

INVARIANCE PRINCIPLE AND MCKEAN–VLASOV LIMIT FOR RANDOMIZED LOAD BALANCING IN HEAVY TRAFFIC

BY RAMI ATAR^{1,a}  AND GERSHON WOLANSKY^{2,b}

¹Viterbi Faculty of Electrical and Computer Engineering, Technion—Israel Institute of Technology, ^arami@technion.ac.il

²Department of Mathematics, Technion—Israel Institute of Technology, ^bgershonw@technion.ac.il

We consider a load balancing model where a Poisson stream of jobs arrive at a system of many servers whose service time distribution possesses a finite second moment. A small fraction of arrivals pass through the power-of-choice algorithm, which assigns a job to the shortest among ℓ , $\ell \geq 2$, randomly chosen queues, and the remaining jobs are assigned to queues chosen uniformly at random. The system is analyzed at critical load in an asymptotic regime where both the number of servers and the usual heavy traffic parameter associated with individual queue lengths grow to infinity. The first main result is a hydrodynamic limit stating that the empirical measure of the diffusively normalized queue lengths converges to a path in measure space whose density is given by the unique solution of a parabolic PDE with nonlocal coefficients. The PDE has a stationary solution expressed explicitly, which in particular provides a quantification of the balance achieved by the algorithm. For fixed ℓ , as the intensity of the load balancing stream varies between its limits, this solution varies from exponential distribution to a Dirac distribution, demonstrating that the result covers the whole range from independence to state space collapse.

Further, two forms of an invariance principle are proved, one under a rather general initial distribution and the other under exchangeability of the initial conditions. In the latter, the limit of individual normalized queue length is given by a McKean–Vlasov SDE, and propagation of chaos holds. The McKean–Vlasov limit is closely related to limit results for Brownian particles on \mathbb{R}_+ interacting through their rank (with a specific interaction). However, an entirely different set of tools is required, as the collection of n prelimit particles does not obey a Markovian evolution on \mathbb{R}_+^n .

1. Introduction. This paper is concerned with a load balancing model where a Poisson stream of arrivals faces a system of n servers working in parallel, with a general common service time distribution having a finite second moment. Following a setting proposed in [6], motivated by maintaining low communication overhead, the stream undergoes thinning, where a small fraction of the arrivals are routed via the *join the shortest of ℓ queues* (abbreviated JSQ(ℓ)) algorithm, while the other arrivals undergo *uniform routing*. Here, JSQ(ℓ), also known as *power of choice*, is a randomized load balancing algorithm that assigns a job to the shortest among ℓ , $\ell \geq 2$, queues chosen uniformly at random, whereas uniform routing refers to assigning a job to a queue chosen uniformly at random.

The focus is on the asymptotics as the number of servers, n , becomes large, while individual queues fluctuate at the diffusion scale. This is achieved by taking n to also serve as the heavy traffic parameter of individual queue lengths and critically loading the system. More precisely, it is assumed that the per-server stream of arrivals resulting from the uniform routing is asymptotic to $\lambda n + \hat{\lambda} n^{1/2}$, whereas the processing rate of each server is asymptotic

Received February 2024; revised September 2024.

MSC2020 subject classifications. 68M20, 90B22, 60F17, 35K20.

Key words and phrases. Randomized load balancing, power of choice, diffusion limit, hydrodynamic limit, McKean–Vlasov limit, propagation of chaos, parabolic initial boundary value problem, viscous scalar conservation law, Daley–Miyazawa semimartingale representation.

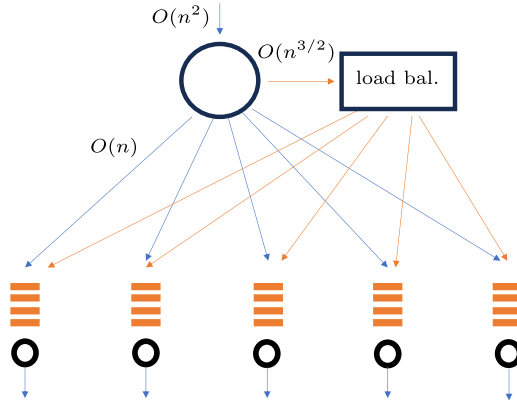


FIG. 1. Orders of magnitude of the various streams.

to $\lambda n + \hat{\mu}n^{1/2}$, where $\lambda > 0$, $\hat{\lambda}, \hat{\mu} \in \mathbb{R}$ are constants. Moreover, the overall intensity of the load balancing stream is taken asymptotic to $bn^{3/2}$, where $b > 0$ is a constant (see Figure 1). Note that the average load on a queue, defined as the rate of arrival of jobs per server divided by a server’s processing rate, is asymptotic to $1 + \rho n^{-1/2}$, where $\rho = \lambda^{-1}(\hat{\lambda} + b - \hat{\mu})$. In particular, $\rho < 0$ is a stability condition.

Let $X_i^n(t)$, $i \in [n] := \{1, \dots, n\}$ denote the i th queue length process at the n th system. Then the diffusion scaled queue length processes are given by $\hat{X}_i^n(t) = n^{-1/2}X_i^n(t)$. Let also $\bar{\xi}_t^n = n^{-1} \sum_{i \in [n]} \delta_{\hat{X}_i^n(t)}$ denote the empirical distribution of normalized queue lengths at time t , where δ_x is the Dirac measure at $x \in \mathbb{R}$. In particular, $\bar{\xi}_t^n[0, x]$ is the fraction of queues i for which $\hat{X}_i^n(t) \in [0, x]$. As argued informally below, taking the load balancing stream intensity to scale like $n^{3/2}$ is the only choice that gives interesting limits for $\bar{\xi}^n$.

The parameters b and ℓ determine the volume of communication between the servers and the dispatcher. They also have dramatic influence on the degree to which the system is load balanced. One of the goals of this work is to quantify the level of achieved load balancing, which could be defined, for example, by the expected empirical variance, namely

$$\begin{aligned}
 \sigma^n(t)^2 &:= \mathbb{E} \left[\frac{1}{n} \sum_{i \in [n]} \hat{X}_i^n(t)^2 - \left(\frac{1}{n} \sum_{i \in [n]} \hat{X}_i^n(t) \right)^2 \right] \\
 (1.1) \quad &= \mathbb{E} \left[\int_{\mathbb{R}_+} x^2 \bar{\xi}_t^n(dx) - \left(\int_{\mathbb{R}_+} x \bar{\xi}_t^n(dx) \right)^2 \right].
 \end{aligned}$$

Whereas it is difficult to estimate $\sigma^n(t)$ for each fixed n , suppose that the limit as $n \rightarrow \infty$ of $\bar{\xi}^n = \{\bar{\xi}_t^n\}$ exists and is given by a deterministic path $\xi = \{\xi_t\}$ in the space of measures on \mathbb{R}_+ . Suppose, moreover, that ξ_t has density $u(\cdot, t)$ for every t , that is, $\xi_t(dx) = u(x, t) dx$. This path could then be regarded as the macroscopic description of the model. Thus one could instead look at the substitute

$$(1.2) \quad \sigma_{\text{mac}}(t)^2 := \int_{\mathbb{R}_+} x^2 u(x, t) dx - \left(\int_{\mathbb{R}_+} x u(x, t) dx \right)^2$$

and obtain an approximation to $\sigma^n(t)$ for large n .

Assuming $\bar{\xi}_0^n \rightarrow \xi_0$ in probability, ξ_0 a Borel probability measure on \mathbb{R}_+ , our first main result (Theorem 2.7) addresses this by providing a hydrodynamic limit, which shows that $\bar{\xi}_t^n$ converges in probability, as a measure-valued process, to a deterministic measure-valued path given by $\xi_t(dx) = u(x, t) dx$, $t > 0$. Here, the density u is the unique classical solution of the

initial-boundary value problem

$$(1.3) \quad \begin{cases} u_t = -[(b\ell v^{\ell-1} - c_1)u]_x + au_{xx} & (x, t) \in (0, \infty)^2, \\ v(x, t) = \int_x^\infty u(y, t) dy & (x, t) \in (0, \infty)^2, \\ (c_1 - b\ell)u(0, t) + au_x(0, t) = 0 & t \in (0, \infty), \\ u(\cdot, t) dx \rightarrow \xi_0(dx) & \text{weakly as } t \downarrow 0. \end{cases}$$

Above, $c_1 = -\lambda\rho + b$ and $a > 0$ is a variance parameter. The first line of (1.3) contains a non-local drift coefficient that captures the algorithm’s instantaneous effect on the macroscopic distribution. Its third line is a boundary condition expressing zero flux through the origin (a Robin boundary condition; see, e.g., [37], p. 109), ensuring that the total mass is preserved. The last line is understood as weak convergence of probability measures. Our results also show that $\sigma^n(t) \rightarrow \sigma_{\text{mac}}(t)$ for each $t > 0$ (Proposition 2.8).

When $\rho < 0$, (1.3) has a unique stationary solution given by

$$(1.4) \quad \begin{aligned} v_{\text{stat}}(x) &= w(x)^{-1/(\ell-1)}, & u_{\text{stat}}(x) &= \frac{c_1}{a}(1 - \alpha)e^{\frac{c_1}{a}(\ell-1)x} w(x)^{-\ell/(\ell-1)}, \\ w(x) &= (1 - \alpha)e^{\frac{c_1}{a}(\ell-1)x} + \alpha, & \alpha &= \frac{b}{c_1} = \left(1 - \lambda \frac{\rho}{b}\right)^{-1} \in (0, 1), \end{aligned}$$

for $x \in \mathbb{R}_+$ (as stated below in Proposition 2.4). While an explicit expression for the stationary distribution is interesting by its own right, it also sheds light on the following aspect. If we fix (ℓ, ρ, λ, a) and let $b \rightarrow 0$ and, respectively, $b \rightarrow \infty$, then the measure $u_{\text{stat}}(x) dx$ converges weakly to the exponential distribution and to $\delta_0(dx)$ (details on this calculation are given in Section 2.3; graphs of u_{stat} for different values of ℓ and b are shown in Figure 3). These two limits indicate that our setting covers the entire range from queues operating independently to a *state space collapse*. The latter term is used in the heavy traffic literature to describe a coordination between the queues to the degree that the queue lengths, when normalized at the diffusion scale, are asymptotically indistinguishable. In addition, this argument confirms the heuristic that taking the load balancing intensity to scale like $n^{3/2}$ is the only choice for obtaining nontrivial limits of $\bar{\xi}^n$. (At the same time, it suggests that working with load balancing intensity of a higher order, the behavior of $\{X_i^n/\eta(n)\}_{i \in [n]}$, for a suitable choice of $1 \ll \eta(n) \ll n^{1/2}$ may exhibit interesting behavior; see [18] for estimates on the order of magnitude of X_i^n when all arrivals are routed via JSQ(ℓ)).

Although we prove that (1.4) is the unique stationary solution of the dynamics given by (1.3), let us emphasize that in this paper we do not prove that every solution to (1.3) converges to u_{stat} as $t \rightarrow \infty$, nor that the interchange of the t and n limits is valid. These important aspects of the model will be the subject of future work.

Our second goal is to provide an invariance principle. This is proved under two different assumptions about initial conditions. The first such result (Theorem 2.9) assumes that, for a fixed k , $(\hat{X}_i^n(0))_{i \in [k]} \Rightarrow (X_i(0))_{i \in [k]}$, and states that $(\hat{X}_i^n)_{i \in [k]} \Rightarrow (X_i)_{i \in [k]}$ where the latter is the solution to the SDE in \mathbb{R}_+^k

$$(1.5) \quad X_i(t) = X_i(0) + b_1t + b_0 \int_0^t v(X_i(s), s)^{\ell-1} ds + \sigma W_i(t) + L_i(t), \quad i \in [k],$$

in which L_i are boundary terms having continuous nondecreasing sample paths, satisfying the condition $\int_{[0, \infty)} X_i(t) dL_i(t) = 0$. Here, $b_0 = b\ell$, $b_1 = -c_1$ and $\sigma^2 = 2a$. Alternatively (Theorem 2.10), if the initial queue lengths are assumed exchangeable, then the limit in distribution of the k -tuple is given by k independent copies of

$$(1.6) \quad X(t) = X(0) + b_1t + b_0 \int_0^t v(X(s), s)^{\ell-1} ds + \sigma W(t) + L(t),$$

with $\int_{[0,\infty)} X(t) dL(t) = 0$. In both cases, the diffusion limits depend on the service time distribution only through its first two moments.

The second form of the invariance principle, where exchangeability is assumed, combined with the characterization of the limiting empirical distribution in terms of (1.3), shows that the pair (X, \bar{F}) , describing the limiting stochastic dynamics of a rescaled queue length and its law, satisfies

$$(1.7) \quad \begin{aligned} X(t) &= X(0) + \int_0^t \mathbf{b}(X(s), \bar{F}(\cdot, s)) ds + \sigma W(t) + L(t), \\ \bar{F}(x, t) &= \mathbb{P}(X(t) > x), \end{aligned}$$

where $\mathbf{b}(x, \bar{F}) = b_1 + b_0 \bar{F}(x)^{\ell-1}$. Viewed this way, the result is closely related to the literature on interacting diffusion models, and in particular to the subject of diffusions interacting through their ranks, where (1.7) is referred to as a McKean–Vlasov SDE. For background on McKean–Vlasov limits and propagation of chaos we refer to [21, 42].

To explain the relation, consider a parametric regime in which the heavy traffic parameter is taken to its limit first. In this case, each normalized queue length behaves as a Brownian particle on \mathbb{R}_+ , with an interaction among the particles caused by the load balancing algorithm. The algorithm effects the dynamics by selecting a particle with probability depending on its rank. In the original model, the selected queue increases by one job. In the particle system this can be imitated by imposing a positive rank-dependent drift. It therefore comes as no surprise that the McKean–Vlasov SDE (1.7) is a special case of the one that arises in the study of rank-dependent diffusions. Indeed, existence and uniqueness of solutions to (1.7), as well as convergence results, follow from work on interacting diffusions, specifically [29], Section 2.4.

Clearly, the above explanation is only heuristic. In the regime under consideration, the prelimit objects are queue lengths, which do not follow a Markovian evolution when considered as an n -tuple with state space \mathbb{R}_+^n . As the basic assumptions of the interacting diffusion setting do not hold, the treatment requires a completely different set of tools.

1.1. Motivation and related work. Motivation: Invariance principle. Invariance principles for queueing models in heavy traffic are diffusion limit results in which the limit processes depend on model ingredients such as service time and inter-arrival distributions only through their first two moments. Their significance stems from the fact that they express robustness of performance to underlying distributions. As far as load balancing algorithms in many-server settings are concerned, invariance principles have not been established before and in fact, diffusion limit results have not been proved beyond the case of exponential service times (for treatments of nonexponential service in other asymptotic regimes, see below). This is a rather restrictive assumption from a modeling viewpoint.

Motivation: Nonperfect balancing, sparse messaging. A large body of work on load balancing has been devoted to subcritically loaded systems, where on average queues are short. Critically loaded systems have also been studied, where either state space collapse or near optimal performance were proved; see references below. The question of quantifying the degree of balancing achieved by the algorithm, when balancing is far from perfect, does not arise in such settings, and has not been addressed before except under a fixed number of servers [6]. As already mentioned, this is one of the aspects of practical significance that motivates this work. In addition, the practical importance of load balancing algorithms that operate under low communication volume is widely acknowledged (see, e.g., [33]) and provides further motivation for our setting.

Related asymptotic settings. The literature on load balancing in asymptotic regimes is vast; see [43] for a recent survey. The present paper is not the first to study diffusion limits in a

regime in which the number of servers scales like n and individual queues are critically loaded, with each queue length scaling like $n^{1/2}$. This precise parameterization was considered before in [9], which studied a multi-agent game rather than a load balancing model.

Closer to the present paper is [18] which studies a system of n exponential servers where all arrivals are routed via $\text{JSQ}(\ell)$, under critical load. Rather than the empirical measure of rescaled queue lengths, [18] is concerned with a rescaling of the empirical measure of the (unscaled) queue lengths, focusing of the short queues, and identifying their dynamics under the limit.

The asymptotic regime where individual queues undergo heavy traffic scaling and $\text{JSQ}(\ell)$ is applied on a small fraction of the arrival stream was introduced in the aforementioned [6], in a setting where the number of servers is fixed. This could model two different scenarios, that are also relevant for our model. One is where a single stream is split by the dispatcher. Another is where each server has a dedicated stream due, for example, to geographical or compatibility constraints, and an additional stream is shared by all servers. Under an exponential service assumption, it was proved in [6] that the collection of rescaled queue lengths converges to a rank-based diffusion.

Asymptotics of $\text{JSQ}(\ell)$. Perhaps the most well-known load balancing algorithm is the *join-the-shortest-queue* (JSQ), which routes jobs to the shortest among all queues. Its variation $\text{JSQ}(\ell)$, that is typically applied with ℓ much smaller than n , significantly reduces the communication overhead while maintaining good performance. It was introduced in [45], where under the exponential service time assumption and subcritical load, the $n \rightarrow \infty$ limit of the length of a typical queue in stationarity was shown to have a doubly exponential tail decay, a dramatic improvement over routing uniformly at random in which case decay is exponential. Related results appeared in [34]. The result was extended in [24] where empirical measures were considered and propagation of chaos was obtained. Functional central limit theorems (CLT) were obtained in [25] and strong approximation results, including law of large numbers (LLN) and CLT, in [32]. (These CLT results were not invariance principles and were not concerned with the asymptotic regime studied in this paper). Results on the mixing rate and the size of the maximal queue length were obtained in [31], and aspects of stability and performance under server heterogeneity were investigated in [36]. The paper [35] studied how $\ell(n)$ should grow so that $\text{JSQ}(\ell(n))$ with exponential servers would perform like JSQ, and thus achieve asymptotic delay optimality.

A series of papers [11–13] extended [45] to several families of general service time distributions, based on an approach that first establishes propagation of chaos and then uses it to compute queue length distribution in equilibrium. Another line of research treating $\text{JSQ}(\ell)$ with general service times is [2, 3] and [1]. In [2, 3], the dynamical behavior was studied under general load and the hydrodynamic limit was shown to be given by an infinite system of coupled PDE. A numerical method was developed to solve these PDE. In [1], an infinite system of PDE was constructed and shown to constitute the invariant state of the aforementioned system of PDE under a subcriticality condition. Our setting, where individual queues undergo heavy traffic scaling, as well as our results, are quite different from both these lines of work.

Asymptotics of JSQ and other load balancing algorithms. JSQ is known to be delay optimal under exponential service and asymptotically delay optimal in various limiting regimes; see, for example, [15]. The diffusion limit in heavy traffic of JSQ in a many-server setting was established in [19], where, assuming exponential servers, a rescaled empirical measure of queue lengths was shown to converge to a process expressed in terms of a 2d diffusion. Convergence of the steady state at the same scale was proved in [14], and properties of the diffusion process were investigated in [7, 8]. Under the regime considered in this line of work, the number of queues that are of length 0, 1 and 2 is of order $n^{1/2}$, n and, respectively, $n^{1/2}$,

and only a negligible fraction exceeds length 2. Thus this regime captures quite a different behavior from what is described in this paper. The papers [26] and [48] considered JSQ in other diffusive regimes, namely the so called nondegenerate slowdown regime and, respectively, the super-Halfin–Whitt regime.

Numerous load balancing algorithms besides JSQ and JSQ(ℓ) have been proposed. Ones that emphasize sparse communication, where the messaging rate often goes far below that of JSQ(ℓ), include [4, 20, 44, 47]. Further asymptotic results on load balancing with nonexponential service times include join-the-idle-queue [20], pull-based load distribution [40, 41], zero waiting algorithms [30], and join-the-shortest-estimated-queue [5].

Rank-based diffusions. McKean–Vlasov limits of diffusions interacting through their ranks were addressed in [39]. Observing that in these models the dependence of each particle’s coefficients on the empirical law is discontinuous, a situation not covered by the classical treatment such as [21], this paper gave convergence results assuming that the coefficients are merely measurable. The paper [29] mentioned above addresses a setting of which rank-based interaction is a special case, and attains convergence results in a stronger topology than weak convergence.

1.2. *Notation.* Denote $\mathbb{R}_+ = [0, \infty)$. Let $\iota : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denote the identity map. In \mathbb{R}^N , denote the Euclidean norm by $\|\cdot\|$. For $(\mathbf{X}, d_{\mathbf{X}})$ a Polish space, let $C(\mathbb{R}_+, \mathbf{X})$ and $D(\mathbb{R}_+, \mathbf{X})$ denote the space of continuous and, respectively, càdlàg paths, endowed with the topology of uniform convergence on compacts and, respectively, the J_1 topology. Denote by \mathcal{M}_1 the space of probability measures on \mathbb{R}_+ equipped with the topology of weak convergence. Denote by C^\uparrow the set of members of $C(\mathbb{R}_+, \mathbb{R}_+)$ that are nondecreasing and start at 0. For $\xi \in D(\mathbb{R}_+, \mathbb{R}^N)$, an interval $I \subset \mathbb{R}_+$, and $0 \leq \delta \leq T$, denote

$$\begin{aligned} \Delta\xi(t) &= \xi(t) - \xi(t-), \quad t > 0, \quad \Delta\xi(0) = \xi(0), \\ \text{osc}(\xi, I) &= \sup\{\|\xi(s) - \xi(t)\| : s, t \in I\}, \\ w_T(\xi, \delta) &= \sup\{\|\xi(t) - \xi(s)\| : s, t \in [0, T], |s - t| \leq \delta\}, \\ \|\xi\|_T^* &= \sup\{\|\xi(t)\| : t \in [0, T]\}, \end{aligned}$$

and by $|\xi|(t)$ the total variation of ξ in $[0, t]$.

