

Large deviations analysis for the $\$M/H_{2/n} + M\$$ queue in the Halfin-Whitt regime

Citation for published version (APA):

Mukherjee, D., Li, Y., & Goldberg, D. A. (2018). Large deviations analysis for the $\$M/H_{2/n} + M\$$ queue in the Halfin-Whitt regime. *arXiv, 2018*, Article 1803.01082v1. <https://doi.org/10.48550/arXiv.1803.01082>

DOI:

[10.48550/arXiv.1803.01082](https://doi.org/10.48550/arXiv.1803.01082)

Document status and date:

Published: 01/03/2018

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Large deviations analysis for the $M/H_2/n + M$ queue in the Halfin-Whitt regime

Debankur Mukherjee

Eindhoven University of Technology, d.mukherjee@tue.nl

Yuan Li

Amazon, yuan.leo.li@gmail.com

David A. Goldberg

Cornell University, dag369@cornell.edu

We consider the FCFS $M/H_2/n + M$ queue in the Halfin-Whitt heavy traffic regime, i.e. the first-come-first-served multi-server queue with abandonments, in which inter-arrival and patience times are Markovian, and service times are hyper-exponentially distributed, i.e. the mixture of two exponential distributions. It is known that if one considers a sequence of such queues under the Halfin-Whitt scaling, the appropriately normalized sequence of steady-state queue-length distributions is tight and converges weakly to a limiting random variable W^∞ . However, those works only describe W^∞ implicitly as the unique invariant measure of a certain complicated diffusion process with no closed form. Although it was proven by Gamarnik and Stolyar (2012) that $P(W^\infty > x)$ has a sub-Gaussian decay, and related Lyapunov-type arguments and bounds also appear in the work of Dieker and Gao (2013), the actual value of $\lim_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty > x) \right)$ was left open. In subsequent work, Dai and He (2013) conjectured an explicit form for this exponent, which depends only on the first two moments of the service distribution, and used this conjectured exponent to devise a numerical method for computing the tail probabilities of W^∞ .

In this work, we explicitly compute the true large deviations exponent for W^∞ , in the parameter regime such that the abandonment rate is strictly less than the minimum exponential service rate. This is the first such result for multi-server queues with abandonments when service times are non-Markovian. As this exponent is strictly less than that conjectured by Dai and He, and further does not only depend on the first two moments of the service distribution, we resolve the conjecture of Dai and He in the negative. Our main approach is to extend the stochastic comparison framework of Gamarnik and Goldberg (2013) to the setting of abandonments, which requires several novel and non-trivial contributions (especially in light of the non-monotonicities introduced by abandonments), and to combine with results from the theory of Gaussian processes. Furthermore, our approach sheds light on several novel ways to think about multi-server queues with abandonments in the Halfin-Whitt regime, e.g. by formally connecting to so-called multi-dimensional Ornstein-Uhlenbeck processes (a family of tractable diffusions), and we believe that these insights should hold in considerable generality and provide a novel set of tools for analyzing these systems.

Key words: many-server queues, Halfin-Whitt regime, stochastic comparison, weak convergence, large deviations, Gaussian process, abandonments

Table of contents.

1	Introduction.	4
1.1	Multi-server queues with abandonments in the Halfin-Whitt regime.	4
1.2	Diffusion approximations.	5
1.3	Diffusion approximations with non-Markovian service times.	5
1.4	Large deviations of the steady-state diffusion approximation.	6
1.5	A conjecture of Dai and He [28].	7
1.6	Stochastic comparison methods for queues with abandonments.	9
1.7	Related diffusion processes.	9
1.8	The one-dimensional (i.e. Markovian) case.	10
1.9	Overview of contribution.	11
1.10	Outline of rest of paper.	11
2	Main results.	11
2.1	Brief review of model, notations, and assumptions.	11
2.2	Main results.	12
2.3	Discussion	13
2.3.1	Implications for Dai and He’s numerical methods.	13
2.3.2	Qualitative insights from explicit exponent.	13
2.3.3	Why no insensitivity?	13
2.3.4	Connection to multi-dimensional OU processes.	14
2.3.5	Overview of insights from proofs.	15
3	Upper bound.	16
3.1	CTMC representation for generic $M/H_2/n + M$ queue Q	17
3.2	CTMC representation for bounding system \bar{Q}	18
3.3	Proof that \bar{W} yields an upper bound for W	18
3.3.1	Expanded Markovian representations Q^0 and \bar{Q}^0 for Q and \bar{Q}	19
3.3.2	Formal statement of dominance and description of coupling.	24
3.3.3	Proof of stochastic ordering.	27
4	Skorokhod representation for upper bound \bar{W}.	29
4.1	Review of generalized drift Skorokhod problem.	29
4.2	Expressing \bar{W} as a solution to the generalized drift Skorokhod problem.	31

5	Asymptotic analysis of \overline{W} and proof of upper bound for Theorem 1.	33
5.1	Asymptotic version of Lemma 8.	34
5.2	Relating SKOR^∞ to an integrated OU process.	35
5.3	Stochastic calculus to analyze Λ	36
5.4	Gaussian process theory to analyze $\sup_{0 \leq s \leq t} \mathcal{G}(s)$	38
5.5	Interchange-of-limits results from [31, 24].	40
5.5.1	Review of weak convergence results of [29].	41
5.5.2	Review of ergodicity and interchange-of-limits results of [31, 24].	41
5.6	Proof of upper bound part of main result Theorem 1.	42
6	Lower Bound.	43
6.1	Further analysis of the limit process Q^∞	44
6.2	Relating Q^∞ to a multi-dimensional OU process.	46
6.3	Additional bounds and asymptotics for multi-dimensional OU processes.	48
6.4	Proof of lower bound part of main result Theorem 1.	49
6.4.1	Slepian's inequality and bounds for a finite-horizon infimum.	50
6.4.2	Proof of lower bound part of main result Theorem 1.	51
7	Conclusion.	51
8	Appendix.	54
8.1	Proof of Lemma 5.	54
8.2	Proof of Lemma 15.	59
8.3	Proof of Lemma 17.	61
8.4	Proof of Lemma 21.	62
8.5	Proof of Lemma 23.	62
8.6	Proof of Lemma 26.	63
8.7	Proof of Corollary 2.	64
8.8	Proof of Lemma 27.	65
8.9	Proof of Lemma 33.	66
8.10	Proof of Lemma 35.	68
8.11	Proof of Lemma 36.	68
8.12	Proof of Lemma 38.	72
8.13	Proof of non-insensitivity.	74
8.14	Proof that $-\frac{\theta(\theta+\mu_1\mu_2)}{2(\theta+\mu_1+\mu_2-1)} < -\frac{\mu_1\mu_2\theta}{2(\mu_1+\mu_2-1)}$	76

1. Introduction.

1.1. Multi-server queues with abandonments in the Halfin-Whitt regime.

It is by now well-known that from a queueing perspective, a service system can operate in several different asymptotic regimes, parametrized by the number of servers and traffic intensity. Recently there has been a significant interest in the so-called Halfin-Whitt (i.e. HW, Quality-and-Efficiency Driven, QED) regime, in which one can easily trade-off between the number of servers and the quality of service experienced by arrivals to the system (as parametrized through metrics including the probability that an arriving job must wait for service). This scaling regime was studied originally by Erlang [37] and Jagerman [55], and formally introduced by Halfin and Whitt [50], who studied the First-come-first-serve (FCFS) $G/M/n$ system (for large n) when the traffic intensity ρ scales like $1 - Bn^{-\frac{1}{2}}$ for some strictly positive excess parameter B . We refer the interested reader to [41, 45, 46, 5, 66] for an overview of recent progress on our understanding of multi-server queues in the HW regime, which we make no attempt to survey here, and to [46] for an overview of how this scaling regime relates to other regimes and related results.

In many such service systems customer abandonment is an important feature [108] (see also several relevant empirical works including [81, 20, 78, 13]) and there is also considerable work studying multi-server queueing models with abandonments in the HW regime. For an overview of work on multi-server queues with abandonments in the HW regime, as well as other scaling regimes, we refer the interested reader to the recent surveys [108, 27]. When studying heavy-traffic limits of queueing systems, there are generally two different levels of granularity at which one can study the associated limits: the fluid scale (fluctuations of order n) and the diffusion scale (fluctuations around the fluid scaling, often of order $n^{\frac{1}{2}}$). For multi-server queues with abandonments, questions regarding convergence at the fluid scale have been studied in several works, including [60, 59, 70, 68, 69, 119, 123], where several of these works allow for systems with time-varying parameters, and we refer the interested reader to [27] for additional references. We will instead focus on the second level of granularity, i.e. the relevant diffusion-scaled process, where we note that these works often implicitly demonstrate a fluid scaling as well to provide an appropriate centering for the relevant diffusion limit. Furthermore, we will focus exclusively on models with abandonments in the HW scaling regime, but note in passing that there is a vast literature (beyond the scope of our work) on closely related models and topics pertaining to abandonments, including (but by no means limited to): non-asymptotic analysis and numerical approaches [2, 91, 82, 63, 102]; other scaling regimes [116, 57, 97]; single-server models [109, 111, 98]; network models [99, 52, 88]; models with more complex dependency structure [121, 38]; control and dynamic staffing [39, 7, 71]; simulation [87]; delay announcements [8, 54]; strategic behavior [77, 10]; and ticket queues [56, 32].

1.2. Diffusion approximations.

It seems the first work to address the question of diffusion-scaled multi-server queues with abandonments in the HW scaling was [40], which used the available closed-form expressions for the $M/M/n + M$ queue to work out the relevant asymptotics for the $M/M/n + M$ case, in the steady-state. In [44], these results were also extended to the transient setting for the $M/M/n + M$ queue, and the results gained renewed popularity. To be slightly more precise about the notion of a diffusion limit in the HW regime, and the relevant concerns, suppose we have fixed an inter-arrival distribution A and service distribution S , both with mean 1, a patience distribution R with continuous density function which is strictly positive and finite at 0, and a spare capacity parameter B (which may be negative, positive, or zero). Let Q^n denote the queueing system with inter-arrival times i.i.d. distributed as $\frac{A}{n+Bn^{\frac{1}{2}}}$, service times drawn i.i.d. from the unscaled distribution S , and patience times drawn i.i.d. from the unscaled distribution R . For this system, let $\{W^n(t), t \geq 0\}$ be the queue-length process (i.e. the number of jobs waiting in queue), where for now we leave the initial distribution unspecified. It is well-known that under mild technical conditions such a system will have a unique steady-state distribution (see e.g. [24]), and we let $W^n(\infty)$ denote a random variable (r.v) with the corresponding steady-state distribution (on occasion we will also use r.v to denote a random vector). In that case, the results of [40] showed that for an appropriate sequence of $M/M/n + M$ queues, $\{n^{-\frac{1}{2}}W^n(\infty), n \geq 1\}$ converges weakly to a limiting r.v $W^\infty(\infty)$, while the results of [44] showed in addition that for each fixed $T \geq 0$, $\{n^{-\frac{1}{2}}W^n(t)_{0 \leq t \leq T}\}$ converges weakly to a limiting process $W^\infty(t)_{0 \leq t \leq T}$ (in the appropriate space and topology, i.e. a process-level convergence). [44] also showed that in this case a so-called interchange-of-limits holds, namely that $\{W^\infty(t), t \geq 0\}$ converges, as $t \rightarrow \infty$, to $W^\infty(\infty)$. Namely, the limit of the steady-states coincides with the steady-state of the limit, which is in general non-trivial to verify. Related higher-order asymptotic expansions were derived in [124]. In [79], these results are extended to show a steady-state convergence in the $M/M/n + G$ setting, building on the non-asymptotic work of [11]. In [118], these results are extended to show a process-level convergence in the $G/M/n + G$ setting. Process-level weak convergence and strong approximation results were proven for time-varying networks of $M/M/n$ queues in [74], with several related results summarized in [75], and results for related periodic Markovian systems proven in [93].

1.3. Diffusion approximations with non-Markovian service times.

Note that all the aforementioned works essentially restrict to the setting of Markovian service times, where we note that in the asymptotic regimes considered generalizing the service distribution is typically far more challenging than generalizing the inter-arrival or patience distributions. However, as empirically studied in [20], in many relevant applications of queueing models in service systems

the assumption of Markovian service times does not hold, and may lead to poor policy choices. Significant progress on this front was made in [29], in which a process-level weak-convergence result was proven for the case of phase-type (i.e. PH) service distributions (which are dense within the set of all distributions), building on a state-space collapse relating the number of abandoning jobs to a certain integral of the queue-length process [26]. The associated limit process was a complex multi-dimensional Markovian diffusion with piece-wise linear drift and no explicit solution, called a piecewise Ornstein-Uhlenbeck (OU) process in later work, and could also be described as a certain continuous mapping of a simpler process. We note that a related state-space collapse result was attained in [106], in which an asymptotic equivalence between the system with balking (i.e. in which jobs that will eventually abandon do so upon arrival) and with reneging (i.e. in which jobs remain in queue until they abandon and are counted towards the queue-length) is proven. Process-level weak-convergence results very similar to those of [29], albeit for even more general service distributions, were attained in [76]. In the same paper, the question of proving analogous results for the associated sequence of steady-state distributions, and a corresponding interchange-of-limits, was explicitly stated as an open problem.

This was resolved in the sequence of papers [31, 24]. In [31], the authors proved that the piecewise OU processes which arose in [29] had unique invariant measures, and were exponentially ergodic (i.e. converged exponentially quickly under an appropriate metric to this invariant measure from general initial conditions), where we note that such exponential ergodicity has since been extended to more general processes [7]. In [24], the results of [31] were combined with several additional arguments to prove a formal interchange-of-limits for $G/PH/n + M$ queues in the HW regime, namely that the sequence of appropriately normalized steady-state distributions is tight, and converges weakly to the unique invariant measure of the associated piecewise OU process identified in [29].

1.4. Large deviations of the steady-state diffusion approximation.

However, outside the case of Markovian service times, all of the aforementioned works left open any concrete understanding of the associated complex limit processes, beyond their exponential ergodicity [31]. There have been several works which attempted to shed some light on these processes, especially their large deviations behavior, as the probability of rare events is typically a quantity of significant interest in such models, and we refer the interested reader to [104] for an overview of such results, and to [41] for an overview of such results for multi-server queues in the HW scaling. In [43], the authors study multi-class Markovian queues with abandonments in the HW regime under general non-idling class-scheduling policies. For such systems, they prove several results regarding existence of exponential moments and sub-Gaussian behaviors. Their results include as a special case FCFS $M/H_k/n + M$ queues, i.e. queues in which service times are hyper-exponentially distributed, i.e. distributed as a finite mixture of exponential distributions. For a

sequence of such FCFS $M/H_k/n+M$ queues, the results of [43] imply that, letting $W^\infty(\infty)$ denote the weak limit of $\{n^{-\frac{1}{2}}W^n(\infty), n \geq 1\}$, $\limsup_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right) < 0$. As noted in [28], such a Gaussian decay is to be expected due to the mean-reverting nature of a system with abandonments, and based on explicit formulas associated with the Markovian case. However, the associated bounds identified in [43] involve both stochastic comparison and non-explicit Lyapunov-function-type arguments, and do not seem to shed much insight on the precise set of $z > 0$ such that (s.t.) $\mathbb{E} \left[\exp \left(z(W^\infty(\infty))^2 \right) \right] < \infty$, other than that this set is non-empty. We note that related Lyapunov-type arguments and bounds also appear in [31], but again it seems unlikely that the approach taken there could be used to identify the exact tail behavior. We will loosely refer to $-\sup \left\{ z : \mathbb{E} \left[\exp \left(z(W^\infty(\infty))^2 \right) \right] < \infty \right\}$ as the associated large deviations exponent, equivalently (should the limit exist) $\lim_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right)$, where we note that this concept is seemingly much more developed in the literature for queues without abandonments, in which the associated r.v typically exhibit exponential (as opposed to Gaussian) decay [101].

1.5. A conjecture of Dai and He [28].

To gain further insight into $W^\infty(\infty)$, and the quality of approximation that it provides to $W^n(\infty)$ for any fixed n , various numerical methods were developed in [28] for evaluating the probabilities of $W^\infty(\infty)$, building on the numerical methods for diffusions developed earlier by Dai and Harrison [25]. As discussed in-depth in [25], an important step in these numerical methods is selecting a reference density which is used to define an inner product, which is then used explicitly by the algorithm to determine the probabilities of $W^\infty(\infty)$ by defining the stationary measure through a so-called basic adjoint relationship with respect to this inner product, and then using certain appropriate approximations. The authors in [28] devote a section to discussing the properties that make certain reference densities more desirable than others, and note that ideally the reference density would have similar tail behavior to the stationary measure itself, where in some sense the greater the similarity the faster their algorithm will converge. However, the problem faced by the authors was that for the diffusion processes they were considering, i.e. the limit processes arising from $G/PH/n+G$ queues in the HW scaling, the tail behavior was unknown. As such, to motivate their choice of reference density, the authors made a formal conjecture regarding the large deviations exponent for $W^\infty(\infty)$. Namely, inspired by known results for the case of Markovian processing times, as well as an analogous result for the setting without abandonments [41], Dai and He made the following formal conjecture in Conjecture 2 of [28]. Throughout we suppose without loss of generality (w.l.o.g) that the expected service time equals 1, and let $\sigma_A^2(\sigma_S^2) \in (0, \infty)$ denote the variance of $A(S)$, where we recall that in the n th system the inter-arrival distribution is $\frac{A}{n+Bn^{\frac{1}{2}}}$ for some fixed real excess parameter B , the service distribution is S , and the patience distribution

R has a density which evaluates to $\theta \in (0, \infty)$ at 0. We note that the precise conjecture of [28] makes some additional technical assumptions on S , and as we are only here interested in the case that S is hyper-exponentially distributed and the inter-arrival and patience distributions are both exponentially distributed, which is easily verified to satisfy all these technical conditions, we simply refer the reader to [28] for the precise assumptions.

CONJECTURE 1 (**Conjectured large deviations exponent** [28]).

$$\lim_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right) = -\frac{\theta}{\sigma_A^2 + \sigma_S^2}.$$

Customized to the $M/H_2/n + M$ setting in which service times are with probability (w.p.) p exponentially distributed with rate μ_1 and with probability $1 - p$ exponentially distributed with rate μ_2 , the canonical example of non-exponential service times explicitly used in [27, 24, 28] to illustrate the piecewise OU process, it is easily verified (since $\mathbb{E}[S] = 1$ implies $p = \frac{\mu_1(\mu_2 - 1)}{\mu_2 - \mu_1}$) that this conjecture is equivalent to

$$\lim_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right) = -\frac{\mu_1 \mu_2 \theta}{2(\mu_1 + \mu_2 - 1)}.$$

Furthermore, the authors of [28] note that if this conjecture is true, it provides an example of a so-called insensitivity result, as the exponent is a function only of θ and the first two moments of the inter-arrival and service distributions.

The authors of [28] implemented their numerical algorithm with a reference density strongly based on Conjecture 1. The demonstrated good numerical performance of their algorithm seemed to provide support for the conjecture. Of course, as the authors of [28] note, their algorithm will in general converge as long as the reference density satisfies certain basic technical conditions, i.e. it need not capture the correct large deviations behavior, i.e. they could have gotten good numerical results even if Conjecture 1 was false. Since publication of [28], [31], and [24], no further progress has been made towards resolving Conjecture 1, and in general the associated limit processes remain poorly understood. We do note the works [48, 49, 19] which use Stein's method and related approaches to further our understanding of how good an approximation the associated limit process provides for the rescaled queue length for any fixed n (i.e. bounds the error in using the diffusion to approximate the pre-limit system), and provides some bounds for certain moments of the limiting diffusions. However, as discussed in [46], developing further qualitative insight into the properties of the limiting processes themselves could shed considerable light on the pre-limit systems, where we do note that some care must be taken when transferring statements proven about the limit process to the pre-limit system, and refer the interested reader to [43] for further discussion along these lines. Prior to this work, Conjecture 1 remained open.

1.6. Stochastic comparison methods for queues with abandonments.

As it will play a central role in our analysis, we note that for systems without abandonments, a novel stochastic comparison approach for multi-server queues was developed in [41], and has since been used to prove results in several settings related to the HW scaling regime without abandonments [45, 47, 46]. Although there is a vast literature on stochastic comparison results for queues without abandonments [105, 21], the results for systems with abandonments are much more limited, seeming to consist of only a few works [15, 58, 18, 26, 43, 89]. One of the reasons for this is that certain basic monotonicities that hold in great generality for queueing systems without abandonments no longer hold in the presence of abandonments. One non-monotonicity of particular relevance to the approach of [41] is that in systems with abandonments, adding extra arrivals may actually cause the queue length to decrease along any given sample path. For example, consider the following simple queueing system. Suppose the queue is initially empty, and that there is a single arrival at time 1, where that single job has service time 10 and patience 4. Note that in this queue, the number of jobs in system at time 10 equals 1. However, consider the modified queueing system in which we add an extra arrival at time 0, with service time 6 and patience 1. In this modified system, the original arrival at time 1 will abandon the system at time 5, and the extra arrival will depart the system at time 6, leaving 0 jobs in system at time 10. Thus in systems with abandonments, adding jobs can actually decrease the number in system, which cannot happen in systems without abandonments. This and related subtleties makes it considerably more challenging to make comparative statements for queues with abandonments.

1.7. Related diffusion processes.

Although the piecewise OU processes analyzed in [29, 31, 24] remain poorly understood, as noted in [31], there is a well-studied family of related processes called multi-dimensional OU processes [103, 83]. Indeed, as noted in [31], one of the main hurdles to understanding the piecewise OU processes arising in [29, 31, 24] is that the infinitesimal drift is a non-linear function of the multi-dimensional state of the system, involving terms of the form $\max(0, \bar{x}_1 + \bar{x}_2)$, with \bar{x} the state of the diffusion. If within the system of stochastic differential equations (SDE) defining these processes, one simply replaces the terms of the form $\max(0, \bar{x}_1 + \bar{x}_2)$ with $\bar{x}_1 + \bar{x}_2$, the drift becomes a linear function of the state, and the unique strong solution of the resulting system is a so-called multi-dimensional OU process. Indeed, these processes are defined as the unique strong solutions to systems of SDE in which the drift is a linear function of the state, and the driving random process is Levy. These processes have essentially closed-form solutions and much is known about them, and we refer the reader to [103, 83] for formal definitions and further details. We also note that there is some work on certain modified OU processes also reminiscent of the piecewise OU processes arising

in [29, 31], including reflected OU processes [110, 122] and OU processes driven by OU processes [14]. There is also a literature on the large deviation properties of reflections and Skorokhod maps of general diffusion processes [36, 30, 61]. We also note that one tool that has proven useful in analyzing such reflected processes is a generalization of the traditional Skorokhod reflection map machinery to processes with state-dependent drift, and we refer the interested reader to [99, 100] for details, especially [99] which proved that a reflected multi-dimensional OU process can be used to model a certain queueing network in heavy-traffic, in a setting very different from the QED regime.

1.8. The one-dimensional (i.e. Markovian) case.

As for the $M/M/n + M$ queue our approach becomes straightforward, we briefly comment in more depth on certain facts either known or straightforward to prove (from known results) in this setting, as they set the stage for our own results and approach. It is well-known that for any fixed n , such a queueing system can be modeled as a birth-death process, and is time-reversible in steady-state. Furthermore, in this setting it is easy to see that if one attempts to “mimic” the stochastic comparison approach of [41], i.e. adding an extra arrival any time a server would have gone idle, one ends up with essentially the same 1-dimensional birth-death process, the only difference being that there is now a reflection at 0. In this case the steady-state distributions of both systems can be computed in essentially closed form using well-known properties of birth-death processes (as in [40, 44]). Furthermore, it is similarly well-known that here the steady-state probabilities for the number in queue in the modified system will be equivalent to those of the first system conditioned on having a positive queue. It follows that: 1. the steady-state number waiting in queue in the modified system stochastically dominates that in the original system (i.e. the non-monotonocities associated with systems with abandonments can be overcome due to the memoryless property) and 2. under the HW scaling both systems will have the same steady-state large-deviations behavior, where this tail exponent is explicitly given in [40, 44] and equals $-\frac{\theta}{2}$, consistent with the conjecture of Dai and He.

We note that in this case, it is easily verified that the resulting limiting 1-dimensional piecewise OU process Q^∞ corresponding to the process-level weak limit of the true queueing system in the HW regime has non-linear drift function $b(x) = B - \theta x - (\theta - 1)(x^+ - x)$, with x the state of the process and $x^+ \triangleq \max(0, x)$ for any $x \in \mathcal{R}$. Illustrating our aforementioned notion of “relaxing” the nonlinearity $(x^+ - x)$, i.e. replacing x^+ with x , we arrive at the defining equations for a one-dimensional OU process with drift function $\bar{b}(x) = B - \theta x$. As it is well-known that this one-dimensional OU process has the same large deviations exponent $-\frac{\theta}{2}$, we find that in the Markovian (i.e. one-dimensional) case, the upper-bounding process analogous to that defined for systems

without abandonments in [41], the true queueing system, and the “relaxed” OU process all have the same large deviations behavior. Of course, when S is non-Markovian, these arguments fundamentally break down and essentially nothing is known regarding the questions of interest.

1.9. Overview of contribution.

In this paper, we resolve Conjecture 1 in the negative. We do so by explicitly computing the large deviations exponent for the setting of $M/H_2/n + M$ queues, in the parameter regime s.t. the abandonment rate is strictly less than the minimum exponential service rate. We find that this true exponent is strictly less than the conjectured exponent, and does not depend only on the first two moments of the service distribution (i.e. does not exhibit an insensitivity-type result). This implies that the probability of seeing a large queue length decays faster than conjectured. Our main approach is to extend the stochastic comparison framework of [41] to the setting of abandonments, which requires several novel and non-trivial contributions, and to combine with results from the theory of Gaussian processes. Furthermore, we believe that our approach sheds light on several novel ways to think about multi-server queues with abandonments in heavy traffic, e.g. by formally connecting to so-called multi-dimensional OU processes (a family of tractable multi-dimensional diffusions), and should extend to much more general settings.

1.10. Outline of rest of paper.

The rest of the paper proceeds as follows. We present our main results, and discuss their context and significance, in Section 2. We present a stochastic comparison result for $M/H_2/n + M$ queues in Section 3, showing that a certain modified queueing system in which an extra arrival is added any time a server would otherwise go idle provides an upper bound. In Section 4, we prove that the upper bound from Section 3 can be expressed as the solution to a certain generalized Skorokhod problem. In Section 5, we provide an asymptotic analysis of the solution to the aforementioned Skorokhod problem in the HW regime, and combine with several results from the theory of Gaussian processes to complete the proof of the upper-bound of our main result. We complete the proof of our main result by demonstrating a matching lower bound in Section 6. We summarize our main results, and provide concluding remarks and some interesting directions for future research, in Section 7. We include a technical appendix containing the proofs of several results from throughout the paper in Section 8. For clarity of exposition, we put an asterisk next to each statement whose proof appears in the appendix. For similar reasons the details of certain straightforward proofs are omitted, and we put a dash next to each statement whose proof is omitted.

2. Main results.

2.1. Brief review of model, notations, and assumptions.

Recall that for $z > 0$, $Expo(z)$ refers to a generic exponentially distributed r.v. with rate z (mean z^{-1}). Suppose that a spare capacity parameter $B \in \mathcal{R}$, exponential rates $\mu_1, \mu_2 \in (0, \infty)$, mixing

probability $p \in (0, 1)$, and abandonment rate $\theta \in (0, \infty)$ are fixed. For $n > B^2$, let $\lambda_n \triangleq n + Bn^{\frac{1}{2}}$, and Q^n denote the $M/H_2/n + M$ queue with inter-arrival times distributed as $Expo(\lambda_n)$, service times hyper-exponentially distributed being $Expo(\mu_1)$ with probability p and $Expo(\mu_2)$ with probability $1 - p$, and patience times distributed as $Expo(\theta)$. W.l.o.g. (by a simple scaling argument), and for clarity of exposition, suppose throughout that the expected service time equals 1, i.e. $\frac{p}{\mu_1} + \frac{1-p}{\mu_2} = 1$. With a slight abuse of notation, we will sometimes refer to sequences such as $\{Q^n, n \geq 1\}$ even though the relevant system is only defined for $n > B^2$, and in all cases this should be read as a sequence starting at some sufficiently large n to ensure all quantities are well-defined. For general feasible initial conditions (i.e. having a finite number of jobs in system, and having positive queue only if all servers are busy), let $W^n(t)$ denote the number of jobs waiting in queue in Q^n at time t . Then it follows from the results of [24] that $\{W^n(t), t \geq 0\}$ converges in distribution (as $t \rightarrow \infty$) to a steady-state r.v. $W^n(\infty)$, independent of initial conditions. Let $\hat{W}^n(\infty) \triangleq n^{-\frac{1}{2}}W^n(\infty)$. Then it further follows from the results of [24] that $\{\hat{W}^n(\infty), n \geq 1\}$ converges in distribution (as $n \rightarrow \infty$) to a limiting non-defective (i.e. almost surely finite) r.v. $W^\infty(\infty)$ with support on \mathcal{R}^+ . Then the conjecture of Dai and He from [28] is equivalent to: $\lim_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right) = -\frac{\mu_1 \mu_2 \theta}{2(\mu_1 + \mu_2 - 1)}$.

2.2. Main results.

We now state our main results. As mentioned in Section 1, our proofs require that we restrict the range of allowed parameters as dictated by the following assumption.

Assumption 1 (Hazard-rate ordering assumption.)

$$0 < \theta < \min(\mu_1, \mu_2).$$

Intuitively, Assumption 1 implies that the infinitesimal rate at which a job leaves the system is always greater while that job is in service, as opposed to while waiting in queue, irregardless of how long that job has already been in service or waiting in queue. As we will discuss in-depth later, this assumption is needed to formally prove that our upper-bounding system is indeed an upper bound, and we leave the question of what happens outside this parameter range as an interesting question for future research. We also note that similar assumptions appeared in several arguments of [43], and that although our results still hold if $\theta = \min(\mu_1, \mu_2)$ (i.e. the second part of the inequality is not strict), for clarity of exposition we assume strict inequality throughout.

Then our main result is as follows.

THEOREM 1. *Under Assumption 1,*

$$\lim_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right) = -\frac{\theta(\theta + \mu_1 \mu_2)}{2(\theta + \mu_1 + \mu_2 - 1)}.$$

Furthermore, for all θ, μ_1, μ_2 s.t. $\mathbb{E}[S] = 1$ and $\mu_1 \neq \mu_2$, it is true that $-\frac{\theta(\theta+\mu_1\mu_2)}{2(\theta+\mu_1+\mu_2-1)}$ is strictly less than the exponent $-\frac{\mu_1\mu_2\theta}{2(\mu_1+\mu_2-1)}$ conjectured by Dai and He in [28]. Furthermore, the true exponent $-\frac{\theta(\theta+\mu_1\mu_2)}{2(\theta+\mu_1+\mu_2-1)}$ is not a function of only θ and the first two moments of the service distribution, and as such the insensitivity phenomena conjectured in [28] does not hold. Namely, there exist two hyper-exponential distributions with the same first two moments for which the true exponent differs (for the same value of θ).

2.3. Discussion

We now present a brief discussion of some implications of our main results, as well as some auxiliary results and aspects of our proof which will be discussed formally in later sections. Our results represent the first concrete insights into the qualitative behavior of the limiting processes which arise from queues with non-Markovian service times in the HW regime, beyond their ergodicity [31], the fact that the tails are in general sub-Gaussian [43], and the rate-of-convergence and moment bounds provided in [19]. We view the main implications to be four-fold.

2.3.1. Implications for Dai and He’s numerical methods. First, the results have direct implications for the selection of a reference density for the numerical algorithms of [28], and indeed suggest that one could get even better performance by using reference densities based on the true exponent which we identify. Our results also strongly suggest that the reference density used in [28] should indeed satisfy the technical requirements laid forth in [28] that any valid reference density should satisfy, which require that certain integrals with respect to the true stationary measure are finite. We note that [28] was not able to formally verify this, as the true scaling of the stationary measure was not known.

2.3.2. Qualitative insights from explicit exponent. Second, the explicit form of the large deviations exponent yields some concrete insights. For example, one finds that as $\theta \downarrow 0$, the large deviations exponent scales to first order like $-\frac{\mu_1\mu_2}{2(\mu_1+\mu_2-1)}\theta$, i.e. converges to 0 linearly in θ at that particular rate. Another interesting regime is the setting in which one of the service rates (say μ_2) diverges to ∞ , in which case the large deviations exponent converges to $-\frac{\theta\mu_1}{2}$, i.e. converges to a constant independent of μ_2 . We note that such a limit is conceptually related to the behavior of the $M/H_2^*/n + M$ queue, in which service times are with some probability 0 and with some probability exponentially distributed, whose limiting behavior in the HW regime was stated as an open question in [118].

2.3.3. Why no insensitivity? Third, the fact that Dai and He’s conjecture is false illustrates a more subtle manner in which abandonments change the qualitative behavior of the limiting processes, beyond the obvious incorporation of some sort of mean-reverting behavior. There are several

ways to understand this difference. One can reason as follows, somewhat heuristically. Without abandonments large deviations are only exponentially unlikely (not having a Gaussian decay). However, as the underlying limit processes are still continuous mappings of Gaussian processes, this exponential behavior is in some sense driven by the following phenomena: when considering the probability of a rare event (e.g. exceeding a large x), one must aggregate these underlying Gaussian processes over very large time intervals whose length depends on x itself, allowing the accumulation of a variance which itself depends on x , which leads (after standard Gaussian calculations) to the exponential (as opposed to Gaussian) decay. But the fact that one must aggregate these Gaussian processes over large intervals to induce this phenomena means that certain universality results from probability theory, e.g. the central limit theorem for renewal processes which plays an important role in the covariance structure of these underlying Gaussian processes, will have washed away many of the particulars of the underlying Gaussian processes, leading to the kind of insensitivity phenomena conjectured by Dai and He. However, in the setting of abandonments large deviations have a Gaussian decay, hence are not driven by such long-range aggregations, and hence do not exhibit such insensitivity. A second reason is as follows. For $z \in \{1, 2\}$, we will call a job assigned an $Exp(\mu_z)$ service time a type- z job. When the large deviations of the queue length have a Gaussian decay, the fluctuations of the number of type-1 and type-2 jobs in service are of the same qualitative magnitude (Gaussian) as the large deviations in the queue length, and hence more subtle interactions between the number of jobs waiting in queue and the type-composition of jobs must be considered when analyzing how rare events occur.

