



**HAL**  
open science

# Density Estimators of the Cumulative Reward up to a Hitting Time to a Rarely Visited Set of a Regenerative System

Marvin K Nakayama, Bruno Tuffin

► **To cite this version:**

Marvin K Nakayama, Bruno Tuffin. Density Estimators of the Cumulative Reward up to a Hitting Time to a Rarely Visited Set of a Regenerative System. WSC 2022 - Winter Simulation Conference, Dec 2022, Singapour, Singapore. pp.1-18. hal-03709329

**HAL Id: hal-03709329**

**<https://inria.hal.science/hal-03709329>**

Submitted on 29 Jun 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Density Estimators of the Cumulative Reward up to a Hitting Time to a Rarely Visited Set of a Regenerative System

Marvin K. Nakayama  
Computer Science Department  
New Jersey Institute of Technology  
Newark, NJ 07102, USA

Bruno Tuffin  
Inria, Univ Rennes, CNRS, IRISA  
Campus de Beaulieu, 263 Avenue Général Leclerc  
35042 Rennes, FRANCE

June 29, 2022

## Abstract

For a regenerative process, we propose various estimators of the density function of the cumulative reward up to hitting a rarely visited set of states. The approaches exploit existing weak-convergence results for the hitting-time distribution, and we apply simulation (often with previously developed importance samplers for estimating the mean) to estimate parameters of the limiting distribution. We also combine these ideas with kernel methods. Numerical results from simulation experiments show the effectiveness of the estimators.

## 1 Introduction

A *regenerative* stochastic process on a state space  $\mathcal{S}$  possesses an increasing sequence of time points at which the process “probabilistically restarts” [9]. An example is the queue-length process (including any customer in service) of a stable GI/GI/1 queue, with the starts of busy periods being regeneration times [9, Example 1.2.2]. Suppose that there is a set  $\mathcal{A} \subset \mathcal{S}$  of rarely visited states (e.g., large queue lengths in a stable queue), and for a given reward function on the state space

$\mathcal{S}$ , consider the cumulative reward  $R$  up to first hitting  $\mathcal{A}$ . Our goal is to estimate the density function  $f$  of  $R$ .

One reason for estimating  $f$  is that an analyst may glean important features of the distribution of  $R$  from a graph of its density, arguably more easily than from a plot of its *cumulative distribution function* (CDF)  $F$ . Another motivation arises in quantile estimation [4]. A central limit theorem for the estimator of the  $q$ -quantile  $\xi = F^{-1}(q)$  for  $0 < q < 1$  has an asymptotic variance that includes  $f(\xi)$ , so constructing a confidence interval for  $\xi$  often entails estimating  $f(\xi)$ .

When sampling independent and identically distributed (i.i.d.) copies of  $R$  from  $F$  via crude simulation (CS), we can estimate the density  $f$  through kernel methods [17], a class of nonparametric techniques for function estimation. But since  $\mathcal{A}$  is rarely visited, generating an observation of  $R$  from  $F$  can be computationally costly as it entails simulating the process for a typically large number of transitions. Rare-event simulation techniques, such as importance sampling (IS), have exploited regenerative structure to obtain efficient estimators of the mean  $\mathbb{E}[R]$  (e.g., [5] and Chapter VI of [1]). Rather than replicating i.i.d. copies of  $R$  from  $F$ , the IS methods instead sample cycles from a different distribution than the original, unbiasing results by multiplying by a correction factor, the likelihood ratio. Compared to their CS counterparts, the IS estimators of  $\mathbb{E}[R]$  can have much smaller variance, with also substantially less computation cost. But ordinary kernel methods are then no longer applicable to estimate  $f$  because i.i.d. replicates of  $R$  from  $F$  are not available.

We now consider estimating  $f$  by employing ideas from [6], which take advantage of known weak-convergence results based on regenerative properties in an asymptotic regime in which visiting the set  $\mathcal{A}$  becomes rarer. These limit theorems [10] generalizes a classical result of Rényi [10, p. 3] establishing that for a *geometric sum* (i.e., the sum of a geometrically distributed number of i.i.d. nonnegative random variables with finite mean), the ratio of the sum and its mean converges weakly to an exponential as the geometric's parameter  $p$  shrinks to 0. Then simulation is applied to estimate the limiting distribution's parameters, where those related to rare events are dealt with via existing IS methods for  $\mathbb{E}[R]$ . Applying this approach, [6] develops so-called exponential and convolution estimators for the CDF  $F$  of  $R$ , along with its  $q$ -quantile  $\xi = F^{-1}(q)$  and the conditional tail expectation  $\gamma = \mathbb{E}[R \mid R > \xi]$ , the latter two being frequently employed risk measures in finance, where  $\xi$  (resp.,  $\gamma$ ) is known as the value-at-risk (resp., conditional value-at-risk); see, e.g., [8]. Our current paper extends the methods of [6] to handle the density  $f$ , and also further combines the approaches with IS kernel techniques [11].

The rest of the paper proceeds as follows. Section 2 reviews (ordinary) kernel density estimation. We then describe the assumed regenerative structure in Section 3, and Section 4 explains the asymptotic regime for the weak-convergence

results upon which our density estimators are based. Sections 5 and 6 extend the exponential and convolution estimators, respectively, of [6] to instead estimate densities. We then combine the convolution estimator with kernels in Section 7. Section 8 presents numerical results, and we give concluding remarks in Section 9.

## 2 Review of kernel density estimation

Consider a stochastic process  $X = [X(t) : t \geq 0]$  evolving on a state space  $\mathcal{S} \subset \mathfrak{R}^d$ . Let  $T = \inf\{t : X(t) \in \mathcal{A}\}$  be the hitting time (or first passage time) to a set  $\mathcal{A} \subset \mathcal{S}$ , and let  $R = \int_0^T r(X(t)) dt$  be the cumulative reward up to  $T$  for a reward function  $r : \mathcal{S} \rightarrow \mathfrak{R}_+$ , where  $\mathfrak{R}_+$  is the set of nonnegative real numbers. Let  $F$  be the CDF of  $R$ ; i.e.,  $F(x) = \mathbb{P}(R \leq x)$ ,  $x \in \mathfrak{R}$ . Under the assumption that  $F$  has a density  $f$  (with respect to Lebesgue measure), the goal is to estimate  $f$ .

We first review estimating  $f$  via kernel methods. Suppose that  $R_1, R_2, \dots, R_n$  are  $n$  i.i.d. observations from CDF  $F$ . We can estimate  $F$  through the *empirical CDF*  $\widehat{F}_n$  defined by

$$\widehat{F}_n(x) = \frac{1}{n} \sum_{i=1}^n \mathcal{I}(R_i \leq x), \quad (1)$$

where  $\mathcal{I}(\cdot)$  is the indicator function, which takes value 1 (resp., 0) when its argument is true (resp., false). Thus,  $\widehat{F}_n$  assigns a mass of size  $1/n$  to each observed  $R_i$ , and  $\widehat{F}_n(x)$  is the fraction of the  $n$  data points that are at most  $x$ . As the derivative of  $F$ , the density  $f$  satisfies  $f(x) = \lim_{\delta \rightarrow 0} [F(x + \delta) - F(x - \delta)] / (2\delta)$ , which suggests estimating  $f$  by  $\widehat{f}_{n,\delta}$  with

$$\widehat{f}_{n,\delta}(x) = \frac{\widehat{F}_n(x + \delta) - \widehat{F}_n(x - \delta)}{2\delta} \quad (2)$$

for a small but fixed  $\delta > 0$  known alternatively as the *bandwidth*, *window width*, or *smoothing parameter*. Sometimes called the *naive density estimator* [17, Section 2.3],  $\widehat{f}_{n,\delta}(x)$  is a (central) finite difference; e.g., see Section VII.1 of [1]. We next give an equivalent representation of  $\widehat{f}_{n,\delta}$  in terms of a uniform density function to motivate the kernel estimators.