For $\mathcal{D} \subset \mathbb{R}^N$, denote by \mathcal{D}° and $\bar{\mathcal{D}}$ its interior and, respectively, closure. Denote by $C(\mathcal{D})$ the set of continuous functions $f : \mathcal{D} \rightarrow \mathbb{R}$ and by $C_0(\mathcal{D})$ the set of $f \in C(\mathcal{D})$ whose support is a compact subset of \mathbb{R}^N . For $k, l \in \mathbb{N}$ and an open set $\mathcal{D} \subset \mathbb{R}$ (respectively, $\mathcal{D} \subset \mathbb{R} \times \mathbb{R}_+$), denote by $C^k(\mathcal{D})$ (respectively, $C^{k,l}(\mathcal{D})$) the set of functions $f : \mathcal{D} \rightarrow \mathbb{R}$ possessing continuous derivatives up to and including k (respectively, (k, l)). For $\mathcal{D} \subset \mathbb{R}$, let $C^k(\mathcal{D})$ denote the set of functions $f : \mathcal{D} \rightarrow \mathbb{R}$ whose restriction to \mathcal{D}° lies in $C^k(\mathcal{D}^\circ)$, and whose derivatives up to and including k have continuous extensions to \mathcal{D} . Define $C^{k,l}(\mathcal{D})$ analogously. For $\mathcal{D} \in \mathbb{R}$ and $\mathcal{D} \in \mathbb{R} \times \mathbb{R}_+$ denote $C^\infty(\mathcal{D}) = \bigcap_{k \in \mathbb{N}} C^k(\mathcal{D})$ and $C^\infty(\mathcal{D}) = \bigcap_{k,l \in \mathbb{N}} C^{k,l}(\mathcal{D})$, respectively. A subscript 0 denotes compact support, for example, $C_0^k(\mathcal{D}) = C^k(\mathcal{D}) \cap C_0(\mathcal{D})$. A subscript b denotes boundedness, for example, $C_b^k(\mathcal{D})$ are functions in $C^k(\mathcal{D})$ whose derivatives of order $0 \leq l \leq k$ are bounded. For f defined on (a subset of) \mathbb{R} , f' denotes derivative, and for f defined on (a subset of) $\mathbb{R} \times \mathbb{R}_+$, f_x and f_t denote spatial and temporal derivatives, respectively.

For an interval $I \subset \mathbb{R}_+$ and a normed space S , denote by $\mathbb{L}^p(I; S)$, $1 \leq p \leq \infty$, the usual \mathbb{L}^p space defined in terms of the Lebesgue measure on I . Denote by $\mathbb{L}_{\text{loc}}^p(\mathbb{R}_+; S)$ the set of $f : \mathbb{R}_+ \rightarrow S$ such that $f \in \mathbb{L}^p([0, T]; S)$ for all finite T . $\mathbb{L}^p(\mathbb{R}_+; \mathbb{R})$ is abbreviated to $\mathbb{L}^p(\mathbb{R}_+)$.

For $f, g : \mathbb{R}_+ \rightarrow \mathbb{R}$ and a measure m on \mathbb{R}_+ , $\langle f, m \rangle = \int_{\mathbb{R}_+} f(x)m(dx)$ and $\langle f, g \rangle = \int_{\mathbb{R}_+} f(x)g(x)dx$.

If $V \in \mathbb{R}^n$ then $V_i, i \in [n]$ denote its components in the standard basis, and vice versa: Given $V_i, i \in [n]$, V denotes the vector (V_1, V_2, \dots, V_n) . Both these conventions hold also for random variables V_i and processes $V_i(\cdot)$. Occasionally, with a slight abuse of standard terminology, a sequence of random elements (random variables or processes) is referred to as *tight* when their laws form a tight sequence of probability measures. The term *with high probability* (w.h.p.) means “holds, for each n , on Ω_n , where $\lim_n \mathbb{P}(\Omega_n) = 1$ ”. Convergence in distribution is denoted by \Rightarrow . c denotes a positive constant whose value may change from one expression to another.

Throughout the paper, superscript n attached to scalars, random variables or processes, denotes dependence on the index n rather than power.

1.3. Paper organization. The load balancing model is presented in Section 2.1. In §2.2, equation (2.16) is introduced, which is a PDE closely related to (1.3), and the main results are stated. A discussion of the results appears in Section 2.3. An outline of the proof is given in Section 2.4. Section 3 is devoted to developing several tools required for the proof. The first tool, provided in Section 3.1, is a crucial PDE uniqueness result for weak solutions of (2.16). Next, in [17], a semimartingale representation for counting processes was introduced, that we find extremely useful. Section 3.2 presents this representation and, based on it, identifies various martingales that are of importance to the model. These martingales are key to the proof of convergence provided later in Section 4. Section 3.3 provides various estimates on the rescaled queueing processes. Tightness of the empirical process $\bar{\xi}^n$ is shown in Section 3.4. With this set of tools, the proof is then carried out in Section 4, which starts by writing down an equation for $\langle \phi, \bar{\xi}^n \rangle$, for ϕ a test function, namely equation (4.4). Thanks to the uniqueness of solutions to (2.16) and the tightness of $\bar{\xi}^n$, the hydrodynamic can be established by showing that subsequential limits of $\bar{\xi}^n$ form weak solutions to this PDE. This is achieved by relating limits of the various terms in (4.4) to weak solutions of (2.16), performed in three steps in Section 4.1, Section 4.2 and Section 4.3 (a detailed description of these steps appears at the end of Section 2.4). Finally, all the main results are proved, based on the above, in Section 4.4.

2. Load balancing in heavy traffic.

2.1. The load balancing model.

2.1.1. Arrivals, queue lengths and basic relations. In the model there are n servers and a queue in front of each. There is a dedicated stream of arrivals into each queue and an additional stream of arrivals, called the *load balancing stream* (LBS), that go through the JSQ(ℓ, n) algorithm. These $n + 1$ arrival streams are modeled as mutually independent Poisson processes. Clearly, this could be recast as a single Poisson stream, out of which $n + 1$ thinned streams are created by means of random selection.

In what follows, the queueing systems will be indexed by $n \in \mathbb{N}$, the number of servers. The processes X_i^n, E_i^n, D_i^n and T_i^n represent the i th queue length process, dedicated arrival process, departure process and busyness process, respectively. Denote by A_0^n the LBS arrival process, and by A_i^n the process counting LBS arrivals routed to server i . For each i , E_i^n is a Poisson process of parameter λ^n , and A_0^n is Poisson of parameter λ_0^n , all having right-continuous sample paths. The $n + 1$ Poisson processes are mutually independent, for each n . We have

$$(2.1) \quad X_i^n(t) = X_i^n(0-) + E_i^n(t) + A_i^n(t) - D_i^n(t), \quad i \in [n], t \in \mathbb{R}_+.$$

Work conservation is assumed, hence

$$(2.2) \quad T_i^n(t) = \int_0^t \mathbf{1}_{\{X_i^n(s) > 0\}} ds.$$

2.1.2. *The load balancing algorithm.* Upon each LBS arrival, ℓ out of the n queues, chosen uniformly at random, are sampled. The arrival is routed to the queue that is shortest among the ℓ ; if there are ties, the queue with the smaller index is preferred. Given $x \in \mathbb{R}^n$, let

$$(2.3) \quad \text{rank}(i; x) = \#\{j : x_j < x_i\} + \#\{j \leq i : x_j = x_i\}.$$

Then the probability that a LBS arrival is routed to the queue whose rank, defined by (2.3), is r is given by

$$(2.4) \quad p_{n,r} = \frac{\binom{n-r}{\ell-1}}{\binom{n}{\ell}}, \quad r \in [n],$$

with $\binom{k}{j} = 0$ when $j > k$. Note that

$$(2.5) \quad p_{n,r} = 0 \quad \text{for } r \geq n - \ell + 2, \quad \max_{r \in [n]} p_{n,r} = p_{n,1} = \frac{\ell}{n}.$$

Thus the randomization mechanism can equivalently be achieved by letting, for each n , θ_k^n , $k \in \mathbb{N}$ be i.i.d. random variables with $\mathbb{P}(\theta_1^n = r) = p_{n,r}$, $r \in [n]$, and routing the k th LBS arrival to the queue whose rank is θ_k^n . That is,

$$(2.6) \quad \begin{aligned} A_i^n(t) &= \int_{[0,t]} 1_{\{\mathcal{R}_i^n(s-) = \theta_{A_0^n(s)}^n\}} dA_0^n(s), \\ \mathcal{R}_i^n(t) &= \text{rank}(i; X^n(t)), \quad i \in [n], t \in \mathbb{R}_+. \end{aligned}$$

Of course, this randomization scheme is not used in practice as it requires keeping track of the queue lengths of the entire system. However, it is convenient for carrying out the analysis.

REMARK 2.1. Although we work with the expression (2.4) that corresponds to sampling without replacement, all our results are valid for sampling with replacement as well. This is stated and explained in Remark 4.5.

2.1.3. *The initial condition.* We allow quite a general initial condition, where residual times of jobs already being processed at time 0 may be dependent and have unspecified distributions. Denote the (random) set of queues that at time 0 contain no jobs and, respectively, at least one job, by $\mathcal{N}^n = \{i \in [n] : X_i^n(0) = 0\}$ and $\mathcal{P}^n = \{i \in [n] : X_i^n(0) > 0\}$. For $i \in \mathcal{P}^n$, let $Z_i^n(0)$ denote the initial residual time of the head-of-the-line job in queue i . For $i \in \mathcal{N}^n$ we add fictitious jobs having zero processing time. To this end, rather than specifying $X^n(0)$ as the initial queue length, $X^n(0-)$ is specified; and for each $i \in \mathcal{N}^n$, the queue length is set to $X_i^n(0-) = 1$ and the residual processing time to $Z_i^n(0) = 0$. Obviously, this results in $X_i^n(0) = 0$. This convention allows us to greatly simplify notation when we later construct counting processes for service and departure. The initial condition is thus a tuple

$$\mathcal{I}^n = (\{X_i^n(0-), Z_i^n(0), i \in [n]\}, \mathcal{N}^n, \mathcal{P}^n),$$

where $(\mathcal{N}^n, \mathcal{P}^n)$ partitions $[n]$, and

$$\begin{aligned} X_i^n(0-) &= 1, & Z_i^n(0) &= 0, & i \in \mathcal{N}^n, \\ X_i^n(0-) &\geq 1, & Z_i^n(0) &> 0, & i \in \mathcal{P}^n. \end{aligned}$$

2.1.4. *Service times.* Let Φ_{ser} be a Borel probability measure on $[0, \infty)$ with mean 1 and standard deviation $\sigma_{\text{ser}} \in (0, \infty)$, such that $\Phi_{\text{ser}}(\{0\}) = 0$. Let Φ_{ser}^n be defined as a scaled version of Φ_{ser} uniquely specified via $\Phi_{\text{ser}}^n[0, x] = \Phi_{\text{ser}}[0, \mu^n x]$, $x \in \mathbb{R}_+$. Here, $\mu^n > 0$ is the service rate in the n th system. For $k \geq 1$, let $Z_i^n(k)$ denote the service time of the k th job to be served by server i after the head-of-the-line job at time 0— there (for $i \in \mathcal{N}^n$ this means the k th job after the fictitious one). For every i , $Z_i^n = (Z_i^n(k), k \geq 1)$ is an i.i.d. sequence with common distribution Φ_{ser}^n .

Next, the potential service process S_i^n , evaluated at t , gives the number of jobs completed by server i by the time it has worked t units of time. It is given, with $\sum_0^{-1} = 0$, by

$$S_i^n(t) = \max \left\{ k \in \mathbb{Z}_+ : \sum_{j=0}^{k-1} Z_i^n(j) \leq t \right\}, \quad t \geq 0.$$

The departure processes are given by $D_i^n(t) = S_i^n(T_i^n(t))$. This is the number of jobs completed by time t by server i . Note that the first departure counted by D_i^n is the one initially processed if $i \in \mathcal{P}^n$ and the fictitious job if $i \in \mathcal{N}^n$.

2.1.5. *Dependence structure.* For each n , the $2n + 3$ stochastic elements

$$(2.7) \quad E_i^n, i \in [n], Z_i^n, i \in [n], \mathcal{I}^n, A_0^n, \text{ and } \{\theta_k^n\} \text{ are mutually independent.}$$

2.1.6. *Scaling and critical load condition.* The arrival and service rates are assumed to satisfy the following. There are constants $\lambda > 0$ and $\hat{\lambda} \in \mathbb{R}$ such that

$$(2.8) \quad \hat{\lambda}^n := n^{-1/2}(\lambda^n - n\lambda) \rightarrow \hat{\lambda} \quad \text{as } n \rightarrow \infty,$$

a constant $b > 0$ such that

$$(2.9) \quad \hat{\lambda}_0^n := n^{-3/2}\lambda_0^n \rightarrow b \quad \text{as } n \rightarrow \infty,$$

and constants $\mu > 0$ and $\hat{\mu} \in \mathbb{R}$ such that

$$(2.10) \quad \hat{\mu}^n := n^{-1/2}(\mu^n - n\mu) \rightarrow \hat{\mu} \quad \text{as } n \rightarrow \infty.$$

The critical load condition is assumed, namely

$$(2.11) \quad \lambda = \mu.$$

Some further notation used throughout is

$$\begin{aligned} \hat{b}_1^n &= \hat{\lambda}^n - \hat{\mu}^n, & b_1 &= \hat{\lambda} - \hat{\mu}, & c_1 &= -b_1, & b_0 &= b\ell, \\ \sigma^2 &= \lambda(1 + \sigma_{\text{ser}}^2), & a &= \frac{\sigma^2}{2}. \end{aligned}$$

Note that the average load on a queue, defined as the rate of arrival of work per server, $\lambda^n + n^{-1}\lambda_0^n$, divided by a server’s processing rate, μ^n , that is,

$$\frac{n\lambda + n^{1/2}\hat{\lambda} + n^{1/2}b + o(n^{1/2})}{n\mu + n^{1/2}\hat{\mu} + o(n^{1/2})},$$

is asymptotic to

$$1 + n^{-1/2}\rho, \quad \rho := \frac{\hat{\lambda} + b - \hat{\mu}}{\lambda}.$$

Hence the *load parameter* ρ is given by $\rho = (b_1 + b)/\lambda$.

Let rescaled versions of the queue length processes and cumulative idle time processes be defined by

$$(2.12) \quad \hat{X}_i^n(t) = n^{-1/2} X_i^n(t), \quad \hat{L}_i^n(t) = n^{-1/2} \mu^n(t - T_i^n(t)),$$

and denote $\bar{\xi}_i^n = n^{-1} \sum_{i \in [n]} \delta_{\hat{X}_i^n(t)}$. Further, assume that $\bar{\xi}_{0-}^n \rightarrow \xi_0$ (equivalently $\bar{\xi}_0^n \rightarrow \xi_0$) in probability, as $n \rightarrow \infty$, where ξ_0 is a deterministic Borel probability measure on \mathbb{R}_+ . For the rescaled residual times $Z_i^n(0)$, assume that for every $\varepsilon > 0$,

$$(2.13) \quad \lim_{n \rightarrow \infty} \max_{i \in [n]} \mathbb{P}(\tilde{Z}_i^n(0) > \varepsilon) = 0 \quad \text{where } \tilde{Z}_i^n(0) := \tilde{\mu}^n Z_i^n(0), \tilde{\mu}^n := n^{-1/2} \mu^n,$$

and

$$(2.14) \quad \sup_n \max_{i \in [n]} \mathbb{E}[\tilde{Z}_i^n(0)^2] < \infty.$$

Finally, assume that

$$(2.15) \quad \sup_n \max_{i \in [n]} \mathbb{E}[\hat{X}_i^n(0-)^2] < \infty.$$

Note that, by Fatou’s lemma, this imposes a condition on ξ_0 , namely ξ_0 necessarily satisfies $\int x^2 \xi_0(dx) < \infty$.

All assumptions made thus far are assumed to hold throughout the remainder of the article, without further mention.

REMARK 2.2. Because $\tilde{\mu}^n$ scales like $n^{1/2}$, conditions (2.13) and (2.14) allow for $Z_i^n(0)$ to be $o(n^{-1/2})$, a larger scale than $O(n^{-1})$ assumed for $Z_i^n(k), k \geq 1$. Thus the assumptions we have on $Z_i^n(0)$ are much weaker than those on $Z_i^n(k), k \geq 1$. Of course, the case where $Z_i^n(0)$ do scale like $Z_i^n(k), k \geq 1$, is covered by these assumptions. For example, if, for all n and i , $Z_i^n(0)$ has the same distribution as $Z_1^n(1)$, then

$$\max_{i \in [n]} \mathbb{P}(\tilde{Z}_i^n(0) > \varepsilon) = \mathbb{P}(\tilde{\mu}^n Z_1^n(1) > \varepsilon) = \mathbb{P}(n^{-1/2} \mu^n Z_1^n(1) > \varepsilon) = \Phi_{\text{ser}}[\varepsilon n^{1/2}, \infty),$$

since $\mu^n Z_1^n(1) \sim \Phi_{\text{ser}}$. Clearly, this shows that (2.13) holds in this case. It can similarly be seen that (2.14) holds.

2.2. *Main results.* First we address well-posedness of the relevant PDE, starting with classical solutions to (1.3).

THEOREM 2.3. *Within the class of functions $u \in C^{2,1}((0, \infty) \times \mathbb{R}_+; \mathbb{R}_+)$ satisfying, for every $T < \infty$, $\sup_{t \in (0, T]} \int_{\mathbb{R}_+} x u(x, t) dx < \infty$, there exists a unique solution to equation (1.3). Moreover, for each $t > 0$, $u(\cdot, t)$ is a probability density.*

A function $u_{\text{stat}} \in C^2(\mathbb{R}_+)$ is said to be a *stationary solution associated with (1.3)*, if it is a probability density possessing a finite second moment, and, setting $\xi_0(dx) = u_{\text{stat}}(x) dx$, the solution to (1.3) is given by $u(x, t) = u_{\text{stat}}(x)$ for all x, t .

PROPOSITION 2.4. *Assume $\rho < 0$. Then there exists a unique stationary solution to (1.3). It is given by (1.4).*

Theorem 2.3 is a consequence of a result stated next, that gives uniqueness of weak solutions to a viscous scalar conservation law, also referred to as a Burgers-type equation, and provides one of the main tools used in this paper. This is the parabolic equation

$$(2.16) \quad \begin{cases} v_t = (\mathfrak{f}(v))_x + a v_{xx}, & (x, t) \in \mathbb{R}_+^2, \\ v(0, t) = 1, & t > 0, \\ v(\cdot, 0) = v_0. \end{cases}$$

A key point is that this result does not require any regularity assumption in the x variable on the class of solutions (beside integrability). The definition of a weak solution is as follows.

DEFINITION 2.5. A function $v \in L^\infty_{\text{loc}}(\mathbb{R}_+; L^\infty(\mathbb{R}_+)) \cap L^\infty_{\text{loc}}(\mathbb{R}_+; L^1(\mathbb{R}_+))$ is a weak solution of (2.16) if for any $t \in (0, \infty)$ and any $\phi \in C^\infty_0(\mathbb{R}_+)$ satisfying $\phi(0) = 0$,

$$(2.17) \quad \langle v(\cdot, t), \phi \rangle - \langle v_0, \phi \rangle = - \int_0^t \langle f(v(\cdot, s)), \phi' \rangle ds + a \int_0^t \langle v(\cdot, s), \phi'' \rangle ds + a\phi'(0)t.$$

The choice of f that will be of interest is $f(z) = c_1z - bz^\ell$. In particular, (1.3) is related (2.16) when the latter takes the special form

$$(2.18) \quad \begin{cases} v_t = (c_1v - bv^\ell)_x + av_{xx}, & (x, t) \in \mathbb{R}_+^2, \\ v(0, t) = 1, & t > 0, \\ v(\cdot, 0) = \xi_0(\cdot, \infty). \end{cases}$$

The existence and uniqueness of a classical solution of (2.16) is well known [23], Lemma 3.2, p. 68, and there is a vast literature on weak solutions in the $W^{1,p}_{\text{loc}}$ sense. However, we found no uniqueness result of a weak solution as given in Definition 2.5 under the mere assumption $v \in L^\infty_{\text{loc}}(\mathbb{R}_+; L^\infty(\mathbb{R}_+)) \cap L^\infty_{\text{loc}}(\mathbb{R}_+; L^1(\mathbb{R}_+))$.

THEOREM 2.6. Assume $f \in C^\infty(\mathbb{R})$ and $v_0 \in L^1(\mathbb{R}_+) \cap L^\infty(\mathbb{R}_+)$. Then there exists a unique weak solution to (2.16) (in particular, to (2.18)) in the sense of Definition 2.5. This solution is in $C^\infty(\mathbb{R}_+ \times (0, \infty))$. If v denotes the weak solution to (2.18) then $u = -v_x$ is the classical solution of (1.3).

The following is our first main result, showing that (1.3) provides a macroscopic description of the model, specifically its hydrodynamic limit. Recall that, in the load balancing model, ξ_0 stands for the limit in probability of $\bar{\xi}_0^n$.

THEOREM 2.7. Let ξ_0 be extended to a trajectory $\xi = \{\xi_t, t \in \mathbb{R}_+\}$ in \mathcal{M}_1 by setting

$$\xi_t(dx) = u(x, t) dx, \quad t > 0,$$

where u is the unique solution to (1.3). Then $\xi \in C(\mathbb{R}_+, \mathcal{M}_1)$, and one has $\bar{\xi}_0^n \rightarrow \xi$ in probability in $D(\mathbb{R}_+, \mathcal{M}_1)$, as $n \rightarrow \infty$.

Under a slightly stronger moment condition, the above result implies

PROPOSITION 2.8. Assume that, for some $\varepsilon > 0$, Φ_{ser} possesses a finite $2 + \varepsilon$ moment and that $\sup_n \max_{i \in [n]} \{\mathbb{E}[\tilde{Z}_i^n(0)^{2+\varepsilon}] \vee \mathbb{E}[\hat{X}_i^n(0-)^{2+\varepsilon}]\} < \infty$ (compare with (2.14) and (2.15)). Let σ_{mac} from (1.2) be defined with u of (1.3). Then $\sigma^n(t) \rightarrow \sigma_{\text{mac}}(t)$ for every $t > 0$.

Our second main result is the conjunction of the following two invariance principles, obtained under two different assumptions on the initial conditions.

