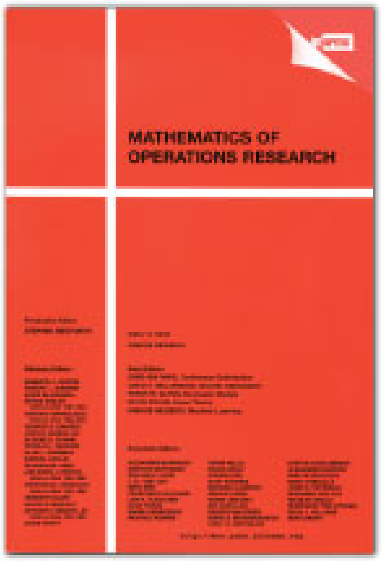

Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Zero-Variance Importance Sampling Estimators for Markov Process Expectations

Hernan P. Awadhttp://moya.bus.miami.edu/~hawad, Peter W. Glynnhttp://www.stanford.edu/~glynn/, Reuven Y. Rubinsteinhttp://ie.technion.ac.il/home/users/ierrr010,

MATHEMATICS OF OPERATIONS RESEARCH
Vol. 38, No. 2, May 2013, pp. 358–388
ISSN 0364-765X (print) | ISSN 1526-5471 (online)

http://dx.doi.org/10.1287/moor.1120.0569

Zero-Variance Importance Sampling Estimators for Markov Process Expectations

Hernan P. Awad
Department of Management Science, University of Miami, Coral Gables, Florida 33124,
h.awad@miami.edu, http://moya.bus.miami.edu/~hawad

Peter W. Glynn
Department of Management Science and Engineering, Stanford University, Stanford, California 94305,
glynn@stanford.edu, http://www.stanford.edu/~glynn/

Reuven Y. Rubinstein
William Davidson Faculty of Industrial Engineering and Management, Technion – Israel Institute of Technology,
Technion City, Haifa 32000, Israel, ierrr01@ie.technion.ac.il, http://ie.technion.ac.il/home/users/ierrr010

We consider the use of importance sampling to compute expectations of functionals of Markov processes. For a class of expectations that can be characterized as positive solutions to a linear system, we show there exists an importance measure that preserves the Markovian nature of the underlying process, and for which a zero-variance estimator can be constructed. The class of expectations considered includes expected infinite horizon discounted rewards as a particular case. In this setting, the zero-variance estimator and associated importance measure can exhibit behavior that is not observed when estimating simpler path functionals (like exit probabilities). The zero-variance estimators are not implementable in practice, but their characterization can guide the design of a good importance measure and associated estimator by trying to approximate the zero-variance ones. We present bounds on the mean-square error of such an *approximate zero-variance* estimator, based on Lyapunov inequalities.

Key words: importance sampling; Markov process; simulation
MSC2000 subject classification: Primary: 65C40, 68U20
OR/MS subject classification: Primary: probability/Markov processes, simulation/efficiency
History: Received September 6, 2006; revised June 23, 2011, and June 4, 2012. Published online in *Articles in Advance* November 28, 2012.

1. Introduction. Importance sampling (IS) is one of the major *variance reduction* and *efficiency improvement* methods used in stochastic simulation, and has enjoyed notable success in certain *rare event* simulation problems. A large literature describes its value in application settings as diverse as dependability modeling (e.g., Goyal et al. [18], Glynn et al. [17], Shahabuddin [33], Nakayama [29], Heidelberger et al. [23]), queueing theory (e.g., Glynn and Iglehart [16], Sadowsky [31], Heidelberger [22], Smith et al. [34]), and computational finance (e.g., Su and Fu [35], Glasserman [14]).

The basic idea behind IS is as follows. One is interested in estimating an expectation of the form $\alpha = \mathrm{E}\, Z$, where Z is a random variable defined on a probability space $(\Omega, \mathcal{F}, \mathrm{P})$. If the restriction of the probability P to the event $\{Z \neq 0\}$ has a density relative to an alternative probability measure Q, so that

$$I(Z \neq 0)\, \mathrm{P}(d\omega) = L(\omega)\, Q(d\omega)$$

for some random variable L (where $I(B)$ denotes the indicator random variable associated with the event B), then α can be alternatively expressed as

$$\begin{aligned}\alpha &= \int_\Omega Z(\omega)\, \mathrm{P}(d\omega)\\ &= \int_\Omega Z(\omega) L(\omega)\, Q(d\omega)\\ &= \mathrm{E}_Q(ZL),\end{aligned}$$

where we now subscript the expectation operator by Q to denote its dependence on the choice of Q. Hence, α can be estimated by averaging the random variables $Z_1L_1, Z_2L_2, \ldots, Z_nL_n$ obtained as independently sampled replicates of ZL, sampled under Q. The probability Q is referred to as the *importance* (*probability*) *measure*, or the *change of measure*, and L is often called the *likelihood ratio*. Of course, the simulationist wants to choose Q so that ZL has small variance (under Q).

It has long been known that if Z is a nonnegative rv, then a zero-variance change of measure for computing $\alpha = \mathrm{E}\,Z$ exists. That is, there exists a unique importance measure Q° and associated likelihood ratio L° such that P is absolutely continuous with respect to Q° on $\{Z > 0\}$, and $ZL^\circ = \alpha$, Q°–a.s. These are given by

$$Q^\circ(d\omega) = \frac{Z(\omega)\,\mathrm{P}(d\omega)}{\int_\Omega Z(\omega')\,\mathrm{P}(d\omega')}, \qquad L^\circ(\omega) = \frac{I(Z>0)}{Z(\omega)}\int_\Omega Z(\omega')\,\mathrm{P}(d\omega'); \tag{1}$$

see, for example, Hammersley and Handscomb [20].

The zero-variance change of measure Q° cannot be literally implemented in practice, because L° and Q° both contain the unknown quantity α. Nevertheless, (1) provides valuable theoretical guidance to the simulationist on how to construct a "good" change of measure: it suggests that an estimator with low variance can be obtained by using an importance measure Q that weights outcomes roughly in proportion to $Z(\omega)\,\mathrm{P}(d\omega)$.

When the above zero-variance change of measure is specialized to the computation of probabilities, so that $Z = I(A)$ for some event A, the zero-variance change of measure requires simulating from the conditional distribution given the rare event, i.e., $Q^\circ(\cdot) = \mathrm{P}(\cdot \mid A)$. Again, this insight is not applied literally. Rather, it suggests that one will get significant variance reduction by building an importance sampler that is a good approximation to sampling from the conditional distribution.

This insight has been successfully applied to rare-event simulation for Markov chains and processes. (We use the term chain for discrete time and process for continuous time.) When used in the context of computing *exit probabilities* for Markov chains and processes, the process retains its Markov dynamics under the zero-variance change-of-measure $Q^\circ(\cdot) = \mathrm{P}(\cdot \mid A)$; this knowledge allows the simulationist to concentrate on so-called *state-dependent* (i.e., Markovian) importance sampling algorithms. Within that class, an importance sampler exhibiting good variance reduction characteristics is obtained by using the problem structure to develop a good approximation to the exit probability function u^* satisfying $u^* = Pu^*$ (often by leveraging off large deviations results). This approach implicitly underlies all the known efficient algorithms for computing rare-event probabilities, both for light-tailed and heavy-tailed models; see, for example, Chang et al. [8], Glasserman et al. [15], Dupuis and Wang [10, 11], Borkar et al. [6], and Blanchet and Glynn [3].

Our goal here is to develop a corresponding theory for computing general expectations for path functionals of Markov chains and processes. In particular, we will study the question of when a zero-variance change of measure exists that is Markovian. As in the rare-event computation setting, the knowledge that the zero-variance change of measure is itself Markov suggests that in building practical importance samplers for expectations, one can restrict the search space to Markov changes of measure. This offers two advantages in terms of algorithm design:

- The search space is vastly reduced. Consider, for example, an expectation that involves a chain on d states over n time steps. With no a priori structure on the appropriate importance distribution, the search space has a dimension of order d^n (that is, the number of paths). With a Markov change of measure, one needs to assign $n-1$ different transition matrices to fully determine the change of measure, so the number of decision variables is reduced to the order of nd^2. Finally, as we shall see later, the particular Markov change of measure that turns out to be zero variance is completely characterized, at each time step, by a specific function defined on the states of the chain. The number of decision variables in developing a good approximation to the zero-variance change of measure is therefore further reduced to the order of nd.
- Effective implementation of a good importance sampler also requires an ability to efficiently simulate sample paths under the proposed importance distribution. This implementation is generally much easier when the joint distribution is Markov than when simulating from an arbitrary joint distribution involving n rvs. For example, if a discrete-event simulation is a GSMP (generalized semi-Markov process) under P, but loses the Markov property under a change of measure, then to simulate the system dynamically it is not enough to know the "physical state" and the state of the "clocks," but rather one needs to keep track of the whole trajectory of the process and compute increasingly complicated conditional laws. Similarly, when simulating a diffusion, a Markovian change of measure will yield another diffusion, which can be simulated using simple Euler-type schemes. In contrast, if the change of measure does not preserve the Markov dynamics, implementing an Euler scheme may involve keeping track of the sample path to compute each increment. Moreover, if the change of measure is not of Girsanov type (which may be the case for Q°), the resulting process may not be amenable to simulation via the Euler method at all.

As we discuss in §2, the class of path functionals for which the zero-variance change-of-measure Q° is Markov is fairly restrictive, consisting only of certain product-form functionals. However, one of our main results, in §3, shows that this class can be greatly expanded if one allows the use of *filtered importance sampling estimators* (i.e., estimators in which each summand of an additive path functional is multiplied by a different likelihood ratio, where the sequence of likelihood ratios is adapted to the Markov process; cf. Glasserman [13]): We show that a wide class of functionals, whose expectations can be characterized as positive solutions to a linear system, are amenable to zero-variance estimation, in the sense that a Markovian change-of-measure Q^* exists under which a filtered estimator has zero variance. Of course Q^* does not coincide with Q° for those functionals for which Q° is not Markov. In fact, it may well be that Q° is not absolutely continuous with respect to Q^*: Q° has the property that, when restricted to the sigma-algebra generated by the Markov process up to a finite time and the event that the path functional of interest is nonzero, P is absolutely continuous with respect to Q°; in contrast, Q^* does not necessarily have this property. Another interesting point is that, in some settings, Q^* may not be unique: more than one zero-variance Markovian change of measure and associated filtered estimator may exist to compute the same expectation.

Our work here is the importance-sampling analog to Henderson and Glynn's [24] work on control variates schemes via *approximating martingales*. Henderson and Glynn [24] consider the use of appropriately chosen martingales as control variates to estimate many different performance measures for a wide class of Markov processes. In their work the martingales are defined in terms of a function u, which, if chosen as the (unknown) solution to the linear system that characterizes the desired performance measure, then the control variate estimator has zero variance; more realistically, u can be chosen as an approximation to the solution of the linear system in question, to obtain a martingale control variate that presumably provides significant variance reduction. Here, we characterize the zero-variance change-of-measure Q^* in terms of the solution to a linear system—a linear integral equation in the Markov chain setting and a linear partial differential equation in the diffusion setting. Knowledge of the solution to this system would allow one to construct a zero-variance filtered IS estimator. In practical situations, a simulationist may be able to use the problem structure to develop a good approximation to the solution to the linear system, and use this approximation to construct an importance sampler: what L'Ecuyer and Tuffin [27] call *approximate zero-variance* simulation. We expect that, in this manner, the characterization of the zero-variance estimators we provide here may help guide the search for an importance sampler with good variance reduction properties, as has been the case in the rare-event simulation setting discussed earlier.

Another way to construct an importance sampler that approximates Q^* is to use the simulation output itself to estimate the solution of the linear system involved and use these estimates to dynamically modify the change of measure used to run the simulation; this leads to adaptive schemes, in which the change of measure gradually converges to Q^*. Ahamed et al. [1] develop and prove convergence of such an adaptive scheme, in the finite-state space Markov chain setting, to estimate the expectation of an additive path functional (a particular case of the ones we consider in §3). Bolia et al. [4] also propose such an adaptive scheme to estimate the price of an option: their setup is a time-inhomogeneous Markov chain living in $\mathbb{R}^d$, and the functional is a "final reward" (a particular case of the ones we consider in §2). Adaptive schemes have also been developed and studied–motivated by applications in particle physics: see for example Booth [5], Kollman et al. [26] (both in the finite state-space setting), and Baggerly et al. [2] (who work with a general state space); they study the convergence rate of their adaptive algorithms to the zero-variance measure (polynomial in Booth [5], geometric in Kollman et al. [26], and Baggerly et al. [2]). (See also Halton [19] for earlier work related to Booth's.) Each of the above studies on adaptive schemes describes the zero-variance importance measure and associated filtered estimator for estimating the expectation of some class of additive path functionals; the path functionals they study, as well as the model frameworks with which they work, are particular cases of those covered in §3. The same is true of L'Ecuyer and Tuffin's [27] work, which is also closely related to ours; they present examples of good importance samplers—to estimate various rare event probabilities—which were obtained by approximating the solution to the linear system involved.

When implementing approximate zero-variance importance sampling, in the sense described above, the resulting estimator will not have zero variance, and a simulationist would like to have some type of guarantee on its efficiency. In §6 we develop bounds on the mean-square error (MSE) of such an estimator. These bounds are based on Lyapunov inequalities, and extend the bounds developed by Blanchet and Glynn [3] for estimators of exit probabilities to the more general class of expectations considered here. Apart from the MSE of the estimator, a simulationist would also be concerned with the length of the simulation run required to compute the estimator. As we discuss below, this can be a very relevant concern in the setting considered here. For example, if the linear system characterizing the desired expectation has multiple nonnegative solutions, and the "wrong" (nonminimal) solution is used as the approximation to implement the approximate zero-variance sampler, then there is necessarily a positive probability that the estimator will need an infinite run length to be computed; moreover, on the event that it is computed in finite time, the estimator is over biased. (See Theorem 2 and discussion in §6.) This concern is also relevant to adaptive schemes like the ones mentioned in the previous paragraph, because it is plausible that an adaptive scheme may converge to the wrong solution of the relevant linear system. One important practice recommendation from our work is, therefore, that when implementing approximate zero-variance importance sampling for expectations of the class considered here, a simulationist should always impose a condition ensuring that the completion time is finite a.s. under the importance measure; we discuss this issue in §6.

2. The zero-variance change of measure for Markov chains. In this section and the next, the setting is a probability space $(\Omega, \mathcal{F}, \mathrm{P})$ supporting a Markov chain $X = (X_j\colon\ j \geq 0)$ in state space $(S, \mathcal{S})$, where S is a Polish (complete, separable, metric) space and $\mathcal{S}$ the Borel sigma-field on S. We denote by $\{\mathcal{F}_j = \sigma(X_0, \dots, X_j), j \geq 0\}$ the filtration generated by X, and by P the one-step transition kernel of X (under P). For any probability measure Q in $(\Omega, \mathcal{F})$ we write E_Q to denote the expectation operator associated with Q, and $Q_x(\cdot) = Q(\cdot \mid X(0) = x)$. (More precisely, $Q_x(\cdot) = \mu_Q(x, \cdot)$, where $(\omega, A) \mapsto \mu_Q(X_0(\omega), A)$ is a regular conditional probability given $\mathcal{F}_0$.) We use E and E_x to denote expectation with respect to P and P_x, respectively.

As discussed in the introduction, given a nonnegative random variable Z with finite mean, there is a unique change-of-measure Q on $\mathcal{F}$ such that the importance sampling estimator $Z(d\mathrm{P}/dQ)$ of $\alpha = \mathrm{E}Z$ has zero variance; namely, Q° satisfying $dQ^\circ/d\mathrm{P} = cZ$, where $c = 1/\mathrm{E}Z$.

In the Markov chain setting, Z will be a path functional of the form $Z = f(X_0, X_1, \dots)$. So, $dQ^\circ/d\mathrm{P} = cf(X_0, X_1, \dots)$. Because X is Markov under P, it will be Markov under a change-of-measure Q iff $dQ/d\mathrm{P}$ is of the form $\prod_i q_i(X_i, X_{i+1})$. Hence, the zero-variance change-of-measure Q° will make X Markov iff Z is P–a.s. of the form

$$Z = f(X_0, X_1, \dots) = \prod_{i=0}^{\infty} q_i(X_i, X_{i+1}), \tag{2}$$

where the infinite product has finite expectation with respect to (w.r.t.) P. (Of course, if f depends on the path of X only up to a finite time n, then $q_i = 1$ for $i > n$.)

When $Z = I_A$ for some event A, then Q° corresponds to $\mathrm{P}(\cdot \mid A)$. As the next three examples illustrate, the indicator of many events of interest in applications can be expressed in product form; hence, the zero-variance change-of-measure Q° associated with them preserves the Markov dynamics.

Example 1 (Tail Probabilities for a Hitting Time). For $x \in S$ and $K \subset S$, put $T(K) = \inf\{j \geq 0\colon X_j \in K\}$, and suppose that we wish to compute $\alpha = \mathrm{P}_x(T(K) > n)$, the probability that K has not been hit by time n, for $n > 0$. Note that $\alpha = \mathrm{E}\,I(T(K) > n)$, and $I(T(K) > n) = \prod_{i=1}^{n} I(X_i \notin K)$, so $Q^\circ = \mathrm{P}(\cdot \mid T(K) > n)$ will induce Markov dynamics on X. In fact, for $0 \leq m \leq m + j \leq n$,

$$Q^\circ(X_{m+j} \in B \mid X_i\colon 0 \leq i \leq m) = \tilde{P}(n - m, j, X_m, B),$$

where $\tilde{P}(i, j, y, B) = \mathrm{P}_y(X_j \in B \mid T(K) > i)$, showing X is indeed a Markov chain (with nonstationary transition probabilities) under Q°.