Let  $k_{\mathcal{U}[-1,1]}$  be the density function of a (continuous) uniform distribution on the interval  $[-1, 1)$ , so  $k_{\mathcal{U}[-1,1]}(u) = \mathcal{I}(-1 \leq u < 1) / 2$ . Using (1) allows us to rewrite  $\widehat{f}_{n,\delta}(x)$  in (2) as

$$\widehat{f}_{n,\delta}(x) = \frac{1}{2n\delta} \sum_{i=1}^n \mathcal{I}(x - \delta < R_i \leq x + \delta) = \frac{1}{n} \sum_{i=1}^n \frac{1}{\delta} k_{\mathcal{U}[-1,1]} \left( \frac{x - R_i}{\delta} \right). \quad (3)$$

We can obtain other estimators of  $f$  by replacing  $k_{\mathcal{U}[-1,1]}$  in (3) with any density or more generally a *kernel*  $k : \mathfrak{X} \rightarrow \mathfrak{R}$ , which is a function with  $\int_{\mathfrak{X}} k(u) du = 1$ ; e.g., see [17, Section 2.4]. This then leads to the *kernel density estimator*  $\widehat{f}_{n,k,\delta}$  (with kernel  $k$ ) defined by

$$\widehat{f}_{n,k,\delta}(x) = \frac{1}{n} \sum_{i=1}^n \frac{1}{\delta} k\left(\frac{x-R_i}{\delta}\right) = \frac{1}{n} \sum_{i=1}^n k_{\delta}(x-R_i), \quad (4)$$

where  $k_{\delta}(u) = (1/\delta)k(u/\delta)$  is the scaled kernel, and  $\delta > 0$  is the bandwidth. When  $k$  is a density, the mean (resp., variance) of  $k_{\delta}$  equals that of  $k$  multiplied by  $\delta$  (resp.,  $\delta^2$ ), and the kernel estimator (4) is a mixture of  $n$  scaled densities  $k_{\delta}$ , each translated by an  $R_i$  and with equal mixture weight  $1/n$ . In practice,  $\delta = \delta_n$  is usually chosen as a function of the sample size  $n$  such that  $\delta_n \rightarrow 0$  and  $n\delta_n \rightarrow \infty$  as  $n \rightarrow \infty$ . Kernel density estimators are generally biased, with the bias shrinking to 0 as  $\delta_n \rightarrow 0$ . (Chapter 5 of [17] also discusses adaptive estimators of  $f$  with bandwidth depending on  $x$  or the distances between the data points, but for simplicity, we do not consider these variants.)

The kernel is often taken as a density function that is symmetric (about the origin). Although these properties are not required generally [17, Section 3.6.1], we will simplify our discussion by assuming that  $k$  is a density (i.e.,  $k(u) \geq 0$  for all  $u \in \mathfrak{X}$  and  $\int_{\mathfrak{X}} k(u) du = 1$ ) but not necessarily symmetric. In addition to  $k_{\mathcal{U}[-1,1]}(u)$ , other common choices for a symmetric kernel  $k$  include the Gaussian kernel  $k_{\mathcal{N}}(u) = \phi(u) = (2\pi)^{-1/2}e^{-u^2/2}$  (with CDF  $\Phi$ ), and the Epanechnikov kernel [17, p. 42].

Using a symmetric kernel  $k$  in (4) can lead to  $\widehat{f}_{n,k,\delta}$  assigning positive mass to sets of negative values (i.e.,  $\int_{-\infty}^0 \widehat{f}_{n,k,\delta}(x) dx > 0$ ) when the bandwidth  $\delta$  is sufficiently large. This may be undesirable when the true density  $f$  of  $R$  has  $f(x) = 0$  for all  $x < 0$ , as is the case under our assumption that the reward function  $r(\cdot) \geq 0$ . One possible way to address this is to truncate  $\widehat{f}_{n,k,\delta}$  so that  $\widehat{f}_{n,k,\delta}(x) = 0$  for  $x < 0$ , but then  $\widehat{f}_{n,k,\delta}$  may not be a density as it may integrate to less than 1; [17], Section 2.10, discusses other issues also arising from truncation. [2] suggests instead choosing what we call a *positive-support kernel*, i.e., a density function  $k$  whose support is contained in  $[0, \infty)$ , such as the uniform $[0, 2)$  kernel  $k_{\mathcal{U}[0,2)}(u) = \mathcal{I}(0 \leq u < 2)/2$  and the exponential kernel  $k_{\mathcal{E}}(u) = e^{-u} \mathcal{I}(u \geq 0)$ . When  $r(\cdot) \geq 0$ , positive-support kernels always result in  $\int_0^{\infty} \widehat{f}_{n,k,\delta}(x) dx = 1$ , so  $\widehat{f}_{n,k,\delta}$  is a density. But the lack of symmetry of  $k$  can lead to a less statistically efficient estimator of  $f$ . For a positive-support kernel  $k$  and sample size  $n$ , [2] derives the bandwidth  $\delta = \delta_n$  (also depending on  $k$  and  $f$ ) that asymptotically (as  $n \rightarrow \infty$ ) minimizes the mean integrated squared error (MISE)  $\mathbb{E}[\int (\widehat{f}_{n,k,\delta}(x) - f(x))^2 dx]$  of the kernel density estimator. The optimal  $\delta_n$  shrinks at a rate of order  $n^{-1/3}$  as  $n \rightarrow \infty$ , leading to the MISE

optimally decreasing at a rate of order  $n^{-2/3}$ , worse than the order  $n^{-4/5}$  optimal MISE rate for symmetric kernels with optimal bandwidth of order  $n^{-1/5}$ ; e.g., see Section 3.3.2 of [17].

### 3 Regenerative process

A drawback of the kernel estimators of Section 2 is that generating i.i.d. copies of  $R$  from  $F$  can be computationally expensive when  $\mathcal{A}$  is rarely visited. To obtain more efficient methods for estimating  $f$ , we will now impose additional structure on the process  $X = [X(t) : t \geq 0]$ , which will enable us to adapt approaches from [6]. Specifically, we first approximate the CDF  $F$  of  $R$  using existing weak-convergence results [10] in an asymptotic regime where visits to the set  $\mathcal{A}$  become rarer (in the sense of (7) below). We then apply simulation to estimate the limiting distribution's unknown parameters, some of which relate to rare events, and we will exploit previously developed IS methods designed to efficiently estimate  $\mathbb{E}[R]$  to handle such parameters.

To accomplish this, we now assume that  $X$  is a (non-delayed) regenerative process [9]. Thus,  $X$  has a sequence of regeneration times  $0 = \Gamma_0 < \Gamma_1 < \Gamma_2 < \dots$ , so the process “probabilistically restarts” at each  $\Gamma_i$ . An example of a regenerative process is a positive-recurrent continuous-time Markov chain (CTMC) on a discrete state space  $\mathcal{S}$  with a fixed starting state  $x_0 \in \mathcal{S}$ , and successive entry times to  $x_0$  form a regeneration sequence [9, Example 2.1 on p. 15].

For each  $i \geq 1$ , define  $\tau_i = \Gamma_i - \Gamma_{i-1}$ , and  $[X(\Gamma_{i-1} + s) : 0 \leq s < \tau_i]$  is called the  $i$ th (regenerative) cycle, which has length  $\tau_i$ . Also, define  $Y_i = \int_{\Gamma_{i-1}}^{\Gamma_i} r(X(t)) dt$  as the reward accrued during cycle  $i$ . As  $X$  is regenerative,  $(\tau_i, [r(X(\Gamma_{i-1} + s)) : 0 \leq s < \tau_i])$ ,  $i \geq 1$ , is a sequence of i.i.d. pairs of cycle lengths and reward processes during cycles. For  $i \geq 1$ , let  $T_i = \inf\{t \geq 0 : X(\Gamma_{i-1} + t) \in \mathcal{A}\}$  be the time elapsing after  $\Gamma_{i-1}$  until the next hit to  $\mathcal{A}$ . For  $x, y \in \mathfrak{R}$ , let  $x \wedge y = \min(x, y)$  and  $x \vee y = \max(x, y)$ . For  $i \geq 1$ , let  $D_i = \int_{\Gamma_{i-1}}^{T_i \wedge \tau_i} r(X(t)) dt$  be the reward accrued during the  $i$ th cycle up to hitting  $\mathcal{A}$  or the end of the cycle, whichever occurs first. The regenerative property ensures that  $(\tau_i, Y_i, D_i, \mathcal{I}(T_i < \tau_i))$ ,  $i = 1, 2, \dots$ , are i.i.d. 4-tuples. To simplify notation, let  $(\tau, Y, D, \mathcal{I}(T < \tau)) = (\tau_1, Y_1, D_1, \mathcal{I}(T_1 < \tau_1))$ .

