

Efficient Rare-event Simulation for Perpetuities

Blanchet, J., Lam, H., and Zwart, B.

Abstract

We consider perpetuities of the form

$$D = B_1 \exp(Y_1) + B_2 \exp(Y_1 + Y_2) + \dots,$$

where the Y_j 's and B_j 's might be i.i.d. or jointly driven by a suitable Markov chain. We assume that the Y_j 's satisfy the so-called Cramér condition with associated root $\theta_* \in (0, \infty)$ and that the tails of the B_j 's are appropriately behaved so that D is regularly varying with index θ_* . We illustrate by means of an example that the natural state-independent importance sampling estimator obtained by exponentially tilting the Y_j 's according to θ_* fails to provide an efficient estimator (in the sense of appropriately controlling the relative mean squared error as the tail probability of interest gets smaller). Then, we construct estimators based on state-dependent importance sampling that are rigorously shown to be efficient.

1 Introduction

We consider the problem of developing efficient rare-event simulation methodology for computing the tail of a perpetuity (also known as infinite horizon discounted reward). Perpetuities arise in the context of ruin problems with investments and in the study of financial time series such as ARCH-type processes (see for example, Embrechts *et al.* (1997) and Nyrhinen (2001)).

In the sequel we let $X = (X_n : n \geq 0)$ be an irreducible finite state-space Markov chain (see Section 2 for precise definitions). In addition, let $((\xi_n, \eta_n) : n \geq 1)$ be a sequence of i.i.d. (independent and identically distributed) two dimensional r.v.'s (random variables) independent of the process X . Given $X_0 = x_0$ and $D_0 = d_0$ the associated (suitably scaled by a parameter $\Delta > 0$) discounted reward at time n takes the form

$$\begin{aligned} D_n(\Delta) &= d_0 + \lambda(X_1, \eta_1) \Delta \exp(S_1) + \lambda(X_2, \eta_2) \Delta \exp(S_2) \\ &\quad + \dots + \lambda(X_n, \eta_n) \Delta \exp(S_n) \end{aligned}$$

where the accumulated rate process $(S_k : k \geq 0)$ satisfies

$$S_{k+1} = S_k + \gamma(X_{k+1}, \xi_{k+1}),$$

given an initial value $S_0 = s_0$. In order to make the notation compact, throughout the rest of the paper we shall often omit the explicit dependence of Δ in $D_n(\Delta)$ and we will simply write D_n . We stress that $\Delta > 0$ has been introduced as a scaling parameter which eventually will be sent to zero. Introducing Δ , as we shall see, will be helpful in the development of the state-dependent importance sampling algorithm that we study here.

The functions $(\gamma(x, z) : x \in \mathcal{S}, z \in \mathbb{R})$ and $(\lambda(x, z) : x \in \mathcal{S}, z \in \mathbb{R})$ are deterministic and represent the discount and reward rates respectively. For simplicity we shall assume that $\lambda(\cdot)$ is non-negative. Define

$$\begin{aligned} \phi_{(s_0, d_0, x_0)}(\Delta) &\triangleq P(D_\infty > 1 | S_0 = s_0, D_0 = d_0, X_0 = x_0) \\ &= P(T_\Delta < \infty | S_0 = s_0, D_0 = d_0, X_0 = x_0), \end{aligned} \tag{1}$$

where $T_\Delta = \inf\{n \geq 0 : D_n(\Delta) > 1\}$.

Throughout this paper the distributions of $\lambda(x, \eta_1)$ and $\gamma(x, \xi_1)$ are assumed to be known both analytically and via simulation, as well as the transition probability of the Markov chain X_i . Our main focus on this paper is on the efficient estimation via Monte Carlo simulation of $\phi(\Delta) \triangleq \phi_{(0,0,x_0)}(\Delta)$ as $\Delta \searrow 0$ under the so-called Cramér condition (to be reviewed in Section 2) which in particular implies (see Theorem 1 below)

$$\phi(\Delta) = c_* \Delta^{\theta_*} (1 + o(1)) \tag{2}$$

for a given pair of constants $c_*, \theta_* \in (0, \infty)$. Note that

$$\phi(\Delta) = P\left(\sum_{k=1}^{\infty} \exp(S_k) \lambda(X_k, \eta_k) > \frac{1}{\Delta}\right),$$

so Δ corresponds to the inverse of the tail parameter of interest.

Although our results will be obtained for $s_0 = 0 = d_0$, it is convenient to introduce the slightly more general notation in (1) to deal with the analysis of the state-dependent algorithms that we will introduce.

Approximation (2) is consistent with well known results in the literature (e.g. Goldie (1991)) and it implies a polynomial rate of decay to zero, in $1/\Delta$, for the tail of the distribution of the perpetuity $\sum_{k=1}^{\infty} \exp(S_k) \lambda(X_k, \eta_k)$. The construction of our efficient Monte Carlo procedures is based on importance sampling, which is a variance reduction technique popular in rare-event simulation (see, for instance, Asmussen and Glynn (2008)). It is important to emphasize that, since our algorithms are based on importance sampling, they allow to efficiently estimate conditional expectations of functions of the sample path of $\{D_n\}$ given that $T_\Delta < \infty$. The computational complexity analysis of the estimation of such conditional expectations is relatively straightforward given the analysis of an importance sampling algorithm based on $\phi(\Delta)$ (see for instance the discussion in Adler, Blanchet and Liu (2010)). Therefore, as it is customary in rare-event simulation, we concentrate solely on the algorithmic analysis of a class of estimators for $\phi(\Delta)$.

Asymptotic approximations related to (2) go back to Kesten (1973) who studied a suitable multidimensional analogue of D_∞ . In the one-dimensional setting a key reference is Goldie (1991). Under i.i.d. assumptions, he gave an expression for the constant c_* which is only explicit if θ_* is an integer. More recent work on this type of asymptotics was conducted by Benoit de Saporta (2005), who assumed that the interest rate process $\{\lambda(X_n, \eta_n)\}$ itself forms a finite state Markov chain, and by Enriquez *et. al* (2009), who obtained a different representation for c_* if the X_i are i.i.d. (we will refer to this important special case as the i.i.d. case). Collamore (2009) studied the case when the sequence $\{\lambda(X_n, \eta_n)\}$ is i.i.d. (not dependent on X_n) and $\{\gamma(X_n, \xi_n)\}$ is modulated by a Harris recurrent Markov chain $\{X_n\}$. Since in our case we need certain Markovian assumptions that apparently have not been considered in the literature, at least in the form that we do here, in Section 2, we establish an asymptotic result of the form (2) that fits our framework.

Our algorithms allow to efficiently compute $\phi(\Delta)$ with arbitrary degree of precision, in contrast to the error implied by asymptotic approximations such as (2). In particular, our algorithms can be used to efficiently evaluate the constant c_* , whose value is actually of importance in the statistical theory of ARCH processes (see for example, Chapter 8 of Embrechts *et al.* (1997)).

The efficiency of our simulation algorithms is tested according to widely applied criteria in the context of rare event simulation. These efficiency criteria requires the relative mean squared error of the associated estimator to be appropriately controlled (see for instance the text of Asmussen and Glynn (2008)). Let us recall some basic notions on these criteria in rare-event simulation. An unbiased estimator Z_Δ is said to be strongly efficient if $E(Z_\Delta^2) = O(\phi(\Delta)^2)$. The estimator is said to be asymptotically optimal if $E(Z_\Delta^2) \leq O(\phi(\Delta)^{2-\varepsilon})$ for every $\varepsilon > 0$. Jensen's inequality yields $E Z_\Delta^2 \geq \phi(\Delta)^2$, so asymptotic optimality requires the best possible rate of decay for the second moment, and hence the variance, of the underlying estimator. Despite being a weaker criterion than strong efficiency, asymptotic optimality is perhaps the most popular efficiency criterion in the rare-event simulation literature given its convenience yet sufficiency to capture the rate of decay.

We shall design both strongly efficient and asymptotically optimal estimators and explain the advantages and disadvantages behind each of them from an implementation standpoint. Some of these points of comparison relate to the infinite horizon nature of D_∞ . We are interested in studying unbiased estimators. In addition, besides the efficiency criteria we just mentioned, at the end we are also interested in being able to estimate the overall running time of the algorithm and argue that the total computational cost scales graciously as $\Delta \rightarrow 0$. Our main contributions are summarized as follows:

- 1) The development of an asymptotically optimal state-dependent importance sampling estimator for $\phi(\Delta)$ (see Theorem 3). The associated estimator is shown to be unbiased and the expected termination time of the algorithm is of order $O(\log(1/\Delta)^p)$ for some $p < \infty$ (see Proposition 2 in Section 6).
- 2) An alternative, state-independent, estimator is also constructed which is strongly effi-

cient, see Proposition 2. The state-independent estimator, however, often will have to be implemented incurring in some bias (which can be reduced by increasing the length of a simulation run).

3) New proof techniques based on Lyapunov inequalities. Although Lyapunov inequalities have been introduced recently for the analysis of importance sampling estimators in Blanchet and Glynn (2008), the current setting demands a different approach for constructing the associated Lyapunov function given that the analysis of $\phi(\Delta)$ involves both light-tailed and heavy-tailed features (see the discussion later this section; also see Proposition 1 in Section 5).

4) A new class of counter-examples showing that applying a very natural state-independent importance sampling strategy can in fact lead to infinite variance (see Section 3.2). This contribution adds to previous work by Glasserman and Kou (1995) and further motivates the advantages of state-dependent importance sampling.

5) The development of an asymptotic result of the form (2) that may be of independent interest (see Theorem 1).

As we mentioned earlier, importance sampling is a variance reduction technique whose appropriate usage leads to an efficient estimator. It consists in sampling according to a suitable distribution in order to appropriately increase the frequency of the rare event of interest. The corresponding estimator is just the indicator function of the event of interest times the likelihood ratio between the nominal distribution and the sampling distribution evaluated at the observed outcome. The sampling distribution used to simulate is said to be the importance sampling distribution or the change-of-measure. Naturally, in order to design efficient estimators one has to mimic the behavior of the zero-variance change-of-measure, which coincides precisely with the conditional distribution of $\{D_n\}$ given $D_\infty > 1$. Now, assume that $S_0 = 0 = D_0$. As is known in the literature on ARCH (it is actually made precise in Enriquez *et. al.* (2009)), the event $T_\Delta < \infty$ is typically caused by the event $E_\Delta = \{\max_{k \geq 0} S_k > \log(1/\Delta)\}$, which is the event that the additive process $\{S_k\}$ hits a large value; see Section 2 for more discussion. In turn, the limiting conditional distribution of the underlying random variables given E_Δ as $\Delta \downarrow 0$ is well understood and strongly efficient estimators based on importance sampling have been developed for computing $P(E_\Delta)$ (see Collamore (2002)). Surprisingly, as we will show by means of a simple example, directly applying the corresponding importance sampling estimator which is strongly efficient for $P(E_\Delta)$ can actually result in infinite variance for any $\Delta > 0$ when estimating $\phi(\Delta)$ (see Section 3.2).

Given the issues raised in the previous paragraph, the development of efficient simulation estimators for computing $\phi(\Delta)$ calls for techniques that go beyond the direct application of standard importance sampling estimators. In particular, our techniques are based on state-dependent importance sampling, which has been substantially studied in recent years (see

for example, Dupuis and Wang (2004 and 2007), Blanchet and Glynn (2008), and Blanchet, Leder and Glynn (2009)). The work of Dupuis and Wang provides a criterion, based on a suitable non-linear partial differential inequality, in order to guarantee asymptotic optimality in light tailed settings. It is crucial for the development of Dupuis and Wang to have an exponential rate of decay in the parameter of interest (in our case $1/\Delta$). Blanchet and Glynn (2008) develop a technique based on Lyapunov inequalities that provides a criterion that can be used to prove asymptotic optimality or strong efficiency beyond the exponential decay rate setting. Such criterion, however, demands the construction of a suitable Lyapunov function whose nature varies depending on the type of large deviations environment considered (light vs heavy-tails). The construction of such Lyapunov functions has been studied in light and heavy-tailed environments (see for instance, Blanchet, Leder and Glynn (2009) and Blanchet and Glynn (2008)).

The situation that we consider here is novel since it has both light and heavy-tailed features. On one hand, the large deviations behavior is caused by the event E_Δ , which involves light-tailed phenomena. On the other hand, the scaling of the probability of interest, namely, $\phi(\Delta)$ is not exponential but polynomial in $1/\Delta$ (i.e. the tail of the underlying perpetuity is heavy-tailed, in particular, Pareto with index θ_*). Consequently, the Lyapunov function required to apply the techniques in Blanchet and Glynn (2008) includes new features relative to what has been studied in the literature.

Finally, we mention that while rare event simulation of risk processes has been considered in the literature (see for instance Asmussen (2000)), such simulation in the setting of potentially negative interest rates has been largely unexplored. A related paper is that of Asmussen and Nielsen (1995) in which deterministic interest rates are considered. A conference proceedings version of this paper (Blanchet and Zwart (2007), without proofs) considers the related problem of estimating the tail of a perpetuity with stochastic discounting, but the discounts are assumed to be i.i.d. and premiums and claims are deterministic. Finally, we also note a paper by Collamore (2002), who considered ruin probability of multidimensional random walks that are modulated by general state space Markov chains.

During the second revision of this paper Collamore et al (2011) proposed an independent algorithm for the tail distribution of fixed point equations, which include perpetuities as a particular case. We shall discuss more about this algorithm in Section 7.

The rest of the paper is organized as follows. In Section 2 we state our assumptions and review some large deviations results for $\phi(\Delta)$. Section 3 focuses on state-independent sampling. The state-dependent sampling algorithm is developed in Section 4, and its efficiency analysis and cost-per-replication are studied in Sections 5 and 6. In Section 7 we include additional extensions and considerations. Section 8 illustrates our results with a numerical example.

2 Some Large Deviations for Perpetuities

As discussed in the Introduction, we shall assume that the process $(S_n : n \geq 0)$ is a Markov random walk. As it is customary in the large deviations analysis of quantities related to these objects, we shall impose some assumptions in order to guarantee the existence of an asymptotic logarithmic moment generating function for S_k .

First we have the following assumption:

Assumption 1: We assume that X is an irreducible Markov chain taking values in a finite state space \mathcal{S} with transition matrix $(K(x, y) : x, y \in \mathcal{S})$. Moreover, we further assume that the ξ_k 's and $\gamma(\cdot)$ satisfy

$$\sup_{x \in \mathcal{S}, \theta \in \mathcal{N}} E \exp(\theta \gamma(x, \xi_1)) < \infty, \quad (3)$$

where \mathcal{N} is a neighborhood of the origin.