2.3.4. Connection to multi-dimensional OU processes. Fourth, and closely related to the third point, is that our proofs ultimately suggest a close connection between the rough behavior (at least up to the large deviations properties) of the piecewise OU process limits arising from queues with abandonments in the HW regime, and the much simpler multi-dimensional OU processes mentioned in Section 1. Although we will defer a formal description of these processes and their central role in our proofs to later sections, we now sketch the connection at a high level. As noted in Section 1, one of the main hurdles to understanding the piecewise OU processes arising in [29, 31, 24] is that the infinitesimal drift is a non-linear function of the multi-dimensional state of the system, involving terms of the form $\max(0, \bar{x}_1 + \bar{x}_2)$, with \bar{x} the state of the diffusion. If within the system of SDE defining these processes, one simply replaces the terms of the form $\max(0, \bar{x}_1 + \bar{x}_2)$ with $\bar{x}_1 + \bar{x}_2$, the drift becomes a linear function of the state, and the unique strong solution of the resulting system is a so-called multi-dimensional OU process, a well-understood family of Markov processes with essentially closed-form solutions. Intuitively, this multi-dimensional OU process corresponds to a relaxation of the true queueing system in which the queue length is

allowed to go negative while maintaining the linear relationships (e.g. as regards abandonments) which typically hold only when the queue is positive. Thus, for example, part of the dynamics are a mean-reversion back to zero at a rate containing θ while the queue is negative, a sort of “negative abandonment”. We note that through a simple linear transformation, we are further able to focus on a conceptually even simpler process in which the first component represents the “net infinitesimal drift” induced by the discrepancy between the number of type-1 and type-2 jobs in service, while the second component represents the total number of jobs in queue (which can be negative), with the dynamics again linear (i.e. leading to another multi-dimensional OU process) but with the time-derivative of the first component a function only of the first component (i.e. the coefficient in the appropriate linear combination dictating the contribution of the second component equals zero). Interestingly, our analysis shows that the large deviations exponent of the true limit process agrees with that of this “relaxed” multi-dimensional OU process.

2.3.5. Overview of insights from proofs. To better understand the connection and our proof approach, we now give a very brief overview of how our proof relates the true limit process to the aforementioned multi-dimensional OU process. Note that when there is a positive queue, the dynamics of the true limit process and the multi-dimensional OU process agree. Thus both our upper and lower bounds relate the true process to this multi-dimensional OU process by conceptually forcing a restriction to states in which the queue is positive. For our upper bound, this is accomplished by considering a sequence of modified queueing systems in which such a positivity is enforced by introducing extra arrivals whenever a server would otherwise go idle. The dynamics of this bounding system, in the limit, are essentially the same as that of the aforementioned multi-dimensional OU process under a certain type of reflection captured through a relevant Skorokhod-type mapping of a more primitive driving Gaussian process. We note that this asymptotic upper bound is actually very explicit, i.e. the supremum of a Gaussian process with completely explicit covariance function.

In our lower bound, this is accomplished by considering the true limit process initialized with its unknown stationary measure, picking a fixed T , and taking as a lower bound the probability that the stationary process exceeds the desired large value at T , AND happened to stay positive on $[0, T]$. As the true piecewise OU limit has the same dynamics as a multi-dimensional OU process over any such interval (as it will have never gone negative), it thus suffices to lower bound the corresponding probability for a multi-dimensional OU process initialized with not its own stationary measure, but with that of the complex piecewise OU limit. Although this measure is non-explicit (and in fact what we are trying to understand), it suffices to use non-explicit mixing-type arguments to reason that as this unknown measure is non-defective, by picking large enough T and

then considering the large-deviations behavior at time T one essentially achieves (modulo some error terms which can be dealt with) as a lower bound the probability that a (nicely initialized) multi-dimensional OU process is positive over an interval $[0, T]$ and exceeds a large value x at T . One then considers the double limit of fixing larger values of T , and for each considering the large deviations behavior (i.e. $x \rightarrow \infty$). We show that ultimately enforcing positivity on $[0, T]$ does not change the lower-bound on the large deviations behavior that one is able to show, which will coincide with that of the associated stationary multi-dimensional OU process, which matches our upper bound (i.e. the reflection considered in our upper bound does not change the large deviations behavior). Intuitively, this follows from the fact that the probability of remaining positive on $[0, T]$ is a function only of T , while the large deviations involve a limit $x \rightarrow \infty$, combined with the fact that enforcing positivity on $[0, T]$ does not decrease the probability of seeing a large deviation at time T (which we show using Slepian's inequality).

Indeed, we prove that both these upper-and-lower bounding processes/approaches lead to the same large deviations exponent, matching that of the aforementioned multi-dimensional OU process one gets by linearizing all relevant non-linearities which appear in the definition of the limit process. Although we only formalize these connections for the case of H_2 services, and for space considerations leave several of these connections implicit in our arguments and proofs, we believe that the aforementioned relationships should be quite robust, holding in much greater generality, and that the study of such multi-dimensional OU relaxations may yield a new set of tools for analyzing multi-server queues with abandonments in the HW regime broadly.

3. Upper bound.

In this section we present an upper bound for a generic $M/H_2/n + M$ queue Q (under Assumption 1), by considering (at least conceptually) a queueing system equivalent to Q but with an extra arrival added any time a server would otherwise have gone idle. By never transitioning to that part of the state-space in which there are idle servers, the non-linearities which make the analysis of the true queueing system difficult fade away. This methodology was pioneered in [41] for systems without abandonments, and has been applied to various queueing models without abandonments in [45, 46, 47]. However, several fundamental challenges arise when attempting to implement this framework in systems with abandonments. In particular, as noted in Section 1 and illustrated with a simple example, in systems with abandonments adding extra arrivals may actually cause the queue length to decrease along any given sample path. Although there is a small literature on stochastic comparison results for systems with abandonments, see e.g [15, 58, 18, 43, 26], none of these results seems strong enough to determine whether or not the modified system which we will consider yields an upper bound, not to mention pin down the precise large deviations behavior of the original queueing system.

3.1. CTMC representation for generic $M/H_2/n + M$ queue Q .

Due to the extra complexity introduced by abandonments, namely that our stochastic comparison results will no longer hold sample-path-wise (and instead require a more sophisticated coupling of various r.v), it will be simplest to formally define our modified upper-bounding system, and formalize our stochastic comparison arguments, in the language of CTMC. We begin by reminding the reader of the natural CTMC descriptor for a generic $M/H_2/n + M$ queue Q (not necessarily in the HW regime), which is standard in the literature (see e.g. [2, 29]). When referring to such a generic queue Q , we suppose the arrival rate is λ ; the service times are hyper-exponentially distributed with mean 1 and parameters μ_1, μ_2, p ; and patience times are $Expo(\theta)$. In particular, the state of the system is a 3-dimensional vector of non-negative integers, with the first component N_1 tracking the number of type-1 jobs in service, the second component N_2 tracking the number of type-2 jobs in service, and the third component W tracking the number of jobs waiting in queue. Here we recall that a type- z job is a job assigned an $Expo(\mu_z)$ service time. We note that under this description, a job's type is assigned only upon its entering service (if it ever actually enters service, as opposed to abandoning the system), which allows us to track only the total number of jobs waiting in queue, and not the types of those jobs. We state the relevant dynamics in terms of an appropriate infinitesimal generator, where for a countable-state CTMC and any two states x, y , the infinitesimal generator q evaluated at (x, y) , i.e. $q(x, y)$, corresponds to the infinitesimal rate at which the CTMC transitions to state y while in state x , as is standard in the literature. Here we denote the infinitesimal generator of Q by q . In that case, q has the following dynamics.

DEFINITION 1 (GENERATOR FOR Q).

- $q((N_1, N_2, W), (N_1 + 1, N_2, W)) = I(N_1 + N_2 \leq n - 1)\lambda p$;
- $q((N_1, N_2, W), (N_1, N_2 + 1, W)) = I(N_1 + N_2 \leq n - 1)\lambda(1 - p)$;
- $q((N_1, N_2, W), (N_1, N_2, W + 1)) = I(N_1 + N_2 = n)\lambda$;
- $q((N_1, N_2, W), (N_1 - 1, N_2, W)) = I(W = 0)\mu_1 N_1$;
- $q((N_1, N_2, W), (N_1, N_2 - 1, W)) = I(W = 0)\mu_2 N_2$;
- $q((N_1, N_2, W), (N_1 - 1, N_2 + 1, W - 1)) = I(W \geq 1)\mu_1 N_1(1 - p)$;
- $q((N_1, N_2, W), (N_1 + 1, N_2 - 1, W - 1)) = I(W \geq 1)\mu_2 N_2 p$;
- $q((N_1, N_2, W), (N_1, N_2, W - 1)) = I(W \geq 1)(\mu_1 N_1 p + \mu_2 N_2(1 - p) + \theta W)$.

Note that in the final transition, the contribution $\mu_1 N_1 p + \mu_2 N_2(1 - p)$ corresponds to the occurrence of a departure when the job at the head of the queue is of the same type as the departing job, while the contribution θW corresponds to an abandonment, with both of these events resulting in N_1 and N_2 remaining unchanged and W being reduced by 1.

3.2. CTMC representation for bounding system \bar{Q} .

We next formalize our upper-bounding system as a different CTMC. For ease of exposition, it will be simpler to introduce our upper-bounding system as its own CTMC, as opposed to initially constructing it on the same probability space as Q by adding extra arrivals. Thus we will never formalize that this bounding CTMC is a type of modified queueing system with abandonments, in which extra arrivals occur any time a server would otherwise have gone idle, as our formal proofs will not require it, and doing so requires a significant amount of additional notation and definitions. However, we will at times informally refer to this modified system as if it were a queueing system, referring to e.g. jobs, departures, abandonments, etc., to help build intuition.

Our upper-bounding system, denoted \bar{Q} , will similarly be a 3-dimensional CTMC on $Z^+ \times Z^+ \times Z^+$, where the components will have the same intuitive meaning as in Q . In particular, we now define the CTMC $\bar{Q} = (\bar{N}_1, \bar{N}_2, \bar{W})$. The dynamics of \bar{Q} as a CTMC are given by the infinitesimal generator \bar{q} as follows.

DEFINITION 2 (GENERATOR FOR \bar{Q}).

- $\bar{q}((\bar{N}_1, \bar{N}_2, \bar{W}), (\bar{N}_1, \bar{N}_2, \bar{W} + 1)) = \lambda$;
- $\bar{q}((\bar{N}_1, \bar{N}_2, \bar{W}), (\bar{N}_1 - 1, \bar{N}_2 + 1, \bar{W})) = I(\bar{W} = 0)\mu_1\bar{N}_1(1 - p)$;
- $\bar{q}((\bar{N}_1, \bar{N}_2, \bar{W}), (\bar{N}_1 + 1, \bar{N}_2 - 1, \bar{W})) = I(\bar{W} = 0)\mu_2\bar{N}_2p$;
- $\bar{q}((\bar{N}_1, \bar{N}_2, \bar{W}), (\bar{N}_1 - 1, \bar{N}_2 + 1, \bar{W} - 1)) = I(\bar{W} \geq 1)\mu_1\bar{N}_1(1 - p)$;
- $\bar{q}((\bar{N}_1, \bar{N}_2, \bar{W}), (\bar{N}_1 + 1, \bar{N}_2 - 1, \bar{W} - 1)) = I(\bar{W} \geq 1)\mu_2\bar{N}_2p$;
- $\bar{q}((\bar{N}_1, \bar{N}_2, \bar{W}), (\bar{N}_1, \bar{N}_2, \bar{W} - 1)) = I(\bar{W} \geq 1)(\mu_1\bar{N}_1p + \mu_2\bar{N}_2(1 - p) + \theta\bar{W})$.

Note that in \bar{Q} , when there are 0 jobs waiting in queue and a departure occurs (say of type-1), the transition which occurs is equivalent to what would happen if at that moment the departing job leaves, but in addition an arrival occurs, being type-1 w.p. p (in which case no transition occurs whatsoever as the state is unchanged), and type-2 w.p. $1 - p$ (in which case \bar{N}_1 decreases by 1 and \bar{N}_2 increases by 1). Otherwise the dynamics are the same as those of Q .

3.3. Proof that \bar{W} yields an upper bound for W .

We now prove that \bar{W} stochastically dominates W . For a general stochastic process $Z = (Z^1, \dots, Z^k)$, let $Z(t) = (Z^1(t), \dots, Z^k(t))$ denote the state of the process at time t .

THEOREM 2. *For all $n \geq 1$, all θ, μ_1, μ_2 satisfying Assumption 1, and all non-negative integer triplets (n_1, n_2, w) s.t. $n_1 + n_2 = n$ and $w = 0$, one may construct Q and \bar{Q} on a common probability space s.t. $Q(0) = \bar{Q}(0) = (n_1, n_2, w)$, and w.p.1 $\bar{W}(t) \geq W(t)$ for all $t \geq 0$.*

We note that under Assumption 1 such a result is intuitive, as the addition of extra arrivals to the system causes certain jobs to wait longer in queue, and assumption 1 ensures that (in an appropriate sense) the longer a job waitings in queue the “slower” it leaves the system. We also note that very similar proofs would allow us to prove an analogous result allowing for more general initial conditions for Q and \bar{Q} , e.g. having strictly positive initial queue-length, but as this is not needed for our main results, for clarity of exposition we restrict to the setting in which there are exactly n jobs initially in system. There are many ways to prove Theorem 2. We proceed by considering a certain expanded Markovian representation for both systems, to be able to keep track of individual jobs, and then formally construct both systems on a common probability space as projections of a larger CTMC. Alternative proofs proceed by e.g. using the virtual waiting time representation for the time-in-system [111]. The complexity with such alternative approaches is that certain “folk-theorem” type results for $M/G/n + M$ queues involving various representations and conditional independences related to Poisson processes and the system dynamics are not immediate for the modified system with extra arrivals, and rederiving these facts for our modified system (which no doubt hold) seems an even more tedious and cumbersome task than formally constructing the appropriate coupling as a CTMC (which, admittedly, is somewhat cumbersome). We note that such alternative approaches to proving the desired stochastic comparison may be important for extending the proof to general service and patience distributions, under an assumption analogous to Assumption 1 for the supremum of the hazard rate of the patience distribution and the infimum of the hazard rate of the service distribution.

3.3.1. Expanded Markovian representations Q^0 and \bar{Q}^0 for Q and \bar{Q} . Before proceeding with the proof of Theorem 2, it will be convenient to provide expanded Markovian representations for both Q and \bar{Q} , which we will denote by Q^0 and \bar{Q}^0 respectively, and which will enable us to explicitly couple the arrivals to both systems, as well as the abandonments from one system to service completions in the other. In these expanded representations, we will keep track of particular customers by identifying each customer with an index equal to either its arrival number (i.e. its position in the total arrival order, if it is a standard arrival), or a dummy index zero (if it is a so-called special/dummy/extra arrival to be defined below) as well as a type (either 1 or 2) specifying whether its exponential rate of service is μ_1 or μ_2 . Jobs with index $i \geq 1$ will be colloquially referred to as standard/regular jobs/arrivals/customers, while jobs with index 0 will be colloquially referred to as special/dummy/extra jobs/arrivals/customers. In summary, a given customer will be specified by a 2-dimensional vector (i.e. tuple), the first component being the index, the second component being the type.

We will present our expanded Markovian framework at a level of generality to encompass both

Q and \overline{Q} , but note that certain features (e.g. the presence of dummy jobs in service) will only be used meaningfully in our expanded representation of \overline{Q} . Our expanded Markovian representation will consist of three components, i.e. the state descriptor will have three components.

First, there will be an unordered multi-set of 2-dimensional vectors, \mathcal{S} , which will correspond to the set of jobs in service. Each of these vectors will have first component a non-negative integer, and second component belonging to $\{1, 2\}$. For any strictly positive integer i , \mathcal{S} will by construction contain at most one 2-dimensional vector with first component i , while \mathcal{S} may contain multiple copies of both $(0, 1)$ and $(0, 2)$. If $(0, 1)$ and/or $(0, 2)$ have multiplicity greater than one, then each copy will correspond to a different job in service. Let $|\mathcal{S}|$ denote the cardinality of this unordered multi-set, where we note that $|\mathcal{S}|$ is not necessarily the number of distinct elements, as the vectors $(0, 1)$ and $(0, 2)$ contribute their respective multiplicities to this cardinality. Our construction will ensure that $|\mathcal{S}| \leq n$. Given such an unordered multi-set \mathcal{S} of tuples, and $i \geq 0$, $z \in \{1, 2\}$, let $N(\mathcal{S}, i, z)$ denote the multiplicity of (i, z) in \mathcal{S} . Given $i \geq 0$ and $z \in \{1, 2\}$, let $\mathcal{S} \setminus (i, z)$ denote the unordered multi-set of tuples equivalent to \mathcal{S} , but with one copy of (i, z) removed. If $N(\mathcal{S}, i, z) = 0$, let $\mathcal{S} \setminus (i, z)$ simply denote \mathcal{S} . Similarly, let $\mathcal{S} \cup (i, z)$ denote the unordered multi-set of tuples equivalent to \mathcal{S} , but with an additional copy of (i, z) added. We note that by construction, all operations we perform on such unordered multi-sets will never attempt to add a tuple (i, z) if $i \geq 1$ and $\sum_{z=1}^2 N(\mathcal{S}, i, z)$ is already strictly positive. In general we allow these operations to be composed, e.g. $(\mathcal{S} \cup (i_1, z_1)) \setminus (i_2, z_2)$.

Second, there will be an ordered set (as opposed to multi-set) of 2-dimensional vectors, \mathcal{W} , which will correspond to the ordered set of jobs waiting in queue, with each vector corresponding to a particular job in queue. Each of these vectors will have first component a strictly positive integer (i.e. by construction dummy customers with index 0 never wait in queue, as they only arrive to the system when a server would otherwise have gone idle), and second component belonging to $\{1, 2\}$. Let $|\mathcal{W}|$ denote the cardinality of this ordered set. The job in position 1 will correspond to the job at the front of the queue (i.e. the next job to enter service). The job in position $|\mathcal{W}|$ will be the job at the back of the queue. Sometimes we will colloquially refer to the jobs as being in left-to-right order, with the job in position 1 being the left-most job, and the job in position $|\mathcal{W}|$ being the right-most job. Our construction will ensure that $|\mathcal{S}| \leq n - 1$ implies $\mathcal{W} = \emptyset$, and $\mathcal{W} \neq \emptyset$ implies $|\mathcal{S}| = n$. Given $i \geq 0$ and $z \in \{1, 2\}$, let $N(\mathcal{W}, i, z)$ denote the multiplicity of (i, z) in \mathcal{W} , where we note that $\sum_{z=1}^2 N(\mathcal{W}, i, z) \in \{0, 1\}$ for all $i \geq 0$, i.e. here the multiplicity reduces to the appropriate indicator. Given such a \mathcal{W} , and $i \geq 1$, $z \in \{1, 2\}$, let $\mathcal{W} \setminus (i, z)$ denote the ordered set of tuples equivalent to \mathcal{W} , but with element (i, z) removed, and the relative order of all remaining tuples the same as in \mathcal{W} , if $N(\mathcal{W}, i, z) = 1$. Otherwise, i.e. if $N(\mathcal{W}, i, z) = 0$, let $\mathcal{W} \setminus (i, z)$ simply denote \mathcal{W} . Also, given $i \geq 1$ and $z \in \{1, 2\}$, let $\mathcal{W} \cup (i, z)$ denote the ordered set of tuples equivalent to \mathcal{W} , but

with the tuple (i, z) appended to the right of all tuples in \mathcal{W} , with the relative order of the tuples already in \mathcal{W} unaltered. We note that by construction, all operations we perform on such ordered sets will never attempt to add a tuple (i, z) if $i \geq 1$ and $\sum_{z=1}^2 N(\mathcal{W}, i, z)$ is already strictly positive, nor attempt to add a tuple with index 0. Also, for such a set \mathcal{W} and $j \in [1, |\mathcal{W}|]$, let \mathcal{W}_j denote the tuple in position j in \mathcal{W} , and let $\mathcal{W}_{j,i}$ denote the i th component of that tuple, $i \in \{1, 2\}$. Given an event A , let $I(A)$ denote the corresponding indicator function. Also, given a 2-dimensional vector z , let $I(z \in \mathcal{W})$ evaluate to 1 if z appears in the ordered set \mathcal{W} and 0 otherwise.

Third, there will be a non-negative integer counter ν which keeps track of how many standard (i.e. non-dummy) arrivals there have been to the system, which will be necessary to properly assign indices to customers.

Thus the state of a queue is such a triplet $(\mathcal{S}, \mathcal{W}, \nu)$. Given such a triplet \mathcal{Q} , $i \geq 0$, and $z \in \{1, 2\}$, let $N(\mathcal{Q}, i, z) \triangleq N(\mathcal{S}, i, z) + N(\mathcal{W}, i, z)$.

We now use the above definitions and operations to define CTMC Q^0 (equivalent to Q in our expanded representation) and \overline{Q}^0 (equivalent to \overline{Q} in our expanded representation), which will provide an upper bound for Q^0 (in an appropriate sense). We first define the CTMC $Q^0 = (\mathcal{S}^0, \mathcal{W}^0, \nu^0)$, denoting the corresponding generator by q^0 . As a notational convenience, we define $p_1 \triangleq p, p_2 \triangleq 1 - p$.

DEFINITION 3 (GENERATOR FOR Q^0).

(i). If $|\mathcal{S}| \leq n - 1$, for $z \in \{1, 2\}$:

$$q^0 \left((\mathcal{S}, \mathcal{W}, \nu), (\mathcal{S} \cup (\nu + 1, z), \mathcal{W}, \nu + 1) \right) = \lambda p_z.$$

(ii). If $|\mathcal{S}| = n$, for $z \in \{1, 2\}$:

$$q^0 \left((\mathcal{S}, \mathcal{W}, \nu), (\mathcal{S}, \mathcal{W} \cup (\nu + 1, z), \nu + 1) \right) = \lambda p_z.$$

(iii). If $\mathcal{W} = \emptyset$:

(a) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}, i, z) = 1$:

$$q^0 \left((\mathcal{S}, \mathcal{W}, \nu), (\mathcal{S} \setminus (i, z), \mathcal{W}, \nu) \right) = \mu_z.$$

(iv). If $\mathcal{W} \neq \emptyset$:

(a) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}, i, z) = 1$:

$$q^0 \left(\left((\mathcal{S}, \mathcal{W}, \nu), \left((\mathcal{S} \setminus (i, z)) \cup \mathcal{W}_1, \mathcal{W} \setminus \mathcal{W}_1, \nu \right) \right) \right) = \mu_z.$$

(b) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{W}, i, z) = 1$:

$$q^0 \left((\mathcal{S}, \mathcal{W}, \nu), (\mathcal{S}, \mathcal{W} \setminus (i, z), \nu) \right) = \theta.$$

We next define the generator of the CTMC $\bar{Q}^0 = (\bar{\mathcal{S}}^0, \bar{\mathcal{W}}^0, \bar{\nu}^0)$, denoting the corresponding generator by \bar{q}^0 .

DEFINITION 4 (GENERATOR FOR \bar{Q}^0).

(i). For all $z \in \{1, 2\}$:

$$\bar{q}^0 \left((\mathcal{S}, \mathcal{W}, \nu), (\mathcal{S}, \mathcal{W} \cup (\nu + 1, z), \nu + 1) \right) = \lambda p_z.$$

(ii). If $\mathcal{W} = \emptyset$:

(a) For all $i \geq 1$ and $z_1, z_2 \in \{1, 2\}$ s.t. $N(\mathcal{S}, i, z_1) = 1$:

$$\bar{q}^0 \left((\mathcal{S}, \mathcal{W}, \nu), \left((\mathcal{S} \setminus (i, z_1)) \cup (0, z_2), \mathcal{W}, \nu \right) \right) = \mu_{z_1} p_{z_2}.$$

(b) For all $z_1, z_2 \in \{1, 2\}$ s.t. $z_1 \neq z_2$:

$$\bar{q}^0 \left((\mathcal{S}, \mathcal{W}, \nu), \left((\mathcal{S} \setminus (0, z_1)) \cup (0, z_2), \mathcal{W}, \nu \right) \right) = N(\mathcal{S}, 0, z_1) \mu_{z_1} p_{z_2}.$$

(iii). If $\mathcal{W} \neq \emptyset$:

(a) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}, i, z) = 1$:

$$\bar{q}^0 \left((\mathcal{S}, \mathcal{W}, \nu), \left((\mathcal{S} \setminus (i, z)) \cup \mathcal{W}_1, \mathcal{W} \setminus \mathcal{W}_1, \nu \right) \right) = \mu_z.$$

(b) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{W}, i, z) = 1$:

$$\bar{q}^0 \left((\mathcal{S}, \mathcal{W}, \nu), (\mathcal{S}, \mathcal{W} \setminus (i, z), \nu) \right) = \theta.$$

(c) For all $z \in \{1, 2\}$:

$$\bar{q}^0 \left((\mathcal{S}, \mathcal{W}, \nu), \left((\mathcal{S} \setminus (0, z)) \cup \mathcal{W}_1, \mathcal{W} \setminus \mathcal{W}_1, \nu \right) \right) = N(\mathcal{S}, 0, z) \mu_z.$$

We note that such non-standard Markovian representations for Markovian queues with abandonments are common in the literature, see e.g. [19]. We now define several notions which will act as a proxy for a state being “reasonable”.

DEFINITION 5 (GOOD STATE, 0-GOOD STATE, $\bar{0}$ -GOOD STATE, VERY-GOOD STATE).

In our expanded Markovian representation, a triplet $\mathcal{Q} = (\mathcal{S}, \mathcal{W}, \nu)$ is good if:

- $I(|\mathcal{S}| \leq n - 1, |\mathcal{W}| \geq 1) = 0$;
- For $i \geq 1$, $\sum_{z=1}^2 N(\mathcal{Q}, i, z) \leq 1$;
- $N(\mathcal{W}, 0) = 0$;
- For $1 \leq i < j \leq |\mathcal{W}|$, $\mathcal{W}_{i,1} < \mathcal{W}_{j,1}$;
- If $|\mathcal{W}| \geq 1$, then for all $(i, z) \in \mathcal{S}$, it holds that $\mathcal{W}_{1,1} > i$;
- For all i s.t. $\sum_{z=1}^2 N(\mathcal{Q}, i, z) = 1$, it holds that $i \leq \nu$.

Let us say that the triplet is:

- 0-good if: in addition to being good, $\sum_{z=1}^2 N(\mathcal{S}, 0, z) = 0$;
- $\bar{0}$ -good if: in addition to being good, $|\mathcal{S}| = n$;
- very-good if: it is both 0-good and $\bar{0}$ -good, and in addition $\nu = n$, $\mathcal{W} = \emptyset$, and $\sum_{z=1}^2 N(\mathcal{S}, i, z) = 1$ for all $i \in [1, n]$ (i.e. initially there is one job in service with index i for each $i \in [1, n]$, there are no jobs waiting in queue, and the initial counter equals n).

Note that the fourth and fifth requirements of being a good state ensure that the FCFS ordering is preserved, where we note that even in the presence of abandonments the relative order of those jobs in the system must obey the FCFS order.

We now formally state the precise connection between Q and Q^0 , and \bar{Q} and \bar{Q}^0 .

LEMMA 1. – For any very-good triplet $(\mathcal{S}, \mathcal{W}, \nu)$, one may construct Q and Q^0 on a common probability space s.t.: $Q^0(0) = (\mathcal{S}, \mathcal{W}, \nu)$, $Q(0) = \left(\sum_{i=1}^n N(\mathcal{S}, i, 1), \sum_{i=1}^n N(\mathcal{S}, i, 2), 0 \right)$, and w.p.1, for all $t \geq 0$, $N_1(t) = \sum_{i=1}^{\infty} N(\mathcal{S}^0(t), i, 1)$, $N_2(t) = \sum_{i=1}^{\infty} N(\mathcal{S}^0(t), i, 2)$, and $W(t) = |\mathcal{W}^0(t)|$.

Similarly, for any very-good triplet $(\mathcal{S}, \mathcal{W}, \nu)$, one may construct \bar{Q} and \bar{Q}^0 on a common probability space s.t.: $\bar{Q}^0(0) = (\mathcal{S}, \mathcal{W}, \nu)$, $\bar{Q}(0) = \left(\sum_{i=1}^n N(\mathcal{S}, i, 1), \sum_{i=1}^n N(\mathcal{S}, i, 2), 0 \right)$, and w.p.1, for all $t \geq 0$, $\bar{N}_1(t) = \sum_{i=0}^{\infty} N(\bar{\mathcal{S}}^0(t), i, 1)$, $\bar{N}_2(t) = \sum_{i=0}^{\infty} N(\bar{\mathcal{S}}^0(t), i, 2)$, and $\bar{W}(t) = |\bar{\mathcal{W}}^0(t)|$.

The proof of Lemma 1 is elementary and follows from a standard and well-known style of coupling argument, using the fact that in both Q^0 and \bar{Q}^0 the job indices do not impact the abandonment or service dynamics, and the fact that there is similarly no impact on the abandonment and service dynamics if each arriving job generates its type upon arrival, as opposed to generating its type upon its beginning service (if it ever begins service). We omit the details, and refer the interested reader to e.g. [107] for an example of related arguments.

Next, it will be useful to note that w.p.1, both Q^0 and \bar{Q}^0 are always in a “reasonable” state, if they are initialized in a “reasonable” state, as dictated by our previous definitions of 0-good and $\bar{0}$ -good states. In all cases the stated results follow from a straightforward induction, and we omit the details.

LEMMA 2. – If $Q^0(0)$ is θ -good, then w.p.1 $Q^0(t)$ is θ -good for all $t \geq 0$. Similarly, if $\bar{Q}^0(0)$ is $\bar{\theta}$ -good, then w.p.1 $\bar{Q}^0(t)$ is $\bar{\theta}$ -good for all $t \geq 0$.

We note that several of our later arguments will use these properties implicitly, especially e.g. the fact that the indices must obey the FCFS ordering and that there is at most one job with any given strictly positive index in the system at any given time, and for clarity of exposition sometimes we will use these properties without explicitly referring back to Lemma 2.

3.3.2. Formal statement of dominance and description of coupling. In light of Lemma 1, to prove the main upper bound result Theorem 2, it suffices to prove an analogous statement for Q^0 and \bar{Q}^0 , which we now formalize.

DEFINITION 6 (\geq COMPARATOR FOR EXPANDED MARKOVIAN REPRESENTATIONS). Given a θ -good triplet $\mathcal{Q}^1 = (\mathcal{S}^1, \mathcal{W}^1, \nu^1)$ and $\bar{\theta}$ -good triplet $\mathcal{Q}^2 = (\mathcal{S}^2, \mathcal{W}^2, \nu^2)$, let us say that $\mathcal{Q}^2 \geq \mathcal{Q}^1$ if:

- For all $i \geq 1$ and $z \in \{1, 2\}$, $N(\mathcal{W}^2, i, z) \geq N(\mathcal{W}^1, i, z)$;
- For all $i \geq 1$ and $z \in \{1, 2\}$, $N(\mathcal{Q}^2, i, z) \geq N(\mathcal{S}^1, i, z)$;
- $\nu^1 = \nu^2$.

Intuitively, $\mathcal{Q}^2 \geq \mathcal{Q}^1$ iff every job waiting in queue in \mathcal{Q}^1 is also waiting in queue in \mathcal{Q}^2 , every job in system in \mathcal{Q}^1 is also in system in \mathcal{Q}^2 (note that a job in service in \mathcal{Q}^1 may be in service or waiting in queue in \mathcal{Q}^2), and the arrival counters are equal. Then we will prove the following dominance result for Q^0 and \bar{Q}^0 .

THEOREM 3. For any very-good triplet $(\mathcal{S}, \mathcal{W}, \nu)$, one may construct Q^0 and \bar{Q}^0 on a common probability space s.t. $Q^0(0) = \bar{Q}^0(0) = (\mathcal{S}, \mathcal{W}, \nu)$, and w.p.1, for all $t \geq 0$, $\bar{Q}^0(t) \geq Q^0(t)$.

Combining with Lemmas 1 and 2, Theorem 3 is easily seen to imply the desired result Theorem 2, so we now focus on proving Theorem 3. We proceed by explicitly exhibiting such a construction/coupling as a 6-dimensional CTMC Q' , with the first 3 components corresponding to Q^0 , and the second 3 components corresponding to \bar{Q}^0 . In addition to coupling the arrival times and types of regular jobs, the key idea of the coupling is as follows. If a given job is in service in both Q^0 and \bar{Q}^0 , its departure is coupled between the two systems. However, if (e.g. due to the extra dummy jobs) a job is in service in Q^0 but still waiting in queue in \bar{Q}^0 , we couple its departure from Q^0 to its abandonment from \bar{Q}^0 by splitting the infinitesimal rate μ_z at which it departs from Q^0 into two parts: an infinitesimal rate of θ (which is synced to the abandonment process of the job from \bar{Q}^0) and an infinitesimal rate $\mu_z - \theta$ (which is independent from \bar{Q}^0). In this way, we are able to ensure that the job cannot depart from \bar{Q}^0 before it departs from Q^0 , preserving the desired dominance.

Of course, here Assumption 1 is needed for this coupling to make sense.

We first define the relevant CTMC $Q' = (\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'})$, denoting the corresponding generator by q' . When describing a generator q , in general we have written out $q(x,y)$ for every possible current state x and next state y . However, note that this is in some sense redundant, as the form of x is generic, so the same state x is written repeatedly. As here the state x is 6-dimensional and somewhat cumbersome, we will instead suppose that the state x is fixed at

$$\left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'} \right),$$

and when we need reference the entire state simply use the shorthand “ \cdot ”. In that case, the generator q' is as follows.

DEFINITION 7 (GENERATOR FOR Q').

