

Empirical Gittins for Data-Driven M/G/1 Scheduling with Arbitrary Job Size Distributions

SHEFALI RAMAKRISHNA*, Cornell University, USA

AMIT HARLEV*, Cornell University, USA

ZIV SCULLY, Cornell University, USA

When scheduling in the M/G/1 queue to minimize mean response time, a classic result is that if job sizes are unknown, then the *Gittins* policy is optimal. Gittins is best described as a policy construction: it takes as input the queue's *job size distribution*, and it outputs a job prioritization rule that optimizes mean response time for that particular distribution. But in practice, instead of knowing the exact job size distribution, one usually only has *samples* from it. We therefore ask: given finitely many samples from the job size distribution, how can one construct a scheduling policy with near-optimal mean response time?

Our main result is that to achieve near-optimal mean response time, it suffices to simply *apply the Gittins construction to the empirical distribution* of the job size samples. We call this policy *empirical Gittins*, and we prove an explicit high-probability bound on its mean response time. Our bound implies convergence to the optimal mean response time as one increases the number of samples. We also show that if one has even vague knowledge of the true distribution's tail asymptotics, one can make empirical Gittins more robust using truncation, resulting in better convergence rates.

It is surprising that empirical Gittins works well even for continuous job size distributions. This is because the Gittins construction is sensitive to the distribution's density, yet the empirical distribution, being discrete, cannot possibly approximate a continuous density. Our main technical contribution is to show that despite its sensitivity to density, the Gittins construction yields a good policy as long as one gives it a distribution with an approximately correct *tail*, even if the density is completely wrong. Underlying this finding are two new extensions of the WINE queueing identity.

CCS Concepts: • **General and reference** → **Performance**; Estimation; • **Mathematics of computing** → *Nonparametric statistics*; **Queueing theory**; • **Computing methodologies** → *Model development and analysis*; • **Software and its engineering** → **Scheduling**.

Additional Key Words and Phrases: scheduling; M/G/1 queue; mean response time; Gittins index; data-driven decision-making; finite-sample analysis; nonparametric methods

ACM Reference Format:

Shefali Ramakrishna, Amit Harlev, and Ziv Scully. 2026. Empirical Gittins for Data-Driven M/G/1 Scheduling with Arbitrary Job Size Distributions. *Proc. ACM Meas. Anal. Comput. Syst.* 10, 1, Article 9 (March 2026), 38 pages. <https://doi.org/10.1145/3788091>

1 Introduction

A central goal of queueing theory is to design scheduling policies that reduce *mean response time*, the average time between each job's arrival and departure, under *job size uncertainty*, namely when each job's size (a.k.a. service time) is unknown to the scheduler. A landmark result is that in the

* Authors contributed equally to this research.

Authors' Contact Information: Shefali Ramakrishna, School of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA; Amit Harlev, Center for Applied Mathematics, Cornell University, Ithaca, NY, USA; Ziv Scully, School of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 2476-1249/2026/3-ART9

<https://doi.org/10.1145/3788091>

preemptive M/G/1 queue, the *Gittins* policy provably minimizes mean response time [1, 12, 13, 41, 48, 51]. The idea behind Gittins is to leverage the *job size distribution*, from which each job's size is (by assumption) sampled i.i.d., to prioritize jobs that are likely to be nearly complete. One can thus view Gittins as a distributional counterpart to shortest remaining processing time (SRPT) [34, 35].

However, in practice, one seldom has access to an exact job size distribution. Instead, it is more likely that one has access to *samples* from it, such as from jobs completed in the past. This motivates the question: how should one use job sizes samples to design a scheduling policy that reduces mean response time? Are samples alone sufficient to approximate Gittins?

We take a first step towards answering these questions, proving that *empirical Gittins*, namely Gittins applied to the (possibly truncated) empirical distribution of job size samples, is sufficient to approximate Gittins applied to the true job size distribution. In the rest of this section, we describe:

- (Section 1.1) How we formulate data-driven scheduling as a one-shot decision problem.
- (Section 1.2) Why the one-shot formulation is a useful problem to study.
- (Section 1.3) Why it is not a priori obvious whether or not empirical Gittins is effective, and why tools from prior work are insufficient to analyze empirical Gittins.
- (Section 1.4) The new technical tools we develop to analyze empirical Gittins, and the results we prove using them.

1.1 Data-driven preemptive scheduling as a one-shot decision problem

We consider scheduling in an M/G/1 queue with job size distribution F (represented as a CDF). In brief, we study the problem of using n i.i.d. samples from F to construct a scheduling policy that has near-optimal mean response time with high probability. This is a one-shot problem because we require the policy constructed from the n samples to be “static” in a certain sense, described more precisely below, which rules out learning from additional job size samples observed.

In this paper, most of the scheduling policies we consider, including the Gittins policy, are *SOAP policies* (Section 2.1), a specific class of preemptive index policies [43]. When specialized to our setting, a SOAP policy is one that assigns each job a numerical priority, called its *rank* (lower is better), as a function of its *age* (a.k.a. attained service), the amount of time the job has been in service so far. This mapping is called the *rank function* of the SOAP policy π , denoted

$$r_\pi : \text{age } a \mapsto \text{rank } r_\pi(a). \quad (1.1)$$

That is, at every moment in time, π preemptively serves the job whose age maps to the lowest (i.e. best) rank under r_π . Even though an individual job's rank may change as the job's age increases, a SOAP policy is “static” in the sense that the rank function itself never changes.

The Gittins policy can be viewed as a *SOAP policy construction*, which we denote by γ , that maps

$$\gamma : \text{distribution } F \mapsto \text{SOAP policy } \gamma(F).$$

In most works on the Gittins policy in the M/G/1, the policy being studied is $\gamma(F)$, namely the Gittins construction γ applied to the true job size distribution F . In this work, we often call $\gamma(F)$ the *true Gittins* policy to distinguish it from γ applied to other distributions. True Gittins is optimal for mean response time: for any nonclairvoyant policy π (including non-SOAP policies), we have

$$\text{cost}_F(\gamma(F)) \leq \text{cost}_F(\pi),$$

where $\text{cost}_F(\pi)$ is the mean response time of policy π in an M/G/1 with job size distribution F and some fixed arrival rate that we leave unspecified for now. (In the notation of Section 2, $\text{cost}_F(\pi) = \mathbb{E}_\pi[T]$.)

Our goal is to achieve performance multiplicatively close to that of $\gamma(F)$, but with access to only n samples from F instead of access to F itself. Specifically, we seek a policy construction

$$\beta : \text{samples } s_1, \dots, s_n \mapsto \text{SOAP policy } \beta(s_1, \dots, s_n)$$

such that when $S_1, \dots, S_n \sim F$ are sampled i.i.d. from F , we have a PAC-type result

$$\mathbb{P}_{S_1, \dots, S_n \sim F} \left[\frac{\text{cost}_F(\beta(S_1, \dots, S_n))}{\text{cost}_F(\gamma(F))} < 1 + \varepsilon \right] \geq 1 - \delta \quad \text{for all job size distributions } F,$$

where we would like $\delta, \varepsilon > 0$ to be small when n is large. This as a one-shot decision problem that can be viewed as having three stages:

- *Sample time*: we receive as input the n samples S_1, \dots, S_n .
- *Design time*: we use the samples to choose a SOAP policy $\beta(S_1, \dots, S_n)$.
- *Run time*: we achieve mean response time $\text{cost}_F(\beta(S_1, \dots, S_n))$, which is a deterministic function of our chosen policy $\beta(S_1, \dots, S_n)$.

In particular, this is not a learning problem, because we commit to a SOAP policy at design time, without updating it from further samples at run time.

Our main finding is that a good choice for β is what we call the *empirical Gittins* policy, namely

$$\beta(S_1, \dots, S_n) = \gamma(\text{mixture of } n \text{ point masses at } S_1, \dots, S_n).$$

We also obtain results for the *truncated empirical Gittins* policy, which throws out a few of the largest samples before applying γ . Truncation is helpful in theory, especially when F has infinite variance (Section 6), but we find in simulation that the untruncated version tends to perform better (Section 7).

1.2 Why study a one-shot formulation?

In a practical system, one would of course want to use all available data to optimize one's scheduling policy, including data collected at run time. Indeed, there is recent work that studies the intersection of online learning and scheduling or dispatching, in both centralized [25, 55, 57] and decentralized [9–11, 47] settings. And looking beyond queues, in the Markovian multi-armed bandit problem [12, 13], there is work on learning the Gittins policy from online samples [6, 7, 33]. Why, then, do we study a problem where decisions are made only once at design time?

In our view, the high-level reason to study data-driven scheduling is that real systems have nonstationary arrival processes. With stationary arrivals, one could simply learn the true job size distribution to arbitrary accuracy over time. But with nonstationary arrivals, there is an incentive to quickly adapt to changes in the job size distribution when they occur. Therefore, the main question studied by much prior work on learning in queues is *how quickly a given policy can adapt* to an unknown but stationary arrival process. This requires *transient analysis*, which is technically challenging because the scheduling or dispatching policy changes over time as it sees more job size samples. Directly studying data-driven scheduling with nonstationary arrivals is harder still, with only Yang et al. [55] doing so, to the best of our knowledge.

Our one-shot formulation is designed to still capture the notion of making the best possible use of limited data, but with *no transient or nonstationary analysis*. There are two specific difficulties that make such analysis very difficult with current techniques.

- Prior work on learning while scheduling studies settings where only parametric estimation is necessary, e.g. learning service rates of multiclass M/M arrivals. In contrast, we make essentially no assumptions on the job size distribution, so we need *nonparametric* estimation.




Fig. 1.1. Rank functions (Sections 1.1 and 2.1) of true Gittins and empirical Gittins policies. Each policy works by using its rank function to map each job’s age to a numerical rank, then serves the job of least rank. The true job size distribution is bounded Pareto (Section 7.1), from which the empirical distribution is generated with 1000 i.i.d. samples. Despite the empirical and true distributions having similar tails (e.g. as measured by Kolmogorov distance), the two rank functions are very different: while true Gittins’s rank function is continuous, empirical Gittins’s rank function has a jump discontinuity from zero to a nonzero value at each job size sample.

- It is often clear what the optimal scheduling policy is with nonstationary M/M arrivals of known time-varying parameters, e.g. prioritizing by service rate. In contrast, with nonstationary M/G arrivals with known time-varying parameters, the *optimal baseline is unknown*, even if only the arrival rate is varying.

The apparent intractability of handling transient or nonstationary analysis is what motivates our one-shot problem formulation, which lets us use the growing set of queueing theoretic tools for analyzing mean response time of $M/G/1$ scheduling policies [16, 18, 22, 36, 49, 50, 54]. However, because mean response time captures only long-run average behavior, we cannot allow arbitrary policies, as one could eventually learn the true job size distribution to arbitrary precision over time. As such, to make the design-time policy decision meaningful, we restrict our attention to a class of index policies that do not learn over time, meaning all the sample-dependent decisions must be made at design time. Our problem thus captures the core task of using a limited number of samples to construct a good scheduling policy, but without requiring transient or nonstationary analysis.

1.3 Obstacles to analyzing empirical Gittins

Having explained that the one-shot problem formulation we study is motivated in part by the numerous queueing theoretic tools available for analyzing mean response time, it is natural to ask: why is analyzing empirical Gittins still challenging? The main obstacle is that empirical Gittins can have a very different priority structure from that of true Gittins, as illustrated in Fig. 1.1, and the two tools best suited to analyzing Gittins are insufficient for comparing such different policies.

One relevant tool is the *SOAP analysis* [43], which gives a mean response time formula for any SOAP policy. The inputs to the formula are the arrival rate, the (true) job size distribution F , and the SOAP policy’s rank function (Sections 1.1 and 2.1). While the SOAP formula is suitable for comparing concrete policies [17, 39, 42], it is difficult to use the formula to compare *policy constructions*, because the formula’s dependence on the rank function is complicated. The only example we are aware of is the work of Scully et al. [44], who compare the Gittins construction to a simpler alternative, but the analysis depends on the two constructions yielding similar rank

functions. Unfortunately, true Gittins and empirical Gittins need not have similar rank functions, as shown in Fig. 1.1, which makes the SOAP approach appear not promising.

Another relevant tool is the *WINE identity* [36], which, when specialized to our setting, expresses the mean number of jobs in a system (and thereby mean response time [27]) in terms of quantities related to the true Gittins policy’s rank function. See Section 3.1 for a statement of WINE as it applies to our setting, and see Scully [36, Section 2.2.3 and Chapter 4] for a broader overview of WINE and its several independent discoveries [3, 14, 15, 32, 37, 41].

The most important fact about WINE for our purposes is that it implies the true Gittins policy has a “multiplicatively robust” rank function: if $r_\pi(a)/r_{\gamma(F)} \in [e^{-\epsilon}, e^\epsilon]$ for all ages a , then $\text{cost}_F(\pi)/\text{cost}_F(\gamma(F)) \leq e^{2\epsilon}$ [36, 39, 46]. Moreover, Ramakrishna et al. [30] show that two distributions F and G induce multiplicatively close Gittins rank functions $\gamma(F)$ and $\gamma(G)$ if they have multiplicatively close *densities*, specifically Radon-Nikodym derivative in $[e^{-\epsilon/2}, e^{\epsilon/2}]$. This is enough to show that empirical Gittins performs well when the true distribution F has finite (discrete) support, in which case the empirical distribution has density close to F ’s with high probability. But when F is continuous, *empirical Gittins and true Gittins do not have multiplicatively close rank functions*, as illustrated in Fig. 1.1. This is because discrete distributions always induce Gittins rank functions which approach zero prior to each atom, whereas continuous distributions with everywhere-nonzero density induce everywhere-nonzero Gittins rank functions [2]. So for general distributions F , prior methods that rely on multiplicative closeness of rank functions are insufficient to compare empirical Gittins to true Gittins.

Having seen that SOAP and WINE seem unlikely to solve the problem on their own, we take a step back and ask: what properties inherently shared by the true and empirical distributions could potentially serve as the basis for comparing true and empirical Gittins? One clear candidate is that they have similar *tail functions* (a.k.a. complementary CDFs), as measured either additively [28?] or multiplicatively [53?]. This is promising, as Moseley et al. [29] recently show that in a batch version of unknown-size job scheduling (i.e. no arrivals), applying the Gittins construction to a distribution whose tail is multiplicatively close to F yields near-optimal performance. However, there are still two obstacles to overcome. First, the approach taken by Moseley et al. [29] is specific to the batch setting and seems difficult to directly generalize to our setting with M/G arrivals. Second, the results of Moseley et al. [29] require multiplicatively close tail *everywhere*, but the empirical distribution’s tail will be multiplicatively close to F ’s only up to a certain threshold [53, Lemmas 1–4]. This limitation is inevitable when F is unbounded, because the empirical distribution always has finite support.

1.4 Contributions and key innovation: two new flavors of WINE

To analyze empirical Gittins, we introduce *two new variants of the WINE identity*. Unlike classical WINE (Proposition 3.6), which is tied specifically to the rank function of $\gamma(F)$, our new variant can be tied to the rank function of $\gamma(G)$ for any distribution G that is close to F in the right sense.

Our first new WINE variant, which we call *multiplicatively close WINE* (Theorem 3.9), says that if F and G have multiplicatively close tails, then, roughly speaking, using $\gamma(G)$ instead of $\gamma(F)$ in standard WINE yields a quantity which is multiplicatively close to the number of jobs in the system. This nearly immediately implies that $\gamma(G)$ ’s mean response time is multiplicatively close to that of $\gamma(F)$ ’s (Theorem 3.1). While the empirical distribution’s tail is not multiplicatively close to F ’s tail everywhere, it is so up to some threshold with high probability. We use this to prove high-probability bounds on *truncated* empirical Gittins (Theorems 4.1, 6.1, and 6.2): multiplicatively close WINE handles jobs with size below the truncation threshold, and we handle larger jobs by choosing a “safe fallback policy” for ages beyond the threshold, specifically *preemptive last-come first-served*

(PLCFS). This enables high-probability bounds even for infinite-variance job size distributions, though our bounds require very many samples to imply near-optimal performance in this case.

However, multiplicatively close WINE alone is not enough to analyze empirical Gittins without truncation. This motivates our second new WINE variant, which we call *WINE for empirical distributions* (Theorem 5.5). As mentioned before, the main obstacle is that if G is the empirical distribution, then the tail ratio $\bar{G}(x)/\bar{F}(x)$ is far from 1 for large x . In general, such a property could cause using $\gamma(G)$ instead of $\gamma(F)$ in standard WINE to yield an arbitrarily poor estimate for the number of jobs. Our insight is to use the fact that G is not just any distribution, but *a mixture of n point masses*, which opens up another path to bounding the error of using $\gamma(G)$ instead of $\gamma(F)$ in standard WINE (Theorem 5.3). The error does grow with n , but only logarithmically, so it is more than canceled out by other terms that decrease polynomially in n . We use this to prove high-probability bounds on untruncated empirical Gittins’s mean response time (Theorems 5.1 and 6.3).

Our mean response time bound for untruncated empirical Gittins is looser than the one we show for the truncated version, and we require finite variance to obtain a high-probability bound.¹ But the untruncated version performs very well in simulation (Section 7), often beating the truncated version, and we view our untruncated bounds as a first step towards explaining why.

To summarize, our main contributions are the *first analyses of the empirical Gittins policy in the M/G/1*, both with and without truncation, as enabled by *two new variants of the WINE identity*. These constitute Sections 3–5 of the paper. Section 6 states the resulting bounds on (truncated) empirical Gittins as convergence rates to optimality as the number of samples n increases. Finally, Section 7 conducts a brief simulation study to evaluate the impact of truncation in practice.

2 Preliminaries

We consider an M/G/1 queue to which jobs arrive according to a Poisson process of rate $\lambda > 0$. Each job has a *size* (a.k.a. service time) sampled i.i.d. from a distribution F , represented as a CDF. We use the random variable S to denote a generic job’s size. By default, $S \sim F$, so $\mathbb{P}[S \leq x] = F(x)$; when we sample S from other distributions, we note it explicitly in a subscript. When we need to disambiguate F from other distributions, we call it the *true job size distribution*. We denote the system *load* by $\rho = \lambda \mathbb{E}[S]$. We assume throughout that $\rho < 1$, ensuring ergodicity and existence of a steady-state distribution.