Example 2 (Probability Mass Function of a Hitting Time). In the same setting as the previous example, suppose that we wish to compute $\alpha = \mathrm{P}_x(T(K) = n)$, the probability that X first hits K at time n. Note that $\alpha = \mathrm{E}\,I(T(k) = n)$, and $I(T(k) = n) = I(X_n \in K) \prod_{i=1}^{n-1} I(X_i \notin K)$, so X is a Markov chain under $Q^\circ(\cdot \mid T(K) = n)$. Indeed, for $0 \leq m \leq m + j \leq n$,

$$Q^\circ(X_{m+j} \in B \mid X_i\colon 0 \leq i \leq m) = \tilde{P}(n - m, j, X_m, B),$$

where $\tilde{P}(i, j, y, B) = \mathrm{P}_y(X_j \in B \mid T(K) = i)$. (Note X has nonstationary transition probabilities under Q°.)

Example 3 (Distribution of X_n). Suppose that we wish to compute $\alpha = \mathrm{P}_x(X_n \in K)$ the probability that X is in K at time n. Note that $\alpha = \mathrm{E}\,I(X_n \in K)$ is trivially of product form, so X is a Markov chain under $Q^\circ = \mathrm{P}(\cdot \mid X_n \in K)$. Indeed, for $0 \leq m \leq m + j \leq n$,

$$Q^\circ(X_{m+j} \in B \mid X_i\colon 0 \leq i \leq m) = \tilde{P}(n - m, j, X_m, B),$$

where $\tilde{P}(i, j, y, B) = \mathrm{P}_y(X_j \in B \mid X_i \in K)$. (Again, X has nonstationary transition probabilities under Q°.)

On the other hand, there are events of interest for which the indicator does not have product form, so the associated zero-variance change-of-measure Q° does not preserve the Markov structure.

Example 4 (Distribution of an Additive Functional of X). For $g\colon S\times S\to\mathbb{R}$, let $\Gamma_n=\sum_{j=0}^{n} g(X_j, X_{j+1})$ be the associated *additive functional* of X. For $x\in S$, put $\alpha=\mathrm{P}_x(\Gamma_n>\gamma)$, and set $Q^\circ(\cdot)=\mathrm{P}_x(\cdot\mid\Gamma_n>\gamma)$. Note $I(\Gamma_n>\gamma)$ cannot be written as a product $\prod_{i=0}^{\infty} q_i(X_i, X_{i+1})$, so that Q° does not induce Markov dynamics on X. In fact, for $0\le m\le m+j\le n$,

$$Q^\circ(X_{m+j}\in B\mid X_i\colon 0\le i\le m)=\tilde{P}(n-m, j, X_m, \gamma-\Gamma_m, B),$$

where $\tilde{P}(i, j, y, r, B)=\mathrm{P}_y(X_j\in B\mid\Gamma_i>r)$.

Here, X is not Markov under Q° because the conditional distribution of X, given its history up to m, depends on both X_m and Γ_m. Of course, if we append Γ_m to X_m as a "supplementary variable," then (X,Γ) is Markov under Q°. We return to this issue in §3.

Moving beyond estimating probabilities, and turning to more general expectations, there are some functionals of interest that have the required product form.

Example 5 (Expected Discounted Terminal Reward). Suppose that we wish to compute $\alpha=\mathrm{E}_x\prod_{i=0}^{n}\beta(X_j, X_{j+1})g(X_n)$. Because this functional has the requisite product form, Q° will preserve the Markov structure of X. We give an explicit expression for the transition kernel of X under Q° in §3.

Example 6 (Moment Generating Function of an Additive Functional of X). Consider again an additive functional of X, $\Gamma_n=\sum_{j=0}^{n} h(X_j, X_{j+1})$, and suppose one is interested in evaluating its moment generating function at the point $\theta\in\mathbb{R}$, i.e., computing $\varphi(\theta)=\mathrm{E}_x\exp(\theta\Gamma_n)$. Because this functional has the requisite product form, the zero variance change of measure Q° will preserve the Markov structure of X. (This is, of course, a particular case of Example 5 with $g=1$ and $\beta=e^h$.)

If Z is a path functional involving a *hitting time* $T=\inf\{j\colon X_j\in K\}$ (where $K\subset S$), so that $Z=\sum_{i=0}^{\infty} f_i(X_0,\ldots,X_i)I(T=i)$, then requiring Z to have the product form (2) seems, at first glance, so restrictive that one cannot expect Q° to preserve the Markov structure of X on any nontrivial example. However, one should note that one only needs to simulate X up to the hitting time T, so it is enough if Q° induces Markov dynamics on X over $\{0,\ldots,T\}$. Hence, one need only consider the restrictions of P and Q° to $\mathcal{F}_T$, the sigma-algebra generated by X up to the stopping time T. As before, the unique change of measure on $\mathcal{F}_T$ such that the importance sampling estimator $Z(d\mathrm{P}/dQ)$ has zero variance is Q° satisfying $dQ^\circ/d\mathrm{P}=cZ$. For X to be Markov under Q°, $dQ^\circ/d\mathrm{P}$ must be of the form $\prod_{i=0}^{T-1} q_i(X_i, X_{i+1})$. Hence, for Q° to make X Markov on $\{0,\ldots,T\}$ it is necessary and sufficient that Z be P–a.s. of the form

$$Z=f_T(X_0, X_1,\ldots,X_T)=\prod_{i=0}^{T-1} q_i(X_i, X_{i+1}),$$

so that the product form criterion arises again in the hitting time setting.

The problems in Examples 1–6 considered functionals that depended on X up to a deterministic time n; they all have counterparts in which the functional depends on X up to a hitting time. For the following examples, fix $K\subset S$ and let $T=\inf\{n\ge 0\colon X_n\in K\}$ be the first hitting time of K. Assume that $\mathrm{P}_x(T<\infty)=1$, $x\in S$.

Example 7 (Exit Probabilities: Distribution of X at a Hitting Time T). For $D\subset K$, our goal here is to compute $\alpha=\mathrm{P}_x(X_T\in D)$. Note that $I(X_T\in D)$ is trivially of the form $\prod_{i=1}^{T-1} q_i(X_i, X_{i+1})$. (Put $q_i(x,y)=0$ for $y\in K\backslash D$ and 1 otherwise.) Hence, $Q^\circ=\mathrm{P}_x(\cdot\mid X_T\in D)$ makes X Markov over $\{0,\ldots,T\}$. In fact, on $\{T>m\}$,

$$Q^\circ(X_{m+1}\in B\mid X_i\colon 0\le i\le m)=\int_B P(X_m, dy)\frac{u(y)}{u(X_m)},$$

where $u(y)\triangleq\mathrm{P}_y(X_T\in D)$, $y\in S$. Note X retains stationary transition probabilities under Q°. (For additional discussion, see Glynn and Iglehart [16].)

Example 8 (Probability of Hitting K Before D). For $D\subset S$ and denoting $T(D)$ the first hitting time of D, our goal here is to compute $\alpha=\mathrm{P}_x(T(D)>T)$, i.e., the probability of hitting K before D. Note that $I(T(D)>T)=\prod_{i=0}^{T} I(X_j\neq D)$, so that $Q^\circ=\mathrm{P}_x(\cdot\mid T(D)>T)$ makes X Markov over $\{0,\ldots,T\}$. In fact, on $\{T>m\}$,

$$Q^\circ(X_{m+1}\in B\mid X_i\colon 0\le i\le m)=\int_B P(X_m, dy)\frac{u(y)}{u(X_m)},$$

where $u(y)\triangleq P_y(T(D)>T)$, $y\in S$. Again, we see X retains stationary transition probabilities under Q°.

Example 9 (Expected Discounted Terminal Reward at a Hitting Time). Suppose that we wish to compute $\alpha = \mathrm{E}_x \prod_{i=0}^{T-1} \beta(X_j, X_{j+1}) g(X_T)$. Because this functional has the requisite product form, Q° will preserve the Markov structure of X over $\{0, \ldots, T\}$. We give an explicit expression for the transition kernel of X under Q° in §3.

3. The zero variance filtered change of measure for Markov chains. In the same setting as in the previous section, consider now a path functional Z with the additive form

$$Z = \sum_i f_i(X_0, \ldots, X_i).$$

Suppose one wants to estimate $\alpha(x) = \mathrm{E}_x Z$, for $x \in S$, using an importance measure Q satisfying $\mathrm{P}(d\omega) = LQ(d\omega)$ on some appropriate sigma algebra. Instead of using the conventional importance sampling estimator ZL, the additive structure gives one the ability to apply a different likelihood ratio to each of the summands, i.e., to construct estimators of the form

$$W = \sum_i f_i(X_0, \ldots, X_i) L_i.$$

(Here $L_i = \mathrm{E}(L \mid X_0, \ldots, X_i)$.) In the spirit of Glasserman [13] we call this a *filtered importance sampling estimator*.

As we will see below, using filtered estimators extends the class of functionals whose expectation can be computed via Markovian zero-variance importance sampling (beyond the product-form functionals discussed in the previous section). In particular, a class of (generalized) expected cumulative discounted rewards can be computed via a zero-variance filtered importance sampling estimator; that is, expectations of the form $\alpha(x) \triangleq \mathrm{E}_x Z$, where

$$Z \triangleq f(X_0) + \sum_{i=1}^{T} \big[f(X_i) + g(X_{i-1}, X_i)\big] \prod_{j=1}^{i} \beta(X_{j-1}, X_j); \tag{3}$$

here T is the first hitting time of a set $K \subset S$ (i.e., $T = \inf\{n \geq 0\colon X_n \in K\}$), $f\colon S \to [0, \infty)$, $\beta\colon S \times S \to (0, \infty)$, and $g\colon S \times S \to [0, \infty)$; we assume for notational ease that $g(x, \cdot) = 0$ for $x \in K$.

The expectations of product-form functionals discussed in the previous section are special cases of (3). In particular, an exit probability at a hitting time T and the expected discounted terminal reward at a hitting time (as in Examples 7–8) are covered by (3). The finite-horizon expectations of product form functionals (as those in Examples 1–3, 5–6) can also be put into the form (3) by appending the time to the state—see Corollary 1. However, (3) is significantly more general, covering expectations beyond those of product form functionals, e.g., the expected sojourn on a set D before hitting a set K, or the expected discounted cumulative reward until hitting a set K. Also, on our next theorem we permit $K = \varnothing$, in which case $T = \infty$, so that the infinite horizon expected discounted reward is also a special case of the expectation (3) (having $\beta(x, y) = \delta$ for some $\delta < 1$).

The Markovian change of measure and associated filtered estimator that we construct will be closely connected with positive solutions to the linear system (4), whose existence we assume:

Assumption 1. *There exists a finite nonnegative solution u to the linear integral equation*

$$u(x) = f(x) + \int_S \beta(x, y)[g(x, y) + u(y)] P(x, dy), \quad x \in K^C, \tag{4}$$

subject (if $K \neq \varnothing$) to the boundary condition $u(x) = f(x)$ for $x \in K$.

We note that the expectation $\alpha(\cdot)$ given in (3) corresponds to the minimal nonnegative solution to (4), which we denote u^*. Hence, if $\mathrm{E}_x Z < \infty$ for $x \in K^C$, Assumption 1 holds.

Given any nonnegative function $u\colon S \to [0, \infty)$, we construct a change of measure based on it as follows. Define

$$w(x) \triangleq f(x) + \int_S \beta(x, y)[g(x, y) + u(y)] P(x, dy), \quad x \in K^C.$$

(Note $w = u$ if u is a solution to (4).) Let Q be the measure on $\mathscr{F}_T$ under which X is a Markov chain with one step transition kernel

$$M(x, dy) = \begin{cases} \dfrac{\beta(x, y)[g(x, y) + u(y)] P(x, dy)}{w(x) - f(x)}, & x \in K^C, \ w(x) > f(x), \\ P(x, dy), & \text{otherwise.} \end{cases}$$

Put

$$W = f(X_0) + \sum_{i=1}^{T}\big[f(X_i) + g(X_{i-1}, X_i)\big] L_i \prod_{j=1}^{i} \beta(X_{j-1}, X_j), \tag{5}$$

where $L_i = \prod_{j=1}^{i} l(X_{i-1}, X_i)$ and

$$l(x, y) = \begin{cases} \dfrac{w(x) - f(x)}{\beta(x, y)[g(x, y) + u(y)]}, & x \in K^C, \quad w(x) > f(x), \quad g(x, y) + u(y) > 0, \\ 1, & \text{else.} \end{cases}$$

In the particular case in which $u = u^*$, the minimal nonnegative solution to (4), we denote the above objects Q^*, M^* and l^*, respectively.

Our next result shows that, in great generality, W is an unbiased estimator of $\alpha(x)$ under Q_x. (Note the result does not use Assumption 1: in particular, u is not necessarily a solution to (4).)

Theorem 1. *Suppose u: $S \to \mathbb{R}$ satisfies $f(x) \le u(x) < \infty$ for $x \in S$, and $u(x) > f(x)$ for $x \in A$, where*

$$A = \big\{x \in K^C\colon\ \mathrm{P}_x(Z > f(x)) > 0\big\}.$$

Also, assume $\int_S u(y)\beta(x, y)P(x, dy) < \infty$ for $x \in K^C$. Then $\mathrm{E}_{Q_x} W = \alpha(x)$.

Proof. For $i \ge 1$ and $x_0 \in A$, put

$$A_i(x_0) \triangleq \big\{(x_1, \ldots, x_i)\colon x_j \in A,\ j < i;\ f(x_i) + g(x_{i-1}, x_i) > 0\big\}.$$

Note that, for $i \ge 1$, $f(X_i) + g(X_{i-1}, X_i) = 0$ on $\{(X_1, \ldots, X_i)\colon \notin A_i(x)\}$, both P_x–a.s. and also Q_x–a.s. (because $Q_x \ll \mathrm{P}_x$ on $\mathcal{F}_{i\wedge T}$). Also, if $x \in A$ then it must be that $\mathrm{P}_x(f(X_1) + g(x, X_1) > 0) > 0$ or $P(x, A) > 0$; in either case, $\int_S \beta(x, y)[g(x, y) + u(y)]P(x, dy) > 0$, so that $w(x) > f(x)$ for $x \in A$. Hence, for $i \ge 1$ and $(x_0, \ldots, x_i) \in A \times A_i(x_0)$,

$$l(x_{j-1}, x_j) = \frac{w(x_{j-1}) - f(x_{j-1})}{\beta(x_{j-1}, x_j)[g(x_{j-1}, x_j) + u(x_j)]}, \quad j \le i.$$

It follows that $\mathrm{P}_x \ll Q_x$ on $\{(X_1, \ldots, X_i) \in A_i(x)\} \cap \mathcal{F}_{i\wedge T}$. Thus, for fixed $x \in A$ and denoting $x_0 = x$,

$$\begin{aligned}
\mathrm{E}_x Z &= f(x) + \sum_{i=1}^{\infty} \mathrm{E}_x[f(X_i) + g(X_{i-1})]\bigg(\prod_{j=1}^{i} \beta(X_{j-1}, X_j)\bigg) I(T \ge i) \\
&= f(x) + \sum_{i=1}^{\infty} \int_{K^C\times\cdots\times K^C\times S} [f(x_i) + g(x_{i-1})]\bigg(\prod_{j=1}^{i} \beta(x_{j-1}, x_j)\bigg) P(x, dx_1)\cdots P(x_{i-1}, dx_i) \\
&= f(x) + \sum_{i=1}^{\infty} \int_{A_i(x)} [f(x_i) + g(x_{i-1})]\bigg(\prod_{j=1}^{i} \beta(x_{j-1}, x_j)\bigg) P(x, dx_1)\cdots P(x_{i-1}, dx_i) \\
&= f(x) + \sum_{i=1}^{\infty} \int_{A_i(x)} [f(x_i) + g(x_{i-1})]\bigg(\prod_{j=1}^{i} \beta(x_{j-1}, x_j) l(x_{j-1}, x_j)\bigg) M(x, dx_1)\cdots M(x_{i-1}, dx_i) \\
&= \mathrm{E}_{Q_x} W. \quad \square
\end{aligned}$$

When u is a solution to (4), one can say more about the behavior of W; and if $u = u^*$, then W has zero-variance under Q^*, as our next result shows.

Theorem 2. *Let u be as in Assumption 1. Then, the following hold*:
(i) *On* $\{T < \infty\}$, $W = u(X_0)$, Q-a.s.
(ii) *If* $u = u^*$, $W = u^*(X_0)$, Q^*-a.s.

Proof. Set $Y_0 = u(X_0)$, and for $n \ge 1$ put

$$\begin{aligned}
Y_n = f(X_0) &+ \sum_{i=1}^{(T\wedge n)-1} \big[f(X_i) + g(X_{i-1}, X_i)\big] L_i \prod_{j=1}^{i} \beta(X_{j-1}, X_j) \\
&+ L_{T\wedge n}\big[g(X_{T\wedge n-1}, X_{T\wedge n}) + u(X_{T\wedge n})\big] \prod_{j=1}^{T\wedge n} \beta(X_{j-1}, X_j).
\end{aligned}$$

Observe that, on $\{T > n\}$,

$$Y_{n+1} = Y_n + \big\{f(X_n) - u(X_n) + \beta(X_n, X_{n+1}) l(X_n, X_{n+1})\big[g(X_n, X_{n+1}) + u(X_{n+1})\big]\big\} L_n \prod_{j=1}^{n} \beta(X_{j-1}, X_j).$$

Note that $X_n \in K^C$ on $\{T > n\}$. Hence, if $u(X_n) > f(X_n)$, it immediately follows that

$$f(X_n) - u(X_n) + \beta(X_n, X_{n+1}) l(X_n, X_{n+1})\big[g\big(X_n, X_{n+1} + u(X_{n+1})\big)\big] = 0, \tag{6}$$

so that $Y_{n+1} = Y_n$. On the other hand, if $u(X_n) = f(X_n)$, then

$$\int_S \beta(X_n, y)\big[g(X_n, y) + u(y)\big] P(X_n, dy) = 0;$$

see (4). Since $M(x, \cdot) = P(x, \cdot)$ when $u(x) = f(x)$, it follows that $E_Q[g(X_n, X_{n+1}) + u(X_{n+1}) \mid X_n] = 0$, implying that $g(X_n, X_{n+1}) + u(X_{n+1}) = 0$ Q-a.s. on $\{u(X_n) = f(X_n)\}$. As a consequence, (6) also holds when $u(X_n) = f(X_n)$. We conclude that $Y_{n+1} = Y_n$ whenever $n < T$. Since $W = Y_T$ on $\{T < \infty\}$ and $Y_0 = u(X_0)$, we find that $W = u(X_0)$ on $\{T < \infty\}$. This proves part (i).

