THE CRITICAL GREEDY SERVER ON THE INTEGERS IS RECURRENT

By James R. Cruise and Andrew R. Wade

Heriot-Watt University and Durham University

Each site of $\mathbb{Z}$ hosts a queue with arrival rate λ. A single server, starting at the origin, serves its current queue at rate μ until that queue is empty, and then moves to the longest neighbouring queue. In the critical case $\lambda = \mu$, we show that the server returns to every site infinitely often. We also give a sharp iterated logarithm result for the server's position. Important ingredients in the proofs are that the times between successive queues being emptied exhibit doubly exponential growth, and that the probability that the server changes its direction is asymptotically equal to 1/4.

1. Introduction and main results. The following continuous-time stochastic model was introduced by Kurkova and Menshikov [13]. Each site of the one-dimensional integer lattice $\mathbb{Z}$ is associated with a queue. Each queue has an independent Poisson arrival stream of rate $\lambda \in (0, \infty)$. The system has a single server, which starts at the origin at time 0. The server serves the queue at its current site exhaustively at rate $\mu \in (0, \infty)$. If the queue at the current site is empty at time $t \geq 0$, the server moves to one of the two neighbouring sites using a greedy policy: it chooses to move to the site with the longest queue (measured at time t), randomly breaking any tie. The server moves (deterministically) at unit speed, and so arrives at the new site at time $t + 1$, at which time it starts to serve the new queue. Of interest is the asymptotic behaviour of $S(t)$, the location of the server at time $t \geq 0$.

There are 3 cases. The least interesting is when $\lambda > \mu$ and any queue under service is transient, so each time that the server starts serving a non-empty queue, there is a (uniformly) positive probability that the server remains at the site for all time. Thus with probability 1, the server changes site only finitely many times, so $S(t)$ converges almost surely (a.s.) See Theorem 1.1 of [13].

The main object of study in [13] is the case $\lambda < \mu$. Theorem 1.2 of [13] shows that in this case the server changes its direction only finitely many times, so that the server eventually moves from site to site in a single direction, and $S(t) \to \pm\infty$ a.s. The intuition behind this result is as follows. Any queue under service is now recurrent, so becomes emptied in finite time. Consider the server's first visit to site $x > 0$ at time t (say). At time t, with very high probability, the queue at $x - 1$ will be essentially empty, while the queues at x and $x + 1$ will have lengths concentrated about λt. It takes time about ct for the server to empty the queue at site x, where c

---

is a positive constant depending on λ and μ, and in this time there will be about $\lambda c t$ new arrivals at the queues at sites $x - 1$ and $x + 1$. So by the time the queue at site x is emptied, the queue lengths at sites $x - 1$ and $x + 1$ are about $\lambda c t$ and $(\lambda + \lambda c)t$ respectively. The fluctuations are of order $t^{1/2}$, and so with very high probability, the server will choose to go to site $x + 1$ next.

In the present paper, we study the critical case $\lambda = \mu$, which was left open in [13]. The analysis of this case seems somewhat more delicate. Again, any queue under service is recurrent. But now an attempt to follow the idea of the argument sketched in the previous case reveals a new issue. Once again, upon the server's first arrival at site x, the queues at x and $x + 1$ will have lengths about λt while the queue at site $x - 1$ will be essentially empty. But now the queue at site x is critically recurrent, and so typically takes time of order t^2 to empty. In this time, the fluctuations in the new arrivals at sites $x - 1$ and $x + 1$ are of order t, i.e., on the same scale as the initial difference in queue lengths. So it seems likely that the server will change direction many times; an understanding of the details of this behaviour seems necessary to obtain the asymptotic behaviour of the server.

In the case $\lambda = \mu$, Menshikov and Kurkova [13] proved that

(1.1) $$\limsup_{t \to \infty} |S(t)| = +\infty, \text{ a.s.,}$$

showing that the server does not get stuck. Our main result is that the server is *recurrent*, in the sense that it returns to every site infinitely many times:

THEOREM 1.1.   *Suppose that $\lambda = \mu \in (0, \infty)$. Then, a.s., for every $x \in \mathbb{R}$, the set $\{t \geq 0 : S(t) = x\}$ is unbounded.*

We also establish the following result on the growth rate of $S(t)$.

THEOREM 1.2.   *Suppose that $\lambda = \mu \in (0, \infty)$. Then, a.s.,*

$$\limsup_{t \to \infty} \frac{S(t)}{\sqrt{\log\log t \log\log\log\log t}} = \sqrt{\frac{6}{\log 2}}, \text{ and}$$
$$\liminf_{t \to \infty} \frac{S(t)}{\sqrt{\log\log t \log\log\log\log t}} = -\sqrt{\frac{6}{\log 2}}.$$

The rest of this paper is devoted to the proofs of Theorems 1.1 and 1.2. We give a concrete construction of the process, via a discrete-time process that is the basic object of study in this paper, in Section 2. In Section 2 we also describe the main steps in the proof and give the outline of the rest of the paper.

The greedy server on $\mathbb{Z}$ is a variant of the *greedy server* problem introduced in [5] and surveyed in [17], in which a server greedily moves from job to job arriving randomly in some space, such as on a line or a circle; see also [9]. Also related is the so-called *greedy walk* problem [2]. These models have received significant

attention over several decades, in part because the dynamics of the server possess features of both self-interacting processes and processes in random environments, which remain very active topics of current research, and because the problems they pose are challenging. A number of open problems remain: see e.g. [17, 2].

A continuum analogue of our problem on $\mathbb{Z}$ is the greedy walk on $\mathbb{R}$, for which it was recently shown that the server escapes to infinity [10], in contrast to our Theorem 1.1.

2. Discrete-time process and paper outline.

For the remainder of the paper we fix $\lambda = \mu \in (0, \infty)$. We write $\mathbb{Z}_+ := \{0, 1, 2, \ldots\}$ and $\mathbb{N} := \{1, 2, 3, \ldots\}$.

We will study the continuous-time process described in Section 1 via a discrete-time process obtained by observing the full process at the nth time at which the server empties a queue. Consider a Markov process $\Psi_n = (Q_n, X_n, T_n)$, $n \in \mathbb{Z}_+$, where $Q_n = (Q_n(x))_{x \in \mathbb{Z}} \in (\mathbb{Z}_+ \cup \{*\})^{\mathbb{Z}}$, $X_n \in \mathbb{Z}$, and $T_n \in \mathbb{R}_+$. If $Q_n(x) \in \mathbb{Z}_+$, that is the number of customers at queue $x \in \mathbb{Z}$; if $Q_n(x) = *$ then this indicates that the queue at $x \in \mathbb{Z}$ has yet to be inspected by the server. The coordinate X_n represents the location of the server when a queue is emptied for an nth time, and T_n represents the total time that has elapsed (i.e. the sum of all the services times plus the travel times up to this point).

Set $Q_0(x) := 0$ for $|x| \leq 3$ and $Q_0(x) := *$ for $|x| \geq 4$; set $X_0 := 0$ and $T_0 := 0$. We describe the law of this process by its Markovian transitions. The random ingredients that go into this description are as follows. Let $\xi_1, \xi_2, \ldots$ be i.i.d. with $\mathbb{P}(\xi_1 = +1) = \mathbb{P}(\xi_1 = -1) = 1/2$ (these will be the tie-breaking variables).

We write $P(\kappa)$ to denote a Poisson random variable with mean $\kappa \in \mathbb{R}_+$; for a random variable W on $\mathbb{R}_+$ we write $P(W)$ to denote a random variable that, conditional on W, has a Poisson distribution with mean W.

Let $(Z_t, t \in \mathbb{R}_+)$ denote an $M/M/1$ queue with arrival rate λ and service rate μ, with $Z_0 = k \in \mathbb{Z}_+$ initial customers; let $\zeta(k)$ denote a random variable distributed as $\inf\{t \geq 0 : Z_t = 0\}$, the time to empty for the queue. Similarly, conditional on a random variable $W \in \mathbb{Z}_+$, $\zeta(W)$ is distributed as the time to empty the queue started from $Z_0 = W$.

Given (Q_n, X_n, T_n) for $n \in \mathbb{Z}_+$, generate $(Q_{n+1}, X_{n+1}, T_{n+1})$ as follows.

- Define
$$
\eta_{n+1} := \begin{cases} +1 \text{ if } Q_n(X_n + 1) > Q_n(X_n - 1), \\ -1 \text{ if } Q_n(X_n - 1) > Q_n(X_n + 1). \end{cases}
$$
If $Q_n(X_n + 1) = Q_n(X_n - 1)$, then take $\eta_{n+1} = \xi_n$.
- Let $X_{n+1} = X_n + \eta_{n+1}$.
- Let τ_{n+1} be distributed as 1 plus $\zeta(Q_n(X_{n+1}) + P(\lambda))$.
- Let $T_{n+1} = T_n + \tau_{n+1}$.
- For every $k \in \mathbb{Z} \setminus \{X_{n+1}\}$ such that $Q_n(k) \neq *$, take $Q_{n+1}(k)$ to be distributed as $Q_n(k) + P(\lambda \tau_{n+1})$, independently for each k.
- If $Q_n(X_{n+1} + 3\eta_n) = *$ then let $Q_{n+1}(X_{n+1} + 3\eta_n)$ be distributed as $P(\lambda T_{n+1})$.

- Set $Q_{n+1}(X_{n+1}) = 0$.

Note that $Q_n(X_n) = 0$ for all n, and $T_n = \sum_{k=1}^n \tau_k$. If $a_n = \min_{0 \leq j \leq n} X_n$ and $b_n = \max_{0 \leq j \leq n} X_n$, then at time T_n all the queues in the interval $[a_n, b_n]$ have been visited at least once, and, by construction, $Q_n(x) \neq *$ for $x \in [a_n - 3, b_n + 3]$, so the process tracks the queue lengths within distance 3 of the queues visited so far.

For $n \in \mathbb{Z}_+$, define the σ-algebra $\mathcal{F}_n := \sigma(\Psi_0, \xi_0, \Psi_1, \xi_1, \ldots, \Psi_n, \xi_n)$. For $n \in \mathbb{N}$, clearly both $\tau_n = T_n - T_{n-1}$ and $\eta_n = X_n - X_{n-1}$ are $\mathcal{F}_n$-measurable, but it is important to observe that X_{n+1} and η_{n+1} are also $\mathcal{F}_n$-measurable. Thus (X_n, η_n) is $\mathcal{F}_{n-1}$ measurable. Note that $\tau_n \geq 1$, a.s., so that $T_n \geq n$, a.s.

Let $N_t := \max\{n \in \mathbb{Z}_+ : T_n \leq t\}$ denote the number of times that a queue has been emptied by time t. Since $T_n \to \infty$ a.s., we have $N_t < \infty$ a.s. for all $t \in \mathbb{R}_+$; indeed, $N_t \leq T_{N_t} \leq t$, a.s. Moreover, N_t is nondecreasing in t. Note $T_{N_t+1} > t$. Thus if $N_t \to N < \infty$, we have $T_{N+1} = \infty$ which contradicts $T_n < \infty$ for all n; hence $N_t \to \infty$ a.s. as $t \to \infty$.

Observe that X_{N_t} is the most recent queue that was emptied prior to time t, and X_{N_t+1} is the next queue to be emptied after time t. Also, $T_{N_t} \leq t$ is the time at which the most recently emptied queue was emptied. It follows that we have the representation

$$(2.1) \qquad S(t) = \begin{cases} X_{N_t+1} & \text{if } t - T_{N_t} \geq 1, \\ X_{N_t} + (t - T_{N_t})(X_{N_t+1} - X_{N_t}) & \text{if } t - T_{N_t} \leq 1. \end{cases}$$

We end this section by outlining the main steps in the proofs of Theorems 1.1 and 1.2, and some of the underlying intuition. The first key ingredient is that τ_n and T_n exhibit *doubly exponential* growth (see Propositions 4.1 and 4.2). This very rapid growth in the time-scales suggests an effective memorylessness for the system: the configuration of the discrete-time system more than one or two time-steps ago is not important. This provides the intuition behind the second key ingredient, which is establishing (in Proposition 5.1) that the *turning probability* converges:

$$(2.2) \qquad \mathbb{P}(\eta_{n+1} \neq \eta_n \mid \mathcal{F}_{n-1}) \to \frac{1}{4}, \text{ a.s.}$$

Some rough heuristics behind this result are as follows. Because of the double-exponential growth, when the server arrives at a queue, the size of the new queue is about the same as the neighbour ahead of it, whereas the queue behind is essentially empty. In order for the server to backtrack, two conditions must be satisfied: the difference between the sizes of the queues ahead and behind must (i) reach zero before the current queue is emptied, and (ii) be negative when the current queue is emptied. The probability of (i) is close to $1/2$, and when (i) occurs the conditional probability of (ii) equals $1/2$, hence the $1/4$.