If Assumption 1 is in force, the Perron-Frobenius theorem for positive and irreducible matrices guarantees the existence of $(u_\theta(x) : x \in \mathcal{S}, \theta \in \mathcal{N})$ and $\exp(\psi(\theta))$ so that

$$u_\theta(x) = E_x[\exp(\theta \gamma(X_1, \xi_1) - \psi(\theta)) u_\theta(X_1)]. \quad (4)$$

The function $u_\theta(\cdot)$ is strictly positive and unique up to constant scalings. Indeed, to see how the Perron-Frobenius theorem is applied, define

$$E \exp(\theta \gamma(x, \xi_1)) = \exp(\chi(x, \theta))$$

and note that (4) is equivalent to the eigenvalue problem

$$(Q_\theta u_\theta)(x) = \exp(\psi(\theta)) u_\theta(x),$$

where $Q_\theta(x, y) = K(x, y) \exp(\chi(y, \theta))$.

We also impose the following assumption that is often known as Cramér's condition.

Assumption 2: Suppose that there exists $\theta_* > 0$ such that $\psi(\theta_*) = 0$. Moreover, assume that there exists $\theta > \theta_*$ such that $\psi(\theta) < \infty$.

In order to better understand the role of Assumption 2, it is useful to note that under Assumption 1, given $X_0 = x_0$, $\tau(x_0) = \inf\{k \geq 1 : X_k = x_0\}$ is finite almost surely and $D = \sum_{k=1}^{\infty} \exp(S_k) \lambda(X_k, \eta_k)$ admits the decomposition

$$D = B + \exp(Y) D', \quad (5)$$

where D' is identical in distribution to D , and

$$Y = S_{\tau(x_0)}, \quad B = \sum_{j=1}^{\tau(x_0)} \lambda(X_j, \eta_j) \exp(S_j), \quad (6)$$

and D' is equal in distribution to D and independent of $(B, S_{\tau(x_0)})$. In other words, D can be represented as a perpetuity with i.i.d. pairs of reward and discount rates. This decomposition will be invoked repeatedly during the course of our development. Now, as we shall see in the proof of Theorem 1 below, it follows that $\theta_* > 0$ appearing in Assumption 2 is the Cramér root associated to Y , that is,

$$E \exp(\theta_* Y) = 1. \tag{7}$$

Note also that since the moment generating function of Y is convex, and $\theta_* > 0$, we must have that $EY < 0$ and therefore by regenerative theory we must have that $E\gamma(X_\infty, \xi_1) < 0$.

An additional final assumption is imposed in our development.

Assumption 3: Assume that $\sup_{x \in \mathcal{S}} E_x \lambda(X_1, \eta_1)^\alpha < \infty$ for each $\alpha \in (0, \infty)$.

The following examples are given to illustrate the flexibility of our framework.

Example 1 ARCH sequences have been widely used in exchange rate and log-return models (see for example, Embrechts *et al.* (1997)). In these models the object of interest, A_n , are the standard deviations of the log-return. The simplest case of ARCH sequences is the ARCH(1) process, which satisfies

$$A_{n+1}^2 = (\alpha_0 + \alpha_1 A_n^2) Z_{n+1}^2.$$

Typically, the Z_n 's are i.i.d. standard Gaussian random variables, and $\alpha_0 > 0$ and $\alpha_1 < 1$. The stationary distribution is a perpetuity. We can directly work with the stationary distribution of this process or transform the problem into one with constant rewards (equal to α_0) by noting that

$$T_{n+1} \triangleq \alpha_0 + \alpha_1 A_{n+1}^2 = \alpha_0 + \alpha_1 (\alpha_0 + \alpha_1 A_n^2) Z_{n+1}^2 = \alpha_0 + \alpha_1 T_n Z_{n+1}^2.$$

We obtain that

$$T_\infty - \alpha_0 \stackrel{D}{=} B_1 \exp(Y_1) + B_2 \exp(Y_1 + Y_2) + \dots$$

where $B_i = \alpha_0$ and $Y_i = \log(\alpha_1 Z_i^2)$ for $i \geq 1$. Assumptions 1 to 3 are in place in this setting.

Example 2 A changing economic environment can be modeled by say, a two-state Markov chain denoting good and bad economic states. We can then model the discounted return of a long-term investment under economic uncertainty as a perpetuity with this underlying Markov modulation. Denoting $X_i \in \{\text{good, bad}\}, i = 1, 2, \dots$ as the Markov chain, our return can be represented as

$$D = B_1 \exp(Y_1) + B_2 \exp(Y_1 + Y_2) + \dots$$

where $B_i = \lambda(X_i, \eta_i)$ and $Y_i = \gamma(X_i, \xi_i)$ are the return and discount rate at time i , and η_i and ξ_i are i.i.d. r.v.'s denoting the individual random fluctuations.

We now state a result that is useful to test the optimality of our algorithms.

Theorem 1 Under Assumptions 1 to 3,

$$\phi(\Delta) = c_* \Delta^{\theta_*} (1 + o(1))$$

as $\Delta \searrow 0$.

The proof of Theorem 1 will be given momentarily, but first let us discuss the intuition behind the asymptotics described in the theorem. It is well known that under our assumptions $\exp(\max\{S_k : k \geq 0\})$ is regularly varying with index $-\theta_* < 0$ (a brief argument indicating the main ideas behind this fact is given in Section 3.1). The principle of the largest jump in heavy-tailed analysis indicates that the large deviations behavior of D is dictated by a large jump of size b/Δ arising from the largest contribution in the sum of the terms defining D . Given that the reward rates are light-tailed, such contribution is likely caused by $\exp(\max\{S_k : k \geq 0\})$, as made explicit by Enriquez *et al.* (2009) in the i.i.d. case. Therefore, the large deviations behavior of D is most likely caused by the same mechanism that causes a large deviations behavior in $\exp(\max\{S_k : k \geq 0\})$. This type of intuition will be useful in developing an efficient importance sampling scheme for estimating the tail of D .

Proof of Theorem 1. Note that equations (5) and (6) allow us to apply Theorem 4.1 from Goldie (1991). In particular, in order to apply Goldie's results we need to show that

$$E_{x_0} \exp(\theta_* S_{\tau(x_0)}) = 1, \quad (8)$$

$$E_{x_0} \exp(\theta S_{\tau(x_0)}) < \infty, \quad (9)$$

for some $\theta > \theta_*$ and that

$$E_{x_0} B^\alpha < \infty \quad (10)$$

for some $\alpha > \theta_*$ (conditions (8), (9) and (10) here correspond to conditions (2.3), (2.4) and (4.2) respectively in Goldie (1991)). First we show (8). Note that equation (4) implies that the process

$$M_n^{\theta_*} = \frac{u_{\theta_*}(X_n)}{u_{\theta_*}(x_0)} \exp(\theta_* S_n)$$

is a positive martingale. Therefore, we have that

$$\begin{aligned} 1 &= E_{x_0} M_{n \wedge \tau(x_0)}^{\theta_*} \\ &= E_{x_0} [\exp(\theta_* S_{\tau(x_0)}) I(\tau(x_0) \leq n)] \\ &\quad + E_{x_0} \left[\frac{u_{\theta_*}(X_n)}{u_{\theta_*}(x_0)} \exp(\theta_* S_n) I(\tau(x_0) > n) \right]. \end{aligned}$$

By the monotone convergence theorem we have that

$$E [\exp(\theta_* S_{\tau(x_0)}) I(\tau(x_0) \leq n)] \longrightarrow E \exp(\theta_* S_{\tau(x_0)})$$

as $n \nearrow \infty$. On the other hand note that (4) implies that the matrix $(K_{\theta_*}(x, y) : x, y \in \mathcal{S})$ defined via

$$K_{\theta_*}(x, y) = K(x, y) \frac{u_{\theta_*}(y) \exp(\chi(y, \theta_*))}{u_{\theta_*}(x)}$$

is an irreducible stochastic matrix. Now let $\left(\tilde{K}_{\theta_*}(x, y) : x, y \in \mathcal{S} \setminus \{x_0\}\right)$ be the submatrix of $K_{\theta_*}(\cdot)$ that is obtained by removing the row and column corresponding to state x_0 . Observe that

$$\left(\tilde{K}_{\theta_*}^n \mathbf{1}\right)(x_0) = E_{x_0} \left[\frac{u_{\theta_*}(X_n)}{u_{\theta_*}(x_0)} \exp(\theta_* S_n) I(\tau(x_0) > n) \right].$$

By irreducibility we have that $\tilde{K}_{\theta_*}(\cdot)$, therefore

$$E_{x_0} \left[\frac{u_{\theta_*}(X_n)}{u_{\theta_*}(x_0)} \exp(\theta_* S_n) I(\tau(x_0) > n) \right] \longrightarrow 0$$

as $n \nearrow \infty$, obtaining (8). The bound (9) follows easily by noting that

$$\begin{aligned} & E_{x_0} [\exp(\theta S_{\tau(x_0)})] \\ &= \sum_{k=1}^{\infty} E_{x_0} [\exp(\theta S_k); \tau(x_0) > k-1, X_k = x_0] \\ &= \sum_{k=1}^{\infty} E_{x_0} [\exp(\theta S_{k-1}) K(X_{k-1}, x_0) \exp(\chi(\theta, x_0)); \tau(x_0) > k-1]. \end{aligned}$$

So, if we define, for $x \neq x_0$

$$v_{\theta}(x) = \exp(\chi(\theta, x_0)) K(x, x_0) \frac{u_{\theta}(x_0)}{u_{\theta}(x)},$$

and

$$R_{\theta}(x, y) = \exp(\chi(\theta, y)) K(x, y) \frac{u_{\theta}(y)}{u_{\theta}(x)}$$

for $x, y \neq x_0$ we see that

$$E_{x_0} [\exp(\theta S_{\tau(x_0)})] = \sum_{k=1}^{\infty} (R_{\theta}^k v_{\theta})(x_0).$$

Note that $R_{\theta_*} = \tilde{K}_{\theta_*}$ is strictly substochastic. So, by continuity there exists $\theta > \theta_*$ for which R_{θ} has a spectral radius which is strictly less than one and therefore (9) holds. Finally, we establish (10). Observe that

$$B \leq \tau(x_0) \max_{1 \leq k \leq \tau(x_0)} \lambda(X_k, \eta_k) \exp(\max\{S_k : 1 \leq k \leq \tau(x_0)\}).$$

Therefore, for $1/p + 1/q + 1/r = 1$ and $p, q, r > 1$ we have that

$$\begin{aligned} EB^{\alpha} &\leq E\tau(x_0)^{\alpha} \max_{1 \leq k \leq \tau(x_0)} \lambda(X_k, \eta_k)^{\alpha} \exp(\max\{\alpha S_k : 1 \leq k \leq \tau(x_0)\}) \\ &\leq (E_{x_0} \tau(x_0)^{p\alpha})^{1/p} \times (E_{x_0} \max_{1 \leq k \leq \tau(x_0)} \lambda(X_k, \eta_k)^{q\alpha})^{1/q} \\ &\quad \times (E_{x_0} \exp(\max\{r\alpha S_k : 1 \leq k \leq \tau(x_0)\}))^{1/r}. \end{aligned}$$

Since $E_{x_0} \tau(x_0)^{p\alpha} + E_{x_0} \max_{1 \leq k \leq \tau(x_0)} \lambda(X_k, \eta_k)^{q\alpha} < \infty$ for all $p, q \in (0, \infty)$ it suffices to show that

$$E_{x_0} \exp(\max\{\theta S_k : 1 \leq k \leq \tau(x_0)\}) < \infty \quad (11)$$

for some $\theta > \theta_*$. In order to do this define $T(b) = \inf\{k \geq 0 : S_k > b\}$ and note that

$$P_{x_0}(\max\{S_k : 1 \leq k \leq \tau(x_0)\} > b) \leq P_{x_0}(T(b) \leq \tau(x_0))$$

Bound (9) implies that $S_{\tau(x_0)}$ decays at an exponential rate that is strictly larger than θ_* . To analyze $P_{x_0}(T(b) \leq \tau(x_0))$, let $P_{x_0}^{\theta_*}(\cdot)$ (respectively $E_{x_0}^{\theta_*}(\cdot)$) be the probability measure (resp. the expectation operator) associated to the change-of-measure induced by the martingale $(M_n^{\theta_*} : n \geq 0)$ introduced earlier. Note that

$$\begin{aligned} P_{x_0}(T(b) \leq \tau(x_0)) &= E_{x_0}^{\theta_*} \left(\exp(-\theta_* S_{T(b)}) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T(b)})} I(T(b) \leq \tau(x_0)) \right) \\ &\leq c \exp(-\theta_* b) P_{x_0}^{\theta_*}(\tau(x_0) \geq T(b)). \end{aligned}$$

The probability $P_{x_0}^{\theta_*}(\tau(x_0) > T(b))$ decays exponentially fast as $b \nearrow \infty$ by using a standard large deviations argument on $T(b)/b$ and because $\tau(x_0)$ has exponentially decaying tails. Therefore we obtain (11), which in turn yields the conclusion of the theorem. ■

3 State-Independent Importance Sampling

In order to design efficient estimators we will apply importance sampling. It is well known (see for example, Liu (2001)) that the zero-variance importance sampling distribution is dictated by the conditional distribution of the process (in this case the triplet $\{(S_n, D_n, X_n)\}$), given the occurrence of the rare event in question. As we discussed in the previous section, the occurrence of the event $T_\Delta < \infty$ is basically driven by the tail behavior of $\max\{S_n : n \geq 0\}$. In turn, the large deviations behavior of Markov random walks (such as S) is well understood from a simulation standpoint and, under our assumptions, there is a natural state-independent change-of-measure that can be shown to be efficient for estimating the tail of $\max\{S_n : n \geq 0\}$.

The present section is organized as follows. First, we shall explain the change-of-measure that is efficient for estimating the tail of $\max\{S_n : n \geq 0\}$ because it will serve as the basis for our change-of-measure in the setting of perpetuities (but some modifications are crucial to guarantee good performance). After that, we shall show by means of an example that this type of importance sampling algorithm can lead to estimators that have infinite variance. We then close this section with a modified state-independent importance sampling algorithm that is strongly efficient but biased.

3.1 The standard approach

As we indicated in the proof of Theorem 1, given θ and $X_0 = x_0$ equation (4) indicates that the process

$$M_n^\theta = \exp(\theta S_n - n\psi(\theta)) \frac{u_\theta(X_n)}{u_\theta(x_0)}, \quad n \geq 0$$

is a positive martingale as long as $\psi(\theta) < \infty$ and therefore it generates a change-of-measure. The probability measure in path-space induced by this martingale is denoted by $P_{x_0}^\theta(\cdot)$ and in order to simulate the process $((D_n, S_n, X_n) : n \geq 0)$ according to $P_{x_0}^\theta(\cdot)$ one proceeds as follows.

1. Generate X_1 according to the transition matrix

$$K_\theta(x, y) = K(x, y) \exp(\chi(y, \theta) - \psi(\theta)) u_\theta(y) / u_\theta(x),$$

which is guaranteed to be a Markov transition matrix by the definition of $u_\theta(\cdot)$ and $\chi(\cdot, \theta)$.