To prove part (ii), all that remains to be shown is that if $u = u^*$ then $W = u(X_0)$ also holds on $\{T = \infty\}$. For this, note that on $\{T = \infty\}$,

$$\begin{aligned} u(X_0) &= f(X_0) + \sum_{i=1}^{n-1} \big[f(X_i) + g(X_{i-1}, X_i)\big] L_i \prod_{j=1}^{i} \beta(X_{j-1}, X_j) \\ &\quad + L_n\big[g(X_{n-1}, X_n) + u(X_n)\big] \prod_{j=1}^{n} \beta(X_{j-1}, X_j) \end{aligned}$$

for $n \geq 0$. Hence,

$$L_n\big[g(X_{n-1}, X_n) + u(X_n)\big] \prod_{j=1}^{n} \beta(X_{j-1}, X_j) \searrow u(X_0) - W \tag{7}$$

Q^*–a.s. as $n \to \infty$ on $\{T = \infty\}$. But, if $u = u^*$,

$$\begin{aligned} &E_{Q_x^*} I(T > n) L_n\big[g(X_{n-1}, X_n) + u(X_n)\big] \prod_{j=1}^{n} \beta(X_{j-1}, X_j) \\ &\quad = \mathrm{E}_x I(T > n)\big[g(X_{n-1}, X_n) + u(X_n)\big] \prod_{j=1}^{n} \beta(X_{j-1}, X_j) \\ &\quad = \mathrm{E}_x I(T > n)\big[g(X_{n-1}, X_n) + f(X_n)\big] \prod_{j=1}^{n} \beta(X_{j-1}, X_j) \\ &\qquad + \mathrm{E}_x \sum_{i=n+1}^{\infty} \big[f(X_i) + g(X_{i-1}, X_i)\big] I(T > i) \prod_{j=1}^{i} \beta(X_{j-1}, X_j) \\ &\quad = \mathrm{E}_x I(T > n) \sum_{i=n}^{T-1} \big[f(X_i) + g(X_{i-1}, X_i)\big] \prod_{j=1}^{i} \beta(X_{j-1}, X_j) \\ &\quad \longrightarrow 0 \end{aligned}$$

by the dominated convergence theorem. (Note the quantity inside the expectation is dominated by Z as in (3), and $\mathrm{E}_x Z = \alpha(x) = u^*(x) = u(x) < \infty$.) Fatou's lemma therefore implies that

$$\liminf_{n \to \infty} \prod_{j=1}^{n} \beta(X_{j-1}, X_j) L_n[g(X_{n-1}, X_n) + u(X_n)] = 0 \quad Q_x^*\text{-a.s.}$$

on $\{T = \infty\}$. Relation (7) then proves that $W = u(X_0)$ on $\{T = \infty\}$ if $u = u^*$. □

Theorem 2 proves that, in significant generality, there exist Markovian importance measures and associated zero-variance filtered estimators for expectations of the form (3), i.e., for the minimal nonnegative solution to the linear integral equation (4). It also provides an explicit formula for the transition kernel under the importance measure. Of course, the kernel M^* requires knowledge of u^* (and in particular of the expectation one is trying to compute), whence the zero-variance estimators are not implementable in practice. But, as discussed earlier, knowledge of the structure of the transition kernel M^* can guide the design of good importance samplers; we discuss this point further in §§5 and 6.

Remark 1. If Z corresponds to an infinite horizon discounted reward and the problem data is appropriately bounded, then it is enough that the solution u to (4) be bounded (because such a bounded u must coincide with u^*).

Remark 2 (Infinite Completion Times when $u \neq u^*$). Note that if u is a finite-valued nonnegative solution to (4), but $u \neq u^*$, then it follows from part (i) in Theorem 2 that, on $\{T < \infty\}$, W is an over-biased estimator of $u^*(x)$ under Q_x (because $u(x) \geq u^*(x)$). Since W is unbiased (by Theorem 1), it follows that, necessarily, $Q(T = \infty) > 0$. That is, there is a positive probability that a simulation run will not return a value for W in finite time. We return to this issue in §6.

Remark 3 (Nonuniqueness of Q^*). In principle one can, without loss of generality, assume that $f = 0$ in (3) and (4), since these "per visit rewards" can be incorporated into the "per transition rewards": if $f > 0$, one can set it to zero by replacing g with $\tilde{g}$, where

$$\tilde{g}(x, y) = f(x)/\beta(x, y) + g(x, y) + I_K(y)f(y).$$

This leaves Z in (3) unchanged. Also, both formulations are equivalent in terms of the linear system (4) they define, in the sense that if u solves it with the first formulation, then $\tilde{u} = uI_{K^C}$ is a solution under the second formulation. However, these two formulations define a different transition kernel M^*, and hence give two different importance measures and associated zero-variance estimators. More generally, if $f > 0$, one can replace f, g in (3) by $\tilde{f}, \tilde{g}$, where $\tilde{f}(x) = (1 - \gamma(x))f(x)$ and $\tilde{g}$ as above but with γf in place of f; changing $\gamma(\cdot) \in (0, 1)$ leaves Z unchanged, but it affects the change of measure and estimator. This is illustrated in Example 10.

Example 10. Consider a Markov chain with two states $S = \{0, 1\}$, and transition kernel $P(1, 0) = q$, $P(1, 1) = 1 - q$, and state 0 absorbing. We want to estimate the expected time until absorbtion, starting from state 1 (i.e., the mean of a geometric(q) rv). That is, $K = \{0\}$, $Z = \sum_{i=1}^{T} 1$, $\alpha(1) = 1/q$ and $\alpha(0) = 0$. This can be mapped to the representation (3) of Z in several ways:

(i) Put $f(1) = 1$, $f(0) = 0$, $g = 0$ and $\beta = 1$. This gives $M^*(1, 1) = 1$, $M^*(1, 0) = 0$ (so that $T = \infty$: Q^*-a.s.), $l^*(1, 1) = 1 - q$, and hence $W = \sum_{j=0}^{\infty}(1 - q)^j = 1/q$.

(ii) Put $f = 0$, $\beta = 1$ and $g(1, 1) = g(1, 0) = 1$. This gives $M^*(1, 1) = ((1 - q)(1 + 1/q))/(1/q) = 1 - q^2$, $M^*(1, 0) = q^2$ (so that $Q^*(T = \infty) = 0$), $l^*(1, 1) = 1/(1 + q)$, $l^*(1, 0) = 1/q$, and hence $W = \sum_{j=0}^{T-1}(1/(1 + q))^j + (1/q)(1/(1 + q))^{T-1} = 1/q$.

(iii) More generally, for $0 < \gamma \leq 1$ put $f(1) = 1 - \gamma$, $f(0) = 0$, $g(1, 1) = g(1, 0) = \gamma$, $\beta = 1$. This gives $M^*(1, 1) = ((1 - q)(\gamma + 1/q))/(1/q - 1 + \gamma)$, $M^*(1, 0) = (q\gamma)/(1/q - 1 + \gamma)$, $l^*(1, 1) = (1 - q + q\gamma)/(1 + q\gamma)$, $l^*(1, 0) = (1 - q + q\gamma)/(q\gamma)$, and hence $W = 1 - \gamma + \sum_{j=1}^{T-1}((1 - q + q\gamma)/(1 + q\gamma))^j + \gamma((1 - q + q\gamma)/(1 + q\gamma))^{T-1}((1 - q + q\gamma)/(q\gamma)) = 1/q$.

We see that even for this extremely simple example there is a multiplicity of zero-variance filtered estimators for $\mathrm{E}\,Z$, each associated to a different Markovian change of measure. Although all of them have zero variance, they differ in terms of the length of the simulation run used to compute W under Q^*: it is a geometric $(q\gamma/(1/q - 1 + \gamma))$ in case (iii), whereas computing W requires simulating a path of infinite length in case (i). This suggests that, when there are alternative ways to specify the functions f and g in a problem of interest, the choice can significantly impact the efficiency of the estimator one constructs.

Note also that Q°, the zero-variance change of measure for the conventional (nonfiltered) estimator, satisfies $Q_1^\circ(T = n) = (n(1 - q)^{n-1}q)/(1/q) = nq^2(1 - q)^{n-1}$, $n \geq 1$. Hence, $Q^\circ(X_{n+1} = 1 \mid X_0 = \cdots = X_n = 1) = (1 - q) \cdot (1 + 1/(n + 1/q - 1))$. Although X is still Markov (with nonstationary dynamics) under Q°, this is only because in this simple example the time period n itself contains all the relevant path information, including the reward accrued up to n.

Remark 4 (Potentially $Q^\circ \not\ll Q^*$ and $\mathrm{P} \not\ll Q^*$). The importance measure Q° discussed in the previous section has the property that P is absolutely continuous with respect to Q° when restricted to $\{Z > 0\}$ and a finite time horizon, i.e., on $\mathcal{F}_n \cap \{Z > 0\}$. This is not necessarily true of Q^*: Note that $M^*(x, D(x)) = 0$ for $D(x) = \{y\colon g(x, y) + u(y) = 0\}$, even if $P(x, D(x)) > 0$. This is not an issue with the product form functionals discussed in the previous section, since $Z = 0$ on a path that includes such a transition from a state x to $y \in D(x)$. But here a path that moves from x to $D(x)$ may have accrued positive reward before that transition occurs, so such a transition may have positive probability under Q° and P, and not under Q^*. This can be observed in Example 10(i), where Q^* assigns probability 0 to paths of finite length, even though $Z > 0$ on such paths.

In some problems of interest the relevant time horizon is not a hitting time, but rather a fixed finite time horizon n. The result above extends to this situation. Suppose that

$$\alpha_n(x) \triangleq \mathrm{E}_x\, f_0(X_0) + \sum_{i=1}^{n} \big[f_i(X_i) + g_i(X_{i-1}, X_i)\big] \prod_{j=1}^{i} \beta_j(X_{j-1}, X_j), \tag{8}$$

where $f_i\colon S \to [0, \infty)$, $g_i\colon S \times S \to [0, \infty)$ and $\beta_i\colon S \times S \to (0, \infty)$.

It is easily verified that $\alpha_n(x) = u(0, x)$, where u solves

$$u(k, x) = f_k(x) + \int_S P(x, dy)\beta_{k+1}(x, y)\big[g_{k+1}(x, y) + u(k+1, y)\big], \tag{9}$$

for $k \ge 0$, subject to the boundary condition $u(n, \cdot) = f_n$.

This can be formulated as a particular case of the problem treated in Theorem 2, by appending the time period to the state variable; that is, by considering the Markov chain $Y = (Y_j\colon j \ge 0)$ in $S \times \mathbb{N}$, where $Y_j = (X_j, j)$. Then (8) becomes of the form (3), where "reward" is accumulated until hitting the set $K = S \times \{n\}$. We have then the following result.

Corollary 1. *Suppose that u solves (9), and satisfies $0 < u(j, x) < \infty$ for $x \in S$ and $0 \le j \le n$. Let Q^* be the importance measure under which X is a (time inhomogeneous) Markov chain with time i transition kernel*

$$M_i^*(x, dy) = \begin{cases} \dfrac{\beta_i(x, y)\big[g_i(x, y) + u(i, y)\big]P(x, dy)}{u(i-1, x) - f_{i-1}(x)}, & x \in S, \quad u(i-1, x) > f_{i-1}(x), \\[2ex] P(x, dy), & \textit{otherwise}, \end{cases}$$

$i \ge 1$. *Put*

$$W = f_0(X_0) + \sum_{i=1}^{n} L_i\big[f_i(X_i) + g_i(X_{i-1}, X_i)\big] \prod_{j=1}^{i} \beta_j(X_{j-1}, X_j),$$

where $L_i = \prod_{j=1}^{i} l_j^*(X_{i-1}, X_i)$ and

$$l_j^*(x, y) = \begin{cases} \dfrac{u(i-1, x) - f_{i-1}(x)}{\beta_i(x, y)\big[g_i(x, y) + u(i, y)\big]}, & u(i-1, x) > f_{i-1}(x), \quad g_i(x, y) + u(i, y) > 0, \\[2ex] 1, & \textit{else}. \end{cases}$$

Then

$$W = u(0, X_0) \quad Q^*\text{-a.s.}$$

Remark 5. Note that the product form functionals considered in §2 are special cases of the theory that we have just developed; in particular, the Markovian conditional distributions arising in several of the examples discussed there, as well as the zero-variance change of measure associated with expected discounted terminal rewards, can be viewed as special cases of the results developed in this section. For such product form functionals, the change-of-measure Q^* described in Theorem 2 coincides with Q°.

The above discussion shows that, for a wide class of Markov process expectations that can be expressed as minimal nonnegative solutions to a linear system of the form (4), a zero-variance filtered importance sampling estimator exists for which the associated change of measure is Markovian. Furthermore, the change of measure induces nonstationary transition probabilities whenever the solution depends explicitly on time. Our next result shows that a partial converse holds: if a filtered estimator of the form considered above has zero variance under some measure that makes the underlying process Markovian, then the quantity it is estimating must be the minimal nonnegative solution to a linear system like (4).

Theorem 3. *Let X be a Markov chain in state-space S, with one-step transition kernel M under some probability measure Q. Let T be the hitting time of $K \subset S$, and*

$$W = f(X_0) + \sum_{i=1}^{T} [f(X_i) + g(X_{i-1}, X_i)] \prod_{j=1}^{i} \beta(X_{j-1}, X_j) l(X_{j-1}, X_j),$$

where $f, g \geq 0$, $\beta > 0$ and $l \geq 0$ is such that

$$\int_S l(x, y) M(x, dy) \leq 1,$$

$x \in S$. Suppose there exists u: $S \to \mathbb{R}$ such that $W = u(x)$ Q_x–a.s., $x \in S$. Then, the following hold:

(i) *The function u solves*

$$u(x) = f(x) + \int_S [g(x, y) + u(y)] \beta(x, y) B(x, dy), \tag{10}$$

where $B(x, dy) = l(x, y) M(x, dy)$, $x \in K^C$, and $B(x, dy) = 0$, $x \in K$.

(ii) *The function u is given by*

$$u(x) = \left(\sum_{j=0}^{\infty} \tilde{B}^j \tilde{f} \right)(x),$$

$x \in S$, where $\tilde{f}(x) = f(x) + \int_S \beta(x, y) g(x, y) B(x, dy)$ and $\tilde{B}(x, dy) = \beta(x, y) B(x, dy)$. Hence, u is the minimal nonnegative solution to (10).

Proof. Note that, for $x \in K^C$,

$$\begin{aligned} u(x) &= W \\ &= f(x) + l(x, X_1)\beta(x, X_1) g(x, X_1) \\ &\quad + l(x, X_1)\beta(x, X_1) \left\{ f(X_1) + \sum_{j=2}^{T} [f(X_j) + g(X_{j-1}, X_j)] \prod_{k=2}^{j} \beta(X_{k-1}, X_k) l(X_{k-1}, X_k) \right\} \\ &= f(x) + l(x, X_1)\beta(x, X_1) g(x, X_1) + l(x, X_1)\beta(x, X_1) u(X_1), \end{aligned}$$

Q_x–a.s. (The last follows using the assumption that $W = u(y)$ Q_y-a.s. on the rv within the braces.) In particular, integrating both sides,

$$\begin{aligned} u(x) &= f(x) + \int_S [g(x, y) + u(y)] l(x, y) \beta(x, y) M(x, dy), \\ &= f(x) + \int_S [g(x, y) + u(y)] \beta(x, y) B(x, dy), \end{aligned}$$

for $x \in K^C$, giving (10).

Also, note that

$$W = f(X_0) + \sum_{j=1}^{\infty} I(T \geq j)[f(X_j) + g(X_{j-1}, X_j)] \prod_{k=1}^{j} l(X_{k-1}, X_k) \beta(X_{k-1}, X_k),$$

so that

$$\begin{aligned} u(x_0) &= \mathrm{E}_{Q_{x_0}} W \\ &= f(x_0) + \sum_{j=1}^{\infty} \int_{K^C \times \cdots \times K^c \times S} M(x_0, dx_1) \cdots M(x_{j-1}, dx_j) [f(x_j) + g(x_{j-1}, x_j)] \prod_{k=1}^{j} l(x_{k-1}, x_k) \beta(x_{k-1}, x_k) \\ &= f(x_0) + \sum_{j=1}^{\infty} \int_{S \times \cdots \times S} \tilde{B}(x_0, dx_1) \tilde{B}(x_1, dx_2) \cdots \tilde{B}(x_{j-1}, dx_j) [f(x_j) + g(x_{j-1}, x_j)] \\ &= \left(\sum_{j=0}^{\infty} \tilde{B}^j \tilde{f} \right)(x_0). \quad \square \end{aligned}$$

Note the above result provides another way to justify one of the observations made in Remark 2: If one builds the change of measure and estimator W in (5) using a solution u to (4), which is not u^*, the minimal nonnegative solution, then the estimator cannot have zero variance.

We conclude this section by discussing how the results above can be applied to estimating steady-state expectations, and how they extend to Markov pure-jump processes.

Remark 6 (Steady-State Expectations). When X is a discrete-time Markov chain possessing positive recurrent regenerative structure, then the steady-state expectation of a nonnegative function $f\colon S \to [0, \infty)$ can be expressed as a ratio of two expectations. Each of the two expectations can, in turn, be computed via zero-variance filtered importance sampling. For example, if X is an irreducible positive recurrent Markov chain on discrete state space S with stationary distribution $\pi = (\pi(x)\colon x \in S)$, then

$$\sum_{x \in S} \pi(x) f(x) = \frac{E_z \sum_{j=0}^{\tau(z)-1} f(X_j)}{E_z \tau(z)}, \tag{11}$$

where $\tau(z) = \inf\{n \geq 1\colon X_n = z\}$ is the first return time to the *regeneration state* z. As a consequence, steady-state performance measures can be handled by separately appealing to Theorem 2 for both the numerator and denominator of (11).