THEOREM 2.9. Fix $k \in \mathbb{N}$ and assume that $(\hat{X}_i^n(0))_{i \in [k]} \Rightarrow (X_i(0))_{i \in [k]}$. Let v be the weak solution to (2.18). Then $(\hat{X}_i^n, \hat{L}_i^n)_{i \in [k]} \Rightarrow (X_i, L_i)_{i \in [k]}$ in $(D(\mathbb{R}_+, \mathbb{R}_+) \times C(\mathbb{R}_+, \mathbb{R}_+))^k$, where the latter tuple is the unique in law solution to the system (1.5), driven by a k -dimensional standard BM $(W_i)_{i \in [n]}$ independent of $(X_i(0))_{i \in [k]}$.

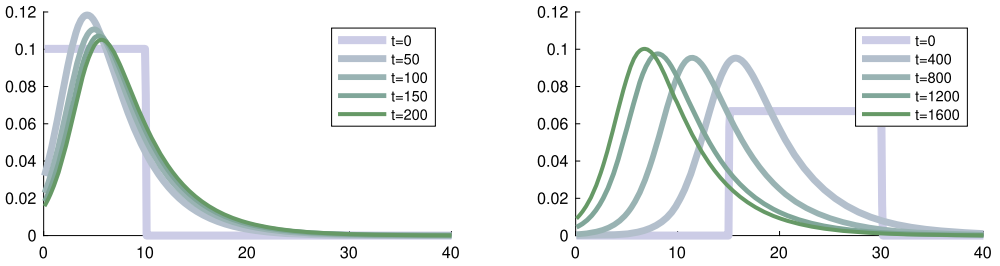


FIG. 2. Solution of (1.3) for initial condition $\text{unif}[0, 10]$ (left) and $\text{unif}[15, 30]$ (right) at different time instances, with $\rho = -0.01, c_1 = 0.21, b = 0.2, \ell = 4, a = 1$.

THEOREM 2.10. Assume that, for each n , the random variables $X_i^n(0-), i \in [n]$, are exchangeable (necessitating that the limit law of each $X_i^n(0-)$ is ξ_0). Let v be the weak solution to (2.18). Then for $k \in \mathbb{N}, (\hat{X}_i^n, \hat{L}_i^n)_{i \in [k]} \Rightarrow (X_i, L_i)_{i \in [k]}$ in $(D(\mathbb{R}_+, \mathbb{R}_+) \times C(\mathbb{R}_+, \mathbb{R}_+))^k$, where the latter tuple is given by k independent copies of the unique in law solution (X, L) to (1.6) driven by a standard BM W independent of $X(0)$, and the latter is distributed according to ξ_0 .

In view of Theorem 2.7, the SDE in Theorem 2.10 is a McKean–Vlasov SDE, as it can be seen to take the form (1.7). Note that we do not assume that the entire initial data is exchangeable, nor that the tie breaking rule is symmetric w.r.t. exchanging i 's. The result implies that these issues have negligible effect upon taking the limit.

2.3. Discussion. The hydrodynamic limit result establishes that the dynamics of the model are given at the macroscopic scale by the PDE (1.3). As argued in the Introduction, one can use this to express basic performance criteria in terms of the solution u . In particular, denoting the macroscopic mean by

$$m_{\text{mac}}(t) := \int_0^\infty x u(x, t) dx,$$

a first order approximation to the mean queue length is given by $m_{\text{mac}}(t)n^{1/2}$. Accordingly, the mean delay is, to first order, given by $\lambda^{-1}m_{\text{mac}}(t)n^{-1/2}$. Similarly, the macroscopic standard deviation, $\sigma_{\text{mac}}(t)$ defined in (1.2), may be taken as an *index of balance*. As a second order parabolic equation in one spatial variable, there exist standard tools for numerically solving (1.3) and evaluating these performance criteria. Some examples of solutions are plotted in Figure 2.

Whereas the macroscopic dynamics can be solved numerically, the macroscopic equilibrium state has an explicit formula, namely (1.4). As already mentioned, it is beyond the scope of this paper to show rigorously that all solutions u converge to u_{stat} as $t \rightarrow \infty$, and further, that the invariant distribution of the stochastic model's dynamics converges to u_{stat} as $n \rightarrow \infty$. However, even without establishing these results, u_{stat} constitutes a legitimate solution to the PDE. In particular, a combination of Theorem 2.7 and Proposition 2.4 implies that $\bar{\xi}^n \rightarrow \xi$ in probability holds provided that $\xi_0(dx) = u_{\text{stat}}(x) dx$. Here, $\xi_t = \xi_0$ for all t . Figure 3 shows graphs of u_{stat} for different values of ℓ and b , based on (1.4).

We can also use u_{stat} in place of $u(\cdot, t)$ in the above macroscopic performance criteria, and define analogously m_{mac} and σ_{mac} which are now time independent quantities. Figure 4 shows their dependence on ℓ and b .

Let us provide details on the statement made in the Introduction regarding the limits, as $b \rightarrow 0$ and $b \rightarrow \infty$, of the stationary distribution. Fix (ℓ, ρ, λ, a) , with $\rho < 0$. According

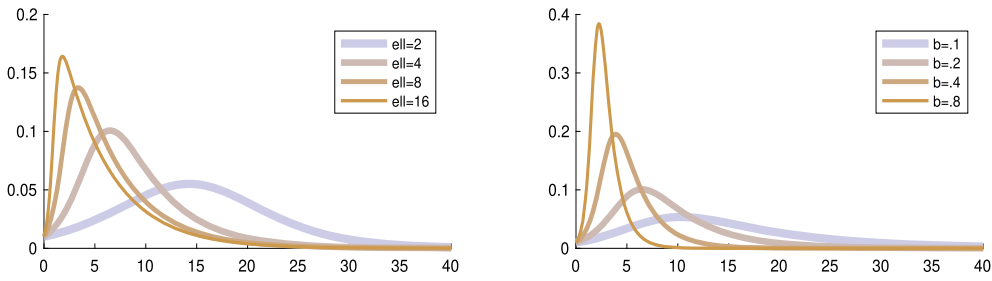


FIG. 3. Stationary solution for different values of ℓ , with $b = 0.2$ (left), and for different values of b , with $\ell = 4$ (right). In all cases, $\rho = -0.01, a = 1, \lambda = 1$.

to the formula in (1.4), as $b \rightarrow 0$, one has $\alpha \rightarrow 0$ and $c_1 \rightarrow -\lambda\rho = \lambda|\rho|$. Hence one has $w(x) \rightarrow e^{\frac{\lambda|\rho|}{a}(\ell-1)x}$ pointwise, and consequently,

$$\lim_{b \rightarrow 0} v_{\text{stat}}(x) = e^{-\frac{\lambda|\rho|}{a}x}, \quad x \geq 0.$$

As $b \rightarrow \infty$, one has $c_1 \rightarrow \infty$ and, although $\alpha \rightarrow 1$, it is obvious that $w(x) \rightarrow \infty$ for every $x > 0$. This gives

$$\lim_{b \rightarrow \infty} v_{\text{stat}}(x) = 1_{\{x=0\}}, \quad x \geq 0.$$

In both cases, the limit holds at every continuity point of the limiting CDF, proving that $v_{\text{stat}}(dx) = u_{\text{stat}}(x) dx \rightarrow \frac{\lambda|\rho|}{a}e^{-\frac{\lambda|\rho|}{a}x} dx$ weakly as $b \rightarrow 0$, and $u_{\text{stat}}(x) dx \rightarrow \delta_0(dx)$ weakly as $b \rightarrow \infty$. The former is an exponential distribution, which agrees with the limiting distribution of a single server queue in heavy traffic. This makes perfect sense since in this limit the queues operate independently. The latter is a distribution that has all its mass at 0, showing that, by choosing b large, the hydrodynamic limit can be made arbitrarily close to a state space collapse, where the diffusion-scale processes \hat{X}_i^n are asymptotically indistinguishable.

We can also provide some details on the $b \rightarrow 0$ and $b \rightarrow \infty$ limits of solutions to the parabolic equation (1.3), although here the discussion is heuristic. Again, fix ℓ and ρ (not necessarily negative). Then $c_1 = -\lambda\rho + b$, and as $b \rightarrow 0$, equation (1.3) converges formally to the linear heat equation $u_t = -\lambda\rho u_x + au_{xx}$ with $-\lambda\rho u(0, t) + au_x(0, t) = 0$. To describe the $b \rightarrow \infty$ limit, observe that the drift coefficient in the SDE (1.5) is given by

$$b_1 + b_0v^{\ell-1} = \lambda\rho + (-1 + \ell v(X_i(t), t))^{\ell-1}b.$$

When $-1 + \ell v(X_i(t), t)^{\ell-1} < -\varepsilon$ (respectively, $> \varepsilon$), the drift tends to $-\infty$ (respectively, ∞), which suggests that the particles concentrate near the point x determined by the equation $-1 + \ell v(x, t)^{\ell-1} = 0$. Thus one expects that $u(x, t) dx \rightarrow \delta_{\gamma(t)}(dx)$ where u is the solution

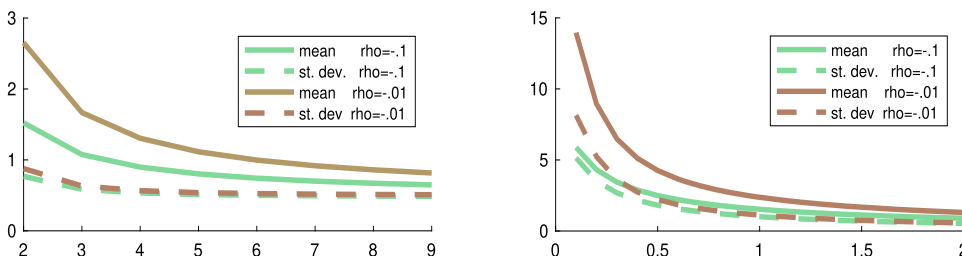


FIG. 4. Stationary mean and standard deviation for $2 \leq \ell \leq 9$, with $b = 2, a = 1$ (left) and for $0.1 \leq b \leq 2$, with $\ell = 4, a = 1$ (right).

to (1.3) and γ is a curve determined by a parabolic free boundary problem. This will be explored in future work.

The nature of the drift in equation (1.3) is akin to the dynamics of a mass distribution subjected to a gravitational field [46], where the force acting on an infinitesimal mass depends on the amount of mass that piles above it. Similarly, in (1.3), the force is determined by a nonlinear function of $v(x, t)$.

2.4. *Outline of the proof.* The approach used to proving the hydrodynamic limit result is based on PDE uniqueness. Thus the first main tool required is the uniqueness of weak solutions to (2.16). The crux of the argument is as follows. If v_1, v_2 are two weak solutions then one can show that $w = v_1 - v_2$ satisfies the equality

$$\langle w(\cdot, t), \psi(\cdot, t) \rangle = \int_0^t \left\langle w(\cdot, s), \left[\psi_s(\cdot, s) - \left[\frac{f(v_1) - f(v_2)}{v_1 - v_2} \right] \psi_x + \psi_{xx} \right] \right\rangle ds$$

for any $t > 0$ and test function ψ satisfying $\psi \in C_0^\infty(\mathbb{R}_+ \times [0, t])$, $\psi(0, s) = 0$ for $s \in [0, t]$. One can then find an \mathbb{L}^∞ approximation of the solution to the linear, backward equation

$$\psi_s(\cdot, s) - \left[\frac{f(v_1) - f(v_2)}{v_1 - v_2} \right] \psi_x + \psi_{xx} = 0$$

on $\mathbb{R}_+ \times [0, t]$ satisfying $\psi(\cdot, t) = \phi(\cdot)$, given any smooth ϕ . As a result, $\langle w(\cdot, t), \phi \rangle = 0$ for any t and any such ϕ , hence $w \equiv 0$.

Next, $\bar{\xi}^n$ is shown to form a tight sequence, and then the next main step is to prove that limit points satisfy (2.16) in weak sense. Achieving this goal relies on two key elements. One is a semimartingale decomposition for point processes, introduced in [17]. This decomposition is particularly convenient in our setting, where the number of point processes involved, given by the departure processes, grows to infinity. The other is an estimate showing C -tightness of \hat{X}_i^n uniformly in $i \in [n]$. With this toolbox we can then show the following. Let $\tilde{\phi}$ be a test function as in Definition 2.5 and ϕ its antiderivative. Then every limit point ξ of $\bar{\xi}^n$ satisfies

$$\begin{aligned} \langle \phi, \xi_t \rangle &= \langle \phi, \xi_0 \rangle + \int_0^t \langle b_1 \phi' + a \phi'', \xi_s \rangle ds \\ (2.19) \quad &+ \frac{b_0}{\ell} \int_0^t \int_{\mathbb{R}_+} \phi'(x) \mathfrak{S}(\xi_s[x, \infty), \xi_s(x, \infty)) \xi_s(dx) ds, \end{aligned}$$

where

$$(2.20) \quad \mathfrak{S}(a, b) = a^{\ell-1} + a^{\ell-2}b + \dots + ab^{\ell-2} + b^{\ell-1}, \quad a, b \in \mathbb{R}.$$

To see the relation to equation (2.17) required by Definition 2.5, let $v(x, t) = \xi_t(x, \infty)$. Then $\langle \phi, \xi_t \rangle = \int_{\mathbb{R}_+} (v(x, t) - 1) \tilde{\phi}(x) dx$. Suppose that ξ_t has no atoms for every $t > 0$. Then, owing to the fact that $\mathfrak{S}(a, a) = \ell a^{\ell-1}$, (2.19) reduces to (2.17) when $f(z) = c_1 z - b z^\ell$, which is the case of interest. However, an a priori proof of the atomless property is not required, as a calculation via integration by parts shows that (2.19) and (2.17) are equivalent even in presence of atoms. As a result, the existence of weak limit of $\bar{\xi}^n$ satisfying the PDE follows, completing the proof of the hydrodynamic limit.

Equation (2.16) is well known to have C^∞ solutions, and as a consequence of the above result the relation of the hydrodynamic limit to the PDE (1.3) for the density follows.

Finally, the two invariance principles follow by combining the C -tightness of individual rescaled queue lengths with the hydrodynamic limit result, where the atomless property, that by now has been proved, gives a control over the interaction term.

We now provide some more details on the structure of Sections 3–4. In Section 3.1 we prove uniqueness of weak solutions to (2.16). The aforementioned semimartingale representations are provided in Section 3.2. In Section 3.3 uniform estimates on the rescaled queueing

processes are derived, including C -tightness and second moment. Tightness of $\bar{\xi}^n$ is shown in Section 3.4.

In Section 4.1 equation (4.4) for $\langle \phi, \bar{\xi}^n \rangle$ is written and it is shown that its linear term converges to the linear term in (2.19). In §4.2, the interaction term is shown to converge to that in (2.19). The reduction, alluded to above, of (2.19) to (2.17), is shown in Section 4.3. In Section 4.4, the convergence of $\bar{\xi}^n$ to a solution of (2.17) is deduced from the above steps, and, based on that, the proof of all main results is then provided.

3. Preliminaries.

3.1. *PDE uniqueness.* Here we prove the uniqueness part of Theorem 2.6, stated as Lemma 3.2 below. In preparation for proving this result, we provide an extension of the class of test functions allowed in Definition 2.5, which is relatively standard but is given for completeness.

LEMMA 3.1. *If v is a weak solution of (2.16) (in the sense of Definition 2.5) then for any $t \in (0, \infty)$ and any $\psi \in C_0^\infty(\mathbb{R}_+ \times [0, t])$ satisfying $\psi(0, s) = 0, 0 \leq s \leq t$, one has*

$$(3.1) \quad \begin{aligned} \langle v(\cdot, t), \psi(\cdot, t) \rangle - \langle v_0, \psi(\cdot, 0) \rangle &= \int_0^t \langle v(\cdot, s), \psi_s(\cdot, s) \rangle ds - \int_0^t \langle f(v(\cdot, s)), \psi_x \rangle ds \\ &+ a \int_0^t \langle v(\cdot, s), \psi_{xx} \rangle ds + a \int_0^t \psi_x(0, s) ds. \end{aligned}$$

PROOF. Observe by (2.17) that for any function ϕ satisfying Definition 2.5 and for any $0 \leq \tau_1 < \tau_2 \leq t$,

$$(3.2) \quad \langle v(\cdot, \tau_2), \phi \rangle = \langle v(\cdot, \tau_1), \phi \rangle - \int_{\tau_1}^{\tau_2} \langle f(v(\cdot, s)), \phi' \rangle ds + a \int_{\tau_1}^{\tau_2} (\langle v(\cdot, s), \phi'' \rangle + \phi'(0)) ds.$$

Since v and $f(v)$ are in $\mathbb{L}^\infty(\mathbb{R}_+)$ uniformly in t , it follows, in particular, that $t \mapsto \langle v(\cdot, t), \phi \rangle$ is continuous and, moreover,

$$(3.3) \quad | \langle v(\cdot, \tau_2) - v(\cdot, \tau_1), \phi \rangle | \leq |\tau_1 - \tau_2| K,$$

where $K = K(\phi, v)$.

Let now $\psi \in C_0^\infty(\mathbb{R}_+ \times [0, t])$ be a function that satisfies the conditions of the lemma. Let $N \in \mathbb{N}$. Let $t_k := tk/N$ where $k = 0, \dots, N$. Define $\phi_k(x) = \psi(t_k, x)$. Then, from (3.2)

$$\langle v(\cdot, t_{k+1}), \phi_k \rangle - \langle v(\cdot, t_k), \phi_k \rangle = - \int_{t_k}^{t_{k+1}} \langle f(v(\cdot, s)), \phi'_k \rangle ds + a \int_{t_k}^{t_{k+1}} (\langle v(\cdot, s), \phi''_k \rangle + \phi'_k(0)) ds.$$

Summing over $k = 0, \dots, N - 1$ and using $\phi_N = \psi(\cdot, t), \phi_0 = \psi(\cdot, 0)$ we get

$$(3.4) \quad \begin{aligned} \langle v(\cdot, t), \psi(\cdot, t(1 - 1/N)) \rangle - \langle v_0, \psi(\cdot, 0) \rangle &= \sum_{k=1}^{N-1} \langle v(\cdot, t_{k+1}), \phi_{k+1} - \phi_k \rangle \\ &= \sum_{k=0}^{N-1} \left[- \int_{t_k}^{t_{k+1}} \langle f(v(\cdot, s)), \phi'_k \rangle ds + a \int_{t_k}^{t_{k+1}} (\langle v(\cdot, s), \phi''_k \rangle + \phi'_k(0)) ds \right]. \end{aligned}$$

Evidently, $\langle v(\cdot, t), \psi(\cdot, t(1 - 1/N)) \rangle \rightarrow \langle v(\cdot, t), \psi(\cdot, t) \rangle$ as $N \rightarrow \infty$. Next, by (3.3)

$$\begin{aligned} \langle v(\cdot, t_{k+1}), \phi_{k+1} - \phi_k \rangle &= \int_{t_k}^{t_{k+1}} \langle v(\cdot, s), \psi_s(\cdot, s) \rangle ds + \int_{t_k}^{t_{k+1}} \langle v(\cdot, t_{k+1}) - v(\cdot, s), \psi_s(\cdot, s) \rangle ds \\ &= \int_{t_k}^{t_{k+1}} \langle v(\cdot, s), \psi_s(\cdot, s) \rangle ds + O(N^{-2}). \end{aligned}$$

Thus

$$\sum_{k=1}^{N-1} \langle v(\cdot, t_{k+1}), \phi_{k+1} - \phi_k \rangle \rightarrow \int_0^t \langle v(\cdot, s), \psi_s(\cdot, s) \rangle ds$$

as $N \rightarrow \infty$. Likewise, the right side of (3.4) converges to

$$- \int_0^t \langle \mathfrak{f}(v(\cdot, s)), \psi_x \rangle ds + a \int_0^t (\langle v(\cdot, s), \psi_{xx} \rangle + \psi_x(0, s)) ds.$$

This proves the lemma. \square

LEMMA 3.2. Assume $\mathfrak{f} \in C^\infty(\mathbb{R})$ and $v_0 \in \mathbb{L}^1(\mathbb{R}_+) \cap \mathbb{L}^\infty(\mathbb{R}_+)$. Then there is at most one weak solution to (2.16) in the sense of Definition 2.5.

PROOF. A change of variables $x \mapsto a^{-1/2}x$ transforms the first line of equation (2.16) into $v_t = (\tilde{\mathfrak{f}}(v))_x + v_{xx}$, where $\tilde{\mathfrak{f}} = a^{-1/2}\mathfrak{f}$. Therefore we may and will assume w.l.o.g. that $a = 1$. Suppose v_1, v_2 are two solutions. The goal is to show that $v_1 = v_2$. Let $w = v_1 - v_2$ and note that $w \in \mathbb{L}_{loc}^\infty(\mathbb{R}_+, \mathbb{L}^\infty(\mathbb{R}_+)) \cap \mathbb{L}_{loc}^\infty(\mathbb{R}_+, \mathbb{L}^1(\mathbb{R}_+))$. Let

$$J(x, t) := \begin{cases} \frac{\mathfrak{f}(v_1(x, t)) - \mathfrak{f}(v_2(x, t))}{v_1(x, t) - v_2(x, t)} & \text{if } w(x, t) \neq 0, \\ \mathfrak{f}'(v_1(x, t)) & \text{if } w(x, t) = 0. \end{cases}$$

Then for any test function ψ satisfying the conditions of Lemma 3.1,

$$(3.5) \quad \langle w(\cdot, t), \psi(\cdot, t) \rangle = \int_0^t \langle w(\cdot, s), [\psi_s(\cdot, s) - J(\cdot, s)\psi_x + \psi_{xx}] \rangle ds.$$

Fix T . By the assumptions on \mathfrak{f} and the definition of a solution, $J \in \mathbb{L}^\infty([0, T]; \mathbb{L}^\infty(\mathbb{R}_+))$. Moreover, for any $x_0 > 0$, $|\mathfrak{f}'(x) - \mathfrak{f}'(0)| \leq cx$ holds provided that $|x| < x_0$, with $c = c(x_0)$. Hence $J - \mathfrak{f}'(0) \in \mathbb{L}^\infty([0, T]; \mathbb{L}^1(\mathbb{R}_+))$.