The convergence of the turning probability at (2.2) means that the server's motion is asymptotically similar to the Gillis–Domb–Fisher *correlated random walk* [3]. In

fact, more than convergence in (2.2) is necessary: we need a bound on the *rate of convergence* with n. Fairly rough estimates are enough, due to the double-exponential growth of the time-scales. The double-exponential growth is also the origin of the iterated logarithm in Theorem 1.2; the precise value of the constant comes in part from the value of (2.2). A technical device central to the proofs of both theorems is the construction of a function $f(X_n, \eta_n)$ of the process that is close to a martingale.

The outline of the rest of the paper is as follows. Section 3 collects some results on the random variables $\zeta(k)$ that we will need in our analysis. Section 4 contains the key estimates on the growth of τ_n and T_n. Section 5 contains the convergence result for the turning probability. Section 6 contains the martingale construction that allows us to complete the proofs of Theorems 1.1 and 1.2. Finally, Appendix A collects a couple of auxiliary results used in the analysis.

3. The critically-loaded queue. Let $(Z_t, t \in \mathbb{R}_+)$ be a continuous-time symmetric simple random walk on $\mathbb{Z}$ with jump rate 2λ, i.e., for every $x \in \mathbb{Z}$, for all $t \in \mathbb{R}_+$,

$$\mathbb{P}(Z_{t+h} = x \pm 1 \mid Z_t = x) = \lambda h + o(h),$$

as $h \to 0$. Suppose that $Z_0 = k \in \mathbb{Z}_+$ and let $\zeta(k) := \inf\{t \geq 0 : Z_t = 0\}$, the time to reach 0 started from k. Note that up until $\zeta(k)$, Z_t is distributed as the number of customers in an $M/M/1$ queue with arrival rate and service rate both equal to $\lambda \in (0, \infty)$; so $\zeta(k)$ is the time to empty such an $M/M/1$ queue, started from $k \in \mathbb{Z}_+$ initial customers, as described in Section 2.

First we collect several straightforward results about $\zeta(k)$ that we need in the rest of the paper.

LEMMA 3.1.

(a) *For $k \geq \ell \geq 0$, $\zeta(k)$ stochastically dominates $\zeta(\ell)$.*
(b) *For any $\alpha \in (0, 2)$ and $c \in (0, \infty)$ there exist $\varepsilon > 0$ and $k_1 \in \mathbb{N}$ such that*

$$\mathbb{P}(\zeta(k) \leq ck^\alpha) \leq \exp\{-k^\varepsilon\}, \text{ for all } k \geq k_1.$$

(c) *For any $\beta \in (2, \infty)$ and $c \in (0, \infty)$ there exist $\varepsilon > 0$ and $k_2 \in \mathbb{N}$ such that*

$$\mathbb{P}(\zeta(k) \geq ck^\beta) \leq k^{-\varepsilon}, \text{ for all } k \geq k_2.$$

Before proving this lemma, we make some observations. Suppose $Z_0 = k \in \mathbb{Z}_+$. Then we can represent $\zeta(k)$ as

$$(3.1) \qquad\qquad \zeta(k) = Y_1 + \cdots + Y_k,$$

where

$$Y_j = \inf\{t \geq 0 : Z_t = j - 1\} - \inf\{t \geq 0 : Z_t = j\}.$$

Note that, by the strong Markov property and the spatial homogeneity of the random walk, the Y_j in (3.1) are i.i.d. copies of $\zeta(1)$.

It is well known that for $k \in \mathbb{N}$, $\zeta(k)$ has density

$$f_k(u) := \frac{k}{u} I_k(2\lambda u) \mathrm{e}^{-2\lambda u}, \text{ for } u > 0;$$

see for example Sections II.7 and XIV.6 of [8]. Here I_k is the modified Bessel function of the first kind:

$$I_k(u) := \sum_{j=0}^{\infty} \frac{(u/2)^{2j+k}}{j!(j+k)!}.$$

In particular, the density of $\zeta(1)$ is

$$(3.2) \qquad f(u) := f_1(u) = \frac{1}{u} I_1(2\lambda u) \mathrm{e}^{-2\lambda u} = \frac{1}{2\sqrt{\pi\lambda}} u^{-3/2} + O(u^{-5/2}),$$

as $u \to \infty$, using the asymptotic expressions of [20, p. 203]. Note that $f(u) > 0$ for all $u \in (0, \infty)$. Let $F(u) := \mathbb{P}(\zeta(1) \leq u)$ and $\bar{F}(u) := \mathbb{P}(\zeta(1) > u)$. Then, by (3.2),

$$(3.3) \qquad \bar{F}(u) = \int_u^{\infty} f(v)\mathrm{d}v = \frac{1}{\sqrt{\pi\lambda}} u^{-1/2} + O(u^{-3/2}),$$

as $u \to \infty$.

PROOF OF LEMMA 3.1. For $0 \leq \ell \leq k$, the representation (3.1) gives $\zeta(k) \geq \zeta(\ell)$ in the obvious coupling, so we get part (a).

For part (b), we use (3.1) and the fact that $Y_j \geq 0$ to write

$$\mathbb{P}(\zeta(k) \leq r) \leq \mathbb{P}\left(\max_{1 \leq j \leq k} Y_j \leq r \right)$$

$$= \mathbb{P}\left(\bigcap_{j=1}^k \{Y_j \leq r\} \right)$$

$$= \left(1 - (\pi\lambda(1 + o(1))r)^{-1/2} \right)^k,$$

as $r \to \infty$, by (3.3). In particular, taking $r = ck^\alpha$ with $c, \alpha > 0$ gives

$$\mathbb{P}(\zeta(k) \leq ck^\alpha) \leq \left(1 - (\pi\lambda c(1 + o(1))k^\alpha)^{-1/2} \right)^k,$$

as $k \to \infty$, so that

$$\log \mathbb{P}(\zeta(k) \leq ck^\alpha) \leq k \log \left(1 - (\pi\lambda c(1 + o(1))k^\alpha)^{-1/2} \right) \sim -(\pi\lambda c)^{-1/2} k^{1-(\alpha/2)},$$

which gives part (b).

For part (c), let $p \in (1/\beta, 1/2)$. Note from (3.3) that the Y_j appearing in (3.1) have $\mathbb{E}(Y_j^p) \leq C$ for some $C < \infty$. Then by subadditivity of the function $y \mapsto y^p$ we have from (3.1) that $\mathbb{E}(\zeta(k)^p) \leq \sum_{j=1}^{k} \mathbb{E}(Y_j^p) \leq Ck$. Hence, by Markov's inequality,

$$\mathbb{P}(\zeta(k) \geq ck^{\beta}) = \mathbb{P}(\zeta(k)^p \geq c^p k^{\beta p}) \leq Cc^{-p} k^{1-\beta p},$$

which, by choice of p, gives part (c). $\qquad\square$

Let Φ be the distribution function of the standard normal distribution, and let $\bar{\Phi}(u) := 1 - \Phi(u)$ for $u \in \mathbb{R}$. We say that S has a *Lévy distribution* with location parameter 0 and scale parameter 1 (see [14, §1.1]) if $S \in \mathbb{R}_+$ has distribution function given by

$$(3.4) \qquad F_S(u) := \mathbb{P}(S \leq u) = 2\bar{\Phi}(u^{-1/2}), \text{ for } u > 0.$$

Note that the density $f_S(u) := F_S'(u)$ corresponding to (3.4) is

$$(3.5) \qquad f_S(u) = \frac{1}{\sqrt{2\pi}} u^{-3/2} e^{-u^{-1}/2}, \text{ for } u > 0.$$

LEMMA 3.2. *Let S be a random variable with the distribution given by (3.4). There exists a constant $C \in \mathbb{R}_+$ such that, for all $k \in \mathbb{N}$,*

$$\sup_{u \in \mathbb{R}_+} \left| \mathbb{P}(k^{-2}\zeta(k) \leq u) - F_S(2\lambda u) \right| \leq Ck^{-1}.$$

PROOF. For the purposes of this proof only, we take $Z_0 = 0$. Let $D := D(\mathbb{R}_+, \mathbb{R})$ denote the space of functions from $\mathbb{R}_+ \to \mathbb{R}$ that are right-continuous and have left limits, endowed with the Skorokhod metric. Define for $m \in \mathbb{N}$,

$$z_m(t) := m^{-1/2} Z_{mt}, \text{ for } t \geq 0.$$

Then $z_m \in D$ for each $m \in \mathbb{N}$. Let $(b(t), t \in \mathbb{R}_+)$ denote standard Brownian motion started at $b_0 = 0$. The invariance principle for continuous-time random walks implies that as $m \to \infty$, $z_m \Rightarrow \sqrt{2\lambda}b$ in the sense of weak convergence on D (one may apply Theorem 7.1.4 of [7, pp. 339–340], for example).

For $z \in D$, let $\sigma(z) := \inf\{t \geq 0 : z(t) > 1\}$. For Brownian motion, we have that $\sigma(b) = \inf\{t \geq 0 : b(t) = 1\}$ a.s., and, for any $\varepsilon > 0$, $\sup_{0 \leq s \leq \sigma(b)-\varepsilon} b(s) < 1$, a.s. Thus the set of discontinuities of the mapping $z \mapsto \sigma(z)$ has measure zero under the measure induced by Brownian motion (see Section 5.7.5 of [21]). So by the mapping theorem (see e.g. Theorem 2.7 of [1]) we get $\sigma((2\lambda)^{-1/2} z_m) \xrightarrow{d} S := \sigma(b)$ (where $\xrightarrow{d}$ denotes convergence in distribution). Here

$$\begin{aligned}
\sigma((2\lambda)^{-1/2} z_m) &= \inf\{t \geq 0 : Z_{mt} > \sqrt{2\lambda m}\} \\
&= m^{-1} \inf\{s \geq 0 : Z_s \geq 1 + \lfloor \sqrt{2\lambda m} \rfloor\} \\
&\overset{d}{=} m^{-1} \zeta(1 + \lfloor \sqrt{2\lambda m} \rfloor),
\end{aligned}$$

where $\stackrel{d}{=}$ denotes equality in distribution. Setting $k = 1 + \lfloor \sqrt{2\lambda m} \rfloor \in \mathbb{N}$ we have that $2\lambda/k^2 = m^{-1} + O(m^{-3/2})$ so that, as $k \to \infty$,

$$(3.6) \qquad \frac{2\lambda}{k^2}\zeta(k) \xrightarrow{d} S.$$

The reflection principle for Brownian motion (see e.g. [6, p. 372]) shows that

$$\mathbb{P}(S > u) = 1 - 2\mathbb{P}(b(u) \geq 1) = 2\Phi(u^{-1/2}) - 1,$$

so S has the distribution given by (3.4).

It remains to estimate the rate of convergence in (3.6). By (3.3) and Theorem 2.6.7 of [11], we have that $\zeta(1)$ is in the normal domain of attraction of a positive stable law with index $1/2$. Indeed, S is stable with index $1/2$ since, by the scaling and strong Markov properties of Brownian motion, for any $m \in \mathbb{N}$,

$$S \stackrel{d}{=} \inf\{t \geq 0 : m^{-1}b(m^2 t) = 1\} = m^{-2}\inf\{t \geq 0 : b(t) = m\} \stackrel{d}{=} m^{-2}(S_1 + \cdots + S_m),$$

where the S_j are independent copies of S. Thus we can apply results on the rate of convergence in the stable central limit theorem for the sum in (3.1). First note that, by Taylor's theorem,

$$\Phi(u^{-1/2}) = \Phi(0) + \Phi'(0)u^{-1/2} + \frac{\Phi''(0)}{2}u^{-1} + O(u^{-3/2})$$

$$= \frac{1}{2} + \frac{1}{\sqrt{2\pi}}u^{-1/2} + O(u^{-3/2}),$$

as $u \to \infty$. Thus if $\bar{F}_S(u) := 1 - F_S(u)$ we have from (3.4) that

$$(3.7) \qquad \bar{F}_S(2\lambda u) = 2\Phi((2\lambda u)^{-1/2}) - 1 = \frac{1}{\sqrt{\pi\lambda}}u^{-1/2} + O(u^{-3/2}),$$

as $u \to \infty$. Combining (3.7) with (3.3) we have that

$$|F(u) - F_S(2\lambda u)| = O(u^{-3/2}).$$

This condition enables one to verify standard 'pseudomoments' conditions for Berry–Esseen bounds in stable limit theorems. Indeed, setting $H(u) = F(u) - F_S(2\lambda u)$ and

$$\mu_\ell = \int_0^\infty u^\ell \mathrm{d}H(u), \text{ and } \nu_\ell = \int_0^\infty |u|^\ell |\mathrm{d}H(u)|,$$

we have that $\nu_1 < \infty$ and $\mu_0 = 0$, so we may apply the results of [18] (which has a statement but no proof), [15] (Corollary 1) or [4] (combine Theorem 3.11 of [4, p. 66] with Lemma 2.5 of [4, p. 27]). This gives the result. $\qquad \square$

4. Time-scale estimates. In this section we study the asymptotics of τ_n and T_n. First we have a lower bound.

PROPOSITION 4.1. *For any $\alpha \in (1,2)$, $T_n \geq \tau_n \geq \mathrm{e}^{\alpha^n}$ for all but finitely many n, a.s.*

We also have the following upper bound.

