

Randomized Dimension Reduction for Monte Carlo Simulations

Nabil Kahalé *

October 29, 2018

Abstract

We present a new unbiased algorithm that estimates the expected value of $f(U)$ via Monte Carlo simulation, where U is a vector of d independent random variables, and f is a function of d variables. We assume that f does not depend equally on all its arguments. Under certain conditions we prove that, for the same computational cost, the variance of our estimator is lower than the variance of the standard Monte Carlo estimator by a factor of order d . Our method can be used to obtain a low-variance unbiased estimator for the expectation of a function of the state of a Markov chain at a given time-step. We study applications to volatility forecasting and time-varying queues. Numerical experiments show that our algorithm dramatically improves upon the standard Monte Carlo method for large values of d , and is highly resilient to discontinuities.

Keywords: dimension reduction; variance reduction; effective dimension; Markov chains; Monte Carlo methods

1 Introduction

Markov chains arise in a variety of fields such as finance, queuing theory, and social networks. While much research has been devoted to the study of steady-states of Markov chains, several practical applications rely on the transient behavior of Markov chains. For example, the volatility of an index can be modelled as a Markov chain using the GARCH model (Hull 2014, Ch. 23). Financial institutions conducting stress tests may need to estimate the probability that the volatility exceeds a given level in a few years from now. Also, due to the nature of human activity, queuing systems in areas such as health-care, manufacturing, telecommunication and transportation networks, have often time-varying features and do not have a steady-state. For instance, empirical data show significant daily variation in traffic in wide-area networks (Paxson 1994, Thompson, Miller and Wilder 1997) and vehicular flow on roads (Nagel, Wagner and Woesler 2003). Estimating the expected delay of packets in a wide-area network at a specific time of the day (12pm, say) could be used to dimension such networks. Similarly, estimating the velocity of cars in a region at 6pm could be used to design transportation networks. In the same vein, consider the problem of estimating the queue-length at the end of a business day in a call center that operates with fixed hours. In such call centers, knowing how many calls would still need to be answered at 5pm could be an important metric that would be needed in estimating their staffing requirements. Methods to determine appropriate staffing levels in call centers and other many-server queueing systems with time-varying arrival rates have been designed in (Feldman, Mandelbaum, Massey and Whitt 2008). Also, approximation tools have been developed to study time-varying queues (see (Whitt 2017) and references therein). However, in many situations, there are no analytical tools, except Monte Carlo simulation, to study accurately systems modeled by a Markov chain. A drawback of Monte Carlo simulation is its

*ESCP Europe, Labex ReFi, Big data research center, 75011 Paris, France; e-mail: nkahale@escpeurope.eu.

high computation cost. This motivates the need to design efficient simulation tools to study the transient behavior of Markov chains, with or without time-varying features.

This paper gives a new unbiased algorithm to estimate $E(f(U))$, where $U = (U_1, \dots, U_d)$ is a vector of d independent random variables U_1, \dots, U_d taking values in a measurable space F , and f is a real-valued Borel-measurable function on F^d such that $f(U)$ is square-integrable. For instance, F can be equal to \mathbb{R} or to any vector space over \mathbb{R} . Under certain conditions, we show that our algorithm yields substantial lower variance than the standard Monte Carlo method for the same computational effort. Our techniques can be used to efficiently estimate the expected value of a function of the state of a Markov chain at a given time-step d , for a class of Markov chains driven by independent random variables. An alternative algorithm for Markov chains estimation, based on Quasi-Monte Carlo sequences, that substantially improves upon standard Monte Carlo in certain numerical examples, is given in (L'Ecuyer, Lécot and Tuffin 2008), with bounds on the variance proven for special situations where the state space of the chain is a subset of the real numbers.

In a standard Monte Carlo scheme, $E(f(U))$ is estimated by simulating n independent vectors in F^d having the same distribution as U , and taking the average of f over the n vectors. In the related Quasi-Monte Carlo method (see (Glasserman 2004, Ch. 5)), f is evaluated at a predetermined deterministic sequence of points. In several applications, the efficiency of Quasi-Monte Carlo algorithms can be improved by reordering the U_i 's and/or making a change of variables, so that the value of $f(U)$ depends mainly on the first few U_i 's. For instance, the Brownian bridge construction and principal components analysis have been used (Caffisch, Morokoff and Owen 1997, Acworth, Broadie and Glasserman 1998, Åkesson and Lehoczky 2000) to reduce the error in the valuation of financial derivatives via Quasi-Monte Carlo methods (see (Caffisch 1998) for related results). The relative importance of the first variables can formally be measured by calculating the effective dimension in the truncation sense, a concept defined in (Caffisch, Morokoff and Owen 1997): when the first variables are important, the effective dimension in the truncation sense is low in comparison to the nominal dimension. It is proven in (Sloan and Woniakowski 1998) that Quasi-Monte Carlo methods are effective for a class of functions where the importance of U_i decreases with i . L'Ecuyer and Lemieux (2000) apply Quasi-Monte Carlo methods to queueing simulation and option pricing, and examine their connection to the effective dimension. The truncation dimension and a related notion, the effective dimension in the superposition sense, are studied in (Sobol 2001, Owen 2003, Liu and Owen 2006). It is shown in (Wang and Fang 2003, Wang and Sloan 2005, Wang 2006) that the Brownian bridge and/or principal components analysis algorithms substantially reduce the truncation dimension of certain financial instruments. Alternative linear transformations have been proposed in (Imai and Tan 2006, Wang and Sloan 2011, Wang and Tan 2013) to reduce the effective dimension of financial derivatives and improve the performance of Quasi-Monte Carlo methods.