Let η denote the standard mollifier in \mathbb{R} , $\eta_N = N\eta(N\cdot)$, and

$$J_N(x, t) = \int_{\mathbb{R}_+} \eta_N(x - y)J(y, t) dy.$$

Then J_N is a sequence bounded uniformly in $\mathbb{L}^\infty(\mathbb{R}_+ \times [0, T])$, and, for every $t \in [0, T]$, $J_N(\cdot, t) \in C^\infty(\mathbb{R}_+)$. Moreover, one has $\|J_N - J\|_1 \rightarrow 0$ as $N \rightarrow \infty$, where we denote $\|\cdot\|_1 = \|\cdot\|_{\mathbb{L}^1(\mathbb{R}_+ \times [0, T])}$ (see, e.g., [22], Lemma 7.2). By reparameterization, we may assume w.l.o.g. that $\|J_N - J\|_1 = O(N^{-1})$.

Given $t \in (0, T]$, let $\tilde{\psi}^N$ be the classical solution of the backward linear problem on the time interval $[0, t]$:

$$(3.6) \quad \begin{aligned} \tilde{\psi}_s^N(\cdot, s) - J_N(\cdot, s)\tilde{\psi}_x^N + \tilde{\psi}_{xx}^N &= 0, & \tilde{\psi}^N(\cdot, t) &= \phi, \\ \tilde{\psi}^N(0, s) &= 0, & 0 \leq s < t, \end{aligned}$$

where $\phi \in C_0^\infty(\mathbb{R}_+)$. Since $\tilde{\psi}^N$ and all its derivatives decay to zero as $x \rightarrow \infty$, uniformly in $s \in [0, t]$, we may replace them by ψ^N which satisfy the conditions of Lemma 3.1 and

$$\|\psi^N - \tilde{\psi}^N\|_\infty + \|\psi_x^N - \tilde{\psi}_x^N\|_\infty + \|\psi_{xx}^N - \tilde{\psi}_{xx}^N\|_\infty + \|\psi_t^N - \tilde{\psi}_t^N\|_\infty = O(N^{-1})$$

uniformly on $[0, t]$. It follows from (3.5) that

$$\langle w(\cdot, t), \phi \rangle = \int_0^t \langle w(\cdot, s), [\tilde{\psi}_s^N(\cdot, s) - J(\cdot, s)\tilde{\psi}_x^N + \tilde{\psi}_{xx}^N] \rangle ds$$

$$\begin{aligned} &= \int_0^t \langle w(\cdot, s), [J_N(\cdot, s) - J]\tilde{\psi}_x^N \rangle ds \\ &= \int_0^t \langle w(\cdot, s), [J_N(\cdot, s) - J]\psi_x^N \rangle ds + O(N^{-1}) \end{aligned}$$

for any such ϕ . Suppose

$$(3.7) \quad \sup_N \sup_{s \in [0, t]} \|\tilde{\psi}_x^N(\cdot, s)\|_\infty < \infty$$

if t is small enough. Since $w \in \mathbb{L}^\infty([0, T]; \mathbb{L}^\infty(\mathbb{R}_+))$ and $\|J_N - J\|_1 = O(N^{-1})$, this implies $\langle w(\cdot, t), \phi \rangle = O(N^{-1})$ for any $\phi \in C_0^\infty(\mathbb{R}_+)$ and any N , hence $w(\cdot, t) = 0$. As a consequence, $w(\cdot, t) = 0$ for all $t \in [0, t_0]$, some $t_0 > 0$. Moreover, t_0 does not depend on the initial condition v_0 . Thus, iterating the argument shows that $w = 0$.

It thus suffices to show (3.7). To this end, denote $M := \sup_N \|J_N\|_\infty < \infty$. Let ψ^0 be the solution of

$$\psi_s^0(\cdot, s) + \psi_{xx}^0 = 0, \quad \psi^0(\cdot, t) = \phi, \quad \psi^0(0, s) = 0, \quad 0 \leq s \leq t.$$

Let $m_N(y, \tau) := J_N(y, \tau)\tilde{\psi}_y^N(y, \tau)$. Since, as mentioned earlier, $\|\tilde{\psi}_y^N\|_\infty < \infty$ for all N , one also has $\|m_N\|_\infty < \infty$ for all N . Moreover, $\tilde{\psi}^N = \psi_0 + \hat{\psi}^N$ where $\hat{\psi}^N$ is the solution of

$$\hat{\psi}_s^N + \hat{\psi}_{xx}^N = m_N, \quad \hat{\psi}^N(\cdot, t) = 0, \quad \hat{\psi}^N(0, s) = 0, \quad 0 \leq s \leq t.$$

The solution $\hat{\psi}^N$ is given by Duhamel’s principle:

$$\begin{aligned} \hat{\psi}^N(x, s) &= \frac{1}{2\sqrt{\pi}} \int_s^t d\tau (t - \tau)^{-1/2} \\ &\quad \times \left[\int_0^\infty m_N(y, t - \tau) e^{-\frac{(x-y)^2}{4(t-\tau)}} dy - \int_0^\infty m_N(y, t - \tau) e^{-\frac{(x+y)^2}{4(t-\tau)}} dy \right]. \end{aligned}$$

Then $\hat{\psi}_x^N(x, s) = -A_N(x, s) + B_N(x, s)$, where

$$A_N(x, s) = \frac{1}{4\sqrt{\pi}} \int_s^t d\tau (t - \tau)^{-3/2} \int_0^\infty m_N(y, t - \tau)(x - y) e^{-\frac{(x-y)^2}{4(t-\tau)}} dy,$$

$$B_N(x, s) = \frac{1}{4\sqrt{\pi}} \int_s^t d\tau (t - \tau)^{-3/2} \int_0^\infty m_N(y, t - \tau)(x + y) e^{-\frac{(x+y)^2}{4(t-\tau)}} dy.$$

Both $|A_N(x, s)|$ and $|B_N(x, s)|$ are bounded by

$$\frac{\|m_N\|_\infty}{4\sqrt{\pi}} \int_s^t d\tau (t - \tau)^{-3/2} \int_{-\infty}^\infty |x + y| e^{-\frac{(x+y)^2}{4(t-\tau)}} dy.$$

Changing variables of integration,

$$|A_N(x, s)| \vee |B_N(x, s)| \leq \frac{\|m_N\|_\infty}{2\sqrt{\pi}} \int_s^t d\tau (t - \tau)^{-1/2} \int_{-\infty}^\infty |z| e^{-\frac{z^2}{4}} dz \leq C \|m_N\|_\infty t^{1/2},$$

for a constant C . Recall that $\|m_N\|_\infty \leq M \|\psi_x^N\|_\infty$ and $\tilde{\psi}^N = \psi^0 + \hat{\psi}^N$. Thus

$$\|\tilde{\psi}_x^N(\cdot, s)\|_\infty \leq \|\psi_x^0\|_\infty + 2C \|m_N\|_\infty t^{1/2} \leq \|\psi_x^0\|_\infty + 2CM \|\psi_x^N(\cdot, s)\|_\infty t^{1/2}, \quad 0 < s \leq t.$$

This shows $\|\tilde{\psi}_x^N(\cdot, s)\|_\infty \leq 2\|\psi_x^0\|_\infty$ for $0 \leq s \leq (2CM)^{-2}/2$ and completes the proof that $v_1 = v_2$. \square

3.2. *Martingale toolbox.* A semimartingale decomposition of counting processes was introduced in [17], which deviates from the classical Doob–Meyer decomposition and is convenient for our purposes. For the renewal process S_i^n , this decomposition is constructed as follows. Denote by

$$R_i^n(t) = \inf\{s > 0 : S_i^n(t + s) > S_i^n(t)\}, \quad t \geq 0,$$

the residual time to the next counting instant (note that it is right-continuous, hence at a time of counting it already shows the time until the next counting). One has

$$(3.8) \quad t + R_i^n(t) = \sum_{k=0}^{S_i^n(t)} Z_i^n(k).$$

Hence with

$$M_i^{\text{ser},n}(t) = \sum_{k=1}^{S_i^n(t)} \zeta_i^n(k), \quad \zeta_i^n(k) = 1 - \mu^n Z_i^n(k), \quad k \geq 1,$$

and $\sum_1^0 = 0$, one has the Daley–Miyazawa semimartingale representation,

$$(3.9) \quad S_i^n(t) = \mu^n(t - Z_i^n(0) + R_i^n(t)) + M_i^{\text{ser},n}(t), \quad t \geq 0.$$

Denoting $\mathcal{F}_i^{\text{ser},n}(t) = \sigma\{S_i^n(s), R_i^n(s), s \leq t\}$, the first term on the right is adapted to the filtration $\{\mathcal{F}_i^{\text{ser},n}(t)\}$ while $M_i^{\text{ser},n}$ a martingale on it. A form of this decomposition that will be useful here is based on a filtration \mathcal{F}_i^n that, for each t , contains all relevant information about the system by time t . It will be obtained once a time change transformation of (3.9) is performed, to represent $D_i^n(t) = S_i^n(T_i^n(t))$.

To this end, some further notation is required. In addition to the delayed renewal processes S_i^n it will also be useful to introduce the corresponding nondelayed renewal processes

$$S_i^{0,n}(t) = \max\left\{k \in \mathbb{Z}_+ : \sum_{j=1}^{k-1} Z_i^n(j) \leq t\right\}, \quad t \geq 0.$$

Note that $S_i^{0,n}(0) = 1$ and that these processes are i.i.d. (unlike S_i^n). Additional rescaled processes are denoted as follows

$$(3.10) \quad \hat{E}_i^n(t) = n^{-1/2}(E_i^n(t) - \lambda^n t), \quad \hat{S}_i^n(t) = n^{-1/2}(S_i^n(t) - \mu^n t),$$

$$\hat{S}_i^{0,n}(t) = n^{-1/2}(S_i^{0,n}(t) - \mu^n t),$$

$$(3.11) \quad \hat{A}_0^n(t) = n^{-1/2}(A_0^n(t) - \lambda_0^n t), \quad \hat{A}_i^n(t) = n^{-1/2}A_i^n(t).$$

Next, let

$$\tilde{R}_i^n(t) = n^{-1/2}\mu^n R_i^n(t) = \tilde{\mu}^n R_i^n(t), \quad \hat{M}_i^{\text{ser},n}(t) = n^{-1/2}M_i^{\text{ser},n}(t),$$

and

$$(3.12) \quad M_i^{\text{dep},n}(t) = \sum_{k=1}^{D_i^n(t)} \zeta_i^n(k), \quad \hat{M}_i^{\text{dep},n}(t) = n^{-1/2}M_i^{\text{dep},n}(t).$$

Then, by (3.9),

$$(3.13) \quad D_i^n(t) = \mu^n(T_i^n(t) - Z_i^n(0) + R_i^n(T_i^n(t))) + M_i^{\text{dep},n}(t),$$

and

$$(3.14) \quad \hat{S}_i^n(T_i^n(t)) = -\tilde{Z}_i^n(0) + \tilde{R}_i^n(T_i^n(t)) + \hat{M}_i^{\text{dep},n}(t).$$

The tuple

$$S^n(t) = (E_i^n(t), A_i^n(t), D_i^n(t), X_i^n(t), T_i^n(t), R_i^n(T_i^n(t)), i \in [n], A_0^n(t), \theta_{A_0^n(t)}^n)$$

is referred to as the *state of the system at time t*. Denote the corresponding filtration by

$$\mathcal{F}_t^n = \sigma\{\mathcal{I}^n, S^n(s), s \in [0, t]\}.$$

Note that at times t when server i is active, $R_i^n(T_i^n(t))$ is the residual time till the completion of service of the job being processed by that server, whereas at times when the server is idle, $R_i^n(T_i^n(t))$ gives the service duration of the job that will be processed next by this server.

LEMMA 3.3. (i) *The processes $\hat{M}_i^{\text{dep},n}$, \hat{E}_i^n , and \hat{A}_0^n are $\{\mathcal{F}_t^n\}$ -martingales, with optional quadratic variations given by*

$$[\hat{M}_i^{\text{dep},n}](t) = n^{-1} \sum_{k=1}^{D_i^n(t)} \zeta_i^n(k)^2, \quad [\hat{E}_i^n](t) = n^{-1} E_i^n(t), \quad [\hat{A}_0^n](t) = n^{-1} A_0^n(t),$$

and one has

$$(3.15) \quad \mathbb{E}\{[\hat{M}_i^{\text{dep},n}](t)\} = n^{-1} \sigma_{\text{ser}}^2 \mathbb{E}\{D_i^n(t)\} < \infty.$$

(ii) *For distinct $i, j \in [n]$,*

$$(3.16) \quad \mathbb{E}\{[\hat{M}_i^{\text{dep},n}, \hat{M}_j^{\text{dep},n}](t)\} = 0.$$

(iii) *The process*

$$\hat{M}_i^{A,n}(t) = \hat{A}_i^n(t) - \hat{C}_i^{A,n}(t) \quad \text{where } \hat{C}_i^{A,n}(t) = \lambda_0^n n^{-1/2} \int_0^t p_{n,\mathcal{R}_i^n(s)} ds,$$

is an $\{\mathcal{F}_t^n\}$ -martingale, whose optional quadratic variation is $[\hat{M}_i^{A,n}](t) = n^{-1} A_i^n(t)$.

PROOF. (i) For adaptedness of $\hat{M}_i^{\text{dep},n}$ it suffices to prove that $Z_i^n(k)1_{k \leq D_i^n(t)} \in \mathcal{F}_t^n$ for all k . This is shown as follows. By (3.8), $T_i^n(s) + R_i^n(T_i^n(s)) = \sum_{k=0}^{D_i^n(s)} Z_i^n(k)$. Hence $\{Z_i^n(k), k \leq D_i^n(t)\}$ can all be recovered from the tuple $\{T_i^n(s), R_i^n(T_i^n(s)), D_i^n(s)\}$, as s varies between 0 and t . Since the latter is $\{\mathcal{F}_t^n\}$ -adapted, this proves the claim.

Next it is shown that $\hat{M}_i^{\text{dep},n}(t) \in L_1(d\mathbb{P})$. Note first that as a renewal process, $S_i^n(t)$ has finite expectation for every t . Since $T_i^n(t) \leq t$ this gives $\mathbb{E}[D_i^n(t)] < \infty$. Let

$$t_i^n(k) = \inf\{t \geq 0 : D_i^n(t) \geq k\}, \quad k = 1, 2, \dots$$

These are clearly stopping times on $\{\mathcal{F}_t^n\}$. Hence

$$(3.17) \quad t_i^n(k) \in \mathcal{F}_{t_i^n(k)-}^n, \quad k \geq 1,$$

where we recall that for a stopping time τ ,

$$(3.18) \quad \mathcal{F}_{\tau-}^n = \mathcal{F}_0^n \vee \sigma\{A \cap \{t < \tau\} : A \in \mathcal{F}_t^n, t \geq 0\}$$

(see [27], I.1.11 and I.1.14). The state of the system up to $t_i^n(k)-$, namely $\{S^n(t), t < t_i^n(k)\}$, can be recovered from the tuple \mathcal{I}^n , $(E_i^n(t), t \in \mathbb{R}_+, i \in [n])$, $(A_0^n(t), t \in \mathbb{R}_+)$, $(\theta_j^n, j \in \mathbb{N})$, $(Z_\ell^n(j), j \in \mathbb{N}), \ell \in [n] \setminus \{i\})$ and $(Z_i^n(j), j < k)$, as follows by the construction of the model. By our assumptions, $Z_i^n(k)$ is independent of this tuple. As a result, it is independent of $\mathcal{F}_{t_i^n(k)-}^n$. It follows from (3.17) that

$$\{D_i^n(t) \leq k\} = \{t_i^n(k) \geq t\} \in \mathcal{F}_{t_i^n(k)-}^n.$$

The structure that we have just proved, where

$$(3.19) \quad Z_i^n(k') \in \mathcal{F}_{t_i^n(k)-}^n, \quad k' < k \text{ whereas } Z_i^n(k) \text{ is independent of } \mathcal{F}_{t_i^n(k)-}^n,$$

the fact that $D_i^n(t)$ is a stopping time on the discrete parameter filtration $\{\mathcal{F}_{t_i^n(k)-}^n, k \in \mathbb{N}\}$, $\mathbb{E}[D_i^n(t)] < \infty$, along with the fact that $Z_i^n(k)$ are i.i.d. with $\mathbb{E}[Z_i^n(k)] < \infty$, allows us to use Wald's identity, showing

$$\mathbb{E} \left[\sum_{k=1}^{D_i^n(t)} Z_i^n(k) \right] = \mathbb{E}\{D_i^n(t)\} \mathbb{E}\{Z_i^n(1)\} = \mathbb{E}\{D_i^n(t)\}(\mu^n)^{-1} < \infty.$$

Since $|\zeta_i^n(k)| \leq 1 + \mu^n Z_i^n(k)$, this shows that $\mathbb{E}\{\hat{M}_i^{\text{dep},n}(t)\} < \infty$.

To show the martingale property, note that by the independence stated in (3.19), we have

$$(3.20) \quad \mathbb{E}[\zeta_i^n(k) | \mathcal{F}_{t_i^n(k)-}^n] = 0, \quad k \geq 1.$$

Arguing now along the lines of the proof of [17], Lemma 2.1,

$$\hat{M}_i^{\text{dep},n}(t) = n^{-1/2} \sum_{k=1}^{D_i^n(t)} \zeta_i^n(k) = n^{-1/2} \sum_{k=1}^{\infty} \zeta_i^n(k) 1_{\{t_i^n(k) \leq t\}}.$$

Hence for $s < t$,

$$\mathbb{E}[\hat{M}_i^{\text{dep},n}(t) | \mathcal{F}_s^n] - \hat{M}_i^{\text{dep},n}(s) = n^{-1/2} \sum_{k=1}^{\infty} \mathbb{E}[\zeta_i^n(k) 1_{\{s < t_i^n(k) \leq t\}} | \mathcal{F}_s^n].$$

To show that each term on the right is zero, let $A \in \mathcal{F}_s^n$. By (3.18), $A \cap \{s < t_i^n(k)\} \in \mathcal{F}_{t_i^n(k)-}^n$. Also, by (3.17), $\{t_i^n(k) \leq t\} \in \mathcal{F}_{t_i^n(k)-}^n$. With (3.20), this gives

$$(3.21) \quad \begin{aligned} \mathbb{E}[\zeta_i^n(k) 1_{\{s < t_i^n(k) \leq t\}} 1_A] &= \mathbb{E}[\mathbb{E}[\zeta_i^n(k) 1_{\{s < t_i^n(k) \leq t\}} 1_A | \mathcal{F}_{t_i^n(k)-}^n]] \\ &= \mathbb{E}[\mathbb{E}[\zeta_i^n(k) | \mathcal{F}_{t_i^n(k)-}^n] 1_{\{s < t_i^n(k) \leq t\}} 1_A] = 0. \end{aligned}$$

Since this holds for all $A \in \mathcal{F}_s^n$, it follows that

$$\mathbb{E}[\hat{M}_i^{\text{dep},n}(t) | \mathcal{F}_s^n] - \hat{M}_i^{\text{dep},n}(s) = 0.$$

For the martingale property of \hat{E}_i^n one only needs to show that $E_i^n(t) - E_i^n(s)$ is independent of \mathcal{F}_s^n when $s < t$. Again, this follows from the fact that all the processes comprising $S_u^n, u \leq s$, can be recovered from the tuple $\mathcal{I}^n, (E_i^n(u), u \in [0, s], i \in [n]), (A_0^n(u), u \in [0, s]), (\theta_k^n, k \in \mathbb{N}), (Z_i^n(k), k \in \mathbb{N}, i \in [n])$; but the increment $E_i^n(t) - E_i^n(s)$ is independent of this tuple.

A similar proof holds for \hat{A}_0^n .

The expressions for the quadratic variation are straightforward.

To show (3.15) we can use Wald's identity as before, now with the i.i.d. sequence $\zeta_i^n(k)^2$, which now gives

$$\mathbb{E} \left[\sum_{k=1}^{D_i^n(t)} \zeta_i^n(k)^2 \right] = \mathbb{E}\{D_i^n(t)\} \mathbb{E}\{\zeta_i^n(1)^2\},$$

and (3.15) follows.

(ii) Similar to the argument following (3.17), one can recover the state of the system up to t_{ijkl}^n , where

$$t_{ijkl}^n = \min(t_i^n(k), t_j^n(l)),$$

namely $\{\mathcal{S}^n(t), t < t_{ijkl}^n\}$, from \mathcal{I}^n , $(E_i^n(t), t \in \mathbb{R}_+, i \in [n])$, $(A_0^n(t), t \in \mathbb{R}_+)$, $(\theta_j^n, j \in \mathbb{N})$, $(Z_i^n(k), k \in \mathbb{N})$, $i' \in [n] \setminus \{i, j\}$, $(Z_i^n(k'), k' < k)$ and $(Z_j^n(l'), l' < l)$. However, by our assumptions, the pair $(Z_i^n(k), Z_j^n(l))$ is independent of this tuple, hence it is independent of $\mathcal{F}_{t_{ijkl}^n}^n$. Since $Z_i^n(k)$ and $Z_j^n(l)$ are mutually independent, this gives

$$(3.22) \quad \mathbb{E}[\zeta_i^n(k)\zeta_j^n(l)|\mathcal{F}_{t_{ijkl}^n}^n] = 0, \quad k, l \geq 1.$$

We have

$$\begin{aligned} [\hat{M}_i^{\text{dep},n}, \hat{M}_j^{\text{dep},n}](t) &= n^{-1} \sum_{k=1}^{D_i^n(t)} \zeta_i^n(k) \sum_{l=1}^{D_j^n(t)} 1_{\{t_j^n(l)=t_i^n(k)\}} \zeta_j^n(l) \\ &= n^{-1} \sum_{k=1}^{\infty} \sum_{l=1}^{\infty} \zeta_i^n(k)\zeta_j^n(l) 1_{\{t_j^n(l)=t_i^n(k) \leq t\}}. \end{aligned}$$

In view of (3.17), $1_{\{t_j^n(l)=t_i^n(k) \leq t\}} \in \mathcal{F}_{t_{ijkl}^n}^n$. Hence by (3.22) and an argument as in (3.21),

$$\mathbb{E}[\zeta_i^n(k)\zeta_j^n(l) 1_{\{t_j^n(l)=t_i^n(k) \leq t\}} | \mathcal{F}_{t_{ijkl}^n}^n] = 0, \quad k, l \geq 1.$$

Thus (3.16) follows.