Remark 7 (Markov Pure Jump Processes). Let $X = (X(t)\colon t \geq 0)$ be a pure jump Markov process (or continuous time Markov chain (CTMC)), on discrete state space S. Assume X is nonexplosive; a sufficient condition is to require that the rate matrix $A = (A(x, y)\colon x, y \in S)$ be uniformizable, so that $\inf\{A(x, x)\colon x \in S\} > -\infty$; see, e.g., Chung [9]. Let $K \subset S$ and set $T = \inf\{t \geq 0\colon X(t) \in K\}$. Suppose one is interested incomputing an expectation of the form

$$\alpha(x) = \mathrm{E}_x \int_0^T f(X(s)) \exp\left(\int_0^s h(X(u))\, du\right) ds \tag{12}$$

for given functions $f\colon S \to [0, \infty)$ and $h\colon S \to \mathbb{R}$.

One can prove an analog to Theorem 2 for the continuous-time formulation (12): That is, there exists a change of measure that allows one to simulate X in continuous time, preserving its Markov property, and under which an appropriately defined filtered estimator computes $\alpha(x)$ with zero variance. However, such a result is not really needed, since the expectation (12) can be rewritten in a way that involves only a discrete-time Markov chain, as we discuss next. (See Fox and Glynn [12] for more on this discrete time conversion and its advantages from an efficiency standpoint.)

Let Γ_n be the time of the n'th jump of X, with $\Gamma_0 \triangleq 0$, and let $J(t)$ be the number of jumps of X in the interval $[0, t]$. Note that $\alpha(x)$ can be computed as the expectation of the random variable

$$\mathrm{E}_x\left[\int_0^T f(X(s)) \exp\left(\int_0^s h(X(u))\, du\right) ds \,\middle|\, X(\Gamma_n)\colon n \geq 0\right]. \tag{13}$$

The above conditional expectation is easily computable in closed form and involves only the embedded discrete-time Markov chain $(X(\Gamma_n)\colon n \geq 0)$. Also, the form of the conditional expectation is a particular case of (3) (with $g = 0$ and appropriate definitions for β and f appearing there). Furthermore, as a conditional expectation, (13) has lower variance and is therefore to be preferred computationally to the continuous-time estimand appearing in (12).

4. The zero-variance filtered change of measure for stochastic differential equations. In this section we study the theory of §3 in the stochastic differential equations (SDEs) setting, where $X = (X(t)\colon t \geq 0)$ is a diffusion in $\mathbb{R}^m$. In an analog to the discrete-time case, we find that a wide class of expectations that can be represented as positive solutions to linear partial differential equations can be computed via a zero-variance filtered importance sampling estimator under a Markovian change of measure, so that X is again a diffusion under the importance measure.

Let $B = (B(t)\colon t \geq 0)$ be standard Brownian motion in $\mathbb{R}^d$, and assume that X is a (strong) solution of the SDE

$$dX(t) = \mu(X(t))\, dt + \sigma(X(t))\, dB(t), \tag{14}$$

where $\mu\colon \mathbb{R}^m \to \mathbb{R}^m$ and $\sigma\colon \mathbb{R}^m \to \mathbb{R}^{m \times d}$ satisfy, for some constant $C > 0$,

$$\|\mu(x) - \mu(y)\| + \|\sigma(x) - \sigma(y)\| \leq C\|x - y\|,$$
$$\|\mu(x)\| + \|\sigma(x)\| \leq C(1 + \|x\|),$$

$x, y \in \mathbb{R}^m$, so that (14) is guaranteed to have a (strongly) unique solution.

As in §3, we first consider the (generalized) expected discounted reward up to the hitting time of a set. Suppose $K \subset \mathbb{R}^m$ is closed, and let $T \triangleq \inf\{t \geq 0: X(t) \in K\}$. An analog of (3) in this setting is

$$\alpha(x) = \mathrm{E}_x \int_0^T f(X(t))\beta(t)\,dt + I(T < \infty)g(X(T))\beta(T), \tag{15}$$

$x \in \mathbb{R}^m$, where $f: \mathbb{R}^m \to [0, \infty)$, $g: \mathbb{R}^m \to [0, \infty)$,

$$\beta(t) \triangleq \exp\left(\int_0^t h(X(s))\,ds\right),$$

and $h: \mathbb{R}^m \to \mathbb{R}$.

The function α can be characterized as the solution of a linear partial differential equation. More specifically, we assume there exists a function u satisfying Assumption 2; then, under appropriate integrability conditions, $\alpha = u$.

Assumption 2. *There exists a function $u: \mathbb{R}^m \to \mathbb{R}$ satisfying*
(i) *u is twice continuously differentiable;*
(ii) *u is a solution to*

$$\begin{aligned} \mu(x)^T \nabla_x u(x) + \tfrac{1}{2}\operatorname{tr}(\sigma(x)\sigma(x)^T \nabla_{xx} u(x)) + h(x)u(x) &= -f(x), \quad x \in K^C, \\ u(x) &= g(x), \quad x \in K, \end{aligned} \tag{16}$$

where $\nabla_x u$ and $\nabla_{xx} u$ denote the gradient and the Hessian of u, respectively.

We consider importance measures Q under which X evolves according to the SDE

$$dX(t) = \tilde{\mu}(X(t))\,dt + \sigma(X(t))\,dB(t),$$

with B being standard Brownian motion under Q. More precisely, we consider Q of the form $Q(\cdot) = \int_S \pi(dx) Q_x(\cdot)$, where π is a distribution on $\mathbb{R}^m$ and $\{Q_x: x \in S\}$ together with X defines a Markov family; see, e.g., Karatzas and Shreve [25, Definition 5.11]. Then, under the conditions stated below, $u(x) = \mathrm{E}_{Q_x} W$, where

$$\begin{aligned} W &= \int_0^T f(X(t))\beta(t)L(t)\,dt + I(T < \infty)g(X(T))\beta(T)L(T), \\ L(t) &= \exp\left(-\int_0^t \eta(X(s))\;dB(s) - \tfrac{1}{2}\int_0^t \|\eta(X(s))\|^2\,ds\right), \end{aligned} \tag{17}$$

and $\eta: \mathbb{R}^m \to \mathbb{R}^d$ satisfies

$$\sigma(x)\eta(x) = \tilde{\mu}(x) - \mu(x).$$

The next result, an analog to Theorem 2, describes a Markovian zero-variance importance distribution to compute u. Before stating the theorem we need to introduce one more assumption:

Assumption 3. *There exists a function u^* satisfying*
(i) *u^* satisfies the conditions in Assumption 2;*
(ii) *$u^*(x) > 0$ for $x \in K^C$ and $u^*(x) < \infty$ for $x \in \mathbb{R}^m$;*
(iii) *for any $\delta > 0$, the function $x \mapsto \sigma(x)^T \nabla_x u^*(x)/u^*(x)$ is bounded on $G_\delta^C \cap K^C$, where $G_\delta \triangleq \{x: u^*(x) \leq \delta\}$;*
(iv) $\mathrm{E}\int_0^t \beta^2(s)\|\sigma(X(s))^T \nabla_x u^*(X(s))\|^2\,ds < \infty$, $t > 0$;
(v) $\mathrm{E}_x \beta(t \wedge T)u^*(X(t \wedge T)) \to \mathrm{E}_x I(T < \infty)\beta(T)g(X(T))$ *as $t \to \infty$.*

(Note that for condition (ii) to hold it is sufficient that h be bounded above and $\|\nabla_x u^*(x)\| \leq C(1 + \|x\|)$, $x \in \mathbb{R}^m$, since $\mathrm{E}\int_0^t \|X(s)\|^2\,ds < \infty$; see, e.g., Øksendal [30, Theorem 5.2.1].)

Theorem 4. *Let u^* be as in Assumption 3. Then $\alpha = u^*$. Furthermore, if Q^* is the importance distribution under which*

$$dX(t) = \mu^*(X(t))\,dt + \sigma(X(t))\,dB(t), \tag{18}$$

where

$$\mu^*(x) = \mu(x) + \frac{\sigma(x)\sigma(x)^T \nabla_x u^*(x)}{u^*(x)},$$

then

$$W = u^*(X(0)) \quad Q^*\text{–a.s.}$$

REMARK 8. Because μ^* does not necessarily satisfy Lipschitz and growth conditions like the ones imposed on μ and σ, (18) is not immediately guaranteed to have a strong solution. However, letting $\tau_\delta \triangleq \inf\{t \geq 0\colon X_t \in G_\delta\}$ and $T_n \triangleq n \wedge T \wedge \tau_{1/n}$, then for each integer n (18) can be solved strongly on $t < T_n$ (because of Assumption 3). In this way one obtains a solution on $t < T = \lim_{n\to\infty} T_n$; cf. Øksendal [30, Exercise 7.14].

PROOF. That $\alpha = u^*$ is easily shown by applying Itô's formula to $(\beta(t\wedge T)u^*(X(t\wedge T))\colon t \geq 0)$ and using optional sampling. Let τ_δ and T_n be as in Remark 8, and put

$$Y(t) = \int_0^{t\wedge T} f(X(s))\beta(s)L(s)\,ds + L(t\wedge T)\beta(t\wedge T)u^*(X(t\wedge T)),$$

where

$$L(t) = \exp\left(-\int_0^t \eta(X(s))\,dB(s) - \tfrac{1}{2}\int_0^t \|\eta(X(s))\|^2\,ds\right)$$

and

$$\eta(x) = \frac{\sigma(x)^T \nabla_x u^*(x)}{u^*(x)}.$$

Note that $Y(0) = u^*(X(0))$ and on $t < T$ it follows from Itô's formula that, under Q^*,

$$\begin{aligned} dY(t) = {} & L(t)\beta(t)\big[f(X(t)) + \mu(X(t))^T\nabla_x u^*(X(t)) + \operatorname{tr}\big(\tfrac{1}{2}\sigma(X(t))\sigma(X(t))^T\nabla_{xx}u^*\big) \\ & + \eta(X(t))^T\sigma(X(t))^T\nabla_x u^*(X(t)) + h(X(t))u^*(X(t)) - \eta(X(t))^T\sigma(X(t))^T\nabla_x u^*(X(t))\big]\,dt \\ & + L(t)\beta(t)\big[-u^*(X(t))\eta(X(t)) + \sigma(X(t))^T\nabla_x u^*(X(t))\big]^T\,dB(t). \end{aligned}$$

The term in dt on the right-hand side equals zero, since u^* satisfies (16), and the term in $dB(t)$ also vanishes because of the definition of η. Thus, $dY(t) = 0$ on $t < T$, i.e., $Y(t\wedge T) = Y(0) = u^*(X(0))$ for $t \geq 0$. Since, on $\{T < \infty\}$, $Y(t\wedge T) \to W$ as $t \to \infty$, we conclude that $W = u^*(X(0))$ on $\{T < \infty\}$. On the other hand, on $\{T = \infty\}$,

$$Y(t) = u^*(X(0))$$

for $t \geq 0$, and $W = \int_0^\infty f(X(s))\beta(s)L(s)\,ds$, whence

$$L(t)\beta(t)u^*(X(t)) \searrow u^*(X(0)) - W \tag{19}$$

as $t \to \infty$ on $\{T = \infty\}$. Hence, it suffices to show that $I(T = \infty)L(t)\beta(t)u^*(X(t)) \to 0$ Q^*–a.s. to conclude that $W = u^*(X(0))$ Q^*–a.s.

For this purpose note that, for each n, $(\eta(X(t\wedge T_n))\colon t \geq 0)$ is bounded, whence $(L(t\wedge T_n)\colon t \geq 0)$ is a square-integrable martingale under Q^*. Hence, we can apply Girsanov's formula—see, e.g., Karatzas and Shreve [25, Section 3.5]—to conclude that, for $x \in \mathbb{R}^m$,

$$\begin{aligned} & \mathrm{E}_{Q_x^*} L(t)\beta(t)u^*(X(t))I(T_n > t) \\ & \quad = \mathrm{E}_x\, \beta(t)u^*(X(t))I(T_n > t) \\ & \quad \leq \mathrm{E}_x\, \beta(t)u^*(X(t))I(T > t) \\ & \quad = \mathrm{E}_x\left[\beta(t)I(T > t)E_{X(t)}\left[\int_0^T e^{\int_0^s h(X(r))\,dr} f(X(s))\,ds + I(T < \infty)g(X(T))e^{\int_0^T h(X(r))\,dr}\right]\right] \\ & \quad \leq \mathrm{E}_x\left[\beta(t)I(T > t)\,\mathrm{E}_x\left[\int_t^T e^{\int_t^s h(X(r))\,dr} f(X(s))\,ds + I(T < \infty)g(X(T))e^{\int_t^T h(X(r))\,dr}\,\middle|\,(X(v)\colon 0 \leq v \leq t)\right]\right] \\ & \quad \leq \mathrm{E}_x\, I(T > t)\left(\int_t^T \beta(s)f(X(s))\,ds + I(T < \infty)g(X(T))\beta(T)\right) \\ & \quad \longrightarrow 0 \end{aligned}$$

as $t \to \infty$, by dominated convergence. Thus, for arbitrary $\varepsilon > 0$, there exists t_0 such that for $t > t_0$

$$\begin{aligned} \varepsilon & > \mathrm{E}_{Q_x^*} L(t)\beta(t)u^*(X(t))I(T_n > t) \\ & \geq \mathrm{E}_{Q_x^*} L(t)\beta(t)u^*(X(t))I(T_n > t)I(T = \infty) \end{aligned}$$

uniformly in n. Since $I(T_n > t)I(T = \infty) \nearrow I(T = \infty)$ as $n \to \infty$, it follows that $\mathrm{E}_{Q^*} L(t)\beta(t)u^*(X(t)) \cdot I(T = \infty) \le \varepsilon$ for $t \ge t_0$ by monotone convergence. Hence,

$$\mathrm{E}_{Q^*} L(t)\beta(t)u^*(X(t))I(T = \infty) \to 0$$

as $t \to \infty$. Fatou's lemma then implies

$$\liminf_{t\to\infty} I(T = \infty)L(t)\beta(t)u^*(X(t)) = 0,$$

Q_x^*–a.s., which together with (19) gives

$$I(T = \infty)L(t)\beta(t)u^*(X(t)) \to 0$$

Q^*–a.s., as remained to be shown. □

Next we study an expectation with a fixed time horizon (rather than a hitting time). For ease of notation, we restrict attention to the case $d = 1$ (one-dimensional diffusions); no additional complications arise for $d > 1$. Consider an expectation of the form

$$\alpha(t, x) = \mathrm{E}_x \exp\left(\int_0^t h(X(s))\, ds\right) g(X(t)) \tag{20}$$

for given functions $g\colon \mathbb{R} \to [1, \infty)$ and $h\colon \mathbb{R} \to \mathbb{R}$, with h bounded. We note that (20) is a particular case of (15): indeed, $\alpha(t, x)$ is the expectation of the final reward when the (degenerate) diffusion $\tilde{X}(s) = (s, X(s))$ hits the set $[t, \infty) \times \mathbb{R}$. Thus, the form of a Markovian zero-variance importance measure follows directly from Theorem 4, provided we can verify the conditions in that theorem. For this purpose, we introduce the following assumption:

Assumption 4. *There exists a function $u\colon [0, t] \times \mathbb{R} \to \mathbb{R}$ satisfying*
(i) *u is continuous on $[0, t] \times \mathbb{R}$ and twice continuously differentiable on $(0, t) \times \mathbb{R}$;*
(ii) *u is a solution to*

$$\mu(x)\frac{\partial}{\partial x}u(s, x) + \frac{\sigma^2(x)}{2}\frac{\partial^2}{\partial x^2}u(s, x) + h(x)u(s, x) - \frac{\partial}{\partial s}u(s, x) = 0,$$

$0 \le s \le t$, $x \in \mathbb{R}$, *subject to the boundary condition* $u(0, x) = g(x)$;
(iii) $\mathrm{E}\int_0^t \beta^2(s)\sigma^2(X(s))((\partial/\partial x)u(t-s, X(s)))^2 ds < \infty$;
(iv) *the function* $x \mapsto (\sigma(x)/u(s, x))(\partial/\partial x)u(s, x)$ *is bounded on* $[0, t] \times \mathbb{R}$.

(Note that for condition (iii) to hold it is sufficient that there exists $C_1 > 0$ such that $|(\partial/\partial x)u(s, x)| \le C(1 + |x|)$, $x \in \mathbb{R}$, $0 \le s \le t$.)

A Markovian zero-variance importance distribution is then given in the following result.

Corollary 2. *Assume $u(s, x) < \infty$, $(s, x) \in [0, t] \times K^C$, and u satisfies the conditions in Assumption 4. Then $\alpha = u$. Furthermore, if Q^* is the importance measure under which*

$$dX(s) = \mu^*(s, X(s))\, ds + \sigma(X(s))\, dB(s), \tag{21}$$

$s \le t$, where

$$\mu^*(s, x) = \mu(x) + \frac{\sigma^2(x)}{u(t-s, x)}\frac{\partial}{\partial x}u(t-s, x),$$

then

$$W = u(t, X(0)) \quad Q^*\text{–a.s.},$$

where

$$W = \exp\left(\int_0^t h(X(u))\, du\right) g(X(t))L(t),$$

$$L(s) = \exp\left(-\int_0^s \eta(s,\ X(s))\, dB(u) - (1/2)\int_0^s \eta^2(s,\ X(s))\, du\right),$$

and $\eta(s, x) = (\sigma(x)/u(t-s, x))(\partial/\partial x)u(t-s, x)$.