The system has a single server that serves one job at a time, with the *scheduling policy* deciding which job to serve at every moment in time. We consider only *nonclairvoyant* policies, meaning those that do not have access to jobs’ sizes before they complete, nor to future arrival times. Instead, the policy knows only each job’s *age* (a.k.a. attained service), meaning the amount of time it has been served so far, and other past information, e.g. past arrival times. We assume fully preemptive scheduling with no overhead.

We use $\mathbb{E}_\pi[\cdot]$ and $\mathbb{P}_\pi[\cdot]$ to denote expectation and probability, respectively, under the stationary distribution under scheduling policy π . These are meant in the long-run average or steady-state senses (which are equivalent by ergodicity). Specifically, the *number of jobs*, denoted N , and *work*, denoted W (i.e. total remaining service time of all jobs), can be understood in a long-run time average sense; and the *response time*, denoted T , can be understood in a long-run average-over-jobs sense, or in a tagged-job sense [18].

For any set $U \subseteq \mathbb{R}_{\geq 0}$, $N(\text{size} \in U)$ denotes the number of jobs in the system whose size is in the set U , and $N(\text{age} \in U)$ denotes the number of jobs in the system whose age is in the set U .

¹One can obtain an in-expectation bound for the untruncated infinite-variance case, but we believe a high-probability bound is not possible with our techniques, so we omit it for brevity.

2.1 SOAP policies

A *SOAP* policy [36, 40, 43, 50] is a scheduling policy that works by assigning each job a priority as a function of its age. A SOAP policy π is specified by its *rank function*, denoted

$$r_\pi : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\},$$

which assigns to each age a a numerical *rank* $r_\pi(a)$, representing the priority of an age- a job. The scheduler *always serves the job of minimal rank*, i.e. lower rank is better priority.²

For the most part, when multiple jobs are tied for minimal rank, we break the tie in first-come first-served (FCFS) order, serving the job among them that arrived least recently. The only exception is that when multiple jobs are tied at rank ∞ , we break the tie using (*preemptive*) *last-come first-served* (PLCFS), serving the job among them that arrived most recently.³ We use this convention to allow the use of PLCFS as a “fallback option” for jobs that reach sufficiently large ages. In particular, for all SOAP policies in this paper, once a job reaches rank ∞ , it remains at rank ∞ thereafter.

2.2 True, empirical, and truncated job size distributions

We refer to any distribution supported on $\mathbb{R}_{\geq 0}$ as a *job size distribution*, which we represent as a cumulative distribution function (CDF). For a job size distribution G , we denote its *tail function* (a.k.a. complementary CDF) by $\bar{G}(x) = 1 - G(x)$. Aside from the true job size distribution F , the main other job size distribution we consider is the empirical distribution of samples from F .

Definition 2.1. Given n i.i.d. samples $S_1, \dots, S_n \sim F$, the *empirical job size distribution* is defined as the empirical CDF

$$x \mapsto \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{S_i \leq x\}.$$

We typically denote the empirical distribution by G when needed, though G also often denotes a generic job size distribution.

We will often compare two distributions that are close in tail behavior (usually the true and empirical distribution). The following definition formalizes this notion.

Definition 2.2. We describe distributions G and H as ε -*multiplicatively close* if for all $t \geq 0$

$$\frac{\bar{G}(t)}{\bar{H}(t)} \in [e^{-\varepsilon}, e^\varepsilon],$$

where we adopt the convention that $\frac{0}{0} = 1$.

If this bound only holds for $t \leq \ell$ for some $\ell > 0$, we say they are ε -*multiplicatively close up to* ℓ .

The notion of ε -multiplicative closeness is not introduced as an external modeling assumption. Rather, it arises naturally when comparing an empirical distribution to the true distribution using standard tools from empirical process theory [28, 53?]. As a result, multiplicative tail comparisons are a natural notion of distributional proximity that “fall out” of the problem when the empirical distribution is the only available approximation to F .

The following concentration result shows that, with high probability, the empirical job size distribution obtained from n samples from F is ε -multiplicatively close to F up to some threshold ℓ .

²Always serving the job of minimal rank can naturally lead to processor sharing in certain situations [43, Appendix B], but for the purposes of this paper, one can safely imagine processor sharing as rapidly swapping service between multiple jobs.

³This is a slight deviation from the usual presentation of SOAP policies, but all the relevant definitions and results easily generalize to this slightly different tie-breaking rule along the lines of Scully et al. [43, Appendix A].

LEMMA 2.3 (SAMPLE COMPLEXITY FOR ε -MULTIPLICATIVE CLOSENESS, [53, LEMMA 1]). *Let F denote the CDF of a job size distribution on \mathbb{R}_+ . Let G be the empirical CDF based on n i.i.d. samples from F . Then for any $\varepsilon \in (0, 0.6)$, we have*

$$\sup_{0 \leq x \leq \ell} \frac{\bar{G}(x)}{\bar{F}(x)} \in [e^{-\varepsilon}, e^{\varepsilon}] \quad \text{with probability at least } 1 - 2 \exp\left(\frac{-n\varepsilon^2 \bar{F}(\ell)}{3}\right).$$

PROOF SKETCH. This result is a corollary of Wellner [53, Lemma 1], rephrased for a general distribution F . For details, see the proof in Appendix D.1. \square

Because the concentration bound in Lemma 2.3 only guarantees closeness between the empirical and true distributions up to a finite threshold (as is the case for any finite-sample bound for empirical CDFs), we will sometimes work with *truncated* versions of these distributions.

Definition 2.4. For any job size distribution G and *truncation* at $\ell > 0$, the ℓ -truncated distribution, denoted G_ℓ , is defined by its tail function as

$$\bar{G}_\ell(x) = \begin{cases} \bar{G}(x) & x < \ell \\ 0 & x \geq \ell. \end{cases}$$

Thus G_ℓ coincides with G below ℓ and treats all jobs larger than ℓ as having size exactly ℓ .

2.3 The Gittins policy

In this paper, we focus specifically on the class of *Gittins* scheduling policies.

Definition 2.5. Let G be a generic job size distribution. The *Gittins policy constructed from G* , denoted by $\gamma(G)$, is a SOAP policy with rank function^{4,5}

$$r_{\gamma(G)}(a) = \begin{cases} \inf_{b>a} \frac{\mathbb{E}_{S \sim G}[(S \wedge b) - a | S > a]}{\mathbb{P}_{S \sim G}[S \leq b | S > a]} & \text{if } \bar{G}(a) > 0, \\ \infty & \text{otherwise.} \end{cases}$$

Note that by construction, for any Gittins policy $\gamma(G_\ell)$, jobs whose age exceed ℓ are subjected to the “PLCFS fallback” as in Section 2.1.

We often consider the Gittins policies constructed from certain specific distributions. In informal discussion, we use the following terminology.

- The Gittins policy constructed from the true distribution F , denoted $\gamma(F)$, is referred to as the *true Gittins* policy.
- The Gittins policy constructed from an empirical distribution is referred to as the *empirical Gittins* policy.
- The Gittins policy constructed from an ℓ -truncated empirical distribution for some ℓ is referred to as the *truncated empirical Gittins* policy.

Throughout this paper, we make use of the fact that the true Gittins policy is the optimal nonclairvoyant policy for mean response time.

THEOREM 2.6 (OPTIMALITY OF THE TRUE GITTINS POLICY, [36, THEOREM 16.1]). *Let F denote the true job size distribution in an $M/G/1$ queue. Among all nonclairvoyant scheduling policies π , the true Gittins policy $\gamma(F)$ minimizes the mean response time; that is:*

$$\mathbb{E}_{\gamma(F)}[T] \leq \mathbb{E}_\pi[T].$$

⁴The reciprocal of the Gittins rank is known as the Gittins index [1, 13].

⁵We use the notation $(x \wedge y)$ throughout the paper to mean $\min(x, y)$.

3 Robustness of Gittins to multiplicative tail perturbations

In this section we prove a new version of WINE, which we refer to as multiplicatively close WINE (Theorem 3.9), that is designed to analyze the performance of the Gittins policy $\gamma(G)$ where G is a job size distribution multiplicatively close to the true job size distribution F . We then use multiplicatively close WINE to prove the following theorem, which states that the response time of the Gittins policy in the M/G/1 is robust to multiplicative perturbations of the job size distribution's tail.

THEOREM 3.1. *If G is a job size distribution that is ε -multiplicatively close to F , then*

$$\frac{\mathbb{E}_{\gamma(G)}[T]}{\mathbb{E}_{\gamma(F)}[T]} \in [1, e^{4\varepsilon}].$$

While Theorem 3.1 is not directly applicable when G is an empirical distribution, for which multiplicative closeness holds only up to a threshold (Lemma 2.3), it serves as a stepping stone towards our analysis of truncated empirical Gittins in Section 4.

3.1 Introduction to WINE

WINE, which stands for *Work Integral Number Equality*, is a queueing identity that can relate the number of jobs in the system to certain rank functions [36, Section 4]. We start by introducing the “classical” WINE from prior work, which relates the number of jobs in the system to the rank function of true Gittins. However, our presentation uses more general notation, which will be necessary to state and prove our new extensions of WINE.

The key concept underlying WINE is *r-work*, which is roughly “work if each job were stopped upon reaching rank r ”. Below, we define *r-work* for individual jobs and the system as a whole.

Definition 3.2. Let $r \geq 0$ and let G be a job size distribution. The *remaining G-Gittins r-work* of a job at age a is the amount of service it needs until it either completes or reaches age

$$b_{\gamma(G)}(a, r) = \inf\{b \geq a : r_{\gamma(G)}(b) \geq r\},$$

the first age $\geq a$ whose rank under $\gamma(G)$ is at least r . Note that if the rank is already at least r at age a (i.e. if $r_{\gamma(G)}(a) \geq r$) then $b_{\gamma(G)}(a, r) = a$, and so the job's remaining *G-Gittins r-work* is 0.

Now let H also be a job size distribution. The *expected remaining G-Gittins r-work under H* for a job at age a is

$$s_{H,G}(a, r) = \mathbb{E}_{S \sim H}[(S \wedge b_{\gamma(G)}(a, r)) - a \mid S > a].$$

Definition 3.3. Let $r \geq 0$ and let G be a job size distribution. The *system G-Gittins r-work*, $W(\gamma(G) < r)$, is the sum of the remaining *G-Gittins r-work* of all jobs in the system. We will often care about the expected system *G-Gittins r-work* with respect to the true job size distribution F conditioned on knowing the ages of all jobs currently in the system. Let A_i be the age of i -th job currently in the system and let $\vec{A} = [A_1, \dots, A_N]$ be the list of all jobs' ages, where N is the number of jobs in the system. Then,

$$\mathbb{E}_{\pi}[W(\gamma(G) < r) \mid \vec{A}] = \sum_{i=1}^N s_{F,G}(A_i, r), \quad (3.1)$$

where π is an arbitrary (nonclairvoyant) policy. Notice π is absent from the right-hand side, so this does not depend on the policy π (beyond requiring nonclairvoyance).

The rough idea behind WINE is the following. First, we observe that a certain integral of a single job's expected *r-work* always yields 1, no matter the age of the job, an observation sometimes called “single-job WINE” [36]. Second, we use linearity of expectation to conclude that if we integrate

the expected r -work of the *entire system*, instead of the expected r -work of a single job, then we obtain the number of jobs in the system, because each job contributes 1. This is formalized in the propositions below.

PROPOSITION 3.4 (CLASSICAL SINGLE-JOB WINE, [36, PROPOSITION 15.9]). *Let G be an arbitrary job size distribution. Then, for all ages $a \geq 0$,*

$$\int_0^\infty \frac{s_{G,G}(a, r)}{r^2} dr = 1.$$

Definition 3.5. Let G be a job size distribution. Let the G -pseudonumber of jobs in the system be

$$\hat{N}_G = \int_0^\infty \frac{\mathbb{E}_\pi[W(\gamma(G) < r) \mid \vec{A}]}{r^2} dr,$$

where $\vec{A} = [A_1, \dots, A_N]$ is the list of all jobs' ages. Similarly to Definition 3.3, \hat{N}_G does not depend on the policy π (beyond requiring nonclairvoyance).

PROPOSITION 3.6 (CLASSICAL WINE [36, THEOREM 15.10]). *The F -pseudonumber of jobs equals the true number of jobs: $\hat{N}_F = N$.*

PROOF SKETCH. Letting $G = F$, expand the integral that defines \hat{N}_F using (3.1) to obtain a sum. By Proposition 3.4, each of the N jobs contributes 1 to the sum. \square

3.2 Multiplicatively close WINE

In Theorems 3.8 and 3.9 below, we give an extension of WINE that uses G -Gittins r -work for distributions G that need not be the true distribution F . The main result is that if G and F are multiplicatively close, then \hat{N}_G/N is close to 1, i.e. the G -pseudonumber of jobs is a good approximation for the true number of jobs. Classical WINE is then the special case where $G = F$, in which case $\hat{N}_F/N = 1$. The proofs are straightforward once we have the following lemma in hand.

LEMMA 3.7. *If G and H are ε -multiplicatively close job size distributions, then for all $b > a \geq 0$,*

$$\frac{\mathbb{E}_{S \sim G}[(S \wedge b) - a \mid S > a]}{\mathbb{E}_{S \sim H}[(S \wedge b) - a \mid S > a]} \in [e^{-2\varepsilon}, e^{2\varepsilon}]$$

PROOF. By the tail integral formula,

$$\mathbb{E}_{S \sim G}[(S \wedge b) - a \mid S > a] = \int_a^b \frac{\bar{G}(t)}{\bar{G}(a)} dt \leq e^{2\varepsilon} \int_a^b \frac{\bar{H}(t)}{\bar{H}(a)} dt = e^{2\varepsilon} \mathbb{E}_{S \sim H}[(S \wedge b) - a \mid S > a].$$

The lower bound follows analogously. \square

THEOREM 3.8 (MULTIPLICATIVELY CLOSE SINGLE-JOB WINE). *If G and H are ε -multiplicatively close job size distributions, then for all $a \geq 0$,*

$$\int_0^\infty \frac{s_{H,G}(a, r)}{r^2} dr \in [e^{-2\varepsilon}, e^{2\varepsilon}].$$

PROOF. By Lemma 3.7,

$$\begin{aligned} \int_0^\infty \frac{s_{H,G}(a, r)}{r^2} dr &= \int_0^\infty \frac{\mathbb{E}_{S \sim G}[(S \wedge b(a, \gamma(H) < r)) - a \mid S > a]}{r^2} dr \\ &\leq e^{2\varepsilon} \int_0^\infty \frac{\mathbb{E}_{S \sim H}[(S \wedge b(a, \gamma(H) < r)) - a \mid S > a]}{r^2} dr = e^{2\varepsilon} \int_0^\infty \frac{s_{H,H}(a, r)}{r^2} dr = e^{2\varepsilon}, \end{aligned}$$

where the last equality follows by Proposition 3.4. The lower bound follows analogously. \square

THEOREM 3.9 (MULTIPLICATIVELY CLOSE WINE). *If G is a job size distribution ε -multiplicatively close to F , then for any policy π ,*

$$\frac{\hat{N}_G}{N} \in [e^{-2\varepsilon}, e^{2\varepsilon}] \quad \text{and} \quad \frac{\mathbb{E}_\pi[\hat{N}_G]}{\mathbb{E}_\pi[N]} \in [e^{-2\varepsilon}, e^{2\varepsilon}].$$

PROOF. Observe that,

$$\hat{N}_G = \int_0^\infty \frac{\mathbb{E}_\pi[W(\gamma(G) < r) \mid \vec{A}]}{r^2} dr = \sum_{i=1}^N \int_0^\infty \frac{s_{F,G}(A_i, r)}{r^2} dr \in [e^{-2\varepsilon}N, e^{2\varepsilon}N],$$

where the last inclusion uses Theorem 3.8. Since the ratio between \hat{N}_G and N holds almost surely (for any π), the ratio between $\mathbb{E}_\pi[\hat{N}_G]$ and $\mathbb{E}_\pi[N]$ follows immediately. \square

3.3 Proving robustness of Gittins to multiplicative tail perturbations

The original application of classical WINE was to prove that $\gamma(F)$ minimizes the mean response time across all possible policies in the M/G/1 queue. In addition to WINE, the other key ingredient needed to prove Gittins's optimality was the following lemma (simplified and restated for our setting) and its corollary. We will need these to prove Theorem 3.1.

LEMMA 3.10 ([36, COROLLARY 8.9]). *Let G be a job size distribution. Then for all policies π and ranks r , system G -Gittins r -work is first-order stochastically minimized by the $\gamma(G)$ policy. That is, for all $x \geq 0$,*

$$\mathbb{P}_{\gamma(G)}[W(\gamma(G) < r) \leq x] \leq \mathbb{P}_\pi[W(\gamma(G) < r) \leq x].$$

COROLLARY 3.11. *Let G be a job size distribution. Then for all policies π , $\mathbb{E}_{\gamma(G)}[\hat{N}_G] \leq \mathbb{E}_\pi[\hat{N}_G]$.*

PROOF. This follows almost immediately from Lemma 3.10 and Tonelli's theorem:

$$\mathbb{E}_{\gamma(G)}[\hat{N}_G] = \int_0^\infty \frac{\mathbb{E}_{\gamma(G)}[W(\gamma(G) < r)]}{r^2} dr \leq \int_0^\infty \frac{\mathbb{E}_\pi[W(\gamma(G) < r)]}{r^2} dr = \mathbb{E}_\pi[\hat{N}_G]. \quad \square$$

PROOF OF THEOREM 3.1. The lower bound holds by optimality of true Gittins. For the upper bound, multiplicatively close WINE (Theorem 3.9) and Corollary 3.11 imply

$$e^{2\varepsilon} \mathbb{E}_{\gamma(F)}[N] \geq \mathbb{E}_{\gamma(F)}[\hat{N}_G] \geq \mathbb{E}_{\gamma(G)}[\hat{N}_G] \geq e^{-2\varepsilon} \mathbb{E}_{\gamma(G)}[N],$$

and so by Little's law, $\frac{\mathbb{E}_{\gamma(G)}[T]}{\mathbb{E}_{\gamma(F)}[T]} = \frac{\mathbb{E}_{\gamma(G)}[N]}{\mathbb{E}_{\gamma(F)}[N]} \leq e^{4\varepsilon}$. \square

4 Truncated empirical Gittins

In this section, we prove that a Gittins policy constructed from a truncated empirical distribution achieves, with high probability, a mean response time close to that of the optimal policy $\gamma(F)$. Our main result is the following.