(iii) Clearly A_i^n , defined in (2.6) is adapted and $A_i^n(t)$ is integrable for all t . Hence the same is true for \hat{M}_i^n . Next, let

$$s^n(k) = \inf\{t \geq 0 : A_0^n(t) \geq k\}, \quad k = 1, 2, \dots$$

As in (i), these are stopping times and $s^n(k) \in \mathcal{F}_{s^n(k)-}^n$. To show the martingale property, we can write, using $\int_0^t p_{n, \mathcal{R}_i^n(s-)} ds = \int_0^t p_{n, \mathcal{R}_i^n(s)} ds$,

$$\begin{aligned} n^{1/2} \hat{M}_i^{A,n}(t) &= \int_{[0,t]} 1_{\{\mathcal{R}_i^n(s-) = \theta_{A_0^n(s)}^n\}} dA_0^n(s) - \lambda_0^n \int_0^t p_{n, \mathcal{R}_i^n(s)} ds \\ &= \int_{[0,t]} (1_{\{\mathcal{R}_i^n(s-) = \theta_{A_0^n(s)}^n\}} - p_{n, \mathcal{R}_i^n(s-)}) dA_0^n(s) + \int_0^t p_{n, \mathcal{R}_i^n(s-)} (dA_0^n(s) - \lambda_0^n ds) \\ &=: M_{i,1}^n(t) + M_{i,2}^n(t). \end{aligned}$$

For $M_{i,1}^n$, write

$$A_i^n(t) = \sum_{k=1}^{A_0^n(t)} 1_{\{\mathcal{R}_i^n(s^n(k)-) = \theta_k^n\}}.$$

An argument similar to the one given before shows that θ_k^n is independent of $\mathcal{F}_{s^n(k)-}^n$. Therefore, for $0 \leq s < t$, we have

$$\begin{aligned} \mathbb{E}[A_i^n(t) | \mathcal{F}_s^n] - A_i^n(s) &= \sum_{k=1}^{\infty} \mathbb{E}[1_{\{\mathcal{R}_i^n(s^n(k)-) = \theta_k^n\}} 1_{\{s < s^n(k) \leq t\}} | \mathcal{F}_s^n] \\ &= \sum_{k=1}^{\infty} \mathbb{E}[\mathbb{E}[1_{\{\mathcal{R}_i^n(s^n(k)-) = \theta_k^n\}} 1_{\{s < s^n(k) \leq t\}} | \mathcal{F}_{s^n(k)-}^n] | \mathcal{F}_s^n] \\ &= \sum_{k=1}^{\infty} \mathbb{E}[p_{n, \mathcal{R}_i^n(s^n(k)-)} 1_{\{s < s^n(k) \leq t\}} | \mathcal{F}_s^n] \\ &= \mathbb{E}[C_i^n(t) | \mathcal{F}_s^n] - C_i^n(s), \end{aligned}$$

where

$$C_i^n(t) = \sum_{k=1}^{A_0^n(t)} P_{n, \mathcal{R}_i^n(s^n(k)-)} = \int_{[0,t]} P_{n, \mathcal{R}_i^n(s-)} dA_0^n(s),$$

showing that $A_i^n - C_i^n = M_{i,1}^n$ is a martingale.

In the expression for $M_{i,2}^n$, the integrand is $\{\mathcal{F}_t^n\}$ -adapted and has LCRL sample paths, while the integrator is a martingale on the filtration. As a result, $M_{i,2}^n$ is a local martingale [38], Theorem II.20; using the estimate $\|M_{i,2}^n\|_t^* \leq A_0^n(t) + c$ shows it is in fact a martingale. As a result, so is $\hat{M}_i^{A,n}$. Finally, the expression for the quadratic variation is straightforward. \square

We will also need the following simple fact.

LEMMA 3.4. *Let $M_N, N \in \mathbb{Z}_+$ be a martingale with $M_0 = 0$, for which the increments $\Delta_N = M_N - M_{N-1}$ satisfy $\mathbb{E}[|\Delta_N|1_{\{|\Delta_N|>a\}}] \leq \bar{r}(a), a \geq 0, N \in \mathbb{N}$, where $\bar{r}(a) \rightarrow 0$ as $a \rightarrow \infty$. Then $N^{-1}\mathbb{E}\|M\|_N^* < A(N) \rightarrow 0$, where $\{A(N)\}$ depend only on \bar{r} .*

PROOF. Let $b_N = \mathbb{E}[\Delta_N 1_{\{|\Delta_N| \leq a\}}]$ and note that $b_N = -\mathbb{E}[\Delta_N 1_{\{|\Delta_N| > a\}}]$ and consequently $|b_N| \leq \bar{r}(a)$. Write

$$M_N = P_N + Q_N, \quad P_N = \sum_{i=1}^N \{\Delta_i 1_{\{|\Delta_i| > a\}} + b_i\}, \quad Q_N = \sum_{i=1}^N \{\Delta_i 1_{\{|\Delta_i| \leq a\}} - b_i\}.$$

The quadratic variation of the martingale Q_N is bounded by $(a + \bar{r}(a))^2 n$. This shows that $\mathbb{E}[\|Q\|_N^*] \leq c(a + \bar{r}(a))n^{1/2}$, by the Burkholder–Davis–Gundy (BDG) inequality with $p = 2$. For P_N ,

$$|P_N| \leq \sum_{i=1}^N |\Delta_i| 1_{\{|\Delta_i| > a\}} + N\bar{r}(a),$$

thus $\mathbb{E}[\|P\|_N^*] \leq 2N\bar{r}(a)$. This gives $n^{-1}\mathbb{E}[\|M\|_N^*] \leq c(a + \bar{r}(a))N^{-1/2} + 2\bar{r}(a)$. Taking $a = N^{1/4}$ completes the proof. \square

3.3. *Uniform estimates on individual queue length processes.* The goal here is to calculate the dynamics of the individual rescaled processes and develop estimates showing that they are C -tight uniformly in $i \in [n]$.

In the following lemma, part (i) provides an equation, (3.23), for the dynamics of individual queue lengths, and part (ii) shows that, at the cost of introducing an error term, one can replace the term $\hat{S}_i^n(T_i^n)$ in that equation by a martingale. Both representations (3.23) and (3.26) are used in this paper, where the former is used for estimates on each \hat{X}_i^n that are uniform in i , and the latter is convenient for representations of $\bar{\xi}^n$, as the martingale terms add up to a martingale. The last part of the lemma uses (3.23) and gives uniform second moment and tightness estimates.

LEMMA 3.5. (i) *One has*

$$(3.23) \quad \hat{X}_i^n(t) = \hat{X}_i^n(0-) + \hat{E}_i^n(t) + \hat{A}_i^n(t) - \hat{S}_i^n(T_i^n(t)) + \hat{b}_1^n t + \hat{L}_i^n(t).$$

Moreover, the sample paths of \hat{L}_i^n are in C^\uparrow and

$$(3.24) \quad \int_0^\infty \hat{X}_i^n(t) d\hat{L}_i^n(t) = 0.$$

(ii) One has

$$(3.25) \quad \hat{X}_i^n(t) = \hat{X}_i^{1,n}(t) + e_i^{1,n}(t),$$

where

$$(3.26) \quad \begin{aligned} \hat{X}_i^{1,n}(t) &= \hat{X}_i^n(0-) + \hat{E}_i^n(t) + \hat{A}_i^n(t) - \hat{M}_i^{\text{dep},n}(t) + \hat{b}_1^n t + \hat{L}_i^n(t), \\ e_i^{1,n}(t) &= \tilde{Z}_i^n(0) - \tilde{R}_i^n(T_i^n(t)) = \tilde{\mu}^n(Z_i^n(0) - R_i^n(T_i^n(t))). \end{aligned}$$

(iii) For $H_i^n = \hat{S}_i^n, \hat{E}_i^n, \hat{A}_i^n, \hat{L}_i^n$ and \hat{X}_i^n , one has

$$(3.27) \quad \sup_n \sup_{i \in [n]} \mathbb{E}[(\|H_i^n\|_t^*)^2] < \infty, \quad t \geq 0,$$

and for every $t, \varepsilon > 0$ and $\eta > 0$ there is $\delta > 0$ such that

$$(3.28) \quad \limsup_n \max_{i \in [n]} \mathbb{P}(w_t(H_i^n, \delta) > \varepsilon) < \eta.$$

PROOF. (i.) By (2.1) and (2.12),

$$\begin{aligned} n^{-1/2} X_i^n(t) &= n^{-1/2} X_i^n(0-) + n^{-1/2}(E_i^n(t) - \lambda^n t) + n^{-1/2}(\lambda^n - n\lambda)t + n^{1/2}\lambda t \\ &\quad + n^{-1/2} A_i^n(t) - n^{-1/2}(S_i^n(T_i^n(t)) - \mu^n T_i^n(t)) - n^{-1/2}\mu^n T_i^n(t), \end{aligned}$$

and

$$-n^{-1/2}\mu^n T_i^n(t) = -n^{-1/2}(\mu^n - n\mu)t - n^{1/2}\mu t + \hat{L}_i^n(t).$$

Using (2.11), (3.10) and (3.11) gives (3.23). The properties of \hat{L}_i^n and (3.24) follow from (2.2).

(ii) Using (3.14) in (3.23) gives (3.25).

(iii) By the central limit theorem for renewal processes ([10], Section 17) and the fact that, by (2.8) and (2.10), $n^{-1}(\lambda^n, \mu^n) \rightarrow (\lambda, \mu)$, for each i , $(\hat{E}_i^n, \hat{S}_i^{0,n})$ converge in law to (E, S) , a pair of mutually independent BM starting at zero, with zero drift, and diffusion coefficients $\lambda^{1/2}$ and $\mu^{1/2}\sigma_{\text{ser}}$, respectively (where we recall $\lambda = \mu$).

We prove that (3.27) and (3.28) hold for \hat{S}_i^n by relating these processes to $\hat{S}_i^{0,n}$, whose laws do not depend on i . To this end, note that

$$S_i^n(t) = S_i^{0,n}((t - Z_i^n(0))^+) - 1_{\{t < Z_i^n(0)\}} = S_i^{0,n}(t - Z_i^{\#,n}) - 1_{\{t < Z_i^n(0)\}},$$

where $Z_i^{\#,n} = Z_i^{\#,n}(t) := Z_i^n(0) \wedge t$. Hence by (3.10),

$$\hat{S}_i^n(t) = \hat{S}_i^{0,n}(t - Z_i^{\#,n}) - \tilde{\mu}^n Z_i^{\#,n} - n^{-1/2} 1_{\{t < Z_i^n(0)\}}.$$

As a result,

$$(3.29) \quad \|\hat{S}_i^n\|_t^* \leq \|\hat{S}_i^{0,n}\|_t^* + \tilde{\mu}^n Z_i^n(0) + n^{-1/2} = \|\hat{S}_i^{0,n}\|_t^* + \tilde{Z}_i^n(0) + n^{-1/2},$$

and

$$w_t(\hat{S}_i^n, \delta) \leq w_t(\hat{S}_i^{0,n}, \delta) + \tilde{\mu}^n Z_i^n(0) = w_t(\hat{S}_i^{0,n}, \delta) + \tilde{Z}_i^n(0).$$

The latter inequality and the fact that the law of $\hat{S}_i^{0,n}$ does not depend on i gives

$$\max_{i \in [n]} \mathbb{P}(w_t(\hat{S}_i^n, \delta) > \varepsilon) \leq \mathbb{P}\left(w_t(\hat{S}_1^{0,n}, \delta) > \frac{\varepsilon}{2}\right) + \max_{i \in [n]} \mathbb{P}\left(\tilde{Z}_i^n(0) > \frac{\varepsilon}{2}\right).$$

Using (2.13) and the C -tightness of $\hat{S}_i^{0,n}$, $n \in \mathbb{N}$, shows that (3.28) is satisfied by \hat{S}_i^n .

Next, under our second moment assumptions on the inter-renewal times, it is well known that the rescaled nondelayed renewal processes satisfy

$$(3.30) \quad \sup_n \mathbb{E}[(\|\hat{S}_i^{0,n}\|_t^*)^2] < \infty,$$

for every t ([28], Appendix 1). Since the law of $\hat{S}_i^{0,n}$ does not depend on i , it follows from (3.29) and our assumption (2.14) that \hat{S}_i^n satisfy (3.27).

It follows now that \hat{E}_i^n satisfies both estimates, for the law of E_i^n is merely a special case of the law of the nondelayed renewal processes $S_i^{0,n}$.

As for the processes A_i^n , recall from Lemma 3.3(iii) that $\hat{M}_i^{A,n}$ is a martingale and that $[\hat{M}_i^{A,n}](t) = n^{-1}A_i^n(t)$. One has

$$\mathbb{E}[n^{-1}A_i^n(t)] = n^{-1}\lambda_0^n \int_0^t \mathbb{E}[p_{n,\mathcal{R}_i^n}(s)] ds \leq cn^{-1}n^{3/2}n^{-1}t,$$

where we used (2.5) and (2.9). Since this bound does not depend on i , it follows by the BDG inequality that

$$(3.31) \quad \lim_n \max_{i \in [n]} \mathbb{E}[(\|\hat{M}_i^{A,n}\|_t^*)^2] = 0.$$

Again by the bound on $p_{n,r}$, the Lipschitz constant of $\hat{C}_i^{A,n}$ is bounded by $cn^{3/2}n^{-1/2}n^{-1} = c$. It follows that both estimates (3.27) and (3.28) are satisfied by \hat{A}_i^n .

To treat \hat{L}_i^n , recall Skorohod’s lemma ([16], Section 8), stating that for a trajectory $y \in D(\mathbb{R}_+, \mathbb{R})$, if $(x, z) \in D(\mathbb{R}_+, \mathbb{R}_+)^2$ are such that $x = y + z$, z is nondecreasing and, with the convention $z(0-) = 0$, $\int_{[0,\infty)} x dz = 0$, then z is given by

$$z(t) = \sup_{s \in [0,t]} y^-(s), \quad t \geq 0.$$

In particular,

$$(3.32) \quad z(t) \leq \|y\|_t, \quad w_t(z, \delta) \leq w_t(y, \delta), \quad t > 0, \delta > 0.$$

In view of (3.23) and (3.24), it follows that

$$(3.33) \quad \hat{L}_i^n(t) \leq \|\hat{U}_i^n\|_t, \quad w_t(\hat{L}_i^n, \delta) \leq w_t(\hat{U}_i^n, \delta),$$

where

$$(3.34) \quad \hat{U}_i^n(t) = \hat{X}_i^n(0-) + \hat{E}_i^n(t) + \hat{A}_i^n(t) - \hat{S}_i^n(T_i^n(t)) + \hat{b}_1^n t,$$

and

$$(3.35) \quad \hat{X}_i^n = \hat{U}_i^n + \hat{L}_i^n.$$

Because T_i^n are 1-Lipschitz, we have that $\hat{S}_i^n(T_i^n(\cdot))$ satisfy (3.27) and (3.28). Hence by the second moment assumption on initial condition (2.15) and the boundedness of \hat{b}_1^n , the same holds for \hat{U}_i^n . Finally, by (3.33) and (3.35), this is also true for \hat{L}_i^n and \hat{X}_i^n . \square

3.4. Empirical process tightness.

LEMMA 3.6. $\bar{\xi}^n$ are C -tight in $D(\mathbb{R}_+, \mathcal{M}_1)$.

PROOF. Denote by C^ε the ε -neighborhood, in \mathbb{R}_+ , of a set $C \in \mathbb{R}_+$, and let

$$d_L(p, q) = \inf\{\varepsilon > 0 : p(C^\varepsilon) + \varepsilon \geq q(C) \text{ and } q(C^\varepsilon) + \varepsilon \geq p(C) \text{ for all } C \in \mathcal{B}(\mathbb{R}_+)\},$$

$p, q \in \mathcal{M}_1$, denote the Levy–Prohorov metric, which induces the topology of weak convergence on \mathcal{M}_1 . Since $\bar{\xi}_0^n$ converge in probability, proving C -tightness of $\bar{\xi}^n$ amounts to showing that for every $\varepsilon > 0$ and $\eta > 0$ there exists $\delta > 0$ such that

$$(3.36) \quad \limsup_n \mathbb{P}(w_T(\bar{\xi}^n, \delta) > \varepsilon) < \eta.$$

Here and below we use the same notation w_T for the metric d_L and for the usual metric on \mathbb{R} .

We show that (3.36) is a consequence of Lemma 3.5(iii). Fix $\varepsilon > 0$ and $\eta > 0$. Given n and $\delta > 0$, consider the event $\Omega_\delta^n = \{w_T(\bar{\xi}^n, \delta) > \varepsilon\}$. On this event there are $0 \leq s \leq t \leq T$, $t - s \leq \delta$, and $C \in \mathcal{B}(\mathbb{R}_+)$, such that

$$\xi_i^n(C^\varepsilon) + n\varepsilon < \xi_s^n(C) \quad \text{or} \quad \xi_s^n(C^\varepsilon) + n\varepsilon < \xi_i^n(C).$$

In both cases, the number of trajectories \hat{X}_i^n whose displacement between times s and t exceeds ε is greater than $n\varepsilon$. Therefore

$$\#\{i \in [n] : w_T(\hat{X}_i^n, \delta) \geq \varepsilon\} > n\varepsilon.$$

Hence, by Chebychev’s inequality,

$$\mathbb{P}(\Omega_\delta^n) \leq (n\varepsilon)^{-1} n \max_i \mathbb{P}(w_T(\hat{X}_i^n, \delta) > \varepsilon).$$

In view of (3.28), given any $\eta_1 > 0$ there is $\delta > 0$ such that for all large n , the maximum over i in the above display is $< \eta_1$. Hence, for such δ , the above is $\leq \varepsilon^{-1} \eta_1$. Choosing η_1 such that $\varepsilon^{-1} \eta_1 < \eta$ and the corresponding δ gives (3.36). \square

4. Proof. In Section 4.1 we derive an equation for the empirical measure, namely (4.4) below, that is a precursor to the parabolic PDE. That is, taking limits in (4.4) gives rise to a weak solution to equation (2.19). Specifically, Section 4.1 and, respectively, Section 4.2, show that the linear and nonlinear terms in (4.4) converge to the corresponding terms in (2.19). In Section 4.3, (2.17) is obtained from (2.19). Based on this convergence, the proof of all the main results is then given in Section 4.4.

4.1. *An equation for the empirical process.* Fix a test function $\tilde{\phi}$ as in Definition 2.5, that is, $\tilde{\phi} \in C_0^\infty(\mathbb{R}_+)$, $\tilde{\phi}(0) = 0$. Let

$$(4.1) \quad \phi(x) = \int_0^x \tilde{\phi}(y) dy, \quad x \in \mathbb{R}_+.$$

Then $\phi \in C_b^\infty(\mathbb{R}_+)$ and $\phi(0) = \phi'(0) = 0$. Moreover, ϕ, ϕ' and ϕ'' are uniformly continuous on \mathbb{R}_+ . Apply ϕ to the dynamics (3.26), noting that, unlike $X_i^n, X_i^{1,n}$ may assume negative values. On the r.h.s. of (3.26), the terms \hat{A}_i^n and $\hat{M}_i^{\text{dep},n}$ are piecewise constant. Thus

$$(4.2) \quad \begin{aligned} \phi(\hat{X}_i^{1,n}(t)) &= \phi(\hat{X}_i^n(0-)) + \int_0^t \phi'(\hat{X}_i^{1,n}(s))(d\hat{E}_i^{n,c}(s) + \hat{b}_1^n ds) + e_i^{2,n}(t) \\ &+ \sum_{s \leq t} (\phi(\hat{X}_i^{1,n}(s)) - \phi(\hat{X}_i^{1,n}(s-))), \end{aligned}$$

where $\hat{E}_i^{n,c}$ is the continuous part of \hat{E}_i^n ,

$$e_i^{2,n}(t) = \int_0^t (\phi'(\hat{X}_i^{1,n}(s)) - \phi'(\hat{X}_i^n(s))) d\hat{L}_i^n(s),$$

and we have used (3.24) and the continuity of the sample paths of \hat{L}_i^n to write

$$\int \phi'(\hat{X}_i^n(s)) d\hat{L}_i^n(s) = \phi'(0)\hat{L}_i^n(t) = 0.$$

In the last term of (4.2), by Taylor’s expansion, jumps according to \hat{E}_i^n can be expressed as

$$\phi'(\hat{X}_i^{1,n}(s-))\Delta\hat{E}_i^n(s) + \frac{1}{2}\phi''(\hat{X}_i^{2,n}(s))\Delta\hat{E}_i^n(s)^2,$$

where $\hat{X}_i^{2,n}(s)$ is an intermediate value between $\hat{X}_i^{1,n}(s-)$ and $\hat{X}_i^{1,n}(s)$ (we leave unspecified the value of the processes $\hat{X}_i^{2,n}$ away from times of jumps of $\hat{X}_i^{1,n}$). Similarly, jumps according to $\hat{M}_i^{\text{dep},n}$ are expressed as

$$-\phi'(\hat{X}_i^{1,n}(s-))\Delta\hat{M}_i^{\text{dep},n}(s) + \frac{1}{2}\phi''(\hat{X}_i^{2,n}(s))\Delta\hat{M}_i^{\text{dep},n}(s)^2,$$

where again $\hat{X}_i^{2,n}(s)$ are intermediate points, and because, a.s., jumps of \hat{E}_i^n and $\hat{M}_i^{\text{dep},n}$ do not occur simultaneously (for the same i), we may express the intermediate values by the same process. Finally, jumps according to \hat{A}_i^n —here we only need first order approximation—are expressed as

$$\phi'(\hat{X}_i^{2,n}(s))\Delta\hat{A}_i^n(s).$$

Note that jumps of \hat{E}_i^n and \hat{A}_i^n are of size $n^{-1/2}$. Hence in all three cases,

$$(4.3) \quad |\hat{X}_i^{2,n}(t) - \hat{X}_i^{1,n}(t-)| \leq |\Delta\hat{X}_i^{1,n}(t)| \leq \max\{|\Delta\hat{M}_i^{\text{dep},n}(t)|, n^{-1/2}\}, \quad t \in \mathcal{J}_i^n,$$

where \mathcal{J}_i^n is the set of jump times of $\hat{X}_i^{1,n}$.

