

ROBUST BOUNDS AND OPTIMIZATION AT THE LARGE DEVIATIONS SCALE FOR QUEUEING MODELS VIA RÉNYI DIVERGENCE

BY RAMI ATAR¹, AMARJIT BUDHIRAJA², PAUL DUPUIS³ AND RUOYU WU⁴

¹*Viterbi Faculty of Electrical Engineering, Technion, rami@technion.ac.il*

²*Department of Statistics and Operations Research, University of North Carolina, amarjit@unc.edu*

³*Division of Applied Mathematics, Brown University, paul_dupuis@brown.edu*

⁴*Department of Mathematics, Iowa State University, ruoyu@iastate.edu*

This paper develops tools to obtain robust probabilistic estimates for queueing models at the large deviations (LD) scale. These tools are based on the recently introduced robust Rényi bounds, which provide LD estimates (and more generally risk-sensitive (RS) cost estimates) that hold uniformly over an uncertainty class of models, provided that the class is defined in terms of Rényi divergence with respect to a reference model and that estimates are available for the reference model. One very attractive quality of the approach is that the class to which the estimates apply may consist of hard models, such as highly non-Markovian models and ones for which the LD principle is not available. Our treatment provides exact expressions as well as bounds on the Rényi divergence rate on families of marked point processes, including as a special case renewal processes. Another contribution is a general result that translates robust RS control problems, where robustness is formulated via Rényi divergence, to finite dimensional convex optimization problems, when the control set is a finite dimensional convex set. The implications to queueing are vast, as they apply in great generality. This is demonstrated on two non-Markovian queueing models. One is the multiclass single-server queue considered as a RS control problem, with scheduling as the control process and exponential weighted queue length as cost. The second is the many-server queue with reneging, with the probability of atypically large reneging count as performance criterion. As far as LD analysis is concerned, no robust estimates or non-Markovian treatment were previously available for either of these models.

1. Introduction. An approach for obtaining robust estimates on probabilistic models at the large deviations (LD) scale, as well as on risk-sensitive (RS) functionals associated with these models, has recently been proposed based on Rényi divergence (RD) estimates [3]. According to this approach, a family of models is considered that is defined in terms of RD with respect to a reference model. A tool, that we call in this paper *robust Rényi bounds* (RRB), is then used to translate LD probability estimates (and more generally, RS cost estimates) on the reference model into ones which hold uniformly within this family. This approach is particularly useful in cases when the reference model is one that is easier to analyze than the collection of models on which robust bounds are desired. This paper applies these ideas to queueing models.

Indeed, queueing forms an ideal domain of applicability of this approach, for two main reasons. First, it is very often the case that Markovian queueing models are considerably easier to handle than non-Markovian ones. Among the many examples that strongly support this assertion we mention (1) the $M/M/n$ model for which the many-server law of large numbers (LLN) limit is trivial as opposed to the $G/G/n$ counterpart for which theory is involved and,

in particular, limit processes lie in the space of measure valued trajectories [27]; similarly, at the central limit theorem (CLT) scale, these two models give rise to merely a diffusion on \mathbb{R} [21] and a considerably more complicated, measure-valued diffusion [28], respectively. (2) Queueing control problems, that in a Markovian setting can be analyzed and solved as Markov decision processes (see numerous examples in [39]), but in a general setting, such as when service times are nonexponential, require an infinite dimensional state descriptor and are less tractable. Second, the robustness of estimates to perturbations in the underlying distributions is important in applications. Exponential service distribution (necessary for Markovity) is often assumed without good statistical evidence or physical reasoning. For example, a detailed statistical study argues that there is a good fit of service time distributions in call centers to lognormal [9], but there are far more papers on many-server scaling limits, aimed at modelling large call centers, in which servers operate with exponential distributions than ones treating more general distributions. In a much broader perspective, uncertainty in the underlying distributions is a central issue in applying probabilistic queueing models to real world systems.

To put LD estimates in a broader context of scaling limits as far as sensitivity to perturbations in the underlying distributions is concerned, it should be mentioned that most LLN and CLT results in the queueing literature are tolerant to such perturbations in the sense that the limits depend only on first or first and second moments of the primitive data (the many-server limit regime alluded to above is an exception). This has made these regimes attractive for approximations and indeed provided motivation to study them. On the other hand, the LD regime does not have obvious robustness properties, as probabilities of rare events are sensitive to the assumed tails of the primitives. Consequently, model uncertainty issues and sensitivity to distributional perturbations are much harder to deal with. As already mentioned, this paper addresses these questions by developing the approach of [3] in the context of queueing models.

The development in this paper, which involves performance measures that are determined by rare events and bounds defined in terms of Rényi divergence, is analogous to prior work that bounds ordinary performance measures in terms of Kullback–Leibler divergence (also known as relative entropy). This approach originated in a robust optimal control framework in [15, 34], and was subsequently rediscovered a number of times and much developed in the literature [23, 30, 32]. The corresponding use in model uncertainty bounds and sensitivity bounds appeared later, as in for example, [11, 16].

The literature on LD estimates for queueing models is rich. A partial list of works dealing with non-Markovian queueing models is as follows. In [1], weak limit theorems are proved for the behavior of a $G/G/1$ queue conditioned to exhibit an atypically large waiting time. In [20] the tail behavior of the waiting time steady state distribution is identified for a large class of single server queues. In [31], multiclass feedforward networks are studied at the moderate deviations scale. In [36], the LD principle (LDP) is established and the rate function is identified for the generalized Jackson network. See further references in these sources as well as in the monographs [10, 19, 40] and the paper [13] for numerous results on a variety of models in both Markovian and non-Markovian settings.

Sample path LDP of queueing models are particularly difficult in network settings, due to the fact that these models have discontinuous statistics. References [13, 19, 31, 36, 40] do succeed in addressing such LDP. Yet, even when tools such as LDP and formulas for the rate function are available, a direct approach for obtaining estimates for an event of interest, uniformly over a given family of models, may be notoriously hard, as it amounts to solving a variational problem for each member in the family. Unlike such a naive approach, under the approach based on RRB, LD estimates have to be studied only for the reference model. In fact, the approach does not even require that the LDP holds for each model in the family.

The rest of this paper is organized as follows. The general approach that uses RRB to get robust estimates on families of models is summarized in Section 2. In the same section, an outline of the use of these bounds for queueing models is provided, showing that for these models the bounds heavily rely on estimating the Rényi divergence rate (an asymptotic normalized version of the Rényi divergence) of a renewal process with respect to a Poisson. This provides a motivation to study such estimates for various families of renewal processes. Results in this direction appear in Section 3. Perhaps surprisingly, it seems that such calculations have not appeared before in the literature. In Section 4 we provide a general development on RS control, and demonstrate it with a queueing example. Estimates on RS control are closely related to LD estimates, and in this section we argue that they can be addressed by RRB. A general result developed in Section 4.1 shows that a dramatic simplification occurs when RRB is used for RS control problems, by which robust control estimates are transformed into finite dimensional convex optimization problems when the control set is a finite dimensional convex set. In Section 4.2, we analyze a queueing control problem using this approach. The model considered is the multiclass $G/G/1$ queue, in which the control corresponds to scheduling jobs from the various classes. As a reference model we use known RS control estimates for the multiclass $M/M/1$. Finally, Section 5 provides a queueing example for our robust approach to LD estimates. The example consists of a queueing model with reneging. Whereas reneging from queues is a very active research field, little is known on LD estimates beyond the Markovian setting. The robust LD estimates provided in this section are on both the $G/G/1 + G$ and the many-server $G/G/n + G$ models. The reference model on which they rely is the $M/M/n + M$, for which the sample path LDP has recently been developed in [2].

2. Robust Rényi bounds. This section introduces the RRB and the approach that uses these bounds to quantify robustness. The RRB are described in Section 2.1 and the form they take under scaling is derived in Section 2.2. In Section 2.3 it is argued that in queueing applications the Rényi divergence of a renewal process w.r.t. a Poisson is key in the use of the approach, and the notion of Rényi divergence rate is introduced.

2.1. *Rényi divergence.* Fix a measurable space (S, \mathcal{F}) and denote by \mathcal{P} the set of probability measures on it. For $P, Q \in \mathcal{P}$, the *relative entropy* is given by

$$R(Q\|P) = \begin{cases} \int \log \frac{dQ}{dP} dQ & \text{if } Q \ll P, \\ +\infty & \text{otherwise.} \end{cases}$$

Introduced in [38] (see [29] for a comprehensive treatment), the *Rényi divergence* of degree $\alpha > 1$, for $P, Q \in \mathcal{P}$, is defined by

$$R_\alpha(Q\|P) = \begin{cases} \frac{1}{\alpha(\alpha - 1)} \log \int \left(\frac{dQ}{dP}\right)^\alpha dP & \text{if } Q \ll P, \\ +\infty & \text{otherwise.} \end{cases}$$

For $\alpha = 1$, one sets $R_1(Q\|P) = R(Q\|P)$. Whereas two different formulas are used for the cases $\alpha = 1$ and $\alpha > 1$, it is a fact that $\alpha \mapsto R_\alpha(Q\|P)$ is continuous on $[1, \alpha^*]$ provided $R_{\alpha^*}(Q\|P) < \infty$ for some $\alpha^* > 1$. To mention a few additional properties, one has that $\alpha \mapsto \alpha R_\alpha$ is nondecreasing on $[1, \infty)$, and given $\alpha \geq 1$, one always has $R_\alpha(Q\|P) \geq 0$, and $R_\alpha(Q\|P) = 0$ if and only if $Q = P$. A property that is of crucial importance in our use of Rényi divergence is its additivity for product measures, in the following sense:

$$(2.1) \quad R_\alpha(Q_1 \times Q_2\|P_1 \times P_2) = R_\alpha(Q_1\|P_1) + R_\alpha(Q_2\|P_2).$$

It is well known that exponential integrals and relative entropy satisfy a convex duality relation, stated as follows. Let $Q \in \mathcal{P}$. Then for any bounded measurable $g : \mathcal{S} \rightarrow \mathbb{R}$,

$$(2.2) \quad \log \int e^g dQ = \sup_{P \in \mathcal{P}} \left[\int g dP - R(P \| Q) \right].$$

An analogous relation has been shown for Rényi divergences ([3]; related calculations first appeared in [18]). Namely, fix $\alpha > 1$. Then

$$(2.3) \quad \frac{1}{\alpha} \log \int e^{\alpha g} dQ = \sup_{P \in \mathcal{P}} \left[\frac{1}{\alpha - 1} \log \int e^{(\alpha - 1)g} dP - R_\alpha(P \| Q) \right].$$

The identity (2.3) may indeed be viewed as an extension of (2.2), as the latter is recovered by taking the formal limit $\alpha \downarrow 1$ in the former.

Given P, Q and α , as well as an event $A \in \mathcal{F}$, it follows from (2.3) by taking $g(x) = 0$ [resp., $-M$] for $x \in A$ [resp., $x \in A^c$] and sending $M \rightarrow \infty$, that

$$(2.4) \quad \begin{aligned} & \frac{\alpha}{\alpha - 1} \log P(A) - \alpha R_\alpha(P \| Q) \\ & \leq \log Q(A) \leq \frac{\alpha - 1}{\alpha} \log P(A) + (\alpha - 1)R_\alpha(Q \| P) \end{aligned}$$

(provided that $P(A) > 0$ and $Q(A) > 0$). The first inequality uses (2.3) as written, and the second reverses the roles of P and Q . In words: the logarithmic probability of an event under Q is estimated in terms of the same event under P and Rényi divergence. It is also a fact that both inequalities in (2.4) are tight, in the sense that given α, Q and A one can find P that makes them hold as equalities (with different P for each equality) [3].

The point of view of [3] is to regard (2.4) as perturbation bounds. Given a nominal model P , (2.4) provides performance bounds on a true model Q in terms of performance under P and divergence terms. The same is true in the more general case of a RS cost, namely

$$(2.5) \quad \log \int e^{(\alpha - 1)g} dQ \leq \frac{\alpha - 1}{\alpha} \log \int e^{\alpha g} dP + (\alpha - 1)R_\alpha(Q \| P).$$

In this paper we refer to (2.5) and its special case (2.4) as robust Rényi bounds (RRB). If one fixes a reference model P and a family \mathcal{Q} of true models Q defined by $\{Q : (\alpha - 1)R_\alpha(Q \| P) < r\}$, some $r > 0$, then for any A (2.4) gives $\sup_{Q \in \mathcal{Q}} \log Q(A) \leq \frac{\alpha - 1}{\alpha} \log P(A) + r$. This expresses a uniform estimate on the performance under Q in \mathcal{Q} in terms of that under P and the size of the family (where the latter term is interpreted in terms of Rényi divergence). Clearly, an analogous statement can be made for RS cost by appealing to (2.5), and similarly for lower bounds, by working with $R_\alpha(P \| Q)$ instead of $R_\alpha(Q \| P)$.

2.2. The RRB under scaling. What makes the RRB particularly useful is that they remain meaningful under standard LD scaling. We first demonstrate this in a setting of IID random variables (RVs), and then extend to a continuous time setting.

IID data. Let $Z_1, Z_2, Z_3 \dots$ be a sequence of RVs, and let P and Q be two probability measures that make this sequence IID. Let P_n and Q_n denote the corresponding laws of $Z^n = (Z_1, \dots, Z_n)$. For each n , let A_n be an event that is measurable on $\sigma\{Z^n\}$, the σ -algebra generated by Z^n . We are interested in

$$\frac{1}{n} \log Q(A_n).$$

By the IID assumption, we may appeal to (2.1), according to which $R_\alpha(Q_n \| P_n) = nR_\alpha(Q_1 \| P_1)$. Thus by (2.4) we obtain the bounds

$$(2.6) \quad \begin{aligned} & \frac{\alpha}{\alpha - 1} \frac{1}{n} \log P(A_n) - \alpha R_\alpha(P_1 \| Q_1) \\ & \leq \frac{1}{n} \log Q(A_n) \leq \frac{\alpha - 1}{\alpha} \frac{1}{n} \log P(A_n) + (\alpha - 1)R_\alpha(Q_1 \| P_1). \end{aligned}$$

In these bounds, the divergence terms remain of order 1 under scaling, and so it is possible to compare the asymptotic behavior of $n^{-1} \log Q(A_n)$ to that of $n^{-1} \log P(A_n)$. Moreover, while standard problems in the theory of LD are concerned with limits of these expressions, we emphasize that the bounds (2.6) are valid for all n .

Regarding the normalized logarithmic probability as a performance measure in this setting is indeed natural for studying probabilities of rare events. Thus our remark from Section 2.1 regarding uniform estimates on logarithmic probabilities is relevant also for exponential decay rates. That is, given $r > 0$, let \mathcal{Q} consist of probability measures Q under which X_1, X_2, \dots are IID and $(\alpha - 1)R_\alpha(Q_1 \| P_1) \leq r$. Then (2.6) gives

$$\sup_{Q \in \mathcal{Q}} \frac{1}{n} \log Q(A_n) \leq \frac{\alpha - 1}{\alpha} \frac{1}{n} \log P(A_n) + r.$$

Again, a similar remark holds for RS cost, and a lower bound is obtained similarly by working with $R_\alpha(P_1 \| Q_1)$.

Beyond IID data. When the model is not based on an IID structure one can still apply the RRB under scaling, but one must address the question whether the normalized Rényi divergence term scales suitably. Let us present this issue in a continuous time setting that better suits the aims of this paper. Let $\{Z_t, t \in \mathbb{R}_+\}$ be a stochastic process on the measurable space $(\mathcal{S}, \mathcal{F})$ and, throughout this paper, for a general probability measure $Q \in \mathcal{P}$ denote $Q_t^Z = Q \circ Z|_{[0,t]}^{-1}$ (when there is no room for confusion, the dependence on the process is omitted from the notation). Then for any $t > 0$, any event A_t measurable on $\sigma\{Z|_{[0,t]}\}$, measure P and collection of measures \mathcal{Q} , we have by (2.4)

$$(2.7) \quad \sup_{Q \in \mathcal{Q}} \frac{1}{t} \log Q(A_t) \leq \frac{\alpha - 1}{\alpha} \frac{1}{t} \log P(A_t) + (\alpha - 1) \sup_{Q \in \mathcal{Q}} \frac{1}{t} R_\alpha(Q_t^Z \| P_t^Z).$$

If the last term remains bounded as $t \rightarrow \infty$ then one obtains uniform LD estimates within the family \mathcal{Q} by LD estimates on the reference model P and the Rényi divergence term. This method then remains effective in cases where the latter term can be computed or estimated.

For statements that involve the limit $t \rightarrow \infty$, we shall need a further piece of notation, used throughout. Given a process Z on $(\mathcal{S}, \mathcal{F})$ and measures $P, Q \in \mathcal{P}$, the *Rényi divergence rate* (RDR) of Q w.r.t. P associated with the process Z is defined by

$$r_\alpha^Z(Q \| P) = \limsup_{t \rightarrow \infty} \frac{1}{t} R_\alpha(Q_t^Z \| P_t^Z).$$

For a family \mathcal{Q} of probability measures, let the RDR of \mathcal{Q} w.r.t. P and of P w.r.t. \mathcal{Q} be defined, respectively, by

$$r_\alpha^Z(\mathcal{Q} \| P) = \limsup_{t \rightarrow \infty} \sup_{Q \in \mathcal{Q}} \frac{1}{t} R_\alpha(Q_t^Z \| P_t^Z), \quad r_\alpha^Z(P \| \mathcal{Q}) = \limsup_{t \rightarrow \infty} \sup_{Q \in \mathcal{Q}} \frac{1}{t} R_\alpha(P_t^Z \| Q_t^Z).$$

Again, the dependence on Z will be omitted when there is no confusion. With this notation, we have

$$(2.8) \quad \limsup_{t \rightarrow \infty} \sup_{Q \in \mathcal{Q}} \frac{1}{t} \log Q(A_t) \leq \frac{\alpha - 1}{\alpha} \limsup_{t \rightarrow \infty} \frac{1}{t} \log P(A_t) + (\alpha - 1)r_\alpha^Z(\mathcal{Q} \| P).$$

Often, the selection of P with which to apply the bound (2.7) or (2.8) is based on considerations of tractability. If, for example, P is a model under which performance can be explicitly computed then one may use the approach in order to obtain guaranteed bounds on a set of possibly intractable models Q . Another consideration, that is especially relevant in engineering applications, is that systems often operate under conditions that are distinct from those they are designed for. For such systems, the bounds provide guarantees on their true performance based on designed performance.

As a final general remark, given a particular event or a sequence of events that are of interest, one can optimize over the parameter α for the tightest upper and lower bounds. Namely, in both (2.7) and (2.8) one may take the infimum over $\alpha > 1$ on the right hand side. This observation will be used in Section 4.

2.3. Queueing models. Queueing models are described in terms of service disciplines and stochastic primitives, where the latter term usually refers to arrival processes, service times, routing and other processes. The way in which we propose using the RRB based approach in the queueing context is by working with Rényi divergence estimates for the underlying primitives rather than directly with the “state” processes that are used in describing performance criteria (such as queue lengths, delay, idleness). This is particularly natural when one views such models as dynamical systems driven by renewal processes or more general counting processes (or yet more generally, as marked counting processes). To demonstrate this point, we provide two examples.