1. If $|\mathcal{S}^{0'}| \leq n - 1$, for $z \in \{1, 2\}$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'} \cup (\nu^{0'} + 1, z), \mathcal{W}^{0'}, \nu^{0'} + 1, \bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'} \cup (\bar{\nu}^{0'} + 1, z), \bar{\nu}^{0'} + 1 \right) \right) = \lambda p_z.$$

2. If $|\mathcal{S}^{0'}| = n$, for $z \in \{1, 2\}$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'} \cup (\nu^{0'} + 1, z), \nu^{0'} + 1, \bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'} \cup (\bar{\nu}^{0'} + 1, z), \bar{\nu}^{0'} + 1 \right) \right) = \lambda p_z.$$

3. If $\mathcal{W}^{0'} = \emptyset$ and $\bar{\mathcal{W}}^{0'} = \emptyset$:

- (a) For all $i \geq 1$ and $z_1, z_2 \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z_1) + N(\mathcal{W}^{0'}, i, z_1) = 0, N(\bar{\mathcal{S}}^{0'}, i, z_1) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\bar{\mathcal{S}}^{0'} \setminus (i, z_1)) \cup (0, z_2), \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'} \right) \right) = \mu_{z_1} p_{z_2}.$$

- (b) For all $i \geq 1$ and $z_1, z_2 \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z_1) = 1, N(\bar{\mathcal{S}}^{0'}, i, z_1) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z_1), \mathcal{W}^{0'}, \nu^{0'}, (\bar{\mathcal{S}}^{0'} \setminus (i, z_1)) \cup (0, z_2), \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'} \right) \right) = \mu_{z_1} p_{z_2}.$$

- (c) For all $z_1, z_2 \in \{1, 2\}$ s.t. $z_1 \neq z_2$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\bar{\mathcal{S}}^{0'} \setminus (0, z_1)) \cup (0, z_2), \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'} \right) \right) = N(\bar{\mathcal{S}}^{0'}, 0, z_1) \mu_{z_1} p_{z_2}.$$

4. If $\mathcal{W}^{0'} = \emptyset$ and $\bar{\mathcal{W}}^{0'} \neq \emptyset$:

- (a) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0, N(\bar{\mathcal{S}}^{0'}, i, z) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\bar{\mathcal{S}}^{0'} \setminus (i, z)) \cup \bar{\mathcal{W}}_1^{0'}, \bar{\mathcal{W}}^{0'} \setminus \bar{\mathcal{W}}_1^{0'}, \bar{\nu}^{0'} \right) \right) = \mu_z.$$

(b) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0, N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) = \theta.$$

(c) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1, N(\overline{\mathcal{S}}^{0'}, i, z) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) = \mu_z.$$

(d) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1, N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

i.

$$q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) = \theta;$$

ii.

$$q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right) \right) = \mu_z - \theta.$$

(e) For all $z \in \{1, 2\}$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (0, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) = N(\overline{\mathcal{S}}^{0'}, 0, z) \mu_z.$$

5. If $\mathcal{W}^{0'} \neq \emptyset$ and $\overline{\mathcal{W}}^{0'} \neq \emptyset$:

(a) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0, N(\overline{\mathcal{S}}^{0'}, i, z) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) = \mu_z.$$

(b) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0, N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) = \theta.$$

(c) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1, N(\overline{\mathcal{S}}^{0'}, i, z) = 1$:

$$q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) = \mu_z.$$

(d) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1, N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

i.

$$q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) = \theta;$$

ii.

$$q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right) \right) = \mu_z - \theta.$$

(e) For all $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{W}^{0'}, i, z) = 1, N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'} \setminus (i, z), \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) = \theta.$$

(f) For all $z \in \{1, 2\}$:

$$q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (0, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) = N(\overline{\mathcal{S}}^{0'}, 0, z) \mu_z.$$

In that case, an argument completely analogous to Lemma 2 leads to the following, and we omit the details.

LEMMA 3. – *If $(\mathcal{S}^{0'}(0), \mathcal{W}^{0'}(0), \nu^{0'}(0))$ is 0-good and $(\overline{\mathcal{S}}^{0'}(0), \overline{\mathcal{W}}^{0'}(0), \overline{\nu}^{0'}(0))$ is $\overline{0}$ -good, then w.p.1 $(\mathcal{S}^{0'}(t), \mathcal{W}^{0'}(t), \nu^{0'}(t))$ is 0-good and $(\overline{\mathcal{S}}^{0'}(t), \overline{\mathcal{W}}^{0'}(t), \overline{\nu}^{0'}(t))$ is $\overline{0}$ -good for all $t \geq 0$.*

3.3.3. Proof of stochastic ordering. It is easily verified that to prove Theorem 3, it suffices to prove the following two lemmas.

LEMMA 4. *If: 1. $(\mathcal{S}^{0'}(0), \mathcal{W}^{0'}(0), \nu^{0'}(0))$ is 0-good, 2. $(\overline{\mathcal{S}}^{0'}(0), \overline{\mathcal{W}}^{0'}(0), \overline{\nu}^{0'}(0))$ is $\overline{0}$ -good, and 3. $(\overline{\mathcal{S}}^{0'}(0), \overline{\mathcal{W}}^{0'}(0), \overline{\nu}^{0'}(0)) \geq (\mathcal{S}^{0'}(0), \mathcal{W}^{0'}(0), \nu^{0'}(0))$, then $(\overline{\mathcal{S}}^{0'}(t), \overline{\mathcal{W}}^{0'}(t), \overline{\nu}^{0'}(t)) \geq (\mathcal{S}^{0'}(t), \mathcal{W}^{0'}(t), \nu^{0'}(t))$ for all $t \geq 0$.*

LEMMA 5. * *Under the same assumptions as Lemma 4, $\{(\mathcal{S}^{0'}(t), \mathcal{W}^{0'}(t), \nu^{0'}(t)), t \geq 0\}$ has the same distribution (at the process level) as $\{(\mathcal{S}^0(t), \mathcal{W}^0(t), \nu^0(t)), t \geq 0\}$; and $\{(\overline{\mathcal{S}}^{0'}(t), \overline{\mathcal{W}}^{0'}(t), \overline{\nu}^{0'}(t)), t \geq 0\}$ has the same distribution (at the process level) as $\{(\overline{\mathcal{S}}^0(t), \overline{\mathcal{W}}^0(t), \overline{\nu}^0(t)), t \geq 0\}$, in both cases supposing that corresponding processes have the same initial conditions.*

Proof of Lemma 4: It suffices to prove that if Q' is in a state $\mathcal{Q}' = \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right)$ s.t.: 1. $(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'})$ is 0-good, 2. $(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'})$ is $\overline{0}$ -good, and 3. $(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'}) \geq (\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'})$, then w.p.1 it can only transition to another state with those properties. We note that throughout, we will use the results of Lemma 3 implicitly. Our assumptions about \mathcal{Q}' are easily seen to imply that

$$|\overline{\mathcal{S}}^{0'}| = n \geq |\mathcal{S}^{0'}| \quad , \quad \text{and} \quad |\overline{\mathcal{W}}^{0'}| \geq |\mathcal{W}^{0'}|; \quad (1)$$

and also that for all $i \geq 1$ and $z \in \{1, 2\}$,

$$N(\overline{\mathcal{S}}^{0'}, i, z) + N(\overline{\mathcal{W}}^{0'}, i, z) \geq N(\mathcal{S}^{0'}, i, z) \quad , \quad N(\overline{\mathcal{W}}^{0'}, i, z) \geq N(\mathcal{W}^{0'}, i, z); \quad (2)$$

which we will also use. We proceed by a case analysis, analyzing each transition. Let $\mathcal{Q}'' = \left(\mathcal{S}^{0''}, \mathcal{W}^{0''}, \nu^{0''}, \overline{\mathcal{S}}^{0''}, \overline{\mathcal{W}}^{0''}, \overline{\nu}^{0''} \right)$ denote the state to which Q' next transitions. First, note that the

requirement involving the equivalence of ν'' and $\bar{\nu}''$ is easily seen to be maintained under all transitions (i.e. both parameters either stay unchanged or increase by 1, together, under every transition), so we may focus on the other requirements.

Let us first consider transitions (1) and (2). Under both transitions, our assumptions about \mathcal{Q}' imply that the required inequalities in the definition of \geq (ensuring \mathcal{Q}'' has the same desired properties) hold for all (i, z) s.t. $N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 1$ with $i \leq \nu^{0'}$. By construction, for the one additional job with index $\nu^{0''}$ (first supposing this job has type 1), it holds that $N(\mathcal{S}^{0''}, \nu^{0''}, 1) + N(\mathcal{W}^{0''}, \nu^{0''}, 1) = 1 = N(\bar{\mathcal{W}}^{0''}, \nu^{0''}, 1)$, which ensures that the required inequalities hold for all (i, z) and hence that \mathcal{Q}'' has the desired properties. The argument is identical if the new job has type 2.

Next, let us consider the transitions for which $\mathcal{W}^{0'} = \emptyset$ and $\bar{\mathcal{W}}^{0'} = \emptyset$. We note that under all cases s.t. $\mathcal{W}^{0'} = \emptyset$, we need not worry about the requirement (on \mathcal{Q}'') requiring that $N(\bar{\mathcal{W}}^{0''}, i, z) \geq N(\mathcal{W}^{0''}, i, z)$, as for the transitions being considered no arrivals occur and thus $\mathcal{W}^{0''}$ remains equal to \emptyset under the analyzed transitions. Under transition (3a), the desired properties for \mathcal{Q}'' are inherited from \mathcal{Q}' , as $N(\mathcal{S}^{0'}, i, z_1) + N(\mathcal{W}^{0'}, i, z_1) = 0$ anyways. Under transition (3b), the desired properties are inherited from \mathcal{Q}' , since the same (i, z) pair is removed from both $\mathcal{S}^{0'}$ and $\bar{\mathcal{S}}^{0'}$. Under transition (3c), the desired properties are inherited from \mathcal{Q}' as the only job removed from $\bar{\mathcal{S}}^{0'}$ has index 0.

Next, let us consider the transitions for which $\mathcal{W}^{0'} = \emptyset$ and $\bar{\mathcal{W}}^{0'} \neq \emptyset$. Under transitions (4a) and (4b), the desired properties are inherited from \mathcal{Q}' for the same reason as transition (3a). Under transitions (4c) and (4(d)i), the desired properties are inherited from \mathcal{Q}' for the same reason as transition (3b). Under transition (4(d)ii), the desired properties are inherited from \mathcal{Q}' as the job (i, z) is only removed from $\mathcal{S}^{0'}$, while remaining in $\bar{\mathcal{W}}^{0'}$, preserving the desired property (as it allows for any given (i, z) to appear only in the dominating system). Under transition (4e), the desired properties are inherited from \mathcal{Q}' for the same reason as transition (3c).

Finally, let us consider the transitions for which $\mathcal{W}^{0'} \neq \emptyset, \bar{\mathcal{W}}^{0'} \neq \emptyset$. Here a bit more care will have to be taken, as we will have to verify that $N(\mathcal{W}^{0''}, i, z) \leq N(\bar{\mathcal{W}}^{0''}, i, z)$, in addition to $N(\mathcal{S}^{0''}, i, z) \leq N(\bar{\mathcal{S}}^{0''}, i, z) + N(\bar{\mathcal{W}}^{0''}, i, z)$. Under transition (5a), the reasoning is the same as transition (3a), with the caveat that we must verify that $N(\mathcal{W}^{0''}, i, z) \leq N(\bar{\mathcal{W}}^{0''}, i, z)$. To do so, we will prove that $\bar{\mathcal{W}}_1^{0'} \notin \mathcal{W}^{0'}$, which (combined with our assumptions about \mathcal{Q}') is easily seen to suffice. Indeed, suppose for contradiction that $\bar{\mathcal{W}}_1^{0'} \in \mathcal{W}^{0'}$. Combined with the fact that $\mathcal{W}^{0'} \neq \emptyset$ and the definition of \geq , we conclude that there exist $(i_1, z_1), \dots, (i_n, z_n) \in \mathcal{S}^{0'}$ s.t.: $i_k < \bar{\mathcal{W}}_{1,1}^{0'}$ and $N(\bar{\mathcal{S}}^{0'}, i_k, z_k) + N(\bar{\mathcal{W}}^{0'}, i_k, z_k) = 1$ for all $k \in [1, n]$, and $\{i_k, k \in [1, n]\}$ is a set of n distinct strictly positive integers. Suppose for contradiction that for some such k , it holds that $N(\bar{\mathcal{W}}^{0'}, i_k, z_k) = 1$. Combined with the definition of \geq , we would conclude that $i_k \geq \bar{\mathcal{W}}_{1,1}^{0'}$, a contradiction. Thus it must be true that $N(\bar{\mathcal{S}}^{0'}, i_k, z_k) = 1$ for all k . But as by construction transition (5a) implies that there exists $(i^*, z^*) \in \bar{\mathcal{S}}^{0'}$ which does not belong to $\mathcal{S}^{0'}$, this (combined with the previous assertion) would imply that $|\bar{\mathcal{S}}^{0'}| = n + 1$, which

contradicts our assumptions on \mathcal{Q}' . Combining the above completes the proof.

The proof for transition (5b) follows for the same reason as (3a). For transition (5c), it is easily verified that the only case of possible concern is that $\overline{\mathcal{W}}_1^{0'} \in \mathcal{W}^{0'}$, yet $\overline{\mathcal{W}}_1^{0'} \neq \mathcal{W}_1^{0'}$. But this is impossible, as it would imply the existence of an $(i, z) \in \mathcal{W}^{0'}$ satisfying $i < \overline{\mathcal{W}}_{1,1}^{0'}$. But then by the assumed properties of \mathcal{Q}' , it would have to be the case that $(i, z) \in \overline{\mathcal{W}}^{0'}$, leading to a contradiction as there can exist no $(i, z) \in \overline{\mathcal{W}}^{0'}$ s.t. $i < \overline{\mathcal{W}}_{1,1}^{0'}$. Combining completes the proof for transition (5c). The proofs for transitions (5(d)i) and (5(d)ii) follows for the same reason as (4(d)i) and (4(d)ii). The proof for transition (5e) follows for the same reason as (3b), but applied to the queue itself.

For the final transition (5f), the reasoning is the same as transition (3c), with the caveat that we must verify that $N(\mathcal{W}^{0''}, i, z) \leq N(\overline{\mathcal{W}}^{0''}, i, z)$. It is easily verified that the only case of possible concern is that $\overline{\mathcal{W}}_1^{0'} \in \mathcal{W}^{0'}$. By the same argument used in analyzing transition (5a), we can conclude that if $\overline{\mathcal{W}}_1^{0'} \in \mathcal{W}^{0'}$, then it must hold that $\overline{\mathcal{W}}_1^{0'} = \mathcal{W}_1^{0'}$. The fact that $\mathcal{W}^{0'} \neq \emptyset$ then implies that there exist $\{(i_k, z_k), k = 1, \dots, n\}$ s.t. $N(\mathcal{S}^{0'}, i_k, z_k) = 1$ for all k , $\{i_k, k = 1, \dots, n\}$ are distinct strictly positive integers, and $i_k < \mathcal{W}_{1,1}^{0'} = \overline{\mathcal{W}}_{1,1}^{0'}$ for all k . But as transition (5f) can only occur if $N(\overline{\mathcal{S}}^{0'}, 0, z) \geq 1$ for some $z \in \{1, 2\}$, it follows that transition (5f) can only occur if $\overline{\mathcal{S}}^{0'}$ contains at most $n - 1$ elements from the set $\{(i_k, z_k), k = 1, \dots, n\}$, and hence there exists at least one such tuple $(i^0, z^0) \in \{(i_k, z_k), k = 1, \dots, n\}$ which does not belong to $\overline{\mathcal{S}}^{0'}$. But our assumptions about \mathcal{Q}' then imply that (i^0, z^0) must belong to $\overline{\mathcal{W}}^{0'}$. But this yields a contradiction, since $i^0 < \overline{\mathcal{W}}_{1,1}^{0'}$, yet all $(i, z) \in \overline{\mathcal{W}}^{0'}$ satisfy $i \geq \overline{\mathcal{W}}_{1,1}^{0'}$. Combining the above completes the proof. *Q.E.D.*

We defer the proof of Lemma 5 to the appendix. Combining the above then completes the proof of Theorem 3.

4. Skorokhod representation for upper bound \overline{W} .

To analyze the distribution of \overline{W} (generally, i.e. not necessarily in the HW regime), we will show that \overline{W} can be expressed as the solution to a certain generalized Skorokhod problem, on an appropriate probability space. We will then use a result of [100] to express the distribution in terms of a certain (asymptotically) tractable supremum. We do note that our analysis also yields an upper bound for any fixed n , which is roughly the supremum of many coupled mean-reverting random walks, and leave as an interesting direction for future research how best to interpret and use this bound for any fixed n (again, asymptotically it becomes a tractable supremum of a Gaussian process).

4.1. Review of generalized drift Skorokhod problem.

We begin by formally reviewing the generalized drift Skorokhod problem (as defined in [100]), equivalently the linearly generalized regulator mapping (as defined in [99, 111]). We only describe the relevant problems in the narrow context needed for our purposes, e.g. restricting to one-dimension

and linear drift. Recall that $D[0, \infty)$ denotes the appropriate space of functions, and we refer the interested reader to [115] for further details. In general, for a mapping Z from $D[0, \infty)$ to $D[0, \infty)$, i.e. from functions to functions, we use the notation $Z(X, t)$ to denote the function $Z(X)$ evaluated at time t , for $X \in D[0, \infty)$. Similarly, for a mapping Z from $D[0, \infty) \times \mathcal{R}$ to $D[0, \infty)$, we use the notation $Z(X, s, t)$ to denote the function $Z(X, s)$ evaluated at time t .

DEFINITION 8 (GENERALIZED DRIFT SKOROKHOD PROBLEM IN 1-DIMENSION [100, 99, 111]). Given a function $X \in D[0, \infty)$ with $X(0) = 0$, there exists a unique pair of functions $\Phi(X) : D[0, \infty) \rightarrow D[0, \infty)$ and $\Psi(X) : D[0, \infty) \rightarrow D[0, \infty)$ s.t.:

- $\Phi(X, t) = X(t) - \theta \int_0^t \Phi(X, s) ds + \Psi(X, t)$ for all $t \geq 0$;
- $\Phi(X, t) \geq 0$ and $\Psi(X, t) \geq 0$ for all $t \geq 0$;
- $\Phi(X, 0) = \Psi(X, 0) = 0$;
- $\Psi(X, t)$ is a non-decreasing function of t ;
- $\int_0^t I(\Phi(X, s) > 0) d\Psi(X, s) = 0$ for all $t \geq 0$.

As is standard in the literature, $\int_0^t I(\Phi(X, s) > 0) d\Psi(X, s)$ should be formally understood as a so-called Lebesgue-Stieltjes integral [51], where all such integrals appearing throughout should be similarly interpreted.

We next review an important representation property of Φ , proven in [100]. We begin by defining an auxiliary map.

DEFINITION 9 ([99]). Given a function $X \in D[0, \infty)$ with $X(0) = 0$, for each $s \geq 0$, there exists a unique function $\zeta(X, s) \in D[0, \infty)$ s.t. for all $t \geq 0$,

$$\zeta(X, s, t) = X(s+t) - X(s) - \theta \int_0^t \zeta(X, s, y) dy. \quad (3)$$

In that case, the following representation theorem is proven in [100].

LEMMA 6 ([100]). For any function $X \in D[0, \infty)$ s.t. $X(0) = 0$, for all $t \geq 0$, $\Phi(X, t) = \sup_{0 \leq s \leq t} \zeta(X, s, t-s)$.

We next review an important continuity property of Φ , proven in e.g. [99, 111].

LEMMA 7. [99, 111] Φ is Lipschitz continuous, under the uniform norm, as a function from $D[0, \infty)$ to $D[0, \infty)$. Namely, for each $T \geq 0$, there exists an absolute constant C_T (depending implicitly on θ) s.t. for all X_1 and $X_2 \in D[0, T]$ with $X_1(0) = X_2(0) = 0$, $\sup_{0 \leq t \leq T} |\Phi(X_1, t) - \Phi(X_2, t)| \leq C_T \sup_{0 \leq t \leq T} |X_1(t) - X_2(t)|$.

4.2. Expressing \overline{W} as a solution to the generalized drift Skorokhod problem.

We now prove that \overline{W} may be expressed as the solution to an appropriate generalized drift Skorokhod problem. Intuitively, the arrival of dummy jobs which prevent any servers from going idle will coincide with an appropriate pushing process (appearing in an appropriate Skorokhod problem) which keeps the number in system above n , i.e. we will interpret \overline{W} as a certain type of reflection. Recall that a pooled equilibrium renewal process (of n processes) with renewal distribution S corresponds to the counting process tracking the number of events which have occurred collectively in n independent renewal processes, each of which has renewal distribution S , and for which the initial residual life is independently drawn from the equilibrium distribution of S for each component renewal process. Here we recall that the equilibrium distribution of S is itself a hyper-exponential distribution, with parameters $\mu_1, \mu_2, \hat{p} = \frac{p}{\mu_1}$. Let REN be such an explicitly constructed pooled equilibrium renewal process. On the same probability space, let RES₁ and RES₂ be the stochastic processes tracking the number of these renewal processes on which the current residual life has an $Expo(\mu_1)$ and $Expo(\mu_2)$ distribution, respectively. It is easy to see (and well-known) that the dynamics of (RES_1, RES_2, REN) can be represented as a CTMC, with dynamics given by the following generator q_{Ren} .

- $q_{Ren}((RES_1, RES_2, REN), (RES_1 - 1, RES_2 + 1, REN + 1)) = \mu_1 RES_1 p_2;$
- $q_{Ren}((RES_1, RES_2, REN), (RES_1 + 1, RES_2 - 1, REN + 1)) = \mu_2 RES_2 p_1;$
- $q_{Ren}((RES_1, RES_2, REN), (RES_1, RES_2, REN + 1)) = \mu_1 RES_1 p_1 + \mu_2 RES_2 p_2.$

We note that under the above construction, letting $Bin(n, \hat{p})$ denote a binomially distributed r.v. with parameters n and \hat{p} , $(RES_1(0), RES_2(0), REN(0))$ is equivalent in distribution to $(Bin(n, \hat{p}), n - Bin(n, \hat{p}), 0)$. Let ARR be a rate λ Poisson process, independent of (RES_1, RES_2, REN) . Furthermore, let CLOCK be a rate 1 Poisson process, whose dependency structure w.r.t. ARR, RES₁, RES₂, REN we will intentionally leave unspecified. Supposing that we have constructed ARR, REN, CLOCK, and \overline{Q} on a common probability space, let us define SKOR (on the same probability space) to be the stochastic process s.t. for all $t \geq 0$,

$$SKOR(t) = n^{-\frac{1}{2}}(ARR(t) - REN(t)) - n^{-\frac{1}{2}} \left(CLOCK(\theta \int_0^t \overline{W}(s) ds) - \theta \int_0^t \overline{W}(s) ds \right).$$

Then our main result expressing \overline{W} as a solution to the generalized drift Skorokhod problem will be as follows.

LEMMA 8. *One may construct $\overline{Q}, (RES_1, RES_2, REN), ARR,$ and CLOCK on the same probability space, with ARR independent of $(RES_1, RES_2, REN),$ s.t. $\overline{Q}(0) = (RES_1(0), RES_2(0), REN(0)),$ and w.p.1, for all $t \geq 0, n^{-\frac{1}{2}}\overline{W}(t) = \Phi(SKOR, t).$*

To prove Lemma 8, it will be helpful to again consider an expanded Markovian representation for \bar{Q} , in which we explicitly keep track of the number of arrivals, departures, abandonments, and dummy job arrivals. We define a 7-dimensional CTMC $\tilde{Q} = (\tilde{R}\tilde{E}S_1, \tilde{R}\tilde{E}S_2, \tilde{W}, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y)$. The dynamics of \tilde{Q} are given by the generator \tilde{q} as follows, where as before we denote the generic current state

$$(\tilde{R}\tilde{E}S_1, \tilde{R}\tilde{E}S_2, \tilde{W}, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y)$$

simply by “ \cdot ”. Intuitively, $\tilde{R}\tilde{E}S_1(\tilde{R}\tilde{E}S_2)$ will correspond to the number of jobs in service with residual life distributed as $Expo(\mu_1)(Expo(\mu_2))$; \tilde{W} will correspond to the number of jobs waiting in queue; $\tilde{A}\tilde{R}R$ will correspond to the number of non-dummy arrivals so far; $\tilde{R}\tilde{E}N$ will correspond to the total number of departures so far; $\tilde{A}\tilde{B}A$ will correspond to the total number of abandonments so far; and $\tilde{D}\tilde{M}Y$ will correspond to the total number of dummy arrivals so far.

DEFINITION 10 (GENERATOR FOR \tilde{Q}).

- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1 - 1, \tilde{R}\tilde{E}S_2 + 1, \tilde{W}, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N + 1, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y + 1)\right) = I(\tilde{W} = 0)\mu_1\tilde{R}\tilde{E}S_1p_2;$
- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1 + 1, \tilde{R}\tilde{E}S_2 - 1, \tilde{W}, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N + 1, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y + 1)\right) = I(\tilde{W} = 0)\mu_2\tilde{R}\tilde{E}S_2p_1;$
- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1, \tilde{R}\tilde{E}S_2, \tilde{W}, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N + 1, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y + 1)\right) = I(\tilde{W} = 0)(\mu_1\tilde{R}\tilde{E}S_1p_1 + \mu_2\tilde{R}\tilde{E}S_2p_2);$
- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1, \tilde{R}\tilde{E}S_2, \tilde{W} + 1, \tilde{A}\tilde{R}R + 1, \tilde{R}\tilde{E}N, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y)\right) = \lambda;$
- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1 - 1, \tilde{R}\tilde{E}S_2 + 1, \tilde{W} - 1, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N + 1, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y)\right) = I(\tilde{W} \geq 1)\mu_1\tilde{R}\tilde{E}S_1p_2;$
- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1 + 1, \tilde{R}\tilde{E}S_2 - 1, \tilde{W} - 1, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N + 1, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y)\right) = I(\tilde{W} \geq 1)\mu_2\tilde{R}\tilde{E}S_2p_1;$
- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1, \tilde{R}\tilde{E}S_2, \tilde{W} - 1, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N + 1, \tilde{A}\tilde{B}A, \tilde{D}\tilde{M}Y)\right) = I(\tilde{W} \geq 1)(\mu_1\tilde{R}\tilde{E}S_1p_1 + \mu_2\tilde{R}\tilde{E}S_2p_2);$
- $\tilde{q}\left(\cdot, (\tilde{R}\tilde{E}S_1, \tilde{R}\tilde{E}S_2, \tilde{W} - 1, \tilde{A}\tilde{R}R, \tilde{R}\tilde{E}N, \tilde{A}\tilde{B}A + 1, \tilde{D}\tilde{M}Y)\right) = \theta\tilde{W}.$

We next make several observations regarding \tilde{Q} , and its relation to \bar{Q} . Then our first observation relates \bar{W} to \tilde{W} , and follows from a straightforward comparison of \bar{q} and \tilde{q} , the details of which we omit.

LEMMA 9. – *For all non-negative integer triplets (n_1, n_2, w) s.t. $n_1 + n_2 = n$ and $w = 0$, one can construct \bar{Q} and \tilde{Q} on a common probability space s.t. $\bar{Q}(0) = (n_1, n_2, w)$, $\tilde{Q}(0) = (n_1, n_2, w, 0, 0, 0, 0)$, and w.p.1, for all $t \geq 0$, $\bar{W}(t) = \tilde{W}(t)$.*

Our second observation relates $(\tilde{R}\tilde{E}S_1, \tilde{R}\tilde{E}S_2, \tilde{R}\tilde{E}N, \tilde{A}\tilde{R}R)$ to (RES_1, RES_2, REN, ARR) , and follows from a straightforward comparison of \tilde{q} and q_{ren} , the details of which we omit.

LEMMA 10. – *One can construct (RES_1, RES_2, REN) and ARR on the same probability space as \tilde{Q} s.t. w.p.1, for all $t \geq 0$, $(\tilde{R}\tilde{E}S_1(t), \tilde{R}\tilde{E}S_2(t), \tilde{R}\tilde{E}N(t), \tilde{A}\tilde{R}R(t))$ equals $(RES_1(t), RES_2(t), REN(t), ARR(t))$.*

Next, we observe that as the infinitesimal rate at which $\tilde{A}\tilde{B}A$ jumps (always up by 1) is at all times equal to $\theta\tilde{W}$, we may conclude the following from standard random-time-change constructions (see e.g. [90]), and we omit the details.

LEMMA 11. – *One may construct \tilde{Q} on the same probability space as $CLOCK$ s.t. w.p.1, for all $t \geq 0$, $\tilde{A}\tilde{B}A(t) = CLOCK(\theta \int_0^t \tilde{W}(s) ds)$, where this integral is w.p.1 well-defined as a Riemann integral for all $t \geq 0$.*

Finally, we note that a straightforward examination of \tilde{q} and proof by induction, the details of which we omit, reveals the following relationships between \tilde{W} , $\tilde{A}\tilde{R}R$, $\tilde{R}\tilde{E}N$, $\tilde{A}\tilde{B}A$, and $\tilde{D}\tilde{M}Y$.

LEMMA 12. – *For all non-negative integer triplets (n_1, n_2, w) s.t. $n_1 + n_2 = n$ and $w = 0$, one can construct \tilde{Q} s.t. $\tilde{Q}(0) = (n_1, n_2, w, 0, 0, 0, 0)$, and w.p.1, for all $t \geq 0$:*

- $\tilde{W}(t) = \tilde{A}\tilde{R}R(t) - \tilde{R}\tilde{E}N(t) - \tilde{A}\tilde{B}A(t) + \tilde{D}\tilde{M}Y(t)$;
- $\int_0^t I(\tilde{W}(s) > 0) d\tilde{D}\tilde{M}Y(s) = 0$.

Combining Lemmas 9 - 12 and noting that all relevant constructions can themselves be implemented together on a common probability space, the fact that w.p.1 $\tilde{D}\tilde{M}Y$ is non-decreasing and initialized at 0 for appropriate initial conditions, and Definition 8 then completes the proof of the desired result Lemma 8 (after scaling through by $n^{-\frac{1}{2}}$).

5. Asymptotic analysis of \overline{W} and proof of upper bound for Theorem 1.

In this section we provide an asymptotic analysis of \overline{W} in the HW regime, and complete the proof of one direction of our main result Theorem 1. We proceed as follows. First, we prove an asymptotic version of Lemma 8 in the HW regime. Second, we prove that the associated limit may be rewritten in terms of an integrated OU process. Third, we use this OU representation, along with the supremum representation for the generalized Skorokhod problem from Lemma 6, to prove that in the HW regime \overline{W} may be expressed as the supremum of an explicit Gaussian process. Finally, we analyze this Gaussian process, and combine with the interchange-of-limits results of [31, 24] mentioned in Section 1 to complete the proof of one direction of our main result Theorem 1.

5.1. Asymptotic version of Lemma 8.

We now prove an asymptotic version of Lemma 8, in the HW regime. We begin by introducing a superscript notation to make the dependence of some of the previously defined stochastic processes on n more explicit. Let ARR^n denote the corresponding Poisson process when $\lambda = \lambda_n = n + Bn^{\frac{1}{2}}$; REN^n denote the corresponding pooling of n independent equilibrium renewal processes; RES_1^n and RES_2^n denote the number of these renewal processes with the corresponding residual life distribution (respectively); \overline{Q}^n and \overline{W}^n denote the corresponding stochastic processes when there are n servers and $\lambda = \lambda_n$; all constructed on a common probability space (with CLOCK) using the construction given in Lemma 8. Note that under this construction, $\overline{Q}^n(0) = (\text{RES}_1^n(0), \text{RES}_2^n(0), \text{REN}^n(0))$, i.e. there are exactly n jobs in system with the initial residual life of each job drawn i.i.d from the equilibrium distribution of S , i.e. $\text{Expo}(\mu_1)$ with probability $\hat{p} = \frac{p}{\mu_1}$ and $\text{Expo}(\mu_2)$ with probability $\frac{1-p}{\mu_2}$. On this same probability space, let us define (in analogy with our previous definition) SKOR^n to be the stochastic process s.t. for all $t \geq 0$,

$$\text{SKOR}^n(t) = n^{-\frac{1}{2}}(\text{ARR}^n(t) - \text{REN}^n(t)) - n^{-\frac{1}{2}}\left(\text{CLOCK}\left(\theta \int_0^t \overline{W}^n(s)ds\right) - \theta \int_0^t \overline{W}^n(s)ds\right).$$

Let REN^∞ denote the unique centered continuous Gaussian process s.t. $\mathbb{E}[\text{REN}^\infty(s)\text{REN}^\infty(t)] = \mathbb{E}[(\text{REN}^1(s) - s)(\text{REN}^1(t) - t)]$ for all $0 \leq s \leq t$, i.e. the covariance structure is the same as that of a single equilibrium renewal process, where existence of such a process is proven in [113]. Let \mathcal{B}^0 denote a standard Brownian motion, independent of REN^∞ . Let SKOR^∞ denote the stochastic process s.t. for all $t \geq 0$, $\text{SKOR}^\infty(t) = \mathcal{B}^0(t) - \text{REN}^\infty(t) + Bt$.

Then the main asymptotic result of this section will be the following.

LEMMA 13. *For any $T \in [0, \infty)$, the sequence of processes $\{n^{-\frac{1}{2}}\overline{W}^n(t)_{0 \leq t \leq T}, n \geq 1\}$ converges weakly to $\Phi(\text{SKOR}^\infty)_{0 \leq t \leq T}$ in the space $D[0, T]$ under the J_1 topology.*

To prove Lemma 13, we begin by recalling a known weak convergence result from [41].

LEMMA 14 ([41]). *For any $T \in [0, \infty)$, the sequence of processes $\{n^{-\frac{1}{2}}(\text{ARR}^n(t) - \text{REN}^n(t))_{0 \leq t \leq T}, n \geq 1\}$ converges weakly to $\text{SKOR}^\infty(t)_{0 \leq t \leq T}$ in the space $D[0, T]$ under the J_1 topology.*

We next prove a result showing that the process corresponding to the second term in the definition of SKOR^n converges weakly to 0. Namely, let ERR^n denote the stochastic process s.t. for all $t \geq 0$,

$$\text{ERR}^n(t) = n^{-\frac{1}{2}}\left(\text{CLOCK}\left(\theta \int_0^t \overline{W}^n(s)ds\right) - \theta \int_0^t \overline{W}^n(s)ds\right).$$

Then we prove the following, deferring the proof to the appendix.

LEMMA 15. * For any $T \in [0, \infty)$, the sequence of processes $\{\text{ERR}^n(t)_{0 \leq t \leq T}, n \geq 1\}$ converges weakly to $0_{0 \leq t \leq T}$ in the space $D[0, T]$ under the J_1 topology.

Combining Lemmas 14 and 15 with Lemmas 8 and 7, the fact that convergence in the J_1 metric to a continuous process implies convergence in the uniform metric, and the continuous mapping theorem (see e.g. [115]) then completes the proof of Lemma 13.

5.2. Relating SKOR^∞ to an integrated OU process.

In this section we prove that SKOR^∞ is equal in distribution to the sum of a so-called time-integrated OU process (to be defined) plus an independent Brownian motion. Let OU denote the unique w.p.1 continuous centered Gaussian process s.t. $\mathbb{E}[\text{OU}(s)\text{OU}(t)] = p(1-p)(\mu_1\mu_2)^{-1} \exp(-\mu_1\mu_2(t-s))$ for $0 \leq s \leq t$, equivalently an appropriate stationary OU process, i.e. the strictly stationary process which is the unique such strong solution to the SDE

$$d\text{OU}(t) = -\mu_1\mu_2\text{OU}(t)dt + (2p(1-p))^{\frac{1}{2}}d\mathcal{B}^0(t), \quad (4)$$

with \mathcal{B}^0 the previously defined standard Brownian motion. Here we refer the interested reader to [33] for further details regarding existence and various properties of this Gaussian process, and to [83] for further review of associated SDEs. As continuous functions are Riemann integrable, it follows that OU is w.p.1 Riemann integrable over compact time intervals. Thus we may define IOU to be the stochastic process s.t. $\text{IOU}(t) = \int_0^t \text{OU}(y)dy$ for all $t \geq 0$. Such a process is known as a time-integrated OU process, and is well-studied, e.g. in [12] where the associated variance is computed explicitly (as stated below). Furthermore, as IOU has stationary increments (see e.g. [41]), $\mathbb{E}[\text{IOU}^2(t-s)] = \mathbb{E}[\text{IOU}^2(t)] + \mathbb{E}[\text{IOU}^2(s)] - 2\mathbb{E}[\text{IOU}(s)\text{IOU}(t)]$ for $0 \leq s \leq t$. It follows that the variance calculation in [12], combined with some straightforward algebra the details of which we omit, yields the following.