THEOREM 4.1. *Let G be an empirical distribution constructed from n samples from F . Let $\varepsilon \in (0, 0.6)$, $\ell > 0$, and $\delta \in (0, 1)$. If $n \geq \frac{3 \log(2/\delta)}{\bar{F}(\ell)\varepsilon^2}$, then with probability at least $1 - \delta$,*

$$\mathbb{E}_{\gamma(G_\ell)}[T] \leq e^{4\varepsilon} \mathbb{E}_{\gamma(F)}[T] + \bar{F}(\ell) \left(\frac{\lambda \mathbb{E}[(S \wedge \ell)^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S \mid S \geq \ell]}{(1-\rho)} \right),$$

where $\gamma(F)$ is the Gittins policy constructed from F and $\gamma(G_\ell)$ is the Gittins policy constructed from the truncated empirical distribution G_ℓ .

Note that this bound depends on the underlying true job size distribution F as well as our choice of truncation ℓ . In Section 6, we discuss how to choose ℓ and characterize the resulting error.

Our strategy for proving near-optimality proceeds as follows:

- (1) Using a finite-sample concentration bound (Lemma 2.3), we show that, with high probability, the tails of the true and empirical distributions (\bar{F} and \bar{G}) are ε -multiplicatively close up to some $\ell > 0$ (Definition 2.2).
- (2) We decompose the expected number of jobs in the system under truncated empirical Gittins into the sum of the number of jobs with ages above and below ℓ .
- (3) Due to the optimality of ℓ -truncated true Gittins on the age $< \ell$ subsystem and the ε -closeness of \bar{F}_ℓ and \bar{G}_ℓ , we are able to use Theorem 3.1 to bound the number of jobs with age $< \ell$ in terms of true Gittins.
- (4) We separately bound the number of age $\geq \ell$ jobs using a tagged job argument (Lemma 4.3).
- (5) Finally, we combine the results for age $< \ell$ and age $\geq \ell$ jobs to get a bound on response time ratio via Little's law.

Before proceeding with the proof of Theorem 4.1, we state some necessary auxiliary lemmas.

LEMMA 4.2. *Suppose G is ε -multiplicatively close to F up to some $\ell > 0$. Then G_ℓ is ε -multiplicatively close to F_ℓ .*

PROOF. For all $0 \leq a < \ell$ we have $\bar{F}_\ell(a) = \bar{F}(a)$ and $\bar{G}_\ell(a) = \bar{G}(a)$, and for all $a > \ell$ we have $\bar{F}_\ell(a) = \bar{G}_\ell(a) = 0$. Thus, the lemma follows from Definition 2.2. \square

Lemma 4.2 shows that multiplicative tail closeness between F and G below ℓ transfers directly to their truncated counterparts F_ℓ and G_ℓ . As a result, Theorem 3.1 allows us to compare the behavior of the Gittins policies $\gamma(G_\ell)$ and $\gamma(F_\ell)$ on jobs that have not yet reached age ℓ . However, this comparison alone is insufficient to bound performance relative to the true Gittins policy $\gamma(F)$, since truncation alters the treatment of jobs whose ages exceed ℓ . The main remaining difficulty is therefore to control the contribution of jobs that reach age at least ℓ , which cannot be handled via distributional closeness and must instead be bounded directly. The following lemma provides such a bound by relating the number of large jobs under truncated Gittins to standard workload quantities using a tagged job argument.

LEMMA 4.3. *Let G be a job size distribution. Then the ℓ -truncated Gittins policy $\gamma(G_\ell)$, satisfies*

$$\mathbb{E}_{\gamma(G_\ell)}[N(\text{size} \geq \ell)] \leq \lambda \bar{F}(\ell) \left(\frac{\lambda \mathbb{E}[(S \wedge \ell)^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S | S \geq \ell]}{(1-\rho)} \right).$$

PROOF SKETCH. This follows from Little's law, standard results regarding the mean work in a stationary M/G/1 queue [18], and reasoning about what work in the system must be completed prior to the completion of a tagged job with size $\geq \ell$. See Appendix D.2 for a complete proof. \square

We are now ready to prove Theorem 4.1 using the above lemmas, Lemma 2.3, and Theorem 3.1.

PROOF OF THEOREM 4.1. The assumption on n is equivalent to $\delta \geq 2 \exp(-n\varepsilon^2 \bar{F}(\ell)/3)$, so Lemmas 2.3 and 4.2 imply that G_ℓ is ε -multiplicatively close to F_ℓ with probability at least $1 - \delta$, so assume this multiplicative closeness hereafter.

We then decompose the expected number of jobs in the system under policy $\gamma(G_\ell)$ into the sum of the expected number of jobs in the system currently at age $< \ell$ and the number of jobs in the system currently at age $\geq \ell$, both still under policy $\gamma(G_\ell)$. That is,

$$\mathbb{E}_{\gamma(G_\ell)}[N] = \mathbb{E}_{\gamma(G_\ell)}[N(\text{age} < \ell)] + \mathbb{E}_{\gamma(G_\ell)}[N(\text{age} \geq \ell)].$$

Observation 4.4. When considering solely the number of jobs of age $< \ell$, note that $N(\text{age} < \ell)$ goes down whenever a job either completes or when a job reaches age ℓ . Therefore, when considering the “subsystem” of age ℓ jobs, the “job size distribution” behaves exactly like the ℓ -truncated true job size distribution, F_ℓ . Therefore, by Theorem 2.6, the Gittins policy constructed from F_ℓ is optimal for the subsystem of age $< \ell$.

Combining this result with our concentration bound, we can apply Theorem 3.1 to the age $< \ell$ subsystem and find that

$$\begin{aligned}\mathbb{E}_{\gamma(G_\ell)}[N] &= \mathbb{E}_{\gamma(G_\ell)}[N(\text{age} < \ell)] + \mathbb{E}_{\gamma(G_\ell)}[N(\text{age} \geq \ell)] \\ &\leq e^{4\ell} \mathbb{E}_{\gamma(F_\ell)}[N(\text{age} < \ell)] + \mathbb{E}_{\gamma(G_\ell)}[N(\text{age} \geq \ell)].\end{aligned}$$

Now note that, once again by Observation 4.4, $\gamma(F_\ell)$ will minimize the number of jobs in the system of age $< \ell$ relative to any other policy, including $\gamma(F)$, and so

$$\mathbb{E}_{\gamma(G_\ell)}[N] \leq e^{4\ell} \mathbb{E}_{\gamma(F)}[N(\text{age} < \ell)] + \mathbb{E}_{\gamma(G_\ell)}[N(\text{age} \geq \ell)].$$

We then subtract $\mathbb{E}_{\gamma(F)}[N]$ from both sides. Because jobs of age $\geq \ell$ must necessarily be of size $\geq \ell$,

$$\begin{aligned}\mathbb{E}_{\gamma(G_\ell)}[N] - \mathbb{E}_{\gamma(F)}[N] &\leq (e^{4\ell} - 1) \mathbb{E}_{\gamma(F)}[N(\text{age} < \ell)] + \mathbb{E}_{\gamma(G_\ell)}[N(\text{age} \geq \ell)] - \mathbb{E}_{\gamma(F)}[N(\text{age} \geq \ell)] \\ &\leq (e^{4\ell} - 1) \mathbb{E}_{\gamma(F)}[N(\text{age} < \ell)] + \mathbb{E}_{\gamma(G_\ell)}[N(\text{size} \geq \ell)] \\ &\leq (e^{4\ell} - 1) \mathbb{E}_{\gamma(F)}[N] + \lambda \bar{F}(\ell) \left(\frac{\lambda \mathbb{E}[(S \wedge \ell)^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S \mid S \geq \ell]}{(1-\rho)} \right),\end{aligned}$$

where the last inequality is due to Lemma 4.3 and the fact that $N(\text{age} < \ell) \leq N$. The result then follows by Little’s law. \square

5 Bounding empirical Gittins’s response time

In this section we extend the response time bound in Section 3 under the weaker condition that G is ε -multiplicatively close to F only up to some $\ell > 0$, as long as G is an empirical distribution generated from samples. We then use this to prove a (with high probability) response time bound as in Section 4, but now for empirical Gittins instead of truncated empirical Gittins. This result is summarized in the following theorem.

THEOREM 5.1. *Let G be an empirical distribution constructed from n samples from F . Let $\varepsilon \in (0, 0.6)$ and $\delta \in (0, 1)$. If $n \geq \frac{3 \log(2/\delta)}{\bar{F}(\ell) \varepsilon^2}$, and $e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)} > 0$, then with probability at least $1 - \delta$,*

$$\mathbb{E}_{\gamma(G)}[T] \leq \inf_{\ell > k > 0} \frac{e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n}{e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)}} \left(\mathbb{E}_{\gamma(F)}[T] + \frac{\bar{F}(k)(1+h_n)}{(1-\rho)} \left(\frac{\lambda \mathbb{E}[S^2]}{2(1-\rho)} + \mathbb{E}[S \mid S > k] \right) \right),$$

where $h_n = \sum_{i=1}^n \frac{1}{i}$ is the n -th harmonic number.

Note that unlike the result in Theorem 4.1, which has only one parameter ℓ , we now have a second parameter $k \in (0, \ell)$ to choose. We introduce this new parameter to handle the fact that the bound has a $h_n \approx \log n$ term out front. If we did not introduce k (which in the proof equates to letting $k = \ell$), the term out front would instead be $e^{2\varepsilon} + h_n \rightarrow \infty$ as $n \rightarrow \infty$. This is clearly undesirable. By choosing the k carefully we can ensure that $\frac{\bar{F}(\ell)}{\bar{F}(k)} h_n \rightarrow 0$ as $n \rightarrow \infty$ and thus that $\mathbb{E}_{\gamma(G)}[T]$ approaches $\mathbb{E}_{\gamma(F)}[T]$ (Theorem 6.3).

There is an important distinction between the role the parameter ℓ plays in Theorem 4.1 vs. the roles k and ℓ play in Theorem 5.1. In Theorem 4.1, ℓ is an *input to the policy*, namely the truncation threshold, meaning ℓ is something we have to choose without knowledge of the true distribution. In contrast, in Theorem 5.1, k and ℓ are *inputs only to the analysis*, not the policy itself, meaning k and ℓ can be optimized with knowledge of the true distribution; this is why we can take an infimum over them.

To prove Theorem 5.1 we will separately bound the contribution of jobs with ages $< k$ and those with ages $\geq k$ to the multiplicatively close WINE integral. The former will be analogous to what we did in the previous sections due to G being ε -multiplicatively close to F up to $l > k$ with high probability. The latter requires a new approach and is the focus of the following section.

5.1 A single-job WINE bound for empirical distributions

In this section we wish to bound the contribution of jobs with ages $\geq k$ to the multiplicatively close WINE integral, which, based on the proof of Theorem 3.9, we know we can do by finding an upper bound on

$$\varphi_{H,G}(a) = \int_0^\infty \frac{s_{H,G}(a, r)}{r^2} dr$$

for all ages $a \geq k$ and job size distributions H , given only the assumption that G is a non-negative empirical distribution constructed from n samples. In particular, we will show that under this assumption,

$$\varphi_{H,G}(a) dr \leq h_n = \sum_{i=1}^n \frac{1}{i},$$

for all ages $a \geq 0$ and size distributions H . We call this *single-job WINE for empirical distributions* (Theorem 5.3), though it relies only on G being a mixture of n point masses. Our approach is:

- (1) choose the distribution H that maximizes $\varphi_{H,G}(a)$, removing the dependence on H ; then
- (2) show that h_n is an upper bound by induction on the number of samples n used to construct G .

The step involving H is straightforward. We introduce a new distribution, referred to as ∞ , under which $S = \infty$ almost surely. Intuitively, choosing $H = \infty$ maximizes the expected remaining G -Gittins r -work under H , and one can easily verify by expanding the definition $s_{H,G}(a, r)$ that

$$\varphi_{H,G}(a) dr \leq \varphi_{\infty,G}(a) = \int_0^\infty \frac{b_{Y(G)}(a, r) - a}{r^2} dr.$$

Remark 5.2. Throughout the remainder of this section, we assume without loss of generality that G is a discrete uniform distribution with atoms $0 = G_0 < G_1 < G_2 < \dots < G_n$, and that $a = 0$ (so we only need to show that $\varphi_{\infty,G}(0) \leq h_n$). The former assumption is justified by Theorem C.1 in the appendix, which shows that we can “split up” atoms with probability mass $\frac{m}{n}$ and $m > 1$ into m atoms of mass $\frac{1}{n}$ while still maintaining an upper bound. The latter assumption is justified by the following argument. Let $G|_{>t}$ denote the distribution of X conditioned on $X > t$, where $X \sim G$, and then observe that

$$r_{Y(G|_{>t})}(a) = r_{Y(G)}(a + t), \quad (5.1)$$

by the definition of $r_{Y(G)}(a)$. This can be seen by comparing Figs. 5.1 and 5.2. It follows that

$$\varphi_{\infty,G}(a) = \varphi_{\infty,G|_{>a}}(0),$$

so proving the upper bound for an arbitrary discrete uniform distribution G at $a = 0$ is sufficient.

We now prove that $\varphi_{\infty,G}(0) \leq h_n$ by induction on n . In this section, we outline the key ideas in some detail, but we defer the formal proof to Appendix A.




Fig. 5.1. The Gittins rank function $r_{\gamma(G)}(a)$ for discrete uniform distribution $G = \text{Unif}\{3, 6, 16, 20\}$. The dashed blue line is $w_{\gamma(G)}(0, a) = \sup_{0 \leq u \leq a} r_{\gamma(G)}(u)$, and the region above it, referred to as V in the text, is shaded blue. The blue shaded region is separated into two sub-regions, V_1 and V_2 , which correspond to the ages before and after the first atom respectively.

The first key idea is to view $\varphi_{\infty, G}(0)$ as an integral over a region:

$$\varphi_{\infty, G}(0) = \int_0^{\infty} \frac{b_{\gamma(G)}(0, r)}{r^2} dr = \int_0^{\infty} \int_0^{b_{\gamma(G)}(0, r)} \frac{1}{r^2} da dr = \iint_V \frac{1}{r^2} da dr,$$

where $V = \{(a, r) \in [0, \infty)^2 : a < b_{\gamma(G)}(0, r)\}$. See Fig. 5.1 for an illustration of V , where it the blue shaded region $V_1 \cup V_2$. We would like to separate V into two regions:

- (1) a region which depends only on the first atom G_1 , and
- (2) a region which depends only on the atoms G_2, \dots, G_n .

The integral over the first region can then be explicitly computed, and the integral over the second region can be bounded using the inductive hypothesis, as it is a shifted copy of a region resulting from a distribution with $n - 1$ atoms.

One might hope that the regions V_1 and V_2 from Fig. 5.1 would satisfy the above requirements, but this is not the case. The problem is that $r_{\gamma(G)}(0) > r_{\gamma(G)}(G_1)$, which causes the lower vertical limit of V_2 at some ages to be $r_{\gamma(G)}(0)$, which depends on G_1 , so requirement (2) fails.⁶ To remedy this, we can imagine shifting all the atoms of G to the left until $r_{\gamma(G)}(0) < r_{\gamma(G)}(G_1)$, and then splitting V into two regions with the desired properties. Equation (5.1) tells us that this shifting operation can be accomplished by conditioning G on being greater than t for some sufficiently large $t < G_1$. We see an example of this in Fig. 5.2 which has distribution $G' = G|_{>1}$, i.e. G' is G shifted left by 1. Additionally, we can see that the regions V'_1 and V'_2 in Fig. 5.2 satisfy requirements (1) and (2)—in particular, it is clear that

$$\iint_{V'_2} \frac{1}{r^2} dr da = \varphi_{\infty, G|_{>G_1}}(0).$$

Since $G|_{>G_1}$ is a discrete uniform distribution with $n - 1$ atoms, the inductive hypothesis guarantees that the above integral is bounded above by h_{n-1} . Lastly, since V'_1 depends only on G_1 , it is straightforward to manually compute that

$$\iint_{V'_1} \frac{1}{r^2} dr da = \frac{1}{n},$$

which means that $\varphi_{\infty, G|_{>1}}(0) \leq \frac{1}{n} + h_{n-1} = h_n$, completing the proof of the inductive step. A similar computation works for the base case of $n = 1$. The last step to complete the proof is then to show

⁶It turns out requirement (1) also fails in this case, as perturbing G_2 would perturb $r_{\gamma(G)}(0)$ and thereby perturb V_1 .




Fig. 5.2. The Gittins rank function $r_{\gamma(G)}(a)$ for discrete uniform distribution $G' = \text{Unif}\{2, 5, 15, 19\}$, which is equivalent to $G|_{>1}$ where G is the distribution from Fig. 5.1. The dashed blue line is $w_{\gamma(G)}(0, a)$. The region above it is shaded blue and split into two sub-regions V'_1 and V'_2 , which correspond to the ages before and after the first atom respectively. Because V'_2 does not depend on the first atom, we can use the inductive hypothesis to conclude that the integral of $1/r^2$ over V'_2 is at most h_3 . We can manually compute that the integral of $1/r^2$ over V'_1 is $1/4$.

that $\varphi_{\infty, G}(0) \leq \varphi_{\infty, G|_{>1}}(0)$, that is, that the “sliding” operation can only increase the integral over the region V .