2. Given $X_1 = y$, sample $\gamma(y, \xi_1)$ according to exponential tilting given by

$$P^\theta(\gamma(y, \xi_1) \in dz) = \exp(\theta z - \chi(y, \theta)) P(\gamma(y, \xi_1) \in dz).$$

3. Simulate $\lambda(y, \eta_1)$ according to the nominal conditional distribution of η_1 given that $X_1 = y$ and $\gamma(y, \xi_1)$.

The previous rules allow to obtain (D_1, S_1, X_1) given (D_0, S_0, X_0) . Subsequent steps are performed in a completely analogous way.

Note that if one selects $\theta = \theta_*$ then we have that $\psi(\theta_*) = 0$ and also $S_n/n \rightarrow \psi'(\theta_*) > 0$ a.s. with respect to $P_{x_0}^{\theta_*}(\cdot)$. If $\sigma(y) = \inf\{n \geq 0 : S_n > y\}$ then

$$\begin{aligned} P_{x_0}(\max\{S_n : n \geq 0\} > y) &= P_{x_0}(\sigma(y) < \infty) \\ &= E_{x_0}^{\theta_*}[1/M_{\sigma(y)}^{\theta_*}; \sigma(y) < \infty] \\ &= E_{x_0}^{\theta_*} \left(\exp(-\theta_* S_{\sigma(y)}) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{\sigma(y)})} \right). \end{aligned}$$

In the last equality we have used that $\sigma(y) < \infty$ a.s. with respect to $P_{x_0}^{\theta_*}(\cdot)$. The importance sampling estimator

$$Z'_y = \exp(-\theta_* S_{\sigma(y)}) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{\sigma(y)})},$$

obtained by sampling according to $P_{x_0}^{\theta_*}(\cdot)$ is unbiased and its second moment satisfies

$$\begin{aligned} E_{x_0}^{\theta_*}(Z_y'^2) &= E_{x_0}^{\theta_*} \exp(-2\theta_* S_{\sigma(y)}) \frac{u_{\theta_*}^2(x_0)}{u_{\theta_*}^2(X_{\sigma(y)})} \\ &\leq c \exp(-2\theta_* y), \end{aligned}$$

for some constant $c \in (0, \infty)$. Since, under our assumptions (see for example, Asmussen (2003), Theorems 5.2 and 5.3 on Page 365) $P_{x_0}(\sigma(y) < \infty) = \gamma_* \exp(-\theta_* y) (1 + o(1))$ as $y \nearrow \infty$ for a suitable $\gamma_* \in (0, \infty)$ we obtain that Z'_y is strongly efficient for estimating $P_{x_0}(\sigma(y) < \infty)$ as $y \nearrow \infty$.

3.2 Standard importance sampling can lead to infinite variance

A natural approach would be to apply directly the previous change-of-measure for estimating $T_\Delta < \infty$. Nevertheless, we will show by means of a simple continuous-time example that can be fit in the discrete time setting (through simple embedding) that such approach does not guarantee efficiency. In fact, the estimator might even have infinite variance.

Example 3: Let $(X(t) : t \geq 0)$ be Brownian motion with negative drift $-\mu$ and unit variance. We are concerned with $\phi(\Delta) = P(\int_0^\infty \exp(X(t)) dt > 1/\Delta)$. In particular, here we have $\psi(\theta) = t^{-1} \log E \exp(\theta X(t)) = -\mu\theta + \theta^2/2$ (in the discrete time setting $\psi(\theta)$ will be $-\mu\theta + \theta^2/2$ multiplied by the discretization time scale). In order to analyze the second moment of the natural estimator described previously we need to evaluate θ_* so that $\psi(\theta_*) = 0$. This yields $\theta_* = 2\mu$ and the resulting importance sampling algorithm proceeds by simulating $X(\cdot)$ according to a Brownian motion with *positive* drift μ and unit variance up to time $T_\Delta = \inf\{t \geq 0 : \int_0^t \exp(X(s)) ds \geq 1/\Delta\}$ and returning the estimator $Z_\Delta = \exp(-\theta_* X(T_\Delta))$. The second moment of the estimator is then given by $E^{\theta_*}[Z_\Delta^2; T_\Delta < \infty] = E[Z_\Delta; T_\Delta < \infty]$, by a change of measure back to the original one (here we have used $E^{\theta_*}(\cdot)$ to denote probability measure under which $X(\cdot)$ follows a Brownian motion with drift μ and unit variance). We will show that if $\theta_* \geq 1$ then $E[\exp(-\theta_* X(T_\Delta)) | T_\Delta < \infty] = \infty$. By the definition of θ_* we will conclude that if $\mu \geq 1/2$ then

$$E[Z_\Delta; T_\Delta < \infty] = E[\exp(-\theta_* X(T_\Delta)) | T_\Delta < \infty] P(T_\Delta < \infty) = \infty$$

which implies infinite variance of the estimator. In order to prove this we will use a result of Pollack and Siegmund (1985) (see also Dufresne (1990)) which yields that $D = \int_0^\infty \exp(X(t)) dt$ is equal in distribution to $1/Z$, where Z is distributed gamma with a density of the form

$$f_Z(z) = \frac{\exp(-\lambda z) \lambda^{\theta_*} z^{\theta_*-1}}{\Gamma(\theta_*)},$$

for some $\lambda > 0$. In particular, a transformation of variable gives the density of D as

$$f_D(y) = \frac{\exp(-\lambda/y) \lambda^{\theta_*}}{\Gamma(\theta_*) y^{\theta_*+1}}$$

and hence

$$P(D > 1/\Delta) = \int_{1/\Delta}^\infty \frac{\exp(-\lambda/y) \lambda^{\theta_*}}{\Gamma(\theta_*) y^{\theta_*+1}} dy = \frac{\exp(-\lambda\Delta) \lambda^{\theta_*} \Delta^{\theta_*}}{\theta_* \Gamma(\theta_*)} + \int_{1/\Delta}^\infty \frac{\exp(-\lambda/y) \lambda^{\theta_*+1}}{\theta_* \Gamma(\theta_*) y^{\theta_*+2}} dy$$

where the second inequality follows from integration by parts. Note that

$$\int_{1/\Delta}^{\infty} \frac{\exp(-\lambda/y)\lambda^{\theta_*+1}}{\theta_*\Gamma(\theta_*)y^{\theta_*+2}} dy \leq \frac{\lambda^{\theta_*+1}}{\theta_*\Gamma(\theta_*)} \int_{1/\Delta}^{\infty} \frac{1}{y^{\theta_*+2}} dy = \frac{\lambda^{\theta_*+1}\Delta^{\theta_*+1}}{\theta_*(\theta_*+1)\Gamma(\theta_*)} = O(\Delta^{\theta_*+1})$$

as $\Delta \searrow 0$. Hence $P(D > 1/\Delta) \sim c\Delta^{\theta_*}$ where $c = \exp(-\lambda\Delta)\lambda^{\theta_*}/(\theta_*\Gamma(\theta_*))$. Now, let W be a random variable equal in law to $\Delta[D - 1/\Delta]$ given that $D > 1/\Delta$. Note that conditional on $T_\Delta < \infty$, we have

$$\begin{aligned} D &= \int_0^\infty \exp(X(t)) dt \\ &= \int_0^{T_\Delta} \exp(X(t)) dt + \int_{T_\Delta}^\infty \exp(X(t)) dt \\ &= \frac{1}{\Delta} + \exp(X(T_\Delta)) \int_{T_\Delta}^\infty \exp(X(t) - X(T_\Delta)) dt \\ &= \frac{1}{\Delta} + \exp(X(T_\Delta)) D' \end{aligned}$$

where D' has the same distribution as $1/Z$ and is independent of $\exp(X(T_\Delta))$ given that $T_\Delta < \infty$. We have used the strong Markov property and stationary increment property of Brownian motion in the fourth equality. Hence the random variable $X(T_\Delta)$ satisfies the equality in distribution

$$W =_d \Delta \exp(X(T_\Delta)) D',$$

Now, it is clear that $E[(D')^{-\theta_*}] = EZ^{\theta_*} < \infty$, so it suffices to show that if $\theta_* \geq 1$ then $E(W^{-\theta_*}) = \infty$. Using the definition of W , transformation of variable gives

$$f_W(w) = \frac{\Delta^{-1}\lambda^{\theta_*}}{\Gamma(\theta_*)P(D > 1/\Delta)} \exp(-\lambda/(w/\Delta + 1/\Delta)) (w/\Delta + 1/\Delta)^{-(\theta_*+1)}.$$

Therefore, we have that there exists a constant $c_0 \in (0, \infty)$ such that

$$\begin{aligned} E(W^{-\theta_*}) &= \int_0^\infty f_W(w) w^{-\theta_*} dw \\ &\geq c_0 \int_0^\infty \Delta^{-1} \exp(-\lambda/(w/\Delta + 1/\Delta)) (w/\Delta + 1/\Delta)^{-(\theta_*+1)} (\Delta w)^{-\theta_*} dw \\ &\geq c_0 \int_0^1 \exp(-\lambda\Delta) 2^{-(\theta_*+1)} w^{-\theta_*} dw = \infty. \end{aligned}$$

The problem behind the natural importance sampling estimator is that one would like the difference $[S_{T_\Delta} - \log(1/\Delta)]$ to stay positive, but unfortunately, this cannot be guaranteed and in fact, this difference will likely be negative. The idea that we shall develop in the next subsection is to apply importance sampling just long enough to induce the rare event.

3.3 A modified algorithm

We select $\theta = \theta_*$ and simulate the process according to the procedure described in Steps 1 to 3 explained in Section 3.1 up to time

$$T_{\Delta/a} = \inf\{n \geq 0 : D_n > a\},$$

for some $a \in (0, 1)$. Subsequent steps of the process $\{(S_k, D_k, X_k)\}$, for $k > T_{\Delta/a}$ are simulated under the nominal (original) dynamics up until T_{Δ} . The resulting estimator takes the form

$$Z_{1,\Delta} = \exp\left(-\theta_* S_{T_{\Delta/a}}\right) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T_{\Delta/a}})} I(D_{\infty} > 1). \quad (12)$$

We will discuss the problem of implementing this estimator in a moment, in particular the problem of sampling $I(D_{\infty} > 1)$ in finite time. First we examine its efficiency properties. We assume for simplicity that the rewards are bounded, we will discuss how to relax this assumption right after the proof of this result.

Theorem 2 *In addition to Assumptions 1 and 2, suppose that there is deterministic constant $m \in (0, \infty)$ such that $\lambda(x, \eta) < m$. Then, $Z_{1,\Delta}$ is a strongly efficient estimator of $\phi(\Delta)$.*

Proof. It is clear that $Z_{1,\Delta}$ is unbiased. Now, define

$$U_n = \exp(-S_n) \{D_n - a\} / \Delta$$

for $n \geq 1$. The process $\{U_n\}$ will be helpful to study the overshoot at time $T_{\Delta/a}$. Note that

$$U_{n+1} = \lambda(X_{n+1}, \eta_{n+1}) + \exp(-\gamma(X_{n+1}, \xi_{n+1})) U_n, \quad (13)$$

and also that we can write $T_{\Delta/a} = \inf\{n \geq 0 : U_n > 0\}$.

It is important to observe that

$$D_{\infty} = a + \exp\left(S_{T_{\Delta/a}}\right) \Delta U_{T_{\Delta/a}} + \exp\left(S_{T_{\Delta/a}}\right) D'_{\infty}, \quad (14)$$

where D'_{∞} is conditionally independent of $S_{T_{\Delta/a}}$, $U_{T_{\Delta/a}}$ given $X_{T_{\Delta/a}}$. In addition, D'_{∞} is obtained from the original / nominal distribution. Decomposition (14) implies that

$$I(D_{\infty} > 1) \leq I(\exp\left(S_{T_{\Delta/a}}\right) D'_{\infty} > (1-a)/2) + I(\exp\left(S_{T_{\Delta/a}}\right) \Delta U_{T_{\Delta/a}} > (1-a)/2),$$

and therefore, by conditioning on (S_n, D_n, X_n) for $n \leq T_{\Delta/a}$ we obtain that the second moment of $Z_{1,\Delta}$ is bounded by

$$E_{(0,0,x_0)}^{\theta_*} \left[\exp\left(-2\theta_* S_{T_{\Delta/a}}\right) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T_{\Delta/a}})} \phi_{(0,0,X_{T_{\Delta/a}})} \left(\exp\left(S_{T_{\Delta/a}}\right) 2\Delta / (1-a) \right) \right] \quad (15)$$

$$+ E_{(0,0,x_0)}^{\theta_*} \left[\exp\left(-2\theta_* S_{T_{\Delta/a}}\right) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T_{\Delta/a}})} I(\exp\left(S_{T_{\Delta/a}}\right) \Delta U_{T_{\Delta/a}} > (1-a)/2) \right]. \quad (16)$$

We will denote by I_1 the term in (15) and by I_2 the term in (16). It suffices to show that both I_1 and I_2 are of order $O(\Delta^{2\theta_*})$.

Theorem 1 guarantees the existence of a constant $c_1 \in (1, \infty)$ so that

$$\phi_{(0,0,x_0)}(\Delta) \leq c_1 \exp\left(\theta_* S_{T_{\Delta/a}}\right) \Delta^{\theta_*} / (1-a)^{\theta_*}.$$

Using this bound inside (15) we obtain that

$$I_1 \leq m_1 \frac{\Delta^{\theta_*}}{(1-a)^{\theta_*}} E_{(0,0,x_0)}^{\theta_*} \left[\exp\left(-2\theta_* S_{T_{\Delta/a}}\right) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T_{\Delta/a}})} \right] \quad (17)$$

for some constant $m_1 > 0$ and thus, since

$$E_{(0,0,x_0)}^{\theta_*} \left[\exp\left(-\theta_* S_{T_{\Delta/a}}\right) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T_{\Delta/a}})} \right] = \phi(\Delta) = O(\Delta^{\theta_*}) \quad (18)$$

we conclude that $I_1 = O(\Delta^{2\theta_*})$.