First, consider the queue length process X_t for a GI/GI/1 queue. In this single server queue, arrivals follow a renewal process, denoted by A_t , and service times are IID. Let S_t denote the potential service process: By the time the server is busy for t units of time, S_t jobs have departed. Assuming here, for simplicity, that at time zero the server has no residual work, S_t is also a renewal process. The queue length satisfies the equations

$$X_t = X_0 + A_t - S(T_t), \quad T_t = \int_0^t 1_{\{X_s > 0\}} ds.$$

For our purpose, the key property is that $X|_{[0,t]}$ is fully determined by its initial condition and the primitives $A|_{[0,t]}$, $S|_{[0,t]}$ (this owes to the fact $T_t \leq t$ for all t). Hence, if such a queue is to be analyzed by comparison to M/M/1, the relevant Rényi divergence term dictating events measurable w.r.t. $X|_{[0,t]}$, is $R_\alpha(Q_t \| P_t)$, where Q_t is the law of $(A, S)|_{[0,t]}$ as a pair of (independent) renewal processes and P_t as a pair of Poisson processes.

Next consider a generalized Jackson network. This is a network of N service stations, each having an external (possibly void) stream of arrivals, and upon departure from a service station, jobs are routed probabilistically, according to a given substochastic $N \times N$ matrix, to one of the service stations or to leave the system. Let $\{\xi_i(k)\}$ be $\{0, e_1, \dots, e_N\}$ -valued RVs according to which these routings are determined: $\xi_{ij}(k) = 1$ dictates that the k th i -departure is routed to station j . All arrival processes and potential service processes are assumed to be mutually independent renewals, that are also independent of the routing decision variables, ξ . Denote by X_i , E_i , S_i and D_i the queue length, external arrival, potential service and departure processes, associated with service station i for $1 \leq i \leq N$. Let also D_{ij} denote the counting process of jobs departing from station i and routed back to station j . Finally, let A_i denote the counting process for total arrivals into station i , including external arrivals and reroutings. Then the following equations are satisfied:

$$X_i = X_i(0) + A_i - D_i = X_i(0) + A_i - S_i \circ T_i,$$

$$A_i = E_i + \sum_j D_{ji}, \quad T_i = \int_0^\cdot 1_{\{X_i > 0\}} dt,$$

$$D_{ij} = R_{ij} \circ D_i, \quad R_{ij} = \sum_{k=1}^{\cdot} \xi_{ij}(k).$$

Thus the dependence of queue length on the stochastic primitives is far more complicated than in the case of a single node. Yet, the key property alluded to above is valid in this complicated scenario. That is, $X_i|_{[0, t]}$, $1 \leq i \leq N$, are dictated by their initial condition and the primitives $E_i|_{[0, t]}$, $S_i|_{[0, t]}$ and $(R_{ij} \circ S_i)|_{[0, t]}$, $1 \leq i, j \leq N$. A similar statement is valid for the busyness processes T_i , the counting processes D_{ij} , etc.

The special case where E_i and S_i are Poisson is referred to as a Jackson network. In this case, the queue length is a Markov process with state space \mathbb{Z}_+^N , and is far easier to analyze than the non-Markovian model. The perturbation that is required for translating results on the Markovian model to the more general one again has to do with a change of measure from a Poisson to a renewal process. Once again, the perturbation can be expressed as a Rényi divergence term, this time for each E_i and S_i , $1 \leq i \leq N$. A term that takes into account the routings $(R_{ij} \circ S_i)|_{[0, t]}$ is required only if perturbations of the routing matrix are considered. This is important as far as robust performance bounds are concerned, but note that it is not required for turning a non-Markovian model into Markovian.

As mentioned above, the Markovian model is easier to handle than the non-Markovian one. As far as LD results are concerned, the full LDP for the former was established via a general approach by Dupuis and Ellis, as a special case of a large class of Markovian queueing models [14]. Building on these results, the identification of the rate function was obtained by [4] and by [22]. Expressions for the rate function in these two references were provided as a finite-dimensional convex optimization problem, and as a recursive formula, respectively. Denoting a rescaled version of X by $X^n = n^{-1}X(n\cdot)$, and letting P stand for the probability measure that makes the primitives E_i and S_i independent Poisson processes, the LDP for the Jackson network provides an upper bound in the form of a variational formula for the asymptotic expression

$$\gamma(P, F) = \limsup_{n \rightarrow \infty} \frac{1}{n} \log P(X^n|_{[0, 1]} \in F),$$

where F is any closed (in the J_1 topology) set of paths mapping $[0, 1]$ to \mathbb{R}_+^N . A typical set of interest is $F = \{\phi : \phi(t) \in M \text{ for some } t \in [0, 1]\}$, $M = \mathbb{R}_+^N - \prod_{i=1}^N [0, b_i)$, expressing the buffer overflow event: one of the queues X_i^n exceeds a threshold b_i some time during $[0, 1]$, an event that is rare for large n provided that the network is stable.

LDP is known also for the generalized (i.e., non-Markovian) Jackson network by [36], where the rate function is identified in terms of an optimization problem, that is not in general a convex optimization problem, and for which a recursive formula such as [22] is not available. Considerably less is known on estimates at this scale which hold for a generalized Jackson network *uniformly* w.r.t. the stochastic primitives within certain set. However, (2.8) addresses precisely this question.

Indeed, to state the readily available corollary of (2.8), set

$$r_\alpha(Q||P) = \sum_{i=1}^N r_\alpha^{E_i}(Q||P) + \sum_{i=1}^N r_\alpha^{S_i}(Q||P)$$

to be the sum of RDR over all primitive processes. Then given a collection Q , we have the following uniform bound on generalized Jackson networks associated with $Q \in \mathcal{Q}$ in terms of performance of the Jackson network P , namely

$$(2.9) \quad \sup_{Q \in \mathcal{Q}} \gamma(Q, F) \leq \inf_{\alpha > 1} \left\{ \frac{\alpha - 1}{\alpha} \gamma(P, F) + (\alpha - 1)r_\alpha(Q||P) \right\}.$$

As already mentioned, by the LDP, an upper bound on $\gamma(P, F)$ is known, in the form of a variational formula. Therefore the usefulness of (2.9) depends on the ability to compute or provide an effective bound also on the last term, that is, the RDR of a renewal process w.r.t. a Poisson.

The case made above for the crucial importance of RDR estimates for the applicability of the approach can be made in any scenario where a queueing model is representable as a dynamical system driven by renewal processes or other counting processes, and in the special case of Poisson driving processes is tractable (due to Markovity or for any other reason). Therefore the usefulness of studying the RDR in relation to the proposed approach is broad.

3. Results on RDR. Calculations and bounds of entropy rate and Rényi entropy rate have been studied for some families of stochastic processes, including Markov chains and hidden Markov models [24, 33]. However, the questions that arise from the above discussion are concerned with the RDR of marked point processes with respect to marked Poisson point processes, also known as a Poisson random measure. To the best of our knowledge, estimates on RDR for such models have not been studied before. In this section we present some results in this direction. The marks of the reference Poisson point process we consider will take values in some Polish space S and will have IID distributions given by some probability measure ζ on $(S, \mathcal{B}(S))$, where $\mathcal{B}(S)$ denotes the Borel σ -field on S . Denote by $\mathcal{M}_F(S)$ the space of finite measures on S equipped with the usual weak convergence topology. A marked point process can be represented as a stochastic process $\{N_t\}$ with sample paths in $\Omega = \mathcal{D}([0, \infty) : \mathcal{M}_F(S))$. A rate λ_0 Poisson marked point process with mark distribution ζ is a stochastic process $\{N_t\}$ such that (i) for all $0 \leq s < t < \infty$ and $A \in \mathcal{B}(S)$, $N_t(A) - N_s(A)$ is a Poisson random variable with mean $\lambda_0(t - s)\zeta(A)$; (ii) if for $k \in \mathbb{N}$, $0 \leq s_i < t_i < \infty$ and $A_i \in \mathcal{B}(S)$, $i = 1, 2, \dots, k$, $\{(s_i, t_i] \times A_i\} \cap \{(s_j, t_j] \times A_j\} = \emptyset$, for all $1 \leq i < j \leq k$, then the Poisson random variables $\{N_{t_i}(A_i) - N_{s_i}(A_i), 1 \leq i \leq k\}$ are mutually independent. We will refer to $\lambda_0 ds \times \zeta(dz)$ as the intensity measure of such a marked Poisson point process.

Let $\mathcal{F} = \mathcal{B}(\Omega)$ and abusing notation, let $\{N_t\}_{t \geq 0}$ be the canonical coordinate process on (Ω, \mathcal{F}) . Fix $\lambda_0 \in (0, \infty)$ and $\zeta \in \mathcal{P}(S)$. Let P be the unique probability measure on (Ω, \mathcal{F}) under which N is a rate λ_0 marked Poisson process with mark distribution ζ . We will consider the canonical filtration on (Ω, \mathcal{F}) which will be denoted as $\{\mathcal{F}_t\}_{t \geq 0}$.

In this section we present two types of results: bounds on $r_\alpha(Q\|P)$ for Q a family of models (namely, probability measures on (Ω, \mathcal{F})) and on $r_\alpha(Q\|P)$ for a single model. This is the content of Sections 3.1 and 3.2, respectively. The proofs of the results stated in these two sections appear in Appendix A.1 and A.2, respectively.

3.1. Bounds on RDR for families of processes. This section provides bounds on $r_\alpha(Q\|P)$ for families Q of probability laws of counting processes where as before P is the probability law of a marked Poisson process with intensity measure $\lambda_0 ds \times \zeta(dz)$.

For $x \geq 0$ and $\alpha > 1$ let $k_\alpha(x)$ denote the Rényi divergence of order α of a Poisson RV with parameter $x \in (0, \infty)$ w.r.t. a Poisson RV with parameter 1. A direct computation gives

$$(3.1) \quad k_\alpha(x) = \frac{x^\alpha - \alpha x + \alpha - 1}{\alpha(\alpha - 1)}.$$

Note that for every α , this function is nonnegative, strictly convex and vanishes uniquely at 1.

Denote by \mathcal{V}_0 the set of mappings $v : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $v(x) \rightarrow 0$ as $x \rightarrow \infty$. Also, denote by $\mathcal{P}\mathcal{F}$ the predictable σ -field on $\Omega \times \mathbb{R}_+$ and let $\lambda : \Omega \times \mathbb{R}_+ \times S \rightarrow (0, \infty)$ be a $\mathcal{P}\mathcal{F} \otimes \mathcal{B}(S) \setminus \mathcal{B}(0, \infty)$ measurable map. We will refer to such a map as a predictable process. We will consider probability measures Q on (Ω, \mathcal{F}) under which $\{N_t\}$ is a marked point

process with intensity process λ . Such a probability measure can be characterized as the unique element of $\mathcal{P}(\Omega)$ under which for every bounded predictable process u and $T < \infty$

$$E_P \left[\int_{[0, T] \times S} u(s, z) N(ds dz) \right] = E_P \left[\int_{[0, T] \times S} u(s, z) \lambda(s, z) ds \zeta(dz) \right],$$

where we view N as a $\mathcal{M}_F([0, T] \times S)$ -valued random variable which is defined by the relation $N((s, t] \times A) \doteq N_t(A) - N_s(A)$ for $0 \leq s < t \leq T, A \in \mathcal{B}(S)$. We will be particularly interested in the case where $\lambda(s, x) = \hat{\lambda}(s)\psi(x)$, where $\hat{\lambda} : \Omega \times \mathbb{R}_+ \rightarrow (0, \infty)$ is a $\mathcal{P}\mathcal{F} \setminus \mathcal{B}(0, \infty)$ -measurable map (also referred to as a predictable process) and $\psi : S \rightarrow (0, \infty)$ is a $\mathcal{B}(S) \setminus \mathcal{B}(0, \infty)$ measurable map satisfying $\int_S \psi(z) \zeta(dz) = 1$. This corresponds to the setting in which, under Q , $\{N_t\}$ is a marked point process with points having IID distribution $\tilde{\zeta}(dz) = \psi(z)\zeta(dz)$. In such a case, we will refer to (N, λ) as a marked Cox process.

In the special case where S is a singleton $\{z^*\}$ (and so $\zeta(dz) = \tilde{\zeta}(dz) = \delta_{z^*}(dz)$), $\{N_t\}$ is simply a Cox process with intensity process $\hat{\lambda}(\cdot)$ (see, e.g., [26], Section 1.1). In such a case we will occasionally also refer to $(N, \hat{\lambda})$ as a Cox process. For a probability measure Q on (Ω, \mathcal{F}) and $t \in [0, \infty)$, let $Q_t^N \doteq Q \circ N_{[0, t]}^{-1}$ be the probability measure induced on $\mathcal{D}([0, t] : \mathcal{M}_F(S))$ by the canonical coordinate process. Motivation for the specific forms of the families \mathcal{Q}_i considered in the theorem appears after the theorem statement.

THEOREM 3.1.

(i) Fix $v \in \mathcal{V}_0$ and $\alpha > 1$. Consider the collection \mathcal{Q}_1 of probability measures Q on (Ω, \mathcal{F}) under which (N, λ) is a marked point process with intensity process λ satisfying

$$(3.2) \quad T^{-1} \int_0^T \int_S k_\alpha \left(\frac{\lambda(t, z)}{\lambda_0} \right) \zeta(dz) dt \leq u + v(T), \quad T > 0,$$

for some constant $u \geq 0$. Then

$$r_\alpha(\mathcal{Q}_1 \| P) = \limsup_{t \rightarrow \infty} \frac{1}{t} \sup_{Q \in \mathcal{Q}_1} R_\alpha(Q_t^N \| P_t^N) = u\lambda_0.$$

(ii) Consider the collection \mathcal{Q}_2 of probability measures Q on (Ω, \mathcal{F}) under which (N, λ) is a marked point process with intensity process λ satisfying

$$(3.3) \quad a \leq \frac{\lambda(\cdot)}{\lambda_0} \leq b$$

for constants $0 \leq a \leq 1 \leq b < \infty$. Then

$$(3.4) \quad r_\alpha(\mathcal{Q}_2 \| P) = (k_\alpha(a) \vee k_\alpha(b))\lambda_0.$$

As a special case, the identity holds for the family of measures under which the marks are IID (and independent of jump instants) with distribution $\psi(z)\zeta(dz)$ and $\{N_t(S)\}$ is a delayed renewal processes with hazard rate h (i.e., $(N, h(s)\psi(z))$ is a marked Cox process), and $a \leq \frac{h(s)\psi(z)}{\lambda_0} \leq b$ for all $(s, z) \in \mathbb{R}_+ \times S$.

(iii) Let $v \in \mathcal{V}_0$. Consider the collection \mathcal{Q}_3 of probability measures Q under which (N, λ) is a marked point process with intensity process λ satisfying (3.3) as well as

$$(3.5) \quad \lambda_0 - v(T) \leq \frac{1}{T} \int_{[0, T] \times S} \lambda(t, z) \zeta(dz) dt \leq \lambda_0 + v(T), \quad T > 0.$$

Then

$$r_\alpha(\mathcal{Q}_3 \| P) = (pk_\alpha(a) + qk_\alpha(b))\lambda_0,$$

where $p = \frac{b-1}{b-a}$ and $q = \frac{1-a}{b-a}$.




FIG. 1. RDR for various families as a function of α . Left: $r_\alpha(Q_2\|P)$ and $r_\alpha(Q_3\|P)$ for different values of a and b . Right: $r_\alpha(Q_4\|P)$ for different constraint pairs (α_0, u) . In all cases, $\lambda_0 = 1$.

(iv) Let $v \in \mathcal{V}_0$ and consider a collection \mathcal{Q}_4 of probability measures under which (N, λ) is a marked point process with intensity process λ satisfying (3.2) for some $\alpha = \alpha_0$, as well as (3.5). Then for all $\alpha \in (1, \alpha_0)$,

$$r_\alpha(Q_4\|P) = \frac{\lambda_0}{\bar{\alpha}} [(\bar{\alpha}_0 u + 1)^{\frac{\alpha-1}{\alpha_0-1}} - 1] = (pk_\alpha(0) + qk_\alpha(c))\lambda_0,$$

where $\bar{\alpha} = \alpha(\alpha - 1)$, $\bar{\alpha}_0 = \alpha_0(\alpha_0 - 1)$, $p = 1 - q$, $q = (\bar{\alpha}_0 u + 1)^{-\frac{1}{\alpha_0-1}}$ and $c = (\bar{\alpha}_0 u + 1)^{\frac{1}{\alpha_0-1}}$.

The proof appears in Section A.1. Figure 1 provides several numerical evaluations of the RDR for families analyzed in Theorem 3.1.

REMARK 3.1. Items (i) and (ii) of the result are concerned with classes of marked point processes for which a certain constraint is put on the size of perturbation of the stochastic intensity. Note that for $\alpha = 2$, the left hand side of (3.2) gives half the second moment centered about 1 of the empirical distribution

$$\frac{1}{T} \int_{[0, T] \times S} \delta_{\lambda(t, z)/\lambda_0} \zeta(dz) dt.$$

For other values of α , it provides different types of level dispersion about 1 that take the form of higher order moments. In the same vein, (3.3) can be seen as a constraint on the L_∞ norm of the same empirical distribution, centered about $\frac{a+b}{2}$.

The motivation behind parts (iii) and (iv) is that the reference value λ_0 may play an additional role. If a parameter is regarded as a first order approximation, it may often mean that over a long period of time it represents the true average. Clearly, this additional constraint makes the class of models smaller than the classes from items (i) and (ii), and leads to tighter bounds.

3.2. *Bounds on RDR for a single renewal process.* Recall that P is the unique probability measure on (Ω, \mathcal{F}) under which the canonical coordinate process N is a marked Poisson process with rate λ_0 and mark distribution ζ . In this section, for simplicity, we take $\lambda_0 = 1$. Let Q be another probability measure on (Ω, \mathcal{F}) under which N is a marked renewal process with mark distribution $\tilde{\zeta}(dz) = \psi(z)\zeta(dz)$ and inter-jump distribution π . Note that in such a process the collection of jump-instants is independent of the collection of marks,

and inter-jump times and marks are IID. We denote such a process as a (π, ψ) -marked renewal process. Assume that π has a density denoted by g , and let h denote the hazard rate, $h(x) \doteq g(x)/\pi[x, \infty)$, with $h(x) = 0$ if $\pi[x, \infty) = 0$. Define, for $x \geq 0$,

$$(3.6) \quad H(x) \doteq \int_0^x (1 - h(s)) ds + \log h(x) = x + \log g(x),$$

with $H(x) = -\infty$ when $g(x) = 0$.

To state the next result, let

$$\begin{aligned} \gamma(s) &\doteq \int e^{sH(y)} \nu(dy), \quad s \in \mathbb{R}, \\ \hat{\gamma}(p, q, \alpha) &\doteq \frac{\gamma(q\alpha)^{p/q} - 1}{p}, \quad p, q \geq 1, \end{aligned}$$

where ν is the standard exponential distribution. Denote

$$G_\alpha^{(1)} \doteq \inf_{p, q \geq 1: p^{-1} + q^{-1} = 1} \hat{\gamma}(p, q, \alpha).$$

Also, let

$$\beta(\lambda) \doteq \log \int e^{\lambda_1 y + \lambda_2 H(y)} \nu(dy), \quad \lambda = (\lambda_1, \lambda_2) \in \mathbb{R}^2,$$

and let β^* be the Legendre–Fenchel transform:

$$\beta^*(x) = \sup_{\lambda} \{ \langle \lambda, x \rangle - \beta(\lambda) \}, \quad x \in \mathbb{R}^2.$$

For $\theta \in (0, \infty)$, denote

$$\begin{aligned} G_\alpha^{(2)}(\theta) &\doteq \theta \sup_{x \in \mathbb{R}^2: x_1 \leq \theta^{-1}} [\alpha x_2 - \beta^*(x)], \\ G_\alpha^{(3)}(\theta) &\doteq \theta \sup_{x_2 \in \mathbb{R}} [\alpha x_2 - \beta^*(\theta^{-1}, x_2)]. \end{aligned}$$

Recall that $r_\alpha^N(Q \| P) = \limsup_{t \rightarrow \infty} \frac{1}{t} R_\alpha(Q_t^N \| P_t^N)$. For $z \in \mathbb{R}$, we denote $z^+ = 0 \vee z$. Then we have the following upper bounds.