We can give a stochastically equivalent representation for the cumulative reward  $R$  up to time  $T$  in terms of independent quantities. To do this, let  $W_1, W_2, \dots$  be i.i.d., each with CDF  $G_W(x) = \mathbb{P}(Y \leq x \mid \tau < T)$ , so  $W_i$  is distributed as the reward in a cycle that does not hit  $\mathcal{A}$ . Independent of the  $W_i$ , further define  $M$  as a geometric random variable with parameter  $p = \mathbb{P}(T < \tau)$  (and support starting from 0), so  $\mathbb{P}(M = l) = (1 - p)^l p$  for  $l = 0, 1, 2, \dots$ . Independent of  $M$  and the  $W_i$ , additionally let  $V$  be a random variable with CDF  $H$  defined by  $H(v) = \mathbb{P}(D \leq v \mid T < \tau)$ , so

$V$  is distributed as the reward in a cycle up to hitting  $\mathcal{A}$  given that  $T < \tau$ . For  $G$  as the CDF of the geometric sum  $S = \sum_{i=1}^M W_i$ , the regenerative property of  $X$  implies

$$R \stackrel{\mathcal{D}}{=} S + V \sim F, \text{ with } S \sim G \text{ independent of } V \sim H, \quad (5)$$

where  $\stackrel{\mathcal{D}}{=}$  denotes equality in distribution, and  $\sim$  means “distributed as”. Thus, the independence of  $S$  and  $V$  ensures that  $F = G \star H$ , where  $\star$  is the convolution operator; i.e.,  $G \star H(x) = \int G(x-y) dH(y)$ .

If  $p > 0$ , then the expectation  $\mu$  of the cumulative reward  $R$  up to  $T$  satisfies

$$\mu = \frac{\mathbb{E}[D]}{p} \equiv \frac{\zeta}{p}, \quad (6)$$

where  $\zeta$  and  $p = \mathbb{E}[\mathcal{I}(T < \tau)]$  are means of “cycle-based random variables” (i.e., measurable with respect to the sigma-field of  $X$  up to  $\tau$ ); e.g., see [7] and [5].

## 4 Asymptotic regime

We will exploit existing generalizations [10] of a classical result of Rényi [10, p. 3] establishing that the ratio of a geometric sum and its mean converges weakly to an exponential as the geometric’s parameter  $p$  shrinks to 0. For a theoretical framework to accommodate this, we parameterize the problem by introducing a *rarity parameter*  $\varepsilon > 0$  and examine the behavior of  $R \equiv R_\varepsilon$  and  $S \equiv S_\varepsilon$  as  $\varepsilon \rightarrow 0$ , where we assume that

$$p \equiv p_\varepsilon \rightarrow 0 \text{ as } \varepsilon \rightarrow 0, \quad (7)$$

which is what we meant before by saying that the set  $\mathcal{A}$  is “rarely visited”. We next describe two stochastic models satisfying (7).

**Example 1** For a stable GI/GI/1 queue, let  $X(t)$  be the total number of customers in the system at time  $t \geq 0$ , where the first customer arrives at time  $t = 0$  to an empty system. Thus, the state space is  $\mathcal{S} = \{0, 1, 2, \dots\}$ , and the process  $X$  is regenerative, regenerations occurring when a customer arrives to an empty system [9, pp. 16–17]. Let  $\mathcal{A} = \{b_\varepsilon, b_\varepsilon + 1, b_\varepsilon + 2, \dots\}$  be the set of states in which there are at least  $b_\varepsilon \equiv \lfloor 1/\varepsilon \rfloor$  customers in the system, so for  $\varepsilon > 0$  small,  $T$  represents the first time when there is a large number  $b_\varepsilon$  of customers in the system. Theorem 1 of [15] shows that (7) then holds.

**Example 2** A *highly reliable Markovian system* (HRMS) comprises a fixed number of components, each with exponentially distributed failure times, and a finite number of repairpersons, who fix each failed component in an exponentially distributed time. Each state in the (discrete) state space  $\mathcal{S}$  keeps track of the set of

currently failed components, as well as any other necessary information, e.g., about queuing of failed components. At time  $t = 0$ , all components are operational, and the resulting process  $X$  on  $\mathcal{S}$  is a positive-recurrent CTMC, making it regenerative. The system fails when certain combinations of components are down, and let  $\mathcal{A}$  denote the set of those states, so  $T$  is the first time to system failure. The rarity of system failures results from having failure rates as positive powers of  $\varepsilon$  and repair rates as constants. Then [16] provides conditions guaranteeing (7).

In Example 1, rarity arises by letting the set  $\mathcal{A} = \mathcal{A}_\varepsilon$  recede from state 0 as  $\varepsilon$  shrinks, but the transition dynamics remain unchanged. In contrast, the set  $\mathcal{A}$  in Example 2 does not change as  $\varepsilon$  shrinks, but rarity instead comes from the failure rates becoming smaller with fixed repair rates. As we are now actually considering a family of models indexed by  $\varepsilon$ , we sometimes (but not always) add a subscript  $\varepsilon$  in our notation (e.g.,  $S = S_\varepsilon$ ,  $W_i = W_{\varepsilon,i}$ , or  $F = F_\varepsilon$ ) to emphasize the dependence on  $\varepsilon$ .

To try to apply Rényi's theorem for geometric sums in our regenerative setting, consider  $S_\varepsilon = \sum_{i=1}^{M_\varepsilon} W_{\varepsilon,i}$  in (5). Rényi's result covers the case when the summands' distribution remains fixed as  $p_\varepsilon$  shrinks, but our Examples 1 and 2 violate this assumption. [10] provides generalizations to address this when (7) holds. For example, the conditions in Theorem 3.2.5 of [10] ensure that the ratio of  $S_\varepsilon$  to its mean  $\eta_\varepsilon = \mathbb{E}_\varepsilon[S_\varepsilon]$  converges weakly to an exponential: for all  $y \in \mathfrak{R}$ ,

$$\mathbb{P}_\varepsilon(S_\varepsilon/\eta_\varepsilon \leq y) \rightarrow 1 - e^{-y^+} \quad \text{as } \varepsilon \rightarrow 0, \quad (8)$$

where  $x^+ = \max(x, 0)$ ,  $x \in \mathfrak{R}$ . If the cumulative reward  $R_\varepsilon = S_\varepsilon + V_\varepsilon$  in (5) further has only a “negligible” contribution from  $V_\varepsilon$  (e.g.,  $p_\varepsilon \mathbb{E}_\varepsilon[V_\varepsilon] \rightarrow 0$  as  $\varepsilon \rightarrow 0$  when  $\mathbb{E}_\varepsilon[W_{\varepsilon,i}] = 1$ ), then we also have

$$\mathbb{P}_\varepsilon(R_\varepsilon/\mu_\varepsilon \leq x) \rightarrow 1 - e^{-x^+} \quad \text{as } \varepsilon \rightarrow 0, \quad (9)$$

where  $\mu_\varepsilon = \mathbb{E}_\varepsilon[R_\varepsilon]$ ; e.g., see Theorems 3.2.5 and 3.4.1 of [10] for specific assumptions. [12] give an example (Section 8 provides a simpler model) where  $V_\varepsilon$  is not negligible relative to  $R_\varepsilon$ , so (8) holds more generally than (9). [6] develops simulation methods that exploit (8) or (9), and we next extend those approaches to construct density estimators.