We now study the term I_2 . Just as in the proof of Markov's inequality, note that for any $\beta > 0$

$$I_2 \leq \Delta^\beta \left(\frac{2}{1-\alpha}\right)^\beta E_{(0,0,x_0)}^{\theta_*} \left[\exp\left(-2\theta_* S_{T_{\Delta/a}}\right) \exp\left(\beta S_{T_{\Delta/a}}\right) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T_{\Delta/a}})} U_{T_{\Delta/a}}^\beta \right]. \quad (19)$$

We could pick, for instance, $\beta = 3\theta_*/2$ and use the fact that $u_{\theta_*}(X_{T_{\Delta/a}}) \geq \delta$ for some $\delta > 0$ to obtain that

$$\begin{aligned} & E_{(0,0,x_0)}^{\theta_*} \left[\exp\left(-2\theta_* S_{T_{\Delta/a}}\right) \exp\left(\beta S_{T_{\Delta/a}}\right) \frac{u_{\theta_*}(x_0)}{u_{\theta_*}(X_{T_{\Delta/a}})} U_{T_{\Delta/a}}^\beta \right] \\ & \leq \frac{u_{\theta_*}(x_0)}{\delta} \left(E_{(0,0,x_0)}^{\theta_*} \left[\exp\left(-\theta_* S_{T_{\Delta/a}}\right) \right] \right)^{1/2} \left(E_{(0,0,x_0)}^{\theta_*} \left[U_{T_{\Delta/a}}^{2\beta} \right] \right)^{1/2}. \end{aligned} \quad (20)$$

If we are able to show that

$$E_{(0,0,x_0)}^{\theta_*} (U_{T_{\Delta/a}}^{2\beta}) = O(1) \quad (21)$$

as $\Delta \searrow 0$, then we can conclude, owing to (18), that the right hand side of (20) is of order $O(\Delta^{\theta_*/2})$. Thus, combining this bound on (20), together with (19) we would conclude that $I_2 = O(\Delta^{2\theta_*})$ as required. It suffices then to verify (21), however, this is immediate since under our current assumptions we clearly have that $U_{T_{\Delta/a}} \leq \lambda(X_{T_{\Delta/a}}, \eta_{T_{\Delta/a}}) \leq m$. ■

We shall comment on two important issues behind this result. First, we have assumed that the rewards are bounded in order to simplify our analysis. Note that the only place that used this assumption is in establishing (21). It is possible to estimate the expectation in (21)

only under Assumption 3 using a Lyapunov bound similar to the one that we will discuss in Lemma 1.

Second, the estimator $Z_{1,\Delta}$ is unbiased only if we can generate D_∞ in a finite time. Generating unbiased samples from D_∞ under our current assumptions is not straightforward (see for example Diaconis and Freedman (1999) on issues related to steady-state distributions for iterated random functions, and Blanchet and Sigman (2011) for algorithms that can be used to sample D_∞ under assumptions close to the ones that we impose here). Alternatively, one might recognize that D_∞ is the steady-state distribution of a suitably defined Markov chain. In the presence of enough regeneration structure, one can replace the indicator in (12) by an estimator for the tail of D_∞ based on the corresponding regenerative ratio representation. Note that this replacement would involve a routine simulation problem as there is no need to estimate any rare event. However, once again after using a regenerative-ratio based estimator one introduces bias.

We shall not pursue more discussion on any of the two issues raised given that the class of estimators that we shall discuss in the next section are not only unbiased but are also asymptotically optimal as $\Delta \rightarrow 0$ and can be rigorously shown to have a running time that grows at most logarithmically in $1/\Delta$.

4 State-Dependent Importance Sampling

An issue that was left open in the previous section was that the estimator that we constructed is biased from a practical standpoint. In this section, we illustrate how to construct an efficient importance sampling estimator that terminates in finite time and is unbiased. The estimator based on applying state-independent importance sampling up until time T_Δ has been seen to be inefficient. Examples of changes-of-measure that look reasonable from a large deviations perspective but at the end turn out to have a poor performance are well known in the rare-event simulation literature (see Glasserman and Kou (1995)). It is interesting that estimating the tail of D_∞ provides yet another such example. These types of examples have motivated the development of the theory behind the design of efficient state-dependent importance sampling estimators, which is the basis behind the construction of our estimator here. We shall explain some of the elements behind this theory next.

We will follow the approach based on Lyapunov inequalities (see Blanchet and Glynn (2008) and Blanchet, Glynn and Liu (2007)). Let us introduce some notation for $W_n = (S_n, D_n, X_n)$. The transition kernel associated to W is denoted by $Q(\cdot)$, so

$$P_{w_0}(W_1 \in A) = P(W_1 \in A | W_0 = w_0) = \int_A Q(w_0, dw).$$

A state-dependent importance sampling distribution for W is described by the Markov transition kernel

$$Q_r(w_0, dw_1) = r(w_0, w_1)^{-1} Q(w_0, dw_1), \tag{22}$$

where $r(\cdot)$ is a positive function properly normalized so that

$$\int Q_r(w_0, dw_1) = \int r(w_0, w_1)^{-1} Q(w_0, dw_1) = 1.$$

The idea behind the Lyapunov method is to introduce a parametric family of changes-of-measure. As we shall see, in our case, this will correspond to suitably defined exponential changes-of-measure. This selection specifies $r(\cdot)$. The associated importance sampling estimator, which is obtained by sampling transitions from $Q_r(\cdot)$, takes the form

$$Z_\Delta = r(W_0, W_1) r(W_1, W_2) \dots r(W_{T_\Delta-1}, W_{T_\Delta}) I(T_\Delta < \infty).$$

Using $P_w^{(r)}(\cdot)$ (resp. $E_w^{(r)}(\cdot)$) to denote the probability measure (resp. the expectation operator) induced by the transition kernel $Q_r(\cdot)$ given that $W_0 = w$, we can express the second moment of Z via

$$v_\Delta(w) = E_w^{(r)} Z_\Delta^2 = E_w Z_\Delta.$$

Note that conditioning on the first transition of the process W one obtains

$$v_\Delta(w) = E_w[r(w, W_1) v_\Delta(W_1)],$$

subject to the boundary condition $v_\Delta(w) = 1$ for $w \in \mathbb{R} \times [1, \infty) \times \mathcal{S}$. We are interested in a suitable upper bound for $v_\Delta(w)$, which can be obtained by taking advantage of the following inequality proved in Blanchet and Glynn (2008).

Lemma 1 *If $h_\Delta(\cdot)$ is non-negative and satisfies*

$$E_w[r(w, W_1) h_\Delta(W_1)] \leq h_\Delta(w) \tag{23}$$

subject to $h_\Delta(w) \geq 1$ for $w \in \mathbb{R} \times [1, \infty) \times \mathcal{S}$, then

$$v_\Delta(w) \leq h_\Delta(w).$$

Our strategy in state-dependent importance sampling is aligned with the intuition behind the failure of the natural state-independent importance sampling strategy described in the previous section; it consists in applying importance sampling only when it is “safe” to apply it. In other words, we wish to induce $T_\Delta < \infty$ by exponentially tilting the increments of S , but we want to be careful and maintain the likelihood ratio appropriately controlled. So, for instance, cases where D_n might be close to the boundary value 1, but S_n is significantly smaller than $\log(1/\Delta)$ are of concern. In those cases, we shall turn off importance sampling to avoid the accumulation of a large likelihood ratio in Z_Δ . In summary, suppose that the current position of the cumulative discount rate process S is given by s and that the position of the discounted process D is d . We shall continue applying exponential tilting as long as (s, d, x) belongs to some region C where it is safe to apply importance sampling. We do not apply importance sampling if $(s, d, x) \notin C$. The precise definition of the set C will be given momentarily.

Using the notation introduced earlier leading to the statement of our Lyapunov inequality in Lemma 1 we can describe the sampler as follows. Let C be an appropriately defined subset of $\mathbb{R} \times \mathbb{R} \times \mathcal{S}$. Assume that the current state of the process W is $w_0 = (s_0, d_0, x_0)$ and let us write $w_1 = (s_1, d_1, y)$ for a given outcome of the next transition. The function $r(\cdot) \triangleq r_{\theta_*}(\cdot)$ introduced in (22) takes the form

$$\begin{aligned} & r_{\theta_*}^{-1}((s_0, d_0, x_0), (s_1, d_1, y)) \\ &= I((s_0, d_0, x_0) \in C) \frac{u_{\theta_*}(y)}{u_{\theta_*}(x_0)} \exp(\theta_*(s_1 - s_0)) \\ &+ I((s_0, d_0, x_0) \notin C). \end{aligned} \tag{24}$$

The construction of an appropriate Lyapunov function $h_\Delta(\cdot)$ involves applying Lemma 1. In turn, the definition of the set C is coupled with the construction of $h_\Delta(\cdot)$. We shall construct $h_\Delta(\cdot)$ so that $h_\Delta(s, d, x) = 1$ implies $(s, d, x) \notin C$. Moreover, we shall impose the condition $h_\Delta(\cdot) \in [0, 1]$. Assuming h_Δ can be constructed in this way we immediately have that the Lyapunov inequality is satisfied outside C . We then need to construct h_Δ on C . We wish to find an asymptotically optimal change-of-measure, so it makes sense to propose

$$h_\Delta(s, d, x) = O(P_{(s,d,x)}(T_\Delta < \infty)^{2-\rho_\Delta}),$$

where $\rho_\Delta \searrow 0$ as $\Delta \searrow 0$ (recall the definition of asymptotic optimality given in the Introduction). On the other hand, we have that

$$\begin{aligned} P_{(s,d,x)}(T_\Delta < \infty) &= P_{(0,0,x)}(d + \exp(s) \Delta D_\infty > 1) \\ &= P_{(0,0,x)}\left(D_\infty > \exp(-s) \left(\frac{1-d}{\Delta}\right)\right) \\ &\approx \exp(s\theta_*) [\Delta/(1-d)]^{\theta_*}. \end{aligned} \tag{25}$$

Motivated by the form of this approximation, which is expected to hold at least in logarithmic sense as $\exp(-s)[(1-d)/\Delta] \rightarrow \infty$, we suggest a Lyapunov function of the form

$$h_\Delta(s, d, x) = \min\{c_\Delta^{2\theta_* - \rho_\Delta} \exp([2\theta_* - \rho_\Delta]s) [\Delta/(1-d)_+]^{2\theta_* - \rho_\Delta} u_{\theta_*}(x) u_{\theta_* - \rho_\Delta}(x), 1\}.$$

The introduction of the function $u_{\theta_*}(x) u_{\theta_* - \rho_\Delta}(x)$ in a multiplicative form as given above is convenient for the purpose of verifying Lyapunov inequalities for importance sampling in the setting of Markov random walks (see Blanchet, Leder and Glynn (2009)). The constant $c_\Delta > 0$, which will be specified in the verification of the Lyapunov inequality, is introduced as an extra degree of freedom to recognize that approximation (25) may not be exact. The exponent on top of c_Δ allows to make the estimates in the verification of the Lyapunov inequality somewhat cleaner. Note that we have $h_\Delta(\cdot) \in [0, 1]$ and the set C is defined via

$$C = \{(s, d, x) : h_\Delta(s, d, x) < 1\}. \tag{26}$$

We do not apply importance sampling whenever we reach a state (s, d, x) satisfying $h_\Delta(s, d, x) = 1$.

We shall close this section with a precise description of our state-dependent algorithm. The following procedure generates one sample of our estimator.

State-Dependent Algorithm

Step 1: Set $\rho_\Delta = 1/\log(1/\Delta)$ and $c_\Delta = (B_2/B_1)\rho_\Delta^{-(1+1/(2\theta^*-\rho_\Delta))}$ with $0 < B_1, B_2 < \infty$ as indicated in Proposition 1 below. Initialize $(s, d, x) \leftarrow (0, 0, x_0)$.

Step 2: Initialize likelihood ratio $L \leftarrow 1$.

Step 3: While at (s, d, x) , do the following:

1. If $(s, d, x) \in C$ defined in (26) i.e. $h_\Delta(s, d, x) < 1$,
 - i. Generate X_1 from the kernel

$$K_{\theta_*}(x, y) = K(x, y) \exp(\chi(y, \theta_*) - \psi(\theta_*)) \frac{u_{\theta_*}(y)}{u_{\theta_*}(x)}.$$

Say we have realization $X_1 = y$.

- ii. Given $X_1 = y$, sample $\gamma(y, \xi_1)$ from the exponential tilting

$$P^{\theta_*}(\gamma(y, \xi_1) \in dz) = \exp(\theta_* z - \chi(y, \theta_*)) P(\gamma(y, \xi_1) \in dz).$$

Say we have $\gamma(y, \xi_1) = z$.

- iii. Sample $\lambda(y, \eta_1)$ from the nominal distribution of η_1 given $X_1 = y$ and $\gamma(y, \xi_1) = z$.

Say we have $\lambda(y, \eta_1) = w$.

- iv. Update

$$L \leftarrow L \times \exp(-\theta_* z) \frac{u_{\theta_*}(x)}{u_{\theta_*}(y)}.$$

Else if $(s, d, x) \notin C$ i.e. $h_\Delta(s, d, x) = 1$,

- i. Sample X_1 from its nominal distribution. Given $X_1 = y$, sample $\gamma(y, \xi_1)$ and $\lambda(y, \eta_1)$ from their nominal distributions. Say the realizations are $\gamma(y, \xi_1) = y$ and $\lambda(y, \eta_1) = w$.

2. Update

$$(s, d, x) \leftarrow (s + z, d + \Delta w \exp(s + z), y).$$

3. If $d > 1$, output L and stop; else repeat the loop.

The variance analysis of the unbiased estimator L as well as the termination time of the algorithm are given in the next sections.

5 Efficiency of State-Dependent Importance Sampling

In order to verify asymptotic optimality of L we first show that $h_\Delta(\cdot)$ satisfies the Lyapunov inequality given in Lemma 1. We have indicated that the inequality is satisfied outside C . On the other hand, one clearly has that $h_\Delta(s, d, x) = 1$ for $d \geq 1$, so the boundary condition given in Lemma 1 is satisfied. Consequently, in order to show that $h_\Delta(\cdot)$ is a valid Lyapunov function and that $v_\Delta(w) \leq h_\Delta(w)$ we just have to show the following proposition.

Proposition 1 *Suppose that Assumptions 1 to 3 are in force and select $b_0 < \infty$ such that $0 < 1/[\inf_{\theta \in (0, \theta_*)} \inf_{x \in \mathcal{S}} u_\theta(x)]^2 \leq b_0$.*

i) *Select $b_1 > 0$ such that for each $\delta \in (0, \theta_*)$*

$$\exp(\psi(\theta_* - \delta)) \leq 1 - \delta\mu + b_1\delta^2$$

where $\mu = d\psi(\theta_*)/d\theta > 0$.

ii) *Pick $b_2 \in (0, \infty)$ such that*

$$\sup_{x \in \mathcal{S}, \delta \in (0, \theta_*)} E_x[\lambda(X_1, \eta_1)^{2\theta_* - \delta} \exp((\theta_* - \delta)\gamma(X_1, \xi_1))] \leq b_2$$

and make the following selection of B_1, B_2, ρ_Δ and c_Δ :

iii) *Select $0 < B_1, B_2 < \infty$ and $\rho_\Delta, c_\Delta > 0$ so that $\rho_\Delta \searrow 0, \rho_\Delta \in (0, \theta_*)$, $c_\Delta = (B_2/B_1)\rho_\Delta^{-(1+1/(2\theta_* - \rho_\Delta))}$, $B_1\rho_\Delta < 1$ and*

$$\frac{b_0 b_2 \rho_\Delta}{B_2^{2\theta_* - \rho_\Delta}} + \frac{(1 - \rho_\Delta \mu + b_1 \rho_\Delta^2)}{(1 - B_1 \rho_\Delta)^{(2\theta_* - \rho_\Delta)}} \leq 1.$$

Then, $h_\Delta(\cdot)$ satisfies the Lyapunov inequality (23) on C assuming that $r(\cdot) = r_{\theta_*}(\cdot)$ is given as in (24).