THEOREM 3.2. *Assume that $\bar{H} \doteq \sup_{x \in \mathbb{R}_+} H(x) < \infty$. Also suppose that $c(\alpha) \doteq \int_S (\psi^\alpha(z) - 1) \zeta(dz) < \infty$. Then the following hold for $\alpha > 1$.*

(a)

$$(3.7) \quad r_\alpha^N(Q \| P) \leq \frac{e^{\alpha \bar{H}} - 1 + c(\alpha)}{\alpha(\alpha - 1)}.$$

(b)

$$(3.8) \quad r_\alpha^N(Q \| P) \leq \frac{[G_\alpha^{(1)}]^+ + c(\alpha)}{\alpha(\alpha - 1)}.$$

(c) *If β is finite in a neighborhood of the origin, then*

$$(3.9) \quad r_\alpha^N(Q \| P) \leq \frac{[\sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta)]^+ + c(\alpha)}{\alpha(\alpha - 1)}.$$

(d) *If β is finite in a neighborhood of the origin and $\gamma(s) < \infty$ for all $s \leq 0$, then*

$$(3.10) \quad r_\alpha^N(Q \| P) \leq \frac{[\sup_{\theta \in (0, \infty)} G_\alpha^{(3)}(\theta)]^+ + c(\alpha)}{\alpha(\alpha - 1)}.$$

The proof of this result appears in Section A.2.

REMARK 3.2. We now make some comments on the assumption and behavior of different bounds in Theorem 3.2.

- For (3.9) and (3.10), the assumption that β is finite in a neighborhood of the origin is needed to apply the strengthened Cramér’s theorem [12], Corollary 6.1.6, and Varadhan’s integral lemma [12], Lemma 4.3.6. Unfortunately, if the support of π is not \mathbb{R}_+ , then H is $-\infty$ at some place and β is not always finite around the origin.
- For (3.10), the assumption that $\gamma(s) < \infty$ for all $s \leq 0$ (together with the requirement that $\bar{H} < \infty$) rules out the possibility that π is Exponential(ρ) for $\rho \neq 1$.
- Since $G_\alpha^{(3)}(\theta) \leq G_\alpha^{(2)}(\theta)$, the bound in (3.10) is clearly better than the bound in (3.9) (though the former requires stronger assumptions). Also, the bounds in (3.8), (3.9) and (3.10) are all better than the rough bound (3.7). This can be seen as follows.

For (3.8), since $\gamma(s) \leq e^{s\bar{H}}$, we have $\hat{\gamma}(p, q, \alpha) \leq \frac{e^{p\alpha\bar{H}} - 1}{p}$. Taking $p \rightarrow 1$ gives $G_\alpha^{(1)} \leq e^{\alpha\bar{H}} - 1$.

For (3.9), note that for fixed $\theta \in (0, \infty)$,

$$\begin{aligned}
 (3.11) \quad G_\alpha^{(2)}(\theta) &= \theta \sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta^{-1}} [\alpha x_2 - \beta^*(x)] \\
 &= \theta \sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta^{-1}} \left[\alpha x_2 - \sup_{\lambda \in \mathbb{R}^2} \{ \langle \lambda, x \rangle - \beta(\lambda) \} \right] \\
 &= \theta \sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta^{-1}} \inf_{\lambda \in \mathbb{R}^2} [\alpha x_2 - \lambda_1 x_1 - \lambda_2 x_2 + \beta(\lambda)] \\
 &\leq \theta \inf_{\lambda \in \mathbb{R}^2} \sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta^{-1}} [\alpha x_2 - \lambda_1 x_1 - \lambda_2 x_2 + \beta(\lambda)] \\
 &= \theta \inf_{\lambda_1 \in \mathbb{R}} \sup_{0 \leq x_1 \leq \theta^{-1}} [-\lambda_1 x_1 + \beta(\lambda_1, \alpha)], \\
 &= \inf_{\lambda_1 \in \mathbb{R}} [(\lambda_1)^- + \theta \beta(\lambda_1, \alpha)],
 \end{aligned}$$

where the fifth line follows on observing that $\sup_{x_2 \in \mathbb{R}} [\alpha x_2 - \lambda_1 x_1 - \lambda_2 x_2 + \beta(\lambda)] = \infty$ when $\lambda_2 \neq \alpha$. If $0 < \theta < 1$, taking $\lambda_1 = 0$ in (3.11) we have

$$G_\alpha^{(2)}(\theta) \leq \theta \beta(0, \alpha) \leq \alpha \bar{H} \leq e^{\alpha \bar{H}} - 1.$$

If $\theta \geq 1$, taking $\lambda_1 = 1 - \theta \leq 0$ in (3.11) we have

$$\begin{aligned}
 G_\alpha^{(2)}(\theta) &\leq -\lambda_1 + \theta \beta(\lambda_1, \alpha) = \theta - 1 + \theta \log E_P [e^{(1-\theta)\Delta + \alpha H(\Delta)}] \\
 &\leq \theta - 1 + \alpha \bar{H} \theta - \theta \log \theta \leq e^{\alpha \bar{H}} - 1,
 \end{aligned}$$

where the last inequality becomes equality when $\theta = e^{\alpha \bar{H}}$. Therefore $\sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta) \leq e^{\alpha \bar{H}} - 1$ and (3.9) is better than (3.7). Finally, since (3.10) is better than (3.9), it is also better than (3.7).

3.3. *Examples.* We now consider a few specific cases of π . For simplicity these examples are concerned with point processes without marks (namely, the case where $\zeta(dz) = \delta_{z^*}(dz)$ and hence $c(\alpha) = 0$ in Theorem 3.2). The first example is that of an exponential distribution.

EXAMPLE 3.1. Suppose $\pi = \text{Exp}(\rho)$ with rate $\rho > 1$, namely $g(x) = \rho e^{-\rho x}$. It turns out that in this case the right sides of (3.9) and (3.10) are the same and in fact the inequalities in both cases can be replaced by equalities. Note that in this example $\gamma(s) = \infty$ for $s \leq -\frac{1}{\rho-1}$, which violates the assumption required for (3.10) in part (d). Actually in (3.9) and (3.10) the inequality can be changed to equality even for the case $\rho \in (0, 1]$. However, we note that $H(x) = -(\rho - 1)x + \log \rho \rightarrow \infty$ as $x \rightarrow \infty$ when $\rho \in (0, 1]$, which violates the assumption $\bar{H} < \infty$ required for Theorem 3.2. Proofs of the above statements are given in Appendix A.3. This example shows that the conditions assumed in Theorem 3.2 are not essential for the result.

The second example is Gamma(k, ρ).

EXAMPLE 3.2. Suppose $\pi = \text{Gamma}(k, \rho)$ with $k \geq 1$ and $\rho > 1$, namely $g(x) = \frac{\rho^k}{\Gamma(k)} x^{k-1} e^{-\rho x}$, for $x \geq 0$. For this example, computing an explicit expression for the Rényi divergence is harder and thus we will make use of the bounds in Theorem 3.2. Since $\rho > 1$ and $k \geq 1$, $e^{H(x)} = g(x)e^x$ is bounded from above. Also for $\lambda = (\lambda_1, \lambda_2)$ in a sufficiently small neighborhood of the origin,

$$\begin{aligned} \beta(\lambda) &= \log \int e^{\lambda_1 x} g(x)^{\lambda_2} e^{\lambda_2 x} e^{-x} dx \\ &= \log \int \left(\frac{\rho^k}{\Gamma(k)}\right)^{\lambda_2} x^{\lambda_2(k-1)} e^{-(\lambda_2 \rho - \lambda_2 - \lambda_1 + 1)x} dx \\ &= \log \left[\left(\frac{\rho^k}{\Gamma(k)}\right)^{\lambda_2} \frac{\Gamma(1 + \lambda_2(k - 1))}{(1 + \lambda_2(\rho - 1) - \lambda_1)^{1 + \lambda_2(k-1)}} \right] < \infty. \end{aligned}$$

So all assumptions for (3.9) hold. Note, however, that assumptions for (3.10) are not satisfied since $\gamma(s) = \infty$ for $s \leq -(\rho - 1)^{-1}$. Using Theorem 3.2(c) one can give the following explicit bound for the Rényi divergence rate by estimating $\sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta)$:

$$\begin{aligned} (3.12) \quad r_\alpha^N(Q \| P) &\leq \frac{1}{\alpha(\alpha - 1)} \sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta) \\ &\leq \frac{1}{\alpha(\alpha - 1)} \left[\left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1 + \alpha(k-1)}} - \alpha(\rho - 1) - 1 \right]. \end{aligned}$$

Details of this calculation are given in Appendix A.3. When $k = 1$, namely when π is $\text{Exp}(\rho)$, the bound on the right side equals $\rho^\alpha - \alpha(\rho - 1) - 1$ and the inequalities in the above display are in fact equalities.

The next example can be used to obtain RDR bounds for certain types of phase-type distributions.

EXAMPLE 3.3. Suppose the density $g(x) \leq C e^{-\sigma x}$ for some $\sigma > 1$. In this case, once again, the assumptions for (3.10) are not satisfied in general. However, as we check below the assumptions for (3.9) hold. First note that $e^{H(x)} = g(x)e^x \leq C e^{-(\sigma-1)x}$ and since $\sigma > 1$, $\bar{H} < \infty$.

Next note that for $\lambda = (\lambda_1, \lambda_2)$ such that $1 + \lambda_2(\sigma - 1) - \lambda_1 > 0$, we have

$$\begin{aligned} \beta(\lambda) &= \log \int e^{\lambda_1 x} g(x)^{\lambda_2} e^{\lambda_2 x} e^{-x} dx \\ &\leq \log \int C^{\lambda_2} e^{-(\lambda_2 \sigma - \lambda_2 - \lambda_1 + 1)x} dx \end{aligned}$$

$$\begin{aligned}
 &= \lambda_2 \log C + \log \frac{1}{1 + \lambda_2(\sigma - 1) - \lambda_1} \\
 &< \infty.
 \end{aligned}$$

Thus we have verified that all the assumptions needed for (3.9) are satisfied. Using Theorem 3.2(c) one can give the following simple form bound for the quantity on the right side of (3.9).

$$(3.13) \quad r_\alpha^N(Q\|P) \leq \frac{1}{\alpha(\alpha - 1)} \sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta)^+ \leq \frac{1}{\alpha(\alpha - 1)} [C^\alpha - 1 - \alpha(\sigma - 1)].$$

Details of this calculation are given in Appendix A.3.

4. Robust control of tail properties for a scheduling problem. When considering ordinary cost structures the variational representation for exponential integrals in terms of relative entropy is the starting point for a formulation of optimization and control design that is robust with respect to model errors, where errors are measured by relative entropy distances. To be precise, one can formulate problems such that their solution gives the tightest possible bounds on a given performance measure for a family of models, where the family is defined by a relative entropy distance to a design model [34, 35]. Alternatively, one can fix a desired performance bound, and find the control which gives the largest possible family of models across which the performance criteria is guaranteed to hold. In this section we investigate analogous situations where in place of ordinary cost structures we use costs that are determined by rare events, that is, risk-sensitive costs. Let $(\mathcal{S}, \mathcal{F})$ and \mathcal{P} be as in Section 2.1.

4.1. *A general approach.* Let $g : \mathcal{S} \rightarrow \mathbb{R}$ be bounded and measurable. From the identity

$$\frac{1}{\alpha} \log \int e^{\alpha g} dP = \sup_{Q \in \mathcal{P}} \left[\frac{1}{\alpha - 1} \log \int e^{(\alpha - 1)g} dQ - R_\alpha(Q\|P) \right],$$

valid for $\alpha > 1$, one can easily obtain, for all $0 < \beta < \gamma$,

$$(4.1) \quad \frac{1}{\gamma} \log \int e^{\gamma g} dP = \sup_{Q \in \mathcal{P}} \left[\frac{1}{\beta} \log \int e^{\beta g} dQ - \frac{1}{\gamma - \beta} R_{\frac{\gamma}{\gamma - \beta}}(Q\|P) \right]$$

(extensions to unbounded g are also possible). Fix some risk sensitivity parameter β and a class of models \mathcal{Q} . Let

$$(4.2) \quad f(\alpha) = \sup\{R_\alpha(Q\|P) : Q \in \mathcal{Q}\}, \quad \alpha \in (1, \infty).$$

By appealing to (4.1) we can show the following.

THEOREM 4.1. *Fix $\beta > 0$, g , P , \mathcal{Q} and f as above. Then*

$$\sup_{Q \in \mathcal{Q}} \frac{1}{\beta} \log E_Q e^{\beta g} \leq \inf_{\gamma > \beta} F(\beta, \gamma) \quad \text{where } F(\beta, \gamma) = \left[\frac{f(\frac{\gamma}{\gamma - \beta})}{\gamma - \beta} + \frac{1}{\gamma} \log E_P e^{\gamma g} \right].$$

Moreover, $\tilde{\gamma} \mapsto F(\beta, 1/\tilde{\gamma})$ is a convex function.

We propose using Theorem 4.1 as the basis for the formulation of optimization and control problems, so that given a class \mathcal{Q} , an upper bound is obtained across \mathcal{Q} for a RS control (or optimization) problem. For certain problems we expect to be able to say more, which is that the bound is tight in some sense. A controlled process X is considered with cost of the form $E[e^{\beta g(X)}]$, where β is the sensitivity parameter. The RS control problem will be concerned with $E^a[e^{\beta g(X)}]$, where a denotes a control or a parameter to be optimized over, and the

robust version of this problem is one where a control a is sought to minimize the RS cost uniformly in the family of models. In this context, Theorem 4.1 gives

$$(4.3) \quad \inf_a \sup_{Q \in \mathcal{Q}} \frac{1}{\beta} \log E_Q^a e^{\beta g} \leq \inf_a \inf_{\gamma > \beta} \left[\frac{f(\frac{\gamma}{\gamma-\beta})}{\gamma-\beta} + \frac{1}{\gamma} \log E_P^a e^{\gamma g} \right].$$

Two highly attractive aspects of this bound are:

- (i) it turns an ∞ -dimensional game into a finite dimensional minimization problem when a is finite dimensional;
- (ii) the minimization over γ is tractable computationally, thanks to the convexity stated in Theorem 4.1.

With regard to the optimization over a , that is of course related to the structure of the particular problem. However, it is worth noting that the difficulty of this problem is often related to the difficulty of the ordinary analogue, that is, $\inf_a E_P^a g$. For the example from queueing presented below we see that the risk-sensitive optimization problem has a structure that is very similar to that of the ordinary analogue. The function f is the element that distinguishes this problem from its relative entropy/ordinary cost analogue, for which f is essentially a constant. In some sense, f captures the critical, distribution dependent properties of the tail behavior of \mathcal{Q} . The only part of Theorem 4.1 that does not follow directly from (4.1) is the last sentence, which we now address.

LEMMA 4.1. *Let X be a nonzero nonnegative random variable. Then the function*

$$m(\theta) \doteq \theta \log EX^{1/\theta}$$

is convex in $\theta > 0$.

PROOF. It suffices to show that

$$m(\theta) \leq \lambda m(\theta_1) + (1 - \lambda)m(\theta_2)$$

for every $\theta_1, \theta_2 \in (0, \infty)$ and $\lambda \in (0, 1)$, where $\theta \doteq \lambda\theta_1 + (1 - \lambda)\theta_2$. Assume without loss of generality that $m(\theta_1) < \infty, m(\theta_2) < \infty$. Applying Hölder’s inequality with $p = \frac{\theta}{\lambda\theta_1}$ and $q = \frac{\theta}{(1-\lambda)\theta_2}$, we have

$$\begin{aligned} m(\theta) &= \theta \log E[X^{\frac{\lambda}{\theta}} X^{\frac{1-\lambda}{\theta}}] \leq \theta \log[(EX^{\frac{\lambda}{\theta} p})^{1/p} (EX^{\frac{1-\lambda}{\theta} q})^{1/q}] \\ &= \frac{\theta}{p} \log EX^{1/\theta_1} + \frac{\theta}{q} \log EX^{1/\theta_2} = \lambda m(\theta_1) + (1 - \lambda)m(\theta_2). \end{aligned}$$

This completes the proof. \square

LEMMA 4.2. *For $0 < \beta < \gamma$, let*

$$h(\gamma) \doteq \frac{f(\frac{\gamma}{\gamma-\beta})}{\gamma-\beta} + \frac{1}{\gamma} \log E_P e^{\gamma g},$$

where $f(\alpha)$ is defined by (4.2). Then the function $\tilde{h}(\tilde{\gamma}) \doteq h(1/\tilde{\gamma})$ is convex in $\tilde{\gamma} \in (0, 1/\beta)$, that is, $h(\gamma)$ is convex in $1/\gamma$.

PROOF. Since

$$R_\alpha(Q \| P) = \frac{1}{\alpha(\alpha - 1)} \log \int \left(\frac{dQ}{dP} \right)^\alpha dP,$$

we can write

$$\begin{aligned} \frac{1}{\gamma - \beta} R_{\frac{\gamma}{\gamma-\beta}}(Q \| P) &= \frac{1}{\gamma - \beta} \frac{1}{\frac{\gamma}{\gamma-\beta} (\frac{\gamma}{\gamma-\beta} - 1)} \log \int \left(\frac{dQ}{dP} \right)^{\frac{\gamma}{\gamma-\beta}} dP \\ &= \frac{1}{\beta} \frac{1}{\frac{\gamma}{\gamma-\beta}} \log \int \left(\frac{dQ}{dP} \right)^{\frac{\gamma}{\gamma-\beta}} dP. \end{aligned}$$

Therefore

$$\tilde{h}(\tilde{\gamma}) = h(1/\tilde{\gamma}) = \sup_{Q \in \mathcal{Q}} \left[\frac{1}{\beta} (1 - \beta \tilde{\gamma}) \log \int \left(\frac{dQ}{dP} \right)^{1/(1-\beta \tilde{\gamma})} dP + \tilde{\gamma} \log E_P([e^g]^{1/\tilde{\gamma}}) \right].$$

From Lemma 4.1 we see that the last term is convex in $\tilde{\gamma}$. Since $1 - \beta \tilde{\gamma}$ is just an affine function of $\tilde{\gamma}$, it follows from Lemma 4.1 again that the first term is also convex in $\tilde{\gamma}$. This completes the proof. \square

4.2. *A risk-sensitive scheduling control problem.* We focus on one out of various RS control problems that are of interest in the multiclass $G/G/1$ setting. In this setting each arrival requires a single service. A recurring theme in the literature is how to schedule service so as to minimize delay or queue length costs. The need to cover general service time distributions has been recognized many times in earlier work on this model. However, under RS cost, this question has only been addressed in the Markovian setting. Our goal here is to show how the perturbation bounds, specifically Theorem 4.1, can be used to yield performance guarantees for the non-Markovian setting.