LEMMA 16 ([12]–). IOU is a continuous centered Gaussian process with stationary increments s.t. for all $0 \leq s \leq t$, $\mathbb{E}[\text{IOU}(s)\text{IOU}(t)]$ equals

$$p(1-p)(\mu_1\mu_2)^{-3} \left(2\mu_1\mu_2s + \exp(-\mu_1\mu_2s) + \exp(-\mu_1\mu_2t) - \exp(-\mu_1\mu_2(t-s)) - 1 \right).$$

Suppose that IOU and OU have been constructed on the same probability space as standard Brownian motions \mathcal{B}^0 and \mathcal{B}^1 , with \mathcal{B}^0 the driving Brownian motion of OU, $\text{IOU}(t) = \int_0^t \text{OU}(s)ds$ for all $t \geq 0$, and OU, IOU, \mathcal{B}^0 independent of \mathcal{B}^1 . On the same probability space, let us define the stochastic process Λ s.t. for all $t \geq 0$,

$$\Lambda(t) = -(\mu_1 - \mu_2)\text{IOU}(t) + 2^{\frac{1}{2}}\mathcal{B}^1(t) + Bt.$$

Then our next result formalizes the relationship between SKOR^∞ and Λ , and we defer the proof to the appendix.

LEMMA 17. * $\{\Lambda(t), t \geq 0\}$ has the same distribution, at the process level, as $\{\text{SKOR}^\infty(t), t \geq 0\}$.

5.3. Stochastic calculus to analyze Λ .

In this section, we use stochastic calculus to rewrite $\zeta(\Lambda, s)$ in a form amenable to analyzing the supremum arising in Lemma 6, ultimately relating $\Phi(\Lambda, t)$ to the supremum of a certain Gaussian process. Our main result will be the following. Let us define \mathcal{G} to be the stochastic process s.t. for all $t \geq 0$,

$$\mathcal{G}(t) = -(\mu_1 - \mu_2) \int_0^t \exp(-\theta y) \text{OU}(y) dy + 2^{\frac{1}{2}} \int_0^t \exp(-\theta y) d\mathcal{B}^1(y) + \frac{B}{\theta} \left(1 - \exp(-\theta t) \right),$$

with OU independent of \mathcal{B}^1 . Then we prove the following.

LEMMA 18. *For all $t \geq 0$, $\Phi(\Lambda, t)$ has the same distribution as $\sup_{0 \leq s \leq t} \mathcal{G}(s)$.*

In the rest of this subsection we prove Lemma 18. We begin by providing some more explicit representations/constructions for several previously defined stochastic processes, which follow immediately from (4), standard and well-known results for OU processes [73], a few simple reindexings, and our definitions (especially Definition 9), and we omit the details. We note that in this and remaining sections, we freely use standard notations, definitions, and basic results from the elementary theory of stochastic calculus, and refer the interested reader to [62] for further details. Also, for a general stochastic process X and $s \geq 0$, we let X_s denote the stochastic process s.t. $X_s(t) = X(s+t) - X(s)$ for all $t \geq 0$. Similarly, for a general stochastic process X and $s \geq 0$, we let $X_{+,s}$ denote the stochastic process s.t. $X_{+,s}(t) = X(s+t)$ for all $t \geq 0$. For a generic r.v. X , let $\text{Var}[X]$ denote the variance of X . Let \mathcal{B}^3 and \mathcal{B}^2 denote two independent Brownian motions, constructed on the same probability space, which satisfy $\text{Var}[\mathcal{B}^3(1)] = 2p(1-p)$, $\text{Var}[\mathcal{B}^2(1)] = 2$. Then the desired representations are as follows, where we remind the reader that the mapping ζ , from $D[0, \infty) \times \mathcal{R}^+$ to $D[0, \infty)$, was previously formalized in Definition 9.

LEMMA 19. *– One may construct Λ and OU on the same probability space as \mathcal{B}^3 and \mathcal{B}^2 , s.t. w.p.1, for all $s, t \geq 0$,*

$$\text{OU}_{+,s}(t) = \exp(-\mu_1 \mu_2 t) \text{OU}_{+,s}(0) + \int_0^t \exp(-\mu_1 \mu_2 (t-y)) d\mathcal{B}_s^3(y); \quad (5)$$

$$\zeta(\Lambda, s, t) = -(\mu_1 - \mu_2) \int_0^t \text{OU}_{+,s}(y) dy + \mathcal{B}_s^2(t) + Bt - \theta \int_0^t \zeta(\Lambda, s, y) dy. \quad (6)$$

From Lemma 19 and basic definitions in stochastic calculus, see e.g. [62], we conclude the following. Let $J \triangleq \begin{bmatrix} \mu_1 \mu_2 & 0 \\ \mu_1 - \mu_2 & \theta \end{bmatrix}$. Also, for two appropriately dimensional matrices M_1, M_2 , let $M_1 \bullet M_2$ denote the standard dot product of the two matrices.

COROLLARY 1. *– Under the same construction as Lemma 19, w.p.1, for all $s, t \geq 0$, $\text{OU}_{+,s}$ and $\zeta(\Lambda, s)$ satisfy the system of SDE*

$$\begin{bmatrix} d\text{OU}_{+,s}(t) \\ d\zeta(\Lambda, s, t) \end{bmatrix} = -J \bullet \begin{bmatrix} \text{OU}_{+,s}(t) \\ \zeta(\Lambda, s, t) \end{bmatrix} + \begin{bmatrix} d\mathcal{B}_s^3(t) \\ d\mathcal{B}_s^2(t) + Bdt \end{bmatrix}, \quad (7)$$

with appropriate boundary conditions.

Fortunately, the system of SDE (7) is well-studied and has an explicit solution, the so-called multi-dimensional OU process (see e.g. [83]). Before recalling the details of such processes, it will be helpful to review some additional matrix-related notations. Recall that for a matrix M and scalar z , zM denotes the same matrix M , but with each component multiplied by z . For a square matrix M , we let M^0 denote the identity matrix of the same dimensions, and M^k denote $M \bullet M^{k-1}$ for $k \geq 1$. For a matrix M , let $\exp(M)$ denote the standard matrix exponential, i.e. the matrix equivalent to $\sum_{k=0}^{\infty} \frac{M^k}{k!}$. Also, for an n by m matrix M and $i \in [1, n], j \in [1, m]$, let $M_{i,j}$ denote the entry in row i , column j of M .

LEMMA 20 ([83]). *Under the same construction as Lemma 19, w.p.1, for all $s, t \geq 0$, the system of SDE (7) has the following unique strong solution:*

$$\begin{aligned} \begin{bmatrix} \text{OU}_{+,s}(t) \\ \zeta(\Lambda, s, t) \end{bmatrix} &= \begin{bmatrix} 0 \\ \frac{B}{\theta}(1 - \exp(-\theta t)) \end{bmatrix} + \exp(-tJ) \bullet \begin{bmatrix} \text{OU}_{+,s}(0) \\ \zeta(\Lambda, s, 0) \end{bmatrix} \\ &+ \begin{bmatrix} \int_0^t \exp(-(t-y)J)_{1,1} d\mathcal{B}_s^3(y) + \int_0^t \exp(-(t-y)J)_{1,2} d\mathcal{B}_s^2(y) \\ \int_0^t \exp(-(t-y)J)_{2,1} d\mathcal{B}_s^3(y) + \int_0^t \exp(-(t-y)J)_{2,2} d\mathcal{B}_s^2(y) \end{bmatrix}. \end{aligned} \quad (8)$$

We now simplify (8) through a series of lemmas. First, we compute $\exp(-zJ)$ explicitly for $z \geq 0$, deferring the proof to the appendix. Here we also note that it is easily verified that our assumptions (i.e. $\mathbb{E}[S] = 1$ and Assumption 1) imply that $\mu_1\mu_2 > \theta$, as $\frac{p}{\mu_1} + \frac{1-p}{\mu_2} = 1$ implies $\mu_1\mu_2 = p\mu_2 + (1-p)\mu_1$.

LEMMA 21. * For all $z \geq 0$,

$$\exp(-zJ) = \begin{bmatrix} \exp(-\mu_1\mu_2 z) & 0 \\ \frac{\mu_1 - \mu_2}{\mu_1\mu_2 - \theta} (\exp(-\mu_1\mu_2 z) - \exp(-\theta z)) & \exp(-\theta z) \end{bmatrix}.$$

Combining Lemmas 20 and 21, we conclude the following.

LEMMA 22. *Under the same construction as Lemma 19, w.p.1, for all $0 \leq s \leq t$,*

$$\text{OU}_{+,s}(t) = \exp(-\mu_1\mu_2 t) \text{OU}_{+,s}(0) + \int_0^t \exp(-\mu_1\mu_2(t-y)) d\mathcal{B}_s^3(y), \quad (9)$$

and

$$\begin{aligned} \zeta(\Lambda, s, t-s) &= \frac{\mu_1 - \mu_2}{\mu_1\mu_2 - \theta} \left(\exp(-\mu_1\mu_2(t-s)) - \exp(-\theta(t-s)) \right) \text{OU}_{+,s}(0) \\ &+ \frac{\mu_1 - \mu_2}{\mu_1\mu_2 - \theta} \int_0^{t-s} \left(\exp(-\mu_1\mu_2(t-s-y)) - \exp(-\theta(t-s-y)) \right) d\mathcal{B}_s^3(y) \\ &+ \int_0^{t-s} \exp(-\theta(t-s-y)) d\mathcal{B}_s^2(y) + \frac{B}{\theta} \left(1 - \exp(-\theta(t-s)) \right). \end{aligned} \quad (10)$$

Next, we will combine Lemma 22 with Ito's lemma to prove the following, deferring the proof to the appendix.

LEMMA 23. * Under the same construction as Lemma 19, w.p.1, for all $0 \leq s \leq t$,

$$\begin{aligned} \zeta(\Lambda, s, t-s) &= -(\mu_1 - \mu_2) \int_s^t \exp(-\theta(t-y)) \text{OU}(y) dy \\ &\quad + \int_s^t \exp(-\theta(t-y)) d\mathcal{B}^2(y) + \frac{B}{\theta} \left(1 - \exp(-\theta(t-s)) \right). \end{aligned}$$

The desired result Lemma 18 then follows from Lemma 6, Lemma 23, the facts that (by reversibility): 1. $\text{OU}(s)_{0 \leq s \leq t}$ has the same distribution (at the process level) as $\text{OU}(t-s)_{0 \leq s \leq t}$ and 2. $\mathcal{B}^2(s)_{0 \leq s \leq t}$ has the same distribution (at the process level) as $2^{\frac{1}{2}}(\mathcal{B}^1(t) - \mathcal{B}^1(t-s))_{0 \leq s \leq t}$, and a simple reindexing, and we omit the details.

5.4. Gaussian process theory to analyze $\sup_{0 \leq s \leq t} \mathcal{G}(s)$.

In this section we use Dudley's celebrated entropy bound to prove that the expected all-time supremum of \mathcal{G} is finite. This will later allow us to relate the large deviations behavior of the all-time supremum of \mathcal{G} to the maximum variance of \mathcal{G} at any one time, which we will explicitly compute, using e.g. Borell's inequality.

We begin by recalling some well-known results from the theory of Gaussian processes, and refer the interested reader to [3, 34, 35, 64] for further details. Let us define $\bar{\mathcal{G}}$ to be the stochastic process s.t. for all $t \geq 0$,

$$\bar{\mathcal{G}}(t) = \mathcal{G}(t) - \frac{B}{\theta} \left(1 - \exp(-\theta t) \right).$$

Then one may easily verify the following.

Observation 1 . *W.p.1, $\bar{\mathcal{G}}$ is a centered, continuous Gaussian process, where here continuity is on the space \mathcal{R}^+ under the standard Euclidean metric, and $\bar{\mathcal{G}}(0) = 0$.*

For the Gaussian process $\bar{\mathcal{G}}$, we can define the so-called canonical metric (technically in general a pseudo-metric, although in our special case it will anyways be a metric) for \mathcal{R}^+ w.r.t. $\bar{\mathcal{G}}$, $d_{\bar{\mathcal{G}}}$, s.t. for $x, y \in \mathcal{R}^+$, $d_{\bar{\mathcal{G}}}(x, y) = \left(\mathbb{E} \left[(\bar{\mathcal{G}}(x) - \bar{\mathcal{G}}(y))^2 \right] \right)^{\frac{1}{2}}$. For a given $x \in \mathcal{R}^+$ and $\epsilon > 0$, let $B_{\bar{\mathcal{G}}}(x, \epsilon) \triangleq \{z \in \mathcal{R}^+ : d_{\bar{\mathcal{G}}}(z, x) \leq \epsilon\}$, i.e. the closed ball of radius ϵ centered at x under the $\bar{\mathcal{G}}$ pseudo-metric. For a given $\epsilon > 0$, let $N_{\bar{\mathcal{G}}}(\epsilon)$ denote the smallest $k \geq 1$ for which there exist $x_1, \dots, x_k \in \mathcal{R}^+$ satisfying $\mathcal{R}^+ \subseteq \bigcup_{i=1}^k B_{\bar{\mathcal{G}}}(x_i, \epsilon)$, i.e. the smallest number of closed ϵ -balls needed to cover \mathcal{R}^+ under the $d_{\bar{\mathcal{G}}}$ pseudo-metric, where we set this value to ∞ if no finite number of such balls suffices. Then the celebrated entropy bound (of Sudakov and Dudley, later generalized by Talagrand and others) for the expected value of the supremum of a centered Gaussian process (as customized to our setting) is as follows. We note that in general the relevant bounds only hold for the supremum over any fixed countable subset, and/or for an appropriate version of the associated stochastic process. However, in light of the continuity ensured by Observation 1, those results imply the following in our setting.

We also note that although \mathcal{R}^+ will not in general be compact under the metrics we consider, as we will see it will be totally bounded under the pseudo-metric $\overline{\mathcal{G}}$, which will suffice for our purposes.

LEMMA 24 ([3, 34, 35, 64]). $\mathbb{E}[\sup_{t \in \mathcal{R}^+} \overline{\mathcal{G}}(t)]$ is finite if $\int_0^\infty \log^{\frac{1}{2}}(N_{\overline{\mathcal{G}}}(\epsilon)) d\epsilon < \infty$.

We now prove that the associated entropy integral is indeed finite, i.e. we prove the following.

LEMMA 25. $\int_0^\infty \log^{\frac{1}{2}}(N_{\overline{\mathcal{G}}}(\epsilon)) d\epsilon < \infty$.

En route to proving Lemma 25, and for use in later proofs, it will be helpful to explicitly compute the covariance of $\overline{\mathcal{G}}$. This is accomplished in the following lemma, whose proof we defer to the appendix.

LEMMA 26. * For all $0 \leq s < t$, $\mathbb{E}[(\overline{\mathcal{G}}(t) - \overline{\mathcal{G}}(s))^2]$ equals

$$\begin{aligned} & \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 + \theta)} + \theta^{-1} \right) \exp(-2\theta s) - \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 - \theta)} + \theta^{-1} \right) \exp(-2\theta t) \\ & + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)(\mu_1\mu_2 + \theta)} \exp((\mu_1\mu_2 - \theta)s) \exp(-(\mu_1\mu_2 + \theta)t). \end{aligned}$$

We now establish some key monotonicity properties for the covariance of $\overline{\mathcal{G}}$, for use in later proofs, deferring all proofs to the appendix.

COROLLARY 2. * For each fixed $s \in \mathcal{R}^+$, $\mathbb{E}[(\overline{\mathcal{G}}(t) - \overline{\mathcal{G}}(s))^2]$ is a strictly increasing function (of t) on $[s, \infty)$. Also, for each fixed $t \in \mathcal{R}^+$, $\mathbb{E}[(\overline{\mathcal{G}}(t) - \overline{\mathcal{G}}(s))^2]$ is a strictly decreasing function (of s) on $[0, t)$.

Note that in light of the strictness established in Corollary 2, we find that $(\mathcal{R}^+, d_{\overline{\mathcal{G}}})$ is in fact a metric space. We also note that Lemma 26 implies that $\mathbb{E}[(\overline{\mathcal{G}}(t) - \overline{\mathcal{G}}(s))^2]$ is jointly continuous in s and t , which is a requirement for some of the results from the theory of Gaussian processes which we will be applying. We now use Corollary 2 to compute the diameter of this metric space, and explicitly bound the associated metric, deferring all proofs to the appendix.

LEMMA 27. * $\lim_{t \rightarrow \infty} \mathbb{E}[\overline{\mathcal{G}}^2(t)] = \frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)}$, where in light of Corollary 2 this further equals $\sup_{t \geq 0} \mathbb{E}[\overline{\mathcal{G}}^2(t)]$. In addition, for all $s, t \in \mathcal{R}^+$, it holds that $d_{\overline{\mathcal{G}}}(s, t) \leq \left(\frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)} \right)^{\frac{1}{2}}$, i.e. the diameter of the metric space $(\mathcal{R}^+, d_{\overline{\mathcal{G}}})$ equals $\left(\frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)} \right)^{\frac{1}{2}}$. In addition, for all $0 \leq s \leq t$,

$$d_{\overline{\mathcal{G}}}(s, t) \leq 4 \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)\theta} + \theta^{-1} \right)^{\frac{1}{2}} \left(1 - \exp(-(\mu_1\mu_2 + \theta)(t - s)) \right)^{\frac{1}{2}} \exp(-\theta s).$$

With Lemma 27 in hand, we now complete the proof of Lemma 25.

Proof of Lemma 25: Let $C_0 \triangleq 4 \left(\frac{p(1-p)(\mu_1-\mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2-\theta)} + \theta^{-1} \right)^{\frac{1}{2}} (1 + \mu_1\mu_2 + \theta)^{\frac{1}{2}}$. For any fixed $\epsilon > 0$, let $n_\epsilon \triangleq \lceil \frac{C_0^3}{\theta\epsilon^3} \rceil$. For $k \in \{0, \dots, n_\epsilon\}$, let $x_{k,\epsilon} \triangleq k \frac{\epsilon^2}{C_0^2}$. We now prove that for all $\epsilon > 0$,

$$\mathcal{R}^+ \subseteq \bigcup_{i=0}^{n_\epsilon} B_{\overline{\mathcal{G}}}(x_{i,\epsilon}, \epsilon). \quad (11)$$

We first prove that for all $t \geq x_{n_\epsilon,\epsilon}$, it holds that $t \in B_{\overline{\mathcal{G}}}(x_{n_\epsilon,\epsilon}, \epsilon)$. Indeed, first note that

$$x_{n_\epsilon,\epsilon} \geq \frac{1}{\theta} \times \frac{C_0}{\epsilon} \geq \frac{1}{\theta} \log\left(\frac{C_0}{\epsilon}\right),$$

the final inequality following from the fact that $\log(x) < x$ for all $x > 0$. It follows that for $t \geq x_{n_\epsilon,\epsilon}$, by Lemma 27,

$$d_{\overline{\mathcal{G}}}(x_{n_\epsilon,\epsilon}, t) \leq C_0 \exp(-\theta x_{n_\epsilon,\epsilon}) \leq \epsilon,$$

completing the proof.

Next, we prove that for all $k \in [0, n_\epsilon]$ and $t \in [x_{k,\epsilon}, x_{k+1,\epsilon}]$, it holds that $t \in B_{\overline{\mathcal{G}}}(x_{k,\epsilon}, \epsilon)$. Indeed, note that as $1 - \exp(-x) \leq x$ for all $x \in \mathcal{R}^+$, it follows from Lemma 27 that for all $0 \leq s \leq t$, $d_{\overline{\mathcal{G}}}(s, t) \leq C_0(t-s)^{\frac{1}{2}}$. Thus for all $k \in [0, n_\epsilon]$ and $t \in [x_{k,\epsilon}, x_{k+1,\epsilon}]$,

$$\begin{aligned} d_{\overline{\mathcal{G}}}(x_{k,\epsilon}, t) &\leq C_0(x_{k+1,\epsilon} - x_{k,\epsilon})^{\frac{1}{2}} \\ &= C_0\left(\frac{\epsilon^2}{C_0^2}\right)^{\frac{1}{2}} = \epsilon, \end{aligned}$$

completing the proof.

Combining the above completes the proof of (11), and implies that $N_{\overline{\mathcal{G}}}(\epsilon) \leq \frac{C_0^3}{\theta\epsilon^3} + 2$ for all $\epsilon > 0$. Combining with the fact that Lemma 27 implies that $N_{\overline{\mathcal{G}}}(\epsilon) = 1$ for all $\epsilon \geq \left(\frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)}\right)^{\frac{1}{2}}$, it follows that to complete the proof of Lemma 25, it suffices to demonstrate that

$$\int_0^{\left(\frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)}\right)^{\frac{1}{2}}} \log^{\frac{1}{2}} \left(\frac{C_0^3}{\theta\epsilon^3} + 2 \right) d\epsilon < \infty.$$

Combining with a straightforward exercise in calculus, the details of which we omit, then completes the proof. *Q.E.D.*

5.5. Interchange-of-limits results from [31, 24].

The final ingredient needed in our proofs is the interchange-of-limits results from [31, 24], which will allow us to translate results for $\sup_{0 \leq s \leq t} \mathcal{G}(s)$ over finite intervals $[0, t]$ to results for the stationary measure of interest. Let $\hat{Q}^n = (\hat{N}_1^n, \hat{N}_2^n, \hat{W}^n)$ denote the rescaled CTMC $n^{-\frac{1}{2}}(Q^n - (\frac{p}{\mu_1}n, \frac{1-p}{\mu_2}n, 0))$, i.e. the normalized and rescaled (by $n^{-\frac{1}{2}}$) CTMC corresponding to the $M/H_2/n + M$ queue with arrival rate $\lambda_n = n + Bn^{\frac{1}{2}}$, service distribution S which is hyper-exponentially distributed with parameters μ_1, μ_2, p , and abandonment rate θ . Furthermore, we remind the reader that $\hat{Q}^n(0)$

is equivalent in distribution to $(\text{Bin}(n, \hat{p}) - n\hat{p}, n - \text{Bin}(n, \hat{p}) - n(1 - \hat{p}), 0)$, where $\hat{p} = \frac{\rho}{\mu_1}$ and $\text{Bin}(n, \hat{p})$ is the corresponding binomially distributed r.v. Note that standard results concerning weak convergence of the binomial to the normal imply that $\{\hat{Q}^n(0), n \geq 1\}$ converges weakly to $\left((p(1-p))^{\frac{1}{2}} \mathcal{B}^0(1), -(p(1-p))^{\frac{1}{2}} \mathcal{B}^0(1), 0 \right)$, where we remind the reader that $\mathcal{B}^0, \mathcal{B}^1$ are two independent standard Brownian motions. Also, it is well-known that for each fixed n , $\{\hat{Q}^n(t), t \geq 0\}$ converges weakly to a r.v. $\hat{Q}^n(\infty) = (\hat{N}_1^n(\infty), \hat{N}_2^n(\infty), \hat{W}^n(\infty))$, corresponding to the steady-state of the corresponding (normalized) $M/H_2/n + M$ queue [24].

5.5.1. Review of weak convergence results of [29]. Before reviewing the relevant interchange-of-limits results, we will have to formalize the weak convergence results of [29] in greater depth. The relevant weak convergence result is as follows, and is proven in [29].

LEMMA 28 ([29]). *There exists a continuous Markov process $Q^\infty = (N_1^\infty, N_2^\infty, W^\infty)$ s.t. for all $T \geq 0$, $\{\hat{Q}^n(t)_{0 \leq t \leq T}, n \geq 1\}$ converges weakly to $Q^\infty(t)_{0 \leq t \leq T}$ in the space $D^3[0, T]$ under the J_1 topology. Here we note that under the stated weak convergence, $Q^\infty(0)$ is distributed as $\left((p(1-p))^{\frac{1}{2}} \mathcal{B}^0(1), -(p(1-p))^{\frac{1}{2}} \mathcal{B}^0(1), 0 \right)$.*

5.5.2. Review of ergodicity and interchange-of-limits results of [31, 24]. Given a result such as Lemma 28, there is the question of whether an interchange-of-limits hold, i.e. whether the Markov process Q^∞ has a unique invariant measure, to which $\{\hat{Q}^n(\infty), n \geq 1\}$ converges. Such an interchange is not automatic, and there is a large literature studying when such an interchange of limits holds, see the discussion in [24, 43, 60]. The sequence of papers [31, 24] resolved this question for the setting we consider here, i.e. multi-server queues in the HW regime, when services have a phase-type distribution and there is a strictly positive abandonment rate. First, in [31], the authors proved that the continuous Markov process Q^∞ is positive recurrent and has a unique invariant measure Υ^∞ , and is exponentially ergodic, where we refer the interested reader to [31] for precise details regarding the formal definition of exponential ergodicity. Let $Q^\infty(\infty) = (N_1^\infty(\infty), N_2^\infty(\infty), W^\infty(\infty))$ denote a generic r.v. distributed as Υ^∞ . Then more formally [31] proves such a result for a certain “transformed” Markov process $Y = (Y^1, Y^2)$, where (under analogous initial conditions and on an appropriate probability space) $N_1^\infty(t) = Y^1(t) - p(Y^1(t) + Y^2(t))^+$, $N_2^\infty(t) = Y^2(t) - p(Y^1(t) + Y^2(t))^+$, $W^\infty(t) = (Y^1(t) + Y^2(t))^+$ for all $t \geq 0$. Here we refer the interested reader to [29] for a further discussion of this and related transformations. We note that the results of [31] are easily seen to imply the analogous uniqueness and ergodicity results for Q^∞ . We also note that although the ergodicity results of [31] are stated only when the relevant Markov processes have deterministic initial conditions, when combined with standard results and definitions regarding ergodicity of general Markov processes [31, 84, 85, 86, 92], they are easily seen

to imply the analogous results under random initial conditions. The authors of [31] then complete the proof of the desired interchange-of-limits in [24], where we note that the result is actually proven for a certain transformation, but which is easily seen to imply the analogous results for the setting of interest. These results, including the desired interchange-of-limits, are summarized in the following lemma.

LEMMA 29 ([31, 24]). *The continuous-time Markov process Q^∞ is positive recurrent, and has a unique invariant measure Υ^∞ . Furthermore, $\{\hat{Q}^n(\infty), n \geq 1\}$ converges in distribution to $Q^\infty(\infty)$, whose associated probability measure is Υ^∞ . Namely, an interchange-of-limits holds, which is summarized as follows (after applying the Portmanteau Theorem and basic definitions and results from the theory of weak convergence). Suppose $Q^\infty(0)$ is distributed as $\left((p(1-p))^{\frac{1}{2}} \mathcal{B}^0(1), -(p(1-p))^{\frac{1}{2}} \mathcal{B}^0(1), 0 \right)$. Then for all $x > 0$ which are continuity points of $W^\infty(\infty)$ (which will be dense in \mathcal{R}^+),*

$$\begin{aligned} \liminf_{n \rightarrow \infty} \liminf_{t \rightarrow \infty} P(\hat{W}^n(t) \geq x) &= \liminf_{n \rightarrow \infty} P(\hat{W}^n(\infty) \geq x) = P(W^\infty(\infty) \geq x) \\ &= \lim_{t \rightarrow \infty} P(W^\infty(t) \geq x) \geq \limsup_{t \rightarrow \infty} \limsup_{n \rightarrow \infty} P(\hat{W}^n(t) \geq x), \end{aligned} \quad (12)$$

and

$$\begin{aligned} \limsup_{n \rightarrow \infty} \limsup_{t \rightarrow \infty} P(\hat{W}^n(t) > x) &= \limsup_{n \rightarrow \infty} P(\hat{W}^n(\infty) > x) = P(W^\infty(\infty) > x) \\ &= \liminf_{t \rightarrow \infty} P(W^\infty(t) > x) \leq \liminf_{t \rightarrow \infty} \liminf_{n \rightarrow \infty} P(\hat{W}^n(t) > x). \end{aligned} \quad (13)$$

We note that the analogous interchange also holds for a general sequence of weakly converging initial conditions (i.e. an analogous interchange-of-limits holds when $Q^\infty(0)$ has a different distribution), but for clarity of exposition we do not state at that level of generality. Also, we note that a study of whether $W^\infty(\infty)$ and other related distributions e.g. admit a density (in which case one could avoid the care taken w.r.t open and closed sets above) will be beyond the scope of our paper, and unnecessary for our results.

5.6. Proof of upper bound part of main result Theorem 1.

We now complete the proof of one direction of our main result Theorem 1. First, we briefly recall a result from the theory of Gaussian processes, namely the celebrated Borell's inequality, which relates the large deviation properties of the supremum of a centered Gaussian process, which is bounded almost surely, to its maximal variance. We only state a very special case of the result, as customized to our setting, and refer the interested reader to [3] for further details.

LEMMA 30 (**Borell's inequality** [3]). *Suppose that Z is a centered Gaussian process, and that $P(\sup_{t \geq 0} Z(t) < \infty) = 1$. Then $\lim_{x \rightarrow \infty} x^{-2} \log \left(P(\sup_{t \geq 0} Z(t) > x) \right) = -(2 \sup_{t \geq 0} \mathbb{E}[Z^2(t)])^{-1}$.*

With Lemma 30 in hand, we now complete the proof of the upper bound of our main large deviations result, as formalized in the following lemma.

LEMMA 31. *Under Assumption 1,*

$$\limsup_{x \rightarrow \infty} x^{-2} \log P(W^\infty(\infty) > x) \leq \frac{-\theta(\theta + \mu_1\mu_2)}{2(\theta + \mu_1 + \mu_2 - 1)}.$$

Proof: It follows from Lemma 29 that for all x which are continuity points of the c.d.f. of $W^\infty(\infty)$,

$$P(W^\infty(\infty) > x) \leq \liminf_{t \rightarrow \infty} \liminf_{n \rightarrow \infty} P(\hat{W}^n(t) > x). \quad (14)$$

By Theorem 2, Lemmas 13, 17, and 18, and the Portmanteau Theorem, for each such continuity point x and $t \geq 0$,

$$\liminf_{n \rightarrow \infty} P(\hat{W}^n(t) > x) \leq P\left(\sup_{0 \leq s \leq t} \mathcal{G}(s) \geq x\right).$$

Furthermore, as from definitions and a straightforward calculation w.p.1 $\sup_{t \geq 0} |\mathcal{G}(t) - \bar{\mathcal{G}}(t)| \leq \frac{|B|}{\theta}$, it follows that for all x which are continuity points of the c.d.f. of $W^\infty(\infty)$ and all $t \geq 0$,

$$\liminf_{n \rightarrow \infty} P(\hat{W}^n(t) > x) \leq P\left(\sup_{0 \leq s \leq t} \bar{\mathcal{G}}(s) > x - \frac{|B|}{\theta}\right),$$

and hence by the monotonicity of the supremum operator that

$$\liminf_{n \rightarrow \infty} P(\hat{W}^n(t) > x) \leq P\left(\sup_{s \geq 0} \bar{\mathcal{G}}(s) > x - \frac{|B|}{\theta}\right). \quad (15)$$

Combining (14) and (15), we conclude that for all such continuity points x ,

$$P(W^\infty(\infty) > x) \leq P\left(\sup_{t \geq 0} \bar{\mathcal{G}}(t) > x - \frac{|B|}{\theta}\right).$$

Combining with Lemmas 24, 25, and 30, it follows that

$$\limsup_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right) \leq -\left(2 \sup_{t \geq 0} \mathbb{E}[\bar{\mathcal{G}}^2(t)]\right)^{-1}.$$

Combining with Lemma 27, which implies that $\sup_{t \geq 0} \mathbb{E}[\bar{\mathcal{G}}^2(t)] = \frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)}$, completes the proof. *Q.E.D.*

6. Lower Bound.

With the upper bound proven, we now focus on proving lower bounds, to complete the proof of Theorem 1. At a high level, the proof of the lower bound proceeds as follows. First, we even more carefully review the limiting stochastic process Q^∞ to which $\{\hat{Q}^n, n \geq 1\}$ converges as dictated in Lemma 28, explicitly describing the system of SDE to which it is the unique strong solution. Next, we observe that so long as W^∞ has remained strictly positive on $[0, t]$, the dynamics of Q^∞ will coincide with those of a much-simpler multi-dimensional OU process, essentially the same as that arising in our previous upper bound. Finally, we combine these observations with several lower bounds and known results for such multi-dimensional OU processes to complete the proof.

6.1. Further analysis of the limit process Q^∞ .

We begin by establishing some more details about the limiting stochastic process Q^∞ to which $\{\hat{Q}^n, n \geq 1\}$ converges as dictated in Lemma 28. As mentioned previously, several of the relevant papers describe a certain transformation of this limiting process instead of the limiting process itself. For completeness, and as we will be using the precise associated system of SDE, we begin by formally defining the relevant transformed limiting process, and more formally clarifying how this limiting process relates to our own. Although originally studied in [29] in a more general form (i.e. for phase-type service distributions), a clear exposition of the case of hyper-exponential service times is given in both [27] and [31], and we now remind the reader of those results. Let $\{\kappa^i, i \in [1, 4]\}$ be four independent standard Brownian motions, constructed on a common probability space. For a general stochastic process which is denoted using a superscript, e.g. X^i , we let $X^{i,+}$ denote the stochastic process s.t. $X^{i,+}(t) = (X^i(t))^+$ for all $t \geq 0$. We now formally define the previously referenced transformed limit processes (i.e. processes which are simple transformations of Q^∞), as studied in [29, 27, 31].

LEMMA 32 ([29, 27, 31]). *The following system of SDE (stated for consistency with the relevant literature as a system of integral equations) has a unique strong solution: For all $t \geq 0$,*

$$Y^1(t) = Y^1(0) + Bpt + p\kappa^3(t) + (p(1-p))^{\frac{1}{2}}\kappa^4(t) - p^{\frac{1}{2}}\kappa^1(t) \\ - \mu_1 \int_0^t \left(Y^1(s) - p(Y^1(s) + Y^2(s))^+ \right) ds - p\theta \int_0^t (Y^1(s) + Y^2(s))^+ ds;$$

$$Y^2(t) = Y^2(0) + B(1-p)t + (1-p)\kappa^3(t) - (p(1-p))^{\frac{1}{2}}\kappa^4(t) - (1-p)^{\frac{1}{2}}\kappa^2(t) \\ - \mu_2 \int_0^t \left(Y^2(s) - (1-p)(Y^1(s) + Y^2(s))^+ \right) ds - (1-p)\theta \int_0^t (Y^1(s) + Y^2(s))^+ ds.$$

For a given measure Υ on \mathcal{R}^2 , we let $Y_\Upsilon = (Y_\Upsilon^1, Y_\Upsilon^2)$ denote the corresponding solution when $(Y^1(0), Y^2(0))$ is distributed as Υ , with $Y^1(0), Y^2(0)$ independent of $\{\kappa^i, i \in [1, 4]\}$. In addition, Y_Υ is a diffusion process, hence also Markovian with continuous sample paths.

