, implies $\eta_k(t) = \tilde{\eta}_k(t)$ for all $t \in [\tilde{t}_i, \tilde{t}_{i+1}]$.
- If $(\tilde{t}_i, \tilde{t}_{i+1})$ is a type 1 interval, then $\tilde{\zeta}_k(\tilde{t}_i) = \zeta_k(\tilde{t}_i)$ for all $k \in \mathbb{N}_0$. Indeed, the only case we need to consider is when $k \leq K - 1$, $(\tilde{t}_{i-1}, \tilde{t}_i)$ is a type 2 interval, and (making the dependence on i explicit) $\tilde{t}_i \in \bar{B}_{i-1}$. In this case $\zeta_k(\tilde{t}_i) = 1$ and since $\tilde{\pi}(\tilde{t}_i) \geq \tilde{\pi}(t) \geq K$ for $t \in (\tilde{t}_i, \tilde{t}_{i+1})$, we have $\tilde{\zeta}_k(\tilde{t}_i) = 1$ as well.

- The first two observations together with (5.15), (5.17), and (5.18) show that ζ_k is absolutely continuous for every $k \in \mathbb{N}_0$. Also by construction, ψ_k and η_k are absolutely continuous as well for every $k \in \mathbb{N}$.
- ζ clearly satisfies parts (i)–(iii) in the definition of \mathcal{C}_T .
- From the definition of B , (5.15), and the third line of (5.19) we see that $\tilde{\eta}_{K-1}(t) = \eta_{K-1}(t)$ for all $t \in [0, T]$.

From the above observations we see that equations (2.7) and (2.9)–(2.10) hold for all $i \geq K$. We now verify that (2.7) and (2.9)–(2.10) hold for $1 \leq i \leq K-1$ as well. From the definition of $\tilde{\eta}$, for $k = 1, \dots, K$,

$$(5.20) \quad \begin{aligned} \tilde{\zeta}_k(t) &= \tilde{\psi}_k(t) + \tilde{\eta}_k(t) - \tilde{\eta}_k(t), \\ \tilde{\eta}_k \text{ is nondecreasing,} \quad \tilde{\eta}_k(0) &= 0, \quad \int_0^t \mathbf{1}_{\{\tilde{\zeta}_k(s) < 1\}} \tilde{\eta}_k(ds) = 0, \quad k = 1, \dots, K. \end{aligned}$$

It follows from (5.19) that for each t in a type 2 interval $(\tilde{t}_i, \tilde{t}_{i+1}]$ and $k \leq K$,

$$(5.21) \quad \begin{aligned} \int_{[\tilde{t}_i, t]} \mathbf{1}_{\{\zeta_k(s) < 1\}} \eta_k(ds) &= \int_{[\tilde{t}_i, t] \cap (A \cup C)} \mathbf{1}_{\{\zeta_k(s) < 1\}} \tilde{\eta}_k(ds) \\ &= \int_{[\tilde{t}_i, t] \cap (A \cup C)} \mathbf{1}_{\{\tilde{\zeta}_k(s) < 1\}} \tilde{\eta}_k(ds) = 0, \end{aligned}$$

where the first equality on the second line is from (5.17). From (5.15) it is clear that the above equality also holds when $(\tilde{t}_i, \tilde{t}_{i+1}]$ is a type 1 interval.

It then remains to show

$$(5.22) \quad \zeta_k(t) = \psi_k(t) + \eta_{k-1}(t) - \eta_k(t), \quad k \leq K-1, t \in (\tilde{t}_i, \tilde{t}_{i+1}).$$

Once again, if $(\tilde{t}_i, \tilde{t}_{i+1})$ is a type 1 interval, from the observation in the second bullet above, $\tilde{\zeta}_k(\tilde{t}_i) = \zeta_k(\tilde{t}_i)$ and thus from (5.15) we see that (5.22) is satisfied in this case. We now show (5.22) for a type 2 interval, given that (5.22) holds for $t = \tilde{t}_i$. For each $k \leq K-1$, we have

$$(5.23a) \quad \begin{aligned} \zeta_k(t) - \zeta_k(\tilde{t}_i) &= \int_{\tilde{t}_i}^t \zeta'_k(s) ds \\ &= \int_{[\tilde{t}_i, t] \cap (A \cup C)} \zeta'_k(s) ds \end{aligned}$$

$$(5.23b) \quad \begin{aligned} &= \int_{[\tilde{t}_i, t] \cap (A \cup C)} \tilde{\zeta}'_k(s) ds \end{aligned}$$

$$(5.23c) \quad \begin{aligned} &= \int_{[\tilde{t}_i, t] \cap (A \cup C)} (\tilde{\psi}'_k(s) + \tilde{\eta}'_{k-1}(s) - \tilde{\eta}'_k(s)) ds \end{aligned}$$

$$(5.23d) \quad \begin{aligned} &= \int_{[\tilde{t}_i, t] \cap (A \cup C)} \psi'_k(s) ds + (\eta_{k-1}(t) - \eta_{k-1}(\tilde{t}_i)) - (\eta_k(t) - \eta_k(\tilde{t}_i)) \end{aligned}$$

$$(5.23e) \quad \begin{aligned} &= \int_{\tilde{t}_i}^t \psi'_k(s) ds + (\eta_{k-1}(t) - \eta_{k-1}(\tilde{t}_i)) - (\eta_k(t) - \eta_k(\tilde{t}_i)) \\ &= (\psi_k(t) - \psi_k(\tilde{t}_i)) + (\eta_{k-1}(t) - \eta_{k-1}(\tilde{t}_i)) - (\eta_k(t) - \eta_k(\tilde{t}_i)), \end{aligned}$$

where line (5.23a) uses (5.18), line (5.23b) uses (5.17), line (5.23c) uses (5.20), line (5.23d) uses (5.17) and (5.19), and line (5.23e) uses (5.18) and (5.19). We have thus shown that (5.22) is satisfied. Next, from (5.15), (5.17), (5.18), (5.19) and (5.21) it is clear that $\int_0^T \mathbf{1}_{\{\zeta_k(s) < 1\}} \eta_k(ds) = 0$. Combining the above observations we now have that $(\zeta, \psi) \in \mathcal{C}_T$ and that $\eta = \Gamma(\psi)$ proving statements (1) and (2). Also from (5.15) and (5.19) it is clear that $\varphi \in \mathcal{S}_T(\zeta, \psi)$, proving statement (3).

Finally we prove statement (4), namely that parts (a)–(c) of the lemma hold. For part (a), note that A and $B \cup C = (\tilde{t}_i, \tilde{t}_{i+1}) \setminus A$ are both finite unions of disjoint intervals. Also, by construction,

$$(5.24) \quad \pi(t) = K - 1 \quad \text{for } t \in B \cup C = (\tilde{t}_i, \tilde{t}_{i+1}) \setminus A \quad \text{and} \quad \pi(t) < K - 1 \quad \text{for } t \in A.$$

Therefore part (a) holds.

We now consider part (b). Fix $1 \leq k \leq K - 1$. For $t \in (\tilde{t}_i, \tilde{t}_{i+1})$, where the latter is a type 2 interval,

$$\begin{aligned} & |\psi_k(t) - \tilde{\psi}_k(t)| \\ & \leq |\psi_k(\tilde{t}_i) - \tilde{\psi}_k(\tilde{t}_i)| + \lambda \int_{B \times [0,1]} |\varphi_0(s, y) - \tilde{\varphi}_0(s, y)| ds dy \\ & \quad + \int_{B \times [0,1]} |\mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \varphi_k(s, y) - \mathbf{1}_{[0, \tilde{\zeta}_k(s) - \tilde{\zeta}_{k+1}(s))}(y) \tilde{\varphi}_k(s, y)| ds dy \\ & \leq |\psi_k(\tilde{t}_i) - \tilde{\psi}_k(\tilde{t}_i)| + \int_{B \times [0,1]} (\lambda \tilde{\varphi}_0(s, y) + \tilde{\varphi}_k(s, y)) ds dy \\ & \leq |\psi_k(\tilde{t}_i) - \tilde{\psi}_k(\tilde{t}_i)| + \int_{B \times [0,1]} (\lambda(\ell(\tilde{\varphi}_0(s, y)) + 2) + \ell(\tilde{\varphi}_k(s, y)) + 2) ds dy \\ & \leq |\psi_k(\tilde{t}_i) - \tilde{\psi}_k(\tilde{t}_i)| + \frac{\sigma}{4\tilde{N}K(\lambda+1)} + 2(\lambda+1)\varepsilon, \end{aligned}$$

where the second inequality uses $\varphi_0 = \varphi_k = 0$ on B , the third uses Lemma 3.2(b), and the last uses (5.14). Also, the above inequality clearly holds for a type 1 interval.

Recalling the definition of ε , we have

$$\|\psi_k - \tilde{\psi}_k\|_\infty \leq \tilde{N} \left(\frac{\sigma}{4\tilde{N}K(\lambda+1)} + 2(\lambda+1)\varepsilon \right) \leq \frac{3\sigma}{4K}.$$

It follows from (5.17) that, on a type 2 interval, $\zeta(t) = \tilde{\zeta}(t)$ for $t \in A \cup C$. While for $t \in B$, we must have $t \in U_j \doteq (u_j, s_j)$ for some $j \in \mathbb{N}$, such that $u_j \in C$ and $|t - u_j| < \varepsilon$. It then follows from (5.18), (5.16), and (5.13) that

$$|\zeta_k(t) - \tilde{\zeta}_k(t)| = |1 - \tilde{\zeta}_k(t)| = |\tilde{\zeta}_k(u_j) - \tilde{\zeta}_k(t)| \leq \frac{\sigma}{4\tilde{N}K(\lambda+1)} \leq \frac{\sigma}{4K}.$$

Once again, on a type 1 interval the above inequality holds trivially. Combining above estimates gives

$$\|(\zeta, \psi) - (\tilde{\zeta}, \tilde{\psi})\|_\infty \leq \sum_{k=1}^{K-1} \|\zeta_k - \tilde{\zeta}_k\|_\infty + \sum_{k=1}^{K-1} \|\psi_k - \tilde{\psi}_k\|_\infty \leq \frac{\sigma}{4} + \frac{3\sigma}{4} = \sigma.$$

This verifies part (b).