## 5 Exponential estimator of the density $f$

The weak convergence in (9) suggests that for fixed but small  $\varepsilon > 0$ ,

$$F_\varepsilon(x) = \mathbb{P}_\varepsilon(R_\varepsilon \leq x) = \mathbb{P}_\varepsilon(R_\varepsilon/\mu_\varepsilon \leq x/\mu_\varepsilon) \approx 1 - e^{-x^+/\mu_\varepsilon} \equiv \tilde{F}_{\text{exp},\varepsilon}(x), \quad (10)$$



which we call the *exponential CDF approximation*. This suggests approximating the density  $f$  of  $F$  by

$$\tilde{f}_{\text{exp},\varepsilon}(x) \equiv (1/\mu_\varepsilon)e^{-x/\mu_\varepsilon}\mathcal{I}(x \geq 0). \quad (11)$$

[6] uses (10) to develop simulation estimators of the CDF  $F$  and associated risk measures, and we now adapt this to estimate the density  $f$ .

Dropping the subscript  $\varepsilon$  to simplify notation, we next construct a simulation estimator  $\hat{\mu}_n$  of  $\mu$  in (11). To do this, we apply *measure-specific importance sampling* (MSIS; [7]), which independently estimates the numerator and denominator in (6) using crude simulation and importance sampling, respectively. Let  $n$  be our computation budget of the total number of cycles to simulate, and for a fixed user-specified parameter  $\beta \in (0, 1)$ , we simulate  $n_{\text{CS}} \equiv \lfloor \beta n \rfloor$  (resp.,  $n_{\text{IS}} \equiv n - n_{\text{CS}}$ ) cycles using CS (resp., IS), where  $\lfloor \cdot \rfloor$  denotes the floor (round-down) function.[7] selects  $\beta$  to minimize the asymptotic variance (in the corresponding central limit theorem as  $n \rightarrow \infty$  for fixed  $\varepsilon > 0$ ) of the resulting overall estimator of  $\mu$ . To estimate the numerator  $\zeta = \mathbb{E}[D]$  in (6), we generate  $D_i, i = 1, 2, \dots, n_{\text{CS}}$ , as  $n_{\text{CS}}$  i.i.d. copies of  $D$  sampled using CS. Generating each  $D_i$  entails simulating a cycle until either it ends or  $\mathcal{A}$  is hit, whichever occurs first. A CS estimator of  $\zeta$  is then

$$\hat{\zeta}_n \equiv \frac{1}{n_{\text{CS}}} \sum_{i=1}^{n_{\text{CS}}} D_i. \quad (12)$$

Independently of the simulation runs employed to construct  $\hat{\zeta}_n$  in (12), we use IS to estimate the denominator  $p = \mathbb{P}(T < \tau)$  in (6) as follows. Applying a change of measure leads to

$$p = \mathbb{E}[\mathcal{I}(T < \tau)] = \int \mathcal{I}(T < \tau) d\mathbb{P} = \int \mathcal{I}(T < \tau) \frac{d\mathbb{P}}{d\mathbb{P}'} d\mathbb{P}' = \mathbb{E}'[\mathcal{I}(T < \tau)L], \quad (13)$$

where  $\mathbb{P}'$  (resp.,  $\mathbb{E}'$ ) denotes the probability measure (resp., expectation) under IS, with  $\mathbb{P}$  absolutely continuous with respect to  $\mathbb{P}'$ , and  $L = d\mathbb{P}/d\mathbb{P}'$  is the resulting likelihood ratio. The representation (13) motivates the following approach to estimate  $p$ . Let  $(\mathcal{I}(T'_i < \tau'_i), D'_i, L'_i), i = 1, 2, \dots, n_{\text{IS}}$ , be i.i.d. copies of the cycle-based random vector  $(\mathcal{I}(T < \tau), D, L)$  generated via IS. Then an IS estimator of  $p$  is

$$\hat{p}_n = \frac{1}{n_{\text{IS}}} \sum_{i=1}^{n_{\text{IS}}} \mathcal{I}(T'_i < \tau'_i)L'_i, \quad (14)$$

which is unbiased. Choosing  $\mathbb{P}'$  appropriately, which is problem-specific, can lead to  $\hat{p}_n$  having much smaller variance than its CS counterpart, but a poorly selected  $\mathbb{P}'$  can lead to larger (or even infinite) variance; see [1], Section V.1 and Chapter VI,

and [14] for further details on IS and particular approaches for various settings, including those in Examples 1 and 2.

We combine the estimators  $\widehat{\zeta}_n$  from (12) and  $\widehat{p}_n$  from (14) to obtain the MSIS estimator of  $\mu$  in (6) as

$$\widehat{\mu}_n = \frac{\widehat{\zeta}_n}{\widehat{p}_n}. \quad (15)$$

Putting  $\widehat{\mu}_n$  into (11) results in the following.

**Proposition 1** For a reward function  $r : \mathcal{S} \rightarrow \mathfrak{R}_+$ , the *exponential density estimator*  $\widehat{f}_{\text{exp},n}$  of  $f$  based on (11) satisfies

$$\widehat{f}_{\text{exp},n}(x) = (1/\widehat{\mu}_n)e^{-x/\widehat{\mu}_n} \mathcal{I}(x \geq 0), \quad (16)$$

where  $\widehat{\mu}_n$  is from (15).

## 6 Convolution estimator of the density $f$

With the rarity parameter  $\varepsilon$  reintroduced, the exponential approximation in (11) essentially assumes that  $V_\varepsilon$  makes a negligible contribution to  $R_\varepsilon$  in (5). But for stochastic models in which this is not the case (e.g., see Section 8 and [12]), [6] gives other estimators for the CDF and risk measures of  $R_\varepsilon$  based on a convolution arising from (5) to more explicitly account for the contribution of  $V_\varepsilon$  to  $R_\varepsilon$ . The resulting so-called convolution estimators can apply more generally than the exponential estimators in Section 5, with the convolution estimators often having less bias, especially when  $\varepsilon$  is not so small, which is useful in practice because actual systems have a fixed  $\varepsilon > 0$ . We next extend this convolution approach to construct a density estimator. Dropping again the subscript  $\varepsilon$  to simplify notation, we have that (5) implies that the density  $f$  of  $S+V \sim F$ , with  $S \sim G$  independent of  $V \sim H$ , satisfies

$$f(x) = \int g(x-y) dH(y), \quad (17)$$

assuming that  $G$  has density  $g$  (with respect to Lebesgue measure); e.g., see [3], eq. (20.37).

To develop an estimator of the density  $f$  based on (17), we will use (8) to approximate the density  $g$  of  $S$  by an exponential with mean  $\eta$ . To construct an estimator of  $H$ , [6] again applies MSIS with total sample size  $n$  and a user-specified allocation parameter  $\beta \in (0, 1)$ , as in Section 5. Recall that  $H$  is also the conditional

CDF of  $D$  given that  $T < \tau$  and  $p = \mathbb{P}(T < \tau)$ , so we have

$$\begin{aligned} H(y) &= \mathbb{P}(D \leq y \mid T < \tau) = \frac{\mathbb{P}(D \leq y, T < \tau)}{p} = \frac{\mathbb{E}[\mathcal{I}(D \leq y, T < \tau)]}{p} \\ &= \frac{\mathbb{E}'[\mathcal{I}(D \leq y, T < \tau)L]}{p} \end{aligned} \quad (18)$$

through a change of measure, where as in Section 5,  $\mathbb{E}'$  is the expectation operator under IS, and  $L$  is the corresponding likelihood ratio. We previously obtained in (14) an IS estimator  $\hat{p}_n$  of the denominator  $p$  of (18). From the same IS data used in (14), we then construct an IS estimator  $\hat{H}_n$  of  $H$  as

$$\hat{H}_n(y) = \frac{1}{\hat{p}_n n_{\text{IS}}} \sum_{i=1}^{n_{\text{IS}}} \mathcal{I}(D'_i \leq y, T'_i < \tau'_i) L'_i. \quad (19)$$

Section 7 will consider instead replacing (19) with an IS kernel estimator of  $H$ .