Let N denote the number of classes, and for $i \in \{1, \dots, N\}$ let X_i, A_i and S_i denote the i th queue length process, arrival process and potential service process. Then for each i , the balance equation holds, namely

$$X_i(t) = X_i(0) + A_i(t) - S_i(U_i(t)),$$

where $U_i(t)$ denotes the cumulative time devoted by the server to class i by time t . In particular, U_i are nondecreasing, Lipschitz continuous with constant 1, and $\sum_i U_i(t) \leq t$ for all t . We regard A and S as *primitive processes*, and call X and U a *state process* and a *control process*, respectively, if U is adapted to the filtration $\mathcal{F}_t = \sigma\{A_i(s), X_i(s), s \leq t, i \leq N\}$.

It is assumed that A_i and S_i are mutually independent renewal processes. In the n th system, A_i and S_i are replaced by $A_i^n = A_i(n \cdot)$ and $S_i^n = S_i(n \cdot)$, and the corresponding control and queue length processes are denoted by U^n and X^n , respectively. Normalized queue length is denoted by $\bar{X}_i^n = n^{-1} X_i^n$.

Fix $T > 0$ and constants $c_i > 0$. For $\beta > 0$, denote

$$J^n(U^n; Q, \beta) = \frac{1}{n} \frac{1}{\beta} \log E_Q e^{\beta \sum_i c_i X_i^n(T)}.$$

There is redundancy in the definition with respect to β and c_i . We use the parameter β to be consistent with Theorem 4.1, but one could let $\beta = 1$ without loss. Denote by P the Markovian model, where A_i and S_i are Poisson processes with parameters λ_i and μ_i , respectively. Denoting

$$V^n(P, \beta) = \inf_{U^n} J^n(U^n; P, \beta),$$

where the infimum ranges over control processes U^n , the limit $V(P, \beta) = \lim_n V^n(P, \beta)$ was shown to exist and was characterized in [5] as the viscosity solution of a HJB equation. In [6] it was proved that for zero initial conditions one has $V \leq \beta^{-1} WT$, where

$$W = W(\beta) = \min_{u \in \mathbf{u}} \sum_{i=1}^N (\hat{\lambda}_i - u_i \hat{\mu}_i)^+,$$

$\mathbf{u} = \{u \in \mathbb{R}_+^N : \sum u_i \leq 1\}$, $\hat{\lambda}_i = \lambda_i(e^{\beta c_i} - 1)$ and $\hat{\mu}_i = \mu_i(1 - e^{-\beta c_i})$. It was also shown in [6] that when $e^{-\beta c_i} < \lambda_i/\mu_i$ for all i the bound is tight, that is, $V(\beta) = \beta^{-1}WT$. In this case, it is asymptotically optimal to prioritize according to the index $\mu_i(1 - e^{-\beta c_i})$, regardless of T , with larger values given priority.

Let P be fixed as above, and consider a family \mathcal{Q} defined via part (ii) of Theorem 3.1. That is, letting $h_{i,1}$ and $h_{i,2}$ stand for the hazard rates for A_i and S_i , respectively, assume that

$$(4.4) \quad a_{i,1} \leq \frac{h_{i,1}(\cdot)}{\lambda_i} \leq b_{i,1}, \quad a_{i,2} \leq \frac{h_{i,2}(\cdot)}{\mu_i} \leq b_{i,2}$$

for some constants $0 < a_{i,j} < b_{i,j}$. Denote by Q_T^n and P_T^n the law of $(A^n, S^n)|_{[0,T]}$ under Q and P . Then by (3.4), for all $Q \in \mathcal{Q}$,

$$R_\alpha(Q_T^n \| P_T^n) \leq nTf_0(\alpha),$$

where

$$f_0(\alpha) = \sum_i [k_\alpha(a_{i,1}) \vee k_\alpha(b_{i,1})] \lambda_i + \sum_i [k_\alpha(a_{i,2}) \vee k_\alpha(b_{i,2})] \mu_i$$

[we recall $k_\alpha(x) = [x^\alpha - \alpha x + \alpha - 1]/\alpha(\alpha - 1)$ introduced in (3.1)]. Thus Theorem 4.1 may be applied with $f(\alpha) = nTf_0(\alpha)$. Denoting

$$V^n(Q, \beta) = \inf_{U^n} \sup_{Q \in \mathcal{Q}} J^n(U^n; Q, \beta),$$

we have by Theorem 4.1 that, for all n ,

$$(4.5) \quad V^n(Q, \beta) \leq \inf_{U^n} \inf_{\gamma > \beta} \left[\frac{Tf_0(\frac{\gamma}{\gamma-\beta})}{\gamma - \beta} + J^n(U^n; P, \gamma) \right] = \inf_{\gamma > \beta} \left[\frac{Tf_0(\frac{\gamma}{\gamma-\beta})}{\gamma - \beta} + V^n(P, \gamma) \right].$$

As mentioned above, for each γ , $\limsup_n V^n(P, \gamma) \leq \gamma^{-1}W(\gamma)T$ (according to Theorem 2.1 of [6]). Hence we obtain the following upper bound.

THEOREM 4.2. *For each β , one has*

$$(4.6) \quad \limsup_n V^n(Q, \beta) \leq B(Q, \beta) \doteq \inf_{\gamma > \beta} \left[\frac{f_0(\frac{\gamma}{\gamma-\beta})}{\gamma - \beta} + \frac{W(\gamma)}{\gamma} \right] T.$$

PROOF. To deduce (4.6) from (4.5), let $\varepsilon > 0$ be given and let γ_0 be ε -optimal for the RHS of (4.6). Then use (4.5) with the infimum over γ replaced by the substitution $\gamma = \gamma_0$. Taking the limit and using the aforementioned limit result for $V^n(P, \gamma_0)$, then sending $\varepsilon \rightarrow 0$, gives (4.6). \square

We now show that using the last result one can identify a policy for which the robust bound (4.6) is valid.

REMARK 4.1. The robust bound derived in Theorem 4.2 can in fact be achieved by an index policy. Given β and $\varepsilon > 0$, let γ^* be an ε -minimizer of the RHS of (4.6). Let us show that prioritizing queues according to the index $\mu_i(1 - e^{-\gamma^* c_i})$ (rather than $\mu_i(1 - e^{-\beta c_i})$) guarantees the uniform bound

$$\limsup_n \sup_{Q \in \mathcal{Q}} J^n(U^{*,n}; Q, \beta) \leq B(Q, \beta) + \varepsilon,$$

where $U^{*,n}$ denotes the control corresponding to the above mentioned index. To see this, note that by Theorem 4.1, when $U^{*,n}$ is implemented one has for all n

$$J^n(U^{*,n}; \mathcal{Q}, \beta) \leq \inf_{\gamma > \beta} \left[\frac{Tf_0(\frac{\gamma}{\gamma-\beta})}{\gamma-\beta} + J^n(U^{*,n}; P, \gamma) \right] \leq \left[\frac{Tf_0(\frac{\gamma^*}{\gamma^*-\beta})}{\gamma^*-\beta} + J^n(U^{*,n}; P, \gamma^*) \right],$$

and since by [6], under the fixed priority $U^{*,n}$ one has that $J^n(U^{*,n}; P, \gamma^*)$ converges to $(\gamma^*)^{-1}TW(\gamma^*)$, the claim follows.

Define $\ell(x) = x \log x - x + 1$ for $x \geq 0$. When β is small a natural assumption to make on the $a_{i,k}$ and $b_{i,k}$, consistent with the fact that the Rényi rate $k_\alpha(x)$ becomes the relative entropy rate $\ell(x)$ as $\alpha \downarrow 1$, is that $a_{i,k} < 1 < b_{i,k}$ and $\ell(a_{i,k}) = \ell(b_{i,k})$ so long as $\ell(b_{i,k}) \leq 1$, and $a_{i,k} = 0$ if $\ell(b_{i,k}) > 1$. In this case one can show that the bounds are also tight in a precise sense, which is that there exists a model in \mathcal{Q} such that the two sides differ by no more than error term that vanishes as $\beta \downarrow 0$ and which can be calculated.

As pointed out above, the robust RS control policy thus obtained prioritizes according to an index that is distinct from that used for the reference model. This illustrates an important aspect of the general approach of using Theorem 4.1 and (4.3), namely that there is more to this approach than directly applying the Rényi bounds to the state process obtained under the optimal RS control for the reference model P . Indeed, the latter approach would give rise to a control for (\mathcal{Q}, β) that agrees with that for (P, β) . Instead, the minimization problem (4.3) allows for the control (and consequently the state process) to differ from the one that is optimal for (P, β) by allowing freedom in choosing the sensitivity parameter γ . Thus γ is selected to best fit the family \mathcal{Q} , which may indeed result in a control policy that is not optimal for the “reference problem” (P, β) .

EXAMPLE 4.1. We evaluate the bound $B(\mathcal{Q}, \beta)$ of (4.6) numerically. We consider an example with five classes, with data $\lambda = (1, 1.5, 1.8, 2, 2)$ and $\mu = (8, 10, 12, 9, 14)$. The overall traffic intensity $\rho = \sum_i \frac{\lambda_i}{\mu_i}$ is $\rho = 0.790$. The relative costs c_i are taken to be $c = (0.3, 0.2, 0.2, 0.1, 0.2)$, and the time horizon $T = 1$.

First the reference model is considered. When \mathcal{Q} is a singleton consisting of the model P , the bound is $B(\mathcal{Q}, \beta) = \beta^{-1}W(\beta) = V(\beta)$. This function is shown in blue in Figure 2 (left), for β in the range $[0, 15]$.

Consider the family determined by (4.4), with $a_{i,1} = a_{i,2} = 1 - \delta$ and $b_{i,1} = b_{i,2} = 1 + \delta$ for all i , for $\delta = 0.65$. Recall that this corresponds to a family of models driven by renewal processes, where the interarrival and service time distributions have hazard rates that deviate from the respective Poisson rates of the reference model by at most 65%. Moreover, according to Theorem 3.1(ii), this may also stand for a family of models where the driving processes are Cox, for which the stochastic intensities deviate from those of the reference model by at most 65%. This family is denoted by \mathcal{Q}_2 (for it corresponds to part (ii) of Theorem 3.1). In Figure 2 (left) the bound $B(\mathcal{Q}_2, \beta)$ the RS cost for this family is shown in solid black line. A dotted black line shows the bound $B(\mathcal{Q}_2, \beta)$ where now the parameter δ is taken as $\delta = 0.15$.

Next, consider the family of models, denoted by \mathcal{Q}_3 , for which the driving processes are as in Theorem 3.1(iii). These are Cox processes which, in addition to bounds on the deviation from the reference Poisson rates, the stochastic intensities satisfy a long run average constraint. For example, the potential service process for class 1 has stochastic intensity that deviates from μ_1 by at most δ , and in addition is constrained to have a long run average equal to μ_1 . In Figure 2 (left), the bound $B(\mathcal{Q}_3, \beta)$ is shown in solid red line and in dotted red line for $\delta = 0.65$ and $\delta = 0.15$, respectively. As expected, the bounds for \mathcal{Q}_3 are smaller than for \mathcal{Q}_2 , and they are smaller for $\delta = 0.15$ than they are for $\delta = 0.65$.




FIG. 2. Robust risk-sensitive bounds for the scheduling problem.

Finally, all five graphs are repeated in Figure 2 (right) with a different λ , namely $\lambda = (0.5, 0.75, 0.9, 1, 1)$ (leaving the remaining parameters unchanged) in which case $\rho = 0.395$. The queueing system is more stable in this case, and the performance guarantees, as measures by the RS cost bounds, are smaller as expected.

5. Robust LD estimates for queueing models with reneging. In this section we study a multiserver queue with reneging, under a scaling where the number of servers n and the arrival process grow proportionally. This scaling has been referred to as a *many-server* scaling, studied for the first time in [21], for CLT asymptotics in the case of exponential servers, and then in the context of general service times in [25, 27, 28, 37] (for LLN and CLT asymptotics). For models that accommodate reneging, it is natural to define performance in terms of the reneging count. This was addressed recently in [2], where the large time, large n asymptotics of the probability of atypically large reneging count was identified precisely. These results were concerned with the Markovian $M/M/1 + M$ and $M/M/n + M$ models. Whereas the results of [2] identify exact LD asymptotics for one particular model, our interest here is in the spirit of robust bounds, in estimates that are uniform within families of models, that are moreover non-Markovian. The results from [2] will serve us as reference for these uniform bounds.

Treating general service time distributions via a Markovian reference model relies, according to our approach, on Rényi divergence estimates of the underlying primitives, which in this case are given by the potential service processes for each server. This is precisely where our results from Section 3 on divergence of various counting processes w.r.t. Poisson become useful. A similar remark holds for other primitives of the model, namely arrival and patience times. Specifying server characteristics by means of a counting process that lies in a given Rényi radius about some nominal Poisson gives room for modelling servers as different from one another. In fact, in this framework there is no benefit to requiring that servers be statistically identical. This gives rise to a set of models much more rich than $G/G/n + G$, that accommodates (a) heterogeneous servers, and (b) time varying processing capacities.

Whereas item (a) above allows for distinct probabilistic characteristics for each server, our approach is to express the degree of uncertainty (w.r.t. service times distributions) on equal terms for all servers. This should not be confused with models such as $G/G/n$ or $G/G/n + G$ where all servers operate under the same distribution. This modelling approach is perhaps

more satisfactory than models like $G/G/n$ in situations where there is no information that distinguishes between servers but at the same time there is no reason to believe that all are identical. An analogous remark is valid for modelling patience of different customers.

5.1. Model and performance measure.

5.1.1. *Model equations.* Customers arrive at the system with service requirement that can be handled by any one of n parallel servers. They are queued if no servers are available upon arrival, and renege if they are still in the queue at the time their patience expires. The priority within the queue is according to FIFO. Determining which available server takes the next customer is according to a fixed ordering of the servers.

Because, on the one hand, servers have different characteristics and, on the other hand, customer renegeing depends on their state (specifically, whether they are in the queue and for how long), the model equations must account for the state of each server as well as the state of each customer. Hence our system of equations will be based on a balance equation for each server and one for each (of the infinitely many) customers.

The model equations are therefore somewhat complicated. However, because our approach is based on the existence of a mapping from primitive processes to the full state of the system, it is necessary to write down these equations so that a concrete mapping is indeed well defined. Measure valued processes are often used for encoding the dynamics, however, it seems less complicated in the current context to write balance equations, as we will. Also, we are careful to write the equations without relying on an assumption that the underlying discrete events occur one at a time; that is, they allow for the possibility of simultaneous arrival and departure, simultaneous departures at different servers, etc. This assures that the mapping is defined on the full path space of the primitive processes.

The customers are indexed by \mathbb{N} , and a marked point process $\sum_{i \in \mathbb{N}} \delta_{(T_i, P_i)}$, with sample paths in $\mathcal{D}([0, \infty) : \mathcal{M}_F(\mathbb{R}_+^2))$ encodes their time of arrival T_i and their patience time P_i . With a slight abuse of notation, in what follows we refer to $\mathcal{A} = (T_i, P_i)$ as the marked point process. It is assumed that $0 \leq T_1 \leq T_2 \leq \dots$ and $P_i > 0$ for all i . Those customers i with $T_i = 0$ are initially in the system. The n servers are indexed by $[n] = \{1, \dots, n\}$, and a counting process S_j is associated with each server $j \in [n]$, representing its potential service process. That is, $S_j(t)$ customers depart server j by the time this server has worked for t units of time.

We start with a balance equation for each customer. For $i \in \mathbb{N}$, we have

$$Q_i(t) = A_i(t) - K_i^{\text{cust}}(t) - R_i(t).$$

Here, the four processes Q_i , A_i , K_i^{cust} and R_i are $\{0, 1\}$ valued, representing queueing, arrival, routing and renegeing, respectively, associated with customer i . Thus Q_i (resp., A_i , K_i^{cust} , R_i) takes the value 1 at time t if customer i in the queue at that time (resp., has arrived prior to or at t , has been routed to service prior to or at t , has renegeed prior to or at t). In particular, we have $A_i(t) = 1_{\{t \geq T_i\}}$.

Next, a balance equation holds for each server $j \in [n]$, in the form

$$B_j(t) = B_j(0) + K_j^{\text{serv}}(t) - D_j(t), \quad D_j(t) = S_j \left(\int_0^t B_j(s) ds \right).$$

Here, B_j , K_j^{serv} and D_j are busyness, routing and departure processes associated with server j , taking values in $\{0, 1\}$, \mathbb{Z}_+ and \mathbb{Z}_+ , respectively. Namely, B_j takes the value 1 at t if the server is busy, and K_j^{serv} and D_j are counting processes for the number customers routed to and, respectively, departing from server j .

The initial conditions are assumed to match and to satisfy a work conservation condition. Namely, the number of customers initially in the system, $X(0) = \max\{i : T_i = 0\}$, and the number of servers initially busy, $B(0) = \sum_j B_j(0)$, satisfy $B(0) = X(0) \wedge n$.

Next we describe how the routing processes are determined so as to keep the aforementioned priority rules. To this end, we denote by

$$AV^{\text{cust}}(t) = \{i \in \mathbb{N} : \text{either } Q_i(t-) = 1 \text{ or } \Delta A_i(t) = 1\}$$

the set of customers available for routing at time t and by

$$AV^{\text{serv}}(t) = \{j \in [n] : \text{either } \Delta D_j(t) = 1 \text{ or } B_j(t-) = 0\}$$

the set of servers available to serve at this time, where for a real valued càdlàg function f on $[0, \infty)$, $\Delta f(s) = f(s) - f(s-)$. The number of customers to be routed at time t is given by

$$\hat{K}(t) = \#AV^{\text{cust}}(t) \wedge \#AV^{\text{serv}}(t).$$

In terms of $\hat{K}(t)$, one determines which customers i are routed to service, and which servers j admit new customers at time t , according to

$$\hat{K}_i^{\text{cust}}(t) = \begin{cases} 1 & \text{if } i \in AV^{\text{cust}}(t), \#\{i' \in AV^{\text{cust}}(t) : i' \leq i\} \leq \hat{K}(t), \\ 0 & \text{otherwise,} \end{cases}$$

$$\hat{K}_j^{\text{serv}}(t) = \begin{cases} 1 & \text{if } j \in AV^{\text{serv}}(t), \#\{j' \in AV^{\text{serv}}(t) : j' \leq j\} \leq \hat{K}(t), \\ 0 & \text{otherwise.} \end{cases}$$

The corresponding counting processes are given by

$$K_i^{\text{cust}}(t) = \sum_{s \leq t} \hat{K}_i^{\text{cust}}(s),$$

$$K_j^{\text{serv}}(t) = \sum_{s \leq t} \hat{K}_j^{\text{serv}}(s).$$

To determine R_i , note that renegeing occurs at time $T_i + P_i$, but only on the event that the customer is in the queue at that time. Thus

$$R_i(t) = \begin{cases} 1 & \text{if } t \geq T_i + P_i \text{ and } K_i^{\text{cust}}(T_i + P_i) = 0, \\ 0 & \text{otherwise.} \end{cases}$$

According to this definition, if routing of a customer to service and renegeing potentially occur at the same time, priority is given to routing.

Finally, the total queue length, number of busy servers, arrival count, departure count, renegeing count and routing count are given, respectively, by

$$Q = \sum_i Q_i, \quad B = \sum_j B_j, \quad A = \sum_i A_i, \quad D = \sum_j D_j,$$

$$R = \sum_i R_i, \quad K = \sum_i K_i^{\text{cust}} = \sum_j K_j^{\text{serv}}.$$

The state of the system is the process $\Sigma = \{A_i, Q_i, K_i^{\text{cust}}, R_i, B_j, K_j^{\text{serv}}, D_j, Q, B, A, D, R, K\}$.