By (4.2), using $\hat{E}_i^n = \hat{E}_i^{n,c} + \Delta\hat{E}_i^n$ and

$$\int_0^t \phi'(\hat{X}_i^{1,n}(s)) d\hat{E}_i^{n,c}(s) = \int_0^t \phi'(\hat{X}_i^{1,n}(s-)) d\hat{E}_i^{n,c}(s),$$

we have

$$(4.4) \quad \langle \phi, \bar{\xi}_t^n \rangle = \langle \phi, \bar{\xi}_{0-}^n \rangle + \int_0^t \left\langle b_1 \phi' + \frac{\sigma^2}{2} \phi'', \bar{\xi}_s^n \right\rangle ds + \Gamma^n(t) + \sum_{j=1}^6 f^{j,n}(t),$$

where

$$(4.5) \quad \Gamma^n(t) = \frac{b_0}{n} \sum_i \int_0^t \phi'(\hat{X}_i^n(s)) \left(\frac{n - \mathcal{R}_i^n(s)}{n} \right)^{\ell-1} ds$$

is the interaction term,

$$\begin{aligned} f^{1,n}(t) &= f_1^{1,n}(t) + f_2^{1,n}(t) \\ &:= \frac{1}{n} \sum_i \int_{[0,t]} \phi'(\hat{X}_i^{1,n}(s-)) d\hat{E}_i^n(s) - \frac{1}{n} \sum_i \int_{[0,t]} \phi'(\hat{X}_i^{1,n}(s-)) d\hat{M}_i^{\text{dep},n}(s), \end{aligned}$$

$$f^{2,n}(t) = \langle \phi, \bar{\xi}_t^n \rangle - \langle \phi, \bar{\xi}_t^{1,n} \rangle,$$

$$f^{3,n}(t) = \frac{1}{n} \sum_i e_i^{2,n}(t),$$

$$f^{4,n}(t) = \hat{b}_1^n \int_0^t \langle \phi', \bar{\xi}_s^{1,n} \rangle ds - b_1 \int_0^t \langle \phi', \bar{\xi}_s^n \rangle ds,$$

$$f^{5,n}(t) = \frac{1}{2n} \sum_i \int_{[0,t]} \phi''(\hat{X}_i^{2,n}(s)) \{d[\hat{E}_i^n](s) + d[\hat{M}_i^{\text{dep},n}](s)\} - \frac{\sigma^2}{2} \int_0^t \langle \phi'', \bar{\xi}_s^n \rangle ds,$$

$$f^{6,n}(t) = \frac{1}{n} \sum_i \int_{[0,t]} \phi'(\hat{X}_i^{2,n}(s)) d\hat{A}_i^n(s) - \Gamma^n(t)$$

are “error” terms, and

$$\bar{\xi}_t^{1,n} = n^{-1} \sum_{i \in [n]} \delta_{\hat{X}_i^{n,1}(t)}.$$

Note that the terms $f^{5,n}$ and $f^{6,n}$, which involve $\hat{X}_i^{2,n}$, are well defined, for their evaluation requires the values of $\hat{X}_i^{2,n}(t)$ only at \mathcal{J}_i^n .

LEMMA 4.1. *With ϕ as in (4.1), for $1 \leq j \leq 6$, $f^{j,n} \rightarrow 0$ in probability in $D(\mathbb{R}_+, \mathbb{R})$.*

PROOF. *Step 1.* Estimating $f^{1,n}$. The cross variation between independent Poisson processes is zero, hence $[\hat{E}_i^n, \hat{E}_j^n](t) = 0$ for $i \neq j$. Hence by Lemma 3.3,

$$[f_1^{1,n}](t) \leq cn^{-2} \sum_i [\hat{E}_i^n](t) \leq cn^{-3} \sum_i E_i^n(t).$$

Since each E_i^n is a Poisson process of intensity $\leq cn$, $\mathbb{E}\{[f_1^{1,n}](t)\} \leq ctn^{-1}$, and it follows by the BDG inequality that $f_1^{1,n} \rightarrow 0$ in probability. As for $f_2^{1,n}$, we have by Lemma 3.3(ii) that $[\hat{M}_i^{\text{dep},n}, \hat{M}_j^{\text{dep},n}](t)$, $i \neq j$, has zero mean. Hence by Lemma 3.3(i),

$$\mathbb{E}\{[f_2^{1,n}](t)\} \leq cn^{-2} \sum_i \mathbb{E}\{[\hat{M}_i^{\text{dep},n}](t)\} \leq cn^{-3} \sum_i \mathbb{E}\{D_i^n(t)\} \leq cn^{-1}t.$$

Thus $f_2^{1,n} \rightarrow 0$ in probability, and we conclude that $f^{1,n} \rightarrow 0$ in probability.

Step 2. We show that for all $t > 0$ and $\varepsilon > 0$,

$$(4.6) \quad \lim_{n \rightarrow \infty} \max_{i \in [n]} \mathbb{P}(\|e_i^{1,n}\|_t^* > \varepsilon) = 0.$$

In accordance with (2.13) we denote $\tilde{Z}_i^n(k) = \tilde{\mu}^n Z_i^n(k)$, $k \geq 1$. By the definition of $e_i^{1,n}(t)$ and $T_i^n(t) \leq t$,

$$(4.7) \quad |e_i^{1,n}(t)| \leq \tilde{Z}_i^n(0) \vee \|\tilde{R}_i^n\|_t^* \leq \max\{\tilde{Z}_i^n(k) : 0 \leq k \leq S_i^n(t)\} \\ \leq \max\{\tilde{Z}_i^n(k) : 0 \leq k \leq S_i^{0,n}(t)\}.$$

Thus, by (2.13) and the fact that the law of $(\tilde{Z}_i^n(k), k \geq 1)$, $S_i^{0,n}(t)$ does not depend on i , it suffices to prove that

$$\tilde{Y}^n := \max\{\tilde{Z}_1^n(k) : 1 \leq k \leq S_1^{0,n}(t)\} \rightarrow 0 \quad \text{in probability.}$$

This can be argued via the C -tightness of $\{\hat{S}_1^{0,n}\}$ as follows. Given n and $t_1 < t_2$, we have

$$(4.8) \quad \text{if } S_1^{0,n}(t_1) = S_1^{0,n}(t_2) \text{ then } \hat{S}_1^{0,n}(t_1) - \hat{S}_1^{0,n}(t_2) = \tilde{\mu}^n(t_2 - t_1).$$

This shows that

$$\max\{Z_1^n(k) : 1 \leq k \leq S_1^{0,n}(t) - 1\} \leq 2(\tilde{\mu}^n)^{-1} \|\hat{S}_1^{0,n}\|_t^*.$$

To include also $k = S_1^{0,n}(t)$, let the residual time process $R_1^{0,n}$ be defined analogously to R_1^n , for $S_1^{0,n}$ in place of S_1^n . Note that if $R_1^{0,n}(t) > 1$ holds then (4.8) holds with $t_1 = t$ and

$t_2 = t + 1$ hence $\hat{S}_1^{0,n}(t) - \hat{S}_1^{0,n}(t + 1) = \tilde{\mu}^n$. Because $\tilde{\mu}^n \rightarrow \infty$ and $\|\hat{S}_1^{0,n}\|_{t+1}^*$ are tight, this shows that w.h.p., $R_1^{0,n}(t) \leq 1$. As a result, again by (4.8), w.h.p.,

$$Y^n := \max\{Z_1^n(k) : 1 \leq k \leq S_1^{0,n}(t)\} \leq 2(\tilde{\mu}^n)^{-1} \|\hat{S}_1^{0,n}\|_{t+1}^*.$$

This shows that $Y^n \rightarrow 0$ in probability. Next, given $\delta > 0$, on the event $Y^n < \delta$, we have by (4.8),

$$\begin{aligned} w_{t+1}(\hat{S}_1^{0,n}, \delta) &\geq \sup\{|\hat{S}_1^{0,n}(t_1) - \hat{S}_1^{0,n}(t_2)| : 0 \leq t_1 < t_2 \leq t + 1, S_1^{0,n}(t_1) = S_1^{0,n}(t_2)\} \\ &\geq \tilde{\mu}^n \sup\{t_2 - t_1 : 0 \leq t_1 < t_2 \leq t + 1, S_1^{0,n}(t_1) = S_1^{0,n}(t_2)\} \\ &\geq \tilde{\mu}^n Y^n. \end{aligned}$$

Hence, w.h.p., $\tilde{Y}^n = \tilde{\mu}^n Y^n \leq w_{t+1}(\hat{S}_1^{0,n}, \delta)$. Using the C -tightness of $\hat{S}_1^{0,n}$, sending $n \rightarrow \infty$ and then $\delta \rightarrow 0$ shows that $\tilde{Y}^n \rightarrow 0$ in probability. This proves (4.6).

Step 3. We can now control $f^{2,n}$ and $f^{4,n}$. Recall (3.25). Then, denoting by $m_\phi(\cdot)$ the modulus of continuity of ϕ , for any $\delta > 0$,

$$\|f^{2,n}\|_t^* \leq \frac{1}{n} \sum_i \|\phi(\hat{X}_i^n) - \phi(\hat{X}_i^{1,n})\|_t^* \leq m_\phi(\delta) + \frac{2\|\phi\|_\infty}{n} \#\{i : \|e_i^{1,n}\|_t^* > \delta\}.$$

Given $\varepsilon > 0$ let $\delta > 0$ be such that $m_\phi(\delta) < \varepsilon/2$. Then

$$\mathbb{P}(\|f^{2,n}\|_t^* > \varepsilon) \leq 4\|\phi\|_\infty \varepsilon^{-1} \max_i \mathbb{P}(\|e_i^{1,n}\|_t^* > \delta).$$

Sending $n \rightarrow \infty$, the above converges to 0 by (4.6), proving that $f^{2,n} \rightarrow 0$ in probability.

Next, similarly, $\langle \phi', \bar{\xi}^{1,n} \rangle - \langle \phi', \xi^n \rangle \rightarrow 0$ in probability, and

$$\|f^{4,n}\|_t^* \leq t|\hat{b}_1^n| \|\langle \phi', \bar{\xi}^{1,n} \rangle - \langle \phi', \xi^n \rangle\|_t^* + |\hat{b}_1^n - b_1| \int_0^t \langle \phi', \bar{\xi}_s^n \rangle ds.$$

The first term on the right converges to 0 in probability by the argument shown for $f^{2,n}$ and the boundedness of \hat{b}_1^n . In the second term, the integral is bounded by $t\|\phi'\|_\infty$, hence the convergence of this term to zero follows from $\hat{b}_1^n \rightarrow b_1$. This proves $f^{4,n} \rightarrow 0$ in probability.

Step 4. To estimate $f^{3,n}$, note that, for any $\delta > 0$,

$$\|f^{3,n}\|_t^* \leq \frac{m_{\phi'}(\delta)}{n} \sum_i \hat{L}_i^n(t) + \frac{2\|\phi'\|_\infty}{n} \sum_i 1_{\{\|e_i^{1,n}\|_t^* > \delta\}} \hat{L}_i^n(t).$$

Hence

$$\mathbb{E}[\|f^{3,n}\|_t^*] \leq \frac{m_{\phi'}(\delta)}{n} \sum_i \mathbb{E}[\hat{L}_i^n(t)] + \frac{2\|\phi'\|_\infty}{n} \sum_i \mathbb{P}(\|e_i^{1,n}\|_t^* > \delta)^{1/2} (\mathbb{E}[\hat{L}_i^n(t)^2])^{1/2}.$$

Since by Lemma 3.5(iii), $\max_{i \in [n]} \mathbb{E}[\hat{L}_i^n(t)^2] < c$, it follows that

$$\mathbb{E}[\|f^{3,n}\|_t^*] \leq cm_{\phi'}(\delta) + c \max_{i \in [n]} \mathbb{P}(\|e_i^{1,n}\|_t^* > \delta).$$

In view of (4.6), if we take $n \rightarrow \infty$ and then $\delta \rightarrow 0$, the expression on the right converges to 0, hence $f^{3,n} \rightarrow 0$ in probability.

Step 5. Recall

$$f^{5,n}(t) = \frac{1}{2n} \sum_i \int_{[0,t]} \phi''(\hat{X}_i^{2,n}(s)) \{d[\hat{E}_i^n](s) + d[\hat{M}_i^{\text{dep},n}](s)\} - \frac{\sigma^2}{2} \int_0^t \langle \phi'', \bar{\xi}_s^n \rangle ds,$$

$$[\hat{E}_i^n](t) = n^{-1} E_i^n(t), \quad [\hat{M}_i^{\text{dep},n}](t) = n^{-1} \sum_{k=1}^{D_i^n(t)} \zeta_i^n(k)^2.$$

Let

$$\begin{aligned}
 f_1^{5,n} &= \frac{1}{n} \sum_i \int_{[0,t]} \phi''(\hat{X}_i^{1,n}(s-)) d[\hat{E}_i^n](s) - \lambda \int_0^t \langle \phi'', \bar{\xi}_s^{1,n} \rangle ds, \\
 f_2^{5,n} &= \frac{1}{n} \sum_i \int_{[0,t]} \phi''(\hat{X}_i^{1,n}(s-)) d[\hat{M}_i^{\text{dep},n}](s) - \lambda \sigma_{\text{ser}}^2 \int_0^t \langle \phi'', \bar{\xi}_s^{1,n} \rangle ds, \\
 f_3^{5,n} &= \int_0^t \langle \phi'', \bar{\xi}_s^{1,n} \rangle ds - \int_0^t \langle \phi'', \bar{\xi}_s^n \rangle ds \\
 f_4^{5,n} &= \frac{1}{n} \sum_i \int_{[0,t]} (\phi''(\hat{X}_i^{2,n}(s)) - \phi''(\hat{X}_i^{1,n}(s-))) \{d[\hat{E}_i^n](s) + d[\hat{M}_i^{\text{dep},n}](s)\}.
 \end{aligned}$$

Then, recalling that $\sigma^2 = \lambda(1 + \sigma_{\text{ser}}^2)$, we have $f^{5,n} = \frac{1}{2}(f_1^{5,n} + f_2^{5,n} + \sigma^2 f_3^{5,n} + f_4^{5,n})$. For $f_1^{5,n}$, write

$$\begin{aligned}
 f_1^{5,n} &= \frac{1}{n} \sum_i \int_{[0,t]} \phi''(\hat{X}_i^{1,n}(s-)) n^{-1} dE_i^n(s) - \lambda \int_0^t \langle \phi'', \bar{\xi}_s^{1,n} \rangle ds \\
 &= \frac{1}{n} \sum_i \int_{[0,t]} \phi''(\hat{X}_i^{1,n}(s-)) n^{-1/2} d\hat{E}_i^n(s) + \frac{1}{n} \sum_i \int_0^t \phi''(\hat{X}_i^{1,n}(s)) n^{-1/2} \hat{\lambda}^n ds.
 \end{aligned}$$

In the first sum, each term is a martingale having expected quadratic variation bounded by $cn^{-2}\lambda^n t \leq cn^{-1}t$ where we use Lemma 3.3(i), and c does not depend on i or n . Hence the first normalized sum is a martingale whose expected quadratic variation is $\leq cn^{-2}t$. In particular, it converges to zero in probability. In the second sum, each term is bounded in absolute value by $ctn^{-1/2}$, hence the second normalized sum converges to zero. This shows $f_1^{5,n} \rightarrow 0$ in probability.

Using Lemma 3.3 for the expression for $[\hat{M}_i^{\text{dep},n}]$ and denoting

$$P_i^n(t) = \phi''(\hat{X}_i^{1,n}(t)), \quad Q_i^n(t) = \sum_{k=1}^{D_i^n(t)} q_i^n(k), \quad q_i^n(k) = \zeta_i^n(k)^2 - \sigma_{\text{ser}}^2,$$

write $f_2^{5,n}$ as

$$(4.9) \quad f_2^{5,n} = \frac{1}{n} \sum_i \int_{[0,t]} P_i^n(s-) n^{-1} dQ_i^n(s) + \frac{\sigma_{\text{ser}}^2}{n} \sum_i \int_{[0,t]} P_i^n(s-) \{n^{-1} dD_i^n(s) - \lambda ds\}.$$

To estimate the first term above, note that, by (3.10),

$$D_i^n(t) = S_i^n(T_i^n(t)) \leq S_i^n(t) = n^{1/2} \hat{S}_i^n(t) + \mu^n t \leq n^{1/2} \hat{S}_i^n(t) + c_1 nt,$$

for a constant c_1 . For any $\beta > 2c_1 t$,

$$\max_{i \in [n]} \mathbb{P}(D_i^n(t) > \beta n) \leq \max_{i \in [n]} \mathbb{P}(\hat{S}_i^n(t) > n^{1/2} \beta / 2) \leq cn^{-1} \beta^{-2},$$

by Lemma 3.5(iii), where $c = c(t)$. Denote by $0 \leq t_i^n(1) < t_i^n(2) < \dots$ the jump times of D_i^n and let $m_i^n(N) = \sum_{k=1}^N P_i^n(t_i^n(k)-) q_i^n(k)$. Then the first term in (4.9) is given by $n^{-2} \sum_i m_i^n(D_i^n(t))$, which we write as

$$(4.10) \quad \frac{1}{n^2} \sum_i m_i^n(D_i^n(t) \wedge (\beta n)) + \frac{1}{n^2} \sum_i \{m_i^n(D_i^n(t)) - m_i^n(D_i^n(t) \wedge (\beta n))\}.$$

Now $m_i^n(\cdot)$ is a discrete parameter martingale with increments bounded by $\|\phi''\|_\infty |q_i^n(k)|$. Moreover, $q_i^n(k)$ are i.i.d. in k and i , and possesses a first moment in view of our second moment assumption on the service times. Hence Lemma 3.4 is applicable, showing that

$$N^{-1} \mathbb{E} \|m_i^n\|_N^* < A(N) \rightarrow 0 \quad \text{as } N \rightarrow \infty,$$

where $A(\cdot)$ does not depend on n or i . Hence

$$\mathbb{E} \left| \frac{1}{n^2} \sum_i m_i^n(D_i^n(t) \wedge (\beta n)) \right| \leq \frac{1}{n} \sum_i n^{-1} \mathbb{E} [\|m_i^n\|_{\beta n}^*] \leq n^{-1} \beta n A(\beta n) = \beta A(\beta n),$$

where throughout this paragraph βn should be read as $\lfloor \beta n \rfloor$. Hence the limit of the above expression is 0 for any β . Next, using $D_i^n \leq S_i^n$,

$$\begin{aligned} \mathbb{E} \left| \frac{1}{n^2} \sum_i \{m_i^n(D_i^n(t) - m_i^n(D_i^n(t) \wedge (\beta n))\} \right| &\leq \|\phi''\|_\infty n^{-2} \sum_i \mathbb{1}_{\{S_i^n(t) > \beta n\}} \sum_{k=\beta n+1}^{S_i^n(t)} |q_i^n(k)| \\ &= cn^{-2} \sum_i \sum_{k=1}^{(S_i^n(t) - \beta n)^+} |q_i^n(k - \beta n)| \\ &= cn^{-2} \sum_i \mathbb{E}[(S_i^n(t) - \beta n)^+] \mathbb{E}[|q_i^n(1)|] \\ &\leq cn^{-1} \max_{i \in [n]} \mathbb{E}[(S_i^n(t) - \beta n)^+], \end{aligned}$$

where Wald’s identity is used on the third line. Since it follows from Lemma 3.5(iii) that $n^{-1} S_i^n(t)$ are uniformly integrable in i and n ,

$$\lim_{\beta \rightarrow \infty} \limsup_n n^{-1} \max_{i \in [n]} \mathbb{E}[(S_i^n(t) - \beta n)^+] = 0.$$

As a result, the expression in (4.10) converges to zero in probability.

As for the second term in (4.9). In view of Lemma 3.5(iii) and (2.8), one has

$$\max_{i \in [n]} \|n^{-1} S_i^n - \lambda t\|_t^* \rightarrow 0 \quad \text{in probability.}$$

Also by Lemma 3.5(iii) and (2.12), which relates \hat{L}_i^n to T_i^n , $\max_{i \in [n]} \|T_i^n - t\|_t^* \rightarrow 0$ in probability. Hence $D_i^n = S_i^n(T_i^n)$ also satisfies $\kappa_n := \max_{i \in [n]} \|n^{-1} D_i^n - \lambda t\|_t^* \rightarrow 0$ in probability.

Denote $\varepsilon = t/j$, $s_k = k\varepsilon$ and $\tilde{P}_i^n(s) = P_i^n(s-)$. Then

$$\begin{aligned} \int_{(0,t]} \tilde{P}_i^n(s) \{n^{-1} dD_i^n(s) - \lambda ds\} &= \sum_{k=0}^{j-1} \left\{ \int_{(s_k, s_{k+1}]} (\tilde{P}_i^n(s) - \tilde{P}_i^n(s_k)) \{n^{-1} dD_i^n(s) - \lambda ds\} \right. \\ &\quad \left. + \tilde{P}_i^n(s_k) \{n^{-1} D_i^n(s_{k+1}) - n^{-1} D_i^n(s_k) - \lambda \varepsilon\} \right\}. \end{aligned}$$

This gives

$$\begin{aligned} (4.11) \quad &\left| \int_{(0,t]} \tilde{P}_i^n(s) \{n^{-1} dD_i^n(s) - \lambda ds\} \right| \\ &\leq w_T(\tilde{P}_i^n, \varepsilon) (n^{-1} D_i^n(t) + \lambda t) + \|\tilde{P}_i^n\|_t^* 2j\kappa_n \\ &\leq w_T(\tilde{P}_i^n, \varepsilon) (2\lambda t + \kappa_n) + \|\phi''\|_\infty 2j\kappa_n. \end{aligned}$$

We recall from Lemma 3.5(iii) that \hat{X}_i^n are uniformly C -tight. Along with the uniform estimate (4.6) on $e_i^{1,n}$, this shows that so are \tilde{P}_i^n . In particular, these processes satisfy (3.28).