Finally we consider part (c). Using (5.17), (5.18), and the definitions of B and ε we have for a type 2 interval

$$\begin{aligned} & \sum_{k=0}^{\infty} \int_{[\tilde{t}_i, \tilde{t}_{i+1}] \times [0,1]} [\vartheta_k \ell(\varphi_k(s, y)) - \vartheta_k \ell(\tilde{\varphi}_k(s, y))] ds dy \\ & = \sum_{k=0}^{K-1} \int_{B \times [0,1]} [\vartheta_k \ell(\varphi_k(s, y)) - \vartheta_k \ell(\tilde{\varphi}_k(s, y))] ds dy \\ & \leq K(\lambda+1)\ell(0) \frac{\sigma}{4\tilde{N}K(\lambda+1)} = \frac{\sigma}{4\tilde{N}}. \end{aligned}$$

The above bound holds clearly for a type 1 interval. From this and (5.12) we have

$$\begin{aligned} I_T(\zeta, \psi) &\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy \\ &\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\tilde{\varphi}_k(s, y)) ds dy + \frac{\sigma}{4} \\ &\leq I_T(\tilde{\zeta}, \tilde{\psi}) + \sigma. \end{aligned}$$

This gives part (c) and completes the proof of the lemma. \square

With Lemma 5.3 in hand, we can now use an inductive argument to prove the following lemma, which will play a key role in the proof of Lemma 5.2.

LEMMA 5.4. *Fix $\sigma \in (0, 1)$ and $(\tilde{\zeta}, \tilde{\psi}) \in \mathcal{C}_T$ with $I_T(\tilde{\zeta}, \tilde{\psi}) < \infty$. There exists $(\zeta, \psi) \in \mathcal{C}_T$ such that:*

- (a) *There exists a $N \in \mathbb{N}$ and a finite partition $0 = t_0 < t_1 < \dots < t_N = T$ such that $\pi(t)$ is constant over each (t_i, t_{i+1}) .*
- (b) $\|(\zeta, \psi) - (\tilde{\zeta}, \tilde{\psi})\|_{\infty} \leq \sigma$.
- (c) $I_T(\zeta, \psi) \leq I_T(\tilde{\zeta}, \tilde{\psi}) + \sigma$.

PROOF. Let $M \in \mathbb{N}$ be the smallest nonnegative integer such that

$$\frac{1}{M+1} \sup_{0 \leq t \leq T} \sum_{k=0}^M \tilde{\zeta}_k(t) < 1$$

and thus

$$\sup_{0 \leq t \leq T} \tilde{\zeta}_M(t) \leq \frac{1}{M+1} \sup_{0 \leq t \leq T} \sum_{k=0}^M \tilde{\zeta}_k(t) < 1.$$

Existence of such a M is a consequence of property (iii) of \mathcal{C}_T . Therefore $\tilde{\zeta}(t) \doteq \max\{k : \tilde{\zeta}_k(t) = 1\} < M$ for every $t \in [0, T]$. Let $(\zeta^M, \psi^M) \doteq (\tilde{\zeta}, \tilde{\psi})$. For $K = M, M-1, \dots, 2$, apply Lemma 5.3 to (ζ^K, ψ^K) recursively, with σ there replaced by σ/M , to get $(\zeta^{K-1}, \psi^{K-1})$. Then $(\zeta, \psi) \doteq (\zeta^1, \psi^1)$ is the desired trajectory, on noting that over each (t_i, t_{i+1}) , $\pi(t)$ is either less than 1, which means $\pi(t) = 0$, or $\pi(t)$ is some constant $c_i \geq 1$. \square

We now prove Lemma 5.2 by further modifying the trajectory in Lemma 5.4 so that it has nice properties when $\pi(t)$ changes.

PROOF OF LEMMA 5.2. Fix $\sigma \in (0, 1)$ and $(\tilde{\zeta}, \tilde{\psi}) \in \mathcal{C}_T$ with $I_T(\tilde{\zeta}, \tilde{\psi}) < \infty$. By Lemma 5.4, there exists $(\bar{\zeta}, \bar{\psi}) \in \mathcal{C}_T$ such that

- (a) There exists a $N \in \mathbb{N}$ and a finite partition $0 = \bar{t}_0 < \bar{t}_1 < \dots < \bar{t}_{\bar{N}} = T$ such that $\bar{\pi}(t) = \bar{c}_i$ is constant over each $(\bar{t}_i, \bar{t}_{i+1})$.
- (b) $\|(\bar{\zeta}, \bar{\psi}) - (\tilde{\zeta}, \tilde{\psi})\|_{\infty} \leq \frac{\sigma}{16}$.
- (c) $I_T(\bar{\zeta}, \bar{\psi}) \leq I_T(\tilde{\zeta}, \tilde{\psi}) + \frac{\sigma}{16}$.

Here $\bar{\pi}(t) \doteq \max\{k : \bar{\xi}_k(t) = 1\}$. Let $\bar{M} \doteq \max\{\bar{\pi}(t) : t \in [0, T]\} < \infty$. Then with $\bar{\eta} = \bar{\Gamma}(\bar{\psi})$, $\bar{\eta}_k(t) = 0$ for all $t \in [0, T]$ and $k > \bar{M}$. Choose $\bar{\varphi} \in \mathcal{S}_T(\bar{\xi}, \bar{\psi})$ that is $\sigma/16$ optimal so that

$$(5.25) \quad \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\bar{\varphi}_k(s, y)) ds dy \leq I_T(\bar{\xi}, \bar{\psi}) + \frac{\sigma}{8}.$$

We will next modify $(\bar{\xi}, \bar{\psi}, \bar{\varphi})$ to get (ξ, ψ, φ) that satisfies properties (ii) and (v) in the statement while preserving properties (i), (iii), and (iv) that are satisfied by $(\bar{\xi}, \bar{\psi})$. Finally we will make one additional modification that will guarantee that property (vi) holds as well.

Since $(\bar{\xi}_k, \bar{\psi}_k)_{k=0}^{\bar{M}}$ are uniformly continuous on $[0, T]$, there exists some $\varepsilon_{\bar{M}} \in (0, \infty)$ such that

$$(5.26) \quad \|(\bar{\xi}_k(s_1), \bar{\psi}_k(s_1))_{k=0}^{\bar{M}} - (\bar{\xi}_k(s_2), \bar{\psi}_k(s_2))_{k=0}^{\bar{M}}\| \leq \frac{\sigma}{8\bar{N}} \quad \text{whenever } |s_1 - s_2| \leq \varepsilon_{\bar{M}}.$$

From the finiteness of the cost in (5.25)

$$\sum_{k=0}^{\bar{M}+1} \int_{[0, T] \times [0, 1]} \vartheta_k \bar{\varphi}_k(t, y) dt dy \leq \sum_{k=0}^{\bar{M}+1} \int_{[0, T] \times [0, 1]} \vartheta_k (\ell(\bar{\varphi}_k(t, y)) + 2) dt dy < \infty,$$

where the first inequality is from Lemma 3.2 (b). Thus we can assume without loss of generality that $\varepsilon_{\bar{M}}$ is small enough so that

$$(5.27) \quad \sum_{k=0}^{\bar{M}+1} \int_{B \times [0, 1]} \vartheta_k \bar{\varphi}_k(t, y) dt dy \leq \frac{\sigma}{16\bar{N}} \quad \text{whenever } \text{Leb}(B) \leq \varepsilon_{\bar{M}}.$$

Consider the interval $[\bar{t}_i, \bar{t}_{i+1})$ for each fixed $i = 0, 1, \dots, \bar{N} - 1$. We will usually suppress the dependence on i (of K, \bar{K}, δ below) for ease of notation.

First modification of the trajectory.

Case I: Either $\bar{\pi}(\bar{t}_i) = \bar{c}_i$ or $\int_{[\bar{t}_i, \bar{t}_i + \Delta \bar{t}_i] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy = 0$ for some $\Delta \bar{t}_i \in (0, \bar{t}_{i+1} - \bar{t}_i)$. In this case define, for $t \in (\bar{t}_i, \bar{t}_{i+1})$

$$(5.28) \quad \begin{aligned} \zeta(t) &= \bar{\zeta}(t), & \psi(t) &= \psi(\bar{t}_i) + \bar{\psi}(t) - \bar{\psi}(\bar{t}_i), \\ \eta(t) &= \eta(\bar{t}_i) + \bar{\eta}(t) - \bar{\eta}(\bar{t}_i), & \varphi(t) &= \bar{\varphi}(t). \end{aligned}$$

Case II: $\bar{\pi}(\bar{t}_i) > \bar{c}_i$ and $\int_{[\bar{t}_i, \bar{t}_i + \Delta \bar{t}_i] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy > 0$ for every $\Delta \bar{t}_i \in (0, \bar{t}_{i+1} - \bar{t}_i)$. In this case we modify $(\bar{\xi}, \bar{\psi}, \bar{\varphi})$ as follows.

Take $\varepsilon = \varepsilon_{\bar{M}} \wedge \frac{\sigma}{16\bar{N}(\lambda+1)} \wedge \min_i (\bar{t}_{i+1} - \bar{t}_i)$. Let $\bar{K} \doteq \bar{\pi}(\bar{t}_i) > \bar{c}_i \doteq K$.

We first claim that we can assume that $(\bar{\xi}, \bar{\psi})$ and $\bar{\varphi} \in \mathcal{S}_T(\bar{\xi}, \bar{\psi})$ are such that $\bar{\varphi}_K(t) = 0$ for $t \in (\bar{t}_i, \bar{t}_i + \varepsilon)$ whenever $K > 0$, and parts (a)–(c) and (5.25) hold with $\sigma/16$ and $\sigma/8$ replaced with $3\sigma/16$.