The primitive processes $\mathcal{A} = (T_i, P_i)$ and S_j determine the state of the system. Proving this amounts to showing that there exists a unique solution to the set of all equations that appear in this subsection; we skip the details of the elementary proof of this fact. An additional important fact that we state without proof is a causality property, namely that for any t ,

$\{\Sigma(s) : s \in [0, t]\}$ is measurable on the sigma field corresponding to the primitive data up to time t , namely $\sigma\{(T_i, P_i)1_{\{T_i \leq t\}}, S_j(s), s \leq t, j \in [n]\}$.

It is assumed throughout that the potential service processes S_j are mutually independent, and that, moreover, the initial data $(\{B_j(0)\}_{j \in \mathbb{N}}, X(0))$, the service primitive $\{S_j\}_{j \in \mathbb{N}}$ and the customer primitive \mathcal{A} are mutually independent, for each model \mathbb{Q} in the family of models \mathcal{Q} to be considered.

5.1.2. LD scaling and performance measure. We now consider a sequence of models indexed by $n \in \mathbb{N}$. It is convenient to assume, as we will, that the primitives S_j are given for all $j \in \mathbb{N}$, and that for the n th system, one takes $S_j^n = S_j, j \in [n]$. Similarly, for the arrival process, it is convenient to start with a single sequence $\mathcal{A} = (T_i, P_i)$ and obtain the arrival process for the n th system, $\mathcal{A}^n = (T_i^n, P_i^n)$, via $T_i^n = n^{-1}T_i$ and $P_i^n = P_i$. The transformation of arrival times reflects acceleration of arrivals, performed in order to keep a constant traffic intensity as n increases by balancing the increase of processing capacity due to the growing number of servers. The patience times, however, are not accelerated. This is in agreement with literature on many server scaling at LLN and CLT regimes, such as [7, 8, 25]. The superscript n is attached to all processes involved in the n th system.

Our main interest is in the LD behavior of the renegeing count R^n . In the special case of Markovian model, the large time average rate of overall renegeing can be obtained by simple LLN considerations. That is, assume that for some $\lambda, \mu, \theta > 0$, the rate of arrivals is given by λn , the total service rate by μn , and the per-customer renegeing rate by θ . Consider an overloaded system, $\lambda > \mu$. Then the renegeing stabilizes the system at an equilibrium around xn for which $\lambda = \mu + \theta x$. Hence the long time average renegeing rate is given by $\gamma_0 = \theta x = \theta \frac{\lambda - \mu}{\theta} = \lambda - \mu$, and thus for $\gamma > \gamma_0$, the event that the long time average renegeing rate exceeds γ is rare.

For a general model \mathbb{Q} and an arbitrary $\gamma > 0$, define the decay rate

$$\chi(\mathbb{Q}, \gamma) = \limsup_{t \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{tn} \log \mathbb{Q} \left(\frac{R^n(t)}{tn} > \gamma \right),$$

and for a collection of models \mathcal{Q} let

$$\chi(\mathcal{Q}, \gamma) = \limsup_{t \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{tn} \log \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{Q} \left(\frac{R^n(t)}{tn} > \gamma \right).$$

Bounds on $\chi(\mathcal{Q}, \gamma)$ will be based on known bounds on $\chi(\mathbb{P}, \gamma)$, where \mathbb{P} stands for the aforementioned Markovian model (i.e., $M/M/n + M$), and $\gamma > \gamma_0$.

THEOREM 5.1 ([2]). Assume $\lambda \geq \mu$. Let $C(\gamma) = \lambda(1 - z^{-1}) + \mu(1 - z) - \gamma \log z$, where

$$z = z(\gamma) = \frac{\sqrt{\gamma^2 + 4\mu\lambda} - \gamma}{2\mu}.$$

Then $\chi(\mathbb{P}, \gamma) = -C(\gamma)$, for $\gamma \geq \gamma_0$.

5.2. Robust bounds.

5.2.1. Robust bounds in general form. For a collection of models \mathcal{Q} , the marked point process $\mathcal{A}^n = (T_i^n, P_i^n)$, with $\mathcal{A}_t^n = \{(T_i^n, P_i^n) : T_i^n \leq t\}$, which encodes arrival and patience processes, is assumed to satisfy the RDR bound

$$(5.1) \quad \limsup_{t \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{nt} \sup_{\mathbb{Q} \in \mathcal{Q}} R_\alpha(\mathbb{Q} \circ \mathcal{A}^n|_{[0,t]}^{-1} \| \mathbb{P} \circ \mathcal{A}^n|_{[0,t]}^{-1}) \leq r_\alpha^{(1)},$$

where $r_\alpha^{(1)}$ is an α -dependent constant. The probabilistic characteristics of the servers are encoded in the service processes S_j , that are taken to satisfy a similar bound, uniform in j ,

$$(5.2) \quad \limsup_{t \rightarrow \infty} \frac{1}{t} \sup_{\mathbb{Q} \in \mathcal{Q}} \sup_{j \in \mathbb{N}} R_\alpha(\mathbb{Q} \circ S_j|_{[0,t]}^{-1} \| \mathbb{P} \circ S_j|_{[0,t]}^{-1}) \leq r_\alpha^{(2)}.$$

THEOREM 5.2. *Assume (5.1) and (5.2). Then we have, for every $\gamma \geq \gamma_0$, the estimate*

$$\chi(\mathcal{Q}, \gamma) \leq B(\mathcal{Q}, \gamma) \doteq \inf_{\alpha > 1} \left[-\frac{\alpha - 1}{\alpha} C(\gamma) + (\alpha - 1)(r_\alpha^{(1)} + r_\alpha^{(2)}) \right].$$

PROOF. Clearly, the event $(tn)^{-1}R^n(t) > \gamma$ is measurable on $\sigma\{\Sigma^n(s) : s \in [0, t]\}$, hence on the sigma field corresponding to the primitives, $\sigma\{\mathcal{A}_s^n, S_j(s) : s \in [0, t], j \in [n]\}$. Hence by (2.7),

$$\begin{aligned} \sup_{\mathbb{Q} \in \mathcal{Q}} \frac{1}{tn} \log \mathbb{Q} \left(\frac{R^n(t)}{tn} > \gamma \right) &\leq \frac{\alpha - 1}{\alpha} \frac{1}{tn} \log \mathbb{P} \left(\frac{R^n(t)}{tn} > \gamma \right) \\ &\quad + (\alpha - 1) \frac{1}{tn} \sup_{\mathbb{Q} \in \mathcal{Q}} \left[R_\alpha(\mathbb{Q} \circ \mathcal{A}^n|_{[0,t]}^{-1} \| \mathbb{P} \circ \mathcal{A}^n|_{[0,t]}^{-1}) \right. \\ &\quad \left. + \sum_{j=1}^n R_\alpha(\mathbb{Q} \circ S_j|_{[0,t]}^{-1} \| \mathbb{P} \circ S_j|_{[0,t]}^{-1}) \right], \end{aligned}$$

using the assumed mutual independence of the different service processes S_j , as well as their independence from \mathcal{A}^n . Taking the limit superior in n , then in t , and using Theorem 5.1 and assumptions (5.1) and (5.2), gives

$$\chi(\mathcal{Q}, \gamma) \leq -\frac{\alpha - 1}{\alpha} C(\gamma) + (\alpha - 1)(r_\alpha^{(1)} + r_\alpha^{(2)}).$$

The result follows on optimizing over α . \square

5.2.2. Examples. First we provide examples where uncertainty classes are in the spirit of Theorem 3.1. In the reference model \mathbb{P} , arrivals are Poisson($n\lambda$), patience times are exponential(θ), and individual service rates are $\mu = 1$.

EXAMPLE 5.1. Consider the family of models, \mathcal{Q}_2 , corresponding to the setting of Theorem 3.1(ii), where all potential service processes are Cox processes with stochastic intensities $a \leq \lambda_j(\cdot) \leq b$, for some constants $0 < a \leq 1 \leq b < \infty$ and all $j \in \mathbb{N}$. Alternatively, (again, see Theorem 3.1(ii)) the service processes are renewals with service time distribution for which the hazard rate satisfies $a \leq h_j(\cdot) \leq b$ for each server $j \in \mathbb{N}$. As already mentioned, the potential service processes are assumed to be mutually independent, but they are not assumed to be identically distributed. However, the distributions associated with all servers are assumed to lie within the same uncertainty class, namely the one determined by the bounds a and b . For arrival and patience distributions, recall that $\mathcal{A}^n = (T_i^n, P_i^n)$ are taken as rescaled versions of $\mathcal{A} = (T_i, P_i)$. Our assumptions are that T_i is a Cox process with stochastic intensity $a_{\text{arr}} \leq \frac{\lambda(\cdot)}{\lambda} \leq b_{\text{arr}}$, and the distributions of the patience times P_i have distributions $\psi_i(z)\zeta(dz)$ satisfying the bound $a_{\text{pat}} \leq \psi_i(\cdot) \leq b_{\text{pat}}$, where ζ is the distribution of an exponential(θ) RV $\zeta(dz) = \theta e^{-\theta z} dz$.

For the potential service processes, consider Theorem 3.1(ii) in the special case of no marks. Then, for each server j ,

$$\sup_{\mathbb{Q} \in \mathcal{Q}_2} R_\alpha(\mathbb{Q} \circ S_j|_{[0,t]}^{-1} \| \mathbb{P} \circ S_j|_{[0,t]}^{-1}) \leq \sup_{\mathbb{Q} \in \mathcal{Q}_2^*} R_\alpha(\mathbb{Q} \circ N|_{[0,t]}^{-1} \| \mathbb{P} \circ N|_{[0,t]}^{-1}),$$

where \mathcal{Q}_2^* denotes the class from Theorem 3.1(ii), of all probability measures that make N a Cox process with stochastic intensity $a \leq \lambda(\cdot) \leq b$ (taking $\lambda_0 = 1$ in (3.3), in line with the assumption that $\mu = 1$ under the reference model). Taking the supremum over j we obtain from Theorem 3.1(ii) that (5.2) holds with $r_\alpha^{(2)} = k_\alpha(a) \vee k_\alpha(b)$.

Next, to use the same theorem for the arrival and patience distributions, note first that $\hat{a} \leq \frac{\lambda(s)\psi(z)}{\lambda} \leq \hat{b}$, where $\hat{a} = a_{\text{arr}}a_{\text{pat}}$ and $\hat{b} = b_{\text{arr}}b_{\text{pat}}$. Consequently, we obtain (5.1) with $r_\alpha^{(1)} = k_\alpha(\hat{a}) \vee k_\alpha(\hat{b})$.

Appealing to Theorem (5.2), we obtain the estimate

$$\chi(\mathcal{Q}_2, \gamma) \leq \inf_{\alpha > 1} \left[-\frac{\alpha - 1}{\alpha} C(\gamma) + (\alpha - 1)(k_\alpha(\hat{a}) \vee k_\alpha(\hat{b}))\lambda + (\alpha - 1)(k_\alpha(a) \vee k_\alpha(b)) \right].$$

As a second and third uncertainty classes, denoted by \mathcal{Q}_3 and \mathcal{Q}_4 , we take families of measures corresponding to Theorem 3.1(iii) and 3.1(iv), respectively. In both \mathcal{Q}_3 and \mathcal{Q}_4 , the arrival and patience are taken as in \mathcal{Q}_2 . As for service time distributions, in \mathcal{Q}_3 consider Cox processes for which the stochastic intensity satisfies the long time average constraint (3.5) (with $\lambda_0 = 1$) and the constraint $a \leq \lambda_j(\cdot) \leq b$. In \mathcal{Q}_4 , the stochastic intensity satisfies (3.5) and the constraint (3.2) for some $\alpha = \alpha_0$ and u . The bounds obtained in these cases are

$$\chi(\mathcal{Q}_3, \gamma) \leq \inf_{\alpha > 1} \left[-\frac{\alpha - 1}{\alpha} C(\gamma) + (\alpha - 1)(k_\alpha(\hat{a}) \vee k_\alpha(\hat{b}))\lambda + (\alpha - 1)(pk_\alpha(a) + qk_\alpha(b)) \right],$$

$p = \frac{b-1}{b-a}$, $q = \frac{1-a}{b-a}$, and

$$\chi(\mathcal{Q}_4, \gamma) \leq \inf_{\alpha \in (1, \alpha_0)} \left[-\frac{\alpha - 1}{\alpha} C(\gamma) + (\alpha - 1)(k_\alpha(\hat{a}) \vee k_\alpha(\hat{b}))\lambda + \frac{1}{\alpha} [(\bar{\alpha}_0 u + 1)^{\frac{\alpha-1}{\alpha_0-1}} - 1] \right],$$

where $\bar{\alpha}_0 = \alpha_0(\alpha_0 - 1)$.

For a numerical example we take the following numerical values. Since the reference service rates are normalized to 1, the reference system will be overloaded in $\lambda > 1$. We take $\lambda = 2$. The corresponding LLN reneging rate is $\gamma_0 = 1$. As bounds on intensities we take $a = \hat{a} = 1 - \delta$ and $b = \hat{b} = 1 + \delta$, where $\delta = 0.3$.

Figure 3 (left) gives graphs of $B(\mathcal{Q}_2, \gamma)$ and $B_3(\mathcal{Q}_3, \gamma)$ corresponding to the families \mathcal{Q}_2 and \mathcal{Q}_3 , as well as $\chi(\mathbb{P}, \gamma) = -C(\gamma)$ for reference (the exact decay rate under \mathbb{P}). In addition to the families \mathcal{Q}_2 and \mathcal{Q}_3 , we consider families \mathcal{Q}'_2 and \mathcal{Q}'_3 defined analogously to \mathcal{Q}_2 and \mathcal{Q}_3 , respectively, but where uncertainty is associated with the service processes only, hence $r_\alpha^{(1)}$ is taken to be 0. The corresponding bounds $B(\mathcal{Q}'_2, \gamma)$ and $B(\mathcal{Q}'_3, \gamma)$ are also shown in Figure 3 (left).

Whereas the above example is based on RDR bounds for families of processes (Theorem 3.1), the following is based, in addition, on our RDR bounds for specific renewal distributions (Theorem 3.2).

EXAMPLE 5.2. We consider a family, denoted \mathcal{Q}_Γ , where all servers operate according to Gamma distributions. More precisely, service time distribution for server j is $\Gamma(k_j, \rho_j)$, $k_j \geq 1$, $\rho_j > 1$, and a subset $F \subset [1, \infty) \times (1, \infty)$ is given for which $(k_j, \rho_j) \in F$ for all j . The assumptions on the arrival and patience process, \mathcal{A} , are as in Example 5.1. Then by the bound on RDR for the Gamma distribution stated in Example 3.2, the bound (5.2) is valid with

$$r_\alpha^{(2)} = r_\alpha^{(2)}(F) \doteq \sup_{(k, \rho) \in F} r_\alpha^{(2)}(k, \rho),$$




FIG. 3. Robust bounds for the reneging problem.

where we denote

$$r_\alpha^{(2)}(k, \rho) = \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1 + \alpha(k - 1)}} - \alpha(\rho - 1) - 1.$$

As a consequence,

$$\chi(Q_\Gamma, \gamma) \leq \inf_{\alpha > 1} \left[-\frac{\alpha - 1}{\alpha} C(\gamma) + (\alpha - 1)(k_\alpha(\hat{a}) \vee k_\alpha(\hat{b}))\lambda + (\alpha - 1)r_\alpha^{(2)}(F) \right].$$

Figure 3 (right) gives graphs of $B(Q_\Gamma)$ for Q_Γ for two parameter ranges, namely $(k, \rho) \in [1, 1.1] \times [1, 1.1]$ and $(k, \rho) \in [1, 1.5] \times [1, 1.5]$. The parameters $\lambda, \hat{a}, \hat{b}$ are taken to be $\lambda = 2, \hat{a} = 1 - \delta, \hat{b} = 1 + \delta$, where $\delta = 0.3$.

Finally, we also consider Q'_Γ defined as Q_Γ but where uncertainty is associated with the service processes only ($r_\alpha^{(1)} = 0$). The corresponding bounds are also shown in Figure 3 (right), with the same ranges of parameters (k, ρ) .

Once again, $\chi(\mathbb{P}, \gamma)$ is also plotted for reference.

6. Concluding remarks. The techniques developed in this paper are not limited to queueing models. The basic bound (2.8) can be used in far broader dynamical system settings. In the most general terms, its usefulness relies on the ability to provide (1) a LD estimate under some reference measure P (the first term on the RHS of (2.8)) and (2) a computation of, or an effective upper bound on, the RDR for a family of models of interest (second term on the RHS of (2.8)). For example, if the dynamical systems are driven by point processes (like in queueing applications), the relevant RDR need not correspond to renewal versus Poisson like in this paper, but between families of point processes relevant to the application. Such RDR estimates need to be developed.

The main viewpoint presented in this paper was to consider a reference model for which computation is possible, and a family of models that need not be tractable. A different perspective, initiated in [17], is to use these bounds for sensitivity analysis of rare event probabilities. This paper introduces new gradient based sensitivity indices that are meaningful at the large deviations scale, and develops sensitivity bounds which do not require a rare event sampler for each rare event. This quality is closely related to the fact that in (2.8) the difference in performance under two measures is bounded solely in terms of the Rényi divergence,

and does not depend on the rare event A . This method of [17] arguably has an advantage over more traditional approaches of direct statistical estimation of rare event sensitivities.

Finally, the robust bounds for RS control developed in Section 4.1 are valid in far greater generality than for queueing applications, as we have indeed emphasized in that section. As long as the function F in Theorem 4.1 can be computed (or estimated) for a given design model P and a specified family of models \mathcal{Q} , the robust bounds established in this result are available.

APPENDIX

A.1. Proofs of results from Section 3.1. PROOF OF THEOREM 3.1. (i) For $T > 0$, the Radon–Nikodym (RN) derivative of Q_T^N w.r.t. P_T^N is given by (see [26], Theorem 2.31)

$$(A.1) \quad \Lambda_T = e^{-\int_{[0,T] \times S} (\lambda(t,z) - \lambda_0) \zeta(dz) dt + \int_{[0,T] \times S} \log \frac{\lambda(t,z)}{\lambda_0} N(dt dz)}$$

Raising this expression to the power α gives

$$(A.2) \quad \begin{aligned} \Lambda_T^\alpha &= e^{-\int_{[0,T] \times S} (\lambda(t,z)^\alpha \lambda_0^{1-\alpha} - \lambda_0) \zeta(dz) dt + \int_{[0,T] \times S} \log \frac{\lambda(t,z)^\alpha \lambda_0^{1-\alpha}}{\lambda_0} N(dt dz)} \\ &\times e^{\int_{[0,T] \times S} (\lambda(t,z)^\alpha \lambda_0^{1-\alpha} - \alpha \lambda(t,z) + (\alpha-1)\lambda_0) \zeta(dz) dt} \\ &= M_T e^{\int_{[0,T] \times S} \lambda_0 \alpha (\alpha-1) k_\alpha \left(\frac{\lambda(t,z)}{\lambda_0}\right) \zeta(dz) dt}, \end{aligned}$$

where the process

$$M_t = e^{-\int_{[0,t] \times S} (\lambda(s,z)^\alpha \lambda_0^{1-\alpha} - \lambda_0) \zeta(dz) ds + \int_{[0,t] \times S} \log \frac{\lambda(s,z)^\alpha \lambda_0^{1-\alpha}}{\lambda_0} N(ds dz)}, \quad 0 \leq t \leq T$$

is a P -(local) martingale. By the hypothesis on $\lambda(\cdot)$, for all $Q \in \mathcal{Q}_1$ and all T ,

$$\int_{[0,T] \times S} k_\alpha(\lambda(t,z)/\lambda_0) \zeta(dz) dt \leq T(u + v(T)).$$

Hence $R_\alpha(Q_T^N \| P_T^N) \leq (u + v(T))T\lambda_0$. Consequently, by the definition of the RDR, and since $v(T) \rightarrow 0$, $r_\alpha(Q_1 \| P)$ is bounded above by $u\lambda_0$. Equality follows on taking $\lambda(t,z)$ to be the constant λ_1 for which $k_\alpha(\lambda_1/\lambda_0) = u$.