Then the following result, proven in [29], formalizes the connection between Y_Υ and Q^∞ . Given a measure Υ on \mathcal{R}^3 , letting $U = (U_1, U_2, U_3)$ denote a r.v distributed as Υ , let Υ^2 denote the measure on \mathcal{R}^2 corresponding to the distribution of $(U_1 + pU_3^+, U_2 + (1-p)U_3^+)$. With a slight abuse of notation, given such a measure Υ on \mathcal{R}^3 , we let $Y_\Upsilon = (Y_\Upsilon^1, Y_\Upsilon^2)$ denote $Y_{\Upsilon^2} = (Y_{\Upsilon^2}^1, Y_{\Upsilon^2}^2)$. Also, for any measure Υ on \mathcal{R}^3 s.t. $\Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 < 0, w > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : w < 0\}\right) = 0$, we let $Q_\Upsilon^\infty = (N_{\Upsilon,1}^\infty, N_{\Upsilon,2}^\infty, W_\Upsilon^\infty)$ denote the corresponding Markov process when the initial conditions are distributed as Υ . Here we remind the reader that previously

we only formally defined Q^∞ for the initial conditions distributed as $\left((p(1-p))^{\frac{1}{2}}\mathcal{B}^0(1), -(p(1-p))^{\frac{1}{2}}\mathcal{B}^0(1), 0 \right)$, and refer the interested reader to [29, 31, 24] for further details regarding this Markov process. Then the connection between Y_Υ and Q_Υ^∞ is as follows, where we note that Y_Υ^i intuitively captures the (rescaled and limiting) total number of type- i customers in system (not just in service or waiting in queue).

THEOREM 4 ([29]). *For any measure Υ on \mathcal{R}^3 s.t. $\Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 < 0, w > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : w < 0\}\right) = 0$, one may construct Y_Υ and Q_Υ^∞ on a common probability space s.t. w.p.1, for all $t \geq 0$, $N_{\Upsilon,1}^\infty(t) = Y_\Upsilon^1(t) - p(Y_\Upsilon^1(t) + Y_\Upsilon^2(t))^+$; $N_{\Upsilon,2}^\infty(t) = Y_\Upsilon^2(t) - (1-p)(Y_\Upsilon^1(t) + Y_\Upsilon^2(t))^+$; $W_\Upsilon^\infty(t) = (Y_\Upsilon^1(t) + Y_\Upsilon^2(t))^+$.*

In our analysis, it will be convenient to work not with Y_Υ , but instead with a relevant transformation. Under this transformation, the first component roughly captures the infinitesimal drift due to service completions, with the second component capturing the rescaled total number in system, where this transformation also has the benefit of resulting in certain transformed terms becoming independent of one-another. In particular, for a given measure Υ on \mathcal{R}^2 , let $\bar{Y}_\Upsilon = (\bar{Y}_\Upsilon^1, \bar{Y}_\Upsilon^2)$ denote the stochastic process s.t. for all $t \geq 0$, $\bar{Y}_\Upsilon^1(t) = (1-p)Y_\Upsilon^1(t) - pY_\Upsilon^2(t)$, and $\bar{Y}_\Upsilon^2(t) = Y_\Upsilon^1(t) + Y_\Upsilon^2(t)$. We note that for such a measure Υ on \mathcal{R}^2 , $\bar{Y}_\Upsilon(0)$ is not distributed as Υ , but as $((1-p)U_1 - pU_2, U_1 + U_2)$ for a r.v. (U_1, U_2) distributed as Υ . As before, with a slight abuse of notation, given a measure Υ on \mathcal{R}^3 , we let \bar{Y}_Υ denote the analogously defined stochastic process (defined in terms of Y_Υ with $Y_\Upsilon(0)$ distributed as Υ^2).

Then we may conclude the following from Lemma 32, by taking appropriate linear combinations of the integral equations appearing in Lemma 32, and some straightforward algebra, and we omit the details.

COROLLARY 3. – *The following system of SDE has a unique strong solution: For all $t \geq 0$,*

$$\begin{aligned} \bar{Y}^1(t) = & \bar{Y}^1(0) + (p(1-p))^{\frac{1}{2}}\kappa^4(t) + \left(p(1-p)^{\frac{1}{2}}\kappa^2(t) - (1-p)p^{\frac{1}{2}}\kappa^1(t) \right) \\ & - \mu_1\mu_2 \int_0^t \bar{Y}^1(s)ds + p(1-p)\mu_1\mu_2 \int_0^t (\bar{Y}^{2,+}(s) - \bar{Y}^2(s))ds; \end{aligned}$$

$$\begin{aligned} \bar{Y}^2(t) = & \bar{Y}^2(0) + Bt + \kappa^3(t) - (p^{\frac{1}{2}}\kappa^1(t) + (1-p)^{\frac{1}{2}}\kappa^2(t)) - (\mu_1 - \mu_2) \int_0^t \bar{Y}^1(s)ds \\ & - \theta \int_0^t \bar{Y}^2(s)ds + (p\mu_1 + (1-p)\mu_2 - \theta) \int_0^t (\bar{Y}^{2,+}(s) - \bar{Y}^2(s))ds. \end{aligned}$$

For any given measure Υ on \mathcal{R}^2 , if $(\bar{Y}^1(0), \bar{Y}^2(0))$ is distributed as $((1-p)U_1 - pU_2, U_1 + U_2)$ for a r.v. (U_1, U_2) distributed as Υ , with $\bar{Y}^1(0), \bar{Y}^2(0)$ independent of $\{\kappa^i, i \in [1, 4]\}$, then this unique strong solution has the same distribution (at the process level) as \bar{Y}_Υ .

We now make an additional simplification. In particular, we note that by a simple covariance calculation, it is easily verified that the Brownian motion $p(1-p)^{\frac{1}{2}}\kappa^2 - (1-p)p^{\frac{1}{2}}\kappa^1$ is independent from the Brownian motion $p^{\frac{1}{2}}\kappa^1 + (1-p)^{\frac{1}{2}}\kappa^2$, essentially for the same reason that the sum and difference of two independent Brownian motions are independent Brownian motions. In that case, it follows from some straightforward variance calculations that Corollary 3 implies the following. Recall that \mathcal{B}^3 and \mathcal{B}^2 are two independent Brownian motions which satisfy $\text{Var}[\mathcal{B}^3(1)] = 2p(1-p)$, $\text{Var}[\mathcal{B}^2(1)] = 2$.

COROLLARY 4. *The following system of SDE has a unique strong solution: For all $t \geq 0$,*

$$\begin{aligned}\bar{Y}^1(t) &= \bar{Y}^1(0) - \mu_1\mu_2 \int_0^t \bar{Y}^1(s)ds + \mathcal{B}^3(t) + p(1-p)\mu_1\mu_2 \int_0^t (\bar{Y}^{2,+}(s) - \bar{Y}^2(s))ds; \\ \bar{Y}^2(t) &= \bar{Y}^2(0) + Bt - (\mu_1 - \mu_2) \int_0^t \bar{Y}^1(s)ds - \theta \int_0^t \bar{Y}^2(s)ds + \mathcal{B}^2(t) \\ &\quad + (p\mu_1 + (1-p)\mu_2 - \theta) \int_0^t (\bar{Y}^{2,+}(s) - \bar{Y}^2(s))ds.\end{aligned}$$

For any given measure Υ on \mathcal{R}^2 , if $(\bar{Y}^1(0), \bar{Y}^2(0))$ is distributed as $((1-p)U_1 - pU_2, U_1 + U_2)$ for a r.v. (U_1, U_2) distributed as Υ , with $\bar{Y}^1(0), \bar{Y}^2(0)$ independent of $\mathcal{B}^3, \mathcal{B}^2$, then this unique strong solution has the same distribution (at the process level) as \bar{Y}_Υ .

6.2. Relating Q^∞ to a multi-dimensional OU process.

At this point, it will be helpful to remind ourselves of an SDE we have already encountered in our upper bound arguments, namely that encountered in Lemma 20, whose solution is the multi-dimensional OU process. In particular, the following is a restatement of Lemma 20, written as an equivalent system of stochastic integral equations (as opposed to differential equations), after applying Lemma 21, where here the results are used to analyze an appropriate relaxation of \bar{Y} as opposed to $\text{OU}_{+,s}$ and $\zeta(\Lambda, s)$, and we omit the details (the logic being nearly identical to that already used in our previous analysis). For a given measure Υ on \mathcal{R}^2 , let $\hat{Y}_\Upsilon = (\hat{Y}_\Upsilon^1, \hat{Y}_\Upsilon^2)$ denote the stochastic process explicitly constructed on the same probability space as $\mathcal{B}^3, \mathcal{B}^2$ s.t. w.p.1, for all $t \geq 0$,

$$\begin{aligned}\hat{Y}_\Upsilon^1(t) &= \exp(-\mu_1\mu_2 t) \hat{Y}_\Upsilon^1(0) + \int_0^t \exp(-\mu_1\mu_2(t-y)) d\mathcal{B}^3(y); \\ \hat{Y}_\Upsilon^2(t) &= \frac{\mu_1 - \mu_2}{\mu_1\mu_2 - \theta} (\exp(-\mu_1\mu_2 t) - \exp(-\theta t)) \hat{Y}_\Upsilon^1(0) + \exp(-\theta t) \hat{Y}_\Upsilon^2(0) \\ &\quad + \frac{\mu_1 - \mu_2}{\mu_1\mu_2 - \theta} \int_0^t \left(\exp(-\mu_1\mu_2(t-y)) - \exp(-\theta(t-y)) \right) d\mathcal{B}^3(y) \\ &\quad + \int_0^t \exp(-\theta(t-y)) d\mathcal{B}^2(y) + \frac{B}{\theta} \left(1 - \exp(-\theta t) \right),\end{aligned}$$

with $(\hat{Y}_\Upsilon^1(0), \hat{Y}_\Upsilon^2(0))$ distributed as Υ , independent of $\mathcal{B}^3, \mathcal{B}^2$.

COROLLARY 5. – *The following system of SDE has a unique strong solution: For all $t \geq 0$,*

$$\hat{Y}^1(t) = \hat{Y}^1(0) - \mu_1 \mu_2 \int_0^t \hat{Y}^1(s) ds + \mathcal{B}^3(t);$$

$$\hat{Y}^2(t) = \hat{Y}^2(0) + Bt - (\mu_1 - \mu_2) \int_0^t \hat{Y}^1(s) ds - \theta \int_0^t \hat{Y}^2(s) ds + \mathcal{B}^2(t).$$

For any given measure Υ on \mathcal{R}^2 , if $(\hat{Y}^1(0), \hat{Y}^2(0))$ is distributed as Υ , with $\hat{Y}^1(0), \hat{Y}^2(0)$ independent of $\mathcal{B}^3, \mathcal{B}^2$, this unique strong solution has the same distribution (at the process level) as $(\hat{Y}_\Upsilon^1, \hat{Y}_\Upsilon^2)$.

Now, we make the simple observation that the system of SDE from Corollary 4 and that of Corollary 5 agree on any interval $[0, T]$ s.t. $\bar{Y}^{2,+}(t) = \bar{Y}^2(t)$ for all $t \in [0, T]$. Combining with Theorem 4 and Corollaries 3 - 5, and the fact that for any measure Υ on \mathcal{R}^3 s.t. $\Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 < 0, w > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : w < 0\}\right) = 0$ it holds that $(Y_\Upsilon^1 + Y_\Upsilon^2)^+, \bar{Y}_\Upsilon^{2,+}$, and W_Υ^∞ all have the same distribution (at the process level), we conclude the following Corollary 6, which relates the probability that $W_\Upsilon^\infty(t)$ exceeds a level x , conditioned on W_Υ^∞ having remained strictly positive on $[0, t]$, to the analogous probability for the second component of the much simpler multi-dimensional OU process similarly conditioned. Given a measure Υ on \mathcal{R}^3 , letting $U = (U_1, U_2, U_3)$ denote a r.v distributed as Υ , let $\Upsilon^{2'}$ denote the measure on \mathcal{R}^2 corresponding to the distribution of $((1-p)U_1 - pU_2, U_1 + U_2 + U_3^+)$. With a slight abuse of notation, given such a measure Υ on \mathcal{R}^3 , we let $\hat{Y}_\Upsilon = (\hat{Y}_\Upsilon^1, \hat{Y}_\Upsilon^2)$ denote $\hat{Y}_{\Upsilon^{2'}} = (\hat{Y}_{\Upsilon^{2'}}^1, \hat{Y}_{\Upsilon^{2'}}^2)$. Given a measure Υ on \mathcal{R}^3 s.t. $\Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 < 0, w > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : w < 0\}\right) = 0$, let $\tau_\Upsilon^{Q^\infty} \triangleq \inf\{t \geq 0 : W_\Upsilon^\infty(t) \leq 0\}$, and $\tau_\Upsilon^{\hat{Y}} \triangleq \inf\{t \geq 0 : \hat{Y}_\Upsilon^2(t) \leq 0\}$.

COROLLARY 6. – *For all measures Υ on \mathcal{R}^3 s.t. $\Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 < 0, w > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : x_1 + x_2 > 0\}\right) = \Upsilon\left(\{(x_1, x_2, w) : w < 0\}\right) = 0$, and all $t, x > 0$,*

$$\mathbb{E}[I(W_\Upsilon^\infty(t) > x, \tau_\Upsilon^{Q^\infty} > t)] = \mathbb{E}[I(\hat{Y}_\Upsilon^2(t) > x, \tau_\Upsilon^{\hat{Y}} > t)].$$

To proceed, we will reason about applying Corollary 6 when Υ is the unique invariant measure of Q^∞ , i.e. the previously defined measure Υ^∞ governing the distribution of $Q^\infty(\infty)$, whose existence was established in Lemma 29. We note that it is easily seen to follow from the results of [31] and the definition of $\bar{Y}_{\Upsilon^\infty}$ that $\{\bar{Y}_{\Upsilon^\infty}(t), t \geq 0\}$ is a stationary process. We also note that $\hat{Y}_{\Upsilon^\infty}$ is equivalent to a multi-dimensional OU process initialized not with its own stationary measure, but with that of a different process.

First, we will need to prove that Υ^∞ is sufficiently non-degenerate, and defer the proof to the appendix, noting that a very similar infinite-server lower bound is used in [43] towards different ends.

LEMMA 33. * Under Assumption 1, there exist $C_1, C_2 \in (4\frac{|B|}{\theta}, \infty)$ s.t. $C_1 < C_2$; $C_3, C_4 \in (-\infty, \infty)$ s.t. $C_3 < C_4$; and $\delta \in (0, 1)$, all depending only on $p, \mu_1, \mu_2, B, \theta$, s.t.

$$P(\hat{Y}_{\Upsilon^\infty}^1(0) \in [C_3, C_4], \hat{Y}_{\Upsilon^\infty}^2(0) \in [C_1, C_2]) \geq \delta.$$

Let Υ^* denote the probability measure governing the distribution of $(\hat{Y}_{\Upsilon^\infty}^1(0), \hat{Y}_{\Upsilon^\infty}^2(0))$ when conditioned on the event $\{\hat{Y}_{\Upsilon^\infty}^1(0) \in [C_3, C_4], \hat{Y}_{\Upsilon^\infty}^2(0) \in [C_1, C_2]\}$, for C_1, C_2, C_3, C_4 as provided by Lemma 33. Then by combining Lemma 29 and Theorem 4 with Corollaries 3 - 6 and Lemma 33, we may relate $P(W^\infty(\infty) > x)$ to the probability that the second component of a much simpler multi-dimensional OU process exceeds x , with the additional requirement that this second component does not hit 0 over an appropriate time interval.

LEMMA 34. Let $\delta \in (0, 1)$ be as in Lemma 33. Then for all $x, t \geq 0$,

$$P(W^\infty(\infty) > x) \geq \delta \times \mathbb{E} \left[I \left(\hat{Y}_{\Upsilon^*}^2(t) > x, \tau_{\Upsilon^*}^{\hat{Y}} > t \right) \right].$$

Proof: By combining Lemma 29 and Theorem 4 with Corollaries 3 - 6 and Lemma 33, we conclude that for all $x > 0$,

$$\begin{aligned} P(W^\infty(\infty) > x) &= P(W_{\Upsilon^\infty}^\infty(t) > x) \\ &= \mathbb{E} \left[I \left(\bar{Y}_{\Upsilon^\infty}^2(t) > x \right) \right] \\ &\geq \mathbb{E} \left[I \left(\bar{Y}_{\Upsilon^\infty}^2(t) > x, \bar{Y}_{\Upsilon^\infty}^1(0) \in [C_3, C_4], \bar{Y}_{\Upsilon^\infty}^2(0) \in [C_1, C_2] \right) \right] \\ &\geq \delta \mathbb{E} [I(W_{\Upsilon^*}^\infty(t) > x)] \\ &\geq \delta \mathbb{E} [I(W_{\Upsilon^*}^\infty(t) > x, \tau_{\Upsilon^*}^{Q^\infty} > t)] = \delta \mathbb{E} \left[I \left(\hat{Y}_{\Upsilon^*}^2(t) > x, \tau_{\Upsilon^*}^{\hat{Y}} > t \right) \right], \end{aligned}$$

completing the proof. *Q.E.D.*

6.3. Additional bounds and asymptotics for multi-dimensional OU processes.

Before proceeding, it will be helpful to apply arguments similar to those used in the proof of Lemma 23 to derive a simplified representation for (a process which lower bounds) $\hat{Y}_{\Upsilon^*}^2$. Intuitively, Lemma 33, especially the fact that $C_1 > 4\frac{|B|}{\theta}$, along with the special structure of the relevant multi-dimensional OU processes, will allow us to rewrite the relevant stochastic processes in a form amenable to lower-bounding the hitting times of interest by several relatively simple events occurring simultaneously, e.g. the hitting times to levels involving B, C_1, C_2, C_3, C_4 of certain Brownian motions.

We begin with some additional definitions. Let $\text{OU}_{0,a}$ denote an explicit construction of the unique w.p.1 continuous centered Gaussian process s.t. $\mathbb{E}[\text{OU}_{0,a}(s)\text{OU}_{0,a}(t)] = \frac{p(1-p)}{\mu_1\mu_2} \left(\exp(-$

$\mu_1\mu_2(t-s) - \exp(-\mu_1\mu_2(t+s))$), equivalently the appropriate non-stationary OU process initialized at 0. Similarly, let $\text{OU}_{0,b}$ denote an explicit construction of the unique w.p.1 continuous centered Gaussian process s.t. $\mathbb{E}[\text{OU}_{0,b}(s)\text{OU}_{0,b}(t)] = \theta^{-1} \left(\exp(-\theta(t-s)) - \exp(-\theta(t+s)) \right)$, again an appropriate non-stationary OU process initialized at 0. Further suppose that $\text{OU}_{0,a}$ and $\text{OU}_{0,b}$ are constructed on the same probability space, and are independent of one-another. Let us define Π to be the stochastic process s.t. for all $t \geq 0$,

$$\Pi(t) = (\mu_1 - \mu_2) \exp(-\theta t) \int_0^t \exp(\theta y) \text{OU}_{0,a}(y) dy + \text{OU}_{0,b}(t).$$

Also, for C_1, C_2, C_3, C_4 as in Lemma 33, let us define

$$\text{BUF}(t) \triangleq \exp(-\theta t) C_1 - (\mu_1 + \mu_2) (|C_3| + |C_4|) t \exp(-\theta t) - \frac{|B|}{\theta}.$$

Then the desired representation is as follows, where we defer all proofs to the appendix.

LEMMA 35. * *One can construct Π and $\hat{Y}_{\Upsilon^*}^2$ on a common probability space s.t. w.p.1, for all $t \geq 0$,*

$$\hat{Y}_{\Upsilon^*}^2(t) \geq \Pi(t) + \text{BUF}(t).$$

Note that BUF is strictly positive on some interval containing zero. This will allow us to bound the probability of certain events of interest by the probability that certain processes which start at 0 don't go very negative over a small interval containing zero, and have already raced up to a large value by the end of said interval.

Next, it will be helpful to prove some structural properties of the covariance of Π , where we again defer all proofs to the appendix.

LEMMA 36. * $\mathbb{E}[\Pi(s)\Pi(t)] \geq 0$ for $0 \leq s \leq t$, and $\lim_{t \rightarrow \infty} \mathbb{E}[\Pi^2(t)] = \frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)}$.

6.4. Proof of lower bound part of main result Theorem 1.

We now proceed as follows to complete the proof of the lower bound of our main result Theorem 1. Lemmas 34 and 35, suggest to fix a time t , and analyze the probability that the explicit process Π stays above a certain explicit boundary, related to three things: 1. being above x at time t , 2. requiring that $\tau_{\Upsilon^*}^{\hat{Y}} > t$, and 3. compensating for the shift involving BUF appearing in Lemma 35. As the asymptotic variance of Π , $\frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1\mu_2)}$, matches that appearing in our upper bound, and large deviations are not affected by any fixed shift, the main hurdle will be to show that requiring $\tau_{\Upsilon^*}^{\hat{Y}} > t$ does not somehow change the large deviations behavior of the process at time t .

6.4.1. Slepian's inequality and bounds for a finite-horizon infimum. Fortunately, the fact that Lemma 36 proves that Π has non-negative covariance will allow us to use the celebrated Slepian's inequality to this end. Let us now recall this well-known result from the theory of Gaussian processes, stating a similar variant to that used in [45] (where we note that this variant is typically referred to as Gordon's inequality).

LEMMA 37 (Slepian's/Gordon's inequality [3]). *Let $\{X_i, i = 1, \dots, m\}$ and $\{Y_i, i = 1, \dots, m\}$ be two centered m -dimensional Gaussian vectors s.t. $\mathbb{E}[X_i^2] = \mathbb{E}[Y_i^2]$ for all i , and $\mathbb{E}[X_i X_j] \leq \mathbb{E}[Y_i Y_j]$ for all $i, j \in \{1, \dots, m\}$. Then for all $z \in \mathcal{R}^m$, $P(\bigcap_{i=1, \dots, m} \{Y_i > z_i\}) \geq P(\bigcap_{i=1, \dots, m} \{X_i > z_i\})$.*

In light of the fact that Π has non-negative covariance by Lemma 36, by fixing $T \in \mathcal{R}^+$ and considering the centered Gaussian process Π' s.t. for all $0 \leq s \leq t < T$, $\mathbb{E}[\Pi'(s)\Pi'(t)] = \mathbb{E}[\Pi(s)\Pi(t)]$, while for all $0 \leq s < T$, $\mathbb{E}[\Pi'(s)\Pi'(T)] = 0$, and for which $\mathbb{E}[\Pi'^2(T)] = \mathbb{E}[\Pi^2(T)]$, i.e. the centered Gaussian process whose covariance structure on $[0, T]$ is the same as that of Π but with the process at time T independent of the process on $[0, T)$, Lemma 37 is easily seen (after combining the continuity of Π with a straightforward approximation argument, the details of which we omit) to imply the following.

COROLLARY 7. – *For all $x \in \mathcal{R}$ and $T \geq 0$,*

$$\mathbb{E}\left[I\left(\inf_{t \in [0, T]} (\Pi(t) + \text{BUF}(t)) > 0, \Pi(T) + \text{BUF}(T) > x\right)\right]$$

is at least

$$\mathbb{E}\left[I\left(\inf_{t \in [0, T]} (\Pi(t) + \text{BUF}(t)) > 0\right)\right] \times \mathbb{E}\left[I\left(\Pi(T) > x - \text{BUF}(T)\right)\right].$$

In light of Lemma 35, Corollary 7 essentially accomplishes the goal of allowing us to consider the large deviations at time T and the probability of remaining positive on $[0, T]$ separately. We now focus on analyzing the two terms appearing in Corollary 7. Our general approach to completing the proof of our main result will be to take a double limit, fixing a large value for the time T considered and then for each such large value computing the large deviations behavior (i.e. letting $x \rightarrow \infty$ for x as in Corollary 7). We note that since $\mathbb{E}\left[I\left(\inf_{t \in [0, T]} (\Pi(t) + \text{BUF}(t)) > 0\right)\right]$ does not depend on the choice of x (again as in Corollary 7), and we are only interested in the relevant large deviations behavior, it will suffice to show that for each fixed T this quantity is strictly positive. Then the desired result is the following lemma, whose proof we defer to the appendix.

LEMMA 38. * *For all $T > 0$, $\mathbb{E}\left[I\left(\inf_{t \in [0, T]} (\Pi(t) + \text{BUF}(t)) > 0\right)\right] > 0$.*

6.4.2. Proof of lower bound part of main result Theorem 1. Finally, we are in a position to complete the proof of our main result Theorem 1.

Proof of Theorem 1: In light of the already proven upper bound from Lemma 31, it suffices to prove the corresponding lower bound. By Lemma 34 and Lemma 35, there exists $\delta \in (0, 1)$ s.t. for all $T, x > 0$, $P(W^\infty(\infty) > x)$ is at least

$$\delta \times \mathbb{E} \left[I \left(\Pi(t) + \text{BUF}(t) > x, \inf_{t \in [0, T]} (\Pi(t) + \text{BUF}(t)) > 0 \right) \right], \quad (16)$$

which by Corollary 7 is at least

$$\delta \times \mathbb{E} \left[I \left(\inf_{t \in [0, T]} (\Pi(t) + \text{BUF}(t)) > 0 \right) \right] \times \mathbb{E} \left[I \left(\Pi(T) + \text{BUF}(T) > x \right) \right]. \quad (17)$$

It follows from Lemma 36 that for all $\epsilon > 0$, there exists T_ϵ s.t.

$$\mathbb{E}[\Pi^2(T_\epsilon)] > (1 - \epsilon) \frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1 \mu_2)}. \quad (18)$$

Combining (16) - (18), we conclude that for all $\epsilon > 0$, $\liminf_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right)$ is at least

$$\begin{aligned} \liminf_{x \rightarrow \infty} x^{-2} \log & \left(\delta \times \mathbb{E} \left[I \left(\inf_{t \in [0, T_\epsilon]} (\Pi(t) + \text{BUF}(t)) > 0 \right) \right] \right. \\ & \left. \times \mathbb{E} \left[I \left(\Pi(T_\epsilon) + \text{BUF}(T_\epsilon) > x \right) \right] \right), \end{aligned}$$

itself equal to

$$\begin{aligned} \liminf_{x \rightarrow \infty} x^{-2} \log & \left(\mathbb{E} \left[I \left(\Pi(T_\epsilon) + \text{BUF}(T_\epsilon) > x \right) \right] \right) \\ = \liminf_{x \rightarrow \infty} x^{-2} \log & \left(P(\Pi(T_\epsilon) > x) \right), \end{aligned}$$

which by the basic properties of the normal distribution (see [3]) equals $-(2\mathbb{E}[\Pi^2(T_\epsilon)])^{-1}$, which by (18) is at least $-\frac{1}{1-\epsilon} \frac{\theta(\theta + \mu_1 \mu_2)}{2(\theta + \mu_1 + \mu_2 - 1)}$. Taking the limit $\epsilon \downarrow 0$ then completes the proof of the desired large deviations result. We defer the proof that the proven exponent does not depend only on the first two moments of S , as well as the fact that $-\frac{\theta(\theta + \mu_1 \mu_2)}{2(\theta + \mu_1 + \mu_2 - 1)} < -\frac{\mu_1 \mu_2 \theta}{2(\mu_1 + \mu_2 - 1)}$, to the appendix. *Q.E.D.*

7. Conclusion.

In this paper, we provided the first precise large-deviations analysis of steady-state multi-server queues with abandonments in the Halfin-Whitt regime when service times are non-Markovian. In particular, for the setting in which inter-arrival times are Markovian (with spare capacity parameter

B); service times are hyper-exponentially distributed with mean 1 and parameters μ_1, μ_2, p ; and patience times are Markovian with rate θ satisfying $0 < \theta < \min(\mu_1, \mu_2)$ (i.e. Assumption 1), we proved that the weak limit W^∞ of the associated sequence of normalized steady-state queue lengths satisfies $\lim_{x \rightarrow \infty} x^{-2} \log \left(P(W^\infty(\infty) > x) \right) = -\frac{\theta(\theta + \mu_1 \mu_2)}{2(\theta + \mu_1 + \mu_2 - 1)}$. This is the first result providing such a fine-grained understanding of the behavior of multi-server queues with abandonments in the Halfin-Whitt regime, when service times are non-Markovian. Furthermore, our results disproved a conjecture of Dai and He, which had posited a particular value for this exponent which we ultimately proved to be incorrect. Our results also show that the related conjecture of Dai and He that this exponent should depend only on the first two moments of the service distribution, i.e. should exhibit an insensitivity, was similarly incorrect. This negative resolution shows that in systems with abandonments the large-deviations behavior is considerably more complex than in the setting without abandonments, as the conjectures of Dai and He were in analogy with known results for queues without abandonments, where the true answer is much simpler and indeed insensitive beyond the first two moments. Our results were also able to provide additional insight into the numerical methods of Dai and He which had inspired their conjecture, as well as how the probability of seeing a large queue length changes as e.g. the abandonment rate goes to 0, or the larger of the two service rates goes to ∞ . Perhaps most importantly, our results shed light on a connection between multi-server queues with abandonments in the Halfin-Whitt regime and certain tractable diffusion processes called multi-dimensional Ornstein-Uhlenbeck processes, which can be viewed as relaxations of the dynamics of the true complex limiting processes. Along the way we also proved several results which may be of independent interest, including a stochastic comparison result for general $M/H_2/n + M$ queues (assuming $\theta < \min(\mu_1, \mu_2)$) which bounds the true system by one with simpler dynamics, in analogy to the stochastic comparison results of [41]. We note that such a comparison may prove useful e.g. in the simulation of multi-server queues with abandonments, in analogy to the utility that the results of [41] had for devising simulation algorithms for multi-server queues without abandonments [16].

Our work leaves several interesting directions for future research. The most natural two are: 1. extending our results to $M/PH/n + M$ queues in the Halfin-Whitt regime under a hazard-rate assumption analogous to our Assumption 1, and 2. extending our results beyond the confines of Assumption 1. Although for space considerations we do not here detail precisely how one might achieve 1., i.e. extending our results to the case of phase-type services, we note that for many steps of our approach it is at least clear what one would like to prove (for phase-type services in analogy with our analysis for H_2 services), and (roughly) how one might go about it. As far as 2., a natural first approach would be to attempt to prove stochastic-comparison results directly for the diffusion limit. Indeed, a careful examination of the relevant diffusion limit \bar{Y} reveals certain

(anti-)monotonicities between various terms and their drifts, around which one may be able to build an appropriate coupling. We note that in the limit certain state-space collapse results have already kicked in [29], which may make constructing such couplings an easier task. Generalizing our pre-limit stochastic-comparison upper bound would seem to be a potentially harder task. Indeed, it is far less clear whether such a result even holds in the pre-limit beyond the confines of Assumption 1. It would also be interesting to develop a deeper understanding of our pre-limit upper bound, whose distribution is framed as the solution to an appropriate Skorokhod problem, and which we only analyze in-depth in the limit. We note that for more general service distributions, the analogy of Assumption 1 would be to require that the supremum of the hazard rate of the abandonment distribution is at most the infimum of the hazard rate of the service distribution. On that note, we also believe that our approach should generalize to different inter-arrival and abandonment distributions, and perhaps even to multi-class systems supposing that a state-space-collapse analogous to that which holds in the single-class setting also occurs in the multi-class setting.

Another interesting set of questions surround the connections between the true limit processes and the more tractable multi-dimensional OU processes brought to light by our work. How general is this phenomena, i.e. that taking a complex and non-linear system of SDE and applying an appropriate “relaxation” (linearization in our case) yields a simpler process which provably shares many of the properties of the true process? This may relate to questions surrounding precisely how different types of reflections change the qualitative properties of a stochastic process. Further understanding when the large-deviations exponent of the true process and its relaxation agree might also yield additional insights along these lines, where we note that the framework of [61] may be especially relevant. It would also be of interest to investigate connections between our own approach and the Lyapunov-function approaches of previous related work. We further note that different “relaxations” may be appropriate for answering different questions, e.g. perhaps the large deviations for the number of idle servers would be amenable to an analysis in which the linear drift appearing in the multi-dimensional OU relaxation corresponded to the drift when some servers are idle.

Taking a broader view, such a relaxation technique can be placed in the same philosophical bucket as stochastic comparison, as in both cases one is taking a complex system, making structural modifications to its dynamics to yield a simpler system, and then provably relating the two. It is natural to ask what are the fundamental limits of such an approach. By considering more sophisticated comparisons and relaxations, can we bridge the gap between the high-complexity of the true limit processes and the error of any particular bound? Is there a principled way to develop hierarchies of such relaxations, yielding tighter-and-tighter bounds at the cost of greater complexity, and formalizing this error-complexity trade-off?

Acknowledgements.

The authors gratefully acknowledge support from NSF grant no. 1333457, as well as several insightful conversations with Jim Dai, Ton Dieker, Shuangchi He, Shane Henderson, and Jamol Pender.

8. Appendix.

8.1. Proof of Lemma 5.

Proof of Lemma 5: In light of Lemma 4 and the basic properties of CTMC, it suffices to prove that for any fixed state $(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'})$, again referred to generically using “.”, s.t.: 1. $(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'})$ is 0-good, 2. $(\bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'})$ is $\bar{0}$ -good, and 3. $(\bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'}) \geq (\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'})$, the following are true. Let TRIP denote the set of all possible triplets $(\mathcal{S}, \mathcal{W}, \nu)$ which can appear as the state of any queue under our construction. First, for all triplets $(\mathcal{S}^{0''}, \mathcal{W}^{0''}, \nu^{0''})$,

$$\sum_{(\bar{\mathcal{S}}^{0''}, \bar{\mathcal{W}}^{0''}, \bar{\nu}^{0''}) \in \text{TRIP}} q' \left(\cdot, \left(\mathcal{S}^{0''}, \mathcal{W}^{0''}, \nu^{0''}, \bar{\mathcal{S}}^{0''}, \bar{\mathcal{W}}^{0''}, \bar{\nu}^{0''} \right) \right) = q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), (\mathcal{S}^{0''}, \mathcal{W}^{0''}, \nu^{0''}) \right). \quad (19)$$

Second, for all $(\bar{\mathcal{S}}^{0''}, \bar{\mathcal{W}}^{0''}, \bar{\nu}^{0''})$,

$$\sum_{(\mathcal{S}^{0''}, \mathcal{W}^{0''}, \nu^{0''}) \in \text{TRIP}} q' \left(\cdot, \left(\mathcal{S}^{0''}, \mathcal{W}^{0''}, \nu^{0''}, \bar{\mathcal{S}}^{0''}, \bar{\mathcal{W}}^{0''}, \bar{\nu}^{0''} \right) \right) = \bar{q}^0 \left((\bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'}, \bar{\nu}^{0'}), (\bar{\mathcal{S}}^{0''}, \bar{\mathcal{W}}^{0''}, \bar{\nu}^{0''}) \right). \quad (20)$$

We begin by proving (19). We proceed by verifying that (19) holds for each possible transition in q^0 , where sometimes we verify (19) by treating several cases based on the characteristics of the current state. For clarity of exposition, we always keep track of both which transition of q^0 is being addressed, as well as which transitions in q' (when summed) yield the appropriate rate within each case.