Proof. First, b_0 is finite because $u_{\theta_*}(\cdot)$ is strictly positive and \mathcal{S} is finite. The fact that the selections in i) and iii) are always possible follows from straightforward Taylor series developments. The selection of b_2 in ii) is possible because of Assumptions 2 and 3 combined with Holder's inequality. We shall prove the concluding statement in the result using the selected values in i), ii) and iii). Assume that (s, d, x) is such that $h_\Delta(s, d, x) < 1$. To ease the notation we write $\gamma_1 = \gamma(X_1, \xi_1)$ and $\lambda_1 = \lambda(X_1, \eta_1)$. We need to show that

$$\begin{aligned} E_x h_\Delta(s + \gamma_1, d + \exp(s + \gamma_1) \Delta \lambda_1, X_1) L_{\theta_*}[X_1, \gamma_1] \\ \leq h_\Delta(s, d, x), \end{aligned} \tag{27}$$

where

$$L_{\theta_*}[X_1, \gamma_1] = \frac{u_{\theta_*}(x)}{u_{\theta_*}(X_1)} \exp(-\theta_* \gamma_1).$$

We divide the expectation in (27) into two parts, namely, transitions that lie in a region that corresponds to the complement of C and transitions that lie within C . To be precise, set $a_\Delta \in (0, 1)$ and put

$$A = \{\exp(\gamma_1) \lambda_1 \geq a_\Delta \exp(-s) (1-d)/\Delta\}$$

and write A^c for the complement of A . The expectation in (27) is then equal to $J_1 + J_2$, where

$$\begin{aligned} J_1 &= E_x(h_\Delta(s + \gamma_1, d + \exp(s + \gamma_1) \Delta \lambda_1, X_1) L_{\theta_*}[X_1, \gamma_1]; A), \\ J_2 &= E_x(h_\Delta(s + \gamma_1, d + \exp(s + \gamma_1) \Delta \lambda_1, X_1) L_{\theta_*}[X_1, \gamma_1]; A^c). \end{aligned}$$

We first analyze $J_1/h_\Delta(s, d, x)$. Note that

$$\begin{aligned} \frac{J_1}{h_\Delta(s, d, x)} &\leq \frac{E_x(L_{\theta_*}[X_1, \gamma_1]; A)}{h_\Delta(s, d, x)} \\ &\leq \frac{u_{\theta_*}(x)}{\inf_{y \in \mathcal{S}} u_{\theta_*}(y)} \times \frac{E_x[\exp(-\theta_* \gamma_1); A]}{h_\Delta(s, d, x)}. \end{aligned}$$

Now we note that

$$\begin{aligned} &E_x[\exp(-\theta_* \gamma_1); A] \\ &= E_x(\exp(-\theta_* \gamma_1); \exp(-\gamma_1) \leq \lambda_1 \Delta \exp(s)/[(1-d)a_\Delta]) \\ &\leq \left(\frac{\Delta}{1-d}\right)^{\theta_*} \frac{\exp(\theta_* s)}{a_\Delta^{\theta_*}} \\ &\quad \times E_x[\lambda_1^{\theta_*}; \lambda_1 \geq \exp(-\gamma_1 - s)a_\Delta(1-d)/\Delta]. \end{aligned}$$

Moreover, applying Markov's inequality we obtain that for each $\beta > 0$

$$\begin{aligned} &E_x[\lambda_1^{\theta_*}; \lambda_1 \Delta \exp(\gamma_1 + s)/[a_\Delta(1-d)] \geq 1] \\ &\leq \frac{\Delta^\beta \exp(\beta s)}{[a_\Delta(1-d)]^\beta} E_x[\lambda_1^{\theta_* + \beta} \exp(\beta \gamma_1)]. \end{aligned}$$

Selecting $\beta = \theta_* - \rho_\Delta$ we obtain (using ii))

$$\begin{aligned} \frac{J_1}{h_\Delta(s, d, x)} &\leq \left(\frac{\Delta}{1-d}\right)^{2\theta_* - \rho_\Delta} \frac{\exp((2\theta_* - \rho_\Delta)s) u_{\theta_*}(x) E_x[\lambda_1^{2\theta_* - \rho_\Delta} \exp((\theta_* - \rho_\Delta)\gamma_1)]}{a_\Delta^{2\theta_* - \rho_\Delta} \inf_{y \in \mathcal{S}} u_{\theta_*}(y) h_\Delta(s, d, x)} \\ &\leq \frac{b_2}{(a_\Delta c_\Delta)^{2\theta_* - \rho_\Delta} \inf_{y \in \mathcal{S}} u_{\theta_*}(y) \inf_{y \in \mathcal{S}} u_{\theta_* - \rho_\Delta}(y)} \leq \frac{b_2 b_0}{(a_\Delta c_\Delta)^{2\theta_* - \rho_\Delta}}. \end{aligned}$$

To analyze J_2 we note that on

$$A^c = \{\exp(\gamma_1 + s) \lambda_1 \Delta / (1-d) < a_\Delta\}$$

we have that

$$\begin{aligned} &\frac{h_\Delta(s + \gamma_1, d + \exp(s + \gamma_1) \Delta \lambda_1, X_1)}{h_\Delta(s, d, x)} \\ &= \exp((2\theta_* - \rho_\Delta)\gamma_1) \left(1 - \frac{\Delta \exp(s + \gamma_1) \lambda_1}{1-d}\right)^{-(2\theta_* - \rho_\Delta)} \frac{u_{\theta_*}(X_1) u_{\theta_* - \rho_\Delta}(X_1)}{u_{\theta_*}(x) u_{\theta_* - \rho_\Delta}(x)} \\ &\leq \exp((2\theta_* - \rho_\Delta)\gamma_1) (1 - a_\Delta)^{-(2\theta_* - \rho_\Delta)} \frac{u_{\theta_*}(X_1) u_{\theta_* - \rho_\Delta}(X_1)}{u_{\theta_*}(x) u_{\theta_* - \rho_\Delta}(x)}. \end{aligned}$$

Therefore, we have that

$$\begin{aligned}
& \frac{J_2}{h(s, d, x)} \\
& \leq (1 - a_\Delta)^{-(2\theta_* - \rho_\Delta)} E_x \left(\exp((2\theta_* - \rho_\Delta)\gamma_1) \frac{u_{\theta_*}(X_1) u_{\theta_* - \rho_\Delta}(X_1)}{u_{\theta_*}(x) u_{\theta_* - \rho_\Delta}(x)} L_{\theta_*}[X_1, \gamma_1] \right) \\
& = (1 - a_\Delta)^{-(2\theta_* - \rho_\Delta)} E_x \left(\exp((\theta_* - \rho_\Delta)\gamma_1) \frac{u_{\theta_* - \rho_\Delta}(X_1)}{u_{\theta_* - \rho_\Delta}(x)} \right) \\
& = (1 - a_\Delta)^{-(2\theta_* - \rho_\Delta)} \exp(\psi(\theta_* - \rho_\Delta)).
\end{aligned}$$

Recall that we have assumed in ii) that b_1 is selected so that

$$\exp(\psi(\theta_* - \rho_\Delta)) \leq 1 - \rho_\Delta \mu + b_1 \rho_\Delta^2.$$

Thus we obtain

$$\frac{J_2}{h(s, d, x)} \leq (1 - a_\Delta)^{-(2\theta_* - \rho_\Delta)} (1 - \rho_\Delta \mu + b_1 \rho_\Delta^2).$$

Combining our estimates for J_1 and J_2 together we arrive at

$$\begin{aligned}
& \frac{J_1}{h(s, d, x)} + \frac{J_2}{h(s, d, x)} \\
& \leq \frac{b_0 b_2}{(a_\Delta c_\Delta)^{2\theta_* - \rho_\Delta}} + \frac{(1 - \rho_\Delta \mu + b_1 \rho_\Delta^2)}{(1 - a_\Delta)^{(2\theta_* - \rho_\Delta)}}.
\end{aligned}$$

Let $a_\Delta = B_1 \rho_\Delta$ and $c_\Delta = (B_2/B_1) \rho_\Delta^{-(1+1/(2\theta_* - \rho_\Delta))}$ for $\rho_\Delta < \theta_*$, substitute in the previous inequality and conclude that

$$\begin{aligned}
& \frac{J_1}{h(s, d, x)} + \frac{J_2}{h(s, d, x)} \\
& \leq \frac{b_0 b_2 \rho_\Delta}{B_2^{2\theta_* - \rho_\Delta}} + \frac{(1 - \rho_\Delta \mu + b_1 \rho_\Delta^2)}{(1 - B_1 \rho_\Delta)^{(2\theta_* - \rho_\Delta)}} \leq 1,
\end{aligned}$$

where the previous inequality follows by the selection of B_1 and B_2 in Assumption iii). This concludes the proof of the proposition. ■

The next result summarizes the asymptotic optimality properties of the algorithm obtained out of the previous development.

Theorem 3 *Select $\rho_\Delta = 1/\log(1/\Delta)$ and $c_\Delta = (B_2/B_1) \rho_\Delta^{-(1+1/(2\theta_* - \rho_\Delta))}$ with $0 < B_1, B_2 < \infty$ as indicated in Proposition 1. Then, the resulting estimator L obtained by the State-dependent Algorithm has a coefficient of variation of order $O(c_\Delta^{2\theta_*})$ and therefore, in particular, it is asymptotically optimal.*

Proof. The result follows as an immediate consequence of the fact that h_Δ is a valid Lyapunov function combined with Theorem 1. ■

6 Unbiasedness and Logarithmic Running Time of State-Dependent Sampler

Throughout the rest of our development, in addition to Assumptions 1 to 3, we impose the following mild technical assumption.

Assumption 4: For each $x \in \mathcal{S}$, $Var(\gamma(x, \xi)) > 0$.

The previous assumption simply says that $\gamma(x, \xi)$ is random. The assumption is immediately satisfied (given Assumption 2) in the i.i.d. case. As we shall explain a major component in our algorithm is the construction of a specific path that leads to termination. Assumption 4 is imposed in order to rule out a cyclic type behavior under the importance sampling distribution.

We shall show that the state-dependent algorithm stops with probability one (and hence avoids artificial termination which causes bias, a potential problem with the algorithm in Section 3.3) and that the expected termination time is of order $(\log(1/\Delta))^p$ for some $p > 0$. To do so let us introduce some convenient notation. Let $Z_n \triangleq (1 - D_n)e^{-S_n}$, and put $Y_n = (X_n, Z_n)$ for $n \geq 0$. The dynamics of the process $Y = (Y_n : n \geq 0)$ are such that

$$Y_{n+1} = (X_{n+1}, Z_n e^{-\gamma(X_{n+1}, \xi_{n+1})} - \Delta \lambda(X_{n+1}, \eta_{n+1})).$$

It is then easy to see that our state-dependent algorithm, which mathematically is described by equations (24) and (26), can be stated in the following way in terms of Y_n : Given Y_n ,

- Apply exponential tilting to $\gamma(X_{n+1}, \xi_{n+1})$ (using the tilting parameter θ_*) if

$$Z_n > c_\Delta \Delta u_{\theta_*}(X_n)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_n)^{1/(2\theta_* - \rho_\Delta)}$$

.

- Transition according to the nominal distribution of the system if

$$0 < Z_n \leq c_\Delta \Delta u_{\theta_*}(X_n)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_n)^{1/(2\theta_* - \rho_\Delta)}.$$

- Terminate if $Z_n \leq 0$.

Note that the region when $Z_n > c_\Delta \Delta u_{\theta_*}(X_n)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_n)^{1/(2\theta_* - \rho_\Delta)}$ corresponds to the region C (recall equation (26) in Section 4). If $0 < Z_n \leq c_\Delta \Delta u_{\theta_*}(X_n)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_n)^{1/(2\theta_* - \rho_\Delta)}$ we say that Y_n is in C' . Finally, we say that Y_n is in B if $Z_n \leq 0$. In other words, the set B is the termination set. Let us write $m = \inf_{x \in \mathcal{S}} \{u_{\theta_*}(x)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(x)^{1/(2\theta_* - \rho_\Delta)}\}$ and $M = \sup_{x \in \mathcal{S}} \{u_{\theta_*}(x)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(x)^{1/(2\theta_* - \rho_\Delta)}\}$. Note that $0 < m < M < \infty$. A key observation is that the set C' is bounded. This will help in providing bounds for the running time of the algorithm as we shall see.

We will obtain an upper bound for the algorithm by bounding the time spent by the process in both regions C and C' . Intuitively, starting from an initial position in C , the process moves to C' in finite number of steps. Then the process moves to either C or B . If the process enters C before B , then from C it again moves back to C' and the iteration between region C' and C repeats until the process finally hits B , which is guaranteed to happen by geometric trial argument. Our proof below will make the intuition rigorous and shows that the time for the process to travel each back-and-forth between C and C' is logarithmic in $1/\Delta$, and that there is a significant probability that the process starting from C' hits B before C . This, overall, will imply a logarithmic running time of the algorithm. More precisely, we will show the following lemmas. The reader should keep in mind the selections

$$\rho_\Delta = 1/\log(1/\Delta) \quad \text{and} \quad c_\Delta = (B_2/B_1)\rho_\Delta^{-(1+1/(2\theta^*-\rho_\Delta))} = O(\log(1/\Delta)^{1+1/(2\theta^*)})$$

given in Theorem 3.

Recall the notations $P^{\theta^*}(\cdot)$ and $E^{\theta^*}[\cdot]$ to denote the probability measure and expectation under the state-dependent importance sampler. Note that under $P^{\theta^*}(\cdot)$ no exponential tilting is performed when the current state Y_n lies in C' .

Lemma 2 Denote $T_{C \cup B} = \inf\{n > 0 : Y_n \in C \cup B\}$. Under Assumptions 1 to 4 we have

$$E_{y_0}^{\theta^*}[T_{C \cup B}] = O(c_\Delta^p \log c_\Delta)$$

uniformly over $y_0 \in C'$ for some constant $p > 0$.

Lemma 3 Let $T_C = \inf\{n > 0 : Y_n \in C\}$ and $T_B = \inf\{n > 0 : Y_n \in B\}$. If Assumptions 1 to 4 hold, then, uniformly over $y_0 = (x_0, z_0) \in C'$,

$$P_{y_0}^{\theta^*}(T_B < T_C) \geq \frac{c_1}{c_\Delta^p}$$

for some constant c_1 and p (the p can be chosen as the same p in Lemma 2).