PROPOSITION 4.2. *For any $\beta \in (2,\infty)$, $\tau_n \leq T_n \leq \mathrm{e}^{\beta^n}$ for all but finitely many n, a.s.*

REMARK. A rough calculation (cf. Lemma 5.2 below) suggests that we may have

$$\lim_{n\to\infty} \frac{\log \tau_n}{2^n} = \gamma, \ \text{a.s.},$$

and the same for T_n. Here $\gamma \in (0,\infty)$ is a random variable with representation $\gamma = \sum_{i=1}^{\infty} 2^{-i} \log(\lambda S_i/2)$, where $S_1, S_2, \ldots$ are independent random variables with distribution given by (3.4). To establish Theorems 1.1 and 1.2 however, the bounds in Propositions 4.1 and 4.2 are sufficient (in fact, for Theorem 1.1, we need only the lower bound).

We work towards the proof of Proposition 4.1. We start with a crude bound. Here and elsewhere, 'i.o.' stands for 'infinitely often'.

LEMMA 4.3. *We have $\tau_n \geq n^2$ i.o., a.s.*

PROOF. We have from the description in Section 2 and Lemma 3.1(a) that τ_n stochastically dominates $\zeta(Q_{n-1})$, where $Q_{n-1} := Q_{n-1}(X_n)$. Moreover, since $X_n \neq X_{n-1}$, we have that the queue at X_n is not being served between times T_{n-2} and T_{n-1}, and in that time accumulates a Poisson number of arrivals with mean $\lambda\tau_{n-1} \geq \lambda$, since $\tau_{n-1} \geq 1$ a.s. Hence Q_{n-1} stochastically dominates a Poisson random variable with mean λ, independently of $\tau_{n-1}, \tau_{n-2}, \ldots$. Hence $(\tau_n)_n$ stochastically dominates $(Z_n)_n$, where $Z_1, Z_2, \ldots$ are i.i.d. with

$$\mathbb{P}(Z_1 > r) = \mathbb{P}(\zeta(P(\lambda)) \geq r) \geq \mathbb{P}(P(\lambda) \geq 1)\mathbb{P}(\zeta(1) > r)$$
$$\geq (1 - \mathrm{e}^{-\lambda})(\pi\lambda(1 + o(1))r)^{-1/2},$$

as $r \to \infty$, by (3.3). In particular, $\sum_{n\geq 1} \mathbb{P}(Z_n > n^2) = \infty$, and so $Z_n \geq n^2$ i.o., by the Borel–Cantelli lemma, and hence the same is true of τ_n. $\square$

The next result gives conditions under which an a.s. lower bound for τ_n that holds infinitely often can be converted into a bound that holds all but finitely often.

LEMMA 4.4. *Let $b_1, b_2, \ldots \in (0,\infty)$ and $\alpha \in (1,2)$ be such that*

(i) $\lim_{n \to \infty} b_n = \infty$;

(ii) $\lim_{n \to \infty} (b_{n+1}/b_n^\alpha) = 0$;

(iii) $\sum_{n=1}^{\infty} \mathrm{e}^{-b_n^\varepsilon} < \infty$ for every $\varepsilon > 0$.

Suppose that $\mathbb{P}(\tau_n > b_n \ i.o.) = 1$. Then $\tau_n > b_n$ for all but finitely many n, a.s.

PROOF. We have from Lemma 3.1(a) that, given $\mathcal{F}_n$, τ_{n+1} stochastically dominates $\zeta(Q_n)$, where $Q_n := Q_n(X_{n+1})$. Thus

$$\mathbb{P}(\tau_{n+1} \leq b_{n+1}, \tau_n > b_n) \leq \mathbb{P}(\zeta(Q_n) \leq b_{n+1}, \tau_n > b_n)$$
$$\leq \mathbb{P}(\zeta(Q_n) \leq b_{n+1}, Q_n \geq \lambda \tau_n/2, \tau_n > b_n)$$
$$+ \mathbb{P}(Q_n \leq \lambda \tau_n/2, \tau_n > b_n).$$

Moreover, given τ_n, Q_n stochastically dominates $P(\lambda \tau_n)$, so that

$$\mathbb{P}(\tau_{n+1} \leq b_{n+1}, \tau_n > b_n) \leq \mathbb{P}(\zeta(Q_n) \leq b_{n+1}, Q_n \geq \lambda b_n/2)$$
$$+ \mathbb{P}(P(\lambda \tau_n) \leq \lambda \tau_n/2, \tau_n > b_n).$$

By another application of Lemma 3.1(a), we have

$$\mathbb{P}(\zeta(Q_n) \leq b_{n+1}, Q_n \geq \lambda b_n/2) \leq \mathbb{P}(\zeta(\lfloor \lambda b_n/2 \rfloor) \leq b_{n+1})$$
$$\leq \mathbb{P}(\zeta(\lfloor \lambda b_n/2 \rfloor) \leq \lfloor \lambda b_n/2 \rfloor^\alpha),$$

for all n sufficiently large, since (i) and (ii) imply that $b_{n+1} < \lfloor \lambda b_n/2 \rfloor^\alpha$ for n large enough. Hence, by Lemma 3.1(b) and the fact that $b_n \to \infty$, for some $\varepsilon > 0$,

$$\mathbb{P}(\zeta(Q_n) \leq b_{n+1}, Q_n \geq \lambda b_n/2) \leq \mathrm{e}^{-b_n^\varepsilon},$$

for all n sufficiently large. On the other hand,

$$\mathbb{P}(P(\lambda \tau_n) \leq \lambda \tau_n/2, \tau_n > b_n) \leq \sup_{s \geq b_n} \mathbb{P}(P(\lambda s) \leq \lambda s/2) \leq \mathrm{e}^{-\delta b_n},$$

for some $\delta > 0$, by standard Poisson tail bounds (see e.g. [16, p. 17]). Combining these estimates and using the fact that $b_n \to \infty$, we get, for some $\varepsilon > 0$,

$$\mathbb{P}(\tau_{n+1} \leq b_{n+1}, \tau_n > b_n) \leq \mathrm{e}^{-b_n^\varepsilon},$$

for all n sufficiently large. Then by (iii) and the Borel–Cantelli lemma, we have that $\{\tau_{n+1} \leq b_{n+1}, \tau_n > b_n\}$ occurs only finitely often, a.s. In other words, for all n sufficiently large we have $\tau_n > b_n$ implies $\tau_{n+1} > b_{n+1}$, and since $\tau_n > b_n$ i.o., the result follows. $\square$

We can now deduce the following lower bound, which, despite being far from best possible, is an important step in proving Proposition 4.1.

COROLLARY 4.5. *Almost surely, for all but finitely many n, $\tau_n \geq n^2$.*

PROOF. We take $b_n = n^2$ in Lemma 4.4 and apply Lemma 4.3. $\qquad\square$

The next result, showing that queues are rarely much shorter than we would expect, will be used a couple of times.

LEMMA 4.6. *Almost surely, for all but finitely many n, $Q_n(X_{n+1}) > \lambda\tau_n - \tau_n^{3/4}$.*

PROOF. Let $Q_n := Q_n(X_{n+1})$. Let $n \in \mathbb{N}$. Then, given $\mathcal{F}_{n-1}$, Q_n stochastically dominates a Poisson random variable with mean $\lambda\tau_n$. Thus

$$\mathbb{P}(Q_n \leq \lambda\tau_n - \tau_n^{3/4} \mid \mathcal{F}_{n-1})$$
$$\leq \mathbb{P}(P(\lambda\tau_n) \leq \lambda\tau_n - \tau_n^{3/4}, \tau_n \geq n \mid \mathcal{F}_{n-1}) + \mathbb{P}(\tau_n \leq n \mid \mathcal{F}_{n-1})$$
$$\leq \sup_{s \geq n} \mathbb{P}(P(\lambda s) \leq \lambda s - s^{3/4}) + \mathbb{P}(\tau_n \leq n \mid \mathcal{F}_{n-1})$$
$$\leq e^{-n^\varepsilon} + \mathbb{P}(\tau_n \leq n \mid \mathcal{F}_{n-1}),$$

for some $\varepsilon > 0$ and all n sufficiently large, by Poisson concentration (see e.g. [16, p. 17]). In particular, since by Corollary 4.5, $\tau_n \leq n$ only finitely often, a.s., Lévy's extension of the Borel–Cantelli lemma (see e.g. [12, Corollary 7.20]) implies that $\sum_{n\geq1} \mathbb{P}(\tau_n \leq n \mid \mathcal{F}_{n-1}) < \infty$, and hence

$$\sum_{n\geq1} \mathbb{P}(Q_n \leq \lambda\tau_n - \tau_n^{3/4} \mid \mathcal{F}_{n-1}) < \infty, \text{ a.s.,}$$

which gives the result. $\qquad\square$

The next result gives the final ingredient in the proof of Proposition 4.1, and a bound that we will use later.

LEMMA 4.7. *Let $\alpha \in (1, 2)$. Then for some $\varepsilon > 0$,*

$$\mathbb{P}(\tau_{n+1} \leq \tau_n^\alpha \mid \mathcal{F}_n) \leq e^{-n^\varepsilon}, \text{ for all but finitely many } n, \text{ a.s.}$$

In particular, a.s., $\tau_{n+1} \geq \tau_n^\alpha$ for all but finitely many n, and $\tau_n/\tau_{n+1} \to 0$, a.s.

PROOF. Let $Q_n := Q_n(X_{n+1})$. Let $\alpha \in (1, 2)$. Given $\mathcal{F}_n$, we have from Lemma 3.1(a) that τ_{n+1} stochastically dominates $\zeta(Q_n)$. Hence

$$\mathbb{P}(\tau_{n+1} \leq \tau_n^\alpha \mid \mathcal{F}_n) \leq \mathbb{P}(Q_n \leq \lambda\tau_n/2 \mid \mathcal{F}_n) + \mathbb{P}(Q_n > \lambda\tau_n/2, \zeta(Q_n) \leq \tau_n^\alpha \mid \mathcal{F}_n).$$

Here by Lemma 3.1(a) once more, we have

$$\mathbb{P}(Q_n > \lambda\tau_n/2, \zeta(Q_n) \leq \tau_n^\alpha \mid \mathcal{F}_n) \leq \mathbb{P}(\zeta(\lfloor\lambda\tau_n/2\rfloor) \leq \tau_n^\alpha \mid \mathcal{F}_n)$$
$$\leq \exp\{-\lfloor\lambda\tau_n/2\rfloor^\varepsilon\},$$

for some $\varepsilon > 0$ and all n sufficiently large, by Lemma 3.1(b). By Corollary 4.5 we have $\tau_n \geq n^2$ for all n sufficiently large, so since Q_n and τ_n are $\mathcal{F}_n$-measurable,

$$\mathbb{P}(\tau_{n+1} \leq \tau_n^\alpha \mid \mathcal{F}_n) \leq \mathbf{1}\{Q_n \leq \lambda\tau_n/2\} + \mathrm{e}^{-n^\varepsilon},$$

for all but finitely many n. It follows from Lemma 4.6 that the indicator here vanishes, a.s., for all but finitely many n. The probability bound in the lemma follows.