To see this, suppose $K > 0$. For $t \in (\bar{t}_i, \bar{t}_i + \varepsilon)$, since $\bar{\pi}(t) = K$, we have

$$\begin{aligned} \bar{\xi}_k(t) &= 1, & k &= 0, \dots, K, \\ \bar{r}_k(t) &= 0, & k &= 0, \dots, K-1, & \bar{r}_K(t) &> 0, \\ \bar{\eta}_k(t) &= \bar{\eta}_k(\bar{t}_i) + \lambda \int_{[\bar{t}_i, t] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy, & k &= 1, \dots, K-1. \end{aligned}$$

Note that the last statement is vacuous when $K = 1$. The first property (with $k = K$) implies that

$$\bar{\eta}_{K-1}(t) + \bar{\psi}_K(t) - \bar{\eta}_{K-1}(s) - \bar{\psi}_K(s) \geq 0 \quad \forall \bar{t}_i < s < t < \bar{t}_i + \varepsilon,$$

and hence for a.e. $t \in (\bar{t}_i, \bar{t}_i + \varepsilon)$

$$(5.29) \quad \bar{\eta}'_{K-1}(t) + \bar{\psi}'_K(t) = \lambda \int_0^1 \bar{\varphi}_0(t, y) dy - \int_0^1 \mathbf{1}_{[0, \bar{t}_K(t)]}(y) \bar{\varphi}_K(t, y) dy \geq 0.$$

We now modify $(\bar{\zeta}, \bar{\psi}, \bar{\varphi})$ as follows. Replace $\bar{\varphi}_K(t, y)$, $\bar{\varphi}_0(t, y)$, and $\bar{\psi}_K$, for $t \in (\bar{t}_i, \bar{t}_i + \varepsilon]$, by

$$\begin{aligned} \bar{\varphi}_K^{\text{new}}(t, y) &= 0, & \bar{\varphi}_0^{\text{new}}(t, y) &= \lambda^{-1}(\bar{\eta}'_{K-1}(t) + \bar{\psi}'_K(t)), & y \in [0, 1], \\ (\bar{\psi}_K^{\text{new}})'(t) &= \mathbf{1}_{\{K=1\}} \lambda \int_0^1 \bar{\varphi}_0^{\text{new}}(t, y) dy. \end{aligned}$$

For $t \in (\bar{t}_i + \varepsilon, \bar{t}_{i+1})$, we set

$$\begin{aligned} \bar{\varphi}_j^{\text{new}}(t, \cdot) &= \bar{\varphi}_j(t, \cdot), & j &= 0, K, \\ \bar{\psi}_K^{\text{new}}(t) - \bar{\psi}_K^{\text{new}}(\bar{t}_i + \varepsilon) &= \bar{\psi}_K(t) - \bar{\psi}_K(\bar{t}_i + \varepsilon). \end{aligned}$$

Also, for $t \in (\bar{t}_i, \bar{t}_{i+1})$, define

$$\begin{aligned} (\bar{\psi}_j^{\text{new}})'(t) &= \bar{\psi}_j'(t), & j &\neq 1, K, \\ \bar{\varphi}_j^{\text{new}}(t, \cdot) &= \bar{\varphi}_j(t, \cdot), & j &\neq 0, K, \\ \bar{\zeta}_j^{\text{new}}(t) &= \bar{\zeta}_j(t), & j &\in \mathbb{N}_0 \end{aligned}$$

and

$$(\bar{\psi}_1^{\text{new}})'(t) = \lambda \int_0^1 \bar{\varphi}_0^{\text{new}}(t, y) dy \quad \text{if } K > 1.$$

It is easy to check that the new trajectory is still in \mathcal{C}_T and has the finite partition property (a) as $(\bar{\zeta}, \bar{\psi})$.

The contribution to the difference between the two trajectories over the interval $(\bar{t}_i, \bar{t}_i + \varepsilon)$ can be estimated as

$$\begin{aligned} (5.30) \quad & \int_{\bar{t}_i}^{\bar{t}_i + \varepsilon} |\bar{\psi}'_K(t) - (\bar{\psi}_K^{\text{new}})'(t)| dt + \int_{\bar{t}_i}^{\bar{t}_i + \varepsilon} |\bar{\psi}_1(t) - (\bar{\psi}_1^{\text{new}})'(t)| dt \\ & \leq 2\lambda \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} \bar{\varphi}_0(t, y) dt dy \\ & \leq \frac{\sigma}{8\bar{N}}, \end{aligned}$$

where the last inequality is due to (5.27).

The additional cost of making such a replacement is

$$\begin{aligned} (5.31) \quad & \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} (\lambda \ell(\bar{\varphi}_0^{\text{new}}(t, y)) + \ell(\bar{\varphi}_K^{\text{new}}(t, y)) - \lambda \ell(\bar{\varphi}_0(t, y)) - \ell(\bar{\varphi}_K(t, y))) dt dy \\ & \leq \int_{\bar{t}_i}^{\bar{t}_i + \varepsilon} \left(\lambda \ell \left(\frac{\bar{\eta}'_{K-1}(t) + \bar{\psi}'_K(t)}{\lambda} \right) + \ell(0) - \lambda \ell \left(\int_0^1 \bar{\varphi}_0(t, y) dy \right) - 0 \right) dt \\ & \leq \int_{\bar{t}_i}^{\bar{t}_i + \varepsilon} (\lambda \ell(0) + \ell(0)) dt = (\lambda + 1)\varepsilon \leq \frac{\sigma}{16\bar{N}}, \end{aligned}$$

where the first inequality follows from the convexity of $\ell(\cdot)$ and the second inequality uses (5.29) and the fact that $\ell(x) - \ell(y) \leq \ell(0)$ for $0 \leq x \leq y$.

Combining the contributions to the errors and cost differences over all intervals $(\bar{t}_i, \bar{t}_{i+1})$, we have

$$\|(\bar{\zeta}^{\text{new}}, \bar{\psi}^{\text{new}}) - (\tilde{\zeta}, \tilde{\psi})\|_{\infty} \leq \|(\bar{\zeta}^{\text{new}}, \bar{\psi}^{\text{new}}) - (\bar{\zeta}, \bar{\psi})\|_{\infty} + \frac{\sigma}{16} \leq \frac{3\sigma}{16}$$

and

$$\begin{aligned} I_T(\bar{\zeta}^{\text{new}}, \bar{\psi}^{\text{new}}) &\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\bar{\varphi}_k^{\text{new}}(s, y)) ds dy \\ &\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\bar{\varphi}_k(s, y)) ds dy + \frac{\sigma}{16} \\ &\leq I_T(\tilde{\zeta}, \tilde{\psi}) + \frac{3\sigma}{16}. \end{aligned}$$

We have thus proved the claim.

Abusing notation, we denote $(\bar{\zeta}^{\text{new}}, \bar{\psi}^{\text{new}}, \bar{\varphi}^{\text{new}})$ once more as $(\bar{\zeta}, \bar{\psi}, \bar{\varphi})$ and recall that parts (a)–(c) and (5.25) hold with $\sigma/16$ and $\sigma/8$ replaced with $3\sigma/16$.

Since $\bar{\varphi}_K(t, y) = 0$ for $t \in (\bar{t}_i, \bar{t}_i + \varepsilon)$ when $K > 0$, we have

$$(5.32) \quad \bar{\eta}_k(t) = \bar{\eta}_k(\bar{t}_i) + \lambda \int_{[\bar{t}_i, t] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy, \quad k = 1, \dots, K.$$

We now return to constructing our modification of $(\bar{\zeta}, \bar{\psi}, \bar{\varphi})$ on $[\bar{t}_i, \bar{t}_{i+1})$ under Case II. Since $\bar{\zeta}_{K+1}(t) < 1$ over $(\bar{t}_i, \bar{t}_{i+1})$ and $\bar{K} > K$ in Case II, we have

$$\begin{aligned} 0 &> \bar{\zeta}_{K+1}(\bar{t}_i + \varepsilon) - 1 = \bar{\zeta}_{K+1}(\bar{t}_i + \varepsilon) - \bar{\zeta}_{K+1}(\bar{t}_i) \\ &= \lambda \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy - \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy. \end{aligned}$$

Let δ be the largest value in $(0, \varepsilon)$ such that

$$(5.33) \quad \lambda \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy = \int_{[\bar{t}_i, \bar{t}_i + \delta] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy.$$

We now modify $(\bar{\zeta}, \bar{\psi}, \bar{\varphi})$ on the time interval $[\bar{t}_i, \bar{t}_{i+1})$ as follows. Let

$$\varphi_0(t, y) = 0, \quad t \in [\bar{t}_i, \bar{t}_i + \varepsilon],$$

$$\varphi_{K+1}(t, y) = 0, \quad t \in [\bar{t}_i, \bar{t}_i + \delta],$$

$$\varphi_{K+1}(t, y) = \bar{\varphi}_{K+1}(t, y) \mathbf{1}_{[0, \bar{r}_{K+1}(t))}(y), \quad t \in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon],$$

$$\varphi_k(t, y) = \bar{\varphi}_k(t, y), \quad t \in [\bar{t}_i, \bar{t}_i + \varepsilon], k \neq 0, K+1.$$

$$\varphi_k(t, y) = \bar{\varphi}_k(t, y), \quad t \in [\bar{t}_i + \varepsilon, \bar{t}_{i+1}) \text{ for all } k \in \mathbb{N}_0.$$

Having made such a modification to the control over each interval $(\bar{t}_i, \bar{t}_{i+1})$, consider (ζ, ψ, η) driven by φ , given (on $(\bar{t}_i, \bar{t}_{i+1})$) as follows:

$$\begin{aligned}
 \eta'_k(t) &= 0, \quad k \geq 1, t \in [\bar{t}_i, \bar{t}_i + \varepsilon], \\
 \zeta_k(t) &= \bar{\zeta}_k(t) = 1, \quad \psi'_k(t) = 0, \quad k = 0, \dots, K, t \in [\bar{t}_i, \bar{t}_i + \varepsilon], \\
 \zeta_k(t) &= \bar{\zeta}_k(t), \quad \psi'_k(t) = \bar{\psi}'_k(t), \quad k \geq K + 2, t \in [\bar{t}_i, \bar{t}_i + \varepsilon], \\
 \zeta_{K+1}(t) &= 1, \quad t \in [\bar{t}_i, \bar{t}_i + \delta], \\
 \zeta_{K+1}(t) &= 1 - \int_{[\bar{t}_i + \delta, t] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy, \\
 (5.34) \quad t &\in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon], \\
 \psi'_{K+1}(t) &= 0, \quad t \in [\bar{t}_i, \bar{t}_i + \delta], \\
 \psi_{K+1}(t) &= \psi_{K+1}(\bar{t}_i + \delta) - \int_{[\bar{t}_i + \delta, t] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy, \\
 t &\in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon], \\
 (\zeta'(t), \psi'(t), \eta'(t)) &= (\bar{\zeta}'(t), \bar{\psi}'(t), \bar{\eta}'(t)), \quad t \in (\bar{t}_i + \varepsilon, \bar{t}_{i+1}).
 \end{aligned}$$

We now check that $(\zeta, \psi) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\zeta, \psi)$. For this it suffices to check the evolution of the K th and $(K+1)$ th coordinates on each $[\bar{t}_i, \bar{t}_i + \varepsilon]$.