In Appendix A, we formalize the above ideas to prove the following.

THEOREM 5.3 (SINGLE-JOB WINE FOR EMPIRICAL DISTRIBUTIONS). *Let G be an empirical job size distribution constructed from n samples and let H be an arbitrary job size distribution. Then for all ages $a \geq 0$,*

$$\int_0^{\infty} \frac{s_{H,G}(a, r)}{r^2} dr \leq h_n,$$

where $h_n = \sum_{i=1}^n \frac{1}{i}$ is the n -th harmonic number.

5.2 Bounding the response time of empirical Gittins

The proof of Theorem 5.1 has two main steps, mimicking the structure of Theorem 3.1:

- (1) We first bound the contribution a job of age $< k$ can make to the multiplicatively close WINE integral (analogous to Theorem 3.8) and then combine that with Theorem 5.3 to bound \hat{N}_G in terms of N (analogous to multiplicatively close WINE, Theorem 3.9). This is all done in Theorem 5.5, which is WINE for empirical distributions.
- (2) We then use the fact that $\gamma(G)$ minimizes \hat{N}_G (Corollary 3.11) along with the bound on \hat{N}_G to bound $\mathbb{E}_{\gamma(G)}[N]$ (analogous to the final step in the proof of Theorem 3.1). This is done in Lemmas D.1 and D.2.

Due to the similarity to the proof of Theorem 3.1, we give only step (1) below, deferring step (2) and the proof of Theorem 5.1 itself to Appendix D.3.

Before proceeding, we prove a quick helper lemma.

LEMMA 5.4. *Let G and H be job size distributions and let $\ell > 0$. Then for all ranks r and ages $a < \ell$,*

$$\int_a^{b_{\gamma(G)}(a, r) \wedge \ell} \frac{\bar{H}(t)}{\bar{H}(a)} dt \leq s_{H,G}(a, r) \leq \int_a^{b_{\gamma(G)}(a, r) \wedge \ell} \frac{\bar{H}(t)}{\bar{H}(a)} dt + \frac{\bar{H}(\ell)}{\bar{H}(a)} s_{H,G}(\ell, r).$$

PROOF SKETCH. This follows from a straightforward computation after observing that for any age $a < \ell$, we can split the service until $b_{Y(G)}(a, r)$ into service until $b_{Y(G)}(a, r) \wedge \ell$ and then any remaining service. A complete proof is presented in Appendix D.3. \square

THEOREM 5.5 (WINE FOR EMPIRICAL DISTRIBUTIONS). *Let G be a mass n empirical job size distribution that is ε -multiplicatively close to F up to $\ell > 0$. Then for all $0 < k < \ell$,*

$$\left(e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)} \right) N(\text{age} < k) \leq \hat{N}_G \leq \left(e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n \right) N(\text{age} < k) + h_n N(\text{age} \geq k).$$

PROOF. Fix a $0 < k < \ell$ and then let $0 \leq a < k$. Using Lemma 5.4 with distributions F and G and then using their ε -multiplicative closeness up to ℓ ,

$$e^{-2\varepsilon} \int_a^{b_{Y(G)}(a,r) \wedge \ell} \frac{\bar{G}(t)}{\bar{G}(a)} dt \leq s_{F,G}(a, r) \leq e^{2\varepsilon} \int_a^{b_{Y(G)}(a,r) \wedge \ell} \frac{\bar{G}(t)}{\bar{G}(a)} dt + \frac{\bar{F}(\ell)}{\bar{F}(a)} s_{F,G}(\ell, r).$$

Using Lemma 5.4 again, but this time letting both distributions be G , we can get that

$$s_{G,G}(a, r) - \frac{\bar{G}(\ell)}{\bar{G}(a)} s_{G,G}(\ell, r) \leq \int_a^{b_{Y(G)}(a,r) \wedge \ell} \frac{\bar{G}(t)}{\bar{G}(a)} dt \leq s_{G,G}(a, r).$$

Combining the above inequalities and once again using the ε -multiplicative closeness of F and G , yields

$$e^{-2\varepsilon} s_{G,G}(a, r) - \frac{\bar{F}(\ell)}{\bar{F}(a)} s_{G,G}(\ell, r) \leq s_{F,G}(a, r) \leq e^{2\varepsilon} s_{G,G} + \frac{\bar{F}(\ell)}{\bar{F}(a)} s_{F,G}(\ell, r).$$

Thus, by Proposition 3.4 (classical single-job WINE) and Theorem 5.3 (single-job WINE for empirical distributions),

$$e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(a)} \leq \int_0^\infty \frac{s_{F,G}(a, r)}{r^2} dr \leq e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(a)} h_n.$$

We would now like to bound \hat{N}_G . We will bound the contribution of all jobs of ages $< k$ to the integral using the above bound, and the contribution of all jobs of ages $\geq k$ using Theorem 5.3:

$$\begin{aligned} \hat{N}_G &= \int_0^\infty \frac{\mathbb{E}_\pi[W(\gamma(G) < r) \mid \vec{A}]}{r^2} dr \\ &= \sum_{i=1}^N \mathbb{1}(A_i < k) \int_0^\infty \frac{s_{F,G}(A_i, r)}{r^2} dr + \sum_{i=1}^N \mathbb{1}(A_i \geq k) \int_0^\infty \frac{s_{F,G}(A_i, r)}{r^2} dr, \end{aligned}$$

and then by the above we have,

$$\left(e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)} \right) N(\text{age} < k) \leq \sum_{i=1}^N \mathbb{1}(A_i < k) \int_0^\infty \frac{s_{F,G}(A_i, r)}{r^2} dr \leq \left(e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n \right) N(\text{age} < k)$$

and by Theorem 5.3 (single-job WINE for empirical distributions),

$$0 \leq \sum_{i=1}^N \mathbb{1}(A_i \geq k) \int_0^\infty \frac{s_{F,G}(A_i, r)}{r^2} dr \leq h_n N(\text{age} \geq k).$$

Putting these together gets the desired bounds. \square

6 Asymptotic analysis of mean response time error

In this section, we characterize the asymptotic dependence of the mean response time ratio between empirical and true Gittins policies on the true job size distribution, load ρ , and number of samples n .

For a given true distribution and load, we wish to understand two quantities: (1) how many samples are required to achieve a given *multiplicative error*, meaning mean response time ratio minus one, with high probability; and (2) when using truncated empirical Gittins, where to truncate to minimize the asymptotic error.

Our results are phrased as rates in the $n \rightarrow \infty$, $\rho \rightarrow 1$, and $\delta \rightarrow 0$ limits, where ρ is the load, n is the number of samples, and δ is the desired upper bound on the failure probability. We use standard asymptotic notation $O(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$, along with $\tilde{O}(\cdot)$ to suppress factors that are logarithmic in n , $1/(1-\rho)$, and $1/\delta$. Constants hidden by the asymptotic notation may depend on the distribution. We sometimes use asymptotic notation with other variables whose limiting values will be clear from context.

Our results depend on a parameter α that captures how heavy the true distribution's tail is. Specifically, we assume throughout that the *upper Matuszewska index* of \bar{F} is less than α [4, Chapter 2]. More concretely, this is equivalent to assuming that there exists $c > 0$ such that for all $y \geq x \geq 0$,

$$\frac{\bar{F}(y)}{\bar{F}(x)} \leq c \left(\frac{y}{x}\right)^{-\alpha}. \quad (6.1)$$

One could easily extend our analysis to handle a weaker assumption, such as $\bar{F}(x) \leq O(x^{-\alpha})$, at the cost of worse convergence rates. But (6.1) turns out to be the natural assumption to make, which is unsurprising given that existing heavy-traffic analyses of M/G/1 scheduling also rely on assumptions on Matuszewska indices [21, 26, 38]. We also emphasize that (6.1) is far from a restrictive assumption: most common examples of parametric distributions (Pareto, log-normal, Weibull, exponential, normal, etc.) that satisfy $\bar{F}(x) \leq O(x^{-\alpha})$ also satisfy (6.1).

We defer all of the proofs to Appendix B, as they follow from Theorems 4.1 and 5.1 and straightforward computation.

THEOREM 6.1 (TRUNCATION, FINITE VARIANCE). *Let $\alpha > 2$ be such that F satisfies (6.1), and let G be the empirical distribution constructed from n samples from F . Consider the truncated empirical Gittins policy $\gamma(G_\ell)$ for some truncation level ℓ , which may depend on G . Choosing ℓ such that*

$$\bar{G}(\ell) = \Theta\left(\min\left(n^{-1/3}(1-\rho)^{\frac{2\alpha}{3(\alpha-1)}}, n^{\frac{-\alpha}{3\alpha-2}}(1-\rho)^{\frac{2\alpha}{(3\alpha-2)(\alpha-1)}}\right)\right)$$

yields, with probability at least $1 - \delta$, multiplicative error bounded by

$$\frac{\mathbb{E}_{\gamma(G_\ell)}[T]}{\mathbb{E}_{\gamma(F)}[T]} - 1 \leq \tilde{O}\left(\max\left(n^{-1/3}(1-\rho)^{\frac{-\alpha}{3(\alpha-1)}}, n^{\frac{-(\alpha-1)}{3\alpha-2}}(1-\rho)^{\frac{-\alpha}{(3\alpha-2)(\alpha-1)}}\right)\right).$$

THEOREM 6.2 (TRUNCATION, POTENTIALLY INFINITE VARIANCE). *Let $\alpha \in (1, 2)$ be such that F satisfies (6.1), and let G be the empirical distribution constructed from n samples from F . Consider the truncated empirical Gittins policy $\gamma(G_\ell)$ for some truncation level ℓ , which may depend on G . Choosing ℓ such that*

$$\bar{G}(\ell) = \Theta\left(\min\left(n^{\frac{-\alpha}{5\alpha-4}}(1-\rho)^{\frac{4\alpha}{5\alpha-4}}, n^{\frac{-\alpha}{3\alpha-2}}(1-\rho)^{\frac{2\alpha}{3\alpha-2}}\right)\right)$$

yields, with probability at least $1 - \delta$, multiplicative error bounded by

$$\frac{\mathbb{E}_{\gamma(G_\ell)}[T]}{\mathbb{E}_{\gamma(F)}[T]} - 1 \leq \tilde{O}\left(\max\left(n^{\frac{-2(\alpha-1)}{5\alpha-4}}(1-\rho)^{\frac{-2\alpha}{5\alpha-4}}, n^{\frac{-(\alpha-1)}{3\alpha-2}}(1-\rho)^{\frac{-\alpha}{3\alpha-2}}\right)\right).$$

THEOREM 6.3 (NO TRUNCATION, FINITE VARIANCE). *Let $\alpha > 2$ be such that F satisfies (6.1), and let G be the empirical distribution constructed from n samples from F . The empirical Gittins policy $\gamma(G)$ achieves, with probability at least $1 - \delta$, multiplicative error bounded by*

$$\frac{\mathbb{E}_{\gamma(G)}[T]}{\mathbb{E}_{\gamma(F)}[T]} - 1 \leq \tilde{O}\left(\max\left(n^{-1/4}(1-\rho)^{\frac{-3\alpha}{4(\alpha-1)}}, n^{\frac{-(\alpha-1)}{4\alpha-3}}(1-\rho)^{\frac{-3\alpha}{(4\alpha-3)(\alpha-1)}}\right)\right).$$

We have not proven a result with no truncation and potentially infinite variance, i.e. for $\alpha \in (1, 2)$. We believe the multiplicative error is still small in expectation, but the upper tail may decay slowly enough that the dependence on $1/\delta$ is polynomial rather than polylogarithmic. Understanding this would require digging into the details of what is known about the asymptotic behavior of true Gittins's rank function at large ages [38, 46] and understanding the extent to which those properties carry over to empirical Gittins.

7 Simulations

The goal of this section is to answer the following two questions:

- (1) How well do the empirical policies proposed in this paper do compared to the optimal policy *in practice*?
- (2) Do simulations show an obvious winner between the truncated empirical Gittins policy $\gamma(G_\ell)$ and (untruncated) empirical Gittins policy $\gamma(G)$?

The short answers are: (1) they do quite well in practice, even with relatively few samples (on the order of 1000), and (2) (untruncated) empirical Gittins seems to do better with fewer samples and when the variance is lower, but truncated empirical Gittins also does well when there are more samples, the variance is higher, and there is high load. The primary takeaway from this section should be that, when the variance is finite, *simply using empirical Gittins will usually be good enough*. If the underlying distribution is suspected to have infinite variance, truncation may help, but we leave a rigorous simulation study of this to future work due to engineering challenges (Appendix E).

7.1 Experiment setup

For each combination of experiment parameters (distribution F , load ρ , number of samples n) we ran 100 *trials* each of which was comprised of the following steps:

- (1) Sample n points from the distribution F and construct an empirical distribution G .
- (2) Compute both $\gamma(G)$ and $\gamma(G_\ell)$.
- (3) For each of the policies, simulate 10,000 busy periods in an M/G/1 queue with job size distribution F and load ρ and record the mean response times.

We ran trials with number of samples equaling 10, 100, or 1000 and load chosen to be 0.8 or 0.98. We considered two different discretized distributions,⁷ which we will refer to as the 1-6-14 distribution and the bounded Pareto distribution.

- The bounded Pareto distribution we considered was a Pareto distribution with scale parameter $x_m = 2$ and shape parameter $\alpha = 1.2$, bounded above at 500 and discretized with step size 0.01.
- The 1-6-14 distribution is a Gaussian mixture model with component means at 1, 6, and 14, each having standard deviation 0.5 and equal mixture weights. It was bounded below at 0 and above at 16 and discretized with step size 0.01.

For the truncated empirical Gittins policy, we truncated at the age ℓ corresponding to $\bar{G}(\ell) = n^{-1/3}(1-\rho)^{2/3}$, which we chose by letting $\alpha \rightarrow \infty$ in the truncation rule in Theorem 6.1.⁸

⁷Both distributions needed to be discretized since the simulation could only compute Gittins ranks for discrete distributions.

⁸Since both of these distributions are bounded, they are light-tailed and have all moments finite, which corresponds to a Pareto distribution with $\alpha \rightarrow \infty$.



(a) 1-6-14 distribution, $\rho = 0.8$.(b) Bounded Pareto distribution, $\rho = 0.8$.

Fig. 7.1. Histograms showing the mean response time of the (untruncated) empirical Gittins policy $\gamma(G)$ in trials with different numbers of samples n .

7.2 Results

Figure 7.1 demonstrates how the distribution of mean response times of the empirical Gittins policy across the 100 trials concentrates as we increase the number of samples. In particular, we see that for both distributions the performance of the empirical Gittins policy concentrates near that of the true Gittins policy $\gamma(F)$ at 1000 samples. Figure 7.2 compares the performance of the empirical Gittins policy $\gamma(G)$ and the truncated empirical Gittins policy $\gamma(G_t)$ across distributions, loads, and number of samples. We see that in all but one case, the empirical Gittins policy outperforms the truncated empirical Gittins policy and in many cases almost matches the performance of the true Gittins policy at 1000 samples. The truncated empirical Gittins policy matches the performance of the empirical Gittins policy for the bounded Pareto distribution with load 0.98, suggesting that the truncation may be of use when variance and load are both high and there are a non-trivial number of samples available (≥ 1000). Figure 7.2 also includes FCFS and PLCFS as baseline policies. For the bounded Pareto distribution, both empirical Gittins and truncated empirical Gittins substantially outperform these baselines with as few as 10 samples. For the 1-6-14 distribution, FCFS (whose performance is better than PLCFS) outperforms empirical Gittins at 10 samples, performs comparably at 100 samples, and is surpassed by it at 1000 samples.


8 Future work

We conclude by outlining a few directions for future work motivated by the results and techniques developed in this paper.


8.1 Remaining open questions about empirical Gittins

Our results still leave open several questions about empirical Gittins, particularly regarding lower bounds. For instance, can we prove lower bounds on the convergence rates of empirical Gittins, with and without truncation, that match the upper bounds we obtain in Section 6? In light of our simulation results in Theorem 6.3, we conjecture that the true convergence rate of the untruncated case is at least as good as that of the truncated case, in which case our current result for the untruncated case would not be tight.


Returning to the problem formulation from Section 1.1, we could also ask about a universal lower bound. Does some version of empirical Gittins achieve the optimal convergence rate among SOAP policy constructions? We conjecture that up to sub-polynomial factors, the answer is “yes” in at least some cases, such as the light-tailed case ($\alpha \rightarrow \infty$) with load fixed and a large number of samples




(a) Results from the experiment with the 1-6-14 distribution, $\rho = 0.8$.



(b) Results from the experiment with the 1-6-14 distribution, $\rho = 0.98$.



(c) Results from the experiment with the bounded Pareto distribution, $\rho = 0.8$.



(d) Results from the experiment with the bounded Pareto distribution, $\rho = 0.98$.

Fig. 7.2. Comparison of the empirical Gittins policy $\gamma(G)$ and the truncated empirical Gittins policy $\gamma(G_\ell)$ across all experiments. The vertical axis is the ratio between the mean response time of a single trial under a given policy and the optimal mean response time of $\gamma(F)$. FCFS and PLCFS are included as baselines for comparison; however, FCFS falls outside the plotted range in two of the experiments.

$n \rightarrow \infty$. That is, we believe $\Omega(n^{-1/3})$ scaling is inevitable, as it arises naturally from Lemma 2.3 by asking for the empirical distribution to be ε -multiplicatively close to the true distribution up to its $(1 - \varepsilon)$ -th quantile.

8.2 Leveraging kernel density estimation

Another approach to sample-based scheduling is to estimate the service-time distribution via kernel density estimation (KDE), and then apply the Gittins construction to the resulting smoothed density function rather than to the raw empirical distribution. Kernel density estimators are known to be consistent under standard conditions, including in settings with data-driven bandwidth choices [5, 8, 20]. Smoothing the empirical distribution may allow us to directly bound the density ratio up to some value, which, using Ramakrishna et al. [30], would directly give us a ratio of response times and may potentially achieve better performance. Alternatively, one could estimate and smooth the hazard rate $h(x) = \frac{f(x)}{F(x)}$, which is related to the Gittins rank function [1]. Kernel-based methods for hazard rate estimation have been studied [31], suggesting scheduling based on estimated hazard rates may also be a promising direction to explore.