A consequence of the probability bound is

$$\sum_{n\geq 1} \mathbb{P}(\tau_{n+1} \leq \tau_n^\alpha \mid \mathcal{F}_n) < \infty, \text{ a.s.}$$

Hence, by Lévy's extension of the Borel–Cantelli lemma, $\tau_{n+1} \geq \tau_n^\alpha$ for all but finitely many n. Moreover, Corollary 4.5 shows that $\tau_n \to \infty$ and hence, for all but finitely many n, $\tau_n/\tau_{n+1} \leq \tau_n^{1-\alpha} \to 0$. $\qquad\square$

We can now complete the proof of Proposition 4.1.

PROOF OF PROPOSITION 4.1. Lemma 4.7 shows that there is some N_1 with $\mathbb{P}(N_1 < \infty) = 1$ such that

$$(4.1) \qquad\qquad\qquad \tau_{n+1} \geq \tau_n^\alpha \text{ for all } n \geq N_1.$$

Since (by Corollary 4.5) $\tau_n \to \infty$, a.s., we have $\tau_N \geq \mathrm{e}$ for some a.s. finite $N \geq N_1$. Then iterating (4.1) we have $\tau_{N+k} \geq \mathrm{e}^{\alpha^k}$ for all $k \geq 0$. Take $\tilde{\alpha} \in (1, \alpha)$. Then

$$\tau_n \geq \mathrm{e}^{\alpha^{n-N}}\mathbf{1}\{n \geq N\} \geq \mathrm{e}^{\tilde{\alpha}^n},$$

for all but finitely many n, giving the result. $\qquad\square$

REMARK. *A posteriori*, armed with Proposition 4.1, one can greatly improve the probability bound in Lemma 4.7; as stated, however, it is adequate for its use later in the paper.

The next result shows that $T_n/\tau_n \to 1$, a.s., and will be useful in the next section as well as in the proof of Proposition 4.2.

LEMMA 4.8. *Almost surely, for all but finitely many n, $T_n/\tau_n \leq 1 + \mathrm{e}^{-6n}$.*

PROOF. Let $\alpha \in (1, 2)$. We have from Lemma 4.7 and Proposition 4.1 that there exists N with $\mathbb{P}(N < \infty) = 1$ such that $\tau_{n+1} \geq \tau_n^\alpha$ and $\tau_n \geq \mathrm{e}^{\alpha^n}$ for all $n \geq N$. Set

$$K := 1 + \max_{1\leq m\leq N-1} \prod_{k=m}^{N-1} \frac{\tau_k}{\tau_{k+1}}.$$

Then since $1 \leq \tau_k < \infty$, a.s., we have that $K < \infty$, a.s. Now, for $n > N$,

$$\max_{1 \leq m \leq n-1} \frac{\tau_m}{\tau_n} = \max_{1 \leq m \leq n-1} \prod_{k=m}^{n-1} \frac{\tau_k}{\tau_{k+1}}$$

$$\leq \max_{1 \leq m \leq n-1} \prod_{k=m}^{N-1} \frac{\tau_k}{\tau_{k+1}} \prod_{k=N}^{n-1} \frac{\tau_k}{\tau_{k+1}} \mathbf{1}\{m \leq N-1\}$$

$$+ \max_{1 \leq m \leq n-1} \prod_{k=m}^{n-1} \frac{\tau_k}{\tau_{k+1}} \mathbf{1}\{m \geq N\}.$$

For $k \geq N$ we have that $\tau_k/\tau_{k+1} \leq \tau_k^{1-\alpha} \leq 1$, so for any $m \in \{N, \ldots, n-1\}$,

$$\prod_{k=N}^{n-1} \frac{\tau_k}{\tau_{k+1}} \leq \prod_{k=m}^{n-1} \frac{\tau_k}{\tau_{k+1}} \leq \frac{\tau_{n-1}}{\tau_n}.$$

Hence, for $n > N$,

$$\max_{1 \leq m \leq n-1} \frac{\tau_m}{\tau_n} \leq K \frac{\tau_{n-1}}{\tau_n} \leq K \tau_{n-1}^{1-\alpha} \leq K \mathrm{e}^{-(\alpha-1)\alpha^{n-1}}.$$

It follows that, a.s.,

$$\max_{1 \leq m \leq n-1} \frac{\tau_m}{\tau_n} \leq K \mathrm{e}^{-8n} \leq \mathrm{e}^{-7n},$$

for all but finitely many n. Now the result follows from the fact that

$$\frac{T_n}{\tau_n} = 1 + \sum_{m=1}^{n-1} \frac{\tau_m}{\tau_n} \leq 1 + n \max_{1 \leq m \leq n-1} \frac{\tau_m}{\tau_n}.$$

This completes the proof. $\qquad\square$

We also need a complementary result to Lemma 4.6.

LEMMA 4.9. *Almost surely, for all but finitely many n,*

$$Q_n(X_{n+1}) < \lambda(1 + \mathrm{e}^{-5n})\tau_n.$$

PROOF. At time T_n, the queue at X_n is emptied and the queues at $X_n \pm 1$ are inspected; let $Q_L = Q_n(X_n - 1)$ and $Q_R = Q_n(X_n + 1)$. Then $Q_n := Q_n(X_{n+1}) = \max\{Q_L, Q_R\}$. Suppose that the queue at $X_n - 1$ was most recently emptied at some time $T_L < T_n$, and that the queue at $X_n + 1$ was most recently emptied at some time $T_R < T_n$. After the time at which it was most recently emptied, each queue has been inspected a finite number of times, and, because the queue was not served at any point after it was last emptied, on each inspection it was found to be no larger than the queue to which it was being compared. Each such

inspection therefore (see Lemma A.2) stochastically reduces the queue length. Thus immediately before the inspection at time T_n, we have that Q_L is stochastically dominated by $P(\lambda(T_n - T_L))$ and Q_R is stochastically dominated by $P(\lambda(T_n - T_R))$. It follows that Q_n is stochastically dominated by the maximum of two $P(\lambda T_n)$ random variables. Thus

$$\mathbb{P}(Q_n \geq \lambda T_n + T_n^{3/4}) \leq 2\mathbb{P}(P(\lambda T_n) \geq \lambda T_n + T_n^{3/4}).$$

Now, since $T_n \geq n$ a.s., we get

$$\mathbb{P}(P(\lambda T_n) \geq \lambda T_n + T_n^{3/4}) \leq \sup_{s \geq n} \mathbb{P}(P(\lambda s) \geq \lambda s + s^{3/4}) \leq e^{-n^\varepsilon},$$

for some $\varepsilon > 0$ and all n sufficiently large, by standard Poisson tail bounds (see e.g. [16, p. 17]). Thus the Borel–Cantelli lemma implies that $Q_n \leq \lambda T_n + T_n^{3/4}$ for all but finitely many n, a.s. Lemma 4.8 shows that $T_n \leq (1 + e^{-6n})\tau_n$ for all but finitely many n, so

$$Q_n \leq \lambda(1 + e^{-6n})\tau_n + 2\tau_n^{3/4}.$$

Now, by Proposition 4.1, for all but finitely many n,

$$2\tau_n^{3/4} = 2\tau_n \cdot \tau_n^{-1/4} \leq \tau_n \cdot e^{-7n} \leq \lambda e^{-6n}\tau_n,$$

which gives the result. $\square$

Now we can complete the proof of Proposition 4.2.

PROOF OF PROPOSITION 4.2. Let $\beta > 2$ and set $Q_n := Q_n(X_{n+1})$. Given $\mathcal{F}_n$, τ_{n+1} is distributed as $1 + \zeta(Q_n + \nu)$ where $\nu \sim P(\lambda)$. Thus

$$\begin{aligned}
\mathbb{P}(\tau_{n+1} > T_n^\beta \mid \mathcal{F}_n) &= \mathbb{P}(1 + \zeta(Q_n + \nu) > T_n^\beta \mid \mathcal{F}_n) \\
&\leq \mathbb{P}(\zeta(2Q_n) > T_n^\beta - 1, \nu \leq Q_n \mid \mathcal{F}_n) + \mathbb{P}(\nu > Q_n \mid \mathcal{F}_n).
\end{aligned}$$

We have by Markov's inequality and Lemma 4.6 that

$$\mathbb{P}(\nu > Q_n \mid \mathcal{F}_n) \leq \frac{\lambda}{Q_n} \leq \frac{2}{\tau_n} \leq e^{-2n},$$

for all but finitely many n, a.s., by Proposition 4.1. On the other hand,

$$\begin{aligned}
\mathbb{P}(\zeta(2Q_n) > T_n^\beta - 1, \nu \leq Q_n \mid \mathcal{F}_n) &\leq \mathbb{P}(\zeta(2Q_n) > T_n^\beta - 1, Q_n \leq 2\lambda T_n \mid \mathcal{F}_n) \\
&\quad + \mathbb{P}(Q_n > 2\lambda T_n \mid \mathcal{F}_n) \\
&\leq \mathbb{P}(\zeta(\lceil 4\lambda T_n \rceil) > T_n^\beta - 1 \mid \mathcal{F}_n) + \mathbf{1}\{Q_n > 2\lambda T_n\} \\
&\leq T_n^{-\varepsilon},
\end{aligned}$$

for some $\varepsilon > 0$ and all but finitely many n, a.s., by Lemma 4.9, Lemma 3.1(c), and the fact that $T_n \geq \tau_n \to \infty$. Hence, since $T_n \geq \tau_n$, we have from Proposition 4.1 that a.s.,

$$\mathbb{P}(\tau_{n+1} > T_n^\beta \mid \mathcal{F}_n) \leq \mathrm{e}^{-n},$$

for all but finitely many n. It follows that $\tau_{n+1} \leq T_n^\beta$ for all but finitely many n, a.s. Let $\tilde{\beta} > \beta$. Then Lemma 4.8 shows that $T_{n+1} \leq 2T_n^\beta \leq T_n^{\tilde{\beta}}$, for all $n \geq N$ with $\mathbb{P}(N < \infty) = 1$. It follows that $T_n \leq T_N^{\tilde{\beta}^n}$ for all n. Since $\tilde{\beta} > 2$ was arbitrary, the result follows. $\qquad\square$

5. Turning probability. For $n \in \mathbb{N}$ define

(5.1) $$q_n := \mathbb{P}(\eta_{n+1} \neq \eta_n \mid \mathcal{F}_{n-1}).$$

The main result of this section is the following.

PROPOSITION 5.1. *Let $q := 1/4$. Then there exists $\varepsilon > 0$ such that, a.s., for all but finitely many n, $|q_n - q| \leq \mathrm{e}^{-n^\varepsilon}$.*

We work towards the proof of Proposition 5.1. We need the following result.

LEMMA 5.2. *Let S be a random variable with the distribution given by (3.4). Then, a.s., for all but finitely many n,*

$$\sup_{u \in [0, \mathrm{e}^n]} \left| \mathbb{P}\left(\frac{\tau_n}{\tau_{n-1}^2} \leq u \;\middle|\; \mathcal{F}_{n-1} \right) - F_S(2u/\lambda) \right| \leq \mathrm{e}^{-2n}.$$

PROOF. We have from Lemma 3.1(a) that, given $\mathcal{F}_{n-1}$, τ_n stochastically dominates $\zeta(Q_{n-1})$, where $Q_{n-1} := Q_{n-1}(X_n)$. It follows that for $u \in \mathbb{R}_+$,

$$\begin{aligned}
\mathbb{P}(\tau_n \leq u\tau_{n-1}^2 \mid \mathcal{F}_{n-1}) &\leq \mathbb{P}(\zeta(Q_{n-1}) \leq u\tau_{n-1}^2 \mid \mathcal{F}_{n-1}) \\
&\leq \mathbb{P}(\zeta(Q_{n-1}) \leq u\tau_{n-1}^2, Q_{n-1} \geq \lambda\tau_{n-1} - \tau_{n-1}^{3/4} \mid \mathcal{F}_{n-1}) \\
&\quad + \mathbb{P}(Q_{n-1} \leq \lambda\tau_{n-1} - \tau_{n-1}^{3/4} \mid \mathcal{F}_{n-1}) \\
&\leq \mathbb{P}(\zeta(\lfloor \lambda\tau_{n-1} - \tau_{n-1}^{3/4} \rfloor) \leq u\tau_{n-1}^2 \mid \mathcal{F}_{n-1}),
\end{aligned}$$

for all but finitely many n, a.s., by Lemma 4.6. Write $\gamma_{n-1} := \tau_{n-1} - \lambda^{-1}\tau_{n-1}^{3/4}$. Then

$$\begin{aligned}
&\mathbb{P}(\zeta(\lfloor \lambda\gamma_{n-1} \rfloor) \leq u\tau_{n-1}^2 \mid \mathcal{F}_{n-1}) \\
&\leq \mathbb{P}(\lfloor \lambda\gamma_{n-1} \rfloor^{-2} \zeta(\lfloor \lambda\gamma_{n-1} \rfloor) \leq u(\lambda\gamma_{n-1} - 1)^{-2}\tau_{n-1}^2 \mid \mathcal{F}_{n-1}) \\
&\leq F_S(2u\lambda(\lambda\gamma_{n-1} - 1)^{-2}\tau_{n-1}^2) + C\gamma_{n-1}^{-1},
\end{aligned}$$

for some $C < \infty$ not depending on u or n, by Lemma 3.2. By Proposition 4.1 we have $\gamma_{n-1}^{-1} \leq \mathrm{e}^{-3n}$ for all but finitely many n, while

$$|(\lambda\gamma_{n-1} - 1)^{-2}\tau_{n-1}^2 - \lambda^{-2}| \leq C\tau_{n-1}^{-1/4} \leq \mathrm{e}^{-4n},$$

for all but finitely many n. Since the density of S in (3.5) is uniformly bounded,

$$F_S(2u\lambda(\lambda\gamma_{n-1} - 1)^{-2}\tau_{n-1}^2) \leq F_S(2\lambda^{-1}u) + Cue^{-4n},$$

for all u. Combining our estimates gives, a.s., for all but finitely many n,

$$(5.2) \qquad \mathbb{P}(\tau_n \leq u\tau_{n-1}^2 \mid \mathcal{F}_{n-1}) \leq F_S(2\lambda^{-1}u) + e^{-2n}, \text{ for all } u \in [0, e^n],$$

which is one half of the required result.