5. Examples. In the last two sections we have noted that the Markovian change-of-measure Q^* associated with filtered estimators of path functionals like (3) and (15) can exhibit behavior that does not arise in the context of estimating rare-event probabilities. Notably, that the original measure may not be absolutely continuous w.r.t. the change-of-measure ($Q^\circ \not\ll Q^*$ or $\mathrm{P} \not\ll Q^*$), even when restricted to paths of finite length on which the desired path functional is positive; that there may be positive probability of nontermination, $Q^*(T=\infty)>0$, and that the zero-variance estimator and associated change of measure may not be unique. In this section we illustrate these issues with some specific expectations that arise in the context of queueing and financial applications.

Because the purpose of these examples is to illustrate properties of the change of measure associated with zero-variance estimators, for the most part we focus on stylized problems for which the desired expectation $\alpha(x)$ is known in closed form (so that there is no real need to estimate it via simulation).

As discussed in the introduction, in practical situations in which $\alpha(x)$ is not known, one would use an approximation to the solution u^* of the linear system to construct a "good" change of measure and estimator. There is no general rule or mechanism to obtain such an approximation: it will typically be constructed on an ad-hoc basis, using the structure of the specific problem under study. In some cases one may use "fluid" or "mean field" heuristics to come up with an approximation. Sometimes one can approximate the process X (using a weak convergence result) by a simpler process for which the expectation can be computed in closed form, and then leverage off this solution as an approximation in the original setting. For example, since a $GI/G/1$ single-server queue in heavy traffic can be approximated by a regulated Brownian motion (RBM), one can use the known solutions u^* for RBM in Examples 13–16 as approximations to the corresponding expectations for a more general $GI/G/1$ queue, and construct a change of measure and estimator based on them; we illustrate such a construction in Example 17. Occasionally one can approximate a time-dependent solution by a time-homogeneous one: for instance, in Example 20 we consider a discrete-time Feynman-Kac-type expectation, and approximate the time-dependent solution to the backward Equations (9) by an (easier to compute) time-homogeneous one, to construct an importance sampling estimator that offers significant variance reduction. Some examples in which approximate zero-variance importance sampling was used to estimate rare event probabilities can be found, for example, in L'Ecuyer and Tuffin [27] and Blanchet and Glynn [3].

Our first six examples involve single-server queues. We start by providing the change of measure to compute the probability of the (rare event) of experiencing very large delays during a busy cycle, both for an $M/M/1$ queue (Example 11) and a "Brownian queue," i.e., a fluid queue modeled by RBM (Example 12). We then compare these to the zero-variance change of measure used to estimate the expectation of other functionals of the queueing process. Specifically, we focus on the expected cumulative "cost" until hitting zero (the end of a busy cycle) for various cost functions. Such expectations can in turn be used to compute the steady-state average cost, by taking advantage of the regenerative structure (cf. the discussion in Remark 6); hence the interest in them.

Example 11 ($M/M/1$ Waiting Times; Probability of Hitting b Before 0). Consider the process $X=(X_n\colon n\ge 0)$, where X_n represents the waiting time of customer n in a single server queue. Let $(V_n\colon n\ge 0)$ and $(\chi_n\colon n\ge 1)$ be two independent sequences of iid exponentially distributed random variables, with $\mathrm{E}\,V_0=\mu^{-1}<\lambda^{-1}=\mathrm{E}\,\chi_1$. Here V_n represents the service requirement of customer n, and χ_{n+1} the interarrival time between customers n and $n+1$ ($n\ge 0$). Assuming a first-come first-served service discipline, the waiting time sequence satisfies the well-known recursion

$$X_{n+1}=\max(0, X_n+Z_{n+1}),$$

where $Z_{n+1}\triangleq V_n-\chi_{n+1}$, $n\ge 0$. Define $T_b\triangleq \inf\{n\ge 1\colon X_n\ge b\}$ and $T_0\triangleq\inf\{n\ge 1\colon X_n=0\}$. For $0\le x\le b$, let $\alpha(x)=\mathrm{P}_x(T_b<T_0)$, the probability that a waiting time of b or more is observed before the buffer becomes empty. This is of the form (3) with $K=\{0\}\cup[b,\infty)$ (so $T=T_b\wedge T_0$), $f(x)=I(x\ge b)$, $g(\cdot,\cdot)=0$ and $\beta(\cdot,\cdot)=1$.

It can be shown that for $0<x<b$,

$$\alpha(x)=\rho\frac{e^{(\mu-\lambda)x}-\rho}{e^{(\mu-\lambda)b}-\rho},$$

where $\rho=\lambda/\mu$. Then, according to Theorem 3, there exists a zero-variance importance measure on $\mathcal{F}_T$ that makes X a Markov chain with one-step transition kernel M^* given by $M^*(x,\{0\})=(1-I(0<x<b))\cdot$

$e^{-\lambda x}\mu/(\lambda+\mu)$ and

$$M^*(x,dy)=\begin{cases}\dfrac{\lambda\mu}{\lambda+\mu}\dfrac{e^{(\mu-\lambda)y}-\rho}{e^{(\mu-\lambda)x}-\rho}[e^{\lambda(y-x)}I(x>y)+e^{-\mu(y-x)}I(y\geq x)]dy, & x,\ y\in(0,b),\\[2ex] \dfrac{\lambda\mu}{\lambda+\mu}\dfrac{e^{(\mu-\lambda)b}-\rho^2}{\rho e^{(\mu-\lambda)x}-\rho^2}e^{-\mu(y-x)}dy, & x\in(0,b),\quad y\geq b,\\[2ex] \dfrac{\lambda\mu}{\lambda+\mu}[e^{\lambda(x-y)}I(x>y)+e^{-\mu(y-x)}I(y\geq x)]dy, & x\in\{0\}\cup[b,\infty),\quad y>0.\end{cases}$$

Equivalently, under the zero-variance importance measure Q^*, the "increment" Z_{n+1} is conditionally independent of $(Z_0,\ldots,Z_n)$ given X_n and for $x\in(0,b)$, $Q^*(Z_{n+1}\in dz\mid X_n=x, T>n)=Q^*_x(Z_1\in dz)$, where

$$Q^*_x(Z_1\in dz)\triangleq\begin{cases}\dfrac{\lambda\mu}{\lambda+\mu}\cdot\dfrac{e^{(\mu-\lambda)x}-\rho e^{-(\mu-\lambda)z}}{e^{(\mu-\lambda)x}-\rho}[e^{\mu z}I(z<0)+e^{-\lambda z}I(z\geq 0)]dz, & z\in(-x,b-x),\\[2ex] \dfrac{\lambda\mu}{\lambda+\mu}\dfrac{e^{(\mu-\lambda)b}-\rho^2}{\rho e^{(\mu-\lambda)x}-\rho^2}\cdot e^{-\mu z}dz, & z\geq b-x.\end{cases}$$

Note that if both b and x are very large, but with x small compared to b, then on $\{X_n=x, T>n\}$ the distribution of Z_{n+1} under Q^* is very close to that of $-Z_1$ under P. That is, when the workload is large, the importance distribution Q^* is close to the distribution under which interarrival times and the service requirements are independent sequences of i.i.d. exponential random variables, but with their means exchanged (compared with those under P). It is well known that the latter corresponds to the distribution of the Z_n's conditioned on the random walk $(S_n=Z_1+\cdots+Z_n\colon n\geq 0)$ eventually hitting level b. The distribution Q^* involves the additional conditioning that S_n not hit zero before hitting b. Note that in this case Q^* coincides with Q°, as the functional $I(T_b<T_0)$ is of the product form discussed in §2. Also, under the importance measure the simulation run terminates in finite time ($Q^*(T<\infty)=1$).

For the next four examples we work with a *Brownian queue*. In this model the buffer content is not measured in discrete units but is rather a continuous quantity (a *fluid queue*), and its evolution is described by regulated Brownian motion (RBM). This model arises naturally as a limit of the $GI/GI/1$ queue in heavy traffic, and is also interesting in its own right. Many of the performance metrics for this model are amenable to closed-form solution, oftentimes yielding very simple formulae. Indeed, it has been argued (Salminen and Norros [32]) that it could be used in textbooks in place of the $M/M/1$ model as a prototype for simple queues.

To be specific, let $Z=(Z(t)\colon t\geq 0)$ be the so-called free process (representing total work arrived minus total work processing capacity by time t), given by

$$Z(t)=-\mu t+\sigma B(t),$$

where $\mu,\sigma>0$, and B is standard Brownian motion. The queueing process X is obtained by applying the regulator mapping to Z:

$$X(t)=Z(t)+X(0)\vee\Big(-\inf_{0\leq s\leq t}Z(s)\Big),\qquad t<0.$$

Throughout this section we denote $T_y=\inf\{t\geq 0\colon X(t)=y\}$ and $\theta^*\triangleq 2\mu/\sigma^2$.

Example 12 (Probability of Hitting b Before 0). Let $\alpha(x)\triangleq \mathrm{P}_x(T_b<T_0)$, $0<x\leq b$, the analog of Example 11 for the Brownian queue. This is of the form (15) with $K=\{0\}\cup[b,\infty)$ (so that $T=T_b\wedge T_0$), $f=0$, $h=0$, and $g(x)=I(x\geq b)$.

It is well known that

$$\alpha(x)=(e^{\theta^*x}-1)(e^{\theta^*b}-1)^{-1};$$

see, e.g., Harrison [21]. Then, under the zero-variance importance measure Q^* on $\mathscr{F}_T$ described in Theorem 4, X has drift $\mu^*(X(t))$, where

$$\mu^*(x)=\mu\cdot\frac{e^{\theta^*x}+1}{e^{\theta^*x}-1},$$

$0<x\leq b$.

Note that if $x=\gamma b$ for some $\gamma\in(0,1)$ and b is very large, then $\mu^*(x)\approx\mu$. It is well known that constant positive drift $\mu(x)=\mu$ arises in the distribution of Brownian motion with negative drift $-\mu$ conditioned on

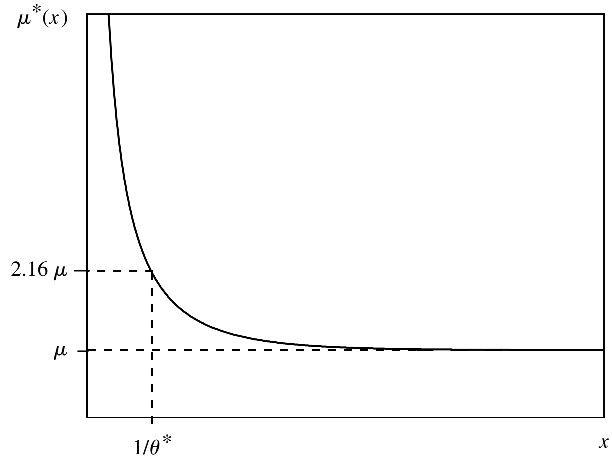

FIGURE 1. Drift as a function of workload under the zero-variance importance measure for $\alpha(x) = \mathrm{P}_x(T_b < T_0)$.

it eventually hitting level b. Here we have the additional conditioning on not hitting zero, which makes the drift increase unboundedly when the workload approaches zero, preventing the buffer from being emptied (see Figure 1). More specifically,

$$\mu^*(x) \sim \frac{\sigma^2}{x}$$

as $x \searrow 0$, a behavior that we will observe in several subsequent examples. As in Example 11, because there is only a "final reward" $I(T_b < T_0)$, we are computing the expectation of a product-form functional, so Q^* coincides with Q°, and $\mathrm{P} \ll Q^*$ when restricted to $\mathscr{F}_T \cap \{T_b < T_0\}$. Also, under the importance measure Q^* the simulation run terminates in finite time ($Q^*(T < \infty) = 1$).

EXAMPLE 13 (EXPECTED TIME TO EMPTY THE BUFFER). Let $\alpha(x) = \mathrm{E}_x T_0$, $x \geq 0$. This is of the form (15) with $K = \{0\}$ (so that $T = T_0$), $f = 1$, $h = 0$, and $g = 0$.

It is well known that

$$\alpha(x) = x/\mu,$$

for $x \geq 0$. Then, under the zero-variance importance measure Q^* of Theorem 4, X has drift $\mu^*(X(t))$, where

$$\mu^*(x) = -\mu + \frac{\sigma^2}{x} I(x > 0).$$

When the workload is small, the dynamics under Q^* are similar to those in Example 12: under Q^* the buffer is prevented from becoming empty, with the drift increasing asymptotically as σ^2/x as $x \searrow 0$. However, for moderate and large workloads the change-of-measure Q^* obtained here differs significantly from the one used in Example 12 to compute $\mathrm{P}_x(T_b < T_0)$: in the previous example the drift remained positive and large (close to μ) even for large workloads, so the process has a steady tendency to increase until hitting level b, at which time

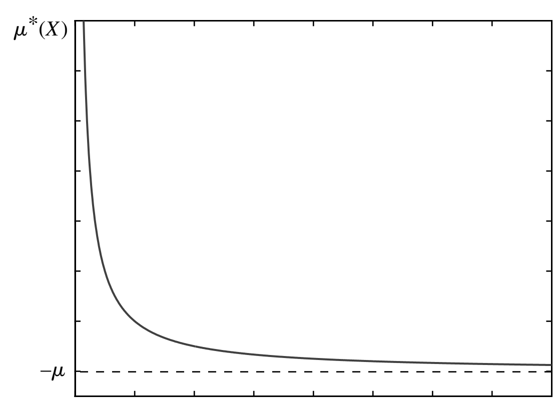

FIGURE 2. Drift as a function of workload under the zero-variance importance measure for $\alpha(x) = \mathrm{E}_x\, T_0$.

the process reverts abruptly to the original dynamics. In contrast, here the drift decreases smoothly to $-\mu$ as the workload increases (see Figure 2), and the importance measure tends to keep the workload fluctuating some moderate distance away from zero, giving it a tendency to revert toward an "equilibrium" value of $2/\theta^*$. Note also that here Q° and P are not absolutely continuous with respect to Q^*, even when restricted to paths of finite length s, i.e., to $\mathcal{F}_s \cap \{T_0 > 0\}$ for some fixed $s > 0$: Q^* assigns probability zero to paths that hit zero in finite time, even though $Z = T_0 > 0$ on such paths, so that Q° would assign positive probability to them.

As just noted, under Q^* the process is prevented from hitting zero, so that $T_0 = \infty$, Q^*–a.s. Thus, although Q^* allows one to construct an estimator with zero variance, computing it requires an infinite run length, which suggests that trying to mimic Q^* may not lead to improved efficiency (even if it does reduce the variance). There is, however, a simple approach to construct a zero-variance estimator for this example that terminates in finite time with probability one: adding a terminal reward when the process exits to zero. Specifically, suppose we make $g(0) = r/\mu$, for some $r > 0$, so that we estimate $\alpha(x) = \mathrm{E}_x T_0 + r/\mu = (x+r)/\mu$ (instead of $\mathrm{E}_x T_0$). With this change, under Q^* the process X has drift $\mu^*(X(t))$, where

$$\mu^*(x) = -\mu + \frac{\sigma^2}{x+r}.$$

This makes the run length T_0 a.s. finite under Q^*. In fact, the expected run length becomes

$$\mathrm{E}_{Q_x^*} T_0 = \frac{x}{\mu} + \frac{2}{\theta\mu}\log(1+x/r) + \frac{2}{\theta^2\mu}\left(\frac{1}{r} - \frac{1}{x+r}\right).$$

Example 14 (Expected Sojourn Over b During a Busy Cycle). Let $\alpha(x) \triangleq \mathrm{E}_x \int_0^{T_0} I_{[b,\infty)}(X(s))\,ds$. This is of the form (15) with $K = \{0\}$ (so $T = T_0$), $f(x) = I_{[b,\infty)}(x)$, $h = 0$, and $g = 0$. Note that the strong Markov property of X implies

$$\alpha(x) = \mathrm{P}_x(T_b < T_0)\alpha(b) + I(x > b)\,\mathrm{E}_x T_b.$$

For the last term, note that for $x > b$, $\mathrm{E}_x T_b = (x-b)/\mu$ (cf. Example 13). In the first term, $\mathrm{P}_x(T_b < T_0) = \min\{1, (e^{\theta^* x} - 1)/(e^{\theta^* b} - 1)\}$ (cf. Example 12); as for $\alpha(b)$, it is known that

$$E_b \exp\left(-\alpha \int_0^{T_0} I_{[b,\infty)}(X(s))\,ds\right) = 2\left(1 + e^{-\theta^* b} + (1 - e^{-\theta^* b})\sqrt{1 + 4\alpha/(\theta^*\mu)}\right)^{-1}$$

for $\alpha > 0$; see, for example, Borodin and Salminen [7, Equation 2.2.4.1]. Hence $\alpha(b) = (1 - e^{-\theta^* b})/(\theta^*\mu)$. It follows that

$$\alpha(x) = \frac{1}{\mu\theta^*}\left[(\theta^*(x-b) + 1 - e^{-\theta^* b})I(x > b) + e^{-\theta^* b}(e^{\theta^* x} - 1)I(x \le b)\right].$$

Thus, under the zero-variance importance distribution of Theorem 4, X has drift $\mu^*(X(t))$ given by

$$\mu^*(x) = \begin{cases} \mu(e^{\theta^* x} + 1)/(e^{\theta^* x} - 1) & x \le b, \\ \mu(1 + e^{-\theta^* b} - \theta^*(x-b))/(1 - e^{-\theta^* b} + \theta^*(x-b)) & x \ge b. \end{cases}$$

Note that on $(0, b)$ the behavior of X under Q^* is the same as that in Example 12. However, the importance measure used in Example 12 would revert to the original drift $-\mu$ immediately upon hitting b, whereas in this case the drift decreases smoothly from μ to $-\mu$ as the workload increases from b to infinity (see Figure 3). Under Q^* the workload has a tendency to revert toward an "equilibrium value" of $b + 1/\theta^*$ (cf. Example 13). Also, once again we find Q° and P are not absolutely continuous with respect to Q^*, even when restricted to $\mathcal{F}_s \cap \{Z > 0\}$ for fixed $s > 0$, for the same reasons as in Example 13. Also, because there are no terminal rewards ($g(0) = 0$), we are once again faced with a situation in which $T_0 = \infty$, Q^*–a.s.; this can be addressed by adding a terminal reward, as in Example 13.