At the same time, there are a few potential limitations to this approach. Existing guarantees for KDE are primarily additive [5, 8], and converting such bounds into multiplicative guarantees still requires restricting attention to intervals where the tail is bounded away from zero, just as in

classical ratio results for the empirical CDF [28, 53?]. Moreover, smoothing introduces bias through bandwidth choice. Further, any benefits from smoothing depend strongly on bandwidth selection, which may itself require additional assumptions on the underlying distribution beyond access to samples alone. Finally, computing the Gittins rank function for a continuous distribution incurs significantly more computational overhead than for the empirical distribution. It remains an open question whether, possibly under additional assumptions, smoothing can improve finite-sample performance while retaining guarantees comparable to ours.

8.3 Empirical Gittins for tail latency

Recent work [19] proposes so-called γ -Gittins policies, which optimize (asymptotic) tail latency when individual job sizes are unknown. These policies, however, assume knowledge of the underlying job size distribution. A natural extension is therefore to study the performance of “empirical γ -Gittins” policies constructed from finite samples. One potential complication is that the construction of γ -Gittins policies requires knowledge of the arrival rate of jobs, which might also need to be estimated from data. An open question is whether the (asymptotic) tail-latency guarantees established for γ -Gittins degrade gracefully under empirical estimation.

8.4 Empirical Gittins for other Markov processes

Gittins policies are known to be optimal not only for the M/G/1, but also for multi-armed bandits [12, 41, 52] and other Markov decision problems [13, 45, 56]. These optimality results assume full knowledge of the underlying Markov dynamics. It is therefore natural to ask how performance is affected when Gittins indices are constructed from empirical estimates of the transition dynamics. Recent work by Charles-Rebuffé et al. [6] takes a step in this direction for finite-state Markovian multi-armed bandits (and their restless generalization). This suggests that empirical Gittins should perform well in Klimov’s problem [13, 23, 24], a case of M/G/1 scheduling in which jobs are finite-state Markov chains. By combining insights from the finite-state case with our work, it is possible one could analyze empirical Gittins in multi-armed bandit and M/G/1 scheduling problems where arms or jobs are piecewise-deterministic Markov processes.

8.5 Beyond one-shot problems: learning while scheduling

Our analysis focuses on a one-shot setting with a fixed underlying job size distribution. In many practical systems, however, the job size distribution may evolve gradually over time, raising the question of how to learn and adapt scheduling policies while continuing to serve jobs. Understanding how empirical Gittins-type policies perform under such distributional drift is therefore an important extension.

One natural approach is to update the scheduling policy at the start of each busy period, constructing an empirical Gittins policy using samples from the most recent P completed jobs. In this view, empirical Gittins acts as a “gray box” scheduler. The system continually re-estimates a policy from recent data and deploys it for the duration of the next busy period. This approach would allow the policy to adapt to slow changes in the job size distribution while retaining the interpretability, and ideally, the near-optimality of empirical Gittins. An open question is how rapidly the underlying distribution can drift while an updating empirical Gittins policy continues to perform near-optimally.

Acknowledgments

This work was supported by the National Science Foundation (NSF) under grant no. CMMI-2307008. Amit Harlev was supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program (<https://ndseg.sysplus.com/>).

Code underlying plots and simulations was prepared in part using generative AI tools.

References

- [1] Samuli Aalto, Urtzi Ayesta, and Rhonda Righter. 2009. On the Gittins Index in the $M/G/1$ Queue. *Queueing Systems* 63, 1-4 (Dec. 2009), 437–458. <https://doi.org/10.1007/s11134-009-9141-x>
- [2] Samuli Aalto, Urtzi Ayesta, and Rhonda Righter. 2011. Properties of the Gittins Index with Application to Optimal Scheduling. *Probability in the Engineering and Informational Sciences* 25, 3 (July 2011), 269–288. <https://doi.org/10.1017/S0269964811000015>
- [3] Sayan Banerjee, Amarjit Budhiraja, and Amber L. Puha. 2022. Heavy Traffic Scaling Limits for Shortest Remaining Processing Time Queues with Heavy Tailed Processing Time Distributions. *The Annals of Applied Probability* 32, 4 (Aug. 2022), 2587–2651. <https://doi.org/10.1214/21-AAP1741>
- [4] Nicholas H. Bingham, Charles M. Goldie, and Jef L. Teugels. 1987. *Regular Variation*. Number 27 in Encyclopedia of Mathematics and Its Applications. Cambridge University Press, Cambridge, UK.
- [5] José E. Chacón and Alberto Rodríguez-Casal. 2010. A Note on the Universal Consistency of the Kernel Distribution Function Estimator. *Statistics & Probability Letters* 80, 17-18 (Sept. 2010), 1414–1419. <https://doi.org/10.1016/j.spl.2010.05.007>
- [6] Joël Charles-Rebuffé, Nicolas Gast, and Bruno Gaujal. 2025. Model-Based Learning of Whittle Indices. <https://doi.org/10.48550/arXiv.2511.20397> arXiv:2511.20397 [cs]
- [7] Harshit Dhankhar, Kshitij Mishra, and Tejas Bodas. 2025. Tabular and Deep Reinforcement Learning for Gittins Index. <https://doi.org/10.48550/arXiv.2405.01157> arXiv:2405.01157 [cs]
- [8] Uwe Einmahl and David M. Mason. 2005. Uniform in Bandwidth Consistency of Kernel-Type Function Estimators. *The Annals of Statistics* 33, 3 (June 2005), 1380–1403. <https://doi.org/10.1214/009053605000000129>
- [9] Daniel Freund, Thodoris Lykouris, and Wentao Weng. 2023. Quantifying the Cost of Learning in Queuing Systems. In *Advances in Neural Information Processing Systems (NeurIPS 2023)*, Vol. 36. Curran Associates, Inc., New Orleans, LA, 6532–6544. <https://doi.org/10.48550/arXiv.2308.07817>
- [10] Daniel Freund, Thodoris Lykouris, and Wentao Weng. 2024. Efficient Decentralized Multi-Agent Learning in Asymmetric Bipartite Queuing Systems. *Operations Research* 72, 3 (May 2024), 1049–1070. <https://doi.org/10.1287/opre.2022.0291>
- [11] Jason Gaitonde and Éva Tardos. 2023. The Price of Anarchy of Strategic Queuing Systems. *J. ACM* 70, 3 (June 2023), 1–63. <https://doi.org/10.1145/3587250>
- [12] John C. Gittins. 1979. Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society: Series B (Methodological)* 41, 2 (Jan. 1979), 148–164. <https://doi.org/10.1111/j.2517-6161.1979.tb01068.x>
- [13] John C. Gittins, Kevin D. Glazebrook, and Richard R. Weber. 2011. *Multi-Armed Bandit Allocation Indices* (2 ed.). Wiley, Chichester, UK. <https://doi.org/10.1002/9780470980033>
- [14] Kevin D. Glazebrook. 2003. An Analysis of Klimov’s Problem with Parallel Servers. *Mathematical Methods of Operations Research* 58, 1 (Sept. 2003), 1–28. <https://doi.org/10.1007/s001860300278>
- [15] Kevin D. Glazebrook and José Niño-Mora. 2001. Parallel Scheduling of Multiclass $M/M/m$ Queues: Approximate and Heavy-Traffic Optimization of Achievable Performance. *Operations Research* 49, 4 (Aug. 2001), 609–623. <https://doi.org/10.1287/opre.49.4.609.11225>
- [16] Isaac Grosf. 2023. *Optimal Scheduling in Multiserver Queues*. Ph.D. Dissertation. Carnegie Mellon University, Pittsburgh, PA. <https://isaacg1.github.io/assets/isaac-thesis.pdf>
- [17] Isaac Grosf and Michael Mitzenmacher. 2022. Incentive Compatible Queues without Money. <https://doi.org/10.48550/arXiv.2202.05747> arXiv:2202.05747 [cs]
- [18] Mor Harchol-Balter. 2013. *Performance Modeling and Design of Computer Systems: Queuing Theory in Action*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/CBO9781139226424>
- [19] Amit Harlev, George Yu, and Ziv Scully. 2024. A Gittins Policy for Optimizing Tail Latency. *ACM SIGMETRICS Performance Evaluation Review* 52, 2 (Sept. 2024), 15–17. <https://doi.org/10.1145/3695411.3695418>
- [20] Nils L. Hjort and M. C. Jones. 1996. Locally Parametric Nonparametric Density Estimation. *The Annals of Statistics* 24, 4 (Aug. 1996), 1619–1647. <https://doi.org/10.1214/aos/1032298288>
- [21] Bart Kamphorst and Bert Zwart. 2020. Heavy-Traffic Analysis of Sojourn Time under the Foreground–Background Scheduling Policy. *Stochastic Systems* 10, 1 (March 2020), 1–28. <https://doi.org/10.1287/stsy.2019.0036>
- [22] Leonard Kleinrock. 1976. *Queueing Systems, Volume 2: Computer Applications*. Wiley, New York, NY.
- [23] Gennadi P. Klimov. 1974. Time-Sharing Service Systems. I. *Theory of Probability & Its Applications* 19, 3 (1974), 532–551. <https://doi.org/10.1137/1119060>
- [24] Gennadi P. Klimov. 1978. Time-Sharing Service Systems. II. *Theory of Probability & Its Applications* 23, 2 (1978), 314–321. <https://doi.org/10.1137/1123034>
- [25] Subhashini Krishnasamy, Rajat Sen, Ramesh Johari, and Sanjay Shakkottai. 2021. Learning Unknown Service Rates in Queues: A Multiarmed Bandit Approach. *Operations Research* 69, 1 (Jan. 2021), 315–330. <https://doi.org/10.1287/opre>

2020.1995

- [26] Minghong Lin, Adam Wierman, and Bert Zwart. 2011. Heavy-Traffic Analysis of Mean Response Time under Shortest Remaining Processing Time. *Performance Evaluation* 68, 10 (Oct. 2011), 955–966. <https://doi.org/10.1016/j.peva.2011.06.001>
- [27] John D. C. Little. 2011. Little’s Law as Viewed on Its 50th Anniversary. *Operations Research* 59, 3 (June 2011), 536–549. <https://doi.org/10.1287/opre.1110.0940>
- [28] Pascal Massart. 1990. The Tight Constant in the Dvoretzky-Kiefer-Wolfowitz Inequality. *The Annals of Probability* 18, 3 (July 1990), 1269–1283. <https://doi.org/10.1214/aop/1176990746>
- [29] Benjamin Moseley, Heather Newman, Kirk Pruhs, and Rudy Zhou. 2025. Robust Gittins for Stochastic Scheduling. <https://doi.org/10.48550/arXiv.2504.10743> arXiv:2504.10743 [cs]
- [30] Shefali Ramakrishna, Amit Harlev, and Ziv Scully. 2025. Empirical Gittins: M/G/1 Scheduling from Job Size Samples. *ACM SIGMETRICS Performance Evaluation Review* 53, 2 (Aug. 2025), 119–121. <https://doi.org/10.1145/3764944.3764974>
- [31] Henrik Ramlau-Hansen. 1983. Smoothing Counting Process Intensities by Means of Kernel Functions. *The Annals of Statistics* 11, 2 (June 1983), 453–466. <https://doi.org/10.1214/aos/1176346152>
- [32] Rhonda Righter, J. George Shanthikumar, and Genji Yamazaki. 1990. On Extremal Service Disciplines in Single-Stage Queueing Systems. *Journal of Applied Probability* 27, 2 (June 1990), 409–416. <https://doi.org/10.2307/3214660>
- [33] Francisco Robledo Relaño, Vivek Borkar, Urtzi Ayesta, and Konstantin Avrachenkov. 2024. Tabular and Deep Learning for the Whittle Index. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 9, 3, Article 11 (Sept. 2024), 21 pages. <https://doi.org/10.1145/3670686>
- [34] Linus E. Schrage. 1968. A Proof of the Optimality of the Shortest Remaining Processing Time Discipline. *Operations Research* 16, 3 (June 1968), 687–690. <https://doi.org/10.1287/opre.16.3.687>
- [35] Linus E. Schrage and Louis W. Miller. 1966. The Queue M/G/1 with the Shortest Remaining Processing Time Discipline. *Operations Research* 14, 4 (Aug. 1966), 670–684. <https://doi.org/10.1287/opre.14.4.670>
- [36] Ziv Scully. 2022. *A New Toolbox for Scheduling Theory*. Ph. D. Dissertation. Carnegie Mellon University, Pittsburgh, PA. <https://ziv.codes/pdf/scully-thesis.pdf>
- [37] Ziv Scully, Isaac Grosf, and Mor Harchol-Balter. 2020. The Gittins Policy Is Nearly Optimal in the M/G/k under Extremely General Conditions. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 3, Article 43 (Dec. 2020), 29 pages. <https://doi.org/10.1145/3428328>
- [38] Ziv Scully, Isaac Grosf, and Mor Harchol-Balter. 2021. Optimal Multiserver Scheduling with Unknown Job Sizes in Heavy Traffic. *Performance Evaluation* 145, Article 102150 (Jan. 2021), 31 pages. <https://doi.org/10.1016/j.peva.2020.102150>
- [39] Ziv Scully, Isaac Grosf, and Michael Mitzenmacher. 2022. Uniform Bounds for Scheduling with Job Size Estimates. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022) (Leibniz International Proceedings in Informatics (LIPIcs))*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Berkeley, CA (virtual), Article 41, 30 pages. <https://doi.org/10.4230/LIPIcs.ITCS.2022.114>
- [40] Ziv Scully and Mor Harchol-Balter. 2018. SOAP Bubbles: Robust Scheduling under Adversarial Noise. In *56th Annual Allerton Conference on Communication, Control, and Computing*. IEEE, Monticello, IL, 144–154. <https://doi.org/10.1109/ALLERTON.2018.8635963>
- [41] Ziv Scully and Mor Harchol-Balter. 2021. The Gittins Policy in the M/G/1 Queue. In *19th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt 2021)*. IEEE, Philadelphia, PA (virtual), 248–255. <https://doi.org/10.23919/WiOpt52861.2021.9589051>
- [42] Ziv Scully and Mor Harchol-Balter. 2021. How to Schedule Near-Optimally under Real-World Constraints. <https://doi.org/10.48550/arXiv.2110.11579> arXiv:2110.11579 [cs]
- [43] Ziv Scully, Mor Harchol-Balter, and Alan Scheller-Wolf. 2018. SOAP: One Clean Analysis of All Age-Based Scheduling Policies. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 2, 1, Article 16 (March 2018), 30 pages. <https://doi.org/10.1145/3179419>
- [44] Ziv Scully, Mor Harchol-Balter, and Alan Scheller-Wolf. 2020. Simple Near-Optimal Scheduling for the M/G/1. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 4, 1, Article 11 (March 2020), 29 pages. <https://doi.org/10.1145/3379477>
- [45] Ziv Scully and Alexander Terenin. 2025. The Gittins Index: A Design Principle for Decision Making under Uncertainty. In *TutORials in Operations Research: Advances in Analytics and Operations Research: Improving Decisions to Secure the Future*, Brian Denton, Nilay Noyan, Yao Xie, David Alderson, and Harvey J. Greenberg (Eds.). INFORMS, Catonsville, MD, 28–70. <https://doi.org/10.1287/educ.2025.0290>
- [46] Ziv Scully and Lucas van Kreveld. 2025. When Does the Gittins Policy Have Asymptotically Optimal Response Time Tail in the M/G/1? *Operations Research* 73, 3 (May 2025), 1412–1429. <https://doi.org/10.1287/opre.2022.0038>
- [47] Flore Sentenac, Etienne Boursier, and Vianney Perchet. 2021. Decentralized Learning in Online Queueing Systems. In *Advances in Neural Information Processing Systems (NeurIPS 2021)*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang,

- and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., virtual event, 18501–18512. <https://doi.org/10.48550/arXiv.2106.04228>
- [48] Kenneth C. Sevick. 1974. Scheduling for Minimum Total Loss Using Service Time Distributions. *J. ACM* 21, 1 (Jan. 1974), 66–75. <https://doi.org/10.1145/321796.321803>
- [49] David A. Stanford, Peter Taylor, and Ilze Ziedins. 2014. Waiting Time Distributions in the Accumulating Priority Queue. *Queueing Systems* 77, 3 (July 2014), 297–330. <https://doi.org/10.1007/s11134-013-9382-6>
- [50] Rocco van Vreumingen. 2019. *Queueing Systems with Non-Standard Service Policies and Server Vacations*. MSc. University of Amsterdam, Amsterdam, The Netherlands. <http://www.mt-support.nl/VanVreumingen/Rocco/wp-content/uploads/2020/01/Master-thesis-Queueing-systems-with-non-standard-service-policies-and-server-vacations.pdf>
- [51] G. von Olivier. 1972. Kostenminimale Prioritäten in Wartesystemen vom Typ M/G/1 [Cost-minimum priorities in queueing systems of type M/G/1]. *Elektronische Rechenanlagen* 14, 6 (Dec. 1972), 262–271. <https://doi.org/10.1524/itit.1972.14.16.262>
- [52] Richard R. Weber. 1992. On the Gittins Index for Multiarmed Bandits. *The Annals of Applied Probability* 2, 4 (Nov. 1992), 1024–1033. <https://doi.org/10.1214/aoap/1177005588>
- [53] Jon A. Wellner. 1978. Limit Theorems for the Ratio of the Empirical Distribution Function to the True Distribution Function. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 45, 1 (March 1978), 73–88. <https://doi.org/10.1007/bf00635964>
- [54] Adam Wierman. 2007. *Scheduling for Today's Computer Systems: Bridging Theory and Practice*. Ph.D. Dissertation. Carnegie Mellon University, Pittsburgh, PA. <https://adamwierman.com/wp-content/uploads/2020/09/thesis.pdf>
- [55] Zixian Yang, R. Srikant, and Lei Ying. 2023. Learning While Scheduling in Multi-Server Systems with Unknown Statistics: MaxWeight with Discounted UCB. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS 2023) (Proceedings of Machine Learning Research, Vol. 206)*, Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (Eds.). PMLR, Valencia, Spain, 4275–4312. <https://proceedings.mlr.press/v206/yang23d.html>
- [56] Qing Zhao. 2020. *Multi-Armed Bandits: Theory and Applications to Online Learning in Networks*. Springer, Cham, Switzerland. <https://doi.org/10.1007/978-3-031-79289-2>
- [57] Yueyang Zhong, John R. Birge, and Amy R. Ward. 2025. Learning to Schedule in Multiclass Many-Server Queues with Abandonment. *Operations Research* 73, 6 (Nov. 2025), 3085–3103. <https://doi.org/10.1287/opre.2022.0197>

A Formal proof of single-job WINE for empirical distributions

In this section we formalize the inductive proof outlined in Section 5.1 to prove Theorem 5.3. Since it will be useful throughout the proof, we begin by first making a few observations about the structure of $r_{Y(G)}(a)$.