If $K = 0$, then $\psi'_0 = \zeta'_0 = \eta' = 0$ by construction, and it is clear that (2.7) holds for the K th coordinate.

If $K > 0$, then $\zeta'_K = \psi'_K = \eta' = \varphi_0 = \varphi_K = \bar{\varphi}_K = 0$ by construction and the claim made below (5.28) (which has been verified), showing that (2.7), (2.9), and (2.10) hold once more for the K th coordinate. Therefore we have the desired evolution of the K th coordinate.

For $K \in \mathbb{N}_0$ and $t \in [\bar{t}_i, \bar{t}_i + \delta]$, since $\zeta'_{K+1} = \psi'_{K+1} = \eta' = \varphi_0 = \varphi_{K+1} = 0$, we have (2.7), (2.9), and (2.10) for the $(K+1)$ th coordinate over this interval. As for $t \in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon]$, note that

$$\begin{aligned}
 &\zeta_{K+1}(t) - \bar{\zeta}_{K+1}(t) \\
 &= \left(1 - \int_{[\bar{t}_i + \delta, t] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy \right) \\
 &\quad - \left(1 + \lambda \int_{[\bar{t}_i, t] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy - \int_{[\bar{t}_i, t] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy \right) \\
 &= \int_{[\bar{t}_i, \bar{t}_i + \delta] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy - \lambda \int_{[\bar{t}_i, t] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy \\
 &= \lambda \int_{[t, \bar{t}_i + \varepsilon] \times [0, 1]} \bar{\varphi}_0(s, y) ds dy \geq 0,
 \end{aligned}$$

where the first equality uses (5.32) and the last equality follows from (5.33). The above equality in particular implies that

$$(5.35) \quad \zeta_{K+1}(\bar{t}_i + \varepsilon) = \bar{\zeta}_{K+1}(\bar{t}_i + \varepsilon),$$

and the inequality $\zeta_{K+1}(t) - \bar{\zeta}_{K+1}(t) \geq 0$ together with the fact that $\zeta_{K+2}(t) = \bar{\zeta}_{K+2}(t)$ for $t \in (\bar{t}_i, \bar{t}_{i+1})$ gives $r_{K+1}(t) \geq \bar{r}_{K+1}(t)$ for $t \in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon]$. From this and the definition of φ_{K+1} , ζ_{K+1} , ψ_{K+1} , we see that for $t \in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon]$,

$$\zeta_{K+1}(t) = 1 - \int_{[\bar{t}_i + \delta, t] \times [0, 1]} \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy$$

$$\begin{aligned}
&= \xi_{K+1}(\bar{t}_i + \delta) - \int_{[\bar{t}_i + \delta, t] \times [0, 1]} \mathbf{1}_{[0, r_{K+1}(s))}(y) \mathbf{1}_{[0, \bar{r}_{K+1}(s))}(y) \bar{\varphi}_{K+1}(s, y) ds dy \\
&= \xi_{K+1}(\bar{t}_i + \delta) - \int_{[\bar{t}_i + \delta, t] \times [0, 1]} \mathbf{1}_{[0, r_{K+1}(s))}(y) \varphi_{K+1}(s, y) ds dy
\end{aligned}$$

and similarly

$$\psi_{K+1}(t) = \psi_{K+1}(\bar{t}_i + \delta) - \int_{[\bar{t}_i + \delta, t] \times [0, 1]} \mathbf{1}_{[0, r_{K+1}(s))}(y) \varphi_{K+1}(s, y) ds dy.$$

From these two displays along with (5.34) and the observation that $\zeta'_{K+1} - \psi'_{K+1} = 0 = \eta'$ we have verified (2.7), (2.9), and (2.10) for the $(K+1)$ th coordinate when $t \in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon]$. This proves that $(\zeta, \psi) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\zeta, \psi)$.

Since δ is the largest value in $(0, \varepsilon)$ such that (5.33) holds, from the definition of ζ_{K+1} we must have $\zeta_{K+1}(t) < 1$ for $t \in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon]$. This implies that $\pi(t) = K$ is constant for $t \in (\bar{t}_i + \delta, \bar{t}_i + \varepsilon]$.

On $[\bar{t}_i, \bar{t}_i + \delta]$, since $\varphi_0(t, y) = 0$, $\pi(t)$ must be nonincreasing, and hence must be a piecewise constant function which can be decomposed into a finite number of intervals. Therefore properties (i) and (ii) hold.

We now estimate the (possible) increase in cost. Note that

$$\begin{aligned}
&\int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} (\lambda \ell(\varphi_0(s, y)) + \ell(\varphi_{K+1}(s, y)) - \lambda \ell(\bar{\varphi}_0(s, y)) - \ell(\bar{\varphi}_{K+1}(s, y))) ds dy \\
&\leq \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} (\lambda + 1) \ell(0) ds dy \leq (\lambda + 1) \varepsilon.
\end{aligned}$$

Therefore the cost after making such modifications to each interval $(\bar{t}_i, \bar{t}_i + \varepsilon)$ can be bounded as

$$\begin{aligned}
(5.36) \quad &\sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy \\
&\leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\bar{\varphi}_k(s, y)) ds dy + (\lambda + 1) \bar{N} \varepsilon \\
&\leq I_T(\tilde{\zeta}, \tilde{\psi}) + \frac{3\sigma}{16} + \frac{\sigma}{16} = I_T(\tilde{\zeta}, \tilde{\psi}) + \frac{\sigma}{4}.
\end{aligned}$$

The difference in the trajectories is estimated as follows. Since $\zeta_{K+1}(t)$ is a monotone decreasing interpolation of $\zeta_{K+1}(\bar{t}_i)$ and $\zeta_{K+1}(\bar{t}_i + \varepsilon)$ for $t \in [\bar{t}_i, \bar{t}_i + \varepsilon]$, we have

$$\begin{aligned}
&\|\zeta_{K+1}(t) - \bar{\zeta}_{K+1}(t)\| \\
&\leq \|\zeta_{K+1}(\bar{t}_i) - \bar{\zeta}_{K+1}(t)\| + \|\zeta_{K+1}(\bar{t}_i + \varepsilon) - \bar{\zeta}_{K+1}(t)\| \\
&\leq \|\zeta_{K+1}(\bar{t}_i) - \bar{\zeta}_{K+1}(\bar{t}_i)\| + \frac{\sigma}{8\bar{N}} + \|\zeta_{K+1}(\bar{t}_i + \varepsilon) - \bar{\zeta}_{K+1}(\bar{t}_i + \varepsilon)\| + \frac{\sigma}{8\bar{N}} \\
&= \frac{\sigma}{4\bar{N}},
\end{aligned}$$

where the second inequality uses (5.26) and the last equality follows from (5.35). Therefore

$$(5.37) \quad \|\zeta - \bar{\zeta}\|_{\infty} \leq \frac{\sigma}{4\bar{N}}.$$

For each $k = 1, \dots, \bar{M} + 1$ and $t \in [\bar{t}_i, \bar{t}_i + \varepsilon]$, using the definition of ψ_k one has

$$|\psi_k(t) - \bar{\psi}_k(t)| \leq |\psi_k(\bar{t}_i) - \bar{\psi}_k(\bar{t}_i)| + \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} (\lambda \bar{\varphi}_0(s, y) + \bar{\varphi}_k(s, y)) ds dy.$$

As a result

$$\|\psi - \bar{\psi}\|_\infty \leq \sum_{k=1}^{\bar{M}+1} \|\psi_k - \bar{\psi}_k\|_\infty \leq \sum_{i=0}^{\bar{N}-1} \sum_{k=1}^{\bar{M}+1} \int_{[\bar{t}_i, \bar{t}_i + \varepsilon] \times [0, 1]} (\lambda \bar{\varphi}_0(s, y) + \bar{\varphi}_k(s, y)) ds dy \leq \frac{\sigma}{8}$$

by (5.27). From this and (5.37) we have

$$\|(\zeta, \psi) - (\bar{\zeta}, \bar{\psi})\|_\infty \leq \frac{\sigma}{16} + \frac{\sigma}{4\bar{N}} \leq \frac{5\sigma}{16}.$$

This combined with $\|(\bar{\zeta}, \bar{\psi}) - (\tilde{\zeta}, \tilde{\psi})\|_\infty \leq 3\sigma/16$ gives property (iii) with σ replaced by $\sigma/2$.

Second modification of the trajectory.