We next develop an estimator of  $g$  in (17) by extending ideas from [6]. In (8), Wald's identity [13, Theorem 3.3.2] yields

$$\eta = \mathbb{E}[S] = \mathbb{E}[M]\mathbb{E}[W] = (1-p)\mathbb{E}[W]/p.$$

The CDF of  $W$  is  $G_W(x) = \mathbb{P}(Y \leq x \mid \tau < T) = \mathbb{P}(Y \leq x, \tau < T)/(1-p)$ , so  $\eta = \mathbb{E}[Y \mathcal{I}(\tau < T)]/p$ . We apply CS to estimate  $\mathbb{E}[Y \mathcal{I}(\tau < T)]$  via

$$\frac{1}{n_{\text{CS}}} \sum_{i=1}^{n_{\text{CS}}} Y_i \mathcal{I}(\tau_i < T_i),$$

where  $(Y_i, \mathcal{I}(\tau_i < T_i))$ ,  $i = 1, 2, \dots, n_{\text{CS}}$ , are i.i.d. copies of the cycle-based random vector  $(Y, \mathcal{I}(\tau < T))$  generated under CS. This then results in

$$\hat{\eta}_n = \frac{1}{\hat{p}_n n_{\text{CS}}} \sum_{i=1}^{n_{\text{CS}}} Y_i \mathcal{I}(\tau_i < T_i) \quad (20)$$

as an MSIS estimator of  $\eta$ . Recalling (8), we obtain a parametric estimator of the exponential approximation  $\tilde{G}_{\text{exp}}(s) = 1 - e^{-s^+/\eta}$  to the CDF  $G$  of  $S$  as  $\hat{G}_{\text{exp},n}(s) = 1 - e^{-s^+/\hat{\eta}_n}$ . The corresponding estimator of the density  $\tilde{g}_{\text{exp}}(s) = (1/\eta)e^{-s/\eta}\mathcal{I}(s \geq 0)$  of  $\tilde{G}_{\text{exp}}$  is

$$\hat{g}_{\text{exp},n}(s) = (1/\hat{\eta}_n)e^{-s/\hat{\eta}_n}\mathcal{I}(s \geq 0). \quad (21)$$

The *convolution estimator* of the density  $f(x)$  in (17) is then

$$\hat{f}_{\star,n}(x) = \int_{y \in [0,x]} \hat{g}_{\text{exp},n}(x-y) d\hat{H}_n(y), \quad (22)$$

and the following works out an expression for  $\hat{f}_{\star,n}(x)$ .

**Proposition 2** For a reward function  $r : \mathcal{S} \rightarrow \mathfrak{R}_+$ , the convolution estimator  $\widehat{f}_{\star,n}$  in (22) of the density  $f$  of  $R$  satisfies

$$\widehat{f}_{\star,n}(x) = \frac{1}{\widehat{p}_n \widehat{\eta}_n n_{\text{IS}}} \sum_{i=1}^{n_{\text{IS}}} \mathcal{I}(D'_i \leq x, T'_i < \tau'_i) L'_i e^{-(x-D'_i)/\widehat{\eta}_n} \quad (23)$$

where  $\widehat{p}_n$  is from (14) and  $\widehat{\eta}_n$  is from (20).

*Proof.* Putting (21) and (19) into (22) leads to

$$\begin{aligned} \widehat{f}_{\star,n}(x) &= \int_{y \in [0,x]} \frac{1}{\widehat{\eta}_n} e^{-(x-y)/\widehat{\eta}_n} \mathcal{I}(y \leq x) \frac{1}{\widehat{p}_n n_{\text{IS}}} \sum_{i=1}^{n_{\text{IS}}} \mathcal{I}(D'_i \in dy) \mathcal{I}(T'_i < \tau'_i) L'_i \\ &= \frac{1}{\widehat{p}_n \widehat{\eta}_n n_{\text{IS}}} \sum_{i=1}^{n_{\text{IS}}} \mathcal{I}(T'_i < \tau'_i) L'_i \int_{y \in [0,x]} e^{-(x-y)/\widehat{\eta}_n} \mathcal{I}(y \leq x) \mathcal{I}(D'_i \in dy), \end{aligned}$$

which equals (23).  $\square$

## 7 Convolution-kernel estimator of the density $f$

The IS estimator  $\widehat{H}_n$  in (19) of the CDF  $H$  of  $V$  in (5) has a point mass at each  $D'_i$ ,  $i = 1, 2, \dots, n_{\text{IS}}$ . Assuming that  $H$  has a density  $h$ , we now devise an IS kernel [11] estimator of  $h$ , which we then convolve with the exponential density estimator  $\widehat{g}_{\text{exp},n}$  in (21) of the geometric sum  $S$  to obtain a convolution-kernel estimator of the density  $f$  of  $R \stackrel{\mathcal{D}}{=} S + V$ .

As  $h(y) = \lim_{\delta \rightarrow 0} [H(y + \delta) - H(y - \delta)] / (2\delta)$ , this suggests using a bandwidth  $\delta > 0$  to estimate  $h$  by the finite difference  $[\widehat{H}_n(y + \delta) - \widehat{H}_n(y - \delta)] / (2\delta) = \frac{1}{2\delta \widehat{p}_n n_{\text{IS}}} \sum_{i=1}^{n_{\text{IS}}} \mathcal{I}(y - \delta < D'_i \leq y + \delta) \mathcal{I}(T'_i < \tau'_i) L'_i$ . Then just as the ordinary kernel estimator (4) generalizes (3), replacing  $\mathcal{I}(y - \delta < D'_i \leq y + \delta) / (2\delta)$  with  $k_\delta(y - D'_i)$  leads to the IS kernel estimator  $\widehat{h}_{n,k,\delta}^\circ$  of  $h$  with

$$\widehat{h}_{n,k,\delta}^\circ(y) = \frac{1}{\widehat{p}_n n_{\text{IS}}} \sum_{i=1}^{n_{\text{IS}}} k_\delta(y - D'_i) \mathcal{I}(T'_i < \tau'_i) L'_i. \quad (24)$$

Recalling the convolution for  $f$  in (17), we next consider a convolution-kernel density estimator

$$\widehat{f}_{\star,n,k,\delta}^\circ(x) = \int_{y=-\infty}^{\infty} \widehat{g}_{\text{exp},n}(x-y) \widehat{h}_{n,k,\delta}^\circ(y) dy \quad (25)$$

of  $f$ , with  $\widehat{g}_{\text{exp},n}$  from (21). Proposition 3 below will derive an expression for  $\widehat{f}_{\star,n,k,\delta}^\circ$ . To do this, let  $\psi_k(\theta) = \int e^{\theta u} k(u) du$ ,  $\theta \in \mathfrak{R}$ , be the *moment generating*

function (MGF) of a kernel  $k$ . Also, for  $z \in \mathfrak{R}$ , define the *lower incomplete* (or partial) *moment-generating function* (LIMGF) [18] of  $k$  as

$$\psi_k^\downarrow(\theta, z) = \int_{u=-\infty}^z e^{\theta u} k(u) du, \quad (26)$$

so  $\psi_k(\theta) = \lim_{z \rightarrow \infty} \psi_k^\downarrow(\theta, z)$ . While the MGF  $\psi_k(\theta)$  for a particular  $k$  may be infinite for some  $\theta$ , the LIMGF always satisfies

$$\psi_k^\downarrow(\theta, z) < \infty, \quad \text{for all } \theta \in \mathfrak{R}_+ \text{ and } z \in \mathfrak{R} \quad (27)$$

because  $\theta \geq 0$  implies  $\psi_k^\downarrow(\theta, z) \leq e^{\theta z} \int_{u=-\infty}^z k(u) du \leq e^{\theta z}$  since  $k$  is a density. We will later give simple expressions for  $\psi_k^\downarrow(\theta, z)$  for particular choices of densities  $k$ .