Observation A.1 (Structure of $r_{Y(G)}(a)$). Recall that

$$r_{Y(G)}(a) = \inf_{b>a} \frac{\mathbb{E}_{S \sim G}[S \wedge b - a \mid S > a]}{\mathbb{P}_{S \sim G}[S \leq b \mid S > a]}.$$

Using this definition and the fact that G is a discrete uniform distribution, it is easy to verify the following.

- (1) The values of b that attain the infimum in the definition of $r_{Y(G)}(a)$ must be atoms of G .
- (2) Therefore, on each interval $[G_i, G_{i+1})$ the rank function $r_{Y(G)}(a)$ is the point-wise minimum of $n - i$ decreasing linear functions. Thus, $r_{Y(G)}(a)$ is piecewise linear and decreasing on each such interval. See Fig. 5.1 for an example illustrating this behavior.

LEMMA A.2 (BASE CASE, $n = 1$). *Let G be the distribution under which $X = G_1 > 0$ almost surely. Then $\varphi_{\infty, G}(0) = 1 = h_1$.*

PROOF. Recall that $b_{Y(G)}(a, r) = \inf\{b \geq a : r_{Y(G)}(b) \geq r\}$, so along with Observation A.1 we know that

$$b_{Y(G)}(0, r) = \begin{cases} 0 & r \leq r_{Y(G)}(0) \\ G_1 & r > r_{Y(G)}(0) \end{cases}.$$

Thus,

$$\varphi_{\infty,G}(0) = \int_0^{\infty} \frac{b_{Y(G)}(0,r)}{r^2} dr = \int_{r_{Y(G)}(0)}^{\infty} \frac{G_1}{r^2} dr = \frac{G_1}{r_{Y(G)}(0)}.$$

Observation A.1 also tells us that

$$r_{Y(G)}(0) = \frac{\mathbb{E}_{S \sim G}[S \wedge G_1]}{\mathbb{P}_{S \sim G}[S \leq G_1]} = G_1$$

so $\varphi_{\infty,G}(0) = 1 = h_1$. □

We separate the inductive step into two cases:

- (1) $r_{Y(G)}(0) \leq r_{Y(G)}(G_1)$, in which case we can apply the inductive hypothesis, and
- (2) $r_{Y(G)}(0) > r_{Y(G)}(G_1)$, in which case we must “slide” the atoms of the distribution to reduce to the first case.

PROPOSITION A.3 (FIRST CASE). *Assume that for any discrete uniform distribution D with $n-1$ atoms, $\varphi_{\infty,D}(0) \leq h_{n-1}$. Let G be a discrete uniform distribution with n atoms and $r_{Y(G)}(0) \leq r_{Y(G)}(G_1)$. Then $\varphi_{\infty,G}(0) \leq h_n$.*

PROOF. This proof requires our earlier observation that $\varphi_{\infty,G}(a)$ can be represented as an integral of $1/r^2$ over some region, which we previously referred to as V :

$$\varphi_{\infty,G}(a) = \int_0^{\infty} \int_a^{b_{Y(G)}(a,r)} \frac{1}{r^2} du dr = \int_a^{\infty} \int_{w_{Y(G)}(a,u)}^{\infty} \frac{1}{r^2} dr du,$$

where $w_{Y(G)}(a,u) = \sup_{a \leq t \leq u} r_{Y(G)}(t)$ is the worst rank attained by a job between ages a and u . Note that we reverse the order of integration to facilitate splitting the region of integration by age in the next step of the proof. The dashed blue line in Fig. 5.1 represents $w_{Y(G)}(0,u)$.

The remainder of the proof splits the region of integration into two parts and deals with each separately:

- (1) ages in $[0, G_1)$, and
- (2) ages in $[G_1, \infty)$.

The contribution from the first region can be computed explicitly, while the second is handled using the inductive hypothesis.

We start by showing how to apply the inductive hypothesis in the latter case. Observe that $r_{Y(G)}(0) \leq r_{Y(G)}(G_1)$ implies $w_{Y(G)}(0,u) = w_{Y(G)}(G_1,u)$ for all $u \geq G_1$.⁹ And then by using (5.1), we can additionally conclude that $w_{Y(G)}(G_1,u) = w_{Y(G|_{>G_1})}(0,u - G_1)$. Thus,

$$\begin{aligned} \int_{G_1}^{\infty} \int_{w_{Y(G)}(0,u)}^{\infty} \frac{1}{r^2} dr du &= \int_{G_1}^{\infty} \int_{w_{Y(G|_{>G_1})}(0,u-G_1)}^{\infty} \frac{1}{r^2} dr du \\ &= \int_0^{\infty} \int_{w_{Y(G|_{>G_1})}(0,u)}^{\infty} \frac{1}{r^2} dr du = \varphi_{\infty,G|_{>G_1}}(0). \end{aligned}$$

Since $G|_{>G_1}$ is a discrete uniform distribution with $n-1$ atoms, the above is at most h_{n-1} ; see Fig. 5.2 for an illustration of this.

We now explicitly compute the contribution from ages in $[0, G_1)$:

$$\int_0^{G_1} \int_{w_{Y(G)}(0,u)}^{\infty} \frac{1}{r^2} dr du = \int_0^{G_1} \int_{r_{Y(G)}(0)}^{\infty} \frac{1}{r^2} dr du = \frac{G_1}{r_{Y(G)}(0)} \quad (\text{A.1})$$

⁹To see why $r_{Y(G)}(0) \leq r_{Y(G)}(G_1)$ is necessary, compare Fig. 5.1 and Fig. 5.2.

where the second equality follows from $r_{Y(G)}(0) \leq r_{Y(G)}(G_1)$ and the fact that $r_{Y(G)}(a)$ decreases on $[0, G_1)$ (Observation A.1). To compute $r_{Y(G)}(0)$ we will make use of the following standard fact about the Gittins policy.

LEMMA A.4 (COROLLARY OF [13, LEMMA 2.2]). *Let G be a discrete size distribution.¹⁰ For all ages $a \geq 0$, the infimum in*

$$r_{Y(G)}(a) = \inf_{b>a} \frac{\mathbb{E}_{S \sim G}[S \wedge b - a \mid S > a]}{\mathbb{P}_{S \sim G}[S \leq b \mid S > a]}$$

is attained at $b = b_{Y(G)}(a, r_{Y(G)}(a))$ (Definition 3.2), the earliest age $b > a$ of equal or greater rank, i.e. such that $r_{(G)}(b) > r_{(G)}(a)$.

Because $r_{Y(G)}(0) \leq r_{Y(G)}(G_1)$, Lemma A.4 implies

$$r_{Y(G)}(0) = \frac{\mathbb{E}_{S \sim G}[S \wedge G_1]}{\mathbb{P}_{S \sim G}[S \leq G_1]} = nG_1,$$

which together with (A.1) yields

$$\int_0^{G_1} \int_{w_{Y(G)}(0,u)}^{\infty} \frac{1}{r^2} dr du = \frac{G_1}{nG_1} = \frac{1}{n}.$$

Putting everything together,

$$\varphi_{\infty,G}(0) = \int_0^{G_1} \int_{w_{Y(G)}(0,u)}^{\infty} \frac{1}{r^2} dr du + \int_{G_1}^{\infty} \int_{w_{Y(G)}(0,u)}^{\infty} \frac{1}{r^2} dr du \leq \frac{1}{n} + h_{n-1} = h_n. \quad \square$$

PROPOSITION A.5 (SECOND CASE). *Assume that for any discrete uniform distribution D with $n - 1$ atoms, $\varphi_{\infty,D}(0) \leq h_{n-1}$. Let G be a discrete uniform distribution with n atoms and $r_{Y(G)}(0) > r_{Y(G)}(G_1)$. Then $\varphi_{\infty,G}(0) \leq h_n$.*

PROOF. We will prove this by reducing to the first case, Proposition A.3. We do this by constructing a new discrete uniform distribution G' that satisfies

$$\varphi_{\infty,G}(0) \leq \varphi_{\infty,G'}(0)$$

and satisfies the requirements of Proposition A.3. In particular, we will show that it suffices to let $G' = G|_{>G_1-\varepsilon}$ for some sufficiently small $\varepsilon > 0$.

To see that there exists a choice of $\varepsilon > 0$ such that $G|_{>G_1-\varepsilon}$ satisfies

$$r_{Y(G|_{>G_1-\varepsilon})}(0) \leq r_{Y(G|_{>G_1-\varepsilon})}(\varepsilon) \tag{A.2}$$

(note that the first atom of $G|_{>G_1-\varepsilon}$ is at ε), observe that $r_{Y(G)}(a) \rightarrow 0$ as $a \uparrow G_1$ because for all $a \in [0, G_1)$, by picking $b = G_1$ in the infimum in Definition 2.5, we have

$$r_{Y(G)}(a) \leq \frac{G_1 - a}{1/n} \rightarrow 0 \quad \text{as } a \rightarrow 0.$$

This means there must be some $\varepsilon > 0$ such that $r_{Y(G)}(G_1 - \varepsilon) \leq r_{Y(G)}(G_1)$, so (A.2) follows by (5.1), as illustrated in Figs. 5.1 and 5.2.

Now we must prove that $\varphi_{\infty,G}(0) \leq \varphi_{\infty,G|_{>G_1-\varepsilon}}(0)$. We do this by showing that, for all $0 \leq t < G_1$,

$$\partial_t^+ \varphi_{\infty,G|_{>t}}(0) \geq 0,$$

¹⁰We restrict to the discrete case only because Gittins et al. [13, Lemma 2.2] is stated for discrete Markov chains, but the same statement is true more generally. For instance, combining two results of Aalto et al. [1, Corollary 1 and Lemma 5] yields nearly the same statement for continuous distributions.

where ∂_t^+ denotes the right-derivative with respect to t . For $0 \leq t < G_1$, let $k(t)$ denote the smallest index $i > 0$ such that the rank at G_i is at least the rank at t :

$$k(t) = \inf\{1 \leq i < n : r_{Y(G)}(G_i) \geq r_{Y(G)}(t)\}.$$

Since $r_{Y(G)}(a)$ is piecewise linear and decreasing on $[0, G_1)$ (Observation A.1), $k(t)$ is monotone decreasing and right-locally constant on $[0, G_1)$.¹¹

Observe that for $0 \leq t < G_1$,

$$\begin{aligned} \varphi_{\infty, G|_{>t}}(0) &= \int_0^\infty \int_{w_{Y(G|_{>t})(0,u)}}^\infty \frac{1}{r^2} dr du \\ &= \int_t^\infty \int_{w_{Y(G)}(t,u)}^\infty \frac{1}{r^2} dr du \\ &= \int_t^{G_{k(t)}} \int_{r_{Y(G)}(t)}^\infty \frac{1}{r^2} dr du + \int_{G_{k(t)}}^\infty \int_{w_{Y(G)}(G_{k(t)},u)}^\infty \frac{1}{r^2} dr du, \end{aligned}$$

where the last equality used the definition of $k(t)$ to rewrite the lower bounds of both integrals with respect to r . Since $k(t)$ is right-locally constant, the second term disappears when we take a right-derivative:

$$\begin{aligned} \partial_t^+ \varphi_{\infty, G|_{>t}}(0) &= \partial_t^+ \int_t^{G_{k(t)}} \int_{r_{Y(G)}(t)}^\infty \frac{1}{r^2} dr du \\ &= \partial_t^+ \frac{G_{k(t)} - t}{r_{Y(G)}(t)} = \frac{-r_{Y(G)}(t) - \partial_t^+ r_{Y(G)}(t) \cdot (G_{k(t)} - t)}{r_{Y(G)}(t)^2}. \end{aligned} \quad (\text{A.3})$$

Additionally, there exists a $\delta > 0$ such that for all $t \leq a < t + \delta$,

$$r_{Y(G)}(a) = \frac{\mathbb{E}_{S \sim G}[S \wedge G_{k(t)} - a]}{\mathbb{P}_{S \sim G}[S \leq G_{k(t)}]},$$

and so

$$\partial_t^+ r_{Y(G)}(t) = -\frac{1}{\mathbb{P}_{S \sim G}[S \leq G_{k(t)}]}.$$

Finally, we get that $\partial_t^+ \varphi_{\infty, G|_{>t}}(0) \geq 0$ by plugging into (A.3) and noting that,

$$r_{Y(G)}(t) = \frac{\mathbb{E}_{S \sim G}[S \wedge G_{k(t)} - t]}{\mathbb{P}_{S \sim G}[S \leq G_{k(t)}]} \leq \frac{G_{k(t)} - t}{\mathbb{P}_{S \sim G}[S \leq G_{k(t)}]}.$$

Since we wish to conclude that $\varphi_{\infty, G}(0) \leq \varphi_{\infty, G|_{>G_1-\varepsilon}}(0)$ from $\partial_t^+ \varphi_{\infty, G|_{>t}}(0) \geq 0$, the last remaining step is to show that $\varphi_{\infty, G|_{>t}}(0)$ is right-continuous with respect to t . However, this follows from the double-integral representation of $\varphi_{\infty, G|_{>t}}(0)$. \square

Combining Lemma A.2 (base case) with Propositions A.3 and A.5 (inductive step) completes the induction and proves Theorem 5.3.

¹¹We say that $k(t)$ is right-locally constant at t if there exists an $\varepsilon > 0$ such that $k(t)$ is constant on $[t, t + \varepsilon)$

B Details for asymptotic analysis

B.1 Implications of Matuszewska index bounds

We summarize below some implications of (6.1) used throughout our proofs of results from Section 6. All of them follow from routine computation. Below, big- O notation refers to the $x \rightarrow \infty$ limit:

$$\begin{aligned}\bar{F}(x) &\leq O(x^{-\alpha}), \\ x &\leq O(\bar{F}(x)^{-1/\alpha}), \\ \mathbb{E}[S \mid S > x] &\leq O(x), \\ \mathbb{E}[(S \wedge x)^2] &\leq O(x^{2-\alpha}) \quad \text{if } \alpha < 2.\end{aligned}$$

Moreover, if G is multiplicatively close to F up to threshold x (or greater), then these still hold if \bar{F} is replaced by \bar{G} . In the contexts where we use these facts, x will be a function of n and ρ such that $\ell \rightarrow 0$ as $n \rightarrow \infty$ and $\rho \rightarrow 1$.

A final property is an implication of a result of Lin et al. [26]: if (6.1) holds with $\alpha > 2$, then the mean response time of true Gittins is bounded below by [26, Theorem 2 and commentary at the end of Section 3.1]

$$\mathbb{E}_{\gamma(F)}[T] \geq \Omega\left(\left(\frac{1}{1-\rho}\right)^{\frac{-(\alpha-2)}{\alpha-1}}\right).$$

Lin et al. [26] actually give lower bounds on SRPT rather than Gittins, but optimality of SRPT among all policies (including clairvoyant policies) implies the lower bound holds for Gittins, too.

Nearly all of the properties above actually hold under the weaker assumption of $\bar{F}(x) \leq O(x^{-\alpha})$. But (6.1) is essential to bound $\mathbb{E}[S \mid S > x] \leq O(x)$, which is important in both the truncated and untruncated cases; and, of course, to bound tail ratios, which appear in the untruncated case.

B.2 Proofs of main asymptotic analysis results

In order to determine the value of \bar{G} that minimizes the leading-order contribution to the multiplicative error, we must rewrite the sample complexity bound using \bar{G} , replacing the unknown $\bar{F}(\ell)$ through the multiplicative-closeness relation from Lemma 2.3.