We now introduce one last modification to (ζ, ψ, φ) so that property (vi) holds. Take $M > \bar{M} + 2$ large enough such that

$$(5.38) \quad \sum_{k=M}^{\infty} \frac{\|(\zeta_k, \psi_k)\|_\infty + 2}{2^k} \leq \frac{\sigma}{2}.$$

Define $(\zeta^{\text{new}}, \psi^{\text{new}}) \in \mathcal{C}_T$ and $\varphi^{\text{new}} \in \mathcal{S}_T(\zeta^{\text{new}}, \psi^{\text{new}})$ by

$$(\zeta_k^{\text{new}}, \psi_k^{\text{new}}, \varphi_k^{\text{new}}) = (\zeta_k, \psi_k, \varphi_k), \quad k < M,$$

$$\varphi_k^{\text{new}} = 1, \quad \zeta_k^{\text{new}}(t) = \psi_k^{\text{new}}(t) = x_k - \int_0^t (\zeta_k^{\text{new}}(s) - \zeta_{k+1}^{\text{new}}(s)) ds, \quad k \geq M.$$

Note that

$$(5.39) \quad \begin{aligned} \|(\zeta^{\text{new}}, \psi^{\text{new}}) - (\zeta, \psi)\|_\infty &= \sum_{k=M}^{\infty} \frac{\|(\zeta_k^{\text{new}}, \psi_k^{\text{new}}) - (\zeta_k, \psi_k)\|_\infty}{2^k} \\ &\leq \sum_{k=M}^{\infty} \frac{\|(\zeta_k, \psi_k)\|_\infty + 2}{2^k} \leq \frac{\sigma}{2}, \end{aligned}$$

$$(5.40) \quad \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k^{\text{new}}(s, y)) ds dy \leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy,$$

where the last line uses $\ell(1) = 0$. Once again, abusing notation, we denote $(\zeta^{\text{new}}, \psi^{\text{new}}, \varphi^{\text{new}})$ as (ζ, ψ, φ) . Clearly, properties (i), (ii), (iii), (iv), and (vi) are satisfied by the above modification.

For property (v), let $\varphi^\sigma \in \mathcal{S}_T(\zeta, \psi)$ be any $\sigma/2$ optimal control. Choose δ^* sufficiently small so that

$$\sum_{k=0}^{\infty} \sum_{i=0}^{\bar{N}-1} \int_{[t_i, t_i + \delta^*] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy \leq \frac{\sigma}{2}.$$

Define, for each i , $\varphi^*(t) = \varphi(t)$ for $t \in (t_i, t_i + \delta^*)$, and for $t \in [t_i + \delta^*, t_{i+1})$, let $\varphi_k^*(t) = 1$ for $k \geq M$, and $\varphi_k^*(t) = \varphi_k^\sigma(t)$ otherwise.

Then φ^* is a σ optimal control, without affecting properties (ii) or (vi). Thus properties (i)–(vi) are satisfied completing the proof of the lemma. \square

6. Compact sub-level sets. In this section we prove the third statement in the proof of Theorem 2.4, namely the property that I_T is a rate function. For this we need to show that for every $M \in \mathbb{N}$, the set $\Gamma_M \doteq \{(\zeta, \psi) \in \mathbb{D}_{\mathbb{R}^\infty \times \mathbb{R}^\infty} : I_T(\zeta, \psi) \leq M\}$ is compact. Now fix such a M and a sequence $\{(\zeta^n, \psi^n)\} \subset \Gamma_M$. It suffices to show that the sequence has a convergent subsequence with the limit in the set Γ_M . From the definition of I_T , it follows that $(\zeta^n, \psi^n) \in \mathcal{C}_T$ and there exists a control $\varphi^n \in \mathcal{S}_T(\zeta^n, \psi^n)$ such that for every n

$$(6.1) \quad \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k^n(s, u)) \, ds \, dy \leq I_T(\zeta^n, \psi^n) + \frac{1}{n} \leq M + \frac{1}{n}.$$

We follow the convention that $\zeta_0^n = \psi_0^n = 1$. We first show pre-compactness of the sequence $\{(\zeta^n, \psi^n, \varphi^n)\}_{n \in \mathbb{N}_0}$. Recall the compact metric spaces S_N , for $N \in \mathbb{N}$, introduced in Section 3.2.

LEMMA 6.1. *The sequence $\{(\zeta^n, \psi^n, \varphi^n)\}_{n \in \mathbb{N}_0}$ is pre-compact in $\mathbb{C}_{\mathbb{R}^\infty \times \mathbb{R}^\infty} \times S_{M+1}$.*

PROOF. Pre-compactness of $\{\varphi^n\}_{n \in \mathbb{N}_0}$ is immediate from the compactness of S_{M+1} .

We next prove pre-compactness of $\{\psi_k^n\}$ for each fixed $k \in \mathbb{N}$. From the definition of ψ^n in (2.9)–(2.10) we have

$$|\psi_k^n(t)| \leq 1 + \int_{[0, T] \times [0, 1]} \lambda \varphi_0^n(t, y) \, dt \, dy + \int_{[0, T] \times [0, 1]} \varphi_k^n(t, y) \, dt \, dy.$$

It then follows from (6.1) and Lemma 3.2 that

$$\sup_{n \in \mathbb{N}} \sup_{t \in [0, T]} |\psi_k^n(t)| < \infty.$$

We now show that $\{\psi_k^n\}$ is equicontinuous. Note that for any $0 < t - s \leq \delta$ and $K > 0$,

$$\begin{aligned} |\psi_k^n(t) - \psi_k^n(s)| &\leq \lambda \int_{[s, t] \times [0, 1]} \varphi_0^n(u, y) \, du \, dy + \int_{[s, t] \times [0, 1]} \varphi_k^n(u, y) \, du \, dy \\ &\leq \lambda \int_{[s, t] \times [0, 1]} \varphi_0^n(u, y) \mathbf{1}_{\{\varphi_0^n(u, y) > K\}} \, du \, dy \\ &\quad + \lambda \int_{[s, t] \times [0, 1]} \varphi_0^n(u, y) \mathbf{1}_{\{\varphi_0^n(u, y) \leq K\}} \, du \, dy \\ &\quad + \int_{[s, t] \times [0, 1]} \varphi_k^n(u, y) \mathbf{1}_{\{\varphi_k^n(u, y) > K\}} \, du \, dy \\ &\quad + \int_{[s, t] \times [0, 1]} \varphi_k^n(u, y) \mathbf{1}_{\{\varphi_k^n(u, y) \leq K\}} \, du \, dy \\ &\leq (\lambda + 1)\gamma(K)(M + 1) + (\lambda + 1)K\delta, \end{aligned}$$

where the final inequality above follows from Lemma 3.2(a) and (6.1). Therefore,

$$\lim_{\delta \rightarrow \infty} \sup_{n \in \mathbb{N}} \sup_{|t - s| \leq \delta} |\psi_k^n(t) - \psi_k^n(s)| \leq (\lambda + 1)\gamma(K)(M + 1)$$

and equicontinuity of $\{\psi_k^n\}$ follows upon sending $K \rightarrow \infty$. Pre-compactness of $\{\psi_k^n\}$ for each k , and therefore of $\{\psi^n\}$, now follows from the Arzela–Ascoli theorem. Pre-compactness of $\{\zeta^n\}_{n \in \mathbb{N}}$ in $\mathbb{C}_{\mathbb{R}^\infty}$ follows immediately from the precompactness of $\{\psi^n\}$ and the Lipschitz property of the Skorokhod map proved in Lemma 2.2. \square

We now characterize the limit points of $(\zeta^n, \psi^n, \varphi^n)$.

LEMMA 6.2. Suppose $(\zeta^n, \psi^n, \varphi^n)$ converges along a subsequence to $(\zeta, \psi, \varphi) \in \mathbb{C}_{\mathbb{R}^\infty \times \mathbb{R}^\infty} \times S_{M+1}$. Then:

- (a) $\sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \vartheta_k \ell(\varphi_k(s, y)) ds dy \leq M$.
- (b) For each $t \in [0, T]$,

$$\psi_1(t) = x_1 + \lambda \int_{[0, t] \times [0, 1]} \varphi_0(s, y) ds dy - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_1(s) - \zeta_2(s))}(y) \varphi_1(s, y) ds dy,$$

$$\psi_k(t) = x_k - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \varphi_k(s, y) ds dy, \quad k \geq 2.$$

- (c) $(\zeta, \psi) \in \mathcal{C}_T$ and $\varphi \in \mathcal{S}_T(\zeta, \psi)$.

PROOF. Part (a) is immediate from [9], Lemma A.1, (6.1), and Fatou's lemma.

We now prove (b). The convergence

$$(6.2) \quad \int_{[0, t] \times [0, 1]} \varphi_0^n(s, y) ds dy \rightarrow \int_{[0, t] \times [0, 1]} \varphi_0(s, y) ds dy$$

is immediate from the definition of the topology on S_{M+1} (see the comment above Lemma 3.3). Consider now the integral on the right side of (2.10). For each $k \in \mathbb{N}$,

$$\begin{aligned} & \left| \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k^n(s) - \zeta_{k+1}^n(s))}(y) \varphi_k^n(s, y) ds dy - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \varphi_k(s, y) ds dy \right| \\ & \leq \int_{[0, t] \times [0, 1]} \left| \mathbf{1}_{[0, \zeta_k^n(s) - \zeta_{k+1}^n(s))}(y) - \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \right| \varphi_k^n(s, y) ds dy \\ & \quad + \left| \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) (\varphi_k^n(s, y) - \varphi_k(s, y)) ds dy \right|. \end{aligned}$$

Using the convergence of

$$\left| \mathbf{1}_{[0, \zeta_k^n(s) - \zeta_{k+1}^n(s))}(y) - \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \right| \rightarrow 0$$

for Leb_t -a.e. $(s, y) \in [0, t] \times [0, 1]$ and the uniform integrability of $(s, y) \mapsto \varphi_k^n(s, y)$ with respect to the normalized Lebesgue measure on $[0, t] \times [0, 1]$, guaranteed by (6.1) and the superlinearity of ℓ , we have

$$\int_{[0, t] \times [0, 1]} \left| \mathbf{1}_{[0, \zeta_k^n(s) - \zeta_{k+1}^n(s))}(y) - \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \right| \varphi_k^n(s, y) ds dy \rightarrow 0.$$

From the convergence of $\varphi^n \rightarrow \varphi$ and recalling the topology on S_{M+1} , we have that

$$\int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) (\varphi_k^n(s, y) - \varphi_k(s, y)) ds dy \rightarrow 0.$$

This gives the convergence:

$$\int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k^n(s) - \zeta_{k+1}^n(s))}(y) \varphi_k^n(s, y) ds dy \rightarrow \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \varphi_k(s, y) ds dy.$$

Combining this with (6.2) gives (b).