(ii) Fix $Q \in \mathcal{Q}_2$. By the convexity of k_α , the property (3.3) implies that for all t , $k_\alpha(\lambda(t,z)/\lambda_0) \leq k_\alpha(a) \vee k_\alpha(b)$, which using (A.2) yields

$$R_\alpha(Q_T \| P_T) \leq [k_\alpha(a) \vee k_\alpha(b)]\lambda_0 T.$$

This shows that $r_\alpha(Q_2 \| P) \leq (k_\alpha(a) \vee k_\alpha(b))\lambda_0$. The equality in (3.4) now follows by taking $\lambda(t,z) = a\lambda_0$ or $\lambda(t,z) = b\lambda_0$.

As for the claim regarding delayed renewal processes, it is well known (see, e.g., [26], Exercise 2.14) that the RN derivative is given by

$$(A.3) \quad \Lambda_T = e^{-\int_{[0,T] \times S} (h(V_t)\psi(z) - \lambda_0) \zeta(dz) dt + \int_{[0,T] \times S} \log \frac{h(V_t)\psi(z)}{\lambda_0} N(dt dz)},$$

where $V_t = t - \tau_{N_t(S)}$ is the backward recurrence time. Hence the process $\lambda(t,z) \equiv h(V_{t-})\psi(z)$ is the intensity process for the marked point process N , satisfying the hypothesis $a \leq \frac{\lambda(\cdot)}{\lambda_0} \leq b$. The result now follows from the first part of (ii).

(iii) Fix $Q \in \mathcal{Q}_3$. Since $k_\alpha(x)$ is convex, we have

$$k_\alpha(x) \leq \frac{x-a}{b-a} k_\alpha(b) + \frac{b-x}{b-a} k_\alpha(a).$$

Therefore

$$\begin{aligned} \frac{1}{T} \int_{[0, T] \times S} k_\alpha \left(\frac{\lambda(t, z)}{\lambda_0} \right) \zeta(dz) dt &\leq \frac{1}{T} \int_{[0, T] \times S} \left(\frac{\frac{\lambda(t, z)}{\lambda_0} - a}{b - a} k_\alpha(b) + \frac{b - \frac{\lambda(t, z)}{\lambda_0}}{b - a} k_\alpha(a) \right) \zeta(dz) dt \\ &\leq \frac{\frac{\lambda_0 + v(T)}{\lambda_0} - a}{b - a} k_\alpha(b) + \frac{b - \frac{\lambda_0 - v(T)}{\lambda_0}}{b - a} k_\alpha(a) \\ &= pk_\alpha(a) + qk_\alpha(b) + \frac{v(T)(k_\alpha(a) + k_\alpha(b))}{\lambda_0(b - a)} \end{aligned}$$

by (3.5). It then follows from (A.2) and the P -(local) martingale property of M that

$$\frac{1}{T} \frac{1}{\alpha(\alpha - 1)} \log E_P[\Lambda_T^\alpha] \leq (pk_\alpha(a) + qk_\alpha(b))\lambda_0 + \frac{v(T)(k_\alpha(a) + k_\alpha(b))}{b - a}.$$

Taking the supremum over $Q \in \mathcal{Q}_3$ and the limsup as $T \rightarrow \infty$, it follows that $r_\alpha(\mathcal{Q}_3 \| P) \leq (pk_\alpha(a) + qk_\alpha(b))\lambda_0$. To obtain the asserted equality, take, for each T , deterministic $\lambda(\cdot)$ that takes the value $a\lambda_0$ (resp., $b\lambda_0$) on $[0, pT]$ (resp., $[pT, T]$).

(iv) Fix $1 < \alpha < \alpha_0$. Let $Q \in \mathcal{Q}_4$ and let $\bar{\lambda}(\cdot)$ be the normalized intensity process $\bar{\lambda}(\cdot) = \lambda(\cdot)/\lambda_0$. Then for all $T > 0$,

$$\begin{aligned} \frac{1}{T} \int_{[0, T] \times S} k_{\alpha_0}(\bar{\lambda}(t, z)) \zeta(dz) dt &\leq u + v(T), \\ \left| \frac{1}{T} \int_{[0, T] \times S} \bar{\lambda}(t, z) \zeta(dz) dt - 1 \right| &\leq \bar{v}(T) = \frac{v(T)}{\lambda_0}. \end{aligned}$$

Denote by \mathcal{G} the collection of all (deterministic) maps $f : \mathbb{R}_+ \times S \rightarrow \mathbb{R}_+$ such that for all $T \in (0, \infty)$

$$\begin{aligned} \frac{1}{T} \int_{[0, T] \times S} k_{\alpha_0}(f(t, z)) \zeta(dz) dt &\leq u + v(T), \\ \left| \frac{1}{T} \int_{[0, T] \times S} f(t, y) \zeta(dy) dt - 1 \right| &\leq \bar{v}(T) \end{aligned} \tag{A.4}$$

and let

$$U_T = \sup \left\{ \int_{[0, T] \times S} k_\alpha(f(t, z)) \zeta(dz) dt : f \in \mathcal{G} \right\}.$$

Since the normalized stochastic intensity $\bar{\lambda}(\cdot)$ is in \mathcal{G} a.s., $\int_{[0, T] \times S} k_\alpha(\bar{\lambda}(t, z)) \zeta(dz) dt \leq U_T$, for every T . Consequently, by (A.2), $\Lambda_T^\alpha \leq M_T e^{\lambda_0 \alpha (\alpha - 1) U_T}$, for every T . Since M is a non-negative local martingale,

$$\frac{1}{T} \frac{1}{\alpha(\alpha - 1)} \log E_P[\Lambda_T^\alpha] \leq \frac{\lambda_0 U_T}{T}. \tag{A.5}$$

We now compute U_T . To this end, for $f \in \mathcal{G}$, let the measure $\mu = \mu_T$ be the corresponding empirical measure on $[0, T] \times S$, namely,

$$\mu(B) = \frac{1}{T} \int_{[0, T] \times S} 1_{\{f(t, z) \in B\}} \zeta(dz) dt, \quad B \in \mathcal{B}([0, \infty)).$$

Let the p th moment of μ be denoted by $m_p(\mu) = \int_{[0, \infty)} x^p d\mu(x)$. Then by (A.4), $\langle k_{\alpha_0}, \mu \rangle \doteq \int k_{\alpha_0} d\mu \leq u + v(T)$ and $|m_1(\mu) - 1| \leq \bar{v}(T)$. The computation proceeds in two steps. First we solve the problem of maximizing $m_\alpha(\mu)$ under the constraints that $m_1(\mu)$ and $m_{\alpha_0}(\mu)$ are given. Then we translate it into the problem of maximizing $\int k_\alpha(x) d\mu(x)$ subject to the constraints (A.4).

Let a, b, k, l be positive constants satisfying $a + b = \alpha$, $ka = 1$, $lb = \alpha_0$ and $k^{-1} + l^{-1} = 1$. Using Hölder inequality,

$$m_\alpha(\mu) = \int x^\alpha d\bar{\mu}(x) = \int x^a x^b d\mu(x) \leq m_1(\mu)^{1/k} m_{\alpha_0}(\mu)^{1/l}.$$

Solving for a, b, k, l gives $k = \frac{\alpha_0 - 1}{\alpha_0 - \alpha}$, $l = \frac{\alpha_0 - 1}{\alpha - 1}$ (and $a = k^{-1}$, $b = \alpha_0 l^{-1}$). Moreover, the inequality is tight, specifically

(A.6)
$$\mu(dx) = p\delta_0(dx) + q\delta_c(dx),$$

satisfies it with equality, with $1 - p = q = m_1 C^{-\frac{1}{\alpha_0 - 1}}$ and $c = C^{\frac{1}{\alpha_0 - 1}}$, where $C = \frac{m_{\alpha_0}}{m_1}$ (note: using the inequality $m_1^{\alpha_0} \leq m_{\alpha_0}$ it is easy to check that $q \leq 1$).

Next, recalling the notation $\bar{\alpha} = \alpha(\alpha - 1)$ and $\bar{\alpha}_0 = \alpha_0(\alpha_0 - 1)$,

$$\begin{aligned} \langle k_\alpha, \mu \rangle &= \frac{1}{\bar{\alpha}} (m_\alpha(\mu) - \alpha m_1(\mu) + \alpha - 1) \\ &\leq \frac{1}{\bar{\alpha}} (m_1(\mu)^{1/k} m_{\alpha_0}(\mu)^{1/l} - \alpha m_1(\mu) + \alpha - 1) \\ &= \frac{1}{\bar{\alpha}} [m_1(\mu)^{1/k} (\bar{\alpha}_0 \langle k_{\alpha_0}, \mu \rangle + \alpha_0 m_1(\mu) - \alpha_0 + 1)^{1/l} - \alpha m_1(\mu) + \alpha - 1]. \end{aligned}$$

We now use the fact that $v \in \mathcal{V}_0$, and the assumed bounds on $m_1(\mu)$ and $\langle k_{\alpha_0}, \mu \rangle$. We obtain

$$\langle k_\alpha, \mu \rangle \leq \frac{1}{\bar{\alpha}} [(\bar{\alpha}_0 u + 1)^{1/l} - 1] + \tilde{v}(T)$$

for suitable $\tilde{v} \in \mathcal{V}_0$, which depends on the parameter but not on μ . Combining with (A.5),

$$\frac{1}{T} \frac{1}{\alpha(\alpha - 1)} \log E[\Lambda_T^\alpha] \leq \frac{\lambda_0}{\bar{\alpha}} [(\bar{\alpha}_0 u + 1)^{\frac{\alpha - 1}{\alpha_0 - 1}} - 1] + \lambda_0 \tilde{v}(T).$$

Taking supremum over $Q \in \mathcal{Q}_4$ and the limsup as $T \rightarrow \infty$ gives

$$r_\alpha(Q_4 \| P) \leq \frac{\lambda_0}{\bar{\alpha}} [(\bar{\alpha}_0 u + 1)^{\frac{\alpha - 1}{\alpha_0 - 1}} - 1].$$

Finally, equality is obtained by selecting, for each T , a deterministic $\lambda(\cdot)$ that agrees with (A.6) in the sense that the empirical measure μ corresponding to $\bar{\lambda} = \lambda/\lambda_0$ is given by (A.6). □

A.2. Proofs of results from Section 3.2. Before presenting the proof of Theorem 3.2, we state and prove the following lemma. Recall the notation P_t^N and Q_t^N from Section 3.1. Let $0 = \tau_0 < \tau_1 < \tau_2 < \dots$ denote the occurrence times of the point process N and let for $i \in \mathbb{N}$, $\Delta_i = \tau_i - \tau_{i-1}$. Write N_t for $N_t(S)$ for short when there is no ambiguity.

LEMMA A.1. *Assume that $\bar{H} \doteq \sup_{x \in \mathbb{R}_+} H(x) < \infty$. Also suppose that $c(\alpha) \doteq \int_S (\psi^\alpha(z) - 1) \zeta(dz) < \infty$. Then for every $\alpha > 1$,*

$$\frac{1}{t} R_\alpha(Q_t^N \| P_t^N) = \frac{1}{\alpha(\alpha - 1)t} \log E_P[e^{\alpha \sum_{i=1}^{N_t} H(\Delta_i)}] + \frac{c(\alpha)}{\alpha(\alpha - 1)} + o_t(1)$$

as $t \rightarrow \infty$.

PROOF. For fixed $t \geq 0$, let $\eta_t \doteq \tau_{N_t+1} = \inf\{s > t : N_s > N_t\}$. Then η_t is a $\{\mathcal{F}_s\}_{s \geq 0}$ -stopping time. Recall the notation $\Lambda_t = \frac{dQ_t^N}{dP_t^N}$ and the expression (A.3). By the optional sampling theorem it follows that $E_P(\Lambda_{\eta_t}) = 1$ for every $t \geq 0$ and for every $s \geq 0$ and $A \in \mathcal{F}_{s \wedge \eta_t}$

$$E_P(1_A \Lambda_{\eta_t}) = E_P(1_A \Lambda_{s \wedge \eta_t}) = E_P(1_A \Lambda_s) = E_Q(1_A).$$

By a monotone class argument we now have that, with $K_s \doteq N_{s \wedge \eta_t}(S)$ for $s \geq 0$ and $\mathcal{G} \doteq \sigma\{K_s : s \geq 0\}$,

$$E_P[1_A \Lambda_{\eta_t}] = E_Q[1_A] \quad \forall A \in \mathcal{G}.$$

Since $\sigma\{N_s : 0 \leq s \leq t\}$ is contained in $\sigma\{K_s : s \geq 0\}$, by the data processing inequality [29], Theorem 1.24 and Corollary 1.29, and [41], Section II, we have

$$\begin{aligned} R_\alpha(Q_t^N \| P_t^N) &= R_\alpha(Q \circ N_{[0,t]}^{-1} \| P \circ N_{[0,t]}^{-1}) \\ &\leq R_\alpha(Q \circ \{K_s : s \geq 0\}^{-1} \| P \circ \{K_s : s \geq 0\}^{-1}) \leq \frac{\log E_P[\Lambda_{\eta_t}^\alpha]}{\alpha(\alpha - 1)}. \end{aligned}$$

Denote by $\{\xi_i\}$ the sequence of marks associated with the point process.

Using the expression of Λ_t in (A.3) and the definition of H in (3.6), we have

$$\begin{aligned} \Lambda_{\eta_t}^\alpha &= \exp \left\{ \alpha \sum_{i=1}^{N_t+1} \left(\int_{\tau_{i-1}}^{\tau_i} (1 - h(V_s)) ds + \log h(\tau_i - \tau_{i-1}) + \log \psi(\xi_i) \right) \right\} \\ &= \exp \left\{ \alpha \sum_{i=1}^{N_t+1} H(\Delta_i) \right\} \exp \left\{ \alpha \sum_{i=1}^{N_t+1} \log \psi(\xi_i) \right\}, \quad t \geq 0. \end{aligned}$$

Since marks are independent of jump instants and H is bounded from above, we have

$$\frac{1}{t} \log E_P \Lambda_{\eta_t}^\alpha = \frac{1}{t} \log E_P [e^{\alpha \sum_{i=1}^{N_t} H(\Delta_i)}] + \frac{1}{t} \log E_P [e^{\alpha \sum_{i=1}^{N_t+1} \log \psi(\xi_i)}] + o_t(1).$$

Also, by standard Laplace transform formulas (see, e.g., [26], Example 1.16)

$$\frac{1}{t} \log E_P [e^{\alpha \sum_{i=1}^{N_t+1} \log \psi(\xi_i)}] = c(\alpha) + o_t(1).$$

The result follows. \square

PROOF OF THEOREM 3.2. (a) From Lemma A.1

$$(A.7) \quad \frac{1}{t} R_\alpha(Q_t^N \| P_t^N) = \frac{1}{\alpha(\alpha - 1)t} \log E_P [e^{\alpha \sum_{i=1}^{N_t} H(\Delta_i)}] + \frac{c(\alpha)}{\alpha(\alpha - 1)} + o_t(1).$$

Bounding $H(\Delta_i)$ by \bar{H} , we have $E_P [e^{\alpha \sum_{i=1}^{N_t} H(\Delta_i)}] \leq E_P [e^{\alpha \bar{H} N_t}] = e^{t(e^{\alpha \bar{H}} - 1)}$ and therefore

$$\limsup_t \frac{1}{t} R_\alpha(Q_t^N \| P_t^N) \leq \frac{e^{\alpha \bar{H}} - 1 + c(\alpha)}{\alpha(\alpha - 1)}.$$

This gives the bound (3.7).

For parts (b), (c) and (d), we will need a more careful analysis of $E_P [e^{\alpha \sum_{i=1}^{N_t} H(\Delta_i)}]$, under different assumptions.

Fix $0 < c_0 < c_1$ and write

$$(A.8) \quad E_P [e^{\alpha \sum_{i=1}^{N_t} H(\Delta_i)}] = \sum_{k=0}^{\infty} E_P [1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}] = Q_0(t) + Q_1(t) + Q_2(t),$$

where

$$Q_0(t) = \sum_{k \leq c_0 t} E_P [1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}],$$

$$Q_1(t) = \sum_{k \geq c_1 t} E_P[1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}],$$

$$Q_2(t) = \sum_{c_0 t < k < c_1 t} E_P[1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}].$$

Using the bound $Q_0(t) \leq e^{\alpha \bar{H} c_0 t}$, we have

$$(A.9) \quad \limsup_t \frac{1}{t} \log Q_0(t) \leq \alpha \bar{H} c_0.$$

For bounding $Q_1(t)$, write

$$Q_1(t) \leq \sum_{k \geq c_1 t} e^{\alpha \bar{H} k} e^{-t} \frac{t^k}{k!} \leq c_{st} e^{-t} \sum_{k \geq c_1 t} e^{\alpha \bar{H} k} t^k e^k k^{-k-\frac{1}{2}},$$

where Stirling's approximation is used, and c_{st} is a universal constant. Using $t/k \leq 1/c_1$ for the summands in the above display, we have

$$Q_1(t) \leq c_{st} e^{-t} \sum_{k \geq c_1 t} e^{\alpha \bar{H} k} e^k c_1^{-k}.$$

For $c_1 \geq 2e^{\alpha \bar{H} + 1}$, we have $c_1^{-1} e^{\alpha \bar{H} + 1} < \frac{1}{2}$ and $Q_1(t) \leq c_{st}$. Hence

$$(A.10) \quad \limsup_t \frac{1}{t} \log Q_1(t) \leq 0 \quad \forall c_1 \geq 2e^{\alpha \bar{H} + 1}.$$

We now estimate $Q_2(t)$, using different approaches under different assumptions in parts (b), (c) and (d).