For the corresponding lower bound, we have that, given $\mathcal{F}_{n-1}$, τ_n is distributed as $1 + \zeta(Q_{n-1} + \nu)$ where $\nu \sim P(\lambda)$. Then, by Lemma 3.1(a),

$$\mathbb{P}(\tau_n \leq u\tau_{n-1}^2 \mid \mathcal{F}_{n-1}) \geq \mathbb{P}(1 + \zeta(Q_{n-1} + \lfloor e^n \rfloor) \leq u\tau_{n-1}^2, \nu \leq e^n \mid \mathcal{F}_{n-1})$$
$$\geq \mathbb{P}(\zeta(Q_{n-1} + \lfloor e^n \rfloor) \leq u\tau_{n-1}^2 - 1 \mid \mathcal{F}_{n-1}) - \mathbb{P}(\nu > e^n \mid \mathcal{F}_{n-1}).$$

By Markov's inequality, $\mathbb{P}(\nu > e^n \mid \mathcal{F}_{n-1}) \leq \lambda e^{-n}$. Let $E_n = \{Q_n \leq \lambda(1 + e^{-5n})\tau_n\}$. Then we have that

$$\mathbb{P}(\zeta(Q_{n-1} + \lfloor e^n \rfloor) \leq u\tau_{n-1}^2 - 1, E_{n-1} \mid \mathcal{F}_{n-1})$$
$$\geq \mathbb{P}(\zeta(\lfloor \lambda(1 + e^{-5(n-1)})\tau_{n-1} \rfloor + \lfloor e^n \rfloor) \leq u\tau_{n-1}^2 - 1 \mid \mathcal{F}_{n-1}) - \mathbf{1}(E_{n-1}^{\mathrm{c}})$$
$$\geq \mathbb{P}(\zeta(\lfloor \lambda(1 + e^{-5(n-1)})\tau_{n-1} + e^n \rfloor) \leq u\tau_{n-1}^2 - 1 \mid \mathcal{F}_{n-1}),$$

for all but finitely many n, a.s., by Lemma 4.9. Set $\gamma_{n-1} = (1 + e^{-5(n-1)})\tau_{n-1} + \lambda^{-1}e^n$. Then

$$\mathbb{P}(\zeta(\lfloor \lambda\gamma_{n-1} \rfloor) \leq u\tau_{n-1}^2 - 1 \mid \mathcal{F}_{n-1})$$
$$\geq \mathbb{P}(\lfloor \lambda\gamma_{n-1} \rfloor^{-2}\zeta(\lfloor \lambda\gamma_{n-1} \rfloor) \leq \lambda^{-2}\gamma_{n-1}^{-2}(u\tau_{n-1}^2 - 1) \mid \mathcal{F}_{n-1})$$
$$\geq F_S(2\lambda^{-1}\gamma_{n-1}^{-2}(u\tau_{n-1}^2 - 1)) - e^{-3n},$$

for all but finitely many n, a.s., by Lemma 3.2 and Proposition 4.1, since $\gamma_{n-1} \geq \tau_{n-1}$. Proposition 4.1 also shows that $\gamma_{n-1} \geq \tau_{n-1} > e^{2n}$ and

$$1 \geq \gamma_{n-1}^{-2}\tau_{n-1}^2 \geq 1 - e^{-4n},$$

for all but finitely many n, a.s. Thus

$$F_S(2\lambda^{-1}\gamma_{n-1}^{-2}(u\tau_{n-1}^2 - 1)) \geq F_S(2\lambda^{-1}u - 2\lambda^{-1}ue^{-4n} - 2\lambda^{-1}e^{-4n}) \geq F_S(2\lambda^{-1}u) - e^{-2n},$$

provided $u \in [0, e^n]$, using the fact that the density of S given by (3.5) is uniformly bounded. Combined with (5.2) this completes the proof. $\qquad \square$

The origin of the value $1/4$ in Proposition 5.1 is the following fact.

LEMMA 5.3. *Let S be a random variable with the distribution given by* (3.4), *and let Z be a standard normal random variable, independent of S. Then*

$$\mathbb{P}(Z\sqrt{S} > 1) = \frac{1}{4}.$$

PROOF. Recall that Φ denotes the distribution function of the standard normal distribution, and that $\bar{\Phi}(u) := 1 - \Phi(u)$. Observe that

$$p := \mathbb{P}(Z > S^{-1/2}) = \mathbb{E}[\mathbb{P}(Z > S^{-1/2} \mid S)] = \mathbb{E}\,\bar{\Phi}(S^{-1/2}).$$

Here $\bar{\Phi} : \mathbb{R} \to (0,1)$ is a strictly decreasing, continuous function with $\lim_{x\to-\infty}\bar{\Phi}(x) = 1$ and $\lim_{x\to+\infty}\bar{\Phi}(x) = 0$, so $\bar{\Phi}$ has a strictly decreasing inverse $\bar{\Phi}^{-1} : (0,1) \to \mathbb{R}$, which is positive on $(0, 1/2)$ and negative on $(1/2, 1)$, and

$$p = \int_0^1 \mathbb{P}(\bar{\Phi}(S^{-1/2}) \geq u)\mathrm{d}u = \int_0^{1/2} \mathbb{P}(S \geq (\bar{\Phi}^{-1}(u))^{-2})\mathrm{d}u$$

$$(5.3) \qquad = \int_0^{1/2} \bar{F}_S((\bar{\Phi}^{-1}(u))^{-2})\mathrm{d}u,$$

where $\bar{F}_S(u) := 1 - F_S(u)$ for $u \in \mathbb{R}_+$. Now applying the formula in (3.4) we get $\bar{F}_S((\bar{\Phi}^{-1}(u))^{-2}) = 1 - 2\bar{\Phi}(\bar{\Phi}^{-1}(u))$, so that

$$p = \int_0^{1/2}(1 - 2u)\mathrm{d}u = \frac{1}{4},$$

as claimed. $\qquad\square$

REMARK. As pointed out by the anonymous referee, a more elegant proof of Lemma 5.3 goes by observing that $Z^{-2} \overset{d}{=} S$, so that, for Z_1, Z_2 independent standard normal random variables, $\mathbb{P}(Z\sqrt{S} > 1) = \mathbb{P}(Z_1 > |Z_2|)$, which is the probability that a bivariate normal random variable lies in the quadrant $\{(z_1, z_2) : |z_2| < z_1\}$. However, we present the argument via (5.3), since it will be used in the proof of Proposition 5.1 below.

Now we can complete the proof of Proposition 5.1.

PROOF OF PROPOSITION 5.1. Given $\mathcal{F}_{n-1}$, we have that $Q_{n-1}(X_{n-1}) = 0$ at time T_{n-1}, and the server then heads to $X_n = X_{n-1} + \eta_n$. At time T_n, after serving the queue at X_n, the server inspects the queues at $X_n \pm 1$.

First we obtain an upper bound on q_n, the probability that the server changes direction. We have

$$q_n \leq \mathbb{P}(Q_n(X_n - \eta_n) \geq Q_n(X_n + \eta_n) \mid \mathcal{F}_{n-1}).$$

Note that $Q_{n-1}(X_n - \eta_n) = Q_{n-1}(X_{n-1}) = 0$, i.e., the queue at $X_n - \eta_n$ was empty at time T_{n-1}. Thus $Q_n(X_n - \eta_n)$ is Poisson with mean λ times $\tau_n = T_n - T_{n-1}$. Set $\nu_n := Q_n(X_n - \eta_n)$.

The queue at $X_n + \eta_n$, which is neither X_n nor $X_{n-1} = X_n - \eta_n$, was, prior to time T_n, last inspected by the server no more recently than at time T_{n-2} (when the

server decided to move to X_{n-1}). Thus the number of customers at the queue at time T_n is at least $\nu'_n + \nu''_n$ where

$$\nu'_n := Q_n(X_n + \eta_n) - Q_{n-1}(X_n + \eta_n), \text{ and } \nu''_n := Q_{n-1}(X_n + \eta_n) - Q_{n-2}(X_n + \eta_n);$$

note that these are well defined since $|X_{n-2} - (X_n + \eta_n)| \leq 3$. So we have

$$q_n \leq \mathbb{P}(\nu_n \geq \nu'_n + \nu''_n \mid \mathcal{F}_{n-1}),$$

where $\nu_n \sim P(\lambda \tau_n)$ and $\nu'_n \sim P(\lambda \tau_n)$ are both $\mathcal{F}_n$-measurable, and are conditionally independent given τ_n, and $\nu''_n \sim P(\lambda \tau_{n-1})$ is $\mathcal{F}_{n-1}$-measurable. Define

$$\begin{aligned}
Z_n &= (\lambda \tau_n)^{-1/2}(\nu_n - \lambda \tau_n); \\
Z'_n &= (\lambda \tau_n)^{-1/2}(\nu'_n - \lambda \tau_n); \\
Z''_n &= (\lambda \tau_{n-1})^{-1/2}(\nu''_n - \lambda \tau_{n-1}).
\end{aligned}$$

Then we get

$$q_n \leq \mathbb{P}(Z_n \geq Z'_n + (\tau_{n-1}/\tau_n)^{1/2} Z''_n + \lambda^{1/2} \tau_n^{-1/2} \tau_{n-1} \mid \mathcal{F}_{n-1}).$$

Hence, writing $W_n := Z_n - Z'_n$, we have

$$\begin{aligned}
q_n &\leq \mathbb{P}(W_n \geq \lambda^{1/2} \tau_n^{-1/2} \tau_{n-1} - \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) \\
&\quad + \mathbb{P}(|Z''_n|(\tau_{n-1}/\tau_n)^{1/2} \geq \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}),
\end{aligned}$$

(5.4)

where, for $p \in (0, 1/2)$,

$$\begin{aligned}
\mathbb{P}(|Z''_n|(\tau_{n-1}/\tau_n)^{1/2} \geq \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) &\leq \mathbb{P}(|Z''_n| \geq \tau_{n-1}^p \mid \mathcal{F}_{n-1}) \\
&\quad + \mathbb{P}(\tau_n \leq \tau_{n-1}^{1+2p} \mathrm{e}^{2n} \mid \mathcal{F}_{n-1}).
\end{aligned}$$