Finally consider part (c). The fact that (ζ, ψ) satisfies property (ii) and (iv) of \mathcal{C}_T is an immediate consequence of the fact that (ζ^n, ψ^n) satisfy these properties and the Lipschitz property of the Skorokhod map proved in Lemma 2.2. Property (i) follows from this and Remark 2.1. For property (iii) note that

$$\begin{aligned} \sup_{0 \leq t \leq T} \sum_{i=1}^{\infty} \zeta_i(t) & \leq \liminf_{n \rightarrow \infty} \sup_{0 \leq t \leq T} \sum_{i=1}^{\infty} \zeta_i^n(t) \\ & \leq \sum_{i=1}^{\infty} x_i + \liminf_{n \rightarrow \infty} \int_{[0, T] \times [0, 1]} \lambda \varphi_0^n(s, y) ds dy < \infty, \end{aligned}$$

where the last inequality is from (6.1) and Remark 2.3. This completes the proof that $(\zeta, \psi) \in \mathcal{C}_T$. The fact that $\varphi \in \mathcal{S}_T(\zeta, \psi)$ is now immediate from part (b). \square

We now return to the proof of compactness of Γ_M . Consider a sequence $\{(\zeta^n, \psi^n)\} \subset \Gamma_M$. Then Lemma 6.1 shows that such a sequence is precompact and Lemma 6.2 shows that any limit point (ζ, ψ) of (ζ^n, ψ^n) is in Γ_M . This establishes the desired compactness.

7. Bounds on probabilities of long queues. In this section we prove Theorem 2.5. Fix $j \geq 3$ and recall the notation G_j, F_j from Section 2.4. From the LDP in Theorem 2.4 and since $G_j \subset F_j$

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n} \log(\mathbb{P}(X^n \in G_j)) &\geq -I_T(G_j), \\ \limsup_{n \rightarrow \infty} \frac{1}{n} \log(\mathbb{P}(X^n \in G_j)) &\leq \limsup_{n \rightarrow \infty} \frac{1}{n} \log(\mathbb{P}(X^n \in F_j)) \leq -I_T(F_j). \end{aligned}$$

In order to prove the first statement in the theorem we first solve for $I_T(F_j)$ and then show that $I_T(F_j) = I_T(G_j)$.

Fix $\varepsilon > 0$, $(\zeta, \psi) \in F_j$, and $\varphi \in \mathcal{S}_T(\zeta, \psi)$ with

$$(7.1) \quad \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \ell(\varphi_k(s, y)) \, ds \, dy \leq I_T(\zeta, \psi) + \varepsilon \leq I_T(F_j) + 2\varepsilon.$$

Define

$$\tau_i \doteq \inf_{t \in [0, T]} \{\zeta_i(t) = 1\}, \quad i \in \mathbb{N}.$$

Since all queues are of length one at time 0, we have that $0 = \tau_1 \leq \tau_2 \leq \dots$ and $\zeta_k(\tau_i) = \mathbf{1}_{\{k \leq i\}}$. We can assume without loss of generality that $\tau_{j-1} = T$. To see this, note that if $\tau_{j-1} < T$, then we can consider the delayed trajectory $(\bar{\zeta}, \bar{\psi})$ defined by

$$\begin{aligned} (\bar{\zeta}(t), \bar{\psi}(t)) &= (\zeta(0), \psi(0)), \quad 0 \leq t \leq T - \tau_{j-1}, \\ (\bar{\zeta}(t), \bar{\psi}(t)) &= (\zeta(t - (T - \tau_{j-1})), \psi(t - (T - \tau_{j-1}))), \quad T - \tau_{j-1} < t \leq T. \end{aligned}$$

Since the cost over time $[0, T - \tau_{j-1}]$ is zero, we have $I_T(\bar{\zeta}, \bar{\psi}) \leq I_T(\zeta, \psi)$. Thus henceforth we assume $\tau_{j-1} = T$.

We can further assume without loss of generality that $\zeta_k(t)$ is nondecreasing in t for each $k \in \mathbb{N}$. To see this, consider the new trajectory $\bar{\zeta}$ defined by

$$\bar{\zeta}_k(t) = \max_{0 \leq u \leq t} \zeta_k(u), \quad k \in \mathbb{N}.$$

Note that this says that for each i and $t \in (\tau_i, \tau_{i+1})$,

$$\bar{\zeta}_k(t) = 1, \quad k \leq i; \quad \bar{\zeta}_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u); \quad \bar{\zeta}_k(t) = 0, \quad k \geq i+2.$$

We claim that $\bar{\zeta}_{i+1}$ is absolutely continuous and $\bar{\zeta}'_{i+1}(t) = \zeta'_{i+1}(t) \mathbf{1}_{\{\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}}$ a.e. $t \in (\tau_i, \tau_{i+1})$, for every i . Absolute continuity is immediate on noting that $0 \leq \zeta_{i+1}(t_2) - \zeta_{i+1}(t_1) \leq \max_{t_1 \leq s \leq t_2} \zeta_{i+1}(s) - \zeta_{i+1}(t_1)$ for $\tau_i \leq t_1 \leq t_2 \leq \tau_{i+1}$, for every i . Note also that for a.e. $t \in (\tau_i, \tau_{i+1})$ such that $\zeta'_{i+1}(t)$ and $\bar{\zeta}'_{i+1}(t)$ exist, if $\zeta_{i+1}(t) < \max_{0 \leq u \leq t} \zeta_{i+1}(u)$, then $\bar{\zeta}'_{i+1}(t) = 0$. On the other hand, if $\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)$, we have two possible cases:

- Case 1: there exists a sequence $t_n \downarrow t$ (i.e., t_n approaches t from above) with $\zeta_{i+1}(t_n) = \max_{0 \leq u \leq t_n} \zeta_{i+1}(u)$. In this case

$$\bar{\zeta}'_{i+1}(t) = \lim_{n \rightarrow \infty} \frac{\bar{\zeta}_{i+1}(t_n) - \bar{\zeta}_{i+1}(t)}{t_n - t} = \lim_{n \rightarrow \infty} \frac{\zeta_{i+1}(t_n) - \zeta_{i+1}(t)}{t_n - t} = \zeta'_{i+1}(t).$$

- Case 2: A sequence t_n as in Case 1 does not exist, namely there exists some $t_0 \in (t, \tau_{i+1})$ such that $\zeta_{i+1}(s) < \max_{0 \leq u \leq s} \zeta_{i+1}(u)$ for all $s \in (t, t_0)$. Then we must have $\bar{\zeta}_{i+1}(s) = \zeta_{i+1}(t)$ for all $s \in (t, t_0)$, and hence $\bar{\zeta}'_{i+1}(t) = 0$. From this and $\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)$ we have $\zeta_{i+1}(s) \leq \zeta_{i+1}(t)$ for all $s \in (0, t_0)$. Therefore $\zeta'_{i+1}(t) = 0 = \bar{\zeta}'_{i+1}(t)$.

This proves the claim.

Define the control $\bar{\varphi}$ over the interval (τ_i, τ_{i+1}) as

$$\bar{\varphi}_k(t, y) = \varphi_k(t, y) \mathbf{1}_{\{\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}} + \mathbf{1}_{\{\zeta_{i+1}(t) < \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}}, \quad k \in \mathbb{N}_0$$

and define the corresponding $\bar{\psi}$ by (2.9) and (2.10) using $\bar{\zeta}$ and $\bar{\varphi}$. Now we show that $(\bar{\zeta}, \bar{\psi}) \in \mathcal{C}_T$. For this it suffices to verify property (iv) of \mathcal{C}_T . Let, for $t \in (\tau_i, \tau_{i+1})$

$$\begin{aligned} \bar{\eta}'_k(t) &= \eta'_k(t) \mathbf{1}_{\{\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}} + \mathbf{1}_{\{\zeta_{i+1}(t) < \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}}, \quad k < i, \\ \bar{\eta}'_i(t) &= \eta'_i(t) \mathbf{1}_{\{\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}} + (1 + \bar{\psi}'_i(t)) \mathbf{1}_{\{\zeta_{i+1}(t) < \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}}, \\ \bar{\eta}'_k(t) &= \eta'_k(t) = 0, \quad k \geq i + 1. \end{aligned}$$

Note that for a.e. t , if $\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)$, then $(\bar{\zeta}', \bar{\psi}', \bar{\eta}') = (\zeta', \psi', \eta')$ and hence

$$\bar{\zeta}'_k(t) = \bar{\psi}'_k(t) + \bar{\eta}'_{k-1}(t) - \bar{\eta}'_k(t), \quad \mathbf{1}_{\{\bar{\zeta}_k(t) < 1\}} \bar{\eta}'_k(t) = 0.$$

If $\zeta_{i+1}(t) < \max_{0 \leq u \leq t} \zeta_{i+1}(u)$, using the definition of $(\bar{\zeta}, \bar{\psi}, \bar{\eta}, \bar{\varphi})$ it can be verified that the above equation still holds. In particular to check the equation for $k = i + 1$ we use the facts $\bar{\psi}'_i(t) + \bar{\psi}'_{i+1}(t) = -1$ and $\bar{\zeta}'_{i+1}(t) = \zeta'_{i+1}(t) \mathbf{1}_{\{\zeta_{i+1}(t) = \max_{0 \leq u \leq t} \zeta_{i+1}(u)\}} = 0$. Therefore (2.7) holds for $(\bar{\zeta}, \bar{\psi}, \bar{\eta})$. Thus we have that $(\bar{\zeta}, \bar{\psi}) \in \mathcal{C}_T$, $\bar{\varphi} \in \mathcal{S}_T(\bar{\zeta}, \bar{\psi})$, and

$$I_T(\bar{\zeta}, \bar{\psi}) \leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \ell(\bar{\varphi}_k(s, y)) ds dy \leq \sum_{k=0}^{\infty} \int_{[0, T] \times [0, 1]} \ell(\varphi_k(s, y)) ds dy.$$

We have therefore shown that one can assume without loss of generality that $\zeta_k(t)$ is non-decreasing in t for each $k \in \mathbb{N}$. Henceforth we will assume that this holds. Note that, in particular this says that $\zeta_1(t) = 1$ for all $t \in [0, T]$.