Example 15 (Polynomial Cost). Let $\alpha(x) \triangleq \mathrm{E}_x \int_0^{T_0} X(s)^p\,ds$, for some integer $p \ge 1$. This is of the form (15) with $K = \{0\}$ (so $T = T_0$), $f(x) = x^p$, $h = 0$ and $g = 0$.

Let u be given by

$$u(x) = \frac{1}{(p+1)\mu\theta^{*p+1}} \sum_{j=1}^{p+1} (\theta^* x)^j (p+1)!/j!.$$

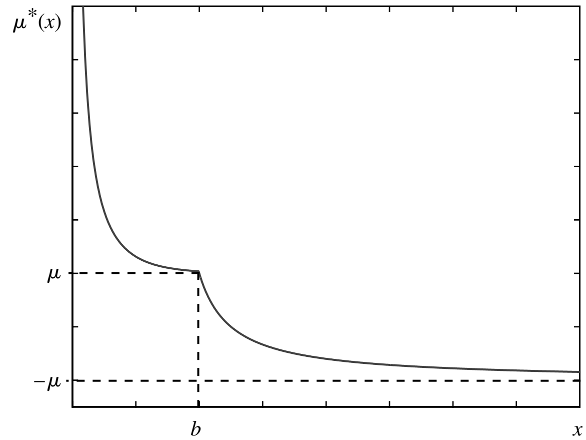

FIGURE 3. Drift as a function of workload under the zero-variance importance measure for $\alpha(x) = \mathrm{E}_x \int_0^{T_0} I_{[b,\infty)}(X(s))\,ds$.

Then $\alpha(x) = u(x)$, $x \geq 0$. To see this, apply Itô's formula to $u(X(t))$, using that u satisfies $-\mu u'(x) + (\sigma^2/2)u''(x) = -x^p$, and that $\sigma \int_0^t u'(X(s))\,dB(s)$ is a square-integrable martingale.

Then, under the change of measure described in Theorem 4, X has drift $\mu^*(X(t))$, where

$$\mu^*(x) = \mu + \frac{2\mu[(p+1)! - (\theta^* x)^{p+1}]}{\sum_{j=1}^{p+1} (\theta^* x)^j (p+1)!/j!} I(x > 0).$$

The evolution of the buffer content under the importance distribution is qualitatively similar to that found in Examples 13 and 14 (and hence, qualitatively different from that under the change of measure used to compute the probability $\mathrm{P}_x(T_b < T_0)$ in Example 12): under Q^* the buffer content has a mean-reverting behavior centered around "moderate" values. Also $\mu^*(x) \searrow -\mu$ as $x \to \infty$, so that for large workloads the dynamics under the importance distribution are similar to those of the original process, whereas at low workload levels $\mu^*(x) \sim \sigma^2/x$ as $x \to 0$, preventing the buffer from ever becoming empty. Thus, $Q^\circ \not\ll Q^*$ on $\mathcal{F}_s \cap \{Z > 0\}$, for the same reasons as in the previous two examples. Also, $T_0 = \infty$ Q^*–a.s., but as in the previous two examples this can be addressed by adding a terminal reward.

EXAMPLE 16 (EXPONENTIAL COST). Let $\alpha(x) \triangleq \mathrm{E}_x \int_0^{T_0} \exp(\theta X(s))\,ds$, for $0 \neq \theta < \theta^*$. This is of the form (15) with $T = T_0$, $f(x) = e^{\theta x}$, $h = 0$ and $g = 0$.

It can be shown that α is given by

$$\alpha(x) = \frac{e^{\theta x} - 1}{(1 - \theta/\theta^*)\theta\mu}.$$

Thus, under the change of measure of Theorem 4, X has drift $\mu^*(X(t))$, where

$$\mu^*(x) = \mu\left[-1 + 2\frac{\theta e^{\theta x}}{\theta^*(e^{\theta x} - 1)}\right].$$

The behavior near the origin under the importance distribution is similar to that of the previous examples, namely, $\mu^*(x) \sim \sigma^2/x$ as $x \searrow 0$. In contrast, for $\theta > 0$, $\lim_{x\to\infty} \mu^*(x) > -\mu$. Hence, the dynamics under the importance distribution differ significantly from the original process, even for large buffer occupancies. In particular, if $\theta > \theta^*/2$, then $\liminf_{x\to\infty} \mu^*(x) > 0$ and in this case, $X(t) \to \infty$ Q^*–a.s. as $t \to \infty$. Thus, the dynamics under Q^* in this case can differ significantly from those seen in all previous examples. Once more we have $Q^\circ \not\ll Q^*$ on $\mathcal{F}_s \cap \{Z > 0\}$, for the same reasons as in the previous three examples. Once more we are faced with a situation in which $T_0 = \infty$ Q^*–a.s., but unlike the previous examples this is not just a consequence of the behavior near the origin, but rather because of the form of the reward rate away from the origin. Adding a terminal reward would not guarantee termination: even if one adds a terminal reward $g(0) > 0$, it would still be the case that $Q^*(T < \infty) < 1$, so there would be a positive probability of nontermination.

Our next example illustrates the construction of a change of measure and filtered estimator based on an approximation to the (unkown) desired solution u^* to the linear system characterizing the expectation. The estimator thus constructed would not have zero variance, but one would expect its variance to be small.

Example 17 ($GI/G/1$ Queue in Heavy Traffic). Consider again the Markov chain $X = (X_n\colon n \geq 0)$ representing the sequence of waiting times in a single-server queue, constructed as in Example 11 but allowing the iid sequences $(V_n\colon n \geq 0)$ and $(\chi_n\colon n \geq 1)$ to have arbitrary distributions (not necessarily exponential). Without loss of generality we assume that $\mathrm{E}\,\chi_i = 1$ and $\mathrm{E}\,V_i = \rho$. We assume $\rho < 1$ (so the queueing process is stable); furthermore, we assume that $\mathrm{E}\,e^{\gamma V_1} < \infty$ for some $\gamma > 0$. As before we define the increment $Z_i \triangleq V_i - \chi_{i+1}$, and introduce the notation $\varphi(\gamma) \triangleq \mathrm{E}\,e^{\gamma Z_1}$, $F(z) \triangleq \mathrm{P}(Z_1 \leq z)$ and $\sigma^2 = \operatorname{var} Z_1$.

Suppose we are interested in estimating the expectation

$$\alpha(x) = \mathrm{E}_x \sum_{j=0}^{T_0} e^{\theta X_j},$$

where $T_0 = \inf\{n \geq 0\colon X_n = 0\}$ and $0 < \theta < \bar{\theta} \triangleq \sup\{\gamma\colon (1-\varphi^2(\gamma))/[\gamma\varphi(\gamma)F(0)(1-\gamma\sigma^2/(2(1-\rho)))] \geq 1-\rho\}$. Note this is of the form (3) with $f(x) = e^{\theta x}$, $g = 0$, $\beta = 1$ and $K = \{0\}$.

We are interested in the case in which ρ is close to one, so the queue is in *heavy traffic*. It is well known that in this setting the waiting time chain X above can be approximated by RBM. The approximation is obtained by embedding X in a sequence $\{X^r\colon r > 0\}$, where $X^r = (X_n^r\colon n \geq 0)$ is the waiting time chain for a $GI/G/1$ queue with traffic intensity ρ_r, satisfying $(1-\rho_r)\sqrt{r} \to \mu > 0$ as $r \to \infty$. Define the centered and scaled process $\hat{X}^r$ by $\hat{X}^r(t) = X^r_{\lfloor rt \rfloor}/\sqrt{r}$. It is known that $\hat{X}^r$ converges weakly to RBM with drift $-\mu$ and variance parameter σ^2 (in the space D of real-valued càdlàg functions on $[0,\infty)$ equipped with the Skorohod J_1 topology); see, e.g., Whitt [36, Theorem 9.3.1].

This weak convergence result suggests that one can approximate the (unknown) desired solution u^* to (4) by the (known) corresponding expectation for RBM. Specifically, let $u^*_{\mathrm{RBM}}(x, \theta, \mu, \sigma^2)$ be the solution for the expected cumulative exponential cost until the end of a busy cycle for RBM, as described in Example 16; that is

$$u^*_{\mathrm{RBM}}(x, \theta, \mu, \sigma^2) = \frac{e^{\theta x} - 1}{(1-\theta/\theta^*)\theta\mu},$$

where $\theta^* = 2\mu/\sigma^2$. Then, if the desired $GI/G/1$ waiting time Markov chain X has traffic intensity ρ, choosing (large) r and μ such that $\mu/\sqrt{r} = 1-\rho$ gives the following approximation:

$$\begin{aligned}
\mathrm{E}_x \sum_{j=0}^{T_0} e^{\theta X_j} &= r\,\mathrm{E}\bigg(\int_0^{T_0/r} \exp\big(\theta\sqrt{r}\hat{X}^r(s)\big)\,ds \mid \hat{X}^r(0) = x/\sqrt{r}\bigg) \\
&\approx u^*_{\mathrm{RBM}}(x/\sqrt{r}, \theta\sqrt{r}, (1-\rho)\sqrt{r}, \sigma^2) \\
&= \frac{e^{\theta x} - 1}{\theta(1-\rho)(1-\sigma^2\theta/(2(1-\rho)))}.
\end{aligned}$$

Note the constant term in the numerator is a consequence of the boundary condition $u^*_{\mathrm{RBM}}(0, \theta, \mu, \sigma^2) = 0$. However in the discrete-time setting $u^*(0) = 1$, so we drop this term from the approximation. Hence, the RBM limit suggests using the following approximation u to the desired solution u^* to (4):

$$u(x) = c_0 e^{\theta x},$$

$x > 0$, and $u(0) = 1$, where $c_0 \triangleq [\theta(1-\rho)(1-\sigma^2\theta/(2(1-\rho)))]^{-1}$. This approximation u can be used to construct the change-of-measure Q and estimator W as in (5). Under Q the increment Z_{n+1} is conditionally independent of $(Z_0, \ldots, Z_n)$ given X_n, and, for $x > 0$, $Q(Z_{n+1} \in dz \mid X_n = x, T_0 > n) = Q_x(Z_1 \in dz)$, which is given by

$$Q_x(Z_1 \in dz) = \begin{cases} (c_0/m_0)\exp(\theta(x+z))F(dz), & z > -x, \\ (1/m_0)F(dz), & z \leq -x, \end{cases}$$

where $m_0 \triangleq c_0 \exp(\theta x) \int_{-x}^{\infty} \exp(\theta z) F(dz) + F(-x)$. Note this is (almost) an (state-dependent) *exponential tilt* of the original distribution of the increment. One would expect the estimator W thus constructed to have low variance. (We compute a bound on its variance in Example 21.)

Our next example illustrates the point raised in Remark 2: if one constructs the filtered estimator and associated change of measure based on a solution u to the linear system that is not u^*, then it is always the case that $Q(T = \infty) > 0$, even if $Q^*(T = \infty) = 0$.

Example 18 (Simple Random Walk with Exponential Cost). Let X be a regulated simple random walk; that is, X is a DTMC in $\mathbb{Z}_+$ with one-step transition matrix P given by $P(x, x+1) = p < 1/2$, $P(x, (x-1)^+) = q = 1-p$ and $P(x, y) = 0$ for $y \notin \{x+1, (x-1)^+\}$. Let $\alpha(x) = \mathrm{E}_x \sum_{j=0}^{T_0} \gamma^{X_j}$, where $T_0 \triangleq \inf\{n \geq 0\colon X_n \leq 0\}$, γ is a constant satisfying $0 < \gamma < q/p$ and $\gamma \neq 1$. This is of the form (3) with $K = \{0\}$ (so $T = T_0$), $f(x) = \gamma^x$, $g = 0$ and $\beta = 1$.

Let $c = \gamma/(\gamma - p\gamma^2 - q)$. For any $a > 0$, the function u given by

$$u(x) = c(\gamma^x - 1) + a((q/p)^x - 1) + 1$$

for $x \geq 0$ and $u(x) = 1$ for $x \leq 0$, is a solution to (4). The expectation α corresponds to the minimal nonnegative solution u^*, which is obtained by setting $a = 0$ above.

If one uses the function u above to build the change-of-measure Q and associated filtered estimator W, then the one-step transition matrix of X under Q is M given by

$$M(x, x+1) = \frac{p[c(\gamma^{x+1} - 1) + a((q/p)^{x+1} - 1) + 1]}{(c-1)(\gamma^x - 1) + a((q/p)^x - 1)},$$

$$M(x, x-1) = \frac{q[c(\gamma^{x-1} - 1) + a((q/p)^{x-1} - 1) + 1]}{(c-1)(\gamma^x - 1) + a((q/p)^x - 1)},$$

for $x > 0$. If one uses u^* to construct the change of measure (i.e., set $a = 0$), then if $\gamma < 1$ one has $M(x, x-1) \to q$ as $x \to \infty$, whereas if $1 < \gamma < q/p$ then $M(x, x-1) \to q/(p\gamma^2 + q)$ as $x \to \infty$. In particular, $Q^*_x(T_0 < \infty) = 1$ for $\gamma < \sqrt{q/p}$, whereas $Q^*_x(T_0 < \infty) < 1$ if $\sqrt{q/p} < \gamma < q/p$.

In contrast, if one uses another solution u to (4), $u \neq u^*$ to construct the change of measure (i.e., $a > 0$), then $M(x, x-1) \to p$ as $x \to \infty$. In particular, $Q(T_0 = \infty) > 0$ (even for $\gamma < 1$), as noted in Remark 2. Also, on paths on which $\{T < \infty\}$, the estimator returns $u(x)$, yielding an arbitrarily large relative error: for $\gamma < 1$ one has $u^*(x) \leq c$, whereas $u(x) \to \infty$ as $x \to \infty$. Also, from the discussion in §3 it follows that $Q_x(T < \infty) \leq u^*(x)/u(x)$, which can be arbitrarily small.

Time-dependent expectations of a Markov process of the type considered in Corollaries 1 and 2 arise in the Kolmogorov backward equations, and are of interest in many financial applications. In particular, in the context of pricing derivatives, the value of an option can often be expressed as an expectation of the above kind (perhaps after space augmentation, for some path-dependent derivatives). In our next example we illustrate the form of the change of measure presented in Corollary 2 when the expectation of interest corresponds to the price of a "plain vanilla" European call option in the well-known Black-Scholes model. Finally, in Example 20 we consider an expected final reward in discrete time for which the solution u to the associated linear system is unknown, and discuss one way to construct an approximation on which to base the change of measure.

Example 19 (Black-Scholes Model). Consider the classical Black-Scholes model of a market consisting of a bond paying deterministic interest rate r and an asset whose price process $X = (X(t)\colon t \geq 0)$ is described by geometric Brownian motion. We are interested in $\alpha(x, t)$, the price of a European call option on the asset with strike price c and maturity t, when the initial price of the asset is $x > 0$. The price of the option can be computed as the expectation

$$\alpha(t, x) = \mathrm{E}_x(e^{-rt}(X_t - c)^+),$$

where E_x is the expectation under the "risk-neutral" probability P_x, under which

$$dX(t) = rX(t)\,dt + \sigma X(t)\,dB(t).$$

This is of the form (20) with $g(x) = (x-c)^+$ and $h(x) = -r$. For this model, $\alpha(t, x)$ is given by the well-known Black-Scholes formula

$$\alpha(t, x) = x\Phi\left(\frac{\log(x/c) + rt + \sigma^2 t/2}{\sigma\sqrt{t}}\right) - ce^{-rt}\Phi\left(\frac{\log(x/c) + rt - \sigma^2 t/2}{\sigma\sqrt{t}}\right),$$

where Φ denotes the standard Gaussian cumulative distribution function.

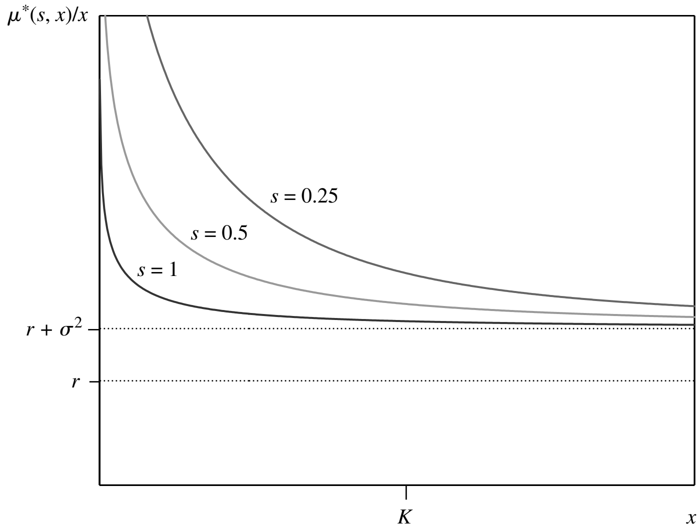

Figure 4. Drift divided by price as a function of price, under the zero-variance importance measure for the Black-Scholes model.
Note. The three curves differ in their time to maturity.