Since Z''_n is $\mathcal{F}_{n-1}$-measurable, we have that $\mathbb{P}(|Z''_n| \geq \tau_{n-1}^p \mid \mathcal{F}_{n-1}) = \mathbf{1}\{|Z''_n| \geq \tau_{n-1}^p\}$. We show that this event occurs only finitely often. Set $\mathcal{F}_{n-2}^+ := \sigma(\mathcal{F}_{n-2}, \tau_{n-1})$; note that ν''_n (and hence Z''_n) only depend on $\mathcal{F}_{n-2}^+$ via τ_{n-1}. We have that

$$\begin{aligned}
\mathbb{P}(|Z''_n| \geq \tau_{n-1}^p \mid \mathcal{F}_{n-2}^+) &\leq \mathbb{P}(|Z''_n| \geq \tau_{n-1}^p, \tau_{n-1} \geq n \mid \mathcal{F}_{n-2}^+) + \mathbf{1}\{\tau_{n-1} \leq n\} \\
&\leq \mathbb{P}(|Z''_n| \geq \tau_{n-1}^p, \tau_{n-1} \geq n \mid \mathcal{F}_{n-2}^+),
\end{aligned}$$

for all but finitely many n, a.s., by Proposition 4.1. Now

$$\begin{aligned}
\mathbb{P}(|Z''_n| \geq \tau_{n-1}^p, \tau_{n-1} \geq n \mid \mathcal{F}_{n-2}^+) &= \mathbb{P}(|\nu''_n - \lambda \tau_{n-1}| \geq \lambda^{1/2} \tau_{n-1}^{(1/2)+p}, \tau_{n-1} \geq n \mid \mathcal{F}_{n-2}^+) \\
&\leq \sup_{s \geq n} \mathbb{P}(|P(\lambda s) - \lambda s| \geq \lambda^{1/2} s^{(1/2)+p}) \\
&\leq \mathrm{e}^{-n^\varepsilon},
\end{aligned}$$

for some $\varepsilon > 0$ and all n sufficiently large, by Poisson concentration (see e.g. [16, p. 17]). So we conclude that $\sum_{n \geq 2} \mathbb{P}(|Z_n''| \geq \tau_{n-1}^p \mid \mathcal{F}_{n-2}^+) < \infty$, a.s., and, since $\{|Z_n''| \geq \tau_{n-1}^p\} \in \mathcal{F}_{n-1}^+$, Lévy's extension of the Borel–Cantelli lemma implies that $|Z_n''| < \tau_{n-1}^p$ for all but finitely many n, a.s. On the other hand, we have from Proposition 4.1 that a.s., for all but finitely many n, $\tau_{n-1}^{1+2p} \mathrm{e}^{2n} \leq \tau_{n-1}^\alpha$, provided that we choose $\alpha \in (1 + 2p, 2)$. Hence by Lemma 4.7 we have for some $\varepsilon > 0$, a.s.,

$$\mathbb{P}(\tau_n \leq \tau_{n-1}^{1+2p} \mathrm{e}^{2n} \mid \mathcal{F}_{n-1}) \leq \mathbb{P}(\tau_n \leq \tau_{n-1}^\alpha \mid \mathcal{F}_{n-1}) \leq \mathrm{e}^{-n^\varepsilon},$$

for all n sufficiently large. Combining these estimates we get that for some $\varepsilon > 0$, a.s., for all but finitely many n,

$$(5.5) \qquad \mathbb{P}(|Z_n''|(\tau_{n-1}/\tau_n)^{1/2} \geq \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) \leq \mathrm{e}^{-n^\varepsilon}.$$

Thus from (5.4) with (5.5) we see that, a.s., for all but finitely many n,

$$(5.6) \qquad q_n \leq \mathbb{P}(W_n \geq \lambda^{1/2} \tau_n^{-1/2} \tau_{n-1} - \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) + \mathrm{e}^{-n^\varepsilon}.$$

For $u \in \mathbb{R}$ write $G_n(u) := \mathbb{P}(W_n \geq u \mid \tau_n)$. Set $S_n := \tau_n/\tau_{n-1}^2$ and, as above, set $\mathcal{F}_{n-1}^+ := \sigma(\mathcal{F}_{n-1}, \tau_n)$. Then we have

$$\mathbb{P}(W_n \geq \lambda^{1/2} S_n^{-1/2} - \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) = \mathbb{E}\left(\mathbb{P}(W_n \geq \lambda^{1/2} S_n^{-1/2} - \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}^+) \mid \mathcal{F}_{n-1}\right).$$

Since W_n depends on $\mathcal{F}_{n-1}^+$ only through τ_n, and S_n is $\mathcal{F}_{n-1}^+$-measurable, we have

$$\mathbb{P}(W_n \geq \lambda^{1/2} S_n^{-1/2} - \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}^+) = G_n(\lambda^{1/2} S_n^{-1/2} - \mathrm{e}^{-n}).$$

Since $W_n = (\lambda \tau_n)^{-1/2}(\nu_n - \nu_n')$ we have from Lemma A.1 and the fact that $\tau_n \geq 1$ that

$$G_n(u) \leq \bar{\Phi}(u/\sqrt{2}) + C\tau_n^{-1/2} \log(1 + \tau_n), \quad \text{a.s.},$$

for all n and all $u \in \mathbb{R}$. It follows that

$$
\begin{aligned}
\mathbb{P}(W_n \geq \lambda^{1/2} S_n^{-1/2} - \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) &= \mathbb{E}(G_n(\lambda^{1/2} S_n^{-1/2} - \mathrm{e}^{-n}) \mid \mathcal{F}_{n-1}) \\
&\leq \mathbb{E}(\bar{\Phi}(\lambda^{1/2} S_n^{-1/2}/\sqrt{2} - \mathrm{e}^{-n}) \mid \mathcal{F}_{n-1}) \\
(5.7) &\qquad + C\,\mathbb{E}(\tau_n^{-1/2} \log(1 + \tau_n) \mid \mathcal{F}_{n-1}),
\end{aligned}
$$

for all n sufficiently large. Here we have, since $\tau_n \geq 1$, for $\alpha \in (1, 2)$, a.s.,

$$(5.8) \quad \mathbb{E}(\tau_n^{-1/2} \log(1 + \tau_n) \mid \mathcal{F}_{n-1}) \leq \mathbb{P}(\tau_n \leq \tau_{n-1}^\alpha \mid \mathcal{F}_{n-1}) + \tau_{n-1}^{-\alpha/3} \leq \mathrm{e}^{-n^\varepsilon} + \mathrm{e}^{-n},$$

for all but finitely many n, by Lemma 4.7 and Proposition 4.1. Moreover, since the standard normal density is uniformly bounded, we have that for some $C < \infty$ and all $u \in \mathbb{R}$, $\bar{\Phi}(u - \mathrm{e}^{-n}) \leq \bar{\Phi}(u) + C\mathrm{e}^{-n}$. Thus, for some $\varepsilon > 0$, a.s.,

$$(5.9) \quad \mathbb{P}(W_n \geq \lambda^{1/2} S_n^{-1/2} - \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) \leq \mathbb{E}(\bar{\Phi}(\lambda^{1/2} S_n^{-1/2}/\sqrt{2}) \mid \mathcal{F}_{n-1}) + \mathrm{e}^{-n^\varepsilon},$$

for all but finitely many n.

Similarly to (5.3), we have

$$\mathbb{E}(\bar{\Phi}(\lambda^{1/2}S_n^{-1/2}/\sqrt{2}) \mid \mathcal{F}_{n-1}) = \int_0^{1/2} \mathbb{P}(S_n \geq \lambda(\bar{\Phi}^{-1}(u))^{-2}/2 \mid \mathcal{F}_{n-1})\mathrm{d}u.$$

Set $a_n := \bar{\Phi}((\lambda/2)^{1/2}\mathrm{e}^{-n/2})$; then $a_n \in (0,1/2)$ with $a_n \to 1/2$, and $\lambda(\bar{\Phi}^{-1}(u))^{-2}/2 \in [0,\mathrm{e}^n]$ for $u \in (0,a_n)$. Thus, by Lemma 5.2,

$$\mathbb{E}(\bar{\Phi}(\lambda^{1/2}S_n^{-1/2}/\sqrt{2}) \mid \mathcal{F}_{n-1}) \leq \int_0^{a_n} \left(\bar{F}_S((\bar{\Phi}^{-1}(u))^{-2}) + \mathrm{e}^{-2n}\right)\mathrm{d}u + \frac{1}{2} - a_n$$

$$(5.10) \qquad\qquad\qquad\qquad\qquad \leq q + \mathrm{e}^{-2n} + C\mathrm{e}^{-n/2},$$

by (5.3) and the fact that the standard normal density is uniformly bounded; here $q = 1/4$ is the probability in (5.3) and Lemma 5.3. Combining (5.6), (5.9), and (5.10) we obtain $q_n \leq q + \mathrm{e}^{-n^\varepsilon}$ for all but finitely many n, a.s.

Now we obtain a lower bound on q_n. In addition to ν_n, ν'_n, ν''_n defined above, also define $\nu'''_n := Q_{n-2}(X_n + \eta_n)$. Then ν'''_n is $\mathcal{F}_{n-1}$-measurable. With W_n and Z''_n as defined above, we have

$$\begin{aligned}
q_n &\geq \mathbb{P}(Q_n(X_n - \eta_n) > Q_n(X_n + \eta_n) \mid \mathcal{F}_{n-1}) \\
&= \mathbb{P}(\nu_n > \nu'_n + \nu''_n + \nu'''_n \mid \mathcal{F}_{n-1}) \\
&= \mathbb{P}(W_n > (\tau_{n-1}/\tau_n)^{1/2}Z''_n + \lambda^{1/2}\tau_n^{-1/2}\tau_{n-1} + \lambda^{-1/2}\tau_n^{-1/2}\nu'''_n \mid \mathcal{F}_{n-1}) \\
&\geq \mathbb{P}(W_n > \lambda^{1/2}\tau_n^{-1/2}\tau_{n-1} + \mathrm{e}^{-n} + \lambda^{-1/2}\tau_n^{-1/2}\nu'''_n \mid \mathcal{F}_{n-1}) \\
&\quad - \mathbb{P}((\tau_{n-1}/\tau_n)^{1/2}|Z''_n| \geq \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}).
\end{aligned}$$

Applying (5.5), we see that a.s., for all but finitely many n,

$$\begin{aligned}
q_n &\geq \mathbb{P}(W_n > \lambda^{1/2}\tau_n^{-1/2}\tau_{n-1} + \mathrm{e}^{-n} + \lambda^{-1/2}\tau_n^{-1/2}\nu'''_n \mid \mathcal{F}_{n-1}) - \mathrm{e}^{-n^\varepsilon} \\
&\geq \mathbb{P}(W_n > \lambda^{1/2}\tau_n^{-1/2}\tau_{n-1} + 2\mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) - \mathbb{P}(\lambda^{-1/2}\tau_n^{-1/2}\nu'''_n \geq \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) - \mathrm{e}^{-n^\varepsilon}.
\end{aligned}$$

Here we have that, a.s., for all but finitely many n,

$$\begin{aligned}
\mathbb{P}(\lambda^{-1/2}\tau_n^{-1/2}\nu'''_n \geq \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) &\leq \mathbb{P}(\nu'''_n \geq \lambda^{1/2}\tau_{n-1}^{3/4}\mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) + \mathbb{P}(\tau_n \leq \tau_{n-1}^{3/2} \mid \mathcal{F}_{n-1}) \\
&\leq \mathbf{1}\{\nu'''_n \geq \lambda^{1/2}\tau_{n-1}^{3/4}\mathrm{e}^{-n}\} + \mathrm{e}^{-n^\varepsilon},
\end{aligned}$$

by Lemma 4.7. By Lemma 4.8 and Lemma 4.7 again, we have that a.s., for all but finitely many n, $T_{n-2} \leq 2\tau_{n-2}$ and $\tau_{n-2} \leq \tau_{n-1}^{2/3}$, so, by Proposition 4.1, since $\tau_{n-1}^{1/12}\mathrm{e}^{-n} \to \infty$,

$$\lambda^{1/2}\tau_{n-1}^{3/4}\mathrm{e}^{-n} = \lambda^{1/2}\tau_{n-1}^{1/12}\mathrm{e}^{-n}\tau_{n-1}^{2/3} \geq 2\lambda T_{n-2}.$$

Thus, a.s., for all but finitely many n,

$$(5.11) \qquad\qquad \mathbf{1}\{\nu'''_n \geq \lambda^{1/2}\tau_{n-1}^{3/4}\mathrm{e}^{-n}\} \leq \mathbf{1}\{\nu'''_n \geq 2\lambda T_{n-2}\}.$$