Lemma 4 Denote $T_{C' \cup B} = \inf\{n > 0 : Y_n \in C' \cup B\}$ and suppose that Assumptions 1 to 4 are in force. For any $y_0 = (x_0, z_0) \in C$, we have $P_{y_0}^{\theta^*}(T_{C' \cup B} < \infty) = 1$ and

$$E_{y_0}^{\theta^*}[T_{C' \cup B}] = O\left(\log\left(\frac{z_0}{mc_\Delta\Delta}\right)\right)$$

The first lemma shows that it takes on average a logarithmic number of steps (in $1/\Delta$) for the process to reach either B or C from C' . The second lemma shows that there is a significant probability, uniformly over the initial positions in C' , that the process reaches B before C . The third lemma states that the time taken from C to C' is also logarithmic in $1/\Delta$. Lemmas 2 and 4 guarantee that each cross-border travel, either from C to C' or from

C' to C , takes on average logarithmic time. On the other hand, Lemma 3 guarantees that a geometric number of iteration, with significant probability of success, will bring the process to B from some state in C' . These will prove the following proposition on the algorithmic running time.

Proposition 2 *Suppose that Assumptions 1 to 4 hold. Then, for any $y_0 = (x_0, z_0)$, we have $P_{y_0}^{\theta_*}(T_B < \infty) = 1$ and*

$$E_{y_0}^{\theta_*}[T_B] = O\left(c_\Delta^p \log c_\Delta + \log\left(\frac{1}{c_\Delta \Delta}\right)\right)$$

for some $p > 0$.

We now give the proofs of the lemmas and Proposition 2.

Proof of Lemma 2. Our strategy to prove Lemma 2 is the following. Given any initial state $y_0 = (x_0, z_0) \in C'$, we first construct explicitly a path that takes y_0 to B within $O(\log c_\Delta)$ steps (which we call event $A(x_0)$ below), and we argue that this path happens with probability $\Omega(c_\Delta^{-p})$ for some $p > 0$. Then we look at the process in blocks of $O(\log c_\Delta)$ steps. For each block, if the process follows the particular path that leads to B , then T_B , and hence $T_{B \cup C}$, is hit; otherwise the process may have hit C , or may continue to the next block starting with some state in C' . In other words, $T_{B \cup C}$ is bounded by the time from the initial position up to the time that the process finishes following exactly the particular path in a block. We note that a successful follow of the particular path is a geometric r.v. with parameter $O(c_\Delta^{-p})$, and hence the mean of $T_{B \cup C}$ is bounded by $O(c_\Delta^{-p}) \times O(\log c_\Delta)$ and therefore the result. Now we make the previous intuition rigorous.

We first prove some elementary probabilistic bounds for $\gamma(\cdot)$ and $\lambda(\cdot)$. As in Section 4 we simplify our notation by writing $\gamma_n = \gamma(X_n, \xi_n)$ and $\lambda_n = \lambda(X_n, \xi_n)$. We first argue that for any $y = (x, z) \in C'$,

$$P_y^{\theta_*} \left(e^{-\gamma_1} \frac{u_{\theta_*}(x)^{1/\theta_*}}{u_{\theta_*}(X_1)^{1/\theta_*}} \leq u_1 \right) > 0 \tag{28}$$

for some $0 < u_1 < 1$. Note that the initial conditioning in the probability in (28) depends on y only through x .

We prove (28) by contradiction. Suppose (28) is not true, then there exists some Markov state w such that

$$P_w^{\theta_*} \left(e^{-\gamma_1} \frac{u_{\theta_*}(w)^{1/\theta_*}}{u_{\theta_*}(X_1)^{1/\theta_*}} \geq 1 \right) = 1.$$

Now if this happens and additionally

$$P_w^{\theta_*} \left(e^{-\gamma_1} \frac{u_{\theta_*}(w)^{1/\theta_*}}{u_{\theta_*}(X_1)^{1/\theta_*}} > 1 \right) > 0,$$

which obviously implies

$$P_w^{\theta_*} \left(e^{\gamma_1} \frac{u_{\theta_*}(X_1)^{1/\theta_*}}{u_{\theta_*}(w)^{1/\theta_*}} < 1 \right) > 0,$$

then

$$E_w^{\theta_*} \left[e^{\theta_* \gamma_1} \frac{u_{\theta_*}(X_1)}{u_{\theta_*}(w)} \right] < 1,$$

which contradicts the definition of θ_* . Hence we are left with the possibility that

$$P_w^{\theta_*} \left(e^{-\gamma_1} \frac{u_{\theta_*}(w)^{1/\theta_*}}{u_{\theta_*}(X_1)^{1/\theta_*}} = 1 \right) = 1,$$

but this contradicts our non-degeneracy assumption, namely, Assumption 4.

Using (28), note that we can pick $u_2 > 0$ small enough such that

$$P_y^{\theta_*} \left(e^{-\gamma_1} \frac{u_{\theta_*}(x)^{1/\theta_*}}{u_{\theta_*}(X_1)^{1/\theta_*}} \leq u_1, e^{-\gamma_1} \geq u_2 \right) > \epsilon_1 > 0 \quad (29)$$

for any $y = (x, z) \in C'$. This follows from a contradiction proof since the non-existence of u_2 would imply $e^{-\gamma_1} = 0$ a.s.

On the other hand, it is easy to see that there exists r_1 and r_2 and a small enough $u_3 > 0$ such that

$$P_{r_1}^{\theta_*}(X_1 = r_2, \lambda(r_2, \eta_1) \geq u_3) > \epsilon_2 > 0 \quad (30)$$

since otherwise $\lambda_1 = 0$ a.s.

We will now construct the path $A(x_0)$ as discussed earlier in the proof. This path will depend on the initial position $y_0 = (x_0, z_0) \in C'$, but it has length $O(\log c_\Delta)$ uniformly over any initial position in C' . The path has the property that whenever the Markov state hits r_1 , it would go to r_2 with $\lambda(r_2, \eta_n) \geq u_3$ in the next state. Moreover, for every step, $e^{-\gamma(X_n, \xi_n)} u_{\theta_*}(X_{n-1})^{1/\theta_*} / u_{\theta_*}(X_n)^{1/\theta_*} \leq u_1$ and $e^{-\gamma_n} \geq u_2$. The path evolves in a periodic way i.e. it hits r_1 in every l steps for N times, where N is a number to be determined later. The existence of l and the occurrence of such periodic cycles with positive probability is guaranteed by the irreducibility of the Markov chain X_n . In other words, consider the event $A(x_0)$ given by

$$A(x_0) = \left\{ \begin{array}{l} \text{for } k = 1, \dots, N, X_{a+kl} = r_1, X_{a+kl+1} = r_2, \lambda(r_2, \eta_{a+kl+1}) \geq u_3; \\ \text{for } i = 1, \dots, a + Nl + b, e^{-\gamma(X_i, \xi_i)} \frac{u_{\theta_*}(X_{i-1})^{1/\theta_*}}{u_{\theta_*}(X_i)^{1/\theta_*}} \leq u_1, e^{-\gamma(X_i, \xi_i)} \geq u_2; \\ X_{a+Nl+b} = x_0 \end{array} \right\}$$

where a is the number of steps for the initial state x_0 to reach r_1 and b is the number of steps for the last hit on r_2 back to state x_0 . Note that a and b all depend on x_0 , but we suppress the dependence for notational convenience. N is an integer that we will pick momentarily.

Under $A(x_0)$ we have

$$\begin{aligned}
Z_{a+Nl+b} &= ze^{-\gamma_1-\dots-\gamma_{a+Nl+b}} - \Delta\lambda_1e^{-\gamma_2-\dots-\gamma_{a+Nl+b}} - \Delta\lambda_2e^{-\gamma_3-\dots-\gamma_{a+Nl+b}} - \dots - \Delta\lambda_{a+Nl+b} \\
&\leq zu_1^{a+Nl+b} - \Delta u_3(u_2^{a+Nl} + u_2^{a+(N-1)l} + \dots + u_2^a) \\
&\leq zu_1^{a+Nl+b} - \Delta u_3 u_2^a \frac{1 - u_2^{(N+1)l}}{1 - u_2}
\end{aligned}$$

Now pick N to be the smallest integer at least as large as

$$\frac{\log((u_1^{a+b} M c_\Delta \Delta (1 - u_2) + \Delta u_3 u_2^{a+l}) / (\Delta u_3 u_2^a))}{l \log(1/u_1)}$$

This implies that

$$N < \frac{\log((z_0 u_1^{a+b} (1 - u_2) + \Delta u_3 u_2^{a+l}) / (\Delta u_3 u_2^a))}{l \log(1/u_1)} + 1$$

(note the definition of C' and M above) and a simple verification reveals that $Z_{a+Nl+b} \leq 0$ on $A(x_0)$. Hence if $A(x_0)$ occurs, then T_B is hit before step $a + Nl + b$. Note that $N = O(\log c_\Delta)$.

Now note that given $y_0 = (x_0, z_0) \in C'$, the probability that $A(x_0)$ happens is larger than q^{a+Nl+b} for some $q > 0$. If we divide the steps of the chain into blocks of size $r + Nl$, where $r = \max_{x \in \mathcal{S}} \{a(x) + b(x)\}$, then the number of blocks required for Z_n to hit 0 (and hence $T_{B \cup C}$ is achieved) is bounded by a geometric r.v. with parameter q^{r+Nl} . Taking also into account the length of the blocks, we have

$$E_{y_0}^{\theta_*} T_{B \cup C} \leq \frac{1}{q^{r+Nl}} (r + Nl) = O(c_\Delta^p \log c_\Delta)$$

for some $p > 0$. ■

Proof of Lemma 3. Given an initial position $y_0 = (x_0, z_0) \in C'$. It suffices to show that the path $A(x_0)$ we have constructed in the proof of Lemma 2 does not hit T_C before T_B i.e. it does not hit T_C for every step up through $a + Nl + b$. The conclusion of Lemma 3 then follows by noting that $P_{y_0}^{\theta_*}(T_B > T_C) \geq P_{y_0}^{\theta_*}(A(x_0))$. To prove T_C is not hit for every step, we show that $Z_n < c_\Delta \Delta u_{\theta_*}(X_n)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_n)^{1/(2\theta_* - \rho_\Delta)}$ i.e. $Z_n \in B \cup C'$, for every $n = 1, \dots, a + Nl + b$ by induction. Suppose $Z_n < c_\Delta \Delta u_{\theta_*}(X_n)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_n)^{1/(2\theta_* - \rho_\Delta)}$, then

$$\begin{aligned}
Z_{n+1} &= Z_n e^{-\gamma_{n+1}} - \Delta \lambda_{n+1} \\
&\leq c_\Delta \Delta u_{\theta_*}(X_n)^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_n)^{1/(2\theta_* - \rho_\Delta)} \cdot u_1 \frac{u_{\theta_*}(X_{n+1})^{1/\theta_*}}{u_{\theta_*}(X_n)^{1/\theta_*}} \\
&< c_\Delta \Delta u_{\theta_*}(X_{n+1})^{1/(2\theta_* - \rho_\Delta)} u_{\theta_* - \rho_\Delta}(X_{n+1})^{1/(2\theta_* - \rho_\Delta)}
\end{aligned}$$

for small enough Δ , where u_1 is defined in (28), by choosing the eigenvectors $u_{\theta_* - \delta}(x)$ that are continuous in δ within a small neighborhood of 0 uniformly over all $x \in \mathcal{S}$. Hence we have proved our claim. ■

Proof of Lemma 4. Suppose we start at $y_0 = (x_0, z_0) \in C$. Consider $\tilde{\gamma}_n = \sum_{i=\tau_{n-1}}^{\tau_n} \gamma_i$ where $\tau_n = \inf\{i > \tau_{n-1} : X_i = x_0\}$. Inside region C the random walk $\tilde{S}_n = \sum_{j=1}^n \tilde{\gamma}_j$ has positive drift i.e. $E\tilde{\gamma}_n > 0$. For the process to hit $C' \cup B$, it suffices to have

$$z_0 e^{-\gamma_1 - \dots - \gamma_n} - \Delta \lambda_1 e^{-\gamma_2 - \dots - \gamma_n} - \dots - \Delta \lambda_n \leq mc_\Delta \Delta$$

or equivalently

$$\Delta(\lambda_n + mc_\Delta \Delta) e^{\gamma_1 + \dots + \gamma_n} + \Delta \lambda_{n-1} e^{\gamma_1 + \dots + \gamma_{n-1}} + \dots + \Delta \lambda_1 e^{\gamma_1} \geq z_0$$

This will be implied by the condition $\Delta(\lambda_n + mc_\Delta \Delta) e^{\gamma_1 + \dots + \gamma_n} \geq z_0$, which in turn can be achieved if $S_n \geq \log(z_0 / (mc_\Delta \Delta))$. Note that if we only consider the steps when $X_n = x_0$, then the condition becomes $\tilde{S}_n \geq \log(z_0 / (mc_\Delta \Delta))$ where \tilde{S}_n is now a positively drifted random walk. This happens with probability one and the expected time for this to happen is $O\left(\log\left(\frac{z_0}{mc_\Delta \Delta}\right)\right)$, which provides an upper bound for $E_{y_0}^{\theta_*}[T_{C' \cup B}]$. ■

Proof of Proposition 2. Consider an initial position at $y_0 = (x_0, z_0) \in C$ (if $y_0 = (x_0, z_0) \in C'$ the same analysis goes through resulting in a shorter mean running time). With probability one Y_n will enter $C' \cup B$ by Lemma 4. If $T_{C'} < T_B$ then by Lemma 3 the process hits B with probability that is bounded away from zero uniformly over $Y_{T_{C'}}$, otherwise it goes back to C . Hence by geometric trial argument the process will hit B eventually. We obtain the first part of the proposition.

We now consider $E_{y_0}^{\theta_*} T_B$. Suppose first that $y_0 = (x_0, z_0) \in C'$. Write

$$E_{y_0}^{\theta_*} T_B = E_{y_0}^{\theta_*}[T_B; T_B < T_C] + E_{y_0}^{\theta_*}[T_B; T_B > T_C]$$

Let $\bar{T}_B = T_B - T_C$ on the set $T_B > T_C$ i.e. \bar{T}_B is the residual time to hit B once T_C is first hit. We can write

$$E_{y_0}^{\theta_*} T_B = E_{y_0}^{\theta_*} T_{B \cup C} + E_{y_0}^{\theta_*}[\bar{T}_B; T_B > T_C] = E_{y_0}^{\theta_*} T_{B \cup C} + E_{y_0}^{\theta_*}[E_{Y_{T_C}(y_0)}^{\theta_*} \bar{T}_B; T_B > T_C]$$

where $Y_{T_C}(y_0)$ is the state at time T_C (the y_0 as a parameter emphasizes the dependence on the initial position y_0). We further write

$$E_{y_0}^{\theta_*} T_B = E_{y_0}^{\theta_*} T_{B \cup C} + E_{y_0}^{\theta_*}[E_{Y_{T_C}(y_0)}^{\theta_*}[\bar{T}_{B \cup C'} + E_{Y_{T_C'}}^{\theta_*}[\bar{\bar{T}}_B; \bar{T}_B > \bar{T}_{C'}]]; T_B > T_C] \quad (31)$$

where $\bar{\bar{T}}_B = \bar{T}_B - \bar{T}_{C'}$ on the set $\bar{T}_B > \bar{T}_{C'}$.