Hence, on sending $n \rightarrow \infty$ and then $\varepsilon \rightarrow 0$, it follows that the right side of (4.11) converges to zero in probability, uniformly in $i \in [n]$. Hence the second term in $f_2^{5,n}$ converges to zero in probability, and we conclude that $f_2^{5,n} \rightarrow 0$ in probability.

Next, recalling the result regarding $f_2^{5,n}$ and replacing ϕ by ϕ'' shows that $f_3^{5,n} \rightarrow 0$ in probability.

To bound $f_4^{5,n}$, fix $\varepsilon > 0$ and let $\delta > 0$ be such that $|x - y| < \delta$ implies $|\phi''(x) - \phi''(y)| < \varepsilon$. Let $\mathcal{J}_i^{n,\delta} = \{s \in \mathcal{J}_i^n : |\hat{X}_i^{2,n}(s) - \hat{X}_i^{1,n}(s-)| > \delta\}$. For n sufficiently large, the size of jumps of \hat{E}^n (which is $n^{-1/2}$) is smaller than δ , hence, by (4.3), the corresponding jump times are not members of $\mathcal{J}_i^{n,\delta}$. Thus the i th term in $f_4^{5,n}$ is bounded, in absolute value, by

$$\begin{aligned} &\varepsilon([\hat{E}_i^n](t) + [\hat{M}_i^{\text{dep},n}](t)) + 2\|\phi''\|_\infty \sum_{s \in \mathcal{J}_i^{n,\delta}, s \leq t} \Delta[\hat{M}_i^{\text{dep},n}](s) \\ &\leq \varepsilon n^{-1} \left(E_i^n(t) + \sum_{k=1}^{D_i^n(t)} \zeta_i^n(k)^2 \right) + cn^{-1} \sum_{k=1}^{D_i^n(t)} \zeta_i^n(k)^2 1_{\{n^{-1/2}|\zeta_i^n(k)| > \delta\}}, \end{aligned}$$

where we used the fact that for $s \in \mathcal{J}_i^{n,\delta}$, $|\Delta \hat{X}_i^{1,n}(s)| = |\Delta \hat{M}_i^{\text{dep},n}(s)|$ by (3.26). Using (3.15), the expected value of the above expression is

$$\leq \varepsilon n^{-1} (\lambda_i^n t + \sigma_{\text{ser}}^2 \mathbb{E}[D_i^n(t)]) + cn^{-1} \mathbb{E}[\zeta_1^n(1)^2 1_{\{n^{-1/2}|\zeta_1^n(1)| > \delta\}}] \mathbb{E}[D_i^n(t)],$$

where for the last term we used Wald’s identity in exactly the same way as in the proof of Lemma 3.3. In view of Lemma 3.5(iii), $\mathbb{E}[S_i^n(t)] < cn$, where c does not depend on i or n . Since $D_i^n(t) \leq S_i^n(t)$, it follows that

$$\mathbb{E}[|f_4^{5,n}|] \leq c\varepsilon + c\mathbb{E}[\zeta_1^n(1)^2 1_{\{n^{-1/2}|\zeta_1^n(1)| > \delta\}}].$$

Taking $n \rightarrow \infty$ then $\varepsilon \rightarrow 0$ shows that $f_4^{5,n} \rightarrow 0$ in probability.

Step 6. Finally,

$$\begin{aligned} f_6^{6,n}(t) &= \frac{1}{n} \sum_i \int_{[0,t]} [\phi'(\hat{X}_i^{2,n}(s)) - \phi'(\hat{X}_i^n(s))] d\hat{A}_i^n(s) \\ &\quad + \frac{1}{n} \sum_i \int_{[0,t]} \phi'(\hat{X}_i^n(s)) d\hat{A}_i^n(s) - \Gamma^n(t). \end{aligned}$$

By Lemma 3.3(iii),

$$\begin{aligned} \frac{1}{n} \sum_i \int_{[0,t]} \phi'(\hat{X}_i^n(s-)) d\hat{A}_i^n(s) &= \lambda_0^n n^{-3/2} \sum_i \int_0^t \phi'(\hat{X}_i^n(s)) p_{n,\mathcal{R}_i^n}(s) ds + \bar{M}^{A,n}(t), \\ \bar{M}^{A,n}(t) &= \frac{1}{n} \sum_i \int_{[0,t]} \phi'(\hat{X}_i^n(s-)) d\hat{M}_i^{A,n}(s). \end{aligned}$$

By Lemma 3.3(iii), $[\bar{M}^{A,n}](t) \leq cn^{-3} \sum_i A_i^n(t) = cn^{-3} A_0^n(t)$. Since A_0^n is a Poisson process of intensity $\lambda_0^n \leq cn^{3/2}$ (by (2.9)), it follows that $\bar{M}^{A,n} \rightarrow 0$ in probability.

Now, writing

$$p_{n,r} = \frac{\ell (n-r)(n-r-1) \cdots (n-r-\ell+2)}{n (n-1)(n-2) \cdots (n-\ell+1)}$$

shows that

$$(4.12) \quad p_{n,r} = \frac{\ell}{n} \left[\left(1 - \frac{r}{n}\right)^{\ell-1} + \alpha_{n,r} \right], \quad \alpha_n := \max_{r \in [n]} |\alpha_{n,r}| \rightarrow 0.$$

Hence,

$$\begin{aligned} & \lambda_0^n n^{-3/2} \sum_i \int_0^t \phi'(\hat{X}_i^n(s)) p_{n, \mathcal{R}_i^n(s)} ds - \Gamma^n(t) \\ &= \sum_i \int_0^t \phi'(\hat{X}_i^n(s)) \left[\lambda_0^n n^{-3/2} p_{n, \mathcal{R}_i^n(s)} - \frac{b_0}{n} \left(\frac{n - \mathcal{R}_i^n(s)}{n} \right)^{\ell-1} \right] ds \\ &= \sum_i \int_0^t \phi'(\hat{X}_i^n(s)) \left(\frac{n - \mathcal{R}_i^n(s)}{n} \right)^{\ell-1} \left[\lambda_0^n n^{-3/2} \frac{\ell}{n} - \frac{b_0}{n} \right] ds \\ & \quad + \sum_i \int_0^t \phi'(\hat{X}_i^n(s)) \lambda_0^n n^{-3/2} \frac{\ell}{n} \alpha_{n, \mathcal{R}_i^n(s)} ds. \end{aligned}$$

We have $\lambda_0^n n^{-3/2} \ell \rightarrow b\ell = b_0$, which shows that the first sum converges to 0. From the fact $\alpha_n \rightarrow 0$ we also have that the last sum converges to 0.

To bound the first term in $f^{6,n}$, note that \hat{A}_i^n are nondecreasing. Also, since the jumps of \hat{A}_i^n are of size $n^{-1/2}$, one has $|\phi'(\hat{X}_i^{2,n}(s)) - \phi'(\hat{X}_i^{1,n}(s))| \leq m_{\phi'}(n^{-1/2})$ at any jump time s of \hat{A}_i^n . Moreover, $\max_{i \in [n]} |\phi'(\hat{X}_i^{1,n}(s)) - \phi'(\hat{X}_i^n(s))| \leq \max_{i \in [n]} m_{\phi'}(\|e_i^{1,n}\|_i^*) \rightarrow 0$ in probability, by (4.6). Hence

$$\max_{i \in [n]} \left| \int_{[0,t]} [\phi'(\hat{X}_i^{2,n}(s)) - \phi'(\hat{X}_i^n(s))] d\hat{A}_i^n(s) \right| \leq \max_{i \in [n]} \{ (m_{\phi'}(n^{-1/2}) + m_{\phi'}(\|e_i^{1,n}\|_i^*)) \hat{A}_i^n(t) \}$$

converges to 0 in probability, where Lemma 3.5(iii) is used for a uniform estimate on $\hat{A}_i^n(t)$. It follows that $f^{6,n} \rightarrow 0$ in probability. \square

REMARK 4.2. One of the steps in the proof above can be simplified if stronger moment conditions are assumed. This is the treatment of the second order term $f_2^{5,n}$, which involves martingales given in terms of squares of the primitive data $\zeta_i^n(k)$. Provided that $\zeta_i^n(k)$ possess a fourth moment, the quadratic variation of the martingale comprising the first part of (4.9) has finite expectation, and so an estimate based on the BDG inequality would show that this term vanishes in the limit.

4.2. Interaction term under limit. In this subsection we prove the following.

LEMMA 4.3. Given a subsequence along which $\bar{\xi}^n \Rightarrow \xi$, one has $\Gamma^n \Rightarrow \Gamma$ in $C(\mathbb{R}_+, \mathbb{R})$ along the subsequence, where

$$\Gamma(t) = \frac{b_0}{\ell} \int_0^t \int_{\mathbb{R}_+} \phi'(x) \mathfrak{S}(\xi_s[x, \infty), \xi_s(x, \infty)) \xi_s(dx) ds.$$

Recall that if ξ_t is atomless for all $t > 0$ then the integrand in the above expression simplifies to $\ell \phi'(x) \xi_s[x, \infty)^{\ell-1}$, in which case Γ is directly related to the form of the interaction term in equations (1.3), (1.5) and (2.16). However, the atomless property has not been established at this stage of the proof.

PROOF. Invoking the Skorohod representation theorem, assume without loss of generality that $\bar{\xi}^n \rightarrow \xi$ a.s. Because $0 \leq \mathcal{R}_i^n \leq n$, the integrands in (4.5) are bounded by $\|\phi'\|_\infty$. Hence by bounded convergence, it suffices to prove that, a.s., for every t ,

$$(4.13) \quad \gamma^n(t) := \frac{1}{n} \sum_i \phi'(\hat{X}_i^n(t)) \bar{\mathcal{R}}_i^{c,n}(t)^{\ell-1} \rightarrow \frac{1}{\ell} \int_{\mathbb{R}_+} \phi'(x) \mathfrak{S}(\xi_t[x, \infty), \xi_t(x, \infty)) \xi_t(dx),$$

where

$$\mathcal{R}_i^{c,n}(t) = n - \mathcal{R}_i^n(t), \quad \bar{\mathcal{R}}_i^{c,n}(t) = n^{-1}\mathcal{R}_i^{c,n}(t).$$

Fix t and $\varepsilon > 0$. The function $x \mapsto \xi_t(x, \infty)$ has at most countably many discontinuities. Hence one can find a finite sequence $x_* = y_0 < y_1 < \dots < y_K = x^*$ such that $[x_*, x^*]$ contains the compact support of ϕ' (recall (4.1)), $y_k - y_{k-1} < \varepsilon, k = 1, \dots, K$, and moreover $\xi_t(\{y_k\}) = 0, k = 0, \dots, K$. Because y_k are not charged, both $\bar{\xi}_t^n(y_{k-1}, y_k)$ and $\bar{\xi}_t^n[y_{k-1}, y_k]$ converge a.s. to $\xi_t(y_{k-1}, y_k)$. Define by $\text{rank}^c(i; x) = n - \text{rank}(i; x), x \in \mathbb{R}^n, i \in [n]$ the complementary rank. Then, $0 \leq \text{rank}^c(i; x) \leq n - 1$, and by (2.3),

$$\text{rank}^c(i; x) = \#\{j : x_j > x_i\} + \#\{j > i : x_j = x_i\}.$$

Denote $x_i = \hat{X}_i^n(t), x = (x_1, \dots, x_n)$, and for $k = 1, \dots, K$,

$$I_k = [y_{k-1}, y_k), \quad V_k = \{i : x_i \in I_k\}, \quad P_k = \xi_t^n(I_k), \quad Q_k = \xi_t^n[y_k, \infty).$$

Fix k , and for $i \in V_k$ let $j(i) = \text{rank}(i; (x_l)_{l \in V_k})$. This is a relabeling of V_k according to the rank of its members within V_k . With this notation we have

$$\text{rank}^c(i; x) = Q_k + P_k - j(i), \quad i \in V_k.$$

Because $j(i)$ take all values between 1 and P_k as i varies in V_k , we have

$$\sum_{i \in V_k} \text{rank}^c(i; x)^{\ell-1} = \sum_{j=1}^{P_k} (Q_k + P_k - j)^{\ell-1} = \sum_{j=0}^{P_k-1} (Q_k + j)^{\ell-1}.$$

If p_n is a sequence s.t. $n^{-1}p_n \rightarrow p \geq 0$ then $n^{-1} \sum_{j=0}^{p_n} (j/n)^{\ell-1} \rightarrow \int_0^p z^{\ell-1} dz = \ell^{-1} p^\ell$. Noting, as mentioned, that $n^{-1}P_k \rightarrow \xi_t(I_k)$, we get

$$\begin{aligned} \frac{1}{n} \sum_{i \in V_k} \bar{\mathcal{R}}_i^{c,n}(t)^{\ell-1} &\rightarrow \frac{1}{\ell} (\xi_t(y_{k-1}, \infty)^\ell - \xi_t(y_k, \infty)^\ell) \\ (4.14) \qquad \qquad \qquad &= \frac{1}{\ell} \xi_t(I_k) \mathfrak{S}(\xi_t(y_{k-1}, \infty), \xi_t(y_k, \infty)). \end{aligned}$$

Fix a sequence $\varepsilon_m > 0, \varepsilon_m \rightarrow 0$. For each m , let points as the above y_k , corresponding to ε_m , be denoted by y_k^m . For $x \in [x_*, x^*]$ let $x_{(m)}$ and $x^{(m)}$ be the unique two points y_{k-1}^m and y_k^m , respectively, for which $x \in [y_{k-1}^m, y_k^m)$. This gives $x^{(m)} \rightarrow x, x^{(m)} > x$. By right-continuity of $x \mapsto \xi_t(x, \infty)$, it follows that $\xi_t(x^{(m)}, \infty) \rightarrow \xi_t(x, \infty)$ for every $x \in [x_*, x^*]$. Similarly, by left continuity of $x \mapsto \xi_t[x, \infty)$, and $x_{(m)} \rightarrow x, x_{(m)} \leq x$, we obtain $\xi_t(x_{(m)}, \infty) \rightarrow \xi_t[x, \infty)$. Let

$$h_{(m)}(x) = \inf\{\phi'(z) : z \in [x_{(m)}, x^{(m)}]\}, \quad h^{(m)}(x) = \sup\{\phi'(z) : z \in [x_{(m)}, x^{(m)}]\}.$$

Summing over k in (4.14), we have for every m ,

$$(4.15) \quad \liminf_n \gamma^n(t) \geq \frac{1}{\ell} \int_{[x_*, x^*]} h_{(m)}(x) \mathfrak{S}(\xi_t(x_{(m)}, \infty), \xi_t(x^{(m)}, \infty)) \xi_t(dx)$$

$$(4.16) \quad \limsup_n \gamma^n(t) \leq \frac{1}{\ell} \int_{[x_*, x^*]} h^{(m)}(x) \mathfrak{S}(\xi_t(x_{(m)}, \infty), \xi_t(x^{(m)}, \infty)) \xi_t(dx).$$

It remains to show the expressions of the r.h.s. of (4.15) and (4.16) both converge to the r.h.s. of (4.13) as $m \rightarrow \infty$. By bounded convergence, it suffices that the integrands in both these integrals converge, for every x , to $\phi'(x) \mathfrak{S}(\xi_t[x, \infty), \xi_t(x, \infty))$. The latter convergence holds because, by continuity of ϕ' , $h_{(m)}(x) \rightarrow \phi'(x)$ and $h^{(m)}(x) \rightarrow \phi'(x)$, and moreover $\xi_t(x_{(m)}, \infty) \rightarrow \xi_t[x, \infty)$ and $\xi_t(x^{(m)}, \infty) \rightarrow \xi_t(x, \infty)$. This completes the proof. \square

4.3. *PDE satisfied by the limit.* We take limits along an arbitrary convergent subsequence. By (4.4), (4.5), Lemma 4.1 and Lemma 4.3, we have now established that every limit point ξ satisfies (2.19).

Again, fix a convergent subsequence and denote its limit by ξ . Some notation used below is as follows. For a right-continuous nondecreasing function $V : \mathbb{R} \rightarrow [0, 1]$ with $V(0-) = 0$, let $V(dx)$ denote the Stieltjes measure induced on \mathbb{R}_+ . Let

$$v(x, t) = \xi_t(x, \infty), \quad \tilde{v}(x, t) = 1 - v(x, t) = \xi_t[0, x].$$

These are right continuous functions for every t . Thus $\tilde{v}(dx, t) = \xi_t(dx)$. Next, for V as above, let the pure jump part and continuous part be denoted, respectively, by

$$V^{\text{jmp}}(x) = \sum_{y \in [0, x]} \Delta V(y), \quad V^{\text{cts}}(x) = V(x) - V^{\text{jmp}}(x).$$

LEMMA 4.4. *v is a weak solution to (2.18).*

PROOF. By equations (4.4) and (4.5) and Lemmas 4.1 and 4.3,

$$\langle \phi, \xi_t \rangle = \langle \phi, \xi_0 \rangle + \int_0^t \langle b_1 \phi' + a \phi'', \xi_s \rangle ds + \Gamma(t).$$

Write $\Gamma(t)$ as $b_0 \ell^{-1} \int_0^t A(s) ds$ where

$$A(t) = \int \phi'(x) \mathfrak{S}(v(x-, t), v(x, t)) \tilde{v}(dx, t).$$

Let $\tilde{v}(x, t) = \tilde{v}^{\text{cts}}(x, t) + \tilde{v}^{\text{jmp}}(x, t)$ denote the decomposition alluded to above. Let also $U(x, t) = 1 - v(x, t)^\ell$. Then

$$U(dx, t) = \ell v(x, t)^{\ell-1} \tilde{v}(dx, t).$$

Hence

$$U^{\text{cts}}(dx, t) = \ell v(x, t)^{\ell-1} \tilde{v}^{\text{cts}}(dx, t), \quad U^{\text{jmp}}(dx, t) = \ell v(x, t)^{\ell-1} \tilde{v}^{\text{jmp}}(dx, t).$$

Now,

$$A(t) = \int \phi'(x) \ell v(x, t)^{\ell-1} \tilde{v}^{\text{cts}}(dx, t) + \int \phi'(x) \frac{v(x-, t)^\ell - v(x, t)^\ell}{v(x-, t) - v(x, t)} \tilde{v}^{\text{jmp}}(dx, t).$$

The second integral on the right is

$$\begin{aligned} \sum_{x \in [0, \infty)} \phi'(x) (v(x-, t)^\ell - v(x, t)^\ell) &= \sum_{x \in [0, \infty)} \phi'(x) (U(x, t) - U(x-, t)) \\ &= \int \phi'(x) U^{\text{jmp}}(dx, t). \end{aligned}$$

Therefore, using $\phi'(0) = 0$,

$$\begin{aligned} A(t) &= \int \phi'(x) (U^{\text{cts}}(dx, t) + U^{\text{jmp}}(dx, t)) = \int \phi'(x) U(dx, t) \\ &= - \int \phi''(x) U(x, t) dx = \int \phi''(x) v(x, t)^\ell dx. \end{aligned}$$

Using integration by parts we have $\langle \phi, \xi_t \rangle = \int_{\mathbb{R}_+} (v(x, t) - 1)\phi'(x) dx$. Recalling $\phi' = \tilde{\phi}$ and combining the above calculations, we obtain

$$\begin{aligned} \int \tilde{\phi}(x)(v(x, t) - 1) dx &= \int \tilde{\phi}(x)(v(x, 0) - 1) dx \\ &+ \int_0^t \int (b_1 \tilde{\phi}'(x) + a \tilde{\phi}''(x))(v(x, s) - 1) dx ds \\ &+ \frac{b_0}{\ell} \int_0^t \int \tilde{\phi}'(x)v(x, t)^\ell dx ds. \end{aligned}$$

Thus

$$\begin{aligned} \int \tilde{\phi}(x)v(x, t) dx &= \int \tilde{\phi}(x)v(x, 0) dx + \int_0^t \int (b_1 \tilde{\phi}'(x) + a \tilde{\phi}''(x))v(x, s) dx ds \\ &+ a \int_0^t \tilde{\phi}'(0) ds + \frac{b_0}{\ell} \int_0^t \int \tilde{\phi}'(x)v(x, t)^\ell dx ds. \end{aligned}$$

According to Definition 2.5, this shows that v is a weak solution of (2.18) once it is verified that $v \in \mathbb{L}^\infty_{loc}(\mathbb{R}_+, \mathbb{L}^1(\mathbb{R}_+))$.

This property can otherwise be stated as $\sup_{t \in (0, T]} \int_0^\infty \xi_t[x, \infty) dx < \infty$ a.s., for every $T < \infty$. Invoking Skorohod’s representation theorem we may assume w.l.o.g. that, along the chosen convergent subsequence, one has $\tilde{\xi}^n \rightarrow \xi$ a.s. Arguing by contradiction, let there exist a sequence $t_N \in (0, T]$ and an event of positive \mathbb{P} measure on which $\int_0^\infty \xi_{t_N}[x, \infty) dx \rightarrow \infty$. Now, for each fixed $t \in (0, T]$, by Fatou’s lemma,

$$\begin{aligned} \mathbb{E} \int_0^\infty \xi_t[x, \infty) dx &\leq \liminf_n \mathbb{E} \int_0^\infty \tilde{\xi}_t^n[x, \infty) dx \\ &\leq \liminf_n \frac{1}{n} \sum_{i \in [n]} \mathbb{E} \hat{X}_i^n(t) \\ &\leq \sup_n \max_{i \in [n]} \mathbb{E} \|\hat{X}_i^n\|_T^* \leq c < \infty, \end{aligned}$$

where $c = c(T)$ and the last assertion follows from Lemma 3.5(iii). This contradicts the assumption, and hence the uniform \mathbb{L}^1 property follows. \square

4.4. *Proof of main results.* We can now prove Theorems 2.3, 2.6, 2.7, 2.9, 2.10 and Propositions 2.4 and 2.8.

PROOF OF THEOREM 2.6. Having proved uniqueness in Lemma 3.2, we proceed to the remaining assertions. Existence of a $C^\infty(\mathbb{R}_+ \times (0, \infty))$ solution is well known; see, for example, [23], Lemma 3.2, p. 68.