(b) Note that

$$\begin{aligned} \frac{1}{t} \log Q_2(t) &\leq \frac{1}{t} \log \left((c_1 - c_0)t + 1 \right) \max_{k \in [c_0 t, c_1 t]} E_P[1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}] \\ &= o_t(1) + \frac{1}{t} \log \max_{k \in [c_0 t, c_1 t]} E_P[1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}]. \end{aligned}$$

Recall that $\gamma(s) \doteq E_P e^{s H(\Delta_1)}$ for $s \in \mathbb{R}$. Let $p > 0$ and $q > 0$ be such that $1/p + 1/q = 1$. Then for each $k \in [c_0 t, c_1 t]$,

$$\begin{aligned} E_P[1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}] &\leq P(N_t = k)^{1/p} \gamma(q\alpha)^{k/q} \\ &= (e^{-t} t^k / k!)^{1/p} \gamma(q\alpha)^{k/q} \\ &\leq (e^{-t} t^k c_{st} e^k k^{-k-\frac{1}{2}})^{1/p} \gamma(q\alpha)^{k/q}. \end{aligned}$$

Letting $\theta \doteq k/t$

$$\begin{aligned} \frac{1}{t} \log E_P[1_{\{N_t=k\}} e^{\alpha \sum_{i=1}^k H(\Delta_i)}] &\leq -\frac{1}{p} + \frac{\theta}{p} \log t + \frac{\theta}{p} - \frac{\theta}{p} \log(\theta t) + \frac{\theta}{q} \log \gamma(q\alpha) + o_t(1) \\ &= \frac{\theta - 1}{p} - \frac{\theta}{p} \log \theta + \frac{\theta}{q} \log \gamma(q\alpha) + o_t(1). \end{aligned}$$

We now maximize the sum of the first three terms on the last line over θ (with p and q fixed). The maximum is attained at $\theta = \gamma(q\alpha)^{p/q}$. If we plug in this value of θ we obtain that the maximum is given by

$$\frac{\gamma(q\alpha)^{p/q} - 1}{p} = \hat{\gamma}(p, q, \alpha).$$

As a result,

$$\limsup_t \frac{1}{t} \log Q_2(t) \leq \inf_{p,q>1:p^{-1}+q^{-1}=1} \hat{\gamma}(p, q, \alpha) = G_\alpha^{(1)}.$$

Combining this with the bounds (A.7), (A.8), (A.9) and (A.10) gives

$$r_\alpha^N(Q\|P) \leq \left[\frac{\bar{H}c_0}{\alpha-1} \vee 0 \vee \frac{G_\alpha^{(1)}}{\alpha(\alpha-1)} \right] + \frac{c(\alpha)}{\alpha(\alpha-1)}.$$

Sending $c_0 \rightarrow 0$ and $c_1 \rightarrow \infty$ gives (3.8).

(c) Given $\varepsilon > 0$ let $c_0 = \theta_0 < \theta_1 < \dots < \theta_J = c_1$ be a finite partition of $[c_0, c_1]$ satisfying $\theta_j - \theta_{j-1} = \varepsilon$ for all $j \leq J$. Then

$$(A.11) \quad \limsup_t \frac{1}{t} \log Q_2(t) \leq \max_{1 \leq j \leq J} \limsup_t \frac{1}{t} \log Q_2^j(t),$$

where

$$Q_2^j(t) = E_P[1_{\{\theta_{j-1}t \leq N_t < \theta_j t\}} e^{\alpha \sum_{i=1}^{N_t} H(\Delta_i)}].$$

Fix $j \leq J$. Denote $n = \lceil \theta_{j-1}t \rceil$. Let $S_n = \sum_{i=1}^n \Delta_i$ and $S_n^H = \sum_{i=1}^n H(\Delta_i)$. Use bar to denote the normalized sum, as in $\bar{S}_n = n^{-1}S_n$. Since β is finite in a neighborhood of the origin, by Cramér’s theorem, (\bar{S}_n, \bar{S}_n^H) has LDP with a good rate function (see, e.g., [12], Corollary 6.1.6) β^* . Note that $\{\theta_{j-1}t \leq N_t < \theta_j t\} \subset \{S_{\lceil \theta_{j-1}t \rceil} \leq t, S_{\lceil \theta_j t \rceil} \geq t\}$. Then

$$(A.12) \quad Q_2^j(t) \leq e^{\alpha \bar{H} \varepsilon t} E_P[1_{\{S_{\lceil \theta_{j-1}t \rceil} \leq t, S_{\lceil \theta_j t \rceil} \geq t\}} e^{\alpha \sum_{i=1}^n H(\Delta_i)}] \leq e^{\alpha \bar{H} \varepsilon t} E_P[1_{\{S_n \leq t\}} e^{\alpha S_n^H}].$$

Let g be the upper semicontinuous function defined as $g(x_1, x_2) = 0$ for $0 \leq x_1 \leq \theta_{j-1}^{-1}$ and $g(x_1, x_2) = -\infty$ otherwise. Then $1_{\{S_n \leq t\}} \leq e^{ng(\bar{S}_n, \bar{S}_n^H)}$. Hence

$$Q_2^j(t) \leq e^{\alpha \bar{H} \varepsilon t} E_P[e^{n(\alpha \bar{S}_n^H + g(\bar{S}_n, \bar{S}_n^H))}].$$

Since $\bar{H} < \infty$ the conditions of Varadhan’s integral lemma (see, e.g., [12], Lemma 4.3.6) are valid and hence

$$(A.13) \quad \begin{aligned} \limsup_t \frac{1}{t} \log Q_2^j(t) &\leq \alpha \bar{H} \varepsilon + \theta_{j-1} \sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta_{j-1}^{-1}} [\alpha x_2 - \beta^*(x)] \\ &\leq \alpha \bar{H} \varepsilon + G_\alpha^{(2)}(\theta_{j-1}). \end{aligned}$$

From this and (A.11) we have

$$\limsup_t \frac{1}{t} \log Q_2(t) \leq \max_{1 \leq j \leq J} \limsup_t \frac{1}{t} \log Q_2^j(t) \leq \alpha \bar{H} \varepsilon + \sup_{\theta \in [c_0, c_1]} G_\alpha^{(2)}(\theta).$$

Combining this with the bounds (A.7), (A.8), (A.9) and (A.10) and sending $\varepsilon \rightarrow 0$ gives

$$\limsup_t \frac{1}{t} R_\alpha(Q_t^N \| P_t^N) \leq \left[\frac{\bar{H}c_0}{\alpha-1} \vee 0 \vee \sup_{\theta \in [c_0, c_1]} \frac{G_\alpha^{(2)}(\theta)}{\alpha(\alpha-1)} \right] + \frac{c(\alpha)}{\alpha(\alpha-1)}.$$

Sending $c_0 \rightarrow 0$ and $c_1 \rightarrow \infty$ gives (3.9).

(d) Let $m = \lceil \theta_j t \rceil$ and $S_m, S_m^H, \bar{S}_m, \bar{S}_m^H$ be defined in a similar manner as in part (c). Once more we use the fact that (\bar{S}_m, \bar{S}_m^H) has a LDP with a rate function β^* . Note that besides the bound (A.12), we also have, with $p^{-1} + q^{-1} = 1, p, q > 1$,

$$(A.14) \quad \begin{aligned} Q_2^j(t) &\leq (E_P[1_{\{S_n \leq t, S_m \geq t\}} e^{p\alpha \sum_{i=1}^m H(\Delta_i)}])^{1/p} \\ &\quad \times (E_P[1_{\{\theta_{j-1}t \leq N_t < \theta_j t\}} e^{-q\alpha \sum_{i=N_t+1}^m H(\Delta_i)}])^{1/q}. \end{aligned}$$

For the first term on the right hand side, we apply Varadhan’s integral lemma as in the proof of part (c) and get

$$\limsup_t \frac{1}{t} \log(E_P[1_{\{S_n \leq t, S_m \geq t\}} e^{p\alpha \sum_{i=1}^m H(\Delta_i)}])^{1/p} \leq \frac{\theta_j}{p} \sup_{x \in \mathbb{R}^2: x_1 \geq \theta_j^{-1}} [p\alpha x_2 - \beta^*(x)].$$

For the second term on the right hand side of (A.14), we have

$$\begin{aligned} & \limsup_t \frac{1}{t} \log(E_P[1_{\{\theta_{j-1}t \leq N_t < \theta_j t\}} e^{-q\alpha \sum_{i=N_t+1}^m H(\Delta_i)}])^{1/q} \\ & \leq \limsup_t \frac{1}{qt} \log\left([\theta_j - \theta_{j-1}]t + 1\right) \max_{\theta_{j-1}t \leq k < \theta_j t} [\gamma(-q\alpha)]^{m-k} \\ & \leq \limsup_t \frac{1}{qt} \log([\varepsilon t + 1][\gamma(-q\alpha)]^{\varepsilon t} \vee 1) \\ & = \frac{\log^+ \gamma(-q\alpha)}{q} \varepsilon. \end{aligned}$$

Combining these two bounds with (A.14) gives

$$\limsup_t \frac{1}{t} \log Q_2^j(t) \leq \theta_j \sup_{x \in \mathbb{R}^2: x_1 \geq \theta_j^{-1}} \left[\alpha x_2 - \frac{1}{p} \beta^*(x) \right] + \frac{\log^+ \gamma(-q\alpha)}{q} \varepsilon.$$

Combining this with (A.13), we have

$$\begin{aligned} \limsup_t \frac{1}{t} \log Q_2^j(t) & \leq \left(\alpha \bar{H} \varepsilon + \theta_{j-1} \sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta_{j-1}^{-1}} [\alpha x_2 - \beta^*(x)] \right) \\ & \quad \wedge \left(\theta_j \sup_{x \in \mathbb{R}^2: x_1 \geq \theta_j^{-1}} \left[\alpha x_2 - \frac{1}{p} \beta^*(x) \right] + \frac{\log^+ \gamma(-q\alpha)}{q} \varepsilon \right). \end{aligned}$$

Since $p > 1$, we have

$$\begin{aligned} & \limsup_t \frac{1}{t} \log Q_2^j(t) \\ & \leq \alpha \bar{H} \varepsilon + \frac{\log^+ \gamma(-q\alpha)}{q} \varepsilon + \theta_j \left(\sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta_{j-1}^{-1}} \left[\alpha x_2 - \frac{1}{p} \beta^*(x) \right] \right. \\ \text{(A.15)} \quad & \quad \left. \wedge \sup_{x \in \mathbb{R}^2: x_1 \geq \theta_j^{-1}} \left[\alpha x_2 - \frac{1}{p} \beta^*(x) \right] \right) \\ & \leq \alpha \bar{H} \varepsilon + \frac{\log^+ \gamma(-q\alpha)}{q} \varepsilon + \theta_j \sup_{\theta \in [\theta_{j-1}, \theta_j]} \left(\sup_{x \in \mathbb{R}^2: 0 \leq x_1 \leq \theta^{-1}} \left[\alpha x_2 - \frac{1}{p} \beta^*(x) \right] \right. \\ & \quad \left. \wedge \sup_{x \in \mathbb{R}^2: x_1 \geq \theta^{-1}} \left[\alpha x_2 - \frac{1}{p} \beta^*(x) \right] \right), \end{aligned}$$

where the second inequality uses the fact that for a function r defined on \mathbb{R} and constants $a < b$, one has

$$\sup_{y \in (-\infty, b]} r(y) \wedge \sup_{y \in [a, \infty)} r(y) \leq \sup_{z \in [a, b]} \left(\sup_{y \in (-\infty, z]} r(y) \wedge \sup_{y \in [z, \infty)} r(y) \right).$$

Since $x \mapsto x_2 - \frac{1}{p}\beta^*(x)$ is a concave function on \mathbb{R}^2 , the last term in (A.15) equals

$$\begin{aligned} & \theta_j \sup_{\theta \in [\theta_{j-1}, \theta_j]} \sup_{x_2 \in \mathbb{R}} \left[\alpha x_2 - \frac{1}{p}\beta^*(\theta^{-1}, x_2) \right] \\ & \leq \frac{\theta_j}{\theta_{j-1}} \sup_{\theta \in [\theta_{j-1}, \theta_j]} \theta \sup_{x_2 \in \mathbb{R}} \left[\alpha x_2 - \frac{1}{p}\beta^*(\theta^{-1}, x_2) \right] \\ & \leq \left(1 + \frac{\varepsilon}{c_0} \right) \sup_{\theta \in [\theta_{j-1}, \theta_j]} \theta \sup_{x_2 \in \mathbb{R}} \left[\alpha x_2 - \frac{1}{p}\beta^*(\theta^{-1}, x_2) \right]. \end{aligned}$$

From this, (A.11) and (A.15), letting

$$G_\alpha^{(3)}(p, \theta) \doteq \theta \sup_{x_2 \in \mathbb{R}} \left[\alpha x_2 - \frac{1}{p}\beta^*(\theta^{-1}, x_2) \right],$$

we have

$$\begin{aligned} \limsup_t \frac{1}{t} \log Q_2(t) & \leq \max_{1 \leq j \leq J} \limsup_t \frac{1}{t} \log Q_2^j(t) \\ & \leq \alpha \bar{H} \varepsilon + \frac{\log^+ \gamma(-q\alpha)}{q} \varepsilon + \left(1 + \frac{\varepsilon}{c_0} \right) \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta). \end{aligned}$$

Combining this with the bounds (A.7), (A.8), (A.9) and (A.10) and sending $\varepsilon \rightarrow 0$ gives

$$\limsup_t \frac{1}{t} R_\alpha(Q_t^N \| P_t^N) \leq \left[\frac{\bar{H}c_0}{\alpha - 1} \vee 0 \vee \sup_{\theta \in [c_0, c_1]} \frac{G_\alpha^{(3)}(p, \theta)}{\alpha(\alpha - 1)} \right] + \frac{c(\alpha)}{\alpha(\alpha - 1)}.$$

Now we claim that

$$(A.16) \quad \lim_{p \rightarrow 1} \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta) = \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(\theta).$$

Once this claim is verified, sending $p \rightarrow 1$, $c_0 \rightarrow 0$ and $c_1 \rightarrow \infty$ gives (3.10).

It remains to prove the claim (A.16). First note that for any $\theta_0 \in [c_0, c_1]$ and $x_2 \in \mathbb{R}$,

$$\liminf_{p \rightarrow 1} \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta) \geq \liminf_{p \rightarrow 1} \theta_0 \left[\alpha x_2 - \frac{1}{p}\beta^*(\theta_0^{-1}, x_2) \right] = \theta_0 [\alpha x_2 - \beta^*(\theta_0^{-1}, x_2)].$$

Taking supremum over $x_2 \in \mathbb{R}$ and $\theta_0 \in [c_0, c_1]$ gives

$$(A.17) \quad \liminf_{p \rightarrow 1} \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta) \geq \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(\theta).$$

Since β^* is a good rate function, we can find $\kappa_0 \in (-\infty, 0)$ such that, for all $p \in [1/2, 2]$, $\sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta) \geq \kappa_0$.

Next we show $\limsup_{p \rightarrow 1} \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta) \leq \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(\theta)$. For $p \in [1/2, 2]$, $p \neq 1$, let $\theta_p \in [c_0, c_1]$ and $x_{2,p} \in \mathbb{R}$ be such that

$$(A.18) \quad \kappa_0 \leq \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta) \leq \theta_p \left[\alpha x_{2,p} - \frac{1}{p}\beta^*(\theta_p^{-1}, x_{2,p}) \right] + |p - 1|.$$

From $\beta^* \geq 0$ we have $x_{2,p} \geq \frac{\kappa_0 - 1}{c_0 \alpha}$. Since $\beta^*(x_1, x_2) \geq \lambda_2 x_2 - \beta(0, \lambda_2)$ for each $\lambda_2 > 0$, and $\beta(0, \lambda_2) < \infty$ for all $\lambda_2 > 0$, we have $\lim_{x_2 \rightarrow \infty} \inf_{x_1 \in \mathbb{R}} \frac{\beta^*(x_1, x_2)}{x_2} = \infty$. This shows that $x_{2,p}$ is bounded from above, since if $x_{2,p} \rightarrow \infty$ as $p \rightarrow 1$ then from (A.18) we must have

$$\limsup_{p \rightarrow 1} \frac{\beta^*(\theta_p^{-1}, x_{2,p})}{p x_{2,p}} \leq \theta_p \alpha$$

which is a contradiction. Hence $x_{2,p} \leq \kappa_1$ for some $\kappa_1 < \infty$ and therefore the sequence $\{(\theta_p, x_2^p)\}$ is bounded. Assume without loss of generality that $(\theta_p, x_2^p) \rightarrow (\bar{\theta}, \bar{x}_2) \in [c_0, c_1] \times \mathbb{R}$ along the whole subsequence. Then

$$\begin{aligned} \limsup_{p \rightarrow 1} \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(p, \theta) &\leq \limsup_{p \rightarrow 1} \left(\theta_p \left[\alpha x_{2,p} - \frac{1}{p} \beta^*(\theta_p^{-1}, x_{2,p}) \right] + |p - 1| \right) \\ &= \bar{\theta} \left[\alpha \bar{x}_2 - \liminf_{p \rightarrow 1} \beta^*(\theta_p^{-1}, x_{2,p}) \right] \\ &\leq \bar{\theta} \left[\alpha \bar{x}_2 - \beta^*(\bar{\theta}^{-1}, \bar{x}_2) \right] \\ &\leq \sup_{\theta \in [c_0, c_1]} G_\alpha^{(3)}(\theta), \end{aligned}$$

where the second inequality follows from the lower semicontinuity of β^* . Combining this with (A.17) gives the claim (A.16). This completes the proof. \square

A.3. Proofs of results from Section 3.3. In this section we provide details of some of the calculations that were omitted from Section 3.3.

PROOFS OF STATEMENTS IN EXAMPLE 3.1. We first consider the case $\rho > 1$ and show that the right sides of (3.9) and (3.10) are the same and the inequalities in both cases can be replaced by equalities.

Note that $h(x) = \rho > 0$, $H(x) = -(\rho - 1)x + \log \rho \leq \log \rho < \infty$ and β is finite in a neighborhood of the origin, namely all assumptions for (3.9) hold. Then, as follows from (A.1), (A.2) in the Appendix,

$$\begin{aligned} \frac{1}{t} R_\alpha(Q_t^N \| P_t^N) &= \frac{1}{\alpha(\alpha - 1)t} \log E_P[\Lambda_t^\alpha] = \frac{1}{\alpha(\alpha - 1)t} \log E_P[e^{\alpha(1-\rho)t + \alpha N_t \log \rho}] \\ &= \frac{1}{\alpha(\alpha - 1)} [\rho^\alpha - 1 - \alpha(\rho - 1)]. \end{aligned}$$

Thus from Theorem 3.2(c) $\rho^\alpha - 1 - \alpha(\rho - 1) \leq \sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta)$. Now we show the reverse inequality.

If $\theta \geq 1 + \alpha(\rho - 1)$, taking $\lambda_1 = 1 + \alpha(\rho - 1) - \theta \leq 0$ in (3.11) gives

$$\begin{aligned} G_\alpha^{(2)}(\theta) &\leq -\lambda_1 + \theta \beta(\lambda_1, \alpha) = -1 - \alpha(\rho - 1) + \theta + \theta \log E_P[e^{(1-\theta)\Delta + \alpha \log \rho}] \\ &= -1 - \alpha(\rho - 1) + \theta + \theta[\alpha \log \rho - \log \theta] \leq \rho^\alpha - 1 - \alpha(\rho - 1), \end{aligned}$$

where the last inequality becomes equality when $\theta = \rho^\alpha$. If $0 < \theta < 1 + \alpha(\rho - 1)$, taking $\lambda_1 = 0$ in (3.11) gives

$$\begin{aligned} G_\alpha^{(2)}(\theta) &\leq \theta \beta(0, \alpha) = \theta \log E_P[e^{-\alpha(\rho-1)\Delta + \alpha \log \rho}] = \theta \log \frac{\rho^\alpha}{1 + \alpha(\rho - 1)} \\ &\leq [1 + \alpha(\rho - 1)] \log \frac{\rho^\alpha}{1 + \alpha(\rho - 1)}, \end{aligned}$$

where the last inequality follows from the fact that $\rho^\alpha = (1 + (\rho - 1))^\alpha \geq 1 + \alpha(\rho - 1)$. Using $\ell(x) \doteq x \log x - x + 1 \geq 0$, the term on the right side of the last display can be written as

$$-\rho^\alpha \ell\left(\frac{1 + \alpha(\rho - 1)}{\rho^\alpha}\right) - [1 + \alpha(\rho - 1)] + \rho^\alpha \leq \rho^\alpha - 1 - \alpha(\rho - 1).$$

Therefore $\sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta) \leq \rho^\alpha - 1 - \alpha(\rho - 1)$. Thus we have shown that the inequality in (3.9) is in fact an equality. From this and the observation that $G_\alpha^{(2)}(\theta) \geq G_\alpha^{(3)}(\theta)$ we see that in (3.10) also the inequality can be replaced with an equality.