- (i): If $|\mathcal{S}^{0'}| \leq n - 1$, for $z \in \{1, 2\}$:

$$\begin{aligned} \lambda p_z &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), (\mathcal{S}^{0'} \cup (\nu^{0'} + 1, z), \mathcal{W}^{0'}, \nu^{0'} + 1) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'} \cup (\nu^{0'} + 1, z), \mathcal{W}^{0'}, \nu^{0'} + 1, \bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'} \cup (\bar{\nu}^{0'} + 1, z), \bar{\nu}^{0'} + 1 \right) \right) \quad (1). \end{aligned}$$

- (ii): If $|\mathcal{S}^{0'}| = n$, for $z \in \{1, 2\}$:

$$\begin{aligned} \lambda p_z &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), (\mathcal{S}^{0'}, \mathcal{W}^{0'} \cup (\nu^{0'} + 1, z), \nu^{0'} + 1) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'} \cup (\nu^{0'} + 1, z), \nu^{0'} + 1, \bar{\mathcal{S}}^{0'}, \bar{\mathcal{W}}^{0'} \cup (\bar{\nu}^{0'} + 1, z), \bar{\nu}^{0'} + 1 \right) \right) \quad (2). \end{aligned}$$

- (iii):

If $\mathcal{W}^{0'} = \emptyset$ and $\overline{\mathcal{W}}^{0'} = \emptyset$, for $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1$ and $N(\overline{\mathcal{S}}^{0'}, i, z) = 1$:

$$\begin{aligned} \mu_z &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), (\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}) \right) \\ &= \sum_{z_2=1}^2 q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup (0, z_2), \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right) \right) \quad (3b). \end{aligned}$$

If $\mathcal{W}^{0'} = \emptyset$ and $\overline{\mathcal{W}}^{0'} \neq \emptyset$, for $i \geq 1$, $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1$ and $N(\overline{\mathcal{S}}^{0'}, i, z) = 1$:

$$\begin{aligned} \mu_z &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), (\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) \quad (4c). \end{aligned}$$

If $\mathcal{W}^{0'} = \emptyset$ and $\overline{\mathcal{W}}^{0'} \neq \emptyset$, for $i \geq 1$, $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1$ and $N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

$$\begin{aligned} \mu_z &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), (\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) \quad (4(d)i) \\ &\quad + q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right) \right) \quad (4(d)ii). \end{aligned}$$

- (iva):

If $\mathcal{W}^{0'} \neq \emptyset$ and $\overline{\mathcal{W}}^{0'} \neq \emptyset$, for $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1$ and $N(\overline{\mathcal{S}}^{0'}, i, z) = 1$:

$$\begin{aligned} \mu_z &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'} \right) \right) \\ &= q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) \quad (5c). \end{aligned}$$

If $\mathcal{W}^{0'} \neq \emptyset$ and $\overline{\mathcal{W}}^{0'} \neq \emptyset$, for $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{S}^{0'}, i, z) = 1$ and $N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

$$\begin{aligned} \mu_z &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'} \right) \right) \\ &= q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) \quad (5(d)i) \\ &\quad + q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right) \right) \quad (5(d)ii). \end{aligned}$$

- (ivb):

If $\mathcal{W}^{0'} \neq \emptyset$ and $\overline{\mathcal{W}}^{0'} \neq \emptyset$, for $i \geq 1$ and $z \in \{1, 2\}$ s.t. $N(\mathcal{W}^{0'}, i, z) = 1$ and $N(\overline{\mathcal{W}}^{0'}, i, z) = 1$:

$$\begin{aligned} \theta &= q^0 \left((\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}), (\mathcal{S}^{0'}, \mathcal{W}^{0'} \setminus (i, z), \nu^{0'}) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'} \setminus (i, z), \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\nu}^{0'} \right) \right) \quad (5e). \end{aligned}$$

Combining with the fact that under those transitions in q' not noted above, i.e. (3a),(3c),(4a),(4b),(4e),(5a),(5b),(5f), it holds that $(\mathcal{S}^{0''}, \mathcal{W}^{0''}, \nu^{0''}) = (\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'})$ (i.e. these are self-loops w.r.t. the first 3 components), and combining with the fact that $(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'}) \geq (\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'})$, combined with (1) and (2), precludes certain configurations, completes the proof of (19).

We now provide a similar analysis for \overline{q}^0 , proving (20).

- (i):

If $|\mathcal{S}^{0'}| \leq n - 1$, for $z \in \{1, 2\}$:

$$\begin{aligned} \lambda p_z &= \overline{q}^0 \left((\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'}), (\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \cup (\overline{\nu}^{0'} + 1, z), \overline{\nu}^{0'} + 1) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'} \cup (\nu^{0'} + 1, z), \mathcal{W}^{0'}, \nu^{0'} + 1, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \cup (\overline{\nu}^{0'} + 1, z), \overline{\nu}^{0'} + 1 \right) \right) \quad (1). \end{aligned}$$

If $|\mathcal{S}^{0'}| = n$, for $z \in \{1, 2\}$:

$$\begin{aligned} \lambda p_z &= \overline{q}^0 \left((\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'}), (\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \cup (\overline{\nu}^{0'} + 1, z), \overline{\nu}^{0'} + 1) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'} \cup (\nu^{0'} + 1, z), \nu^{0'} + 1, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \cup (\overline{\nu}^{0'} + 1, z), \overline{\nu}^{0'} + 1 \right) \right) \quad (2). \end{aligned}$$

- (ia):

If $\overline{\mathcal{W}}^{0'} = \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $i \geq 1$ and $z_1, z_2 \in \{1, 2\}$ s.t. $N(\overline{\mathcal{S}}^{0'}, i, z_1) = 1$, $N(\mathcal{S}^{0'}, i, z_1) + N(\mathcal{W}^{0'}, i, z_1) = 0$:

$$\begin{aligned} \mu_{z_1} p_{z_2} &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (i, z_1)) \cup (0, z_2), \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z_1)) \cup (0, z_2), \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right) \right) \quad (3a). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} = \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $i \geq 1$ and $z_1, z_2 \in \{1, 2\}$ s.t. $N(\overline{\mathcal{S}}^{0'}, i, z_1) = 1$, $N(\mathcal{S}^{0'}, i, z_1) = 1$:

$$\begin{aligned} \mu_{z_1 p_{z_2}} &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (i, z_1)) \cup (0, z_2), \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z_1), \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z_1)) \cup (0, z_2), \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \quad (3b). \end{aligned}$$

- (iib):

If $\overline{\mathcal{W}}^{0'} = \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $z_1, z_2 \in \{1, 2\}$ s.t. $z_1 \neq z_2$:

$$\begin{aligned} N(\overline{\mathcal{S}}^{0'}, 0, z_1) \mu_{z_1 p_{z_2}} &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (0, z_1)) \cup (0, z_2), \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (0, z_1)) \cup (0, z_2), \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \quad (3c). \end{aligned}$$

- (iia):

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{S}}^{0'}, i, z) = 1$, $N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0$:

$$\begin{aligned} \mu_z &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \quad (4a). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{S}}^{0'}, i, z) = 1$, $N(\mathcal{S}^{0'}, i, z) = 1$:

$$\begin{aligned} \mu_z &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \quad (4c). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} \neq \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{S}}^{0'}, i, z) = 1$, $N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0$:

$$\begin{aligned} \mu_z &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \quad (5a). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} \neq \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{S}}^{0'}, i, z) = 1$, $N(\mathcal{S}^{0'}, i, z) = 1$:

$$\begin{aligned} \mu_z &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (i, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \quad (5c). \end{aligned}$$

- (iiib):

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{W}}^{0'}, i, z) = 1, N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0$:

$$\begin{aligned} \theta &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \quad (4b). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{W}}^{0'}, i, z) = 1, N(\mathcal{S}^{0'}, i, z) = 1$:

$$\begin{aligned} \theta &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'} \setminus (i, z), \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \quad (4(d)i). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} \neq \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{W}}^{0'}, i, z) = 1, N(\mathcal{S}^{0'}, i, z) + N(\mathcal{W}^{0'}, i, z) = 0$:

$$\begin{aligned} \theta &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \quad (5b). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} \neq \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{W}}^{0'}, i, z) = 1, N(\mathcal{S}^{0'}, i, z) = 1$:

$$\begin{aligned} \theta &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left((\mathcal{S}^{0'} \setminus (i, z)) \cup \mathcal{W}_1^{0'}, \mathcal{W}^{0'} \setminus \mathcal{W}_1^{0'}, \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \quad (5(d)i). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} \neq \emptyset$, for $i \geq 1, z \in \{1, 2\}$ s.t. $N(\overline{\mathcal{W}}^{0'}, i, z) = 1, N(\mathcal{W}^{0'}, i, z) = 1$:

$$\begin{aligned} \theta &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'} \setminus (i, z), \nu^{0'}, \overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'} \setminus (i, z), \overline{\mathcal{D}}^{0'} \right) \right) \quad (5e). \end{aligned}$$

- (iiic):

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} = \emptyset$, for $z \in \{1, 2\}$:

$$\begin{aligned} N(\overline{\mathcal{S}}^{0'}, 0, z) \mu_z &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\mathcal{D}}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (0, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \nu \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (0, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{D}}^{0'} \right) \right) \quad (4e). \end{aligned}$$

If $\overline{\mathcal{W}}^{0'} \neq \emptyset$ and $\mathcal{W}^{0'} \neq \emptyset$, for $z \in \{1, 2\}$:

$$\begin{aligned} N(\overline{\mathcal{S}}^{0'}, 0, z)\mu_z &= \overline{q}^0 \left(\left(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'} \right), \left((\overline{\mathcal{S}}^{0'} \setminus (0, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \nu \right) \right) \\ &= q' \left(\cdot, \left(\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'}, (\overline{\mathcal{S}}^{0'} \setminus (0, z)) \cup \overline{\mathcal{W}}_1^{0'}, \overline{\mathcal{W}}^{0'} \setminus \overline{\mathcal{W}}_1^{0'}, \overline{\nu}^{0'} \right) \right) \quad (5f). \end{aligned}$$

Combining with the fact that under those transitions in q' not noted above, i.e. (4(d)ii),(5(d)ii), it holds that $(\overline{\mathcal{S}}^{0''}, \overline{\mathcal{W}}^{0''}, \overline{\nu}^{0''}) = (\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'})$ (i.e. these are self-loops w.r.t. the latter 3 components), and combining with the fact that $(\overline{\mathcal{S}}^{0'}, \overline{\mathcal{W}}^{0'}, \overline{\nu}^{0'}) \geq (\mathcal{S}^{0'}, \mathcal{W}^{0'}, \nu^{0'})$, combined with (1) and (2), precludes certain configurations, completes the proof of (20). Combining the above completes the proof. *Q.E.D.*

8.2. Proof of Lemma 15.

Proof of Lemma 15: Let $\tilde{Q}^n = (\tilde{\text{RES}}_1^n, \tilde{\text{RES}}_2^n, \tilde{W}^n, \tilde{\text{ARR}}^n, \tilde{\text{REN}}^n, \tilde{\text{ABAN}}^n, \tilde{\text{DMY}}^n)$ denote the CTMC \tilde{Q} for the given n , when $\lambda = \lambda_n = n + Bn^{\frac{1}{2}}$, and $\tilde{Q}^n(0) = (\overline{N}_1^n(0), \overline{N}_2^n(0), \overline{W}^n(0), 0, 0, 0, 0)$. It follows from a proof essentially identical to that of Lemma 8, the details of which we omit, that one may construct \tilde{Q}^n and CLOCK on a common probability space s.t. w.p.1, for all $t \geq 0$,

- $\tilde{\text{ABAN}}^n(t) = \text{CLOCK}(\theta \int_0^t \tilde{W}^n(s) ds)$;
- $\tilde{W}^n(t) = \tilde{\text{ARR}}^n(t) - \tilde{\text{REN}}^n(t) - \tilde{\text{ABAN}}^n(t) + \tilde{\text{DMY}}^n(t)$;
- $\int_0^t I(\tilde{W}^n(s) > 0) d\tilde{\text{DMY}}^n(s) = 0$.

On the same probability space, let $\tilde{\text{ERR}}^n$ denote the stochastic process s.t. for all $t \geq 0$,

$$\tilde{\text{ERR}}^n(t) = n^{-\frac{1}{2}} \left(\text{CLOCK}(\theta \int_0^t \tilde{W}^n(s) ds) - \theta \int_0^t \tilde{W}^n(s) ds \right).$$

By Lemma 9, it suffices to prove the desired weak convergence for $\{\tilde{\text{ERR}}^n, n \geq 1\}$. Note that

$$\begin{aligned} |\tilde{\text{ERR}}^n(y)| &= n^{-\frac{1}{2}} \left| \theta \int_0^y \tilde{W}^n(s) ds - \text{CLOCK}(\theta \int_0^y \tilde{W}^n(s) ds) \right| \\ &\leq n^{-\frac{1}{2}} \sup_{0 \leq x \leq \int_0^y \tilde{W}^n(s) ds} |\theta x - \text{CLOCK}(\theta x)| \\ &\leq n^{-\frac{1}{2}} \sup_{0 \leq x \leq y \times \sup_{0 \leq s \leq y} \tilde{W}^n(s)} |\theta x - \text{CLOCK}(\theta x)|. \quad (21) \end{aligned}$$

For each fixed $y > 0$, Let $T(y, n) \triangleq \sup\{s \in [0, y] : \tilde{W}^n(s) = 0\}$. Recall that for a stochastic process Z , $Z(t^-)$ denotes the appropriate left limit (when it exists). Note that the definition of $T(y, n)$ and the defining properties of $\tilde{\text{DMY}}^n$ imply that $\tilde{\text{DMY}}^n(T(y, n)^-) = \tilde{\text{DMY}}^n(y)$. Also, it follows from the monotonicity of $\tilde{\text{ABAN}}^n$ that $\tilde{\text{ABAN}}^n(T(y, n)^-) \leq \tilde{\text{ABAN}}^n(y)$. Combining with Lemma 12,

we conclude that we may construct all relevant processes on a common probability space s.t. w.p.1, for all $y \geq 0$,

$$\begin{aligned}
\tilde{W}^n(y) &= \tilde{W}^n(T(y, n)^-) + \left(\tilde{W}^n(y) - \tilde{W}^n(T(y, n)^-) \right) \\
&= \left(\tilde{A}\tilde{R}\tilde{R}^n(y) - \tilde{R}\tilde{E}\tilde{N}^n(y) - \tilde{A}\tilde{B}\tilde{A}^n(y) + \tilde{D}\tilde{M}\tilde{Y}^n(y) \right) \\
&\quad - \left(\tilde{A}\tilde{R}\tilde{R}^n(T(y, n)^-) - \tilde{R}\tilde{E}\tilde{N}^n(T(y, n)^-) - \tilde{A}\tilde{B}\tilde{A}^n(T(y, n)^-) + \tilde{D}\tilde{M}\tilde{Y}^n(T(y, n)^-) \right) \\
&= \left(\tilde{A}\tilde{R}\tilde{R}^n(y) - \tilde{R}\tilde{E}\tilde{N}^n(y) - \tilde{A}\tilde{B}\tilde{A}^n(y) \right) \\
&\quad - \left(\tilde{A}\tilde{R}\tilde{R}^n(T(y, n)^-) - \tilde{R}\tilde{E}\tilde{N}^n(T(y, n)^-) - \tilde{A}\tilde{B}\tilde{A}^n(T(y, n)^-) \right) \\
&\leq \left(\tilde{A}\tilde{R}\tilde{R}^n(y) - \tilde{R}\tilde{E}\tilde{N}^n(y) \right) - \left(\tilde{A}\tilde{R}\tilde{R}^n(T(y, n)^-) - \tilde{R}\tilde{E}\tilde{N}^n(T(y, n)^-) \right) \\
&\leq \sup_{0 \leq s \leq y} \left((\tilde{A}\tilde{R}\tilde{R}^n(y) - \tilde{R}\tilde{E}\tilde{N}^n(y)) - (\tilde{A}\tilde{R}\tilde{R}^n(s) - \tilde{R}\tilde{E}\tilde{N}^n(s)) \right),
\end{aligned}$$

implying that for $0 \leq y \leq t$, $\sup_{0 \leq s \leq y} \tilde{W}^n(s)$ is at most

$$\begin{aligned}
&\sup_{0 \leq r \leq s \leq y} \left((\tilde{A}\tilde{R}\tilde{R}^n(s) - \tilde{R}\tilde{E}\tilde{N}^n(s)) - (\tilde{A}\tilde{R}\tilde{R}^n(r) - \tilde{R}\tilde{E}\tilde{N}^n(r)) \right) \\
&\leq 2 \sup_{0 \leq s \leq t} |\tilde{A}\tilde{R}\tilde{R}^n(s) - \tilde{R}\tilde{E}\tilde{N}^n(s)|.
\end{aligned}$$

Combining the above, we conclude that for all $t \geq 0$ and $y \in [0, t]$,

$$|\tilde{E}\tilde{R}\tilde{R}^n(y)| \leq n^{-\frac{1}{2}} \sup_{0 \leq x \leq 2t \times \sup_{0 \leq s \leq t} |\tilde{A}\tilde{R}\tilde{R}^n(s) - \tilde{R}\tilde{E}\tilde{N}^n(s)|} |\theta x - \text{CLOCK}(\theta x)|,$$

and hence

$$\sup_{0 \leq y \leq t} |\tilde{E}\tilde{R}\tilde{R}^n(y)| \leq n^{-\frac{1}{2}} \sup_{0 \leq x \leq 2t \times \sup_{0 \leq s \leq t} |\tilde{A}\tilde{R}\tilde{R}^n(s) - \tilde{R}\tilde{E}\tilde{N}^n(s)|} |\theta x - \text{CLOCK}(\theta x)|.$$

It then follows from a union bound that for all $T \geq 0$, $P(\sup_{0 \leq y \leq T} |\tilde{E}\tilde{R}\tilde{R}^n(y)| > n^{-\frac{1}{5}})$ is at most

$$\begin{aligned}
&P\left(2T \sup_{0 \leq s \leq T} |\tilde{A}\tilde{R}\tilde{R}^n(s) - \tilde{R}\tilde{E}\tilde{N}^n(s)| \geq n^{\frac{11}{20}}\right) \\
&\quad + P\left(n^{-\frac{1}{2}} \sup_{0 \leq x \leq n^{\frac{11}{20}}} |\theta x - \text{CLOCK}(\theta x)| \geq n^{-\frac{1}{5}}\right) \\
&= P\left(2T \sup_{0 \leq s \leq T} |n^{-\frac{1}{2}}(\tilde{A}\tilde{R}\tilde{R}^n(s) - \tilde{R}\tilde{E}\tilde{N}^n(s))| \geq n^{\frac{11}{20}}\right) \tag{22} \\
&\quad + P\left(\sup_{0 \leq x \leq n^{\frac{11}{20}}} |\theta x - \text{CLOCK}(\theta x)| \geq n^{\frac{3}{10}}\right). \tag{23}
\end{aligned}$$

It follows from Lemma 14, the continuity of the supremum and absolute value maps, and the continuous mapping theorem that for all $T \geq 0$,

$$\lim_{n \rightarrow \infty} P \left(2T \sup_{0 \leq s \leq T} |n^{-\frac{1}{2}} (\text{ARR}^n(s) - \text{REN}^n(s))| \geq n^{\frac{1}{20}} \right) = 0. \quad (24)$$

It follows from the well-known Doob's L_2 inequality for continuous time martingales, see e.g. [22], that (23) is at most $4\theta n^{\frac{11}{20}} (n^{\frac{3}{10}})^{-2} = 4\theta n^{-\frac{1}{20}}$, and hence

$$\lim_{n \rightarrow \infty} P \left(\sup_{0 \leq x \leq n^{\frac{11}{20}}} |\theta x - \text{CLOCK}(\theta x)| \geq n^{\frac{3}{10}} \right) = 0. \quad (25)$$

Combining the above, we conclude that for all $T \geq 0$, $\lim_{n \rightarrow \infty} P(\sup_{0 \leq y \leq T} |\text{ERR}^n(y)| > n^{-\frac{1}{5}}) = 0$. The desired result then follows from the basic properties and definitions associated with weak convergence. *Q.E.D.*

8.3. Proof of Lemma 17.

Proof of Lemma 17: Let $M(t)$ denote the renewal function associated with renewal distribution S evaluated at t , i.e. the expected number of renewals up to time t in the corresponding ordinary renewal process. It follows from a straightforward computation, the details of which we omit and which can be found in e.g. [72], that

$$M(t) = t + p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-2}(1 - \exp(-\mu_1\mu_2 t)). \quad (26)$$

Recall that for a generic r.v X , $\text{Var}[X]$ denotes the variance of X . It then follows from (26), the well-known representation for the variance of an equilibrium renewal process in terms of the renewal function (see e.g. [115] Theorem 7.2.4), and a straightforward calculation, that

$$\begin{aligned} \text{Var}[\text{REN}^\infty(t)] &= 2 \int_0^t (M(s) - s + \frac{1}{2}) ds \\ &= 2p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-2} \int_0^t (1 - \exp(-\mu_1\mu_2 s)) ds + t \\ &= (2p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-2} + 1)t - 2p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-3}(1 - \exp(-\mu_1\mu_2 t)) \\ &= (2p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-2} + 1)t - 2p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-3} \\ &\quad + 2p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-3} \exp(-\mu_1\mu_2 t). \end{aligned}$$

Furthermore, as REN^∞ has stationary increments (see e.g. [41]), $\mathbb{E}[(\text{REN}^\infty(t-s))^2] = \mathbb{E}[(\text{REN}^\infty(s))^2] + \mathbb{E}[(\text{REN}^\infty(t))^2] - 2\mathbb{E}[\text{REN}^\infty(s)\text{REN}^\infty(t)]$ for $0 \leq s \leq t$. We conclude from our calculation of $\text{Var}[\text{REN}^\infty(t)]$, after a straightforward calculation the details of which we omit, that for all $0 \leq s \leq t$, $\mathbb{E}[\text{REN}^\infty(s)\text{REN}^\infty(t)]$ equals

$$\begin{aligned} &(2p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-2} + 1)s - p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-3} \\ &\quad + p(1-p)(\mu_1 - \mu_2)^2(\mu_1\mu_2)^{-3} \left(\exp(-\mu_1\mu_2 s) + \exp(-\mu_1\mu_2 t) - \exp(-\mu_1\mu_2(t-s)) \right). \end{aligned}$$

Combining the above with Lemma 16, definitions, a straightforward calculation the details of which we omit, and the fact that continuous Gaussian processes are uniquely determined by their mean and covariance functions completes the proof. *Q.E.D.*

8.4. Proof of Lemma 21.

Proof of Lemma 21: We begin by proving that for all $n \geq 0$, $J^n = \begin{bmatrix} (\mu_1\mu_2)^n & 0 \\ \frac{\mu_1-\mu_2}{\mu_1\mu_2-\theta}((\mu_1\mu_2)^n - \theta^n) & \theta^n \end{bmatrix}$, and proceed by induction. The base case $n = 0$ is trivial, yielding the identity matrix. Thus, suppose the induction holds for all $n \leq N$ for some $N \geq 0$. Then

$$\begin{aligned} J^{N+1} &= J^N \bullet J \\ &= \begin{bmatrix} (\mu_1\mu_2)^N & 0 \\ \frac{\mu_1-\mu_2}{\mu_1\mu_2-\theta}((\mu_1\mu_2)^N - \theta^N) & \theta^N \end{bmatrix} \bullet \begin{bmatrix} \mu_1\mu_2 & 0 \\ \mu_1 - \mu_2 & \theta \end{bmatrix} \\ &= \begin{bmatrix} (\mu_1\mu_2)^{N+1} & 0 \\ \frac{\mu_1-\mu_2}{\mu_1\mu_2-\theta}((\mu_1\mu_2)^N - \theta^N)(\mu_1\mu_2) + \theta^N(\mu_1 - \mu_2) & \theta^{N+1} \end{bmatrix} \\ &= \begin{bmatrix} (\mu_1\mu_2)^{N+1} & 0 \\ \frac{\mu_1-\mu_2}{\mu_1\mu_2-\theta}((\mu_1\mu_2)^{N+1} - \theta^{N+1}) & \theta^{N+1} \end{bmatrix}, \end{aligned}$$

completing the proof. The desired lemma then follows from the definition of the matrix exponential and a straightforward calculation, the details of which we omit. *Q.E.D.*

8.5. Proof of Lemma 23.

Proof of Lemma 23: First, note that by a simple reindexing,

$$\int_s^t \exp(-\theta(t-y)) \text{OU}(y) dy = \exp(-\theta(t-s)) \int_0^{t-s} \exp(\theta y) \text{OU}_{+,s}(y) dy. \quad (27)$$

Applying the celebrated Ito's lemma (see [62]), we conclude that

$$\exp(\theta(t-s)) \text{OU}_{+,s}(t-s) - \text{OU}_{+,s}(0) = \int_0^{t-s} \theta \exp(\theta y) \text{OU}_{+,s}(y) dy + \int_0^{t-s} \exp(\theta y) d\text{OU}_{+,s}(y). \quad (28)$$

Combining with (7) and standard techniques from stochastic calculus, we conclude that $\int_0^{t-s} \exp(\theta y) \text{OU}_{+,s}(y) dy$ equals

$$\theta^{-1} \left(\exp(\theta(t-s)) \text{OU}_{+,s}(t-s) - \text{OU}_{+,s}(0) - \int_0^{t-s} \exp(\theta y) \left(-\mu_1\mu_2 \text{OU}_{+,s}(y) dy + d\mathcal{B}_s^3(y) \right) \right),$$

from which it follows that $\int_0^{t-s} \exp(\theta y) \text{OU}_{+,s}(y) dy$ equals

$$(\theta - \mu_1\mu_2)^{-1} \left(\exp(\theta(t-s)) \text{OU}_{+,s}(t-s) - \text{OU}_{+,s}(0) - \int_0^{t-s} \exp(\theta y) d\mathcal{B}_s^3(y) \right). \quad (29)$$

Combining with (9), we conclude that $\int_0^{t-s} \exp(\theta y) \text{OU}_{+,s}(y) dy$ equals

$$\begin{aligned} &(\theta - \mu_1\mu_2)^{-1} \exp(\theta(t-s)) \left(\exp(-\mu_1\mu_2(t-s)) \text{OU}_{+,s}(0) + \int_0^{t-s} \exp(-\mu_1\mu_2(t-s-y)) d\mathcal{B}_s^3(y) \right) \\ &- (\theta - \mu_1\mu_2)^{-1} \text{OU}_{+,s}(0) - (\theta - \mu_1\mu_2)^{-1} \int_0^{t-s} \exp(\theta y) d\mathcal{B}_s^3(y). \end{aligned}$$

Combining with (27), we find that

$$\int_s^t \exp(-\theta(t-y)) \text{OU}(y) dy = \int_0^{t-s} \exp(-\theta(t-s-y)) \text{OU}_{+,s}(y) dy,$$

itself equal to

$$\begin{aligned} & (\theta - \mu_1\mu_2)^{-1} \left(\exp(-\mu_1\mu_2(t-s)) \text{OU}_{+,s}(0) + \int_0^{t-s} \exp(-\mu_1\mu_2(t-s-y)) d\mathcal{B}_s^3(y) \right) \\ & - (\theta - \mu_1\mu_2)^{-1} \exp(-\theta(t-s)) \text{OU}_{+,s}(0) - (\theta - \mu_1\mu_2)^{-1} \int_0^{t-s} \exp(-\theta(t-s-y)) d\mathcal{B}_s^3(y), \\ = & (\theta - \mu_1\mu_2)^{-1} \left(\exp(-\mu_1\mu_2(t-s)) - \exp(-\theta(t-s)) \right) \text{OU}_{+,s}(0) \\ & + (\theta - \mu_1\mu_2)^{-1} \int_0^{t-s} \left(\exp(-\mu_1\mu_2(t-s-y)) - \exp(-\theta(t-s-y)) \right) d\mathcal{B}_s^3(y). \end{aligned}$$

Combining with (10) and another simple reindexing completes the proof. *Q.E.D.*

8.6. Proof of Lemma 26.

Proof of Lemma 26: From definitions and independence, $\mathbb{E}[(\overline{\mathcal{G}}(t) - \overline{\mathcal{G}}(s))^2]$ equals

$$(\mu_1 - \mu_2)^2 \mathbb{E} \left[\left(\int_s^t \exp(-\theta y) \text{OU}(y) dy \right)^2 \right] \quad (30)$$

$$+ 2 \mathbb{E} \left[\left(\int_s^t \exp(-\theta y) d\mathcal{B}^1(y) \right)^2 \right]. \quad (31)$$

Furthermore,

$$\begin{aligned} \mathbb{E} \left[\left(\int_s^t \exp(-\theta y) \text{OU}(y) dy \right)^2 \right] &= \mathbb{E} \left[\int_s^t \exp(-\theta y_1) \text{OU}(y_1) dy_1 \times \int_s^t \exp(-\theta y_2) \text{OU}(y_2) dy_2 \right] \\ &= \mathbb{E} \left[\int_s^t \int_s^t \exp(-\theta(y_1 + y_2)) \text{OU}(y_1) \text{OU}(y_2) dy_1 dy_2 \right]. \end{aligned}$$

It is easily verified from the basic properties of the OU-process, see e.g [33], that one may apply Fubini's theorem to the above expectation, and conclude (after a simple reindexing) that (30) equals

$$2(\mu_1 - \mu_2)^2 \int_s^t \int_s^{y_2} \exp(-\theta(y_1 + y_2)) \mathbb{E}[\text{OU}(y_1) \text{OU}(y_2)] dy_1 dy_2,$$

itself equal to (after substituting in the known covariance of OU)

$$\begin{aligned} & 2p(1-p)(\mu_1 - \mu_2)^2 (\mu_1\mu_2)^{-1} \int_s^t \int_s^{y_2} \exp(-\theta(y_1 + y_2)) \exp(-\mu_1\mu_2(y_2 - y_1)) dy_1 dy_2 \\ &= \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \int_s^t \exp(-(\mu_1\mu_2 + \theta)y_2) \left(\exp((\mu_1\mu_2 - \theta)y_2) - \exp((\mu_1\mu_2 - \theta)s) \right) dy_2 \\ &= \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \int_s^t \exp(-2\theta y_2) dy_2 \\ & - \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp((\mu_1\mu_2 - \theta)s) \int_s^t \exp(-(\mu_1\mu_2 + \theta)y_2) dy_2 \end{aligned}$$

$$\begin{aligned}
&= \frac{2p(1-p)(\mu_1 - \mu_2)^2}{2\theta\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp(-2\theta s) - \frac{2p(1-p)(\mu_1 - \mu_2)^2}{2\theta\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp(-2\theta t) \\
&\quad - \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)(\mu_1\mu_2 + \theta)} \exp(-2\theta s) \\
&\quad + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)(\mu_1\mu_2 + \theta)} \exp((\mu_1\mu_2 - \theta)s) \exp(-(\mu_1\mu_2 + \theta)t).
\end{aligned}$$

Combining with some straightforward algebra, we conclude that (30) equals

$$\begin{aligned}
&\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 + \theta)} \exp(-2\theta s) - \frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp(-2\theta t) \\
&\quad + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)(\mu_1\mu_2 + \theta)} \exp((\mu_1\mu_2 - \theta)s) \exp(-(\mu_1\mu_2 + \theta)t).
\end{aligned} \tag{32}$$

We next analyze (31). By a straightforward application of the Ito isometry, see e.g. [62], (31) equals $2 \int_s^t \exp(-2\theta y) dy = \theta^{-1} (\exp(-2\theta s) - \exp(-2\theta t))$. Combining with some straightforward algebra completes the proof. *Q.E.D.*

8.7. Proof of Corollary 2.

Proof of Corollary 2: It follows from Lemma 26 that

$$\begin{aligned}
\partial_t \mathbb{E}[(\bar{\mathcal{G}}(t) - \bar{\mathcal{G}}(s))^2] &= \left(\frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} + 2 \right) \exp(-2\theta t) \\
&\quad - \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp((\mu_1\mu_2 - \theta)s) \exp(-(\mu_1\mu_2 + \theta)t) \\
&\geq \left(\frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} + 2 \right) \exp(-2\theta t) \\
&\quad - \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp((\mu_1\mu_2 - \theta)t) \exp(-(\mu_1\mu_2 + \theta)t) \\
&= \left(\left(\frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} + 2 \right) - \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \right) \exp(-2\theta t) \\
&= 2 \exp(-2\theta t) > 0.
\end{aligned}$$

Similarly,

$$\begin{aligned}
\partial_s \mathbb{E}[(\bar{\mathcal{G}}(t) - \bar{\mathcal{G}}(s))^2] &= -2\theta \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 + \theta)} + \theta^{-1} \right) \exp(-2\theta s) \\
&\quad + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 + \theta)} \exp((\mu_1\mu_2 - \theta)s) \exp(-(\mu_1\mu_2 + \theta)t) \\
&\leq -2\theta \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 + \theta)} + \theta^{-1} \right) \exp(-2\theta s) \\
&\quad + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 + \theta)} \exp((\mu_1\mu_2 - \theta)s) \exp(-(\mu_1\mu_2 + \theta)s) \\
&= \left(-2\theta \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 + \theta)} + \theta^{-1} \right) + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 + \theta)} \right) \exp(-2\theta s) \\
&= -2 \exp(-2\theta s) < 0.
\end{aligned}$$

Combining the above completes the proof. *Q.E.D.*

8.8. Proof of Lemma 27.

Proof of Lemma 27: In light of the monotonicities established in Corollary 2, the first two parts of the lemma would follow if we could prove that $\lim_{t \rightarrow \infty} \mathbb{E}[\bar{\mathcal{G}}^2(t)] = \frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1 \mu_2)}$. A straightforward asymptotic analysis of Lemma 26 demonstrates that

$$\lim_{t \rightarrow \infty} \mathbb{E}[\bar{\mathcal{G}}^2(t)] = \frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)} + \theta^{-1}.$$