**Proposition 3** Consider a reward function  $r : \mathcal{S} \rightarrow \mathfrak{R}_+$ , a kernel function  $k$  that is a density, and any bandwidth  $\delta > 0$ . The convolution-kernel estimator  $\hat{f}_{*,n,k,\delta}^\circ$  in (25) of the density  $f$  of  $R$  satisfies

$$\hat{f}_{*,n,k,\delta}^\circ(x) = \frac{1}{\hat{p}_n \hat{\eta}_n n_{\text{IS}}} \sum_{i=1}^{\text{mIS}} \mathcal{I}(T_i' < \tau_i') L_i' e^{(D_i' - x)/\hat{\eta}_n} \psi_k^\downarrow\left(\frac{\delta}{\hat{\eta}_n}, \frac{x - D_i'}{\delta}\right) \quad (28)$$

for LIMGF  $\psi_k^\downarrow(\cdot, \cdot)$  in (26), where  $\hat{p}_n$  is from (14) and  $\hat{\eta}_n$  is from (20).

*Proof.* Putting (21) and (24) into the right side of (25) leads to

$$\begin{aligned} \hat{f}_{*,n,k,\delta}^\circ(x) &= \frac{1}{\hat{p}_n \hat{\eta}_n n_{\text{IS}}} \sum_{i=1}^{\text{mIS}} \mathcal{I}(T_i' < \tau_i') L_i' \int_{y=-\infty}^{\infty} e^{-(x-y)/\hat{\eta}_n} \mathcal{I}(y \leq x) k_\delta(y - D_i') dy \\ &= \frac{1}{\hat{p}_n \hat{\eta}_n n_{\text{IS}}} \sum_{i=1}^{\text{mIS}} \mathcal{I}(T_i' < \tau_i') L_i' e^{-x/\hat{\eta}_n} \int_{y=-\infty}^x e^{y/\hat{\eta}_n} k_\delta(y - D_i') dy. \end{aligned} \quad (29)$$

In (29), substituting  $u = (y - D_i')/\delta$  gives  $k_\delta(y - D_i') dy = k(u) du$ , so the integral in (29) satisfies

$$\int_{y=-\infty}^x e^{y/\hat{\eta}_n} k_\delta(y - D_i') dy = e^{D_i'/\hat{\eta}_n} \int_{u=-\infty}^{(x-D_i')/\delta} e^{(\delta/\hat{\eta}_n)u} k(u) du = e^{D_i'/\hat{\eta}_n} \psi_k^\downarrow\left(\frac{\delta}{\hat{\eta}_n}, \frac{x - D_i'}{\delta}\right), \quad (30)$$

which (27) ensures is finite because  $r(\cdot) \geq 0$  implies  $\delta/\hat{\eta}_n \geq 0$ . Putting (30) into (29) establishes (28).  $\square$

Evaluating  $\hat{f}_{*,n,k,\delta}^\circ(x)$  in (28) requires computing  $\psi_k^\downarrow(\theta, z)$ , which we need to do for only  $\theta \geq 0$  because  $r(\cdot) \geq 0$  implies  $\delta/\hat{\eta}_n \geq 0$  in (28). We now give the

LIMFG  $\psi_k^\downarrow(\theta, z)$  for some particular choices of kernel  $k$ . The uniform $[-1, 1)$  kernel  $k_{\mathcal{U}[-1,1]}$  has a finite MGF  $\psi_{k_{\mathcal{U}[-1,1]}}(\theta) = [e^\theta - e^{-\theta}]/(2\theta)$  for all  $\theta \in \mathfrak{R}$ , and LIMGF  $\psi_{k_{\mathcal{U}[-1,1]}}^\downarrow(\theta, z) = [e^{\theta((z \wedge 1) \vee -1)} - e^{-\theta}]/(2\theta)$ . The Gaussian kernel  $k_{\mathcal{N}}$  has a finite MGF  $\psi_{k_{\mathcal{N}}}(\theta) = e^{\theta^2/2}$  for all  $\theta \in \mathfrak{R}$ , and LIMGF  $\psi_{k_{\mathcal{N}}}^\downarrow(\theta, z) = e^{\theta^2/2} \Phi(z - \theta)$ . The uniform $[0, 2)$  kernel  $k_{\mathcal{U}[0,2]}$  has a finite MGF  $\psi_{k_{\mathcal{U}[0,2]}}(\theta) = [e^{2\theta} - 1]/(2\theta)$  for all  $\theta \in \mathfrak{R}$ , and LIMGF  $\psi_{k_{\mathcal{U}[0,2]}}^\downarrow(\theta, z) = [e^{\theta((z \wedge 2) \vee 0)} - 1]/(2\theta)$ . While the exponential kernel  $k_{\mathcal{E}}$  has finite MGF for only  $\theta < 1$ , in which case  $\psi_{k_{\mathcal{E}}}(\theta) = 1/(1 - \theta)$ , its LIMGF  $\psi_{k_{\mathcal{E}}}^\downarrow(\theta, z) = \frac{1}{\theta - 1} [e^{(\theta - 1)z} - 1] \mathcal{I}(z \geq 0)$  is finite for any  $\theta \in \mathfrak{R}$ ; also see (27).

Similar to the discussion in the last paragraph of Section 2, choosing a symmetric kernel  $k$  can lead to  $\hat{h}_{n,k,\delta}^\circ$  in (24) assigning positive mass to sets of negative values (i.e.,  $\int_{-\infty}^0 \hat{h}_{n,k,\delta}^\circ(y) dy > 0$ ) when the bandwidth  $\delta$  is sufficiently large. Thus, (25) then shows that  $\int_{x=-\infty}^0 \hat{f}_{*,n,k,\delta}^\circ(x) dx > 0$  is possible. This may be undesirable when the true density  $f$  of  $R$  has  $f(x) = 0$  for all  $x < 0$ , as is the case under our assumption that the reward function  $r(\cdot) \geq 0$ . Selecting instead a positive-support kernel avoids these issues, but can lead to a statistically less-efficient estimator of  $f$ ; see the last paragraph of Section 2.

## 8 Numerical results

We next provide numerical results comparing the previous sections' estimators of the density  $f$  of the cumulative reward  $R$  up to a hitting time  $T$  to a set  $\mathcal{A}$  for a simple tractable model. A variation of a 3-state CTMC of [7], our model is a semi-Markov process (SMP; Section 4.8 of [13])  $X = [X(t) : t \geq 0]$  on state space  $\mathcal{S} = \{0, 1, 2\}$  and  $\mathcal{A} = \{2\}$ . The SMP starts in state 0, and returns to state 0 are regenerations of  $X$ . Figure 1 shows the transition probabilities of the embedded discrete-time Markov chain (DTMC)  $Z = [Z_j : j = 0, 1, 2, \dots]$  of  $X$ . From state 0, the embedded DTMC moves to state 1 with probability  $p \in (0, 1)$ , and returns to state 0 (regeneration) with probability  $1 - p$ . From state 1, the DTMC goes to state 2 with probability 1, so  $\mathbb{P}(T < \tau) = p$ . (To make the embedded DTMC regenerative, we also specify the probability of going from state 2 to state 0 to be 1.) Let  $\mathcal{E}(\lambda)$  denote an exponential distribution with rate  $\lambda > 0$ , and  $\mathcal{U}(a, b)$  be a uniform distribution on an interval  $(a, b)$  with  $a < b$ . Given  $Z$ , the holding times in each state  $Z_j$  visited are independent, where successive holding times for visits to state 0 (resp., 1) are  $\mathcal{E}(\lambda_0)$  for some  $\lambda_0 > 0$  (resp.,  $\mathcal{U}(a, b)$  for some  $0 \leq a < b < \infty$ ). (While the holding times in state 2 do not matter for the process up to  $T$ , we define its distribution as  $\mathcal{E}(\lambda_2)$  for some  $\lambda_2 > 0$ .)

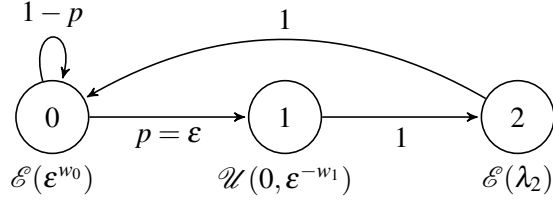


Figure 1: The edge labels are the transition probabilities of the embedded DTMC  $Z$  of an SMP  $X$ , with the holding-time distribution to each visit to a state below that state, where  $w_0$ ,  $w_1$ , and  $\lambda_2$  are constants.