LEMMA B.1. *Let G denote the empirical distribution based on n samples from job size distribution F . Under the finite-sample ε -multiplicative-closeness guarantee of Lemma 2.3, the smallest ε for which the bound of Theorem 4.1 holds satisfies*

$$\varepsilon = \tilde{O}\left((n\bar{G}(\ell))^{-1/2}\right).$$

PROOF. Note that, from Theorem 4.1, if we have a fixed number of samples n and fix $\bar{F}(\ell)$, then ε must satisfy,

$$\varepsilon \geq \sqrt{\frac{3 \log(2/\delta)}{\bar{F}(\ell) n}}, \tag{B.1}$$

for our error bound to hold. Lemma 2.3 tells us that F and G are ε -multiplicatively close, and so

$$\frac{e^{-\varepsilon}}{\bar{G}(\ell)} \leq \frac{1}{\bar{F}(\ell)} \leq \frac{e^{\varepsilon}}{\bar{G}(\ell)}.$$

Substituting the rightmost inequality into (B.1) yields a sufficient condition:

$$\varepsilon e^{-\varepsilon/2} \geq \sqrt{\frac{3 \log(2/\delta)}{\bar{G}(\ell) n}}. \tag{B.2}$$

Equation (B.2) implicitly defines the smallest ε that can satisfy the finite-sample guarantee for a given $\bar{G}(\ell)$. Although the expression cannot be inverted exactly, the exponential term only slightly perturbs the scaling for small ε . In particular, $e^{-\varepsilon/2} > 0.6$ for all $\varepsilon < 1$, so (B.2) implies the simpler bound

$$\varepsilon \geq \frac{5}{3} \sqrt{\frac{3 \log(2/\delta)}{\bar{G}(\ell) n}}. \quad \square$$

THEOREM 6.1 (TRUNCATION, FINITE VARIANCE). *Let $\alpha > 2$ be such that F satisfies (6.1), and let G be the empirical distribution constructed from n samples from F . Consider the truncated empirical Gittins policy $\gamma(G_\ell)$ for some truncation level ℓ , which may depend on G . Choosing ℓ such that*

$$\bar{G}(\ell) = \Theta\left(\min\left(n^{-1/3}(1-\rho)^{\frac{2\alpha}{3(\alpha-1)}}, n^{-\frac{\alpha}{3\alpha-2}}(1-\rho)^{\frac{2\alpha}{(3\alpha-2)(\alpha-1)}}\right)\right)$$

yields, with probability at least $1 - \delta$, multiplicative error bounded by

$$\frac{\mathbb{E}_{\gamma(G_\ell)}[T]}{\mathbb{E}_{\gamma(F)}[T]} - 1 \leq \tilde{O}\left(\max\left(n^{-1/3}(1-\rho)^{\frac{-\alpha}{3(\alpha-1)}}, n^{-\frac{-(\alpha-1)}{3\alpha-2}}(1-\rho)^{\frac{-\alpha}{(3\alpha-2)(\alpha-1)}}\right)\right).$$

PROOF. Substituting the values from Appendix B.1 and Lemma B.1 into the multiplicative error from Theorem 4.1 yields, with probability at least $1 - \delta$,

$$\begin{aligned} \frac{\mathbb{E}_{\gamma(G_\ell)}[T]}{\mathbb{E}_{\gamma(F)}[T]} - 1 &\leq e^{4\varepsilon} - 1 + \frac{\bar{F}(\ell)}{\mathbb{E}_{\gamma(F)}[T]} \left(\frac{\lambda \mathbb{E}[S^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S | S \geq \ell]}{(1-\rho)} \right) \\ &\leq O\left(\varepsilon + (1-\rho)^{\frac{\alpha-2}{\alpha-1}} \left(\frac{\bar{G}(\ell)}{(1-\rho)^2} + \frac{\bar{G}(\ell) \cdot \ell}{1-\rho} \right)\right) \\ &\leq \tilde{O}\left(n^{-1/2} \bar{G}(\ell)^{-1/2} + (1-\rho)^{\frac{\alpha-2}{\alpha-1}} \left(\frac{\bar{G}(\ell)}{(1-\rho)^2} + \frac{\bar{G}(\ell)^{\frac{\alpha-1}{\alpha}}}{1-\rho} \right)\right), \end{aligned}$$

from which the result follows by a routine computation, given in Lemma D.3(a) with $q = 1/2$. \square

THEOREM 6.2 (TRUNCATION, POTENTIALLY INFINITE VARIANCE). *Let $\alpha \in (1, 2)$ be such that F satisfies (6.1), and let G be the empirical distribution constructed from n samples from F . Consider the truncated empirical Gittins policy $\gamma(G_\ell)$ for some truncation level ℓ , which may depend on G . Choosing ℓ such that*

$$\bar{G}(\ell) = \Theta\left(\min\left(n^{-\frac{\alpha}{5\alpha-4}}(1-\rho)^{\frac{4\alpha}{5\alpha-4}}, n^{-\frac{\alpha}{3\alpha-2}}(1-\rho)^{\frac{2\alpha}{3\alpha-2}}\right)\right)$$

yields, with probability at least $1 - \delta$, multiplicative error bounded by

$$\frac{\mathbb{E}_{\gamma(G_\ell)}[T]}{\mathbb{E}_{\gamma(F)}[T]} - 1 \leq \tilde{O}\left(\max\left(n^{-\frac{2(\alpha-1)}{5\alpha-4}}(1-\rho)^{\frac{-2\alpha}{5\alpha-4}}, n^{-\frac{-(\alpha-1)}{3\alpha-2}}(1-\rho)^{\frac{-\alpha}{3\alpha-2}}\right)\right).$$

PROOF. Substituting the values from Appendix B.1 and Lemma B.1 into the multiplicative error from Theorem 4.1 yields, with probability at least $1 - \delta$,

$$\begin{aligned} \frac{\mathbb{E}_{Y(G)}[T]}{\mathbb{E}_{Y(F)}[T]} - 1 &\leq e^{4\varepsilon} - 1 + \frac{\bar{F}(\ell)}{\mathbb{E}_{Y(F)}[T]} \left(\frac{\lambda \mathbb{E}[(S \wedge \ell)^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S \mid S \geq \ell]}{(1-\rho)} \right) \\ &\leq O \left(\varepsilon + \left(\frac{\bar{G}(\ell) \cdot \ell^{2-\alpha}}{(1-\rho)^2} + \frac{\bar{G}(\ell) \cdot \ell}{1-\rho} \right) \right) \\ &\leq \tilde{O} \left(n^{-1/2} \bar{G}(\ell)^{-1/2} + \frac{\bar{G}(\ell)^{\frac{2(\alpha-1)}{\alpha}}}{(1-\rho)^2} + \frac{\bar{G}(\ell)^{\frac{\alpha-1}{\alpha}}}{1-\rho} \right), \end{aligned}$$

from which the result follows by a routine computation, given in Lemma D.3(b) with $q = 1/2$. \square

THEOREM 6.3 (NO TRUNCATION, FINITE VARIANCE). *Let $\alpha > 2$ be such that F satisfies (6.1), and let G be the empirical distribution constructed from n samples from F . The empirical Gittins policy $\gamma(G)$ achieves, with probability at least $1 - \delta$, multiplicative error bounded by*

$$\frac{\mathbb{E}_{Y(G)}[T]}{\mathbb{E}_{Y(F)}[T]} - 1 \leq \tilde{O} \left(\max \left(n^{-1/4} (1-\rho)^{\frac{-3\alpha}{4(\alpha-1)}}, n^{\frac{-(\alpha-1)}{4\alpha-3}} (1-\rho)^{\frac{-3\alpha}{(4\alpha-3)(\alpha-1)}} \right) \right).$$

PROOF. Substituting the values from (B.1) and (6.1) and Appendix B.1 into the multiplicative error from Theorem 5.1 yields, with probability at least $1 - \delta$,

$$\begin{aligned} \frac{\mathbb{E}_{Y(G)}[T]}{\mathbb{E}_{Y(F)}[T]} - 1 &\leq \frac{e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n}{e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)}} \left(1 + \frac{\bar{F}(k)(1+h_n)}{(1-\rho)\mathbb{E}_{Y(F)}[T]} \left(\frac{\lambda \mathbb{E}[S^2]}{2(1-\rho)} + \mathbb{E}[S \mid S > k] \right) \right) - 1 \\ &\leq \tilde{O} \left(\varepsilon + \frac{\bar{F}(\ell)}{\bar{F}(k)} + (1-\rho)^{\frac{\alpha-2}{\alpha-1}} \left(\frac{\bar{F}(k)}{(1-\rho)^2} + \bar{F}(k) \cdot k \right) \right) \\ &\leq \tilde{O} \left(n^{-1/2} \bar{F}(\ell)^{-1/2} + \frac{\bar{F}(\ell)}{\bar{F}(k)} + (1-\rho)^{\frac{\alpha-2}{\alpha-1}} \left(\frac{\bar{F}(k)}{(1-\rho)^2} + \bar{F}(k)^{\frac{\alpha-1}{\alpha}} \right) \right). \end{aligned}$$

Choosing ℓ such that $\bar{F}(\ell) = \Theta(n^{-1/3} \bar{F}(k)^{2/3})$ equalizes the first two terms up to constants, yielding

$$\frac{\mathbb{E}_{Y(G)}[T]}{\mathbb{E}_{Y(F)}[T]} - 1 \leq \tilde{O} \left(n^{-1/3} \bar{F}(k)^{-1/3} + (1-\rho)^{\frac{\alpha-2}{\alpha-1}} \left(\frac{\bar{F}(k)}{(1-\rho)^2} + \bar{F}(k)^{\frac{\alpha-1}{\alpha}} \right) \right)$$

from which the result follows by a routine computation, given in Lemma D.3(a) with $q = 1/3$. \square

C Reducing empirical distributions to discrete uniform distributions

In this section, we justify the “without loss of generality” assumption made in Section 5.1, under which the empirical distribution G is assumed to be discrete uniform.

THEOREM C.1. *Let G be an empirical job size distribution constructed from n samples. Then there exists a choice of $0 < G_1 < G_2 < \dots < G_n$ such that the discrete uniform distribution*

$$G' = \text{Unif}\{G_1, \dots, G_n\}$$

satisfies $\varphi_{\infty, G}(0) \leq \varphi_{\infty, G'}(0)$.

Let G be an empirical job size distribution generated from n samples that has an atom at $v > 0$ with probability mass $\frac{m}{n}$ for some $m > 1$. Let $G(v, \delta)$ be a distribution identical to G except that:

- (1) the probability mass of the atom at v has been reduced to $\frac{1}{n}$,

(2) there is a new atom at $v - \delta$ with probability mass $\frac{m-1}{n}$. See Example C.2 for an explicit example of this construction.

Example C.2. If $X \sim G$ is given by

$$X = \begin{cases} 1 & \text{with probability } \frac{1}{4}, \\ 2 & \text{with probability } \frac{3}{4}, \\ 3 & \text{with probability } \frac{1}{4}, \end{cases}$$

then $X \sim G(2, 0.5)$ is given by

$$X = \begin{cases} 1 & \text{with probability } \frac{1}{4}, \\ 1.5 & \text{with probability } \frac{2}{4}, \\ 2 & \text{with probability } \frac{1}{4}, \\ 3 & \text{with probability } \frac{1}{4}. \end{cases}$$

LEMMA C.3. *Let G be an empirical job size distribution constructed from n samples. Then either, (a) G is a discrete uniform distribution on n values, or (b) there exist $v, \delta > 0$ such that*

$$\varphi_{\infty, G}(0) \leq \varphi_{\infty, G(v, \delta)}(0).$$

PROOF. If G is an empirical job size distribution constructed from n samples and is not discrete uniform, then some atom of G has probability mass $\frac{m}{n}$ with $m > 1$. Denote the location of this atom by $v > 0$. Let $u \geq 0$ denote the largest atom strictly below v , and set $u = 0$ if no such atom exists. We claim that choosing

$$\delta < \min\left(\frac{r_{\gamma(G)}(u)}{n}, v - u\right)$$

ensures that $G(v, \delta)$ has the desired property. To prove this, we will show that for all $a \geq 0$,

$$w_{\gamma(G(v, \delta))}(a) \leq w_{\gamma(G)}(a)$$

and thus,

$$\varphi_{\infty, G}(0) = \int_0^{\infty} \int_{w_{\gamma(G)}(0, u)}^{\infty} \frac{1}{r^2} dr du \leq \int_0^{\infty} \int_{w_{\gamma(G(v, \delta))}(0, u)}^{\infty} \frac{1}{r^2} dr du = \varphi_{\infty, G(v, \delta)}(0).$$

We prove this by considering three cases:

(1) Fix $0 < a < v - \delta$. There exists a $b^* > a$ such that

$$r_{\gamma(G)}(a) = \frac{\mathbb{E}_{S \sim G}[S \wedge b^* - a \mid S > a]}{\mathbb{P}_{S \sim G}[S \leq b^* \mid S > a]}.$$

It is clear from the construction of $G(v, \delta)$ that

$$\mathbb{E}_{S \sim G(v, \delta)}[S \wedge b^* - a \mid S > a] \leq \mathbb{E}_{S \sim G}[S \wedge b^* - a \mid S > a]$$

and that

$$\mathbb{P}_{S \sim G(v, \delta)}[S \leq b^* \mid S > a] \geq \mathbb{P}_{S \sim G}[S \leq b^* \mid S > a].$$

Thus,

$$\begin{aligned} r_{\gamma(G(v, \delta))}(a) &= \inf_{b > a} \frac{\mathbb{E}_{S \sim G(v, \delta)}[S \wedge b - a \mid S > a]}{\mathbb{P}_{S \sim G(v, \delta)}[S \leq b \mid S > a]} \\ &\leq \frac{\mathbb{E}_{S \sim G(v, \delta)}[S \wedge b^* - a \mid S > a]}{\mathbb{P}_{S \sim G(v, \delta)}[S \leq b^* \mid S > a]} \leq r_{\gamma(G)}(a), \end{aligned}$$

and so $w_{Y(G(v,\delta))}(a) \leq w_{Y(G)}(a)$ for all $0 \leq a < v - \delta$.

(2) Fix $v - \delta \leq a < v$. Then since $u < v - \delta < a < v$ and $\delta < \frac{r_{Y(G)}(u)}{n}$,

$$r_{Y(G(v,\delta))}(a) \leq \frac{\mathbb{E}_{S \sim G(v,\delta)}[S \wedge v - a \mid S > a]}{\mathbb{P}_{S \sim G(v,\delta)}[S \leq v \mid S > a]} = \frac{v - a}{\frac{1}{n}} \leq \delta n < r_{Y(G)}(u).$$

Since $u < a$ by assumption, $w_{Y(G(v,\delta))}(a) \leq w_{Y(G)}(a)$ for all $v - \delta \leq a < v$.

(3) Fix $a \geq v$. Because $G(v, \delta)$ agrees with G on $[v, \infty)$, their conditional distributions above v are identical. Thus, $r_{Y(G(v,\delta))}(a) = r_{Y(G)}(a)$. Combined with the fact that

$$w_{Y(G(v,\delta))}(a) \leq w_{Y(G(v,\delta))}(a)w_{Y(G)}(a)$$

for all $0 \leq a < v$, we get that this also holds for all $v \leq a < \infty$, which completes the proof. \square

PROOF OF THEOREM C.1. Assume that G has $m < n$ atoms. Applying Lemma C.3 successively $n - m$, each time feeding the resulting distribution back into the lemma, produces a new distribution G' that is discrete uniform with n atoms and satisfies $\varphi_{\infty, G}(0) \leq \varphi_{\infty, G'}(0)$. \square

D Deferred computational proofs

D.1 Computations for preliminaries

LEMMA 2.3 (SAMPLE COMPLEXITY FOR ε -MULTIPLICATIVE CLOSENESS, [53, LEMMA 1]). *Let F denote the CDF of a job size distribution on \mathbb{R}_+ . Let G be the empirical CDF based on n i.i.d. samples from F . Then for any $\varepsilon \in (0, 0.6)$, we have*

$$\sup_{0 \leq x \leq \ell} \frac{\bar{G}(x)}{\bar{F}(x)} \in [e^{-\varepsilon}, e^{\varepsilon}] \quad \text{with probability at least } 1 - 2 \exp\left(\frac{-n\varepsilon^2 \bar{F}(\ell)}{3}\right).$$

PROOF. The result is a corollary of Wellner [53, Lemma 1], rephrased for a general F . Define, for $u \in (0, 1]$,

$$\bar{F}^{\leftarrow}(u) := \min\{y \geq 0 \mid \bar{F}(y) \leq u\}.$$

Since \bar{F} is nonincreasing, right continuous, and $\lim_{y \rightarrow \infty} \bar{F}(y) = 0$, the set over which we are taking the minimum is nonempty and closed for every $u > 0$, meaning the minimum exists.

Let $U_i \sim \text{Unif}(0, 1)$ i.i.d. Let $X_i = \bar{F}^{\leftarrow}(U_i)$. Then X_i are equivalently n i.i.d. samples from a distribution with CDF F .

Define H as the empirical CDF of U_i , and define H^- as its left-continuous version, meaning

$$H^-(u) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{U_i < u\}.$$

Notice that \bar{F}^{\leftarrow} is an antitone Galois connection [?]. Therefore, $\bar{F}^{\leftarrow}(u) > x$ if and only if $u < \bar{F}(x)$ for $x \in [0, \infty)$ and $u \in (0, 1]$. Therefore, it follows that the event that $X_i > x$ is equivalent to the event that $\bar{F}^{\leftarrow}(U_i) > x$, which is in turn equivalent to the event that $\bar{F}(x) > U_i$. Therefore, $\bar{G}(x) = H^-(\bar{F}(x))$.