But note that since $\mathbb{E}[S] = 1$ implies that $\frac{p}{\mu_1} + \frac{1-p}{\mu_2} = 1$, and thus

$$p = \frac{\mu_1(\mu_2 - 1)}{\mu_2 - \mu_1} \quad \text{and} \quad 1 - p = \frac{\mu_2(1 - \mu_1)}{\mu_2 - \mu_1}, \quad (33)$$

it follows that

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[\bar{\mathcal{G}}^2(t)] &= \frac{\mu_1(\mu_2 - 1)\mu_2(1 - \mu_1)}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)} + \theta^{-1} \\ &= \frac{(\mu_2 - 1)(1 - \mu_1)}{\theta(\mu_1 \mu_2 + \theta)} + \frac{1}{\theta} = \frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1 \mu_2)}, \end{aligned}$$

completing the proof. For the final part of the lemma, i.e. the explicit bound on $d_{\bar{\mathcal{G}}}(s, t)$, note that by Lemma 26, for $0 \leq s \leq t$, $\mathbb{E}[(\bar{\mathcal{G}}(t) - \bar{\mathcal{G}}(s))^2]$ equals

$$\begin{aligned} &\left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)} + \theta^{-1} \right) \exp(-2\theta s) - \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 - \theta)} + \theta^{-1} \right) \exp(-2\theta s) \exp(-2\theta(t-s)) \\ &\quad + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1 \mu_2 (\mu_1 \mu_2 - \theta)(\mu_1 \mu_2 + \theta)} \exp((\mu_1 \mu_2 - \theta)s) \exp(-(\mu_1 \mu_2 + \theta)s) \exp(-(\mu_1 \mu_2 + \theta)(t-s)) \\ &= \left(\left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)} + \theta^{-1} \right) - \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 - \theta)} + \theta^{-1} \right) \exp(-2\theta(t-s)) \right) \exp(-2\theta s) \\ &\quad + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1 \mu_2 (\mu_1 \mu_2 - \theta)(\mu_1 \mu_2 + \theta)} \exp((\mu_1 \mu_2 - \theta)s) \exp(-(\mu_1 \mu_2 + \theta)s) \exp(-(\mu_1 \mu_2 + \theta)(t-s)) \\ &= \left(\left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)} + \theta^{-1} \right) - \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 - \theta)} + \theta^{-1} \right) \exp(-2\theta(t-s)) \right. \\ &\quad \left. + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1 \mu_2 (\mu_1 \mu_2 - \theta)(\mu_1 \mu_2 + \theta)} \exp(-(\mu_1 \mu_2 + \theta)(t-s)) \right) \exp(-2\theta s). \end{aligned}$$

Noting that

$$\begin{aligned} &\left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)} + \theta^{-1} \right) - \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta \mu_1 \mu_2 (\mu_1 \mu_2 - \theta)} + \theta^{-1} \right) + \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1 \mu_2 (\mu_1 \mu_2 - \theta)(\mu_1 \mu_2 + \theta)} \\ &= \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1 \mu_2} \left(\frac{1}{2\theta(\mu_1 \mu_2 + \theta)} - \frac{1}{2\theta(\mu_1 \mu_2 - \theta)} + \frac{1}{(\mu_1 \mu_2 - \theta)(\mu_1 \mu_2 + \theta)} \right) \\ &= \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1 \mu_2} \frac{(\mu_2 \mu_2 - \theta) - (\mu_2 \mu_2 + \theta) + 2\theta}{2\theta(\mu_1 \mu_2 - \theta)(\mu_1 \mu_2 + \theta)} = 0, \end{aligned}$$

we conclude that $\mathbb{E}[(\bar{\mathcal{G}}(t) - \bar{\mathcal{G}}(s))^2]$ equals

$$\begin{aligned} & \left(\left(1 - \exp(-2\theta(t-s)) \right) \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\theta\mu_1\mu_2(\mu_1\mu_2 - \theta)} + \theta^{-1} \right) \right. \\ & \quad \left. - \left(1 - \exp(-(\mu_1\mu_2 + \theta)(t-s)) \right) \frac{2p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)(\mu_1\mu_2 + \theta)} \right) \exp(-2\theta s), \end{aligned}$$

itself at most

$$8 \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)\theta} + \theta^{-1} \right) \max \left(1 - \exp(-2\theta(t-s)), 1 - \exp(-(\mu_1\mu_2 + \theta)(t-s)) \right) \exp(-2\theta s).$$

Recalling that as our assumptions imply that $\mu_1\mu_2 > \theta$, it follows that $2\theta < \mu_1\mu_2 + \theta$, and hence

$$\mathbb{E}[(\bar{\mathcal{G}}(t) - \bar{\mathcal{G}}(s))^2] \leq 8 \left(\frac{p(1-p)(\mu_1 - \mu_2)^2}{\mu_1\mu_2(\mu_1\mu_2 - \theta)\theta} + \theta^{-1} \right) \left(1 - \exp(-(\mu_1\mu_2 + \theta)(t-s)) \right) \exp(-2\theta s).$$

Combining with definitions completes the proof. *Q.E.D.*

8.9. Proof of Lemma 33.

Proof of Lemma 33: Using the results of [31, 24], it is easily verified that to prove the lemma, it suffices to demonstrate that for all $x > 0$, $P(W^\infty(\infty) > x) > 0$. To establish this fact, we will appeal to a straightforward infinite-server lower bound. Indeed, for $n \geq 1$, let \ddot{Q}^n denote the $M/G/\infty$ queue in which arrivals come as a Poisson process with rate λ_n , services are distributed as S , and the system is initially empty; and $\{\ddot{W}^n(t), t \geq 0\}$ the corresponding total number of jobs in system, with $\ddot{W}^n(\infty)$ a r.v. with the corresponding steady-state distribution (simply a Poisson r.v.). Let \dot{Q}^n denote the $M/G/n + M$ queue in which arrivals come as a Poisson process with rate λ_n , services are distributed as S , the abandonment rate is θ , and the system is initially empty; with $\{\dot{W}^n(t), t \geq 0\}$ the corresponding total number of jobs in system, and $\dot{W}^n(\infty)$ a r.v. with the corresponding steady-state distribution (see e.g. [24] for a discussion of ergodicity for $M/G/n + M$ queues). Recall that in an $M/G/n + M$ queue, the virtual waiting time of an arriving job J is the time that job would have to wait in system until it began service, if that job was infinitely patient, and is a measurable function of the inter-arrival, service, and patience times of those jobs which arrived to the system prior to J , as well as the inter-arrival time of job J . Our arguments will rely implicitly on the well-known virtual-waiting-time representation for the time-in-system of an arriving job to a $M/G/n + M$ queue, and we refer the interested reader to e.g. [106, 111] for further details. We now provide an explicit coupling / construction for \ddot{Q}^n and \dot{Q}^n , demonstrating that one may construct both queueing systems on a common probability space s.t. both systems see the same arrival stream, with common arrivals having the same type, in such a way that w.p.1 every job J departs later (i.e. no sooner) from \dot{Q}^n than from \ddot{Q}^n . Intuitively, this follows from the simple fact that (supposing the arrival processes are coupled in the two systems), the infinitesimal

rate at which a type- z job J leaves the system (either due to abandonment or service completion) in \dot{Q}^n is either θ or μ_z , while in \ddot{Q}^n this rate always equals μ_z , in which case the desired dominance follows from Assumption 1.

More formally, let $P_{1,\lambda}$ be a rate- $(p\lambda_n)$ Poisson process, and $P_{2,\lambda}$ be a rate- $((1-p)\lambda_n)$ Poisson process. For $z \in \{1, 2\}$ and $k \geq 1$, let $P_{z,k}$ be a rate μ_z Poisson process. For $z \in \{1, 2\}$ and $k \geq 1$, let $B_{z,k}$ be a geometrically distributed r.v., with support on the strictly positive integers, s.t. $P(B_{z,k} = 1) = \frac{\theta}{\mu_z}$; $E_{z,k}$ be an exponentially distributed r.v. with rate θ ; and $S_{z,k}$ be an exponentially distributed r.v. with rate μ_z . Suppose that all these processes and r.v are constructed together on a common probability space s.t. they are mutually independent. We proceed to generate the inter-arrival, service, and patience times for the arrivals to both systems as follows. For $z \in \{1, 2\}$, both queueing systems witness a type- z arrival at each jump in $P_{z,\lambda}$, and see no other arrivals. We determine the service and patience times of the arrivals to both systems inductively.

For an arriving job J , let us denote the patience time of J in \dot{Q}^n by $\mathring{\text{P}}\dot{\text{A}}\text{T}_J$, service time in \dot{Q}^n by \mathring{S}_J , and service time in \ddot{Q}^n by $\mathring{\ddot{S}}_J$. Also, let us denote the virtual waiting time (in \dot{Q}^n) of job J by $\mathring{\text{V}}\dot{\text{R}}\text{W}_J$ (our inductive definition will ensure this quantity is well-defined, i.e. not circular). In addition, for a job J , let z_J denote the type of J . Also, if job J is the k th type- z_J arrival to the system, we set $k_J = k$. Thus (for example) if job J is type 1 and there have been exactly 3 previous type-1 arrivals to the system, we set $z_J = 1$ and $k_J = 4$. For the job J which is the first arrival to the system, note that $\mathring{\text{V}}\dot{\text{R}}\text{W}_J = 0$, as both systems are initially empty, and $k_J = 1$. Set both \mathring{S}_J and $\mathring{\ddot{S}}_J$ equal to the time of the first jump in P_{z_J, k_J} , and set $\mathring{\text{P}}\dot{\text{A}}\text{T}_J$ equal to E_{z_J, k_J} . Now, suppose for induction that we have constructed the patience and service times for the first i arrivals to the system. Let job J be the $(i+1)$ st arrival to the system. By our description of the arrival processes, our inductive hypothesis, and the measurability properties of the virtual waiting time, $\mathring{\text{V}}\dot{\text{R}}\text{W}_J$ has thus already been defined (implicitly by the basic definitions associated with the relevant queueing systems). Set $\mathring{\ddot{S}}_J$ equal to the time of the first jump in P_{z_J, k_J} . If the time of the B_{z_J, k_J} -th jump in P_{z_J, k_J} occurs before time $\mathring{\ddot{W}}_J$, set $\mathring{\text{P}}\dot{\text{A}}\text{T}_J$ equal to the time of this jump, and set \mathring{S}_J equal to S_{z_J, k_J} . Otherwise, set $\mathring{\text{P}}\dot{\text{A}}\text{T}_J$ equal to $\mathring{\text{V}}\dot{\text{R}}\text{W}_J + E_{z_J, k_J}$, and set \mathring{S}_J equal to the difference between the time of the first jump in P_{z_J, k_J} which occurs after time $\mathring{\text{V}}\dot{\text{R}}\text{W}_J$, and $\mathring{\text{V}}\dot{\text{R}}\text{W}_J$ itself.

That this construction is correct, i.e. yielding sequences of inter-arrival, patience, and service times with the correct joint distributions, is straightforward and follows from the basic properties of Poisson processes and the relevant queueing systems, and we omit the details. It follows that we may construct both systems on a common probability space s.t. every job arrives to both systems at the same time, but departs later in \dot{Q}^n . The well-known ergodicity properties of both systems then imply that for all $x \geq 0$, $P(\mathring{\dot{W}}^n(\infty) > x) \geq P(\mathring{\ddot{W}}^n(\infty) > x)$. Rescaling both systems, and applying the well-known convergence of the appropriately scaled sequence of Poisson r.v.s (i.e.

$\{\tilde{W}^n(\infty), n \geq 1\}$) to a limiting normal distribution then completes the proof, and we refer the interested reader to [41] for further details regarding the relevant limiting arguments. *Q.E.D.*

8.10. Proof of Lemma 35.

Proof of Lemma 35: It follows from an argument nearly identical to that used in the proof of Lemma 23, the details of which we omit, that $\int_0^t \exp(-\theta(t-y)) \hat{Y}_{\Upsilon^*}^1(y) dy$ equals

$$\begin{aligned} & (\theta - \mu_1\mu_2)^{-1} \left(\exp(-\mu_1\mu_2 t) - \exp(-\theta t) \right) \hat{Y}_{\Upsilon^*}^1(0) \\ & + (\theta - \mu_1\mu_2)^{-1} \int_0^t \left(\exp(-\mu_1\mu_2(t-y)) - \exp(-\theta(t-y)) \right) d\mathcal{B}^3(y). \end{aligned}$$

Combining with definitions, we conclude that w.p.1, for all $t \geq 0$,

$$\begin{aligned} \hat{Y}_{\Upsilon^*}^2(t) &= -(\mu_1 - \mu_2) \int_0^t \exp(-\theta(t-y)) \hat{Y}_{\Upsilon^*}^1(y) dy + \exp(-\theta t) \hat{Y}_{\Upsilon^*}^2(0) \\ &+ \int_0^t \exp(-\theta(t-y)) d\mathcal{B}^2(y) + \frac{B}{\theta} \left(1 - \exp(-\theta t) \right), \end{aligned}$$

from which it follows (by Lemma 34) that

$$\hat{Y}_{\Upsilon^*}^2(t) \geq -(\mu_1 - \mu_2) \int_0^t \exp(-\theta(t-y)) \hat{Y}_{\Upsilon^*}^1(y) dy + \int_0^t \exp(-\theta(t-y)) d\mathcal{B}^2(y) + \exp(-\theta t) C_1 - \frac{|B|}{\theta}.$$

Noting from definitions that w.p.1, for all $y \geq 0$,

$$-(\mu_1 - \mu_2) \hat{Y}_{\Upsilon^*}^1(y) \geq -(\mu_1 + \mu_2) (|C_3| + |C_4|) \exp(-\mu_1\mu_2 y) - (\mu_1 - \mu_2) \int_0^y \exp(-\mu_1\mu_2(y-s)) d\mathcal{B}^3(s),$$

we further conclude after some straightforward algebra that w.p.1, for all $t \geq 0$, $\hat{Y}_{\Upsilon^*}^2(t)$ is at least

$$\begin{aligned} & -(\mu_1 - \mu_2) \int_0^t \exp(-\theta(t-y)) \int_0^y \exp(-\mu_1\mu_2(y-s)) d\mathcal{B}^3(s) \\ & + \int_0^t \exp(-\theta(t-y)) d\mathcal{B}^2(y) + \exp(-\theta t) C_1 - (\mu_1 + \mu_2) (|C_3| + |C_4|) t \exp(-\theta t) - \frac{|B|}{\theta}. \end{aligned}$$

The lemma then follows from the well-known stochastic integral representation for the OU process, and the symmetries of the centered OU process, and we refer the reader to [73] for details. *Q.E.D.*

8.11. Proof of Lemma 36.

Proof of Lemma 36: We first compute $\mathbb{E}[\Pi(s)\Pi(t)]$. Let us define

$$C(s, t) \triangleq \mathbb{E} \left[\int_0^t \exp(\theta y_2) \text{OU}_{0,a}(y_2) dy_2 \times \int_0^s \exp(\theta y_1) \text{OU}_{0,a}(y_1) dy_1 \right].$$

First, note that by applying Fubini's theorem as in the proof of Lemma 26, and substituting in the known covariance of the OU process, we conclude that $C(s, t)$ equals

$$\mathbb{E} \left[\int_0^t \int_0^s \exp(\theta(y_1 + y_2)) \text{OU}_{0,a}(y_1) \text{OU}_{0,a}(y_2) dy_1 dy_2 \right]$$

$$\begin{aligned}
&= \int_0^t \int_0^s \exp(\theta(y_1 + y_2)) \mathbb{E}[\text{OU}_{0,a}(y_1)\text{OU}_{0,a}(y_2)] dy_1 dy_2 \\
&= 2 \int_0^s \int_0^{y_2} \exp(\theta(y_1 + y_2)) \mathbb{E}[\text{OU}_{0,a}(y_1)\text{OU}_{0,a}(y_2)] dy_1 dy_2 \tag{34}
\end{aligned}$$

$$+ \int_s^t \int_0^s \exp(\theta(y_1 + y_2)) \mathbb{E}[\text{OU}_{0,a}(y_1)\text{OU}_{0,a}(y_2)] dy_1 dy_2. \tag{35}$$

(34) equals

$$\begin{aligned}
&2 \frac{p(1-p)}{\mu_1 \mu_2} \int_0^s \int_0^{y_2} \exp(\theta(y_1 + y_2)) \left(\exp(-\mu_1 \mu_2(y_2 - y_1)) - \exp(-\mu_1 \mu_2(y_2 + y_1)) \right) dy_1 dy_2 \\
&= 2 \frac{p(1-p)}{\mu_1 \mu_2} \int_0^s \int_0^{y_2} \exp(\theta(y_1 + y_2)) \exp(-\mu_1 \mu_2(y_2 - y_1)) dy_1 dy_2 \\
&\quad - 2 \frac{p(1-p)}{\mu_1 \mu_2} \int_0^s \int_0^{y_2} \exp(\theta(y_1 + y_2)) \exp(-\mu_1 \mu_2(y_2 + y_1)) dy_1 dy_2 \\
&= 2 \frac{p(1-p)}{\mu_1 \mu_2} \int_0^s \exp((\theta - \mu_1 \mu_2)y_2) \int_0^{y_2} \exp((\theta + \mu_1 \mu_2)y_1) dy_1 dy_2 \\
&\quad - 2 \frac{p(1-p)}{\mu_1 \mu_2} \int_0^s \exp((\theta - \mu_1 \mu_2)y_2) \int_0^{y_2} \exp((\theta - \mu_1 \mu_2)y_1) dy_1 dy_2 \\
&= 2 \frac{p(1-p)}{\mu_1 \mu_2(\theta + \mu_1 \mu_2)} \int_0^s \exp((\theta - \mu_1 \mu_2)y_2) \left(\exp((\theta + \mu_1 \mu_2)y_2) - 1 \right) dy_2 \\
&\quad - 2 \frac{p(1-p)}{\mu_1 \mu_2(\mu_1 \mu_2 - \theta)} \int_0^s \exp((\theta - \mu_1 \mu_2)y_2) \left(1 - \exp(-(\mu_1 \mu_2 - \theta)y_2) \right) dy_2 \\
&= 2 \frac{p(1-p)}{\mu_1 \mu_2(\theta + \mu_1 \mu_2)} \int_0^s \exp(2\theta y_2) dy_2 \\
&\quad - 2 \frac{p(1-p)}{\mu_1 \mu_2(\theta + \mu_1 \mu_2)} \int_0^s \exp(-(\mu_1 \mu_2 - \theta)y_2) dy_2 \\
&\quad - 2 \frac{p(1-p)}{\mu_1 \mu_2(\mu_1 \mu_2 - \theta)} \int_0^s \exp(-(\mu_1 \mu_2 - \theta)y_2) dy_2 \\
&\quad + 2 \frac{p(1-p)}{\mu_1 \mu_2(\mu_1 \mu_2 - \theta)} \int_0^s \exp(-2(\mu_1 \mu_2 - \theta)y_2) dy_2 \\
&= \frac{p(1-p)}{\mu_1 \mu_2(\theta + \mu_1 \mu_2)\theta} (\exp(2\theta s) - 1) \\
&\quad - 2 \frac{p(1-p)}{\mu_1 \mu_2(\theta + \mu_1 \mu_2)(\mu_1 \mu_2 - \theta)} \left(1 - \exp(-(\mu_1 \mu_2 - \theta)s) \right) \\
&\quad - 2 \frac{p(1-p)}{\mu_1 \mu_2(\mu_1 \mu_2 - \theta)^2} \left(1 - \exp(-(\mu_1 \mu_2 - \theta)s) \right) \\
&\quad + \frac{p(1-p)}{\mu_1 \mu_2(\mu_1 \mu_2 - \theta)^2} \left(1 - \exp(-2(\mu_1 \mu_2 - \theta)s) \right) \\
&= \frac{p(1-p)}{\mu_1 \mu_2(\mu_1 \mu_2 + \theta)\theta} \exp(2\theta s) \\
&\quad + \frac{4p(1-p)}{(\mu_1 \mu_2 + \theta)(\mu_1 \mu_2 - \theta)^2} \exp(-(\mu_1 \mu_2 - \theta)s) \\
&\quad - \frac{p(1-p)}{\mu_1 \mu_2(\mu_1 \mu_2 - \theta)^2} \exp(-2(\mu_1 \mu_2 - \theta)s)
\end{aligned}$$

$$- \frac{p(1-p)}{\theta(\mu_1\mu_2 - \theta)^2}.$$

Similarly, (35) equals

$$\begin{aligned}
& \frac{p(1-p)}{\mu_1\mu_2} \int_s^t \int_0^s \exp(\theta(y_1 + y_2)) \left(\exp(-\mu_1\mu_2(y_2 - y_1)) - \exp(-\mu_1\mu_2(y_2 + y_1)) \right) dy_1 dy_2 \\
&= \frac{p(1-p)}{\mu_1\mu_2} \int_s^t \int_0^s \exp(\theta(y_1 + y_2)) \exp(-\mu_1\mu_2(y_2 - y_1)) dy_1 dy_2 \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2} \int_s^t \int_0^s \exp(\theta(y_1 + y_2)) \exp(-\mu_1\mu_2(y_2 + y_1)) dy_1 dy_2 \\
&= \frac{p(1-p)}{\mu_1\mu_2} \int_s^t \exp((\theta - \mu_1\mu_2)y_2) \int_0^s \exp((\theta + \mu_1\mu_2)y_1) dy_1 dy_2 \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2} \int_s^t \exp((\theta - \mu_1\mu_2)y_2) \int_0^s \exp((\theta - \mu_1\mu_2)y_1) dy_1 dy_2 \\
&= \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)} \int_s^t \exp((\theta - \mu_1\mu_2)y_2) \left(\exp((\mu_1\mu_2 + \theta)s) - 1 \right) dy_2 \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \int_s^t \exp((\theta - \mu_1\mu_2)y_2) \left(1 - \exp(-(\mu_1\mu_2 - \theta)s) \right) dy_2 \\
&= \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} \left(\exp((\mu_1\mu_2 + \theta)s) - 1 \right) \left(\exp(-(\mu_1\mu_2 - \theta)s) - \exp(-(\mu_1\mu_2 - \theta)t) \right) \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)^2} \left(1 - \exp(-(\mu_1\mu_2 - \theta)s) \right) \left(\exp(-(\mu_1\mu_2 - \theta)s) - \exp(-(\mu_1\mu_2 - \theta)t) \right) \\
&= \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} \exp(2\theta s) \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} \exp((\mu_1\mu_2 + \theta)s - (\mu_1\mu_2 - \theta)t) \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} \exp(-(\mu_1\mu_2 - \theta)s) \\
&\quad + \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} \exp(-(\mu_1\mu_2 - \theta)t) \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)^2} \exp(-(\mu_1\mu_2 - \theta)s) \\
&\quad + \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)^2} \exp(-(\mu_1\mu_2 - \theta)t) \\
&\quad + \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)^2} \exp(-2(\mu_1\mu_2 - \theta)s) \\
&\quad - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)^2} \exp(-(\mu_1\mu_2 - \theta)(s + t)).
\end{aligned}$$

We conclude, after some straightforward algebra, that $C(s, t)$ equals

$$\begin{aligned}
& \frac{p(1-p)}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)\theta} \exp(2\theta s) \\
& + \frac{2p(1-p)}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)^2} \exp(-(\mu_1\mu_2 - \theta)s)
\end{aligned} \tag{36}$$

$$\begin{aligned}
& - \frac{p(1-p)}{\theta(\mu_1\mu_2 - \theta)^2} \\
& - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} \exp((\mu_1\mu_2 + \theta)s - (\mu_1\mu_2 - \theta)t) \\
& + \frac{2p(1-p)}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)^2} \exp(-(\mu_1\mu_2 - \theta)t) \\
& - \frac{p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)^2} \exp(-(\mu_1\mu_2 - \theta)(s+t)).
\end{aligned}$$

It then follows from definitions and the known covariance of $\text{OU}_{0,b}$, and the independence of $\text{OU}_{0,a}$ and $\text{OU}_{0,b}$, that for all $0 \leq s \leq t$, $\mathbb{E}[\Pi(s)\Pi(t)]$ equals

$$\begin{aligned}
& \left(\frac{(\mu_1 - \mu_2)^2 p(1-p)}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)\theta} + \theta^{-1} \right) \exp(-\theta(t-s)) \\
& + \frac{(\mu_1 - \mu_2)^2 2p(1-p)}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)^2} \exp(-\mu_1\mu_2 s - \theta t) \\
& - \left(\frac{(\mu_1 - \mu_2)^2 p(1-p)}{\theta(\mu_1\mu_2 - \theta)^2} + \theta^{-1} \right) \exp(-\theta(s+t)) \\
& - \frac{(\mu_1 - \mu_2)^2 p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} \exp(-\mu_1\mu_2(t-s)) \\
& + \frac{(\mu_1 - \mu_2)^2 2p(1-p)}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)^2} \exp(-\theta s - \mu_1\mu_2 t) \\
& - \frac{(\mu_1 - \mu_2)^2 p(1-p)}{\mu_1\mu_2(\mu_1\mu_2 - \theta)^2} \exp(-\mu_1\mu_2(s+t)).
\end{aligned} \tag{37}$$

Next, let us prove that $\mathbb{E}[\Pi(s)\Pi(t)] \geq 0$. In light of the independence of $\text{OU}_{0,a}$ and $\text{OU}_{0,b}$ and known non-negative covariance of $\text{OU}_{0,b}$, it suffices to prove that $C(s, t)$ is non-negative. It follows from definitions that $C(s, s) \geq 0$ for all $s \geq 0$. Thus it suffices to prove that for all $s \geq 0$, $\frac{d}{dt}C(s, t) \geq 0$ on $[s, \infty)$. Thus let us fix an $s \geq 0$. Then from (36), we compute that $\frac{d}{dt}C(s, t)$ equals

$$\begin{aligned}
& -(\mu_1\mu_2 - \theta) \frac{p(1-p)}{\mu_1\mu_2 - \theta} \left(\frac{2}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} - \frac{1}{\mu_1\mu_2(\mu_1\mu_2 + \theta)} \exp((\mu_1\mu_2 + \theta)s) \right. \\
& \quad \left. - \frac{1}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp(-(\mu_1\mu_2 - \theta)s) \right) \exp(-(\mu_1\mu_2 - \theta)t).
\end{aligned}$$

Thus it suffices to prove that

$$\ddot{g}(s) \triangleq \frac{2}{(\mu_1\mu_2 + \theta)(\mu_1\mu_2 - \theta)} - \frac{1}{\mu_1\mu_2(\mu_1\mu_2 + \theta)} \exp((\mu_1\mu_2 + \theta)s) - \frac{1}{\mu_1\mu_2(\mu_1\mu_2 - \theta)} \exp(-(\mu_1\mu_2 - \theta)s)$$

is non-positive for all $s \geq 0$. Noting that a straightforward calculation shows that $\ddot{g}(0) = 0$, we now accomplish this by proving that $\frac{d}{ds}\ddot{g}(s) \leq 0$ for all $s \geq 0$. Indeed, a straightforward calculation shows that

$$\frac{d}{ds}\ddot{g}(s) = (\mu_1\mu_2)^{-1} \left(\exp(-(\mu_1\mu_2 - \theta)s) - \exp((\mu_1\mu_2 + \theta)s) \right),$$

which is non-positive since the fact that $\mu_1\mu_2 > \theta$ implies that $\exp(-(\mu_1\mu_2 - \theta)s) \leq \exp((\mu_1\mu_2 + \theta)s)$ for all $s \geq 0$. Combining the above then completes the proof that $\mathbb{E}[\Pi(s)\Pi(t)] \geq 0$ for all

$0 \leq s \leq t$.

Finally, let us compute $\lim_{t \rightarrow \infty} \mathbb{E}[\Pi^2(t)]$. It follows from (37), a straightforward asymptotic analysis, and (33) that the desired limit equals

$$\begin{aligned}
& \left(\frac{(\mu_1 - \mu_2)^2 p(1-p)}{(\mu_1 \mu_2 + \theta)(\mu_1 \mu_2 - \theta)\theta} + \theta^{-1} \right) - \frac{(\mu_1 - \mu_2)^2 p(1-p)}{\mu_1 \mu_2 (\mu_1 \mu_2 + \theta)(\mu_1 \mu_2 - \theta)} \\
&= \frac{(\mu_1 - \mu_2)^2 p(1-p)}{\mu_1 \mu_2 (\mu_1 \mu_2 + \theta)(\mu_1 \mu_2 - \theta)\theta} (\mu_1 \mu_2 - \theta) + \frac{1}{\theta} \\
&= \frac{(\mu_1 - \mu_2)^2 p(1-p) + \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)}{\mu_1 \mu_2 (\mu_1 \mu_2 + \theta)\theta} \\
&= \frac{\mu_1(\mu_2 - 1)\mu_2(1 - \mu_1) + \mu_1 \mu_2 (\mu_1 \mu_2 + \theta)}{\mu_1 \mu_2 (\mu_1 \mu_2 + \theta)\theta} \\
&= \frac{\theta + \mu_1 + \mu_2 - 1}{\theta(\theta + \mu_1 \mu_2)},
\end{aligned}$$

completing the proof. *Q.E.D.*

8.12. Proof of Lemma 38.

Proof of Lemma 38: Let $t_1 \triangleq \frac{1}{4\theta} \min\left(\frac{|B|}{(\mu_1 + \mu_2)(|C_3| + |C_4|)}, 1\right)$. Since $C_1 > 4\frac{|B|}{\theta}$, and $\exp(-\theta t) > 1 - \theta t$, it follows that for all $t \in [0, t_1]$,

$$\begin{aligned}
\exp(-\theta t)C_1 &\geq (1 - \theta t)\frac{4|B|}{\theta} \\
&\geq \left(1 - \theta \times \frac{1}{4\theta}\right) \times \frac{4|B|}{\theta} \geq 3\frac{|B|}{\theta};
\end{aligned}$$

and

$$\begin{aligned}
(\mu_1 + \mu_2)(|C_3| + |C_4|)t \exp(-\theta t) &\leq (\mu_1 + \mu_2)(|C_3| + |C_4|) \times \frac{|B|}{4\theta(\mu_1 + \mu_2)(|C_3| + |C_4|)} \\
&< \frac{|B|}{\theta};
\end{aligned}$$

from which we conclude that for all $t \in [0, t_1]$,

$$\text{BUF}(t) > \frac{|B|}{\theta}. \quad (38)$$

It is also easily verified from the definition of BUF and the fact that $C_1 > 0$ and $t \exp(-\theta t) \leq \frac{1}{\theta}$ for all $t \geq 0$ that

$$\inf_{t \geq 0} \text{BUF}(t) \geq -\frac{(\mu_1 + \mu_2)(|C_3| + |C_4|) + |B|}{\theta}. \quad (39)$$

Recall that κ^1, κ^2 are two independent standard Brownian motions. It follows from the well-known representation of a non-stationary OU process as a time-changed Brownian motion [73] that we may construct $\text{OU}_{0,a}, \text{OU}_{0,b}, \kappa^1, \kappa^2, \Pi$ on a common probability space s.t. w.p.1, for all $t \geq 0$,

$$\text{OU}_{0,a}(t) = \left(\frac{p(1-p)}{\mu_1 \mu_2}\right)^{\frac{1}{2}} \exp(-\mu_1 \mu_2 t) \kappa^1(\exp(2\mu_1 \mu_2 t) - 1);$$

$$\text{OU}_{0,b}(t) = \theta^{-\frac{1}{2}} \exp(-\theta t) \kappa^2(\exp(2\theta t) - 1);$$

and $\Pi(t)$ equals

$$\begin{aligned} & (\mu_1 - \mu_2) \left(\frac{p(1-p)}{\mu_1\mu_2} \right)^{\frac{1}{2}} \exp(-(\mu_1\mu_2 + \theta)t) \int_0^t \exp(\theta y) \kappa^1(\exp(2\mu_1\mu_2 y) - 1) dy \\ & + \theta^{-\frac{1}{2}} \exp(-\theta t) \kappa^2(\exp(2\theta t) - 1). \end{aligned} \quad (40)$$

W.l.o.g, suppose $\mu_1 > \mu_2$ (if not an identical argument works but with the roles of μ_1, μ_2 interchanged). It follows from (40) that for $t \in [0, t_1]$,

$$\begin{aligned} \Pi(t) & \geq \left(\inf_{y \in [0, t_1]} \kappa^1(\exp(2\mu_1\mu_2 y) - 1) \right) (\mu_1 - \mu_2) \left(\frac{p(1-p)}{\mu_1\mu_2} \right)^{\frac{1}{2}} \exp(-(\mu_1\mu_2 + \theta)t) \int_0^t \exp(\theta y) dy \\ & + \theta^{-\frac{1}{2}} \exp(-\theta t) \left(\inf_{y \in [0, t_1]} \kappa^2(\exp(2\theta y) - 1) \right), \end{aligned}$$

itself at least

$$\mathcal{H}_1 \triangleq \left(1 + (\mu_1 - \mu_2) \left(\frac{p(1-p)}{\mu_1\mu_2} \right)^{\frac{1}{2}} t_1 \right) \left(\inf_{y \in [0, t_1]} \kappa^1(\exp(2\mu_1\mu_2 y) - 1) \right) + (1 + \theta^{-\frac{1}{2}}) \left(\inf_{y \in [0, t_1]} \kappa^2(\exp(2\theta y) - 1) \right).$$

Let us also define

$$\mathcal{H}_2 \triangleq \min \left(\inf_{y \in [t_1, T]} \kappa^1(\exp(2\mu_1\mu_2 y) - 1), \inf_{y \in [t_1, T]} \kappa^2(\exp(2\theta y) - 1) \right).$$

Here we recall that the basic properties of continuous Gaussian processes ensure that $\inf_{y \in [0, t_1]} \kappa^2(\exp(2\theta y) - 1) < 0, \inf_{y \in [0, t_1]} \kappa^1(\exp(2\mu_1\mu_2 y) - 1) < 0$. Note that for $\epsilon, M > 0$, the event $\{\mathcal{H}_1 > -\epsilon, \mathcal{H}_2 > M\}$ implies the event $\{\inf_{t \in [0, t_1]} \Pi(t) > -\epsilon\}$, as well as the event (for all $t \in [t_1, T]$)

$$\begin{aligned} \int_0^t \exp(\theta y) \kappa^1(\exp(2\mu_1\mu_2 y) - 1) dy & \geq - \int_0^{t_1} \exp(\theta y) \epsilon dy + \int_{t_1}^t \exp(\theta y) \kappa^1(\exp(2\mu_1\mu_2 y) - 1) dy \\ & \geq -t_1 \exp(\theta t_1) \epsilon + (t - t_1) M \geq -t_1 \exp(\theta t_1) \epsilon; \end{aligned}$$

and also the event (for all $t \in [t_1, T]$)

$$\theta^{-\frac{1}{2}} \exp(-\theta t) \kappa^2(\exp(2\theta t) - 1) \geq \theta^{-\frac{1}{2}} \exp(-\theta T) M.$$

Summarizing the above, and combining with (38) - (40), we conclude that for $\epsilon, M > 0$, the event $\{\mathcal{H}_1 > -\epsilon, \mathcal{H}_2 > M\}$ implies the event

$$\left\{ \inf_{t \in [0, t_1]} (\Pi(t) + \text{BUF}(t)) \geq -\epsilon + \frac{|B|}{\theta} \right\}, \quad (41)$$

as well as the event

$$\begin{aligned} & \left\{ \inf_{t \in [t_1, T]} (\Pi(t) + \text{BUF}(t)) \geq \theta^{-\frac{1}{2}} \exp(-\theta T) M \right. \\ & \left. - (\mu_1 - \mu_2) \left(\frac{p(1-p)}{\mu_1\mu_2} \right)^{\frac{1}{2}} t_1 \exp(\theta t_1) \epsilon - \frac{(\mu_1 + \mu_2)(|C_3| + |C_4|) + |B|}{\theta} \right\}. \end{aligned} \quad (42)$$

Setting $\epsilon = \frac{|B|}{2\theta}$ and

$$M = 2\theta^{\frac{1}{2}} \exp(\theta T) \left((\mu_1 - \mu_2) \left(\frac{p(1-p)}{\mu_1\mu_2} \right)^{\frac{1}{2}} t_1 \exp(\theta t_1) \epsilon + \frac{(\mu_1 + \mu_2)(|C_3| + |C_4|) + |B|}{\theta} \right),$$

it follows that

$$\mathbb{E} \left[I \left(\inf_{t \in [0, T]} (\Pi(t) + \text{BUF}(t)) > 0 \right) \right] \geq P(\{\mathcal{H}_1 > -\epsilon, \mathcal{H}_2 > M\}).$$

That for this choice of M and ϵ (and in fact any strictly positive choice for M, ϵ) it holds that $P(\{\mathcal{H}_1 > -\epsilon, \mathcal{H}_2 > M\}) > 0$ follows from well-known properties of Brownian motion and associated standard arguments (see e.g. [17]), and we omit the details. Combining the above completes the proof. *Q.E.D.*

8.13. Proof of non-insensitivity.

To prove that the true large deviations exponent $-\frac{\theta(\theta + \mu_1\mu_2)}{2(\theta + \mu_1 + \mu_2 - 1)}$ is not insensitive, it suffices to exhibit two hyper-exponential distributions, with respective parameters $p^a, \mu_1^a, \mu_2^a, p^b, \mu_1^b, \mu_2^b$, s.t. these two hyper-exponential distributions have the same first two moments (with first moment equal to one), yet there exists strictly positive $\theta < \min(\mu_1^a, \mu_2^a, \mu_1^b, \mu_2^b)$ s.t. $\frac{\theta(\theta + \mu_1^a\mu_2^a)}{2(\theta + \mu_1^a + \mu_2^a - 1)} \neq \frac{\theta(\theta + \mu_1^b\mu_2^b)}{2(\theta + \mu_1^b + \mu_2^b - 1)}$. The existence of such a pair of hyper-exponential distributions is explicitly given in the following lemma.