Let $f(y) = E_y^{\theta_*} T_B$. (31) leads to

$$\begin{aligned} f(y_0) &\leq E_{y_0}^{\theta_*} T_{B \cup C} + E_{y_0}^{\theta_*}[E_{Y_{T_C}(y_0)}^{\theta_*} \bar{T}_{B \cup C'}; T_B > T_C] + \sup_{w \in C'} f(w) P_{y_0}^{\theta_*}(T_B > T_C) \\ &\leq E_{y_0}^{\theta_*} T_{B \cup C} + c E_{y_0}^{\theta_*} \left[\log\left(\frac{Y_{T_C}(y_0)}{mc_\Delta \Delta}\right); T_B > T_C \right] + \sup_{w \in C'} f(w) P_{y_0}^{\theta_*}(T_B > T_C) \end{aligned} \quad (32)$$

where $c > 0$ is a constant, using Lemma 4. Now consider

$$Y_{T_C}(y_0) = z_0 e^{-\gamma_1 - \dots - \gamma_{T_C}} - \Delta \lambda_1 e^{-\gamma_2 - \dots - \gamma_{T_C}} - \dots - \Delta \lambda_{T_C} \leq z_0 e^{-\gamma_1 - \dots - \gamma_{T_C}}$$

and hence

$$\log \left(\frac{Y_{T_C}(y_0)}{m c_\Delta \Delta} \right) \leq \log z_0 - \gamma_1 - \dots - \gamma_{T_C} - \log(m c_\Delta \Delta)$$

Now

$$\begin{aligned} E_{y_0}^{\theta_*}[-\gamma_1 - \dots - \gamma_{T_C}; T_B > T_C] &\leq E_{y_0}^{\theta_*}[|\gamma_1| + \dots + |\gamma_{T_{B \cup C}}|; T_B > T_C] \\ &\leq E_{y_0}^{\theta_*}[|\gamma_1| + \dots + |\gamma_{T_{B \cup C}}|] \\ &\leq \tilde{c} E_{y_0}^{\theta_*} T_{B \cup C} \end{aligned}$$

where $\tilde{c} = \sup_{x \in \mathcal{S}} E^{\theta_*}[|\gamma(X_1, \xi_1)| | X_1 = x] < \infty$, by Wald's identity and Assumption 1 in Section 2. This gives

$$E_{y_0}^{\theta_*} \left[\log \left(\frac{Y_{T_C}(y_0)}{m c_\Delta \Delta} \right); T_B > T_C \right] \leq \log z_0 P_{y_0}^{\theta_*}(T_B > T_C) + \tilde{c} E_{y_0}^{\theta_*} T_{B \cup C} - \log(m c_\Delta \Delta) P_{y_0}^{\theta_*}(T_B > T_C) \quad (33)$$

Putting (33) into (32) and using the fact that $z_0 \leq M c_\Delta \Delta$ for $y_0 = (x_0, z_0) \in C'$ yields

$$\begin{aligned} f(y_0) &\leq E_{y_0}^{\theta_*} T_{B \cup C} + c P_{y_0}^{\theta_*}(T_B > T_C) \log(M c_\Delta \Delta) + \tilde{c} E_{y_0}^{\theta_*} T_{B \cup C} - c \log(m c_\Delta \Delta) P_{y_0}^{\theta_*}(T_B > T_C) \\ &\quad + \sup_{w \in C'} f(w) P_{y_0}^{\theta_*}(T_B > T_C) \end{aligned}$$

Now taking supremum on both sides and using Lemma 2 and 3, we get

$$\sup_{w \in C'} f(w) \leq O \left(c_\Delta^p \log c_\Delta + \log \left(\frac{1}{c_\Delta \Delta} \right) \frac{1}{c_\Delta^p} \right)$$

Suppose we start with $y_0 \in C$, then combining with Lemma 4 concludes the proposition. ■

7 Extensions and Remarks

As noted in equation (5) the perpetuity D satisfies the distributional fixed point equation

$$D =_d B + AD',$$

where $A = \exp(Y)$, D' has the same distribution as D , and (A, B) is independent of D' . We wish to compare our development to that of Collamore et al (2011), which state their results in the absence of Markov modulation, so to make the comparison more transparent we will omit the Markov chain $\{X_n\}$ from our discussion here. We have assumed that B is non-negative but we believe that this is not a strong assumption. We can typically reduce to the case of non-negative B . Indeed, if the B_i 's can take negative values, we can let $\tilde{B}_i = |B_i|$ and define

$$\tilde{D} = \tilde{B}_1 + \exp(Y_1) \tilde{B}_2 + \dots \quad (34)$$

Note that the tail of \tilde{D} is in great generality equivalent (up to a constant) to that of D . Since $\{D > 1/\Delta\} \subseteq \{\tilde{D} > 1/\Delta\}$ we can use the likelihood ratio constructed to estimate the tail of \tilde{D} but apply it to the event $I(D > 1/\Delta)$. The same efficiency analysis applies automatically. So, throughout our discussion we shall keep assuming that B is non-negative.

Now, one could consider more general fixed point equations, for instance,

$$D =_d B + A \max(D', C) \tag{35}$$

where D' has the same distribution as D , (A, B, C) are independent of D' . This equation is the focus of Collamore et al (2011). Their assumptions are similar to ours, primarily that $\theta_* > 0$ satisfying $EA^{\theta_*} = 1$ can be computed; that suitable moment conditions are satisfied for B and C , and that the associated exponentially tilted distributions can be simulated. Their estimator for the tail of D is biased, but it enjoys asymptotic optimality properties parallel to strong efficiency. In this sense this estimator is close in spirit to our state-independent importance sampler. However, their construction is completely different to ours, as we shall explain now.

Equation (35) characterizes the steady-state distribution (if it exists) of the Markov chain, $\{V_n : n \geq 0\}$, defined via $V_0 = v_0$ and

$$V_{n+1} = B_{n+1} + A_{n+1} \max(V_n, C_{n+1}), \tag{36}$$

where $\{(A_n, B_n, C_n) : n \geq 1\}$ is an i.i.d. sequence. Collamore et al (2011) uses a regenerative ratio representation for the steady-state distribution of $\{V_n\}$, assuming suitable minorization conditions required for regeneration are in place. Clearly, there are advantages to simulating V_n (which is Markovian) as opposed to the "backward" process, which corresponds to the discounted reward process (whose limit is the perpetuity and it is not Markovian but requires keeping track of $S_n = Y_1 + \dots + Y_n$). Nevertheless, one of the important features of importance sampling is that it can be applied to estimate conditional expectations of sample path functions given the event of interest (in this case $D > 1/\Delta$). This is also why we wanted our algorithms to be developed under the presence of Markovian modulation, without resorting to a decomposition such as (5).

This feature, we believe, is quite attractive specially in some of the applications behind our motivation to study discounted process, such as insurance and finance. The problem with using the "forward" representation (i.e. $\{V_n\}$ and the associated regenerative ratio) is that, while the tail estimation of D is preserved, it is difficult to use the associated algorithms for estimation of conditional sample path expectations.

Finally, we point out that a similar coupling idea to the one used in the construction of (34) can be applied to reduce the analysis of (35) to the case of standard perpetuities. In particular, define $\tilde{B}_{n+1} = B_{n+1} + A_{n+1}C_{n+1}$, and plug this definition into (34) to define \tilde{D} . Then we have that $\tilde{D} \geq D$, where D is the limit of the "backward" representation associated to (35). Since \tilde{D} and D are typically tail equivalent (except for a constant), again we can proceed as indicated earlier. Because $\{D > 1/\Delta\} \subseteq \{\tilde{D} > 1/\Delta\}$ we can use the likelihood ratio constructed to estimate the tail of \tilde{D} but apply it to the event $I(D > 1/\Delta)$. Again,

bias is introduced because of the infinite horizon nature of D , but the rare-event simulation problem has been removed by the importance sampling strategy constructed based on \tilde{D} .

8 Numerical Experiments

We run our algorithm for the ARCH(1) sequence in Example 1 with $\alpha_0 = 1, 2$ and $\alpha_1 = 3/4, 4/5$. In this example there is no Markov modulation. Using the transformation into T_n as shown in Example 1, the tail probability of the steady-state distribution of the ARCH(1) process with target level $1/\Delta$ is equivalent to the tail probability of a perpetuity with $\lambda(X_i, \eta_i) = \alpha_0, \gamma(X_i, \xi_i) = \log \alpha_1 + \log \chi_i^2$ where χ_i^2 are i.i.d. chi-square r.v.'s, and target level α_1/Δ . One can compute easily that

$$Ee^{\theta\gamma(X_i, \xi_i)} = (2\alpha_1)^\theta \frac{\Gamma(\theta + 1/2)}{\Gamma(1/2)}.$$

and hence verify that Assumptions 1, 2 and 3 are satisfied. Moreover, the conditions in Proposition 1 are also satisfied with appropriate selection of parameters (see the discussion below). Our choices of α_1 would correspond to θ_* with values 1.68, 1.46 and 1.34 respectively. This implies a tail of the steady-state ARCH(1) model that has finite third moment but not the fourth, which frequently arises in the financial context (see, for example, Mikosch and Starica (2000)).

We test the performance of both our state-independent and state-dependent sampler proposed in Sections 3 and 4 by comparing with crude Monte Carlo. To gauge the performance of our algorithms as Δ becomes small, we tune Δ from 0.1 to 0.00001 to see the effect of the magnitude of Δ to the output performance.

For crude Monte Carlo, we truncate the maximum number of steps to be 1000 (so that the sequence does not iterate indefinitely; note that this would certainly cause bias in the sample).

In the case of the state-independent importance sampler, we use $a = 9/10$ and $n_* = 10 \log(1/\Delta)$ (where a is the proportion of the barrier that upon touching would lead to the stop of importance sampling and n_* is the number of steps we continue to simulate after $T_{\Delta/a}$).

For the state-dependent sampler, we can verify that $b_0 = 1, b_1 = \sup_{0 \leq \zeta \leq \theta_*} (\psi''(\zeta) + (\psi'^2(\zeta)))/2$ and $b_2 = \max\{\alpha_0^{2\theta_*}, 1\}$ satisfy the conditions in Proposition 2. To ensure that the Lyapunov inequality holds for small Δ one can choose B_1 and B_2 in Proposition 2 to satisfy $B_2 \geq 1$ and

$$\mu - 2B_1\theta_* - \frac{b_0b_2}{B_2^{\theta_*}} > 0$$

In particular we can choose $B_1 = 0.45\mu/(2\theta_*)$ and $B_2 = \max\{(b_0b_2/(0.45\mu))^{1/\theta_*}, 1\}$.

For each set of input parameters (i.e. α_0 , α_1 and Δ) we simulate using crude Monte Carlo, the state-independent sampler, and the state-dependent sampler. For each method we fix the running time to be five minutes for comparison. In the following tables we show the estimate, empirical coefficient of variation (standard deviation divided by the estimate), and 95% confidence interval for our simulation. Tables are deferred to the appendix below.

Next we also run the algorithm for a Markov modulated perpetuity. The modulating Markov chain lies in state space $\{1, 2\}$ and has transition matrix

$$K = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ 1 & 0 \end{bmatrix}$$

We use $\lambda(1, \eta_1) = \lambda(1) = 1$ and $\lambda(2, \eta_1) = \lambda(2) = 2$ a.s.. Also, $\gamma(1, \xi_1) = \log(2/3) + \log \chi_i^2$ and $\gamma(2, \xi_1) = \log(3/4) + \log \chi_i^2$, where again χ_i^2 are i.i.d. chi-square random variables.

Again we experiment using crude Monte Carlo and both state-independent and dependent importance samplers. Similar to the ARCH(1) setup, a simple calculation reveals that $e^{\chi(1, \theta)} = (2 \times (2/3))^{2\theta} \Gamma(\theta + 1/2) / \Gamma(1/2)$ and $e^{\chi(2, \theta)} = (2 \times (3/4))^{2\theta} \Gamma(\theta + 1/2) / \Gamma(1/2)$. Moreover, $e^{\psi(\theta)} = (e^{\chi(1, \theta)} + \sqrt{e^{2\chi(1, \theta)} + 8e^{\chi(1, \theta) + \chi(2, \theta)}}) / 4$, so $\theta_* = 1.60$. We take $u_\theta(1) = (\sqrt{e^{2\chi(1, \theta)} + 8e^{\chi(1, \theta) + \chi(2, \theta)}} + e^{\chi(1, \theta)}) / (4e^{\chi(1, \theta)})$ and $u_\theta(2) = 0$.

The computational effort needed for this Markov-modulated problem appears to be substantially heavier than the case of ARCH model, and hence we perform a longer and more extensive simulation study. We tune Δ from 0.1 to 0.002, and for each scenario we run the simulation for one hour for each method. For crude Monte Carlo, we use 100,000 as our step truncation. For state-independent importance sampler we use $a = 9/10$ and $n_* = 1000 \log(1/\Delta)$. For state-dependent importance sampler, we choose $b_0 = 1$, $b_1 = \sup_{\zeta \in (0, \theta_*)} (\psi''(\zeta) + \psi'^2(\zeta)) / 2$, and $b_2 = \max\{\sup_{x \in \mathcal{S}} \lambda(x)^{2\theta_*}, 1\} \sup_{x \in \mathcal{S}} E_x e^{\chi(X_1, \theta_*)}$. Then taking $B_1 = 0.45\mu / (2\theta_*)$ and $B_2 = \max\{(b_0 b_2 / (0.45\mu))^{1/\theta_*}, 1\}$ will satisfy the Lyapunov inequality for small enough Δ . The numerical outputs are shown in the appendix below.

For the ARCH model, it is notable from the coefficient of variation and confidence interval that both state-independent and state-dependent samplers perform better than crude Monte Carlo starting from $\Delta = 0.001$. Crude Monte Carlo has much larger coefficient of variation when Δ is 0.0005, and it merely fails (i.e. does not generate any positive sample) when Δ is 0.00001. On the other hand, the coefficient of variation for state-independent sampler remains at around 1 to 2 and that for state-dependent sampler remains under 50 for all the cases we considered. The state-independent sampler appears to perform better than state-dependent sampler for our range of Δ , although one should keep in mind there is bias issue in that algorithm. Similar results hold for the Markov-modulated perpetuity, where crude Monte Carlo fails completely when Δ is 0.005 or larger while the importance samplers still perform reasonably well.