From (2.7) we have

$$\sum_{k=1}^{\infty} \zeta_k(t) = \sum_{k=1}^{\infty} \psi_k(t).$$

From the above display, (2.9), (2.10), and the assumption that $\pi(t) \doteq \max\{k : \zeta_k(t) = 1\} \leq j - 1$ we have

$$\begin{aligned} 1 &= \sum_{k=1}^{\infty} (\zeta_k(\tau_{i+1}) - \zeta_k(\tau_i)) = \sum_{k=1}^{j-1} (\psi_k(\tau_{i+1}) - \psi_k(\tau_i)) \\ &= \int_{[\tau_i, \tau_{i+1}] \times [0, 1]} \varphi_0(s, y) ds dy - \sum_{k=1}^{j-1} \int_{[\tau_i, \tau_{i+1}] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))} (y) \varphi_k(s, y) ds dy. \end{aligned}$$

Let $\theta_i \doteq (\tau_{i+1} - \tau_i)^{-1}$. Note that, since $\zeta_1(t) = 1$ for every t ,

$$\theta_i \sum_{k=1}^{j-1} \int_{[\tau_i, \tau_{i+1}] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))} (y) ds dy = 1.$$

It then follows from Jensen's inequality and the convexity of ℓ that

$$\begin{aligned}
 & \sum_{k=0}^{\infty} \int_{[\tau_i, \tau_{i+1}] \times [0,1]} \ell(\varphi_k(s, y)) \, ds \, dy \\
 & \geq \int_{[\tau_i, \tau_{i+1}] \times [0,1]} \ell(\varphi_0(s, y)) \, ds \, dy \\
 (7.2) \quad & + \sum_{k=1}^{j-1} \int_{[\tau_i, \tau_{i+1}] \times [0,1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \ell(\varphi_k(s, y)) \, ds \, dy \\
 & \geq \theta_i^{-1} \ell\left(\theta_i \int_{[\tau_i, \tau_{i+1}] \times [0,1]} \varphi_0(s, y) \, ds \, dy\right) \\
 & + \theta_i^{-1} \ell\left(\theta_i \sum_{k=1}^{j-1} \int_{[\tau_i, \tau_{i+1}] \times [0,1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \varphi_k(s, y) \, ds \, dy\right).
 \end{aligned}$$

This quantity can be further bounded from below by

$$\frac{1}{\theta_i} \inf\{\ell(a) + \ell(b) : a, b \geq 0, a - b = c > 0\},$$

where $c = \theta_i$. Using Lagrange multipliers one finds that the above infimum is achieved at

$$a = \frac{c + \sqrt{c^2 + 4}}{2}, \quad b = \frac{-c + \sqrt{c^2 + 4}}{2}.$$

Plugging this back into (7.2) gives

$$\sum_{k=0}^{\infty} \int_{[\tau_i, \tau_{i+1}] \times [0,1]} \ell(\varphi_k(s, y)) \, ds \, dy \geq \theta_i^{-1} \ell\left(\frac{\theta_i + \sqrt{\theta_i^2 + 4}}{2}\right) + \theta_i^{-1} \ell\left(\frac{-\theta_i + \sqrt{\theta_i^2 + 4}}{2}\right).$$

From this, Jensen's inequality, the convexity of ℓ , and the fact that $\tau_1 = 0$, $\tau_{j-1} = T$, $\sum_{i=1}^{j-2} \theta_i^{-1} = T$, we have, letting $a_j = (j-2)/T$,

$$\begin{aligned}
 & \sum_{k=0}^{\infty} \int_{[0, T] \times [0,1]} \ell(\varphi_k(s, y)) \, ds \, dy \\
 & \geq \sum_{i=1}^{j-2} \left[\theta_i^{-1} \ell\left(\frac{\theta_i + \sqrt{\theta_i^2 + 4}}{2}\right) + \theta_i^{-1} \ell\left(\frac{-\theta_i + \sqrt{\theta_i^2 + 4}}{2}\right) \right] \\
 & \geq T \ell\left(\frac{(j-2) + \sum_{i=1}^{j-2} \sqrt{1 + 4\theta_i^{-2}}}{2T}\right) + T \ell\left(\frac{-(j-2) + \sum_{i=1}^{j-2} \sqrt{1 + 4\theta_i^{-2}}}{2T}\right) \\
 & \geq T \ell\left(\frac{(j-2) + (j-2)\sqrt{1 + 4(a_j)^{-2}}}{2T}\right) + T \ell\left(\frac{-(j-2) + (j-2)\sqrt{1 + 4(a_j)^{-2}}}{2T}\right) \\
 & = T \ell\left(\frac{a_j + \sqrt{4 + (a_j)^2}}{2}\right) + T \ell\left(\frac{-a_j + \sqrt{4 + (a_j)^2}}{2}\right),
 \end{aligned}$$

where the last inequality is obtained by the convexity of the function $f(x) \doteq \sqrt{1 + 4x^2}$ and monotonicity of the functions $x \mapsto \ell(x + c) + \ell(x - c)$ for $c \geq 0$ and $x > c$. As $\varepsilon > 0$ in (7.1) is arbitrary, this shows

$$I_T(F_j) \geq T \ell\left(\frac{a_j + \sqrt{4 + (a_j)^2}}{2}\right) + T \ell\left(\frac{-a_j + \sqrt{4 + (a_j)^2}}{2}\right).$$

On the other hand, note that the above lower bound can be achieved by a $(\zeta, \psi) \in I_T(F_j)$ given for $k \in \mathbb{N}$, $t \in [0, T]$, and $y \in [0, 1]$ as

$$\begin{aligned}
 \zeta_k(t) &= 0 \vee (a_j t - (k-2)) \wedge 1, \\
 \varphi_0(t, y) &= \frac{a_j + \sqrt{4 + (a_j)^2}}{2}, \\
 (7.3) \quad \varphi_k(t, y) &= \frac{-a_j + \sqrt{4 + (a_j)^2}}{2} \mathbf{1}_{[0, \zeta_k(t) - \zeta_{k+1}(t))}(y) + \mathbf{1}_{[\zeta_k(t) - \zeta_{k+1}(t), 1]}(y), \\
 \psi_k(t) &= \mathbf{1}_{\{k=1\}} + \mathbf{1}_{\{k=1\}} \int_{[0, t] \times [0, 1]} \varphi_0(s, y) \, ds \, dy \\
 &\quad - \int_{[0, t] \times [0, 1]} \mathbf{1}_{[0, \zeta_k(s) - \zeta_{k+1}(s))}(y) \varphi_k(s, y) \, ds \, dy.
 \end{aligned}$$

Therefore

$$(7.4) \quad I_T(F_j) = T \ell\left(\frac{a_j + \sqrt{4 + (a_j)^2}}{2}\right) + T \ell\left(\frac{-a_j + \sqrt{4 + (a_j)^2}}{2}\right).$$

Next we show that $I_T(F_j) = I_T(G_j)$. Since $G_j \subset F_j$, we clearly have $I_T(F_j) \leq I_T(G_j)$.

Now we show the reverse inequality. Fix $\varepsilon > 0$. Consider the modification of the trajectory (7.3), defined over $[0, T - \varepsilon]$, by replacing T throughout by $T - \varepsilon$. Then $\zeta_k(T - \varepsilon) = \mathbf{1}_{\{k \leq j-1\}}$ and the cost of this trajectory over the interval $[0, T - \varepsilon]$ is given by (7.4) with T replaced by $T - \varepsilon$. Define, for $t \in [T - \varepsilon, T]$ and $y \in [0, 1]$,

$$\begin{aligned}
 \zeta_k(t) &\doteq \zeta_k(T - \varepsilon) \mathbf{1}_{\{k \neq j\}} + (t - (T - \varepsilon)) \mathbf{1}_{\{k=j\}}, \\
 \varphi_k(t, y) &\doteq \mathbf{1}_{\{k=0 \text{ or } k \geq j+1\}},
 \end{aligned}$$

so that $\zeta_j(T) > 0$ and

$$\sum_{k=0}^{\infty} \int_{[T - \varepsilon, T] \times [0, 1]} \ell(\varphi_k(s, y)) \, ds \, dy = j \ell(0) \varepsilon.$$

This trajectory clearly lies in G_j and its cost (over $[0, T]$) converges to (7.4) as $\varepsilon \rightarrow 0$. This implies $I_T(G_j) \leq I_T(F_j)$ proving the reverse inequality and hence showing that $I_T(F_j) = I_T(G_j)$. This completes the proof of the first part of the theorem.

Finally we consider the second part of the theorem. Since $\sqrt{4 + x^2} = 2 + o(x)$ and $\ell(x) = \frac{(x-1)^2}{2} + o((x-1)^2)$, as $x \rightarrow 0$ and $x \rightarrow 1$, respectively, we have

$$-T \ell\left(\frac{a_j + \sqrt{4 + (a_j)^2}}{2}\right) - T \ell\left(\frac{-a_j + \sqrt{4 + (a_j)^2}}{2}\right) = -T \left(\frac{(a_j)^2}{4} + o((a_j)^2)\right)$$

as $T \rightarrow \infty$. Sending $T \rightarrow \infty$, we have

$$\lim_{T \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{T}{n} \log(\mathbb{P}(X^n \in G_j)) = \lim_{T \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{T}{n} \log(\mathbb{P}(X^n \in F_j)) = -\frac{(j-2)^2}{4}.$$

Funding. The first author was supported in part by the National Science Foundation Grants DMS-1814894 and DMS-1853968.