Consider now the case $\rho \in (0, 1]$. We show that once more the right sides of (3.9) and (3.10) are the same and the inequalities in both cases can be replaced by equalities. The proof of $\sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta) \leq \rho^\alpha - 1 - \alpha(\rho - 1)$ for (3.9) is exactly as before. For (3.10), observe first that

$$\beta(\lambda_1, \lambda_2) = \lambda_2 \log \rho - \log[1 - \lambda_1 + \lambda_2(\rho - 1)], \quad 1 - \lambda_1 + \lambda_2(\rho - 1) > 0,$$

$$\beta^*(x_1, \log \rho - (\rho - 1)x_1) = x_1 - 1 - \log x_1, \quad x_1 > 0.$$

Therefore

$$\begin{aligned} \sup_{\theta \in (0, \infty)} G_\alpha^{(3)}(\theta) &\geq G_\alpha^{(3)}(\rho^\alpha) = \rho^\alpha \sup_{x_2 \in \mathbb{R}} \left[\alpha x_2 - \beta^*\left(\frac{1}{\rho^\alpha}, x_2\right) \right] \\ &\geq \rho^\alpha \left[\alpha \left(\log \rho - \frac{\rho - 1}{\rho^\alpha} \right) - \beta^*\left(\frac{1}{\rho^\alpha}, \log \rho - \frac{\rho - 1}{\rho^\alpha}\right) \right] \\ &= \rho^\alpha \left[\alpha \left(\log \rho - \frac{\rho - 1}{\rho^\alpha} \right) - \left(\frac{1}{\rho^\alpha} - 1 + \log \rho^\alpha \right) \right] = -\alpha(\rho - 1) - 1 + \rho^\alpha, \end{aligned}$$

and hence

$$\text{LHS(3.10)} = \frac{\rho^\alpha - 1 - \alpha(\rho - 1)}{\alpha(\alpha - 1)} = \text{RHS(3.10)} = \text{RHS(3.9)}. \quad \square$$

PROOF OF (3.12) IN EXAMPLE 3.2. We will use Theorem 3.2(c) and establish (3.12) by estimating $\sup_{\theta \in (0, \infty)} G_\alpha^{(2)}(\theta)$.

If $\theta \geq \frac{1 + \alpha(\rho - 1)}{1 + \alpha(k - 1)}$, taking $\lambda_1 = 1 + \alpha(\rho - 1) - \theta[1 + \alpha(k - 1)] \leq 0$ in (3.11) gives

$$\begin{aligned} G_\alpha^{(2)}(\theta) &\leq -\lambda_1 + \theta\beta(\lambda_1, \alpha) \\ &= -1 - \alpha(\rho - 1) + \theta[1 + \alpha(k - 1)] + \theta \left\{ \alpha k \log \rho + \log \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \right) \right. \\ &\quad \left. - (1 + \alpha(k - 1)) \log[\theta(1 + \alpha(k - 1))] \right\} \\ &\leq \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1 + \alpha(k - 1)}} - \alpha(\rho - 1) - 1, \end{aligned}$$

where the last inequality is attained for $\theta = \theta^*$ that satisfies

$$(A.19) \quad \alpha k \log \rho + \log \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \right) - (1 + \alpha(k - 1)) \log[\theta^*(1 + \alpha(k - 1))] = 0.$$

One can check that θ^* is indeed greater than or equal to $\frac{1 + \alpha(\rho - 1)}{1 + \alpha(k - 1)}$. To see this, note that the left hand side in (A.19) is decreasing in θ . So it suffices to check

$$\alpha k \log \rho + \log \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \right) - (1 + \alpha(k - 1)) \log[(1 + \alpha(\rho - 1))] \geq 0.$$

This is equivalent to checking $\beta(0, \alpha) \geq 0$. But this is immediate since

$$\beta(0, \alpha) = \log \int (g(x)e^x)^\alpha e^{-x} dx \geq \log \left(\int g(x)e^x e^{-x} dx \right)^\alpha = 0$$

by Holder's inequality.

If $0 < \theta < \frac{1+\alpha(\rho-1)}{1+\alpha(k-1)}$, taking $\lambda_1 = 0$ in (3.11) gives

$$\begin{aligned}
 G_\alpha^{(2)}(\theta) &\leq \theta\beta(0, \alpha) \\
 &= \theta \left[\alpha k \log \rho + \log \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \right) \right. \\
 &\quad \left. - (1 + \alpha(k - 1)) \log(1 + \alpha(\rho - 1)) \right].
 \end{aligned}
 \tag{A.20}$$

Since the left hand side in (A.19) is decreasing in θ , the expression obtained by replacing θ by $\frac{1+\alpha(\rho-1)}{1+\alpha(k-1)}$ in this term is nonnegative, which shows that the term on the right side of (A.20) is nonnegative. Therefore

$$\begin{aligned}
 G_\alpha^{(2)}(\theta) &\leq \frac{1 + \alpha(\rho - 1)}{1 + \alpha(k - 1)} \left[\alpha k \log \rho + \log \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \right) \right. \\
 &\quad \left. - (1 + \alpha(k - 1)) \log(1 + \alpha(\rho - 1)) \right] \\
 &= -[1 + \alpha(\rho - 1)] \log \left\{ [1 + \alpha(\rho - 1)] / \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1+\alpha(k-1)}} \right\} \\
 &= - \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1+\alpha(k-1)}} \\
 &\quad \times \ell \left([1 + \alpha(\rho - 1)] / \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1+\alpha(k-1)}} \right) \\
 &\quad - [1 + \alpha(\rho - 1)] + \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1+\alpha(k-1)}} \\
 &\leq \left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1+\alpha(k-1)}} - \alpha(\rho - 1) - 1,
 \end{aligned}$$

where the third line uses the equality

$$-a \log(a/b) = -b \ell(a/b) - a + b.
 \tag{A.21}$$

Combining the above estimates with Theorem 3.2(c) we now have that when $\pi = \text{Gamma}(k, \rho)$ with $k \geq 1$ and $\rho > 1$

$$\begin{aligned}
 r_\alpha^N(Q \| P) &\leq \sup_{\theta \in (0, \infty)} \frac{1}{\alpha(\alpha - 1)} G_\alpha^{(2)}(\theta) \\
 &\leq \frac{1}{\alpha(\alpha - 1)} \left[\left(\frac{\Gamma(1 + \alpha(k - 1))}{(\Gamma(k))^\alpha} \rho^{\alpha k} \right)^{\frac{1}{1+\alpha(k-1)}} - \alpha(\rho - 1) - 1 \right]. \quad \square
 \end{aligned}$$

PROOF OF (3.13) IN EXAMPLE 3.3. Consider first $\theta \geq 1 + \alpha(\sigma - 1)$. In this case, taking $\lambda_1 = 1 + \alpha(\sigma - 1) - \theta \leq 0$ in (3.11), we have

$$G_\alpha^{(2)}(\theta) \leq -\lambda_1 + \theta\beta(\lambda_1, \alpha) \leq -1 - \alpha(\sigma - 1) + \theta + \theta[\alpha \log C - \log \theta] \leq C^\alpha - 1 - \alpha(\sigma - 1),$$

where the last inequality is attained when $\theta = C^\alpha$. Note that C^α is indeed in the range $[1 + \alpha(\sigma - 1), \infty)$. To see this note that

$$1 = \int g(x) dx \leq \int C e^{-\sigma x} dx = \frac{C}{\sigma}$$

which shows that

$$(A.22) \quad C^\alpha \geq \sigma^\alpha = (1 + \sigma - 1)^\alpha \geq 1 + \alpha(\sigma - 1).$$

Now consider the case $0 < \theta < 1 + \alpha(\sigma - 1)$. In this case, taking $\lambda_1 = 0$ in (3.11), we have

$$G_\alpha^{(2)}(\theta) \leq \theta \beta(0, \alpha) \leq \theta \log \frac{C^\alpha}{1 + \alpha(\sigma - 1)} \leq [1 + \alpha(\sigma - 1)] \log \frac{C^\alpha}{1 + \alpha(\sigma - 1)},$$

where the last inequality follows from (A.22). Recalling $\ell(x) = x \log x - x + 1 \geq 0$ and using the equality in (A.21) once more, from the last display we have for all $0 < \theta < 1 + \alpha(\sigma - 1)$

$$G_\alpha^{(2)}(\theta) \leq -C^\alpha \ell\left(\frac{1 + \alpha(\sigma - 1)}{C^\alpha}\right) - [1 + \alpha(\sigma - 1)] + C^\alpha \leq C^\alpha - 1 - \alpha(\sigma - 1).$$

Combining the above estimates with Theorem 3.2(c) we have the bound (3.13) on RDR for this class of models. \square

Funding. Research of RA was supported in part by ISF Grant 1184/16.

Research of AB was supported in part by NSF Grants DMS-1305120, DMS-1814894 and DMS-1853968.

Research of PD was supported in part by NSF Grant DMS-1904992 and AFOSR Grant FA-9550-18-1-0214.

Research of RW was supported in part by DARPA Grant W911NF-15-2-0122.

REFERENCES

- [1] ANANTHARAM, V. (1989). How large delays build up in a $GI/G/1$ queue. *Queueing Syst.* **5** 345–367. [MR1030475 https://doi.org/10.1007/BF01225324](https://doi.org/10.1007/BF01225324)
- [2] ATAR, R., BUDHIRAJA, A., DUPUIS, P. and WU, R. (2019). Large deviations for the single server queue and the reneging paradox. Preprint. Available at [arXiv:1903.06870](https://arxiv.org/abs/1903.06870).
- [3] ATAR, R., CHOWDHARY, K. and DUPUIS, P. (2015). Robust bounds on risk-sensitive functionals via Rényi divergence. *SIAM/ASA J. Uncertain. Quantificat.* **3** 18–33. [MR3299141 https://doi.org/10.1137/130939730](https://doi.org/10.1137/130939730)
- [4] ATAR, R. and DUPUIS, P. (1999). Large deviations and queueing networks: Methods for rate function identification. *Stochastic Process. Appl.* **84** 255–296. [MR1719274 https://doi.org/10.1016/S0304-4149\(99\)00051-4](https://doi.org/10.1016/S0304-4149(99)00051-4)
- [5] ATAR, R., GOSWAMI, A. and SHWARTZ, A. (2013). Risk-sensitive control for the parallel server model. *SIAM J. Control Optim.* **51** 4363–4386. [MR3134419 https://doi.org/10.1137/130926705](https://doi.org/10.1137/130926705)
- [6] ATAR, R., GOSWAMI, A. and SHWARTZ, A. (2014). On the risk-sensitive cost for a Markovian multiclass queue with priority. *Electron. Commun. Probab.* **19** Art. ID 11. [MR3174829 https://doi.org/10.1214/ECP.v19-2905](https://doi.org/10.1214/ECP.v19-2905)
- [7] ATAR, R., KASPI, H. and SHIMKIN, N. (2014). Fluid limits for many-server systems with reneging under a priority policy. *Math. Oper. Res.* **39** 672–696. [MR3246997 https://doi.org/10.1287/moor.2013.0630](https://doi.org/10.1287/moor.2013.0630)
- [8] ATAR, R., MANDELBAUM, A. and REIMAN, M. I. (2004). Scheduling a multi class queue with many exponential servers: Asymptotic optimality in heavy traffic. *Ann. Appl. Probab.* **14** 1084–1134. [MR2071417 https://doi.org/10.1214/105051604000000233](https://doi.org/10.1214/105051604000000233)
- [9] BROWN, L., GANS, N., MANDELBAUM, A., SAKOV, A., SHEN, H., ZELTYN, S. and ZHAO, L. (2005). Statistical analysis of a telephone call center: A queueing-science perspective. *J. Amer. Statist. Assoc.* **100** 36–50. [MR2166068 https://doi.org/10.1198/016214504000001808](https://doi.org/10.1198/016214504000001808)
- [10] BUDHIRAJA, A. and DUPUIS, P. (2019). *Analysis and Approximation of Rare Events: Representations and Weak Convergence Methods. Probability Theory and Stochastic Modelling* **94**. Springer, New York. [MR3967100 https://doi.org/10.1007/978-1-4939-9579-0](https://doi.org/10.1007/978-1-4939-9579-0)
- [11] CHOWDHARY, K. and DUPUIS, P. (2013). Distinguishing and integrating aleatoric and epistemic variation in uncertainty quantification. *ESAIM Math. Model. Numer. Anal.* **47** 635–662. [MR3056403 https://doi.org/10.1051/m2an/2012038](https://doi.org/10.1051/m2an/2012038)
- [12] DEMBO, A. and ZEITOUNI, O. (1993). *Large Deviations Techniques and Applications*. Jones and Bartlett Publishers, Boston, MA. [MR1202429](https://doi.org/10.1007/978-1-4612-0000-0)

- [13] DUPUIS, P. and ELLIS, R. S. (1995). Large deviation analysis of queueing systems. In *Stochastic Networks. IMA Vol. Math. Appl.* **71** 347–365. Springer, New York. MR1381019 https://doi.org/10.1007/978-1-4757-2418-9_17
- [14] DUPUIS, P. and ELLIS, R. S. (1995). The large deviation principle for a general class of queueing systems. *I. Trans. Amer. Math. Soc.* **347** 2689–2751. MR1290716 <https://doi.org/10.2307/2154753>
- [15] DUPUIS, P., JAMES, M. R. and PETERSEN, I. (2000). Robust properties of risk-sensitive control. *Math. Control Signals Systems* **13** 318–332. MR1804350 <https://doi.org/10.1007/PL00009872>
- [16] DUPUIS, P., KATSOLAKIS, M. A., PANTAZIS, Y. and PLECHÁČ, P. (2016). Path-space information bounds for uncertainty quantification and sensitivity analysis of stochastic dynamics. *SIAM/ASA J. Uncertain. Quantificat.* **4** 80–111. MR3455143 <https://doi.org/10.1137/15M1025645>
- [17] DUPUIS, P., KATSOLAKIS, M. A., PANTAZIS, Y. and REY-BELLET, L. (2020). Sensitivity analysis for rare events based on Rényi divergence. *Ann. Appl. Probab.* **30** 1507–1533. MR4132633 <https://doi.org/10.1214/19-AAP1468>
- [18] DVIJOTHAM, K. and TODOROV, E. (2011). A unified theory of linearly solvable optimal control. In *Uncertainty in Artificial Intelligence (UAI)*.
- [19] GANESH, A., O’CONNELL, N. and WISCHIK, D. (2004). *Big Queues. Lecture Notes in Math.* **1838**. Springer, Berlin. MR2045489 <https://doi.org/10.1007/b95197>
- [20] GLYNN, P. W. and WHITT, W. (1994). Logarithmic asymptotics for steady-state tail probabilities in a single-server queue. *J. Appl. Probab.* **31A** 131–156. MR1274722 <https://doi.org/10.2307/3214953>
- [21] HALFIN, S. and WHITT, W. (1981). Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* **29** 567–588. MR0629195 <https://doi.org/10.1287/opre.29.3.567>
- [22] IGNATIOUK-ROBERT, I. (2000). Large deviations of Jackson networks. *Ann. Appl. Probab.* **10** 962–1001. MR1789985 <https://doi.org/10.1214/aoap/1019487515>
- [23] IYENGAR, G. N. (2005). Robust dynamic programming. *Math. Oper. Res.* **30** 257–280. MR2142033 <https://doi.org/10.1287/moor.1040.0129>
- [24] JACQUET, P., SEROUSSI, G. and SZPANKOWSKI, W. (2008). On the entropy of a hidden Markov process. *Theoret. Comput. Sci.* **395** 203–219. MR2424508 <https://doi.org/10.1016/j.tcs.2008.01.012>
- [25] KANG, W. and RAMANAN, K. (2010). Fluid limits of many-server queues with reneging. *Ann. Appl. Probab.* **20** 2204–2260. MR2759733 <https://doi.org/10.1214/10-AAP683>
- [26] KARR, A. F. (1986). *Point Processes and Their Statistical Inference. Probability: Pure and Applied* **2**. Dekker, New York. MR0851982
- [27] KASPI, H. and RAMANAN, K. (2011). Law of large numbers limits for many-server queues. *Ann. Appl. Probab.* **21** 33–114. MR2759196 <https://doi.org/10.1214/09-AAP662>
- [28] KASPI, H. and RAMANAN, K. (2013). SPDE limits of many-server queues. *Ann. Appl. Probab.* **23** 145–229. MR3059233 <https://doi.org/10.1214/11-AAP821>
- [29] LIESE, F. and VAJDA, I. (1987). *Convex Statistical Distances. Teubner-Texte zur Mathematik [Teubner Texts in Mathematics]* **95**. BSB B. G. Teubner Verlagsgesellschaft, Leipzig. MR0926905
- [30] LIM, A. E. B. and SHANTHIKUMAR, J. G. (2007). Relative entropy, exponential utility, and robust dynamic pricing. *Oper. Res.* **55** 198–214. MR2316254 <https://doi.org/10.1287/opre.1070.0385>
- [31] MAJEWSKI, K. (2006). Sample path large deviations for multiclass feedforward queueing networks in critical loading. *Ann. Appl. Probab.* **16** 1893–1924. MR2288708 <https://doi.org/10.1214/105051606000000439>
- [32] NILIM, A. and EL GHAOUI, L. (2005). Robust control of Markov decision processes with uncertain transition matrices. *Oper. Res.* **53** 780–798. MR2171651 <https://doi.org/10.1287/opre.1050.0216>
- [33] ORDENTLICH, E. and WEISSMAN, T. (2004). New bounds on the entropy rate of hidden Markov processes. In *IEEE Information Theory Workshop* 117–122. IEEE, Piscataway, NJ.
- [34] PETERSEN, I. R., JAMES, M. R. and DUPUIS, P. (2000). Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Trans. Automat. Control* **45** 398–412. MR1762853 <https://doi.org/10.1109/9.847720>
- [35] PETERSEN, I. R., UGRINOVSKII, V. A. and SAVKIN, A. V. (2000). *Robust Control Design Using H^∞ -Methods. Communications and Control Engineering Series*. Springer, London. MR1834840 <https://doi.org/10.1007/978-1-4471-0447-6>
- [36] PUHALSKII, A. A. (2007). The action functional for the Jackson network. *Markov Process. Related Fields* **13** 99–136. MR2321753
- [37] REED, J. (2009). The $G/GI/N$ queue in the Halfin–Whitt regime. *Ann. Appl. Probab.* **19** 2211–2269. MR2588244 <https://doi.org/10.1214/09-AAP609>
- [38] RÉNYI, A. (1961). On measures of entropy and information. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. I* 547–561. Univ. California Press, Berkeley, CA. MR0132570
- [39] SENNOTT, L. I. (1999). *Stochastic Dynamic Programming and the Control of Queueing Systems. Wiley Series in Probability and Statistics: Applied Probability and Statistics*. Wiley, New York. MR1645435

- [40] SHWARTZ, A. and WEISS, A. (1995). *Large Deviations for Performance Analysis: Queues, Communications, and Computing. Stochastic Modeling Series*. CRC Press, London. [MR1335456](#)
- [41] VAN ERVEN, T. and HARREMOËS, P. (2010). Rényi divergence and majorization. In *2010 IEEE International Symposium on Information Theory Proceedings (ISIT)* 1335–1339. IEEE, Piscataway, NJ.