Under the change of measure described in Corollary 2, X has drift $\mu^*(s, X(s))$, where

$$\mu^*(s,x)=xr+\frac{x\sigma^2\Phi\Big(\dfrac{\log(x/c)+r(t-s)+\sigma^2(t-s)/2}{\sigma\sqrt{t-s}}\Big)}{\Phi\Big(\dfrac{\log(x/c)+r(t-s)+\sigma^2(t-s)/2}{\sigma\sqrt{t-s}}\Big)-(c/x)e^{-r(t-s)}\Phi\Big(\dfrac{\log(x/c)+r(t-s)-\sigma^2(t-s)/2}{\sigma\sqrt{t-s}}\Big)},$$

or, putting $x=ce^{-r(t-s)}e^y$,

$$\mu^*(s,x)=xr+x\sigma^2\frac{\Phi\Big(\dfrac{y+\sigma^2(t-s)/2}{\sigma\sqrt{t-s}}\Big)}{\Phi\Big(\dfrac{y+\sigma^2(t-s)/2}{\sigma\sqrt{t-s}}\Big)-e^{-y}\Phi\Big(\dfrac{y-\sigma^2(t-s)/2}{\sigma\sqrt{t-s}}\Big)}.$$

Observe that for $x<c$, $\lim_{s\nearrow t}\mu^*(s,x)/x=\infty$, i.e., the importance measure tries to ensure that the option is "in the money" at maturity. In this sense, the change of measure is similar to the one that we would use to compute the probability that the option is exercised (which is a rare event if $x<c$ and t small). However, the two importance distributions differ significantly away from this boundary: Note that, for fixed $s<t$, $\mu^*(s,x)/x\searrow r+\sigma^2$ as $x\nearrow\infty$ (see Figure 4). Thus, when the stock price is large, X behaves (locally) as a geometric Brownian motion with growth rate close to $r+\sigma^2$. We see that, under Q^*, X has a growth rate that is always significantly greater than r (the growth rate under the risk-neutral distribution). In contrast, the change of measure used to estimate the probability that the option is exercised would make X have a drift very close to r when the stock price is large.

Example 20 (Approximate Zero-Variance Simulation of Feynman-Kac-Type Expectations). Suppose one is interested in estimating an expected final reward with state-dependent discounting, of the form

$$\alpha_n(x)=\mathrm{E}_x\bigg(\exp\bigg(\sum_{j=0}^{n-1}h(X_j)\bigg)f(X_n)\bigg).$$

Note α_n is of the form (8) where, for $0\le j\le n-1$, $f_j=0$, $g_{j+1}=0$, $\beta_{j+1}(x,y)=e^{h(x)}$, and $f_n=f$. For simplicity, we assume X lives in discrete state space. As given by Corollary 1, a zero-variance filtered estimator exists under the change-of-measure Q^*, which makes X evolve, at time j, according to the transition kernel

$$M_j^*(x,y)=\begin{cases}e^{h(x)}P(x,y)\alpha_{n-j}(y)/\alpha_{n-j+1}(x) & \text{if } \alpha_{n-j+1}(x)>0\\ P(x,y) & \text{if } \alpha_{n-j+1}(x)=0.\end{cases}$$

Suppose n is large, so that finding the time-dependent solution $(u(j,\cdot)\colon j\le n)$ to the backward Equations (9) is computationally expensive. For fixed j and n large, one can intuitively expect the ratio $u_{n-j}(y)/u_{n-j+1}(x)$ to be roughly independent of j, and this suggests using an importance measure with a time-homogeneous kernel, based on a "large n" time-independent approximation to u. More specifically, put $K(x,y)=e^{h(x)}P(x,y)$, and let λ, v be the Perron-Frobenius eigenvalue and eigenvector such that

$$Kv=\lambda v.$$

Choose the importance measure Q under which X has transition kernel M, where

$$M(x,y)=\frac{K(x,y)v(y)}{\lambda v(x)}=e^{h(x)}\frac{P(x,y)v(y)}{\lambda v(x)}.$$

Note that

$$\alpha_n(x)=\mathrm{E}_x\left(\exp\left(\sum_{j=0}^{n-1}h(X_j)\right)f(X_n)\right)=\mathrm{E}_x^Q\,\frac{v(X_0)}{v(X_n)}f(X_n)\lambda^n.$$

Thus, approximating $u(j,\cdot)$ by $v(\cdot)$, $j\le n$, leads to an importance sampling estimator based on simulating replicates of $W=(v(X_0)/v(X_n))f(X_n)\lambda^n$ under Q. Note that the variance of the naive Monte-Carlo estimator is

$$\mathrm{var}_P\left(\exp\left(\sum_{j=0}^{n-1}h(X_j)\right)f(X_n)\right)=\tilde{\lambda}^n\,\mathrm{E}_{\tilde{Q}_x}\,\frac{\tilde{v}(X_0)}{\tilde{v}(X_n)}f^2(X_n)-\lambda^{2n}\left(\mathrm{E}_{Q_x}\,\frac{v(X_0)}{v(X_n)}f(X_n)\right)^2,$$

where $\tilde{\lambda}$ and $\tilde{v}$ are the Perron-Frobenius eigenvalue and eigenvector of the kernel $\tilde{K}=(e^{2h(x)}P(x,y)\colon x,y\in S)$, solving

$$\tilde{K}\tilde{v}=\tilde{\lambda}\tilde{v},$$

and $\tilde{Q}$ is the measure under which X has transition kernel $\tilde{M}(x,y)=e^{2h(x)}((P(x,y)\tilde{v}(y))/(\tilde{\lambda}\tilde{v}(x)))$. In contrast,

$$\mathrm{var}_Q\,W=\lambda^{2n}\left[\mathrm{E}_{Q_x}\left(\frac{v(X_0)}{v(X_n)}f(X_n)\right)^2-\left(\mathrm{E}_{Q_x}\,\frac{v(X_0)}{v(X_n)}f(X_n)\right)^2\right].$$

Thus, for large n, using the proposed importance sampling scheme provides several advantages: First, one does not need to compute the time-dependent solutions $(u(j,\cdot)\colon j\le n)$, but rather solve only for the approximation v, which is significantly less computationally intensive. Second, the scheme provides exponential (in n) variance reduction compared to naive Monte Carlo. Additionally, one can estimate α_n parametrically in n (because for large n the estimator takes the form λ^n times an expectation that converges to a steady-state value). Note also that for this problem there is no obvious rare event associated with the expectation.

6. Performance of approximate zero-variance importance sampling. As noted in previous sections, the zero-variance importance sampling estimators and corresponding change of measure cannot be directly implemented in practice, because they require knowledge of the solution to the linear system that characterizes the desired expectation, such as (4) or (16), and hence, in particular, knowledge of the expectation one wants to estimate. Nevertheless, the results in the previous sections can offer guidance to the simulationist on how to construct a good importance sampler: if an approximation u to the desired solution of the linear system is available, then one can use such an approximation to construct the change of measure and related estimator W as in (5), that is, do *approximate zero-variance* importance sampling—cf. L'Ecuyer and Tuffin [27]. When studying the performance of such an estimator, there are two aspects related to the efficiency of the estimator on which one would like to have guarantees:

- Completion time–how long does it take to compute the estimator W? In particular, does the algorithm return an estimate in finite time (i.e., is it the case that $Q(T<\infty)=1$?)?
- What is the MSE of the estimator? (It would be zero if using u^* to construct the estimator; by how much has it increased because of using an approximation?)

We address the first of these issues in §6.1 and the second in §§6.2 and 6.3.

6.1. The need to enforce finite completion time. It is apparent from the discussion and examples in previous sections that, for some problems, computing the estimators presented earlier may take an infinite amount of (simulated) time under the importance measure Q; that is, for some problems, one has $Q(T=\infty)>0$. This can happen even in situations in which $\mathrm{P}(T=\infty)=0$, and even when using the desired solution u^* to build the estimator. This problem often arises in situations in which $u^*=0$ on K; in those cases one may get around this issue by putting the "rewards" in the transitions, as in Example 10, or adding an artificial final reward, as in Example 13; however, this behavior can also be unrelated to the boundary condition, but rather be a consequence of the form of the rewards within K^C, as in Examples 16 and 18 (and in such case adding terminal rewards will not guarantee finite completion time). The same issues arise when using an approximation to build the change or measure.

Having a positive probability of an infinite completion time is, of course, undesirable in itself. But it also brings an additional problem: if one restricts attention to $\{T<\infty\}$, then the estimator may be biased.

When doing approximate zero-variance importance sampling, the approximation to u^* that one is using may well be (or be close to) another nonnegative solution u to (4), $u\neq u^*$. It follows from the analysis in §3 that, when using such a solution u to construct Q and W, it is always the case that $Q(T=\infty)>0$. Moreover, on $\{T<\infty\}$, $W=u(x)\geq u^*(x)$, Q_x–a.s. Hence, if running multiple replications, a positive fraction of them will not return a solution in finite time, whereas those that do terminate will return an answer that is always over biased.

The above point is relevant also when implementing adaptive algorithms like those in Ahamed et al. [1], Kollman et al. [26], and Baggerly et al. [2]. Such algorithms are constructed to converge to a change of measure that corresponds to a solution u to (4). If they converge to the wrong solution, then a biased estimator may result.

The above discussion leads to one important policy recommendation of this study: the necessity to explicitly check that the algorithm constructed terminates in finite time a.s. A natural way to do this is by verifying a Lyapunov condition: find a function $v\geq 0$ and $\varepsilon>0$ such that

$$\int v(y)M(x,dy)\leq v(x)-\varepsilon,$$

$x\in K^C$; see, e.g., Meyn and Tweedie [28, Theorem 11.3.4]. When doing approximate zero-variance importance sampling, not verifying such condition puts the simulationist at a real risk of having a nonterminating algorithm, and of having a biased estimator when restricted to paths that terminate in finite time.

6.2. Lyapunov bounds on the variance in the DTMC setting. When doing approximate zero-variance importance sampling, the resulting estimator W will not have zero variance; one expects it to have low variance, but would like to have a way to assess how large its variance and MSE are. In this section we develop Lyapunov bounds for the MSE of the estimator, in the same spirit as those developed in Blanchet and Glynn [3] for estimators of rare-event probabilities. We work in the same setting as in §§2 and 3.

Suppose u is an approximation to u^* satisfying the conditions in Theorem 1. Because W is an unbiased estimator of $u^*(x)$, it follows that its MSE under Q_x is

$$\mathrm{E}_{Q_x}(W-u^*(x))^2=\mathrm{var}_{Q_x}W=m(x)-u^{*2}(x),$$

where

$$m(x)\triangleq \mathrm{E}_{Q_x}W^2.$$

Because only the second moment $m(x)$ depends on the approximation u, we focus on bounding this term.

Theorem 5. *The function m is the minimal nonnegative solution to*

$$m(x)=r(x)+\int_S H(x,dy)m(y), \tag{22}$$

where

$$H(x,dy)\triangleq\frac{w(x)-f(x)}{g(x,y)+u(y)}\beta(x,y)P(x,dy),$$

for $x\in K^C$, $y\in S$; $H(x,dy)=0$ for $x\in K$; and $r(x)\triangleq 2f(x)u^*(x)-f^2(x)+\int_S H(x,dy)[g^2(x,y)+2g(x,y)u^*(y)]$. (Note $r(x)=f^2(x)$ for $x\in K$.)

This solution can be characterized as

$$m = \sum_{j=0}^{\infty} H^j r, \tag{23}$$

where $H^0(x, dy) = \delta_x(dy)$, $H^j(x, D) = \int_S H^{j-1}(x, dy)H(y, D)$, *for* $D \in \mathscr{S}$, *and* $(H^j r)(x) = \int_S H^j(x, dy)r(y)$.

Proof. Let $A, A_i(\cdot)$ be as in the proof of Theorem 1; the argument there shows that $\mathrm{P}_x \ll Q_x$ on $\{(X_1, \ldots, X_i) \in A_i(x)\} \cap \mathscr{F}_{i\wedge T}$. Let

$$B_i = \prod_{k=1}^{i} \beta(X_{k-1}, X_k), \quad B_{i,j} = \prod_{k=i+1}^{j} \beta(X_{k-1}, X_k) \quad \text{and} \quad L_{i,j} = \prod_{k=i+1}^{j} l(X_{k-1}, X_k).$$

Note

$$W = f(X_0) + \sum_{i=1}^{\infty}[f(X_i) + g(X_{i-1}, X_i)]B_i L_i I(T \geq i),$$

and, because every term is nonnegative,

$$\begin{aligned} W^2 = {} & f^2(X_0) + 2f(X_0)\sum_{i=1}^{\infty}\big[f(X_i) + g(X_{i-1}, X_i)\big]B_i L_i I(T \geq i) \\ & + \sum_{i=1}^{\infty}\big[f(X_i) + g(X_{i-1}, X_i)\big]^2 B_i^2 L_i^2 I(T \geq i) \\ & + 2\sum_{i=1}^{\infty}\big[f(X_i) + g(X_{i-1}, X_i)\big]B_i L_i I(T \geq i)\sum_{j>i}[f(X_j) + g(X_{j-1}, X_j)]B_j L_j I(T \geq j). \end{aligned}$$

Hence,

$$\begin{aligned} \mathrm{E}_{Q_x} W^2 = {} & f^2(x) + 2f(x)\,\mathrm{E}_{Q_x}\sum_{i=1}^{\infty}\big[f(X_i) + g(X_{i-1}, X_i)\big]B_i L_i I(T \geq i) \\ & + \mathrm{E}_{Q_x}\sum_{i=1}^{\infty}\big[f(X_i) + g(X_{i-1}, X_i)\big]^2 B_i^2 L_i^2 I(T \geq i) \\ & + 2\sum_{i=1}^{\infty}\mathrm{E}_{Q_x}\left\{\big[f(X_i) + g(X_{i-1}, X_i)\big]B_i^2 L_i^2 I(T \geq i)\,\mathrm{E}_{Q_x}\left[\sum_{j>i}[f(X_j) + g(X_{j-1}, X_j)]B_{i,j} L_{i,j} I(T \geq j) \mid \mathscr{F}_i\right]\right\} \\ = {} & f^2(x) + 2f(x)(u^*(x) - f(x)) + \mathrm{E}_{Q_x}\sum_{i=1}^{\infty}\big[f(X_i) + g(X_{i-1}, X_i)\big]^2 B_i^2 L_i^2 I(T \geq i) \\ & + 2\,\mathrm{E}_{Q_x}\sum_{i=1}^{\infty}\big[f(X_i) + g(X_{i-1}, X_i)\big]B_i^2 L_i^2 I(T \geq i)(u^*(X_i) - f(X_i)), \end{aligned}$$

where the last follows from Theorem 1. Rearranging,

$$\begin{aligned} \mathrm{E}_{Q_x} W^2 = {} & 2f(x)u^*(x) - f^2(x) + \sum_{i=1}^{\infty}\mathrm{E}_{Q_x}\big[2f(X_i)u^*(X_i) - f^2(X_i)\big]B_i^2 L_i^2 I(T \geq i) \\ & + \sum_{i=1}^{\infty}\mathrm{E}_{Q_x}\big[g^2(X_{i-1}, X_i) + 2g(X_{i-1}, X_i)u^*(X_i)\big]B_i^2 L_i^2 I(T \geq i) \\ = {} & 2f(x)u^*(x) - f^2(x) + \mathrm{E}_{Q_x}\sum_{i=1}^{\infty}\big[2f(X_i)u^*(X_i) - f^2(X_i)\big]B_i^2 L_i^2 I(T \geq i) \\ & + \sum_{i=1}^{\infty}\mathrm{E}_{Q_x}\,\tilde{g}(X_{i-1})B_{i-1}^2 L_{i-1}^2 I(T \geq i) \end{aligned}$$

(since $\{T \geq i\} \in \mathscr{F}_{i-1}$), where

$$\tilde{g}(z) \triangleq \int_S \big[g^2(z, y) + 2g(z, y)u^*(y)\big]\beta^2(z, y)l^2(z, y)M(z, dy),$$

$z \in K^C$. Note $g^2(z, y) + 2g(z, y)u^*(y) = 0$ if $z \notin A$ or $y \notin A_1(z)$. Hence, $\tilde{g}(z) = 0$ for $z \notin A$, and

$$\begin{aligned}\tilde{g}(z) &= \int_{A_1(z)} [g^2(z, y) + 2g(z, y)u^*(y)]\beta^2(z, y)l^2(z, y)M(z, dy)\\ &= \int_{A_1(z)} [g^2(z, y) + 2g(z, y)u^*(y)]\beta^2(z, y)l(z, y)P(z, dy)\\ &= \int_S [g^2(z, y) + 2g(z, y)u^*(y)]\beta^2(z, y)l(z, y)P(z, dy)\end{aligned}$$

for $z \in K^C$. Note also that $\tilde{g}(X_i)I(T > i) = \tilde{g}(X_i)I(T \geq i)$, since $\tilde{g}(z) = 0$ for $z \in K$. It follows that

$$\begin{aligned}\mathrm{E}_{Q_x} W^2 &= 2f(x)u^*(x) - f^2(x) + \tilde{g}(x) + \sum_{i=1}^{\infty} \mathrm{E}_{Q_x}[2f(X_i)u^*(X_i) - f^2(X_i) + \tilde{g}(X_i)]B_i^2 L_i^2 I(T \geq i)\\ &= r(x) + \sum_{i=1}^{\infty} \mathrm{E}_{Q_x} r(X_i) B_i^2 L_i^2 I(T \geq i)\\ &= r(x) + \sum_{i=1}^{\infty} \int_{K^C \times \cdots \times K^C \times S} M(x, dx_1) \cdots M(dx_{i-1}, dx_i) r(x_i) \prod_{j=1}^{i} \beta^2(x_{j-1}, x_j) l^2(x_{j-1}, x_j)\\ &= r(x) + \sum_{i=1}^{\infty} \int_{A_i(x)} M(x, dx_1) \cdots M(dx_{i-1}, dx_i) r(x_i) \prod_{j=1}^{i} \beta^2(x_{j-1}, x_j) l^2(x_{j-1}, x_j)\\ &= r(x) + \sum_{i=1}^{\infty} \int_{A_i(x)} P(x, dx_1) \cdots P(dx_{i-1}, dx_i) r(x_i) \prod_{j=1}^{i} \beta^2(x_{j-1}, x_j) l(x_{j-1}, x_j)\\ &= r(x) + \sum_{i=1}^{\infty} \int_{A_i(x)} H(x, dx_1) \cdots H(dx_{i-1}, dx_i) r(x_i)\\ &= r(x) + \sum_{i=1}^{\infty} \int_{S \times \cdots \times S} H(x, dx_1) \cdots H(dx_{i-1}, dx_i) r(x_i)\\ &= r(x) + \sum_{i=1}^{\infty} (H^i r)(x),\end{aligned}$$

giving (23). Using (23), it is easy to show that m must solve (22). □

The corollary below provides Lyapunov bounds for the second moment $m(x)$. We need the following well-known lemma.