Define the reward function  $r(\cdot) \equiv 1$ , so the cumulative reward  $R$  is the hitting time  $T$  to state 2. The SMP  $X$  has only two possible types of paths for  $T \wedge \tau$ , corresponding to paths of  $Z$  with  $0 \rightarrow 0$  and  $0 \rightarrow 1 \rightarrow 2$ . Let  $M$  be the number of cycles completed before the first visit to  $\mathcal{A}$ , so  $M$  is geometric with parameter  $p$  and support  $\{0, 1, 2, \dots\}$ ; i.e.,  $\mathbb{P}(M = m) = (1 - p)^m p$  for  $m \geq 0$ . Let  $A_1, A_2, \dots$  be the successive holding times in state 0, where each  $A_i \sim \mathcal{E}(\lambda_0)$ . Also, let  $B \sim F_B = \mathcal{U}(a, b)$  be the holding time in the first visit to state 1, so  $R = S + V$  with  $S = \sum_{i=1}^M A_i$  and  $V = A_{M+1} + B$  in (5).

To work out the CDF  $F$  of  $R$ , rewrite  $R$  as  $R = S' + B$ , where  $S' = \sum_{i=1}^{M+1} A_i \sim \mathcal{E}(p\lambda_0)$  [10, p. 7]. As a consequence,  $F$  is the convolution of  $\mathcal{E}(p\lambda_0)$  and  $F_B$ . The density  $f_B$  of  $F_B = \mathcal{U}(a, b)$  is given by  $f_B(y) = \mathcal{I}(y \in (a, b)) / (b - a)$ , so the density  $f$  of  $R$  is

$$f(x) = \mathcal{I}(x \geq 0) \int_0^x p\lambda_0 e^{-p\lambda_0(x-y)} f_B(y) dy = \frac{\mathcal{I}(x \geq 0)}{b - a} e^{-p\lambda_0 x} \left[ e^{p\lambda_0(b \wedge x)} - e^{p\lambda_0(a \wedge x)} \right]. \quad (31)$$

To incorporate the asymptotic regime of (7) into the model, we make visiting  $\mathcal{A}$  before regenerating rare by having  $p = \varepsilon$  for small  $\varepsilon > 0$ . We further set the exponential holding time in state 0 to have rate  $\lambda_0 = \varepsilon^{w_0}$  for a constant  $w_0 \geq 0$ , and the holding time in state 1 is uniform on  $(a, b) = (0, \varepsilon^{-w_1})$  for  $w_1 \geq 0$ . Varying  $(w_0, w_1)$  allows investigating the viability of the various estimators. The mean of  $S'$  (resp.,  $S = S' - A_{M+1}$ ) is  $1/(p\lambda_0) = \varepsilon^{-(w_0+1)}$  (resp.,  $1/(p\lambda_0) - (1/\lambda_0) = \varepsilon^{-(w_0+1)} - \varepsilon^{-w_0}$ ), so both  $\mathbb{E}[S']$  and  $\mathbb{E}[S]$  are of order  $\varepsilon^{-(w_0+1)}$  as  $\varepsilon \rightarrow 0$ . The mean of  $B$  (resp.,  $V = A_{M+1} + B$ ) is  $b/2 = \varepsilon^{-w_1}/2$  (resp.,  $(1/\lambda_0) + (b/2) = \varepsilon^{-w_0} + (\varepsilon^{-w_1}/2)$ ), so  $\mathbb{E}[B]$  and  $\mathbb{E}[V]$  are of respective orders  $\varepsilon^{-w_1}$  and  $\varepsilon^{-\max(w_0, w_1)}$  as  $\varepsilon \rightarrow 0$ . Hence, if  $w_0 + 1 > w_1$  (resp.,  $w_0 + 1 < w_1$ ),  $S'$  and  $S$  typically make the dominant (resp., a negligible) contribution to  $R = S + V = S' + B$  compared to  $B$  and  $V$ . If  $w_0 + 1 = w_1$ , then  $S$ ,  $S'$ ,  $V$ , and  $B$  each typically contribute comparably to  $R$ . Thus, as we will

see in Figure 2, the weak convergence in (9) does not hold when  $w_0 + 1 \leq w_1$ , but (8) is still valid. The exponential estimator of Section 5 is then only appropriate for  $w_0 + 1 > w_1$ , but the convolution estimators of Sections 6 and 7 are applicable for all  $w_0$  and  $w_1$  as  $\varepsilon \rightarrow 0$ .

For  $w_0 = 1$  and  $w_1 \in \{1, 2\}$ , Figure 2 plots the exact density  $f$  in (31), the exponential estimator (“Exp. est.”)  $\hat{f}_{\text{exp},n}$  in (16), the convolution estimator (“Conv”)  $\hat{f}_{*,n}$  in (23), and the convolution-kernel estimator (“Conv:k”)  $\hat{f}_{*,n,k,\delta}^\circ$  in (28), for  $k$  as the Gaussian kernel  $k_{\mathcal{N}}$ , the  $\mathcal{U}[0, 2)$  kernel  $k_{\mathcal{U}[0,2)}$ , or the exponential kernel  $k_{\mathcal{E}}$  (Section 2). All of the estimators apply MSIS (Section 5), where for IS, we change  $p$  to  $p' = 0.8$  for  $Z$  (as in [7]) and use the original CDFs (given  $Z$ ) for the holding times. All estimators are based on simulating a total of  $n = n_{\text{CS}} + n_{\text{IS}} = 10^5$  independent cycles, where we determined the MSIS allocation parameter  $\beta$  (used in  $n_{\text{CS}} = \lfloor \beta n \rfloor$ ) to approximately minimize the variance of the MSIS estimator  $\hat{\mu}_n$  in (15) of  $\mu$ , the variance estimated from a small pilot run. While the end of Section 2 mentions asymptotically optimal bandwidths  $\delta = \delta_n$  of order  $n^{-1/5}$  (resp.,  $n^{-1/3}$ ) for symmetric (resp., positive-support) kernels to minimize MISE, such asymptotics may take effect in practice for only enormous sample sizes, especially when kernels are combined with IS. Instead, our kernel estimators used a bandwidth  $\delta = 500/\sqrt{n_{\text{IS}}}$ , chosen based on ad-hoc experiments to avoid oversmoothing.

The top panels of Figure 2 have  $w_1 = 1 < w_0 + 1$ , so the exponential and convolution (both without and with kernels) estimators are all valid: the estimators get closer to the exact density as  $\varepsilon$  shrinks from 0.3 (left panel) to  $10^{-2}$  (right panel). For the larger  $\varepsilon$ , the true density has a mode at a positive  $x$ , which the exponential estimator does not capture, but the convolution estimators do, although placing the mode somewhat to the right of the true mode. For the smaller  $\varepsilon$ , the true mode’s location moves left, making it virtually indistinguishable from the origin, which is the mode of the exponential estimator.

Figure 2’s bottom panels have  $w_1 = 2 = w_0 + 1$ , invalidating the exponential estimator, but the convolution estimators (without and with kernels) are all still valid as  $\varepsilon \rightarrow 0$ . The convolution estimators get closer to  $f$  as  $\varepsilon$  shrinks from 0.3 (left panel) to  $10^{-2}$  (right panel), but the exponential estimator does not.

Zooming in on the left panels of Figure 2 reveals that the Conv: $k_{\mathcal{N}}$  estimators have been truncated at the origin as they assign positive mass to negative values (see the last paragraph of Section 7). In contrast, the estimators with positive-support kernels  $k_{\mathcal{U}[0,2)}$  and  $k_{\mathcal{E}}$  do not require truncation.