The queue at $X_n + \eta_n$, which is neither X_n nor $X_{n-1} = X_n - \eta_n$, was, prior to time T_n, last inspected by the server no more recently than at time T_{n-2}, at which point the server decided to move to X_{n-1} (and not $X_n + \eta_n$). Thus ν_n''' is stochastically dominated by $P(\lambda T_{n-2})$, so

$$\mathbb{P}(\nu_n''' \geq 2\lambda T_{n-2}) \leq \mathbb{P}(P(\lambda T_{n-2}) \geq 2\lambda T_{n-2})$$
$$\leq \sup_{s \geq n-2} \mathbb{P}(P(\lambda s) \geq 2\lambda s),$$

using the fact that $T_{n-2} \geq n - 2$, a.s. Then by standard Poisson tail bounds (see e.g. [16, p. 17]) we have that this last quantity is bounded by $\mathrm{e}^{-\delta n}$ for some $\delta > 0$ and all n sufficiently large. Hence the Borel–Cantelli lemma shows that the indicator random variable in (5.11) is a.s. equal to 0 for all but finitely many n. Thus, a.s., for all but finitely many n,

$$\mathbb{P}(\lambda^{-1/2}\tau_n^{-1/2}\nu_n''' \geq \mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) \leq \mathrm{e}^{-n^\varepsilon}.$$

It follows that, for some $\varepsilon > 0$, a.s., for all but finitely many n,

$$q_n \geq \mathbb{P}(W_n > \lambda^{1/2}\tau_n^{-1/2}\tau_{n-1} + 2\mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) - \mathrm{e}^{-n^\varepsilon}.$$

The estimation of the main term here proceeds in a similar way to in the upper bound. Similarly to (5.7) and (5.8), we have that

$$\mathbb{P}(W_n > \lambda^{1/2}\tau_n^{-1/2}\tau_{n-1} + 2\mathrm{e}^{-n} \mid \mathcal{F}_{n-1}) \geq \mathbb{E}(\bar{\Phi}(\lambda^{1/2}S_n^{-1/2}/\sqrt{2} + 2\mathrm{e}^{-n}) \mid \mathcal{F}_{n-1}) - \mathrm{e}^{-n^\varepsilon}$$
$$\geq \mathbb{E}(\bar{\Phi}(\lambda^{1/2}S_n^{-1/2}/\sqrt{2}) \mid \mathcal{F}_{n-1}) - C\mathrm{e}^{-n} - \mathrm{e}^{-n^\varepsilon}.$$

Finally, similarly to (5.10), we have

$$\mathbb{E}(\bar{\Phi}(\lambda^{1/2}S_n^{-1/2}/\sqrt{2}) \mid \mathcal{F}_{n-1}) \geq \int_0^{a_n} (\bar{F}_S((\bar{\Phi}^{-1}(u))^{-2}) - \mathrm{e}^{-2n}) \geq q - \mathrm{e}^{-2n} - C\mathrm{e}^{-n/2},$$

and this gives $q_n \geq q - \mathrm{e}^{-n^\varepsilon}$, as required. $\qquad\square$

6. Proofs of theorems. With $q = 1/4$ as appearing in Proposition 5.1, set $a := \frac{1-2q}{q} = 2$. To prove Theorem 1.1 we consider the function defined for $x \in \mathbb{Z}$ and $i \in \{-1, +1\}$ by

$$f(x, i) := x + a\mathbf{1}\{i = 1\}.$$

We consider $Y_n := f(X_n, \eta_n)$ for $n \in \mathbb{N}$; recall that (X_n, η_n) is $\mathcal{F}_{n-1}$-measurable. Note that, for all $n \in \mathbb{N}$, $|X_n - Y_n| \leq a$. The next result describes the increments of Y_n, and, in particular, shows that it is close to a martingale.

LEMMA 6.1. *Let q_n be the $\mathcal{F}_{n-1}$-measurable random variable defined in (5.1).*

(a) We have that, for all $n \in \mathbb{N}$,

$$(6.1) \qquad |Y_{n+1} - Y_n| \leq 3, \ \text{a.s.}$$

(b) *There is a sequence δ_n of non-negative $\mathcal{F}_{n-1}$-adapted random variables such that,*

$$(6.2) \qquad |\mathbb{E}(Y_{n+1} - Y_n \mid \mathcal{F}_{n-1})| \leq \delta_n, \ \ a.s.,$$

for all $n \geq 1$, and, for some $\varepsilon > 0$, $\delta_n \leq \mathrm{e}^{-n^\varepsilon}$ for all but finitely many n, a.s. In particular, $\sum_{n \geq 1} \delta_n < \infty$, a.s.

(c) *We have*

$$(6.3) \qquad \mathbb{E}((Y_{n+1} - Y_n)^2 \mid \mathcal{F}_{n-1}) = 1 + 8q_n, \ \ a.s.$$

PROOF. For $x \in \mathbb{Z}$ and $i \in \{-1, +1\}$, define

$$\Delta^+(x, i) := f(x + i, i) - f(x, i), \ \text{and} \ \Delta^-(x, i) := f(x - i, -i) - f(x, i).$$

Then since $X_{n+1} = X_n + \eta_{n+1}$, we have that

$$(6.4) \qquad Y_{n+1} - Y_n = \Delta^+(X_n, \eta_n)\mathbf{1}\{\eta_{n+1} = \eta_n\} + \Delta^-(X_n, \eta_n)\mathbf{1}\{\eta_{n+1} \neq \eta_n\}.$$

Note that $\Delta^+(x, i) = i$ and

$$\Delta^-(x, i) = -i + a\mathbf{1}\{i = -1\} - a\mathbf{1}\{i = 1\} = -i - ai.$$

Thus from (6.4) we have $|Y_{n+1} - Y_n| = 1 + a\mathbf{1}\{\eta_{n+1} \neq \eta_n\} \leq 3$, a.s., giving (6.1). For q_n the $\mathcal{F}_{n-1}$-measurable random variable defined in (5.1), we have from (6.4) that

$$\mathbb{E}(Y_{n+1} - Y_n \mid \mathcal{F}_{n-1}) = (1 - q_n)\Delta^+(X_n, \eta_n) + q_n\Delta^-(X_n, \eta_n).$$

Since $\Delta^\pm(x, i)$ are uniformly bounded, we have from Proposition 5.1 that there is an $\mathcal{F}_{n-1}$-adapted sequence ε_n with $\delta_n := |\varepsilon_n| \leq \mathrm{e}^{-n^\varepsilon}$ for all but finitely many n, such that

$$(1 - q_n)\Delta^+(x, i) + q_n\Delta^-(x, i) = (1 - q)\Delta^+(x, i) + q\Delta^-(x, i) + \varepsilon_n.$$

Here we have that

$$(1 - q)\Delta^+(x, i) + q\Delta^-(x, i) = (1 - q)i + q(-i - ai)$$
$$= (1 - 2q)i - aqi = 0,$$

for all x and all i, by choice of a. This gives (6.2).

For the second moment, note that, by (6.4),

$$\mathbb{E}((Y_{n+1} - Y_n)^2 \mid \mathcal{F}_{n-1}) = (1 - q_n)(\Delta^+(X_n, \eta_n))^2 + q_n(\Delta^-(X_n, \eta_n))^2.$$

Here $(\Delta^+(x, i))^2 = 1$ and $(\Delta^-(x, i))^2 = (1 + a)^2 = 9$, and (6.3) follows. $\qquad \square$

The proofs of our two main theorems will use the following martingale decomposition. Set $\theta_n := \mathbb{E}(Y_{n+1} - Y_n \mid \mathcal{F}_{n-1})$ for $n \in \mathbb{N}$. Note that, by (6.1), $|\theta_n| \le 3$, a.s. As in Doob's decomposition, for $n \ge 1$ let $A_n := \sum_{i=1}^{n-1} \theta_i$, and set $M_n := Y_n - A_n$, so that

$$\mathbb{E}(M_{n+1} - M_n \mid \mathcal{F}_{n-1}) = \mathbb{E}(Y_{n+1} - Y_n \mid \mathcal{F}_{n-1}) - (A_{n+1} - A_n) = \theta_n - \theta_n = 0.$$

Thus M_n is an $\mathcal{F}_{n-1}$-adapted martingale ($n \ge 1$), and

(6.5) $$|M_{n+1} - M_n| \le |Y_{n+1} - Y_n| + |\theta_n| \le 6, \text{ a.s.}$$

Note that $|\theta_n| \le \delta_n$, a.s., where $\sum_{n \ge 1} \delta_n < \infty$, a.s., by Lemma 6.1(b).

PROOF OF THEOREM 1.1. We have that M_n is a martingale with uniformly bounded increments, by (6.5). Theorem 5.3.1 of [6] says that M_n either oscillates $(\liminf_{n \to \infty} M_n = -\infty, \limsup_{n \to \infty} M_n = +\infty)$, or converges $(\lim_{n \to \infty} M_n \to M_\infty \in \mathbb{R})$. Suppose that $M_n \to M_\infty$. Then

$$\limsup_{n \to \infty} |Y_n| \le |M_\infty| + \limsup_{n \to \infty} \sum_{i=1}^{n-1} |\theta_i| \le |M_\infty| + \sum_{i=1}^{\infty} \delta_i < \infty, \text{ a.s.}$$

But since $|Y_n| \ge |X_n| - a$, this contradicts Kurkova and Menshikov's result (1.1), which says that $\limsup_{n \to \infty} |X_n| = \infty$, a.s. Thus we must have that M_n oscillates, a.s. Then since $\sup_{n \ge 1} |A_n| < \infty$, it follows that Y_n oscillates, and hence X_n oscillates. Since $X_n \in \mathbb{Z}$ satisfies $|X_{n+1} - X_n| = 1$, and it oscillates between $-\infty$ and $+\infty$, we must have $X_n = x$ i.o. for every $x \in \mathbb{Z}$. Hence, a.s., for every $x \in \mathbb{Z}$, the set $\{t \ge 0 : S(t) = x\}$ is unbounded. The result extends to all $x \in \mathbb{R}$ by continuity of the server's trajectory. □

Now we turn to the proof of Theorem 1.2. The next result is essentially an inversion of Propositions 4.1 and 4.2.

LEMMA 6.2. *We have that*

$$\lim_{t \to \infty} \frac{N_t}{\log \log t} = \frac{1}{\log 2}, \text{ a.s.}$$

PROOF. Let $\alpha \in (1, 2)$. Since $N_t \to \infty$ a.s., we have from Proposition 4.1 that a.s., for all t sufficiently large,

$$t \ge T_{N_t} \ge e^{\alpha^{N_t}}.$$

It follows that $\log \log t \ge N_t \log \alpha$ for all t sufficiently large. Hence

$$\limsup_{t \to \infty} \frac{N_t}{\log \log t} \le \frac{1}{\log \alpha}, \text{ a.s.}$$

Since $\alpha \in (1, 2)$ was arbitrary, we get

$$\limsup_{t \to \infty} \frac{N_t}{\log \log t} \leq \frac{1}{\log 2}, \text{ a.s.}$$

On the other hand, for $\beta > 2$ we have from Proposition 4.2 that a.s., for all t sufficiently large,

$$t \leq T_{N_t + 1} \leq e^{\beta^{N_t + 1}}.$$

It follows that $\log \log t \leq (N_t + 1) \log \beta$ for all t sufficiently large. Hence

$$\liminf_{t \to \infty} \frac{N_t}{\log \log t} \geq \frac{1}{\log \beta}, \text{ a.s.}$$

Since $\beta > 2$ was arbitrary, we get

$$\liminf_{t \to \infty} \frac{N_t}{\log \log t} \geq \frac{1}{\log 2}, \text{ a.s.}$$

Combined with the lim sup result, this gives the statement in the lemma. $\qquad\square$

Next we have an iterated logarithm law for X_n.