REFERENCES

- [1] AGAZZI, A., DEMBO, A. and ECKMANN, J.-P. (2018). Large deviations theory for Markov jump models of chemical reaction networks. *Ann. Appl. Probab.* **28** 1821–1855. MR3809478 <https://doi.org/10.1214/17-AAP1344>
- [2] ALANYALI, M. and HAJEK, B. (1998). On large deviations of Markov processes with discontinuous statistics. *Ann. Appl. Probab.* **8** 45–66. MR1620409 <https://doi.org/10.1214/aoap/1027961033>
- [3] ATAR, R. and DUPUIS, P. (1999). Large deviations and queueing networks: Methods for rate function identification. *Stochastic Process. Appl.* **84** 255–296. MR1719274 [https://doi.org/10.1016/S0304-4149\(99\)00051-4](https://doi.org/10.1016/S0304-4149(99)00051-4)
- [4] BANERJEE, S. and MUKHERJEE, D. (2019). Join-the-shortest queue diffusion limit in Halfin–Whitt regime: Tail asymptotics and scaling of extrema. *Ann. Appl. Probab.* **29** 1262–1309. MR3910028 <https://doi.org/10.1214/18-AAP1436>
- [5] BHAMIDI, S., BUDHIRAJA, A., DUPUIS, P. and WU, R. (2019). Rare event asymptotics for exploration processes for random graphs. arXiv preprint [arXiv:1912.04714](https://arxiv.org/abs/1912.04714).
- [6] BILLINGSLEY, P. (1999). *Convergence of Probability Measures*, 2nd ed. Wiley Series in Probability and Statistics: Probability and Statistics. Wiley, New York. A Wiley-Interscience Publication. MR1700749 <https://doi.org/10.1002/9780470316962>
- [7] BRAMSON, M., LU, Y. and PRABHAKAR, B. (2012). Asymptotic independence of queues under randomized load balancing. *Queueing Syst.* **71** 247–292. MR2943660 <https://doi.org/10.1007/s11134-012-9311-0>
- [8] BRAVERMAN, A. (2020). Steady-state analysis of the join-the-shortest-queue model in the Halfin–Whitt regime. *Math. Oper. Res.* **45** 1069–1103. MR4135843 <https://doi.org/10.1287/moor.2019.1023>
- [9] BUDHIRAJA, A., CHEN, J. and DUPUIS, P. (2013). Large deviations for stochastic partial differential equations driven by a Poisson random measure. *Stochastic Process. Appl.* **123** 523–560. MR3003362 <https://doi.org/10.1016/j.spa.2012.09.010>
- [10] BUDHIRAJA, A. and DUPUIS, P. (2019). *Analysis and Approximation of Rare Events: Representations and Weak Convergence Methods. Probability Theory and Stochastic Modelling* **94**. Springer, New York. MR3967100 <https://doi.org/10.1007/978-1-4939-9579-0>
- [11] BUDHIRAJA, A., DUPUIS, P. and GANGULY, A. (2016). Moderate deviation principles for stochastic differential equations with jumps. *Ann. Probab.* **44** 1723–1775. MR3502593 <https://doi.org/10.1214/15-AOP1007>
- [12] BUDHIRAJA, A., DUPUIS, P. and MAROULAS, V. (2011). Variational representations for continuous time processes. *Ann. Inst. Henri Poincaré Probab. Stat.* **47** 725–747. MR2841073 <https://doi.org/10.1214/10-AIHP382>
- [13] BUDHIRAJA, A. and FRIEDLANDER, E. (2019). Diffusion approximations for load balancing mechanisms in cloud storage systems. *Adv. in Appl. Probab.* **51** 41–86. MR3984010 <https://doi.org/10.1017/apr.2019.3>
- [14] BUDHIRAJA, A. and WU, R. (2017). Moderate deviation principles for weakly interacting particle systems. *Probab. Theory Related Fields* **168** 721–771. MR3663630 <https://doi.org/10.1007/s00440-016-0723-3>
- [15] DUPUIS, P. and ELLIS, R. S. (1992). Large deviations for Markov processes with discontinuous statistics. II. Random walks. *Probab. Theory Related Fields* **91** 153–194. MR1147614 <https://doi.org/10.1007/BF01291423>
- [16] DUPUIS, P. and ELLIS, R. S. (1995). The large deviation principle for a general class of queueing systems. I. *Trans. Amer. Math. Soc.* **347** 2689–2751. MR1290716 <https://doi.org/10.2307/2154753>
- [17] DUPUIS, P. and ELLIS, R. S. (2011). *A Weak Convergence Approach to the Theory of Large Deviations* **902**. John Wiley & Sons. MR1431744 <https://doi.org/10.1002/9781118165904>
- [18] DUPUIS, P., ELLIS, R. S. and WEISS, A. (1991). Large deviations for Markov processes with discontinuous statistics. I. General upper bounds. *Ann. Probab.* **19** 1280–1297. MR1112416
- [19] DUPUIS, P. and ISHII, H. (1991). On Lipschitz continuity of the solution mapping to the Skorokhod problem, with applications. *Stoch. Stoch. Rep.* **35** 31–62. MR1110990 <https://doi.org/10.1080/17442509108833688>
- [20] DUPUIS, P., RAMANAN, K. and WU, W. (2016). Large deviation principle for finite-state mean field interacting particle systems. arXiv preprint [arXiv:1601.06219](https://arxiv.org/abs/1601.06219).
- [21] ESCHENFELDT, P. and GAMARNIK, D. (2018). Join the shortest queue with many servers. The heavy-traffic asymptotics. *Math. Oper. Res.* **43** 867–886. MR3846076 <https://doi.org/10.1287/moor.2017.0887>
- [22] GRAHAM, C. (2000). Chaoticity on path space for a queueing network with selection of the shortest queue among several. *J. Appl. Probab.* **37** 198–211. MR1761670 <https://doi.org/10.1017/s0021900200015345>

- [23] GUPTA, V. and WALTON, N. (2019). Load balancing in the nondegenerate slowdown regime. *Oper. Res.* **67** 281–294. [MR3919870](https://doi.org/10.1287/opre.2018.1768) <https://doi.org/10.1287/opre.2018.1768>
- [24] HALFIN, S. and WHITT, W. (1981). Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* **29** 567–588. [MR0629195](https://doi.org/10.1287/opre.29.3.567) <https://doi.org/10.1287/opre.29.3.567>
- [25] HARRISON, J. M. and REIMAN, M. I. (1981). Reflected Brownian motion on an orthant. *Ann. Probab.* **9** 302–308. [MR0606992](https://doi.org/10.1214/aop/1019487515)
- [26] IGNATIOUK-ROBERT, I. (2000). Large deviations of Jackson networks. *Ann. Appl. Probab.* **10** 962–1001. [MR1789985](https://doi.org/10.1214/aoap/1019487515) <https://doi.org/10.1214/aoap/1019487515>
- [27] IGNATIOUK-ROBERT, I. (2005). Large deviations for processes with discontinuous statistics. *Ann. Probab.* **33** 1479–1508. [MR2150196](https://doi.org/10.1214/009117905000000189) <https://doi.org/10.1214/009117905000000189>
- [28] IKEDA, N. and WATANABE, S. (1989). *Stochastic Differential Equations and Diffusion Processes*, 2nd ed. *North-Holland Mathematical Library* **24**. North-Holland, Amsterdam. [MR1011252](https://doi.org/10.1011/1252)
- [29] JOFFE, A. and MÉTIVIER, M. (1986). Weak convergence of sequences of semimartingales with applications to multitype branching processes. *Adv. in Appl. Probab.* **18** 20–65. [MR0827331](https://doi.org/10.2307/1427238) <https://doi.org/10.2307/1427238>
- [30] LÉONARD, C. (1995). Large deviations for long range interacting particle systems with jumps. *Ann. Inst. Henri Poincaré Probab. Stat.* **31** 289–323. [MR1324810](https://doi.org/10.1324810)
- [31] MITZENMACHER, M. (2001). The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.* **12** 1094–1104.
- [32] MUKHERJEE, D., BORST, S. C., VAN LEEUWAARDEN, J. S. H. and WHITING, P. A. (2016). Universality of load balancing schemes on the diffusion scale. *J. Appl. Probab.* **53** 1111–1124. [MR3581245](https://doi.org/10.1017/jpr.2016.68) <https://doi.org/10.1017/jpr.2016.68>
- [33] MUKHERJEE, D., BORST, S. C., VAN LEEUWAARDEN, J. S. H. and WHITING, P. A. (2018). Universality of power-of- d load balancing in many-server systems. *Stoch. Syst.* **8** 265–292. [MR3899726](https://doi.org/10.1287/stsy.2018.0016) <https://doi.org/10.1287/stsy.2018.0016>
- [34] PUHALSKII, A. A. and VLADIMIROV, A. A. (2007). A large deviation principle for join the shortest queue. *Math. Oper. Res.* **32** 700–710. [MR2348243](https://doi.org/10.1287/moor.1070.0263) <https://doi.org/10.1287/moor.1070.0263>
- [35] RIDDER, A. and SHWARTZ, A. (2005). Large deviations without principle: Join the shortest queue. *Math. Methods Oper. Res.* **62** 467–483. [MR2229703](https://doi.org/10.1007/s00186-005-0037-1) <https://doi.org/10.1007/s00186-005-0037-1>
- [36] STOLYAR, A. L. (2015). Pull-based load distribution in large-scale heterogeneous service systems. *Queueing Syst.* **80** 341–361. [MR3367704](https://doi.org/10.1007/s11134-015-9448-8) <https://doi.org/10.1007/s11134-015-9448-8>
- [37] TIBI, D. (2010). Metastability in communication networks. arXiv preprint [arXiv:1002.0796](https://arxiv.org/abs/1002.0796).
- [38] VAN DER BOOR, M., BORST, S. C., VAN LEEUWAARDEN, J. S. H. and MUKHERJEE, D. (2018). Scalable load balancing in networked systems: Universality properties and stochastic coupling methods. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited Lectures* 3893–3923. World Sci. Publ., Hackensack, NJ. [MR3966556](https://doi.org/10.1007/s11134-015-9448-8)
- [39] VVEDENSKAYA, N. D., DOBRUSHIN, R. L. and KARPELEVICH, F. I. (1996). A queueing system with a choice of the shorter of two queues—an asymptotic approach. *Problemy Peredachi Informatsii* **32** 20–34. [MR1384927](https://doi.org/10.1384927)