In the case $f(z) = c_1 z - b z^\ell$ and $v_0(x) = \xi_0(x, \infty)$, we have proved in Lemma 4.4 that the function v induced by any subsequential limit ξ is a weak solution. By uniqueness it must be equal to the smooth function mentioned above. Since $x \mapsto v(x, t)$ is automatically nonincreasing, the function $u := -v_x$ is nonnegative and smooth. Using this in (2.17), an integration by parts gives the first 3 parts of (1.3) as well as that u integrates to 1 for each $t > 0$.

As for the initial condition stated as $u(\cdot, t) dx \rightarrow \xi_0(dx)$ in (1.3), again using (2.17) with $-v_x = u$ gives $\int \phi'(x)u(x, t) dx = \int \phi'(x)\xi_0(dx) + O(t)$ as $t \rightarrow 0$, which proves that the initial condition is satisfied. This shows that u thus defined forms a classical solution to (1.3). \square

PROOF OF PROPOSITION 2.4. The fact that u_{stat} given in (1.4) is a stationary solution is verified directly. As for uniqueness, for any solution u , the function $v = \int_0^\infty u(x) dx$ takes values in $[0, 1]$, is monotone, and satisfies

$$f(v)' + av'' = 0, \quad v(0) = 1, \quad \lim_{x \rightarrow \infty} v(x) = 0.$$

Integrating we obtain $f(v) + av' = c_1$, where c_1 is a constant. Since $f(0) = 0$, we obtain by the condition at infinity that $\lim_{x \rightarrow \infty} v'(x) = c_1/a$, which can only hold if $c_1 = 0$. Thus v must satisfy the ODE

$$v' = -a^{-1}f(v) \quad \text{on } \mathbb{R}_+, \quad v(0) = 1.$$

Since f is Lipschitz on $[0, 1]$, the result follows by ODE uniqueness. \square

PROOF OF THEOREM 2.3. The existence of a classical solution has been shown in the proof of Theorem 2.6 above. The property $\sup_{t \in (0, T]} \int xu(x, t) dx < \infty$ follows from the fact that $v \in \mathbb{L}_{\text{loc}}^\infty(\mathbb{R}_+; \mathbb{L}^1(\mathbb{R}_+))$. The uniqueness in the class of $C^{2,1}$ functions satisfying the above uniform integrability follows directly from uniqueness of weak solutions to (2.16), via integration by parts. \square

PROOF OF THEOREM 2.7. The C -tightness of $\bar{\xi}^n$ stated in Lemma 3.6, the fact stated in Lemma 4.4 that for any limit point ξ , $v(x, t) := \xi_t(x, \infty)$ is a weak solution of (2.18), and the uniqueness stated in Lemma 3.1 imply that $\bar{\xi}^n$ possesses a deterministic weak limit ξ , and that $\xi \in C(\mathbb{R}_+, \mathcal{M}_1)$. The assertion that ξ is given by $\xi_t(dx) = u(x, t) dx$, where the latter is the unique classical solution to (1.3), follows now from Lemma 4.4 and Theorem 2.6. \square

PROOF OF PROPOSITION 2.8. First we show that, under the strengthened moment condition, for t fixed,

$$(4.17) \quad \sup_n \max_{i \in [n]} \mathbb{E}[\hat{X}_i^n(t)^{2+\varepsilon}] < \infty.$$

We do this by reviewing the proof of (3.27) in Lemma 3.5(iii). Thanks to the assumption that Φ_{ser} has a finite $2 + \varepsilon$ moment, the moment estimate (3.30) can now be strengthened to

$$\sup_n \mathbb{E}[(\|\hat{S}_1^{0,n}\|_t^*)^{2+\varepsilon}] < \infty,$$

again by [28], Appendix 1. With this and the further $2 + \varepsilon$ moment assumptions on the initial condition, the proof of (4.17) now follows very closely to that of (3.27). We omit the details.

Next, (4.17) implies, using Minkowski’s inequality, that $n^{-1} \sum_i \hat{X}_i^n(t)^2 = \int x^2 \bar{\xi}_t^n(dx)$ are uniformly integrable. Hence it suffices to show $\int x^2 \bar{\xi}_t^n(dx) \rightarrow \int x^2 \xi_t(dx)$ in probability, and a similar statement for the empirical first moment. We have from Theorem 2.7 that $\bar{\xi}_t^n \rightarrow \xi_t$ in probability in \mathcal{M}_1 . Using the Skorohod representation theorem, we assume without loss of generality that this convergence holds a.s. Then $\int x^2 \bar{\xi}_t^n(dx) \rightarrow \int x^2 \xi_t(dx)$ a.s. will follow once the uniform integrability condition $\sup_n \int x^{2+\varepsilon} \bar{\xi}_t^n(dx) < \infty$ a.s. is verified. But this follows from Fatou’s lemma in view of (4.17). \square

PROOF OF THEOREM 2.9. By Lemma 3.3(iii) and Lemma 3.5(i), for $i \in [k]$,

$$(4.18) \quad \begin{aligned} \hat{X}_i^n(t) &= \hat{X}_i^n(0-) + \hat{E}_i^n(t) - \hat{S}_i^n(T_i^n(t)) + \hat{b}_1^n t \\ &\quad + \hat{\lambda}_0^n n^{-1/2} \int_0^t p_{n, \mathcal{R}_i^n(s)} ds + \hat{M}_i^{A,n}(t) + \hat{L}_i^n(t), \end{aligned}$$

and $\int \hat{X}_i^n(t) d\hat{L}_i^n(t) = 0$. From Step 5 in the proof of Lemma 4.1 we have that $T_i^n \rightarrow \iota$ in probability. Recall that $(\hat{E}_i^n, \hat{S}_i^n) \Rightarrow (E_i, S_i)$ and that $E_i - S_i$ are equal in law to $\sigma_i W_i$, where W_i

are mutually independent standard BMs. Since $(\hat{X}_i^n(0-))_{i \in [k]} \Rightarrow (X_i(0))_{i \in [k]}$ by assumption, using the dependence structure (2.7), we obtain

$$(\hat{X}_i^n(0-) + \hat{E}_i^n - \hat{S}_i^n(T_i^n))_{i \in [k]} \Rightarrow (X_i(0) + \sigma_i W_i)_{i \in [k]},$$

where $(X_i(0))_{i \in [k]}$ is independent of $(W_i)_{i \in [k]}$.

From (3.31) in the proof of Lemma 3.5 we have that $\hat{M}_i^{A,n} \rightarrow 0$ in probability. From Lemma 3.5(iii), we have C -tightness of $(\hat{X}_i^n, \hat{L}_i^n)_{i \in [k]}$. If we denote the integral term in (4.18) by \hat{K}_i^n , we have tightness of the tuple $(\hat{X}_i^n, \hat{X}_i^n(0-) + \hat{E}_i^n - \hat{S}_i^n(T_i^n), \hat{K}_i^n, \hat{L}_i^n)_{i \in [k]}$, and denoting a subsequential weak limit point by $(X_i, X_i(0) + \sigma W_i, K_i, L_i)_{i \in [k]}$, one has, for $i \in [k]$,

$$X_i(t) = X_i(0) + \sigma W_i(t) + b_1 t + K_i(t) + L_i(t), \quad \int X_i dL_i = 0.$$

By uniqueness in law of the system of SDE (1.5), the result will be proved once it is shown that $K_i(t) = b_0 \int_0^t v(X_i(s))^{\ell-1} ds$. Using Skorohod’s representation theorem we assume without loss that the convergence along the subsequence is a.s. Thus, in view of (4.12), it suffices to show that, along the subsequence, one has

$$(4.19) \quad (\bar{\mathcal{R}}_i^{c,n})_{i \in [k]} \rightarrow (v(X_i(\cdot), \cdot))_{i \in [k]},$$

a.s., in the uniform topology on $[t_0, T]$, for any $0 < t_0 < T$.

To this end, recall that Theorem 2.7 establishes, in particular, that $\bar{\xi}^n \rightarrow \xi$ in $D(\mathbb{R}_+, \mathcal{M}_1)$, where $\xi \in C(\mathbb{R}_+, \mathcal{M}_1)$, and ξ_t is atomless for every $t > 0$. Again it may be assumed that the convergence is a.s. Hence, with d_L the Levy–Prohorov metric, $\|d_L(\bar{\xi}^n, \xi)\|_T^* \rightarrow 0$ a.s. Because for any $t_0 > 0$ and $x_0 > 0$, v is uniformly continuous on $[0, x_0] \times [t_0, T]$, this gives

$$(4.20) \quad \sup_{t_0 \leq t \leq T} \sup_{x \leq x_0} |\bar{\xi}_t^n[x, \infty) - \xi_t[x, \infty)| \rightarrow 0, \quad \text{a.s.}$$

By the atomless property of ξ_t , $t > 0$, we know that

$$\bar{k}_n := \sup\{\bar{\xi}_t^n(\{x\}) : (x, t) \in [0, x_0] \times [t_0, T]\} \rightarrow 0, \quad \text{a.s.}$$

Now $\bar{\mathcal{R}}_i^{c,n} = n^{-1}(n - \mathcal{R}_i^n)$ satisfies

$$\bar{\xi}_t^n[\hat{X}_i^n(t), \infty) - \bar{k}_n \leq \bar{\mathcal{R}}_i^{c,n}(t) \leq \bar{\xi}_t^n[\hat{X}_i^n(t), \infty) + \bar{k}_n, \quad (x, t) \in [0, x_0] \times [t_0, T], i \in [k].$$

Using this in (4.20) shows

$$\sup_{t_0 \leq t \leq T} |\bar{\mathcal{R}}_i^{c,n}(t) - v(\hat{X}_i^n(t), t)| 1_{\{\hat{X}_i^n(t) \leq x_0\}} \rightarrow 0, \quad \text{a.s.}$$

Hence (4.19) follows by the a.s. convergence $\hat{X}_i^n \rightarrow X_i$ and a.s. finiteness of $\|X_i\|_T^*$, $i \in [k]$, completing the proof. \square

PROOF OF THEOREM 2.10. The assumed exchangeability of $\hat{X}_i^n(0-)$ for every n , along with the convergence in probability of their empirical law $\bar{\xi}_{0-}^n$ to the deterministic limit ξ_0 , imply the convergence $(\hat{X}_i^n(0-))_{i \in [k]} \Rightarrow (X_i(0))_{i \in [k]}$ where the latter are mutually independent and each is ξ_0 -distributed [42], Proposition 2.2. Hence the hypotheses of Theorem 2.9 hold. Because Theorem 2.9 asserts that $(\hat{X}_i^n, \hat{L}_i^n)_{i \in [k]}$ converge in law to the solution of the system (1.5), in which (W_i) are mutually independent and are independent of $(X_i(0))$, the additional mutual independence of $(X_i(0))$ that has just been shown completes the proof. \square

REMARK 4.5 (Sampling with replacement). If instead of sampling without replacement the algorithm employs sampling with replacement, one has the expression

$$\tilde{p}_{n,r} = \left(\frac{n-r+1}{n}\right)^\ell - \left(\frac{n-r}{n}\right)^\ell$$

instead of $p_{n,r}$ from (2.4). The proofs remain valid because the asymptotic of $\tilde{p}_{n,r}$ is as that of $p_{n,r}$. More precisely, analogously to (2.5), we have $\max_{r \in [n]} \tilde{p}_{n,r} = \frac{\ell}{n} + o\left(\frac{1}{n}\right)$. Moreover, it is easy to check that (4.12) holds true for $\tilde{p}_{n,r}$. As a result, the proof of Lemma 3.5(iii), where (2.5) is used, and of Lemma 4.1 (Step 6), where (4.12) is used, hold verbatim. These are the only places where the expression for $p_{n,r}$ is used in the proofs.

Acknowledgments. We thank the referees for many suggestions that helped improve the exposition considerably.

Funding. This research was supported by ISF (grant 1035/20).

REFERENCES

- [1] AGARWAL, P. and RAMANAN, K. (2025). Invariant states of hydrodynamic limits of randomized load balancing networks. *Math. Oper. Res.* arXiv preprint. Available at [arXiv:2008.08510](https://arxiv.org/abs/2008.08510).
- [2] AGHAJANI, R., LI, X. and RAMANAN, K. (2017). The PDE method for the analysis of randomized load balancing networks. *Proc. ACM Meas. Anal. Comput. Syst.* **1** 1–28.
- [3] AGHAJANI, R. and RAMANAN, K. (2019). The hydrodynamic limit of a randomized load balancing network. *Ann. Appl. Probab.* **29** 2114–2174. [MR3984253 https://doi.org/10.1214/18-AAP1444](https://doi.org/10.1214/18-AAP1444)
- [4] ATAR, R., KESLASSY, I., MENDELSON, G., ORDA, A. and VARGAFI, S. (2020). Persistent-idle load-distribution. *Stoch. Syst.* **10** 152–169. [MR4119248 https://doi.org/10.1287/stsy.2019.0054](https://doi.org/10.1287/stsy.2019.0054)
- [5] ATAR, R. and LIPSHUTZ, D. (2021). Heavy traffic limits for join-the-shortest-estimated-queue policy using delayed information. *Math. Oper. Res.* **46** 268–300. [MR4224428 https://doi.org/10.1287/moor.2020.1056](https://doi.org/10.1287/moor.2020.1056)
- [6] BANERJEE, S., BUDHIRAJA, A. and ESTEVEZ, B. (2025). Load balancing in parallel queues and rank-based diffusions. *Math. Oper. Res.* arXiv preprint. Available at [arXiv:2302.10317](https://arxiv.org/abs/2302.10317).
- [7] 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](https://doi.org/10.1214/18-AAP1436)
- [8] BANERJEE, S. and MUKHERJEE, D. (2020). Join-the-shortest queue diffusion limit in Halfin–Whitt regime: Sensitivity on the heavy-traffic parameter. *Ann. Appl. Probab.* **30** 80–144. [MR4068307 https://doi.org/10.1214/19-AAP1496](https://doi.org/10.1214/19-AAP1496)
- [9] BAYRAKTAR, E., BUDHIRAJA, A. and COHEN, A. (2019). Rate control under heavy traffic with strategic servers. *Ann. Appl. Probab.* **29** 1–35. [MR3909999 https://doi.org/10.1214/17-AAP1349](https://doi.org/10.1214/17-AAP1349)
- [10] BILLINGSLEY, P. (2013). *Convergence of Probability Measures*. Wiley, New York.
- [11] BRAMSON, M., LU, Y. and PRABHAKAR, B. (2010). Randomized load balancing with general service time distributions. *ACM SIGMETRICS Perform. Eval. Rev.* **38** 275–286.
- [12] 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](https://doi.org/10.1007/s11134-012-9311-0)
- [13] BRAMSON, M., LU, Y. and PRABHAKAR, B. (2013). Decay of tails at equilibrium for FIFO join the shortest queue networks. *Ann. Appl. Probab.* **23** 1841–1878. [MR3114919 https://doi.org/10.1214/12-AAP888](https://doi.org/10.1214/12-AAP888)
- [14] 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](https://doi.org/10.1287/moor.2019.1023)
- [15] CHEN, H. and YE, H.-Q. (2012). Asymptotic optimality of balanced routing. *Oper. Res.* **60** 163–179. [MR2911665 https://doi.org/10.1287/opre.1110.0998](https://doi.org/10.1287/opre.1110.0998)
- [16] CHUNG, K. L. and WILLIAMS, R. J. (1990). *Introduction to Stochastic Integration*, 2nd ed. *Probability and Its Applications*. Birkhäuser, Boston, MA. [MR1102676 https://doi.org/10.1007/978-1-4612-4480-6](https://doi.org/10.1007/978-1-4612-4480-6)
- [17] DALEY, D. J. and MIYAZAWA, M. (2019). A martingale view of Blackwell’s renewal theorem and its extensions to a general counting process. *J. Appl. Probab.* **56** 602–623. [MR3986955 https://doi.org/10.1017/jpr.2019.35](https://doi.org/10.1017/jpr.2019.35)
- [18] ESCHENFELDT, P. and GAMARNIK, D. (2016). Supermarket queueing system in the heavy traffic regime. Short queue dynamics. arXiv preprint. Available at [arXiv:1610.03522](https://arxiv.org/abs/1610.03522).

- [19] 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>
- [20] FOSS, S. and STOLYAR, A. L. (2017). Large-scale join-idle-queue system with general service times. *J. Appl. Probab.* **54** 995–1007. MR3731281 <https://doi.org/10.1017/jpr.2017.49>
- [21] GÄRTNER, J. (1988). On the McKean–Vlasov limit for interacting diffusions. *Math. Nachr.* **137** 197–248. MR0968996 <https://doi.org/10.1002/mana.19881370116>
- [22] GILBARG, D. and TRUDINGER, N. S. (1983). *Elliptic Partial Differential Equations of Second Order*, 2nd ed. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]* **224**. Springer, Berlin. MR0737190 <https://doi.org/10.1007/978-3-642-61798-0>
- [23] GODLEWSKI, E. and RAVIART, P.-A. (1991). *Hyperbolic Systems of Conservation Laws. Mathématiques & Applications (Paris) [Mathematics and Applications]* **3/4**. Ellipses, Paris. MR1304494
- [24] 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>
- [25] GRAHAM, C. (2005). Functional central limit theorems for a large network in which customers join the shortest of several queues. *Probab. Theory Related Fields* **131** 97–120. MR2105045 <https://doi.org/10.1007/s00440-004-0372-9>
- [26] 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>
- [27] JACOD, J. and SHIRYAEV, A. N. (1987). *Limit Theorems for Stochastic Processes. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]* **288**. Springer, Berlin. MR0959133 <https://doi.org/10.1007/978-3-662-02514-7>
- [28] KRICHAGINA, E. V. and TAKSAR, M. I. (1992). Diffusion approximation for $GI/G/1$ controlled queues. *Queueing Syst. Theory Appl.* **12** 333–367. MR1200872 <https://doi.org/10.1007/BF01158808>
- [29] LACKER, D. (2018). On a strong form of propagation of chaos for McKean–Vlasov equations. *Electron. Commun. Probab.* **23** Paper No. 45. MR3841406 <https://doi.org/10.1214/18-ECP150>
- [30] LIU, X., GONG, K. and YING, L. (2022). Large-system insensitivity of zero-waiting load balancing algorithms. *ACM SIGMETRICS Perform. Eval. Rev.* **50** 101–102.
- [31] LUCZAK, M. J. and MCDIARMID, C. (2006). On the maximum queue length in the supermarket model. *Ann. Probab.* **34** 493–527. MR2223949 <https://doi.org/10.1214/00911790500000710>
- [32] LUCZAK, M. J. and NORRIS, J. (2005). Strong approximation for the supermarket model. *Ann. Appl. Probab.* **15** 2038–2061. MR2152252 <https://doi.org/10.1214/105051605000000368>
- [33] MENDELSON, G. and XU, K. (2025). Care: Resource allocation using sparse communication. *Oper. Res.* arXiv preprint. Available at [arXiv:2206.02410](https://arxiv.org/abs/2206.02410).
- [34] MITZENMACHER, M. (2001). The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distrib. Syst.* **12** 1094–1104.
- [35] 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>
- [36] MUKHOPADHYAY, A. and MAZUMDAR, R. R. (2016). Analysis of randomized join-the-shortest-queue (JSQ) schemes in large heterogeneous processor-sharing systems. *IEEE Trans. Control Netw. Syst.* **3** 116–126. MR3514587 <https://doi.org/10.1109/TCNS.2015.2428331>
- [37] PINCHOVER, Y. and RUBINSTEIN, J. (2005). *An Introduction to Partial Differential Equations*. Cambridge Univ. Press, Cambridge. MR2164768 <https://doi.org/10.1017/CBO9780511801228>
- [38] PROTTER, P. (1990). *Stochastic Integration and Differential Equations: A New Approach. Applications of Mathematics (New York)* **21**. Springer, Berlin. MR1037262 <https://doi.org/10.1007/978-3-662-02619-9>
- [39] SHKOLNIKOV, M. (2012). Large systems of diffusions interacting through their ranks. *Stochastic Process. Appl.* **122** 1730–1747. MR2914770 <https://doi.org/10.1016/j.spa.2012.01.011>
- [40] 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>
- [41] STOLYAR, A. L. (2017). Pull-based load distribution among heterogeneous parallel servers: The case of multiple routers. *Queueing Syst.* **85** 31–65. MR3604117 <https://doi.org/10.1007/s11134-016-9508-8>
- [42] SZNITMAN, A.-S. (1991). Topics in propagation of chaos. In *École D’Été de Probabilités de Saint-Flour XIX—1989. Lecture Notes in Math.* **1464** 165–251. Springer, Berlin. MR1108185 <https://doi.org/10.1007/BFb0085169>
- [43] VAN DER BOOR, M., BORST, S. C., VAN LEEUWAARDEN, J. S. H. and MUKHERJEE, D. (2022). Scalable load balancing in networked systems: A survey of recent advances. *SIAM Rev.* **64** 554–622. MR4461562 <https://doi.org/10.1137/20M1323746>
- [44] VAN DER BOOR, M., ZUBELDIA, M. and BORST, S. (2020). Zero-wait load balancing with sparse messaging. *Oper. Res. Lett.* **48** 368–375. MR4093088 <https://doi.org/10.1016/j.orl.2020.04.006>

- [45] 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](#)
- [46] WOLANSKY, G. (1995). Comparison between two models of self-gravitating clusters: Conditions for gravitational collapse. *Nonlinear Anal.* **24** 1119–1129. [MR1321749](#) [https://doi.org/10.1016/0362-546X\(94\)E0028-F](https://doi.org/10.1016/0362-546X(94)E0028-F)
- [47] YING, L., SRIKANT, R. and KANG, X. (2017). The power of slightly more than one sample in randomized load balancing. *Math. Oper. Res.* **42** 692–722. [MR3685262](#) <https://doi.org/10.1287/moor.2016.0823>
- [48] ZHAO, Z., BANERJEE, S. and MUKHERJEE, D. (2025). Many-server asymptotics for join-the-shortest queue in the super-Halfin–Whitt scaling window. *Math. Oper. Res.* arXiv preprint. Available at [arXiv:2106.00121](#).