Lemma 1. *Let $\tilde{m}$: $S \to \mathbb{R}^+$ be defined as*

$$\tilde{m} = \sum_{j=0}^{\infty} \tilde{H}^j \rho,$$

where ρ: $S \to [0, \infty)$ and $\tilde{H}$ is a kernel function such that $\tilde{H}(x, \cdot)$ is a finite measure on $(S, \mathscr{S})$ for $x \in S$. Suppose there exists a finite-valued function $v \geq 0$ satisfying the Lyapunov inequality

$$v(x) \geq \rho(x) + (\tilde{H}v)(x), \tag{24}$$

$x \in S$. Then, $\tilde{m} \leq v$.

Proof. Because v is finite valued, it follows from (24) that $\tilde{H}v$ is finite valued. Hence,

$$\rho \leq v - \tilde{H}v,$$

whence $\tilde{H}\rho \leq \tilde{H}v \leq v$, and using induction one concludes $\tilde{H}^j\rho \leq \tilde{H}^j v \leq v$, so $\tilde{H}^j\rho$ and $\tilde{H}^j v$ are finite valued for all $j \geq 0$. Applying H^j through the inequality above gives $\tilde{H}^j\rho \leq \tilde{H}^j v - \tilde{H}^{j+1}v$, and summing over j we obtain $\sum_{j=0}^{n} \tilde{H}^j\rho \leq v - \tilde{H}^{n+1}v \leq v$. Sending $n \to \infty$ gives the desired conclusion. □

Corollary 3. *Put $B(x, dy) = \beta(x, y)P(x, dy)$ for $x \in K^C$ and $B(x, dy) = 0$ for $x \in K$. Suppose there exist finite-valued nonnegative functions v_1, v_2: $S \to \mathbb{R}$ satisfying*

$$Bv_1 \le v_1 - \tilde{f},$$

$$Hv_2 \le v_2 - \tilde{r},$$

where $\tilde{r}(x) = 2f(x)v_1(x) - f^2(x) + \int_S H(x, dy)[g^2(x, y) + 2g(x, y)v_1(y)]$, $\tilde{f}(x) \triangleq f(x) + \int_S B(x, dy)g(x, y)$, *and H is as in Theorem 5. Then, $m(x) \le v_2(x)$, $x \in K^C$.*

Proof. Let r and H be as in Theorem 5. It can be easily shown that $u^* = \sum_{j=0}^{\infty} B^j \tilde{f}$. It follows from Lemma 1 that $v_1 \ge u^*$. Hence, $\tilde{r} \ge r$, and it follows from Lemma 1 and Theorem 5 that

$$v_2 \ge \sum_{j=0}^{\infty} H^j \tilde{r} \ge \sum_{j=0}^{\infty} H^j r = m. \quad \square$$

Remark 9. If one has a good approximation u to u^*, then one would expect the resulting estimator W to have small variance, whence its second moment would be close to $u^{*2}(x)$. Thus, a natural first guess for the Lyapunov functions v_1, v_2 is to set $v_1 = c_1 u$ and $v_2 = c_2 u^2$ for constants $c_1, c_2 > 0$.

Example 21 ($GI/G/1$ Queue in Heavy Traffic). Consider again the waiting time chain for a $GI/G/1$ queue in heavy traffic discussed in Example 17, for which we want to estimate

$$\alpha(x) = \mathrm{E}_x \sum_{j=0}^{T_0} e^{\theta X_j},$$

where $T_0 = \inf\{n \ge 0\colon X_n = 0\}$ and $0 < \theta < \bar{\theta} \triangleq \sup\{\gamma\colon (1 - \varphi^2(\gamma))/[\gamma\varphi(\gamma)F(0)(1 - \gamma\sigma^2/(2(1-\rho)))] \ge 1 - \rho\}$. In that example we suggested using a change-of-measure Q and filtered estimator W based on the approximation u to u^* given by $u(x) = c_0 e^{\theta x}$ for $x > 0$, and $u(0) = 1$, where $c_0 \triangleq [\theta(1-\rho)(1 - \sigma^2\theta/(2(1-\rho)))]^{-1}$. Here we obtain a bound on the variance of the estimator W under Q.

We start with candidates for the functions v_1 and v_2 in Corollary 3. Let $v_1(0) = v_2(0) = 1$, and for $x > 0$ put $v_1(x) = c_1 u(x)$ and $v_2(x) = c_2 u^2(x)$, for some constants $c_1, c_2 > 0$ to be determined.

For $x > 0$

$$Bv_1(x) = c_1 \mathrm{E}(c_0 e^{\theta(x+Z_1)};\ Z_1 > -x) + F(-x) \le c_1 u(x)\varphi(\theta) + F(-x).$$

It follows that, for $x > 0$, $Bv_1(x) \le v_1(x) - e^{\theta x}$ as long as $c_1 \ge (1 + F(-x)\varphi(\theta))/c_0(1 - \varphi(\theta))$. This is satisfied for all $x > 0$ by setting $c_1 = (1 + F(0))/c_0(1 - \varphi(\theta))$.

Note that, for $x > 0$, $\tilde{r}(x) = e^{2\theta x}(2c_1c_0 - 1)$, and

$$w(x) - f(x) = \int_S P(x, dy)u(y) = \mathrm{E}(e^{\theta(x+Z_1)};\ Z_1 > -x) + F(-x) \le u(x)\varphi(\theta) + F(-x).$$

Hence, for $x > 0$,

$$\begin{aligned} Hv_2(x) &= (w(x) - f(x))\left[\mathrm{E}\left(\frac{c_2 u^2(x+Z_1)}{u(x+Z_1)};\ Z_1 > -x\right) + F(-x)\right] \\ &\le (u(x)\varphi(\theta) + F(-x))(c_2 u(x)\varphi(\theta) + F(-x)). \end{aligned}$$

It follows that $Hv_2(x) \le v_2(x) - \tilde{r}(x)$ as long as

$$c_2 u(x)[u(x)(1 - \varphi^2(\theta)) - \varphi(\theta)F(-x)] \ge e^{2\theta x}(2c_1c_0 - 1) + u(x)\varphi(\theta)F(-x) + F^2(-x).$$

The above holds for all $x > 0$ if one sets

$$c_2 = \frac{2c_1 + \varphi(\theta)F(0) - (1 - F^2(0))/c_0}{c_0(1 - \varphi^2(\theta)) - \varphi(\theta)F(0)}.$$

Thus, with c_1 and c_2 as above, the functions v_1 and v_2 satisfy the conditions in Corollary 3, whence $\mathrm{E}_{Q_x} W^2 \le c_2 c_0 e^{2\theta x}$.

6.3. Lyapunov bounds on the variance in the SDE setting. In this section we present Lyapunov bounds for the variance in the SDE setting, analogous to those for Markov chains in the previous section.

We use the same setting of §4. Suppose $\tilde{u}$ is an approximation to u^*, the function in Assumption 3. We consider a change of measure and corresponding estimator W as in (17), with

$$\eta(x) = \frac{\sigma(x)^T \nabla_x \tilde{u}(x)}{\tilde{u}(x)},$$

so that, under Q, X has drift

$$\tilde{\mu}(x) = \mu(x) + \frac{\sigma(x)\sigma(x)^T \nabla_x \tilde{u}(x)}{\tilde{u}(x)}.$$

Let $m(x)$ denote the second moment of the estimator W,

$$m(x) = \mathrm{E}_{Q_x} W^2.$$

The following result presents a Lyapunov condition that allows one to bound m.

Theorem 6. *Suppose Assumption 3 holds. Additionally, assume $\tilde{u}$ also satisfies conditions* (iii)–(v) *in Assumption 3. Suppose there exist nonnegative functions v_1, v_2: $\mathbb{R}^m \to \mathbb{R}_+$ that are twice-continuously differentiable and solve*

$$\mu(x)^T \nabla_x v_1(x) + \tfrac{1}{2}\operatorname{tr}(\sigma(x)\sigma(x)^T \nabla_{xx} v_1(x)) + h(x)v_1(x) + f(x) \le 0,$$
$$2f(x)v_1(x) + v_2(x)\big[2h(x) + \|\eta(x)\|^2\big] + \nabla_x v_2(x) \cdot [\mu(x) - \sigma(x)\eta(x)] + \tfrac{1}{2}\operatorname{tr}(\sigma(x)\sigma(x)^T \nabla_{xx} v_1(x)) \le 0,$$

$x \in K^C$, with boundary conditions $v_1(x) \ge g(x)$ and $v_2(x) \ge g^2(x)$, $x \in K$. Then $v_2(x) \ge m(x)$, $x \in K^C$.

Proof. Consider the processes V_1 and V_2 given by

$$V_1(t) \triangleq S(t) + \beta(t)L(t)v_1(X(t)),$$
$$V_2(t) \triangleq S^2(t) + 2\beta(t)L(t)S(t)v_1(X(t)) + \beta^2(t)L^2(t)v_2(X(t)),$$

where $S(t) \triangleq \int_0^t f(X(s))\beta(s)L(s)\,ds$. Using Itô's formula and the PDE satisfied by v_1 one can verify that the Itô representation of V_1 has nonpositive term in dt, whence V_1 is a local supermartingale. Because V_1 is nonnegative and $\mathrm{E}_{Q_x} V_1(0) = v_1(x) < \infty$, it follows that $(V_1(t \wedge T)\colon t \ge 0)$ is in fact a supermartingale, and is L_1 bounded. By the martingale convergence theorem $V_1(t \wedge T)$ converges Q–a.s. to a rv $V(\infty)$ satisfying $\mathrm{E}_{Q_x} V_1(\infty) \le v_1(x)$. Note that on $\{T < \infty\}$

$$\begin{aligned} V_1(\infty) &= \int_0^T f(X(s))\beta(s)L(s)\,ds + \beta(T)L(T)v_1(X(T)) \\ &\ge \int_0^T f(X(s))\beta(s)L(s)\,ds + \beta(T)L(T)g(X(T)) \\ &= W, \end{aligned}$$

while on $\{T = \infty\}$

$$V_1(\infty) = \lim_{t\to\infty} V_1(t) \ge \lim_{t\to\infty} S(t) = W,$$

Q_x–a.s. Hence,

$$v_1(x) \ge \mathrm{E}_{Q_x} V_1(\infty) \ge \mathrm{E}_{Q_x} W = u^*(x).$$

Similarly, using Itô's formula and the PDEs satisfied by v_1 and v_2 one can verify that V_2 is a local supermartingale, and because V_2 is also nonnegative it follows that $(V_2(t \wedge T)\colon t \ge 0)$ is in fact an L_1-bounded supermartingale under Q_x. By the martingale convergence theorem $V_2(t \wedge T)$ converges a.s. to a rv $V_2(\infty)$ satisfying $\mathrm{E}_{Q_x} V_2(\infty) \le v_2(x)$. Note that on $\{T < \infty\}$

$$\begin{aligned} V_2(\infty) &= S^2(T) + 2\beta(T)L(T)S(T)v_1(X(T)) + \beta(T)L(T)v_2(X(T)) \\ &\ge S^2(T) + 2\beta(T)L(T)S(T)g(X(T)) + \beta^2(T)L^2(T)g^2(X(T)) \\ &= W^2, \end{aligned}$$

while on $\{T=\infty\}$

$$V_2(\infty)=\lim_{t\to\infty}V_2(t)\geq\lim_{t\to\infty}S^2(t)=W^2,$$

Q_x–a.s. Hence,

$$v_2(x)\geq \mathrm{E}_{Q_x}V_2(\infty)\geq \mathrm{E}_{Q_x}W^2=m(x).\quad\square$$

Remark 10. As in the discrete case, a good starting guess for v_1 and v_2 is $v_1=c_1\tilde{u}$ and $v_2=c_2\tilde{u}^2$, for some constants c_1, c_2, where $\tilde{u}$ is the approximation to u^*.

References

[1] Ahamed TPI, Borkar VS, Juneja S (2006) Adaptive importance sampling technique for Markov chains using stochastic approximation. *Oper. Res.* 54(3):489–504.
[2] Baggerly K, Cox D, Picard R (2000) Exponential convergence of adaptive importance sampling for Markov chains. *J. Appl. Probab.* 37(2):342–358.
[3] Blanchet J, Glynn P (2008) Efficient rare event simulation for the maximum of heavy-tailed random walks. *Ann. Appl. Probab.* 18(4):1351–1378.
[4] Bolia N, Juneja S, Glasserman P (2004) Function-approximation-based importance sampling for pricing American options. Ingalls RG, Rosetti MD, Smith JS, Peters BA, eds. *Proc. 2004 Winter Simulation Conf.* (IEEE Computer Society, Washington, DC), 604–611.
[5] Booth TE (1989) Zero-variance solutions for linear Monte-Carlo. *Nuclear Sci. Engrg.* 102(4):332–340.
[6] Borkar VS, Juneja S, Kherani AA (2004) Performance analysis conditioned on rare events: An adaptive simulation scheme. *Comm. Inform. Systems* 3(4):259–278.
[7] Borodin AN, Salminen P (2002) *Handbook of Brownian Motion : Facts and Formulae*, 2nd ed. (Birkhäuser, Basel, Boston).
[8] Chang C-S, Heidelberger P, Juneja S, Shahabuddin P (1994) Effective bandwidth and fast simulation of ATM intree networks. *Performance Eval.* 20(1):45–66.
[9] Chung KL (1967) *Markov Chains with Stationary Transition Probabilities*, 2nd ed. (Springer, Berlin).
[10] Dupuis P, Wang H (2004) Importance sampling, large deviations, and differential games. *Stochastic and Stochastics Rep.* 76(6):481–508.
[11] Dupuis P, Wang H (2005) Dynamic importance sampling for uniformly recurrent Markov chains. *Ann. Appl. Probab.* 15(1A):1–38.
[12] Fox BL, Glynn PW (1990) Discrete-time conversion for simulating finite-horizon Markov processes. *SIAM J. Appl. Math.* 50(5):1457–1473.
[13] Glasserman P (1993) Filtered Monte Carlo. *Math. Oper. Res.* 18(3):610–634.
[14] Glasserman P (2004) *Monte Carlo Methods in Financial Engineering* (Springer, New York).
[15] Glasserman P, Heidelberger P, Shahabuddin P (1999) Asymptotically optimal importance sampling and stratification for pricing path-dependent options. *Math. Finance* 9(2):117–152.
[16] Glynn PW, Iglehart DL (1989) Importance sampling for stochastic simulations. *Management Sci.* 35(11):1367–1392.
[17] Glynn PW, Heidelberger P, Nicola VF, Shahabuddin P (1993) Efficient estimation of the mean time between failures in non-regenerative dependability models. Evans GW, Mollaghasemi M, Russell EC, Biles WE, eds. *Proc. 1993 Winter Simulation Conf.* (IEEE Computer Society, Washington, DC), 361–366.
[18] Goyal A, Shahabuddin P, Heidelberger P, Nicola VF, Glynn PW (1992) A unified framework for simulating Markovian models of highly dependable systems. *IEEE Trans. Comput.* 41(1):36–51.
[19] Halton JH (1962) Sequential Monte Carlo. *Proc. Cambridge Philos. Soc.* 58:57–78.
[20] Hammersley JM, Handscomb DC (1964) *Monte Carlo Methods*, Methuen's Monographs on Applied Probability and Statistics, (Methuen, London).
[21] Harrison JM (1985) *Brownian Motion and Stochastic Flow Systems* (Wiley, New York).
[22] Heidelberger P (1995) Fast simulation of rare events in queueing and reliability models. *ACM Trans. Modeling Comput. Simulation* 5(1):43–85.
[23] Heidelberger P, Shahabuddin P, Nicola VF (1994) Bounded relative error in estimating transient measures of highly dependable non-Markovian systems. *ACM Trans. Modeling Comput. Simulation* 4(2):137–164.
[24] Henderson SG, Glynn PW (2002) Approximating martingales for variance reduction in Markov process simulation. *Math. Oper. Res.* 27(2):253–271.
[25] Karatzas I, Shreve S (1991) *Brownian Motion and Stochastic Calculus*, 2nd ed. (Springer-Verlag, New York).
[26] Kollman C, Baggerly K, Cox D, Picard R (1999) Adaptive importance sampling on discrete Markov chains. *Ann. Appl. Probab.* 9(2):391–412.
[27] L'Ecuyer P, Tuffin B (2008) Approximate zero-variance simulation. Mason SJ, Hill RR, Mönch L, Rose O, Jefferson T, Fowler JW, eds. *Proc. 2008 Winter Simulation Conf.* (IEEE Computer Society, Washington, DC), 170–181.
[28] Meyn S, Tweedie RL (2009) *Markov Chains and Stochastic Stability*, 2nd ed. (Cambridge University Press, New York).
[29] Nakayama MK (1996) General conditions for bounded relative error in simulations of highly reliable Markovian systems. *Adv. Appl. Probab.* 28(3):687–727.
[30] Øksendal B (2000) *Stochastic Differential Equations: An Introduction with Applications*, 5th ed. (Springer, Berlin).
[31] Sadowsky JS (1991) Large deviations theory and efficient simulation of excessive backlogs in a $GI/GI/m$ queue. *IEEE Trans. Automatic Control* 36(12):1383–1394.

[32] Salminen P, Norros I (2001) On busy periods of the unbounded Brownian storage. *Queueing Systems* 39(4):317–333.
[33] Shahabuddin P (1994) Importance sampling for the simulation of highly reliable Markovian systems. *Management Sci.* 40(3):333–352.
[34] Smith PJ, Shafi M, Gao H (1997) Quick simulation: A review of importance sampling techniques in communications systems. *IEEE J. Selected Areas Comm.* 15(4):597–613.
[35] Su Y, Fu MC (2000) Importance sampling in derivative securities pricing. Joines JA, Barton RR, Kang K, Fishwick PA, eds. *Proc. 2000 Winter Simulation Conf.* (IEEE Computer Society, Washington, DC), 587–596.
[36] Whitt W (2002) *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues*, Springer Series in Operations Research, (Springer-Verlag, New York).