LEMMA 6.3.   *We have that*

$$\limsup_{n \to \infty} \frac{\pm X_n}{\sqrt{6n \log \log n}} = 1, \text{ a.s.}$$

PROOF. First note that

$$\mathbb{E}((M_{n+1} - M_n)^2 \mid \mathcal{F}_{n-1}) = \mathbb{E}((Y_{n+1} - Y_n - \theta_n)^2 \mid \mathcal{F}_{n-1})$$
$$= \mathbb{E}((Y_{n+1} - Y_n)^2 \mid \mathcal{F}_{n-1}) - \theta_n^2,$$

where, by Lemma 6.1, $|\theta_n| \leq \delta_n$ a.s., and both δ_n and δ_n^2 are a.s. summable. Thus from (6.3) and Proposition 5.1 we have that

$$\mathbb{E}((M_{n+1} - M_n)^2 \mid \mathcal{F}_{n-1}) = 1 + 8q + \varepsilon_n, \text{ a.s.,}$$

where $\varepsilon := \sum_{n \geq 1} \varepsilon_n$ has $|\varepsilon| < \infty$, a.s. Since $q = 1/4$, it follows that

$$(6.6) \qquad s_n^2 := \sum_{i=1}^{n} \mathbb{E}((M_{i+1} - M_i)^2 \mid \mathcal{F}_{i-1}) = 3n + \varepsilon + o(1), \text{ a.s.}$$

The conditions (6.5) and (6.6) show that M_n and $-M_n$ each satisfy the martingale law of the iterated logarithm [19], yielding

$$\limsup_{n \to \infty} \frac{\pm M_n}{\sqrt{6n \log \log n}} = 1, \text{ a.s.}$$

Since $|X_n - Y_n| \leq a$ and $|\theta_n| \leq \delta_n$, we have that

$$\sup_n |M_n - X_n| \leq a + \sup_n |M_n - Y_n| \leq a + \sum_{n=1}^{\infty} \delta_n < \infty, \text{ a.s.}$$

Thus the iterated logarithm law for M_n transfers to X_n. $\qquad\square$

PROOF OF THEOREM 1.2. Since $N_t \to \infty$ as $t \to \infty$, Lemma 6.3 shows that

$$\limsup_{t \to \infty} \frac{\pm X_{N_t}}{\sqrt{N_t \log \log N_t}} = \sqrt{6}, \text{ a.s.}$$

Combining this with Lemma 6.2 gives

$$\limsup_{t \to \infty} \frac{\pm X_{N_t}}{\sqrt{\log \log t \log \log \log \log t}} = \sqrt{\frac{6}{\log 2}}, \text{ a.s.}$$

We have from (2.1) and the fact that $|X_{n+1} - X_n| = 1$ that

$$S(t) \geq \min\{X_{N_t}, X_{N_t+1}\} \geq X_{N_t} - 1, \text{ and } S(t) \leq \max\{X_{N_t}, X_{N_t+1}\} \leq X_{N_t} + 1.$$

The result follows. $\qquad\square$

APPENDIX A: AUXILIARY LEMMAS

Recall that Φ denotes the distribution function of the standard normal distribution.

LEMMA A.1. *Let $\kappa \geq 0$ and let $\nu \sim P(\kappa)$ and $\nu' \sim P(\kappa)$ be independent. Then there exists $C \in \mathbb{R}_+$ such that, for all $\kappa > 0$,*

$$\sup_{u \in \mathbb{R}} \left| \mathbb{P}(\kappa^{-1/2}(\nu - \nu') \leq u) - \Phi(u/\sqrt{2}) \right| \leq C(1 + \kappa)^{-1/2} \log(2 + \kappa).$$

PROOF. Let $u_\kappa := \kappa^{1/2} \log \kappa$. Then, by symmetry,

$$\sup_{u:|u|>u_\kappa} \left| \mathbb{P}(\kappa^{-1/2}(\nu - \nu') \leq u) - \Phi(u/\sqrt{2}) \right| = \sup_{u:u>u_\kappa} \left| \mathbb{P}(\kappa^{-1/2}(\nu - \nu') > u) - \bar{\Phi}(u/\sqrt{2}) \right|$$

$$\leq \mathbb{P}(\kappa^{-1/2}(\nu - \nu') > u_\kappa) + \bar{\Phi}(u_\kappa/\sqrt{2}).$$

Here we have from standard Gaussian tail bounds (see e.g. Theorem 1.2.3 of [6]) that $\bar{\Phi}(u_\kappa/\sqrt{2}) = O(e^{-\kappa})$, say, while, since $\nu' \geq 0$,

$$\mathbb{P}(\kappa^{-1/2}(\nu - \nu') > u_\kappa) \leq \mathbb{P}(\nu > \kappa^{1/2} u_\kappa) = O(e^{-\kappa}),$$

by Poisson large deviations bounds (see e.g. [16, p. 17]). The result in the lemma will thus follow from the claim that there exists $C \in \mathbb{R}_+$ for which

$$(A.1) \qquad \sup_{u:|u|\leq u_\kappa} \left| \mathbb{P}(\kappa^{-1/2}(\nu - \nu') \leq u) - \Phi(u/\sqrt{2}) \right| \leq C\kappa^{-1/2} \log \kappa,$$

for all $\kappa \geq 2$. It remains to prove (A.1).

By Poisson additivity, we can write

$$\nu - \nu' = \sum_{j=1}^{\lfloor \kappa \rfloor} \xi_j + \gamma - \gamma',$$

where $\gamma, \gamma', \xi_1, \xi_2, \ldots$ are independent random variables, each ξ_j being the difference of two independent $P(1)$ random variables, and γ, γ' being Poisson with mean $\kappa - \lfloor \kappa \rfloor < 1$. Let $S_\kappa := \sum_{j=1}^{\lfloor \kappa \rfloor} \xi_j$. Then $\mathbb{E}\,\xi_j = 0$, $\mathbb{E}(\xi_j^2) = 2$, and $\mathbb{E}(|\xi_j|^3) < \infty$, so the Berry–Esseen theorem (see e.g. Theorem 3.4.9 of [6, p. 137]) implies that

$$(A.2) \qquad \sup_{u \in \mathbb{R}} \left| \mathbb{P}\left(\lfloor \kappa \rfloor^{-1/2} S_\kappa \leq u \right) - \Phi(u/\sqrt{2}) \right| \leq C \kappa^{-1/2},$$

for all $\kappa \geq 1$.

First we prove one half of (A.1). Since $\gamma \geq 0$,

$$\mathbb{P}(\nu - \nu' \leq u\kappa^{1/2}) \leq \mathbb{P}(S_\kappa - \gamma' \leq u\kappa^{1/2})$$
$$\leq \mathbb{P}(S_\kappa \leq u\kappa^{1/2} + r) + \mathbb{P}(\gamma' \geq r),$$

for any $r > 0$. Here we have that

$$\mathbb{P}(S_\kappa \leq u\kappa^{1/2} + r) = \mathbb{P}\left(\lfloor \kappa \rfloor^{-1/2} S_\kappa \leq \left(\frac{\kappa}{\lfloor \kappa \rfloor} \right)^{1/2} u + \lfloor \kappa \rfloor^{-1/2} r \right).$$

Note that, by Taylor's theorem,

$$\left(\frac{\kappa}{\lfloor \kappa \rfloor} \right)^{1/2} \leq \left(\frac{\lfloor \kappa \rfloor + 1}{\lfloor \kappa \rfloor} \right)^{1/2} \leq 1 + C\kappa^{-1},$$

for some $C \in \mathbb{R}_+$ and all $\kappa \geq 1$. Thus, by (A.2) and the fact that the standard normal density is uniformly bounded,

$$\mathbb{P}(S_\kappa \leq u\kappa^{1/2} + r) \leq \Phi(u/\sqrt{2}) + C\kappa^{-1}|u| + C\kappa^{-1/2}r + C\kappa^{-1/2},$$

for all u and all $\kappa \geq 1$. In particular, taking $r = \log \kappa$ we have

$$\mathbb{P}(S_\kappa \leq u\kappa^{1/2} + \log \kappa) \leq \Phi(u/\sqrt{2}) + C\kappa^{-1/2}\log \kappa, \text{ for all } u \in [-u_\kappa, u_\kappa],$$

where $C < \infty$ does not depend on u or κ. On the other hand, $\mathbb{P}(\gamma' \geq \log \kappa) \leq \mathbb{P}(P(1) \geq \log \kappa) = O(\kappa^{-1})$ by Poisson large deviations bounds (see e.g. [16, p. 17]). This establishes one half of (A.1).

For the other direction, we have that

$$\mathbb{P}(\nu - \nu' \leq u\kappa^{1/2}) \geq \mathbb{P}(S_\kappa + \gamma \leq u\kappa^{1/2})$$
$$\geq \mathbb{P}(S_\kappa \leq u\kappa^{1/2} - r) - \mathbb{P}(\gamma \geq r).$$

Taking $r = \log \kappa$ we get $\mathbb{P}(\gamma \geq \log \kappa) = O(\kappa^{-1})$ and, similarly to above, we get

$$\mathbb{P}(S_\kappa \leq u\kappa^{1/2} - \log \kappa) \geq \Phi(u/\sqrt{2}) - C\kappa^{-1/2}\log \kappa, \text{ for all } u \in [-u_\kappa, u_\kappa],$$

completing the proof of (A.1). $\qquad\square$

Finally, we record the following elementary result.

LEMMA A.2. *Let X, Y be random variables. Then for any $x \in \mathbb{R}$,*

$$\mathbb{P}(X > x \mid X \leq Y) \leq \mathbb{P}(X > x).$$

PROOF. For $x, y \in \mathbb{R}$ we have

$$\mathbb{P}(X \leq x \mid X \leq y) = \frac{\mathbb{P}(X \leq \min\{x,y\})}{\mathbb{P}(X \leq y)} = \begin{cases} 1 & \text{if } x \geq y, \\ \frac{\mathbb{P}(X \leq x)}{\mathbb{P}(X \leq y)} & \text{if } x \leq y. \end{cases}$$

In any case, we have $\mathbb{P}(X \leq x \mid X \leq y) \geq \mathbb{P}(X \leq x)$, and the result follows. $\qquad\square$

ACKNOWLEDGEMENTS

REFERENCES

[1] BILLINGSLEY, P. (1999). *Convergence of Probability Measures.* 2nd ed., Wiley, New York. MR1700749

[2] BORDENAVE, C., FOSS, S., AND LAST, G. (2011). On the greedy walk problem. *Queueing Syst.* **68** 333–338. MR2834204

[3] CHEN, A. AND RENSHAW, E. (1992). The Gillis–Domb–Fisher correlated random walk. *J. Appl. Probab.* **29** 792–813. MR1188536

[4] CHRISTOPH, G. AND WOLF, W. (1992). *Convergence Theorems with a Stable Limit Law.* Akademie Verlag, Berlin. MR1202035

[5] COFFMAN, E.G. AND GILBERT, E.N. (1987). Polling and greedy servers on a line. *Queueing Syst.* **2** 115–145. MR0905435

[6] DURRETT, R. (2010). *Probability: Theory and Examples.* 4th ed., Cambridge University Press, Cambridge. MR2722836

[7] ETHIER, S.N. AND KURTZ, T.G. (1986). *Markov Processes. Characterization and Convergence.* John Wiley & Sons, Inc., New York. MR0838085

[8] FELLER, W. (1971). *An Introduction to Probability Theory and Its Applications, Vol. II.* 2nd ed., Wiley, New York. MR0270403

[9] FOSS, S. AND LAST, G. (1996) Stability of polling systems with exhaustive service policies and state-dependent routing. *Ann. Appl. Probab.* **6** 116–137.

[10] Foss, S., Rolla, L.T., and Sidoravicius, V. (2015). Greedy walk on the real line. *Ann. Probab.* **43** 1399–1418. MR3342666

[11] Ibragimov, I.A. and Linnik, Y.V. (1971). *Independent and Stationary Sequences of Random Variables.* Wolters-Noordhoff, Groningen, The Netherlands. MR0322926

[12] Kallenberg, O. (2002). *Foundations of Modern Probability.* 2nd ed., Springer, New York. MR1876169

[13] Kurkova, I.A. and Menshikov, M.V. (1997). Greedy algorithm: Z^1 case. *Markov Process. Relat. Fields* **3** 243–259. MR1468176

[14] Nolan, J.P. (2018). *Stable Distributions—Models for Heavy Tailed Data.* Birkhäuser, Boston.

[15] Paulauskas, V.I. (1974). Estimates of the remainder term in limit theorems in the case of stable limit law. *Lith. Math. J.* **14** 127–146. cf. MR0380943

[16] Penrose, M. (2003). *Random Geometric Graphs.* Oxford University Press, Oxford. MR1986198

[17] Rojas-Nandayapa, L., Foss, S., and Kroese, D.P. (2011). Stability and performance of greedy server systems. A review and open problems. *Queueing Syst.* **68** 221–227. MR2834192

[18] Satybaldina, K.I. (1973). Absolute estimates of the rate of convergence to stable laws. *Theory Probab. Appl.* **17** 726–728. cf. MR0310949

[19] Stout, W.F. (1970). A martingale analogue of Kolmogorov's law of the iterated logarithm. *Z. Wahr. verw. Geb.* **15** 279–290. MR0293701

[20] Watson, G.N. (1944). *A Treatise on the Theory of Bessel Functions.* 2nd ed., Cambridge University Press. MR0010746

[21] Whitt, W. (2002). *Stochastic-Process Limits.* Springer, New York. MR1876437

Department of Actuarial Mathematics & Statistics
Heriot-Watt University
Riccarton
Edinburgh EH14 4AS
E-mail: r.cruise@hw.ac.uk

Department of Mathematical Sciences
Durham University
South Road
Durham DH1 3LE
E-mail: andrew.wade@durham.ac.uk