8.1 Appendix: Numerical Output

ARCH model, parameter values: $\alpha_0 = 1$, $\alpha_1 = 3/4$

Crude Monte Carlo

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	6.65×10^{-2}	3.75	$[6.43 \times 10^{-2}, 6.86 \times 10^{-2}]$
$\Delta = 0.05$	2.89×10^{-2}	5.80	$[2.74 \times 10^{-2}, 3.04 \times 10^{-2}]$
$\Delta = 0.001$	1.11×10^{-4}	95.05	$[1.37 \times 10^{-5}, 2.08 \times 10^{-4}]$
$\Delta = 0.0005$	2.11×10^{-5}	217.9	$[-2.02 \times 10^{-5}, 6.24 \times 10^{-5}]$
$\Delta = 0.00001$	0	N/A	N/A

State-Independent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	6.84×10^{-2}	1.79	$[6.82 \times 10^{-2}, 6.86 \times 10^{-2}]$
$\Delta = 0.05$	2.84×10^{-2}	1.76	$[2.83 \times 10^{-2}, 2.85 \times 10^{-2}]$
$\Delta = 0.001$	1.10×10^{-4}	1.75	$[1.09 \times 10^{-4}, 1.10 \times 10^{-4}]$
$\Delta = 0.0005$	4.01×10^{-5}	1.79	$[3.99 \times 10^{-5}, 4.02 \times 10^{-5}]$
$\Delta = 0.00001$	1.34×10^{-7}	1.72	$[1.33 \times 10^{-7}, 1.35 \times 10^{-7}]$

State-Dependent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	6.80×10^{-2}	3.68	$[6.63 \times 10^{-2}, 6.96 \times 10^{-2}]$
$\Delta = 0.05$	2.82×10^{-2}	5.81	$[2.67 \times 10^{-2}, 2.96 \times 10^{-2}]$
$\Delta = 0.001$	1.45×10^{-4}	24.77	$[7.70 \times 10^{-5}, 2.14 \times 10^{-4}]$
$\Delta = 0.0005$	4.22×10^{-5}	29.93	$[1.49 \times 10^{-5}, 6.95 \times 10^{-5}]$
$\Delta = 0.00001$	2.48×10^{-7}	37.71	$[-5.43 \times 10^{-8}, 5.49 \times 10^{-7}]$

ARCH model, parameter values: $\alpha_0 = 2$, $\alpha_1 = 3/4$

Crude Monte Carlo

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	1.51×10^{-1}	2.37	$[1.48 \times 10^{-1}, 1.54 \times 10^{-1}]$
$\Delta = 0.05$	6.64×10^{-2}	3.75	$[6.40 \times 10^{-2}, 6.88 \times 10^{-2}]$
$\Delta = 0.001$	2.61×10^{-4}	61.92	$[1.13 \times 10^{-4}, 4.08 \times 10^{-4}]$
$\Delta = 0.0005$	8.19×10^{-5}	110.5	$[1.64 \times 10^{-6}, 1.62 \times 10^{-4}]$
$\Delta = 0.00001$	0	N/A	N/A

State-Independent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	1.50×10^{-1}	1.92	$[1.495 \times 10^{-1}, 1.503 \times 10^{-1}]$
$\Delta = 0.05$	6.85×10^{-2}	2.48	$[6.83 \times 10^{-2}, 6.88 \times 10^{-2}]$
$\Delta = 0.001$	3.00×10^{-4}	1.87	$[2.99 \times 10^{-4}, 3.02 \times 10^{-4}]$
$\Delta = 0.0005$	1.09×10^{-4}	1.69	$[1.09 \times 10^{-4}, 1.10 \times 10^{-4}]$
$\Delta = 0.00001$	3.69×10^{-7}	1.69	$[3.67 \times 10^{-7}, 3.71 \times 10^{-7}]$

State-Dependent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	1.50×10^{-1}	2.38	$[1.46 \times 10^{-1}, 1.54 \times 10^{-1}]$
$\Delta = 0.05$	6.92×10^{-2}	3.66	$[6.57 \times 10^{-2}, 7.27 \times 10^{-2}]$
$\Delta = 0.001$	5.54×10^{-4}	42.46	$[-2.14 \times 10^{-4}, 1.32 \times 10^{-3}]$
$\Delta = 0.0005$	1.61×10^{-5}	42.72	$[-8.85 \times 10^{-6}, 4.11 \times 10^{-5}]$
$\Delta = 0.00001$	9.13×10^{-9}	34.65	$[-8.78 \times 10^{-9}, 2.70 \times 10^{-8}]$

ARCH model, parameter values: $\alpha_0 = 1$, $\alpha_1 = 4/5$

Crude Monte Carlo

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	7.79×10^{-2}	3.44	$[7.45 \times 10^{-2}, 8.14 \times 10^{-2}]$
$\Delta = 0.05$	3.53×10^{-2}	5.23	$[3.29 \times 10^{-2}, 3.76 \times 10^{-2}]$
$\Delta = 0.001$	2.27×10^{-4}	66.39	$[2.80 \times 10^{-5}, 4.26 \times 10^{-4}]$
$\Delta = 0.0005$	9.13×10^{-5}	104.6	$[-3.52 \times 10^{-5}, 2.18 \times 10^{-4}]$
$\Delta = 0.00001$	0	N/A	N/A

State-Independent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	7.78×10^{-2}	1.72	$[7.75 \times 10^{-2}, 7.81 \times 10^{-2}]$
$\Delta = 0.05$	3.43×10^{-2}	1.56	$[3.41 \times 10^{-2}, 3.44 \times 10^{-2}]$
$\Delta = 0.001$	2.02×10^{-4}	1.57	$[2.01 \times 10^{-4}, 2.03 \times 10^{-4}]$
$\Delta = 0.0005$	8.00×10^{-5}	1.53	$[7.96 \times 10^{-5}, 8.05 \times 10^{-5}]$
$\Delta = 0.00001$	4.21×10^{-7}	1.55	$[4.18 \times 10^{-7}, 4.24 \times 10^{-7}]$

State-Dependent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	7.84×10^{-2}	3.39	$[7.62 \times 10^{-2}, 8.07 \times 10^{-2}]$
$\Delta = 0.05$	3.47×10^{-2}	5.21	$[3.27 \times 10^{-2}, 3.67 \times 10^{-2}]$
$\Delta = 0.001$	1.31×10^{-4}	27.97	$[4.61 \times 10^{-5}, 2.16 \times 10^{-4}]$
$\Delta = 0.0005$	6.52×10^{-5}	23.15	$[2.71 \times 10^{-5}, 1.03 \times 10^{-4}]$
$\Delta = 0.00001$	8.09×10^{-7}	26.51	$[-6.15 \times 10^{-9}, 1.62 \times 10^{-6}]$

ARCH model, parameter values: $\alpha_0 = 2$, $\alpha_1 = 4/5$

Crude Monte Carlo

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	1.59×10^{-1}	2.30	$[1.55 \times 10^{-1}, 1.64 \times 10^{-1}]$
$\Delta = 0.05$	7.96×10^{-2}	3.40	$[7.62 \times 10^{-2}, 8.30 \times 10^{-2}]$
$\Delta = 0.001$	4.49×10^{-4}	47.20	$[1.71 \times 10^{-4}, 7.27 \times 10^{-4}]$
$\Delta = 0.0005$	2.69×10^{-4}	60.92	$[5.39 \times 10^{-5}, 4.85 \times 10^{-4}]$
$\Delta = 0.00001$	0	N/A	N/A

State-Independent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	1.62×10^{-1}	1.67	$[1.61 \times 10^{-1}, 1.62 \times 10^{-1}]$
$\Delta = 0.05$	7.74×10^{-2}	1.57	$[7.71 \times 10^{-2}, 7.77 \times 10^{-2}]$
$\Delta = 0.001$	5.12×10^{-4}	1.57	$[5.10 \times 10^{-4}, 5.15 \times 10^{-4}]$
$\Delta = 0.0005$	2.03×10^{-4}	1.74	$[2.01 \times 10^{-4}, 2.04 \times 10^{-4}]$
$\Delta = 0.00001$	1.07×10^{-6}	1.59	$[1.06 \times 10^{-6}, 1.08 \times 10^{-6}]$

State-Dependent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	1.61×10^{-1}	2.28	$[1.56 \times 10^{-1}, 1.65 \times 10^{-1}]$
$\Delta = 0.05$	7.77×10^{-2}	3.44	$[7.33 \times 10^{-2}, 8.22 \times 10^{-2}]$
$\Delta = 0.001$	4.32×10^{-4}	44.24	$[-2.64 \times 10^{-4}, 1.13 \times 10^{-3}]$
$\Delta = 0.0005$	1.39×10^{-4}	23.65	$[6.72 \times 10^{-6}, 2.71 \times 10^{-4}]$
$\Delta = 0.00001$	5.93×10^{-7}	32.41	$[-5.70 \times 10^{-7}, 1.76 \times 10^{-6}]$

Markov-modulated perpetuity

Crude Monte Carlo

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	7.23×10^{-2}	3.58	$[5.64 \times 10^{-2}, 8.82 \times 10^{-2}]$
$\Delta = 0.05$	2.61×10^{-2}	6.12	$[1.60 \times 10^{-2}, 3.62 \times 10^{-2}]$
$\Delta = 0.02$	3.09×10^{-3}	17.98	$[-4.07 \times 10^{-4}, 6.58 \times 10^{-3}]$
$\Delta = 0.005$	0	N/A	N/A
$\Delta = 0.002$	0	N/A	N/A

State-Independent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	5.82×10^{-2}	45.89	$[5.55 \times 10^{-2}, 6.10 \times 10^{-2}]$
$\Delta = 0.05$	2.08×10^{-2}	10.45	$[2.04 \times 10^{-2}, 2.11 \times 10^{-2}]$
$\Delta = 0.02$	5.24×10^{-3}	14.67	$[5.13 \times 10^{-3}, 5.34 \times 10^{-3}]$
$\Delta = 0.005$	5.92×10^{-4}	11.25	$[5.73 \times 10^{-4}, 6.10 \times 10^{-4}]$
$\Delta = 0.002$	1.37×10^{-4}	9.25	$[1.34 \times 10^{-4}, 1.40 \times 10^{-4}]$

State-Dependent Sampler

	Estimate	C.V.	95% C.I.
$\Delta = 0.1$	5.73×10^{-2}	4.05	$[5.35 \times 10^{-2}, 6.10 \times 10^{-2}]$
$\Delta = 0.05$	2.23×10^{-2}	6.62	$[1.88 \times 10^{-2}, 2.58 \times 10^{-2}]$
$\Delta = 0.02$	3.51×10^{-3}	16.83	$[1.74 \times 10^{-3}, 5.27 \times 10^{-3}]$
$\Delta = 0.005$	4.40×10^{-4}	47.68	$[-4.23 \times 10^{-4}, 1.30 \times 10^{-3}]$
$\Delta = 0.002$	2.35×10^{-5}	44.40	$[-2.26 \times 10^{-5}, 6.96 \times 10^{-5}]$

References

1. Adler, R., Blanchet, J., and Liu, J. C. (2010) Fast Simulation of Gaussian Random Fields with High Excursions. *Preprint*.
2. Asmussen (2000), *Ruin Probabilities*, World Scientific.
3. Asmussen (2003), *Applied Probability and Queues*, Second Edition, Springer.
4. Asmussen, S. and Glynn, P. (2008), *Stochastic Simulation: Algorithms and Analysis*, Springer-Verlag.
5. Asmussen, S. and Nielsen, H. M. (1995), Ruin probabilities via local adjustment coefficients, *J. Appl. Prob.*, **32**, 736-755.
6. Benoite de Saporta (2005). Tail of the stationary solution of the stochastic equation $Y(n+1) = a(n)Y(n) + b(n)$ with Markovian coefficients, *Stoc. Proc. Appl.* **115**(12), 1954–1978.
7. Blanchet, J. and Glynn, P. (2008), Efficient rare-event simulation for the maximum of heavy-tailed random walks, *Ann. Appl. Prob.*, **18**(4), 1351-1378.
8. Blanchet, J., Glynn, P., and Liu, J. C. (2007), Fluid heuristics, Lyapunov bounds and efficient importance sampling for a heavy-tailed G/G/1 queue, *QUESTA*, **57** (2-3), 99–113.
9. Blanchet, J., Leder, K., and Glynn, P. (2009), Lyapunov functions and subsolutions for rare event simulation, *Preprint*.
10. Blanchet, J. and Sigman, K. (2011), Perfect Sampling of Perpetuities, *Journal of Applied Probability*, Special Vol. **48A**, 165-183.
11. Blanchet, J. and Zwart, B. (2007), Importance sampling of compounding processes, *WSC '07: Proceedings of the 39th conference on Winter simulation*, 372-379.
12. Collamore, J. F. (2002), Importance sampling techniques for the multidimensional ruin problem for general Markov additive sequences of random vectors, *Ann. Appl. Prob.*, **12**(1), 382-421.

13. Collamore, J. F. (2009), Random recurrence equations and ruin in a Markov-dependent stochastic economic environment, *Ann. Appl. Prob.*, **19**(4), 1404-1458.
14. Collamore, J. F., Diao, G., and Vidyashankar A. N. (2011), Rare event simulation for processes generated via stochastic fixed point equations. *Preprint*.
15. Diaconis, P. and Freedman, D. (1999), Iterated random functions, *SIAM Review*, **41**(1), 45-76.
16. Dufresne, D. (1990), The distribution of a perpetuity, with applications to risk theory and pension funding, *Scandinavian Actuarial Journal*, **1**(2), 39-79.
17. Dupuis, P. and Wang, H. (2004), Importance sampling, large deviations, and differential games, *Stoc. and Stoc. Reports*, **76**, 481-508.
18. Dupuis, P. and Wang, H. (2007), Subsolutions of an Isaacs equation and efficient schemes for importance sampling: Convergence analysis, *Math. Oper. Res.* **32**, 723-757.
19. Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997), *Modelling Extremal Events for Insurance and Finance*, Springer.
20. Enriquez, N., Sabot, C., Zindy, O. (2009). A probabilistic representation of constants in Kestens renewal theorem. *Prob. Theory and Related Fields*, **144**, 581-613.
21. Glasserman, P. and Kou, S. (1995), Analysis of an importance sampling estimator for tandem queues, *ACM Transactions on Modeling and Computer*, **5**(1), 22-42.
22. Goldie, C. M. (1991), Implicit renewal theory and tails of solutions of random equations, *Ann. Appl. Prob.*, **1**, 126-166.
23. Kesten, H. (1973), Random difference equations and Renewal theory for products of random matrices , *Acta Mathematica*, **131**(1), 207-248.
24. Liu, J. (2001), *Monte Carlo Strategies in Scientific Computing*, Springer.
25. Mikosch, T. and Starica, C. (2000), Limit theory for the sample autocorrelations and extremes of a Garch(1,1) process, *Ann. Stat.*, **28**, 1427-1451.
26. Nyrhinen, H. (2001), Finite and infinite time ruin probabilities in a stochastic economic environment, *Stoc. Proc. and Appl.*, **92**, 265-285.
27. Pollack, M. and Siegmund, D. (1985), A diffusion process and its applications to detecting a change of sign of Brownian motion, *Biometrika*, **72**(2), 267-280.