LEMMA 39. *Let S_a be a hyper-exponentially distributed r.v. with parameters $\mu_1^a = \frac{1}{31}(15 - 2 \times \sqrt{2}), \mu_2^a = 2 \times \sqrt{2}, p^a = \frac{1}{257}(121 - 32 \times \sqrt{2})$; and S_b be a hyper-exponentially distributed r.v. with parameters $\mu_1^b = \frac{1}{7}(3 - \sqrt{2}), \mu_2^b = \sqrt{2}, p^b = \frac{1}{17}(7 - 4 \times \sqrt{2})$. Let $\theta_0 = \frac{1}{5}$. Then $\mu_1^a, \mu_2^a, \mu_1^b, \mu_2^b \in (0, \infty); \theta_0 \in \left(0, \min(\mu_1^a, \mu_2^a, \mu_1^b, \mu_2^b)\right)$; $\mathbb{E}[S_a] = \mathbb{E}[S_b] = 1$; $\mathbb{E}[S_a^2] = \mathbb{E}[S_b^2] = 4$; but $\frac{\theta_0(\theta_0 + \mu_1^a\mu_2^a)}{2(\theta_0 + \mu_1^a + \mu_2^a - 1)} \neq \frac{\theta_0(\theta_0 + \mu_1^b\mu_2^b)}{2(\theta_0 + \mu_1^b + \mu_2^b - 1)}$.*

Proof: That $\mu_1^a, \mu_2^a, \mu_1^b, \mu_2^b \in (0, \infty)$, and $p^a, p^b \in (0, 1)$ is easily verified. That $\mu_1^a > \theta_0$ follows from the fact that $\frac{1}{31} \times (15 - 2 \times \sqrt{2}) > \frac{1}{36} \times (15 - 4) = \frac{1}{4} > \frac{1}{5}$. That $\mu_1^b > \theta_0$ follows from the fact that $\frac{1}{7} \times (3 - \sqrt{2}) > \frac{1}{7} \times (3 - 1.5) = \frac{3}{14} > \frac{1}{5}$. That $\theta_0 < \min(\mu_2^a, \mu_2^b)$ is trivial. Furthermore, we claim that $\mathbb{E}[S_a] = \mathbb{E}[S_b] = 1$ and $\mathbb{E}[S_a^2] = \mathbb{E}[S_b^2] = 4$. Indeed, a straightforward calculation yields the following:

$$\begin{aligned} \mathbb{E}[S_a] &= \frac{\frac{1}{257}(121 - 32 \times \sqrt{2})}{\frac{1}{31}(15 - 2 \times \sqrt{2})} + \frac{1 - \frac{1}{257}(121 - 32 \times \sqrt{2})}{2 \times \sqrt{2}} \\ &= \frac{2 \times \sqrt{2} \times \frac{1}{257}(121 - 32 \times \sqrt{2}) + \frac{1}{31}(15 - 2 \times \sqrt{2}) \times \left(1 - \frac{1}{257}(121 - 32 \times \sqrt{2})\right)}{\frac{1}{31}(15 - 2 \times \sqrt{2}) \times 2 \times \sqrt{2}} \\ &= \frac{62 \times \sqrt{2} \times (121 - 32 \times \sqrt{2}) + (15 - 2 \times \sqrt{2}) \times (136 + 32 \times \sqrt{2})}{257 \times (15 - 2 \times \sqrt{2}) \times 2 \times \sqrt{2}} \end{aligned}$$

$$\begin{aligned}
&= \frac{(-62 \times 32 \times 2 + 15 \times 136 - 32 \times 2 \times 2) + \sqrt{2} \times (62 \times 121 + 15 \times 32 - 2 \times 136)}{-257 \times 8 + 514 \times 15 \times \sqrt{2}} \\
&= \frac{-2056 + 7710 \times \sqrt{2}}{-2056 + 7710 \times \sqrt{2}} = 1.
\end{aligned}$$

$$\begin{aligned}
\mathbb{E}[S_a^2] &= 2 \times \frac{\frac{1}{257}(121 - 32 \times \sqrt{2})}{\left(\frac{1}{31}(15 - 2 \times \sqrt{2})\right)^2} + 2 \times \frac{1 - \frac{1}{257}(121 - 32 \times \sqrt{2})}{(2 \times \sqrt{2})^2} \\
&= \frac{2 \times \frac{1}{257}(121 - 32 \times \sqrt{2}) \times (2 \times \sqrt{2})^2 + 2 \times \left(\frac{1}{31}(15 - 2 \times \sqrt{2})\right)^2 \times \left(1 - \frac{1}{257}(121 - 32 \times \sqrt{2})\right)}{\left(\frac{1}{31}(15 - 2 \times \sqrt{2})\right)^2 \times (2 \times \sqrt{2})^2} \\
&= \frac{16 \times 31^2 \times (121 - 32 \times \sqrt{2}) + 2 \times (15 - 2 \times \sqrt{2})^2 \times (257 - (121 - 32 \times \sqrt{2}))}{257 \times 8 \times (15 - 2 \times \sqrt{2})^2} \\
&= \frac{16 \times 31^2 \times (121 - 32 \times \sqrt{2}) + (466 - 120 \times \sqrt{2}) \times (136 + 32 \times \sqrt{2})}{479048 - 123360 \times \sqrt{2}} \\
&= \frac{\left(16 \times 31^2 \times 121 + 466 \times 136 - 120 \times 32 \times 2\right) + \left(-16 \times 31^2 \times 32 + 32 \times 466 - 120 \times 136\right) \times \sqrt{2}}{479048 - 123360 \times \sqrt{2}} \\
&= \frac{1916192 - 493440 \times \sqrt{2}}{479048 - 123360 \times \sqrt{2}} = 4;
\end{aligned}$$

$$\begin{aligned}
\mathbb{E}[S_b] &= \frac{\frac{1}{17}(7 - 4 \times \sqrt{2})}{\frac{1}{7}(3 - \sqrt{2})} + \frac{1 - \frac{1}{17}(7 - 4 \times \sqrt{2})}{\sqrt{2}} \\
&= \frac{\frac{1}{17} \times (7 - 4\sqrt{2}) \times \sqrt{2} + \left(1 - \frac{1}{17}(7 - 4 \times \sqrt{2})\right) \times \frac{1}{7} \times (3 - \sqrt{2})}{\frac{1}{7} \times (3 - \sqrt{2}) \times \sqrt{2}} \\
&= \frac{(7 \times 7 \times \sqrt{2} - 7 \times 4 \times 2) + (10 + 4\sqrt{2}) \times (3 - \sqrt{2})}{17 \times 3 \times \sqrt{2} - 17 \times 2} \\
&= \frac{(-7 \times 4 \times 2 + 10 \times 3 - 8) + (7 \times 7 - 10 + 12)\sqrt{2}}{-34 + 51\sqrt{2}} = 1;
\end{aligned}$$

$$\begin{aligned}
\mathbb{E}[S_b^2] &= 2 \times \frac{\frac{1}{17}(7 - 4 \times \sqrt{2})}{\left(\frac{1}{7}(3 - \sqrt{2})\right)^2} + 2 \times \frac{1 - \frac{1}{17}(7 - 4 \times \sqrt{2})}{(\sqrt{2})^2} \\
&= \frac{2 \times \frac{1}{17} \times (7 - 4\sqrt{2}) + \left(1 - \frac{1}{17}(7 - 4\sqrt{2})\right) \left(\frac{1}{7}(3 - \sqrt{2})\right)^2}{\left(\frac{1}{7}(3 - \sqrt{2})\right)^2} \\
&= \frac{(98 \times 7 - 98 \times 4 \times \sqrt{2}) + (10 + 4\sqrt{2}) \times (11 - 6\sqrt{2})}{17 \times (11 - 6\sqrt{2})} \\
&= \frac{(98 \times 7 + 10 \times 11 - 6 \times 4 \times 2) + (-98 \times 4 - 6 \times 10 + 4 \times 11)\sqrt{2}}{17 \times 11 - 17 \times 6 \times \sqrt{2}} = 4.
\end{aligned}$$

Next, we show that although S_a and S_b have the same first two moments, they lead to different large deviations exponents. Indeed, note that it suffices to demonstrate that

$$(\theta_0 + \mu_1^a \mu_2^a) \times (\theta_0 + \mu_1^b + \mu_2^b - 1) \neq (\theta_0 + \mu_1^b \mu_2^b) \times (\theta_0 + \mu_1^a + \mu_2^a - 1).$$

Then a straightforward calculation yields

$$\begin{aligned}
& (\theta_0 + \mu_1^a \mu_2^a) \times (\theta_0 + \mu_1^b + \mu_2^b - 1) - (\theta_0 + \mu_1^b \mu_2^b) \times (\theta_0 + \mu_1^a + \mu_2^a - 1) \\
&= \left(\frac{1}{5} + \frac{30}{31} \sqrt{2} - \frac{8}{31} \right) \times \left(\frac{1}{5} + \frac{3}{7} + \frac{6}{7} \sqrt{2} - 1 \right) \\
&\quad - \left(\frac{1}{5} + \frac{3}{7} \sqrt{2} - \frac{2}{7} \right) \times \left(\frac{1}{5} + \frac{15}{31} + \left(2 - \frac{2}{31} \right) \sqrt{2} - 1 \right) \\
&= (5 \times 7 \times 31)^{-1} \times \left((-63 + 1050\sqrt{2}) \times (-403 + 930\sqrt{2}) - (-93 + 465\sqrt{2}) \times (-343 + 2100\sqrt{2}) \right) \\
&= -(5 \times 7 \times 31)^{-1} \times (6510 + 126945\sqrt{2}) \neq 0,
\end{aligned}$$

completing the proof. *Q.E.D.*

8.14. Proof that $-\frac{\theta(\theta+\mu_1\mu_2)}{2(\theta+\mu_1+\mu_2-1)} < -\frac{\mu_1\mu_2\theta}{2(\mu_1+\mu_2-1)}$.

Since $\mathbb{E}[S] = 1$, and hence $\mu_1 + \mu_2 - 1 > 0$, to prove that $-\frac{\theta(\theta+\mu_1\mu_2)}{2(\theta+\mu_1+\mu_2-1)} < -\frac{\mu_1\mu_2\theta}{2(\mu_1+\mu_2-1)}$, it suffices to demonstrate that

$$(\theta + \mu_1\mu_2)(\mu_1 + \mu_2 - 1) - (\theta + \mu_1 + \mu_2 - 1)\mu_1\mu_2 > 0,$$

equivalently that $\theta(\mu_1 + \mu_2 - 1) - \theta\mu_1\mu_2 > 0$. It thus suffices to demonstrate that $\mu_1 + \mu_2 - 1 > \mu_1\mu_2$, equivalently that $\mu_1(1 - \mu_2) > 1 - \mu_2$. First, suppose $\mu_2 < 1$. Then the fact that $\mathbb{E}[S] = 1$ implies that $\mu_1 > 1$, and hence by the implied non-negativity of $1 - \mu_2$ we conclude that $\mu_1(1 - \mu_2) > 1 - \mu_2$, completing the proof. Alternatively, suppose that $\mu_2 > 1$, in which case $\mu_1 < 1$. It then follows from a nearly identical argument that $\mu_1(\mu_2 - 1) < \mu_2 - 1$, completing the proof in this case as well. *Q.E.D.*

References

- [1] Abramowitz, M., I. Stegun. "Handbook of mathematical functions: with formulas, graphs, and mathematical tables." Vol. 55. Courier Corporation, 1964.
- [2] Adan, I., B. Hathaway, and V. Kulkarni. "On First Come, First Served Queues with Two Classes of Impatient Customers." arXiv preprint arXiv:1710.03798 (2017).
- [3] Adler, R. "An introduction to continuity, extrema, and related topics for general Gaussian processes." Lecture Notes-Monograph Series 12 (1990): i-155.
- [4] Aghajani, R., K. Ramanan. "Ergodicity of an spde associated with a many-server queue." arXiv preprint arXiv:1512.02929, 2015.
- [5] Aghajani, R., K. Ramanan. "The Limit of Stationary Distributions of Many-Server Queues in the Halfin-Whitt Regime." arXiv preprint arXiv:1610.01118, 2016.
- [6] Arapostathis, A., and G. Pang. "Ergodic diffusion control of multiclass multi-pool networks in the Halfin-Whitt regime." The Annals of Applied Probability 26.5 (2016): 3110-3153.

-
- [7] Arapostathis, A., G. Pang, and N. Sandric. “Ergodicity of Levy-driven SDEs arising from multiclass many-server queues.” arXiv preprint arXiv:1707.09674 (2017).
- [8] Armony, M., N. Shimkin, and W. Whitt. “The impact of delay announcements in many-server queues with abandonment.” *Operations Research* 57.1 (2009): 66-81.
- [9] Asmussen, S. “Applied probability and queues.” Vol. 51. Springer Science and Business Media, 2008.
- [10] Ata, B., P. Glynn, and X. Peng. “An equilibrium analysis of a discrete-time Markovian queue with endogenous abandonments.” *Queueing Systems* 86.1-2 (2017): 141-212.
- [11] Baccelli, F., and G. Hebuterne. “On queues with impatient customers.” Diss. INRIA, 1981.
- [12] Barndorff-Nielsen, Ole E., and N. Shephard. “Integrated OU Processes and Non-Gaussian OU-based Stochastic Volatility Models.” *Scandinavian Journal of statistics* 30.2 (2003): 277-295.
- [13] Batt, R., and C. Terwiesch. “Waiting patiently: An empirical study of queue abandonment in an emergency department.” *Management Science* 61.1 (2015): 39-59.
- [14] Bercu, B., F. Proia, and N. Savy. “On Ornstein-Uhlenbeck driven by Ornstein-Uhlenbeck processes.” *Statistics and Probability Letters* 85 (2014): 36-44.
- [15] Bhattachary, P., and A. Ephremides. “Stochastic monotonicity properties of multi-server queues with impatient customers.” *Journal of Applied Probability* 28.3 (1991): 673-682.
- [16] Blanchet, J., and J. Dong. “Rare-event simulation for multi-server queues in the Halfin-Whitt regime.” *ACM SIGMETRICS Performance Evaluation Review* 39.4 (2012): 35-35.
- [17] Borodin, A., and P. Salminen. “Handbook of Brownian motion-facts and formulae.” Birkhuser, 2012.
- [18] Boxma, O., and P. Rosalia Waal. “Multi-server queues with impatient customers.” Centrum voor Wiskunde en Informatica, Department of Operations Research, Statistics, and System Theory, 1993.
- [19] Braverman, A., and J. G. Dai. “Stein’s method for steady-state diffusion approximations of $M/Ph/n+M$ systems.” *The Annals of Applied Probability* 27.1 (2017): 550-581.
- [20] Brown, L., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, L. Zhao. “Statistical analysis of a telephone call center: A queueing-science perspective.” *Journal of the American statistical association* 100, no. 469 (2005): 36-50.
- [21] Chen, H., and D. Yao. “Fundamentals of queueing networks: Performance, asymptotics, and optimization.” Vol. 46. Springer Science and Business Media, 2013.
- [22] Cohen, S., and R. Elliott. “Stochastic calculus and applications.” Birkhauser, 2015.
- [23] Cox, David Roxbee, et al. “Renewal theory.” Vol. 1. London: Methuen, 1967.
- [24] Dai, J., A. B. Dieker, X. Gao. “Validity of heavy-traffic steady-state approximations in many-server queues with abandonment.” *Queueing Systems* 78.1 (2014): 1-29.
- [25] Dai, J. G., and J.M. Harrison. “Reflected Brownian motion in an orthant: numerical methods for steady-state analysis.” *The Annals of Applied Probability* (1992): 65-86.

- [26] Dai, J. G., and S. He. “Customer abandonment in many-server queues.” *Mathematics of Operations Research* 35.2 (2010): 347-362.
- [27] Dai, J. G., and S. He. “Many-server queues with customer abandonment: A survey of diffusion and fluid approximations.” *Journal of Systems Science and Systems Engineering* 21.1 (2012): 1-36.
- [28] Dai, J. G., and S. He. “Many-server queues with customer abandonment: Numerical analysis of their diffusion model.” *Stochastic Systems* 3.1 (2013): 96-146.
- [29] Dai, J. G., S. He, and T. Tezcan. “Many-server diffusion limits for $G/Ph/n + GI$ queues.” *The Annals of Applied Probability* 20.5 (2010): 1854-1890.
- [30] Dai, J.G., and M. Miyazawa. “Reflecting Brownian motion in two dimensions: Exact asymptotics for the stationary distribution.” *Stochastic Systems* 1.1 (2011): 146-208.
- [31] Dieker, A.B., and X. Gao. “Positive recurrence of piecewise Ornstein-Uhlenbeck processes and common quadratic Lyapunov functions.” *The Annals of Applied Probability* 23.4 (2013): 1291-1317.
- [32] Ding, D., J. Ou, and J. Ang. “Analysis of ticket queues with renegeing customers.” *Journal of the Operational Research Society* 66.2 (2015): 231-246.
- [33] Doob, J.L. “The Brownian movement and stochastic equations.” *The Annals of Mathematics*, 43(2):351-369, 1942.
- [34] Dudley, R. (1967). “The sizes of compact subsets of Hilbert space and continuity of Gaussian processes”. *J. Functional Analysis*. 1: 290-330. doi:10.1016/0022-1236(67)90017-1. MR 0220340.
- [35] Dudley, R. “VN Sudakov’s work on expected suprema of Gaussian processes.” *High Dimensional Probability VII*. Springer International Publishing, 2016. 37-43.
- [36] Dupuis, P. “Large deviations analysis of reflected diffusions and constrained stochastic approximation algorithms in convex sets.” *Stochastics: An International Journal of Probability and Stochastic Processes* 21.1 (1987): 63-96.
- [37] Erlang, A. K. “On the rational determination of the number of circuits.” *The life and works of AK Erlang* (1948): 216-221.
- [38] Evdokimova, E., K. De Turck, and D. Fiems. “Coupled queues with customer impatience.” *Performance Evaluation* 118 (2018): 33-47.
- [39] Feldman, Z., A. Mandelbaum, W. Massey, and W. Whitt. (2008). “Staffing of time-varying queues to achieve time-stable performance.” *Management Science*, 54(2), 324-338.
- [40] Fleming P.J., A. Stolyar, and B. Simon. “Heavy traffic limit for a mobile phone system loss model.” (1994) *Proc. 2nd Internat. Conf. Telecommunication Systems, Modeling, and Anal* TN:158176
- [41] Gamarnik, D., D.A. Goldberg. “Steady-state $GI/G/n$ queue in the Halfin-Whitt regime.” *The Annals of Applied Probability* 23.6 (2013): 2382-2419.

-
- [42] Gamarnik, D., P. Momcilovic. “Steady-state analysis of a multiserver queue in the Halfin-Whitt regime.” *Advances in Applied Probability* 40.2 (2008): 548-577.
- [43] Gamarnik, D., A. Stolyar. “Multiclass multiserver queueing system in the HalfinWhitt heavy traffic regime: asymptotics of the stationary distribution.” *Queueing Systems* 71.1-2 (2012): 25-51.
- [44] Garnett, O., A. Mandelbaum, and M. Reiman. “Designing a call center with impatient customers.” *Manufacturing and Service Operations Management* 4.3 (2002): 208-227.
- [45] Goldberg, D.A. “On the steady-state probability of delay and large negative deviations for the $GI/GI/n$ queue in the Halfin-Whitt regime.” arXiv preprint arXiv:1307.0241 (2016).
- [46] Goldberg, D.A., Y. Li. “Simple and explicit bounds for multi-server queues with universal $\frac{1}{1-\rho}$ scaling.” arXiv preprint arXiv:1706.04628 (2017).
- [47] Goldberg, D.A., and Y. Li. “Heavy-tailed queues in the Halfin-Whitt regime.” arXiv preprint arXiv:1707.07775 (2017).
- [48] Gurvich, I., J. Huang, A. Mandelbaum. “Excursion-based universal approximations for the Erlang-A queue in steady-state.” *Mathematics of Operations Research* 39.2 (2013): 325-373.
- [49] Gurvich, I. “Diffusion models and steady-state approximations for exponentially ergodic Markovian queues.” *The Annals of Applied Probability* 24.6 (2014): 2527-2559.
- [50] Halfin, S., W. Whitt. “Heavy-traffic limits for queues with many exponential servers.” *Operations research* 29.3 (1981): 567-588.
- [51] Hildebrandt, T.H. “Introduction to the Theory of Integration.” Vol. 13. New York: Academic Press, 1963.
- [52] Huang, J., and H. Zhang. “Diffusion approximations for open Jackson networks with reneging.” *Queueing Systems* 74.4 (2013): 445-476.
- [53] Huang, J., H. Zhang, and J. Zhang. “A unified approach to diffusion analysis of queues with general patience-time distributions.” *Mathematics of Operations Research* 41.3 (2016): 1135-1160.
- [54] Ibrahim, R., and W. Whitt. “Real-time delay estimation in overloaded multiserver queues with abandonments.” *Management Science* 55.10 (2009): 1729-1742.
- [55] Jagerman, D. “Some properties of the Erlang loss function.” *Bell System Technical Journal* 53.3 (1974): 525-551.
- [56] Jennings, O., and J. Pender. “Comparisons of ticket and standard queues.” *Queueing Systems* 84.1-2 (2016): 145-202.
- [57] Jennings, O., and J. Reed. “An overloaded multiclass FIFO queue with abandonments.” *Operations research* 60.5 (2012): 1282-1295.
- [58] Jouini, O., and Y. Dallery. “Monotonicity properties for multiserver queues with reneging and finite waiting lines.” *Probability in the Engineering and Informational Sciences* 21.3 (2007): 335-360.

- [59] Kang, W., and G. Pang. “Equivalence of Fluid Models for $G_t/GI/N + GI$ Queues.” arXiv preprint arXiv:1502.00346 (2015).
- [60] Kang, W., and K. Ramanan. “Asymptotic approximations for stationary distributions of many-server queues with abandonment.” *The Annals of Applied Probability* 22.2 (2012): 477-521.
- [61] Kang, W., and K. Ramanan. “Characterization of stationary distributions of reflected diffusions.” *The Annals of Applied Probability* 24.4 (2014): 1329-1374.
- [62] Karatzas, I., and S. Shreve. “Brownian motion and stochastic calculus.” Vol. 113. Springer Science and Business Media, 2012.
- [63] Katsuda, T. “General hazard-type scaling of abandonment time distribution for a $G/Ph/n + GI$ queue in the Halfin-Whitt heavy-traffic regime.” *Queueing Systems* 80.1-2 (2015): 155-195.
- [64] Ledoux, M. and M. Talagrand (1991). “Probability in Banach spaces.” Berlin: Springer-Verlag. pp. xii+480. ISBN 3-540-52013-9. MR 1102015. (See chapter 11)
- [65] van Leeuwaarden, J., and C. Knessl. “Transient behavior of the Halfin-Whitt diffusion.” *Stochastic Processes and their Applications* 121.7 (2011): 1524-1545.
- [66] van Leeuwaarden, J., B. Mathijssen, and B. Zwart. “Economies-of-scale in resource sharing systems: tutorial and partial review of the QED heavy-traffic regime.” arXiv preprint arXiv:1706.05397 (2017).
- [67] Lingenbrink, D., and K. Iyer. “Optimal Signaling Mechanisms in Unobservable Queues with Strategic Customers.” In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 347-347. ACM, 2017.
- [68] Liu, Y., and W. Whitt. “Large-time asymptotics for the $G_t/M_t/s_t + GI_t$ many-server fluid queue with abandonment.” *Queueing systems* 67.2 (2011): 145-182.
- [69] Liu, Y., and W. Whitt. “Nearly periodic behavior in the overloaded $G/D/s + GI$ queue.” *Stochastic Systems* 1.2 (2011): 340-410.
- [70] Liu, Y., and W. Whitt. “The $G_t/GI/s_t + GI$ many-server fluid queue.” *Queueing Systems* 71.4 (2012): 405-444.
- [71] Liu, Y., and W. Whitt. “Stabilizing performance in a service system with time-varying arrivals and customer feedback.” *European Journal of Operational Research* 256.2 (2017): 473-486.
- [72] Low, Y., and S. Ong. “Count distribution for mixture of two exponentials as renewal process duration with applications.” *AIP Conference Proceedings*. Vol. 1739. No. 1. AIP Publishing, 2016.
- [73] Maller, R., G. Muller, and A. Szimayer. “Ornstein-Uhlenbeck processes and extensions.” *Handbook of Financial Time Series* (2009): 421-437.
- [74] Mandelbaum, A., W.A. Massey, and M.I. Reiman, “Strong approximations for Markovian service networks.” *Queueing Systems* 30 (1998) 149201.

-
- [75] Mandelbaum, A., Massey, W. A., Reiman, M. I., Stolyar, A., and Rider, B. (2002). “Queue lengths and waiting times for multi-server queues with abandonment and retrials.” *Telecommunication Systems*, 21(2), 149-171.
- [76] Mandelbaum, A., and P. Momcilovic. “Queues with many servers and impatient customers.” *Mathematics of Operations Research* 37.1 (2012): 41-65.
- [77] Mandelbaum, A., and N. Shimkin. “A model for rational abandonments from invisible queues.” *Queueing Systems* 36.1 (2000): 141-173.
- [78] Mandelbaum, A., and S. Zeltyn. “The impact of customers patience on delay and abandonment: some empirically-driven experiments with the $M/M/n+G$ queue.” *OR Spectrum* 26.3 (2004): 377-411.
- [79] Mandelbaum, A., and S. Zeltyn. “Call centers with impatient customers: many-server asymptotics of the $M/M/n+G$ queue.” *Queueing Systems* 51.3 (2005): 361-402.
- [80] Mandelbaum, A., and S. Zeltyn. “Staffing many-server queues with impatient customers: constraint satisfaction in call centers.” *Operations Research* 57.5 (2009): 1189-1205.
- [81] Mandelbaum, A., and S. Zeltyn. “Data-stories about (im) patient customers in tele-queues.” *Queueing Systems* 75.2-4 (2013): 115-146.
- [82] Massey, W., and J. Pender. “Gaussian skewness approximation for dynamic rate multi-server queues with abandonment.” *Queueing Systems* 75.2-4 (2013): 243-277.
- [83] Masuda, H. “On multi-dimensional Ornstein-Uhlenbeck processes driven by a general Levy process.” *Bernoulli* 10.1 (2004): 97-120.
- [84] Meyn, S., and R. L. Tweedie. “Stability of Markovian processes. I. Criteria for discrete-time chains.” *Adv. in Appl. Probab.*, 24(3):542574, 1992
- [85] Meyn, S., and R.L. Tweedie. “Stability of Markovian processes II: Continuous-time processes and sampled chains.” *Advances in Applied Probability* 25.3 (1993): 487-517.
- [86] Meyn, S., and R. L. Tweedie. “Stability of Markovian processes III: Foster-Lyapunov criteria for continuous-time processes.” *Advances in Applied Probability* 25.3 (1993): 518-548.
- [87] Ni, E., and S. Henderson. “How Hard are Steady-State Queueing Simulations?.” *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 25.4 (2015): 27.
- [88] Pender, J., and W. Massey. “Approximating and stabilizing dynamic rate Jackson networks with abandonment.” *Probability in the Engineering and Informational Sciences* 31.1 (2017): 1-42.
- [89] Moyal, P. “Coupling in the queue with impatience: case of several servers.” *arXiv preprint arXiv:1711.02931* (2017).
- [90] Pang, G., R. Talreja, and W. Whitt. “Martingale proofs of many-server heavy-traffic limits for Markovian queues.” *Probability Surveys* 4 (2007): 193-267.

- [91] Pender, J. “Gram charlier expansion for time varying multiserver queues with abandonment.” *SIAM Journal on Applied Mathematics* 74.4 (2014): 1238-1265.
- [92] Peszat, S., and J. Zabczyk. “Strong Feller property and irreducibility for diffusions on Hilbert spaces.” *The Annals of Probability* (1995): 157-172.
- [93] Puhalskii, A.A. “On the $M_t/M_t/K_t + M_t$ queue in heavy traffic.” *Mathematical Methods of Operations Research* 78.1 (2013): 119-148.
- [94] Puhalskii, A.A., and J. Reed. “On many-server queues in heavy traffic.” *The Annals of Applied probability* 20.1 (2010): 129-195.
- [95] Puhalskii, A.A., and M. Reiman. “The multiclass GI/PH/N queue in the Halfin-Whitt regime.” *Advances in Applied Probability* 32.2 (2000): 564-595.
- [96] Reed, J. “The G/GI/N queue in the Halfin-Whitt regime.” *The Annals of Applied Probability* 19.6 (2009): 2211-2269.
- [97] Reed, J., and T. Tezcan. “Hazard rate scaling of the abandonment distribution for the GI/M/n+ GI queue in heavy traffic.” *Operations Research* 60.4 (2012): 981-995.
- [98] Reed, J. E., and A. Ward. “Approximating the GI/GI/1+ GI queue with a nonlinear drift diffusion: Hazard rate scaling in heavy traffic.” *Mathematics of Operations Research* 33.3 (2008): 606-644.
- [99] Reed, J., and A. Ward. “A diffusion approximation for a generalized Jackson network with reneging.” *Proceedings of the 42nd Allerton Conference on Communication, Control, and Computing*. 2004.
- [100] Reed, J., A. Ward, and D. Zhan. “On the generalized drift Skorokhod problem in one dimension.” *Journal of Applied Probability* 50.1 (2013): 16-28.
- [101] Sadowsky, J. “Large deviations theory and efficient simulation of excessive backlogs in a GI/GI/m queue.” *IEEE Transactions on Automatic Control* 36.12 (1991): 1383-1394.
- [102] Sarhangian, V., and B. Balcioglu. “Waiting time analysis of multi-class queues with impatient customers.” *Probability in the Engineering and Informational Sciences* 27.3 (2013): 333-352.
- [103] Sato, K., T. Watanabe, and M. Yamazato. “Recurrence conditions for multi-dimensional processes of Ornstein-Uhlenbeck type.” *Journal of the Mathematical Society of Japan* 46.2 (1994): 245-265.
- [104] Shwartz, A., and A. Weiss. “Large deviations for performance analysis: queues, communication and computing.” Vol. 5. CRC Press, 1995.
- [105] Stoyan, D., and D. Daley (1983). “Comparison methods for queues and other stochastic models.” John Wiley and Sons, Inc., 605 Third Avenue, New York, NY 10158, USA, 1983.
- [106] Talreja, R., and W. Whitt. “Heavy-traffic limits for waiting times in many-server queues with abandonment.” *The Annals of Applied Probability* (2009): 2137-2175.
- [107] Tezcan, T. “State space collapse in many-server diffusion limits of parallel server systems and applications.” Diss. Georgia Institute of Technology, 2006.

-
- [108] Ward, A. “Asymptotic analysis of queueing systems with reneging: A survey of results for FIFO, single class models.” *Surveys in Operations Research and Management Science* 17.1 (2012): 1-14.
- [109] Ward, A., and P. Glynn. “A diffusion approximation for a Markovian queue with reneging.” *Queueing Systems* 43.1 (2003): 103-128.
- [110] Ward, A., and P. Glynn. “Properties of the reflected Ornstein-Uhlenbeck process.” *Queueing Systems* 44.2 (2003): 109-123.
- [111] Ward, A., and P. Glynn. “A diffusion approximation for a GI/GI/1 queue with balking or reneging.” *Queueing Systems* 50.4 (2005): 371-400.
- [112] Weerasinghe, A. “Diffusion approximations for G/M/n+ GI queues with state-dependent service rates.” *Mathematics of Operations Research* 39.1 (2013): 207-228.
- [113] Whitt, W. “Queues with superposition arrival processes in heavy traffic.” *Stochastic processes and their applications* 21.1 (1985): 81-91.
- [114] Whitt, W. “The impact of a heavy-tailed service-time distribution upon the M/GI/s waiting-time distribution.” *Queueing Systems* 36.1-3 (2000): 71-87.
- [115] Whitt, W. “Stochastic-process limits: an introduction to stochastic-process limits and their application to queues.” Springer Science and Business Media, 2002.
- [116] Whitt, W. “Efficiency-driven heavy-traffic approximations for many-server queues with abandonments.” *Management Science* 50.10 (2004): 1449-1461.
- [117] Whitt, W. “Two fluid approximations for multi-server queues with abandonments.” *Operations Research Letters* 33.4 (2005): 363-372.
- [118] Whitt, W. “Heavy-traffic limits for the $G/H_2^*/n/m$ queue.” *Mathematics of Operations Research* 30.1 (2005): 1-27.
- [119] Whitt, W. “Fluid models for multi-server queues with abandonments.” *Operations research* 54.1 (2006): 37-54.
- [120] Whitt, W. “Sensitivity of performance in the Erlang-A queueing model to changes in the model parameters.” *Operations Research* 54.2 (2006): 247-260.
- [121] Wu, C., A. Bassamboo, and O. Perry. “Service Systems with Dependent Service and Patience Times.” (2017).
- [122] Xing, X., W. Zhang, and Y. Wang. “The stationary distributions of two classes of reflected Ornstein-Uhlenbeck processes.” *Journal of Applied Probability* 46.3 (2009): 709-720.
- [123] Zhang, J. “Fluid models of many-server queues with abandonment.” *Queueing Systems* 73.2 (2013): 147-193.
- [124] Zhang, B., J. van Leeuwen, and B. Zwart. “Staffing call centers with impatient customers: refinements to many-server asymptotics.” *Operations Research* 60.2 (2012): 461-474.