Let us now fix some $\ell \geq 0$, and define $b := \bar{F}(\ell)$. Then $\{x \leq \ell\} \subseteq \{\bar{F}(x) \geq \bar{F}(\ell)\} = \{\bar{F}(x) \geq b\}$. Using $\bar{G}(x) = H^-(\bar{F}(x))$, it follows that

$$\sup_{0 \leq x \leq \ell} \frac{\bar{G}(x)}{\bar{F}(x)} = \sup_{0 \leq x \leq \ell} \frac{H^-(\bar{F}(x))}{\bar{F}(x)} \leq \sup_{b \leq t \leq 1} \frac{H^-(t)}{t}; \quad (\text{D.1})$$

and likewise,

$$\inf_{0 \leq x \leq \ell} \frac{\bar{G}(x)}{\bar{F}(x)} = \inf_{0 \leq x \leq \ell} \frac{H^-(\bar{F}(x))}{\bar{F}(x)} \geq \inf_{b \leq t \leq 1} \frac{H^-(t)}{t}. \quad (\text{D.2})$$

Because $H^-(t) \leq H(t)$, it follows that

$$\left\{ \sup_{b \leq t \leq 1} \frac{H^-(t)}{t} > \eta \right\} \subseteq \left\{ \sup_{b \leq t \leq 1} \frac{H(t)}{t} > \eta \right\}.$$

The empirical CDF H is a step function with finitely many jumps, and $H(t)/t$ is decreasing on each interval between jumps. As a result, the infimum of $H(t)/t$ on $[b, 1]$ is at either $t = 1$ or at the left limit of jump points (which would equal that of H^-), meaning that

$$\left\{ \inf_{b \leq t \leq 1} \frac{H^-(t)}{t} < \eta \right\} = \left\{ \inf_{b \leq t \leq 1} \frac{H(t)}{t} < \eta \right\}.$$

Wellner [53, Lemma 1] states that for $0 < a < 1$ and $\lambda > 1$,

$$\begin{aligned} \mathbb{P} \left[\sup_{a \leq t \leq 1} \frac{H(t)}{t} \geq \lambda \right] &\leq \exp(-nah(\lambda)), \\ \mathbb{P} \left[\inf_{a \leq t \leq 1} \frac{H(t)}{t} \leq \frac{1}{\lambda} \right] &\leq \exp\left(-nah\left(\frac{1}{\lambda}\right)\right), \end{aligned}$$

where $h(u) = u(\log u - 1) + 1$. Setting $a = b$, $\lambda = e^\varepsilon$, using the union bound, and using (D.1) and (D.2) gives

$$\mathbb{P} \left[\sup_{0 \leq x \leq \ell} \frac{\bar{G}(x)}{\bar{F}(x)} \notin [e^{-\varepsilon}, e^\varepsilon] \right] \leq \exp(-nbh(e^{-\varepsilon})) + \exp(-nbh(e^\varepsilon)).$$

For $0 < \varepsilon < 0.6$ we can safely bound both $h(e^{-\varepsilon}), h(e^\varepsilon) \geq \frac{\varepsilon^2}{3}$. This means

$$\mathbb{P} \left[\sup_{0 \leq x \leq \ell} \frac{\bar{G}(x)}{\bar{F}(x)} \notin [e^{-\varepsilon}, e^\varepsilon] \right] \leq 2 \exp\left(\frac{-n\bar{F}(\ell)\varepsilon^2}{3}\right). \quad \square$$

D.2 Computations for truncated empirical Gittins

LEMMA 4.3. *Let G be a job size distribution. Then the ℓ -truncated Gittins policy $\gamma(G_\ell)$, satisfies*

$$\mathbb{E}_{\gamma(G_\ell)}[N(\text{size} \geq \ell)] \leq \lambda \bar{F}(\ell) \left(\frac{\lambda \mathbb{E}[(S \wedge \ell)^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S | S \geq \ell]}{(1-\rho)} \right).$$

PROOF. By Little's law and Poisson splitting,

$$\mathbb{E}_\pi[N(\text{size}) \geq \ell] = \lambda \bar{F}(\ell) \mathbb{E}_\pi[T | S \geq \ell].$$

Thus, our goal will be to upper bound $\mathbb{E}_{\gamma(G_\ell)}[T]$. We claim

$$\mathbb{E}_{\gamma(G_\ell)}[T | S \geq \ell] \leq \left(\frac{\mathbb{E}[W_{<\ell}] + \mathbb{E}[S | S \geq \ell]}{1-\rho} \right),$$

where $W_{<\ell}$ denotes the stationary distribution of work in a system with job size distribution F_ℓ .

To prove this bound, imagine a tagged job with size $\geq \ell$ enters the system. Under any ℓ -truncated Gittins policy, prior to completion, the tagged job must wait for every job currently in the system with age $< \ell$ to be served up to age ℓ (or completion). In fact, the tagged job must wait for the busy period started by this work, because starting at age ℓ , all jobs are served in PLCFS order and only after the system contains no jobs of age $< \ell$. Furthermore, for the same reason, the tagged job will also need to wait for the busy period started by its own size. The first term is exactly $\frac{\mathbb{E}[W_{<\ell}]}{1-\rho}$ in

expectation, and the second term is $\frac{\mathbb{E}[S|S \geq \ell]}{1-\rho}$ in expectation. Using standard results for the expected stationary work in an M/G/1 queue [18], we have,

$$\mathbb{E}_{Y(G_r)}[T | S \geq \ell] \leq \frac{\lambda \mathbb{E}[(S \wedge \ell)^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S | S \geq \ell]}{1-\rho}. \quad \square$$

D.3 Computations for empirical Gittins

LEMMA 5.4. *Let G and H be job size distributions and let $\ell > 0$. Then for all ranks r and ages $a < \ell$,*

$$\int_a^{b_{Y(G)}(a,r) \wedge \ell} \frac{\bar{H}(t)}{\bar{H}(a)} dt \leq s_{H,G}(a,r) \leq \int_a^{b_{Y(G)}(a,r) \wedge \ell} \frac{\bar{H}(t)}{\bar{H}(a)} dt + \frac{\bar{H}(\ell)}{\bar{H}(a)} s_{H,G}(\ell,r).$$

PROOF. For any age $a < \ell$, we can split the service until $b_{Y(G)}(a,r)$ into service until $b_{Y(G)}(a,r) \wedge \ell$ and any remaining service:

$$\begin{aligned} s_{H,G}(a,r) &= \mathbb{E}_{S \sim H}[S \wedge b_{Y(G)}(a,r) - a | S > a] \\ &= \mathbb{E}_{S \sim H}[S \wedge b_{Y(G)}(a,r) \wedge \ell - a | S > a] + \mathbb{E}_{S \sim H}[(S \wedge b_{Y(G)}(a,r) - \ell)^+ | S > a] \\ &= \int_a^{b_{Y(G)}(a,r) \wedge \ell} \frac{\bar{H}(t)}{\bar{H}(a)} dt + \mathbb{E}_{S \sim H}[(S \wedge b_{Y(G)}(a,r) - \ell)^+ | S > a], \end{aligned}$$

where for any x we define $x^+ = \max(x, 0)$ as the positive part of x . This immediately gives us the lower bound. For the upper bound, note that $b_{Y(G)}(\ell,r) \geq b_{Y(G)}(a,r)$ and then observe that,

$$\begin{aligned} \mathbb{E}_{S \sim H}[(S \wedge b_{Y(G)}(a,r) - \ell)^+ | S > a] &= \frac{1}{\bar{H}(a)} \mathbb{E}_{S \sim H}[\mathbb{1}(S > a)(S \wedge b_{Y(G)}(a,r) - \ell)^+] \\ &= \frac{1}{\bar{H}(a)} \mathbb{E}_{S \sim H}[\mathbb{1}(S > \ell)(S \wedge b_{Y(G)}(a,r) - \ell)^+] \\ &\leq \frac{1}{\bar{H}(a)} \mathbb{E}_{S \sim H}[\mathbb{1}(S > \ell)(S \wedge b_{Y(G)}(\ell,r) - \ell)^+] \\ &= \frac{\bar{H}(\ell)}{\bar{H}(a)} \mathbb{E}_{S \sim H}[(S \wedge b_{Y(G)}(\ell,r) - \ell)^+ | S > \ell] \\ &= \frac{\bar{H}(\ell)}{\bar{H}(a)} s_{H,G}(\ell,r). \quad \square \end{aligned}$$

LEMMA D.1. *Let G be an empirical job size distribution generated by n samples, and suppose it is ε -multiplicatively close to F up to $\ell > 0$. Then for all $0 < k < \ell$ such that $e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)} > 0$,*

$$\mathbb{E}_{Y(G)}[N] \leq \frac{e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n}{e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)}} (\mathbb{E}_{Y(F)}[N] + h_n \mathbb{E}_{Y(F)}[N(\text{age} \geq k)]) + \mathbb{E}_{Y(G)}[N(\text{age} \geq k)].$$

PROOF. First, by Corollary 3.11, $\mathbb{E}_{Y(G)}[\hat{N}_G] \leq \mathbb{E}_\pi[\hat{N}_G]$ for all policies π . Then letting $\pi = \gamma(F)$ and using Theorem 5.5,

$$\begin{aligned} \left(e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)} \right) \mathbb{E}_{Y(G)}[N(\text{age} < k)] &\leq \mathbb{E}_{Y(G)}[\hat{N}_G] \leq \mathbb{E}_{Y(F)}[\hat{N}_G] \\ &\leq \left(e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n \right) \mathbb{E}_{Y(F)}[N(\text{age} < k)] + h_n \mathbb{E}_{Y(F)}[N(\text{age} \geq k)]. \end{aligned}$$

Thus, we have

$$\begin{aligned} \mathbb{E}_{Y(G)}[N] &\leq \frac{c_2}{c_1} \mathbb{E}_{Y(F)}[N(\text{age} < k)] + \frac{h_n}{c_1} \mathbb{E}_{Y(F)}[N(\text{age} \geq k)] + \mathbb{E}_{Y(G)}[N(\text{age} \geq k)] \\ &\leq \frac{c_2}{c_1} \left(\mathbb{E}_{Y(F)}[N] + \frac{h_n - c_2}{c_2} \mathbb{E}_{Y(F)}[N(\text{age} \geq k)] \right) + \mathbb{E}_{Y(G)}[N(\text{age} \geq k)] \\ &\leq \frac{c_2}{c_1} (\mathbb{E}_{Y(F)}[N] + h_n \mathbb{E}_{Y(F)}[N(\text{age} \geq k)]) + \mathbb{E}_{Y(G)}[N(\text{age} \geq k)] \end{aligned}$$

where $c_1 = e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)}$, and $c_2 = e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n$. \square

LEMMA D.2. For all policies π ,

$$\mathbb{E}_\pi[N(\text{size} \geq k)] \leq \lambda \bar{F}(k) \left(\frac{\lambda \mathbb{E}[S^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S | S \geq k]}{1-\rho} \right).$$

PROOF SKETCH. The proof is essentially the same as that of Lemma 4.3, presented in Appendix D.2, if one lets $\ell \rightarrow \infty$ (i.e. all other work gets completed before the tagged job). We omit the details as the necessary changes are straightforward. \square

THEOREM 5.1. Let G be an empirical distribution constructed from n samples from F . Let $\varepsilon \in (0, 0.6)$ and $\delta \in (0, 1)$. If $n \geq \frac{3 \log(2/\delta)}{\bar{F}(\ell) \varepsilon^2}$, and $e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)} > 0$, then with probability at least $1 - \delta$,

$$\mathbb{E}_{Y(G)}[T] \leq \inf_{\ell > k > 0} \frac{e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n}{e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)}} \left(\mathbb{E}_{Y(F)}[T] + \frac{\bar{F}(k)(1+h_n)}{(1-\rho)} \left(\frac{\lambda \mathbb{E}[S^2]}{2(1-\rho)} + \mathbb{E}[S | S > k] \right) \right),$$

where $h_n = \sum_{i=1}^n \frac{1}{i}$ is the n -th harmonic number.

PROOF. Lemma 2.3 implies that F and G are ε -multiplicatively close up ℓ with probability at least $1 - \delta$. Replacing both relevant terms in Lemma D.1 with the bound from Lemma D.2 and then applying a straightforward upper bound,

$$\mathbb{E}_{Y(G)}[N] \leq \frac{e^{2\varepsilon} + \frac{\bar{F}(\ell)}{\bar{F}(k)} h_n}{e^{-2\varepsilon} - \frac{\bar{F}(\ell)}{\bar{F}(k)}} \left(\mathbb{E}_{Y(F)}[N] + (1+h_n) \lambda \bar{F}(k) \left(\frac{\lambda \mathbb{E}[S^2]}{2(1-\rho)^2} + \frac{\mathbb{E}[S | S \geq k]}{1-\rho} \right) \right).$$

The response time bound then follows from another application of Little's law. \square

D.4 Computations for asymptotic analysis

LEMMA D.3. Let $q > 0$ be fixed and $n, r > 0$.

(a) For fixed $\alpha > 2$, define

$$f(\zeta) := (n\zeta)^{-q} + r^{\frac{\alpha}{\alpha-1}} \zeta + r^{\frac{1}{\alpha-1}} \zeta^{\frac{\alpha-1}{\alpha}}.$$

Then, as $n, r \rightarrow \infty$, the value of ζ that minimizes $f(\zeta)$ up to constant factors satisfies

$$\zeta^* = \Theta \left(\min \left\{ n^{-\frac{q}{1+q}} r^{-\frac{\alpha}{(1+q)(\alpha-1)}}, n^{-\frac{\alpha q}{\alpha q + \alpha - 1}} r^{-\frac{\alpha}{(\alpha-1)(\alpha q + \alpha - 1)}} \right\} \right),$$

and the corresponding minimal value of $f(\zeta)$, up to constant factors, scales as

$$\min_{\zeta > 0} f(\zeta) = \Theta \left(\max \left\{ n^{-\frac{q}{1+q}} r^{-\frac{\alpha}{(1+q)(\alpha-1)}}, n^{-\frac{q(\alpha-1)}{\alpha q + \alpha - 1}} r^{-\frac{\alpha}{(\alpha-1)(\alpha q + \alpha - 1)}} \right\} \right).$$

(b) For $\alpha \in (1, 2]$, define

$$g(\zeta) := (n\zeta)^{-q} + (r\zeta^{\frac{\alpha-1}{\alpha}})^2 + r\zeta^{\frac{\alpha-1}{\alpha}}.$$

Then, as $n, r \rightarrow \infty$, the value of ζ that minimizes $g(\zeta)$ up to constant factors satisfies

$$\zeta^* = \Theta \left(\min \left\{ n^{\frac{-q}{q+\frac{2(\alpha-1)}{\alpha}}} r^{\frac{-2}{q+\frac{2(\alpha-1)}{\alpha}}}, n^{\frac{-q}{q+\frac{\alpha-1}{\alpha}}} r^{\frac{-1}{q+\frac{\alpha-1}{\alpha}}} \right\} \right),$$

and the corresponding minimal value of $g(\zeta)$, up to constant factors, scales as

$$\min_{\zeta>0} g(\zeta) = \Theta \left(\max \left\{ n^{\frac{-q}{q+\frac{2(\alpha-1)}{\alpha}}} r^{\frac{2q}{q+\frac{2(\alpha-1)}{\alpha}}}, n^{\frac{-q}{q+\frac{\alpha-1}{\alpha}}} r^{\frac{q}{q+\frac{\alpha-1}{\alpha}}} \right\} \right).$$

PROOF. We derive the scaling of the minimizer by balancing dominant terms. The notation $f(x) \asymp g(x)$ means that $f(x)$ and $g(x)$ are of the same order, namely that there exist positive constants a, b such that $a g(x) \leq f(x) \leq b g(x)$ for all sufficiently large x .

(a) Let $q > 0, \alpha > 2$, and

$$f(\zeta) = (n\zeta)^{-q} + r^{\frac{\alpha}{\alpha-1}} \zeta + r^{\frac{1}{\alpha-1}} \zeta^{\frac{\alpha-1}{\alpha}}.$$

The first term decreases in ζ , while the last two increase in ζ . Consequently, $f(\zeta)$ has a single minimum, and its minimizer occurs when the decreasing term $(n\zeta)^{-q}$ is of the same order as one of the increasing terms. Among the two possible values of ζ obtained by balancing the first term with either the second or the third, the smaller ζ yields the smaller function value, since all terms are positive and increase for larger ζ at the minimum.

Setting the first and second terms equal to each other gives

$$\zeta_1 = n^{-\frac{q}{1+q}} r^{-\frac{\alpha}{(1+q)(\alpha-1)}}.$$

Setting the first and third terms equal to each other gives

$$\zeta_2 = n^{-\frac{\alpha q}{\alpha q + \alpha - 1}} r^{-\frac{\alpha}{(\alpha-1)(\alpha q + \alpha - 1)}}.$$

The smaller ζ achieves the asymptotic minimum:

$$\zeta^* = \Theta(\min\{\zeta_1, \zeta_2\}).$$

The two terms set equal to each other will dominate $f(\zeta)$, giving

$$\min_{\zeta>0} f(\zeta) = \Theta(\max\{f(\zeta_1), f(\zeta_2)\}),$$

where

$$f(\zeta_1) \asymp (n\zeta_1)^{-q} = n^{-\frac{q}{1+q}} r^{\frac{\alpha}{(1+q)(\alpha-1)}},$$

$$f(\zeta_2) \asymp (n\zeta_2)^{-q} = n^{-\frac{q(\alpha-1)}{\alpha q + \alpha - 1}} r^{\frac{\alpha}{(\alpha-1)(\alpha q + \alpha - 1)}}.$$

(b) The reasoning for this part closely follows that of part (a). Let $q > 0, \alpha \in (1, 2]$, and

$$g(\zeta) = (n\zeta)^{-q} + r^2 \zeta^{\frac{2(\alpha-1)}{\alpha}} + r\zeta^{\frac{\alpha-1}{\alpha}}.$$

Setting the first and second terms equal to each other gives

$$\zeta_1 = n^{-\frac{q}{q+\frac{2(\alpha-1)}{\alpha}}} r^{-\frac{2}{q+\frac{2(\alpha-1)}{\alpha}}}.$$

Setting the first and third terms equal to each other gives

$$\zeta_2 = n^{-\frac{q}{q+\frac{\alpha-1}{\alpha}}} r^{-\frac{1}{q+\frac{\alpha-1}{\alpha}}}.$$

Hence

$$\zeta^* = \Theta(\min\{\zeta_1, \zeta_2\}),$$

and

$$\min_{\zeta > 0} g(\zeta) = \Theta(\max\{g(\zeta_1), g(\zeta_2)\}),$$

where

$$g(\zeta_1) \asymp (n\zeta_1)^{-q} = n^{-\frac{q}{q+\frac{2(\alpha-1)}{\alpha}}} r^{\frac{2q}{q+\frac{2(\alpha-1)}{\alpha}}},$$

$$g(\zeta_2) \asymp (n\zeta_2)^{-q} = n^{-\frac{q}{q+\frac{\alpha-1}{\alpha}}} r^{\frac{q}{q+\frac{\alpha-1}{\alpha}}}. \quad \square$$

E Challenges to simulating empirical Gittins with infinite-variance true distributions

We omit the infinite-variance case from our simulations in Section 7 because simulating empirical Gittins with samples from infinite-variance F requires solving some nontrivial software engineering challenges, which are outside the scope of the simulator we created for this work. For example, our simulator is currently completely discrete but with very small discretization increments, which is necessary for (approximately) computing and simulating the true Gittins policy for continuous distributions. But the empirical distributions sampled from infinite-variance F tend to have a few samples at very large values, which makes using small discretization increments impractical. One could possibly overcome this with adaptive discretization, but this would involve substantial engineering effort. Another possibility is using discrete event simulation for empirical Gittins, but this makes it hard to get an apples-to-apples comparison with true Gittins. A final idea is to directly compute the SOAP mean response time formulas, but there is no publicly available software package for doing so.