Figure 3 examines the convolution estimators (without and with kernels) in more detail. The left panel shows the errors (i.e., differences of estimator and  $f$ ) for different kernels, where the Gaussian kernel  $k_{\mathcal{N}}$  does about the same as the non-kernel convolution estimator. The kernels  $k_{\mathcal{U}[0,2)}$  and  $k_{\mathcal{E}}$  have larger errors,



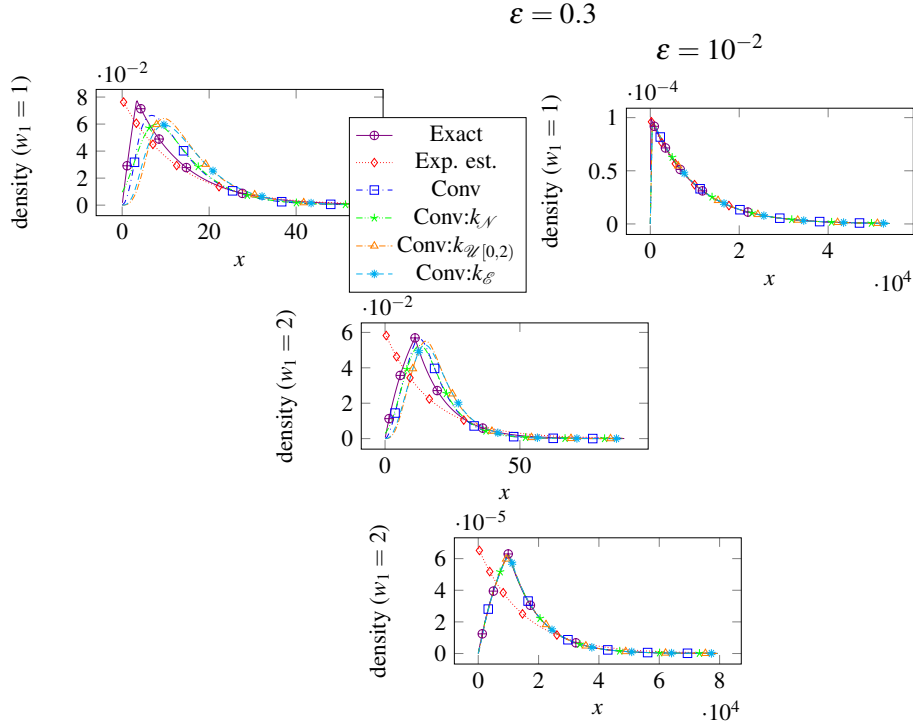


Figure 2: The exact density and estimators for the 3-state SMP with  $w_0 = 1$  are plotted for  $w_1 = 1$  (top panels) and for  $w_1 = 2$  (bottom panels), with  $\varepsilon = 0.3$  (left panels) and  $\varepsilon = 10^{-2}$  (right panels).

in line with the theory (last paragraph of Section 2) for ordinary (non-IS) kernel estimators. The right panel shows that for  $k_{\mathcal{N}}$ , varying the bandwidth  $\delta$  can lead to different errors, where  $\delta = 500$  results in some difference (with some improvement for some small  $x$ ) compared to the non-kernel convolution estimator.

## 9 Concluding remarks

We provided density estimators of the cumulative reward  $R = S + V$  until hitting a rare set  $\mathcal{A}$  of a regenerative process  $X$ . The estimators exploit known weak-convergence results that show that  $S$  or  $R$  divided by its mean converges to an exponential, which [6] uses to construct exponential and convolution estimators of the CDF of  $R$  and related risk measures. We extended these approaches to estimate the density and further incorporated kernels. We introduced a simple model for which numerical results show that our density estimators can do well, and also

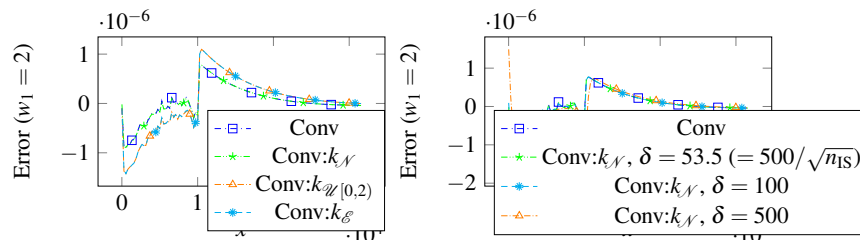


Figure 3: When  $\epsilon = 10^{-2}$  and  $(w_0, w_1) = (1, 2)$ , errors are plotted for different kernels (left panel) and for different bandwidths for the Gaussian kernel (right panel).

demonstrate when the exponential estimator can lead to poor results. ([12] considers another more complicated model that could also be used.) Further work is needed to determine how to choose the bandwidth  $\delta$  for the convolution-kernel estimator.

## References

- [1] S. Asmussen and P. Glynn. *Stochastic Simulation: Algorithms and Analysis*. Springer, New York, 2007.
- [2] I. Bagai and B. L. S. Prakasa Rao. Kernel type density estimates for positive valued random variables. *Sankhya*, 57(1):56–67, 1995.
- [3] P. Billingsley. *Probability and Measure*. John Wiley and Sons, New York, 3rd edition, 1995.
- [4] H. Dong and M. K. Nakayama. A tutorial on quantile estimation via Monte Carlo. In P. L’Ecuyer and B. Tuffin, editors, *Monte Carlo and Quasi-Monte Carlo Methods: MCQMC 2018*. Springer Proceedings in Mathematics and Statistics, 2020.
- [5] P. W. Glynn, M. K. Nakayama, and B. Tuffin. On the estimation of the mean time to failure by simulation. In W. K. V. Chan, A.D’Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, editors, *Proceedings of the 2017 Winter Simulation Conference*, pages 1844–1855, Piscataway, NJ, 2017. Institute of Electrical and Electronics Engineers.
- [6] P. W. Glynn, M. K. Nakayama, and B. Tuffin. Using simulation to calibrate exponential approximations to tail-distribution measures of hitting times to rarely visited sets. In M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, editors, *Proceedings of the 2018 Winter Simulation Conference*, pages 1802–1813, Piscataway, NJ, 2018. Institute of Electrical and Electronics Engineers.
- [7] A. Goyal, P. Shahabuddin, P. Heidelberger, V. Nicola, and P. W. Glynn. A unified framework for simulating Markovian models of highly dependable systems. *IEEE Transactions on Computers*, C-41(1):36–51, 1992.
- [8] L. J. Hong, Z. Hu, and G. Liu. Monte Carlo methods for value-at-risk and conditional value-at-risk: A review. *ACM Transactions on Modeling and Computer Simulataion*, 24(4):22:1–22:37, 2014.

- [9] V. Kalashnikov. *Topics on Regenerative Processes*. CRC Press, Boca Raton, 1994.
- [10] V. Kalashnikov. *Geometric Sums: Bounds for Rare Events with Applications*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997.
- [11] M. K. Nakayama. Asymptotic properties of kernel density estimators when applying importance sampling. In S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu, editors, *Proceedings of the 2011 Winter Simulation Conference*, pages 556–568, Piscataway, New Jersey, 2011. Institute of Electrical and Electronics Engineers.
- [12] M. K. Nakayama and B. Tuffin. Efficient estimation of the mean hitting time to a set of a regenerative system. In N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, editors, *Proceedings of the 2019 Winter Simulation Conference*, pages 416–427, Piscataway, New Jersey, 2019. Institute of Electrical and Electronics Engineers.
- [13] S. M. Ross. *Stochastic Processes*. Wiley, New York, second edition, 1995.
- [14] G. Rubino and B. Tuffin, editors. *Rare Event Simulation Using Monte Carlo Methods*. John Wiley, Chichester, UK, 2009.
- [15] J. S. Sadowsky. Large deviations theory and efficient simulation of excessive backlogs in a  $g_i/g_i/m$  queue. *IEEE Transactions on Automatic Control*, 36:1383–1394, 1991.
- [16] P. Shahabuddin. Importance sampling for highly reliable Markovian systems. *Management Science*, 40(3):333–352, 1994.
- [17] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, London, 1986.
- [18] R. L. Winkler, G. M. Roodman, and R. R. Britney. The determination of partial moments. *Management Science*, 19(3):290–296, 1972.