Other previously known variance reduction techniques have exploited the importance of certain variables or states. For instance, stratified sampling along important directions is used in pricing path-dependent options (Glasserman, Heidelberger and Shahabuddin 1999, Glasserman 2004, Section 4.3.2). Importance sampling methods aim to increase the number of samples that hit an important set via a change of measure technique (Asmussen and Glynn 2007, Section V.5). When $d = 2$ and $f(U_1, U_2)$ is more influenced by U_1 than by U_2 , and the expected time to generate U_1 is much lower than the expected time to generate U_2 , the splitting technique (Asmussen and Glynn 2007, Section V.5) simulates several independent copies of U_1 for each copy of U_2 . Asmussen and Glynn (2007, Section V.5) give the variance of the splitting estimator and the optimal number of copies of U_1 , and show that the splitting technique is related to the conditional Monte Carlo method. Multilevel splitting techniques are often used for variance reduction in the estimation of rare event probabilities (Asmussen and Glynn 2007, VI.9). The idea is to split each path that reaches an important region into a number of subpaths

in order to produce more paths that hit the rare event set. The rare event probability is then evaluated via a telescoping product. Ermakov and Melas (1995) analyse multilevel splitting techniques that estimate functionals of Markov chains with a discrete state space and of ergodic Markov chains in their steady state. Glasserman, Heidelberger, Shahabuddin and Zajic (1999) analyse the performance of multilevel splitting techniques for rare event estimation and give, under certain conditions, the optimal degree of splitting as the probability of the event goes to 0. Multilevel splitting methods have had many applications, such as the estimation of network reliability (Botev, L'Ecuyer, Rubino, Simard and Tuffin 2013) and of rare events in Jackson networks (Blanchet, Leder and Shi 2011). Multilevel splitting techniques for rare event simulation with finite time constraints are analysed in (Jiang and Fu 2017). A comprehensive survey on multilevel splitting techniques with applications to rare event simulations, sampling from complicated distributions, Monte Carlo counting, and randomized optimization, can be found in (Rubinstein and Kroese 2016, Ch. 9).

Another technique, the multilevel Monte Carlo (MLMC) method introduced in (Giles 2008), which relies on low dimensional approximations of the function to be estimated, dramatically reduces the computational complexity of estimating an expected value arising from a stochastic differential equation. Related randomized multilevel methods that produce unbiased estimators for equilibrium expectations of functionals defined on homogeneous Markov chains have been provided in (Glynn and Rhee 2014). These methods apply to the class of positive Harris recurrent Markov chains, and to chains that are contracting on average. It is shown in (Rhee and Glynn 2015) that similar randomized multilevel methods can be used to efficiently compute unbiased estimators for expectations of functionals of solutions to stochastic differential equations. The MLMC method has had numerous other applications (e.g., (Rosenbaum and Staum 2017)).

The basic idea behind our algorithm is that, if f does not depend equally on all its arguments, the standard Monte Carlo method can be inefficient because it simulates all d arguments of f at each iteration. Assuming that the expected time needed to simulate $f(U)$ is of order d and that the variance of $f(U)$ is upper and lower-bounded by constants, the expected time needed to achieve variance ϵ^2 by standard Monte Carlo simulation is $\Theta(d\epsilon^{-2})$. In contrast, our algorithm simulates at each iteration a random subset of arguments of f , and reuses the remaining arguments from the previous iteration. Under certain conditions, we show that our algorithm estimates $E(f(U))$ with variance ϵ^2 in $O(d + \epsilon^{-2})$ expected time. We also establish central limit theorems on the statistical error of our algorithm. When $d = 2$, our method is very similar to the splitting technique in (Asmussen and Glynn 2007, Section V.5). Our approach can thus be viewed as a multidimensional version of this technique. However, in contrast with existing multilevel splitting algorithms where splitting decisions typically depend on the current state, in our method, the arguments of f to be redrawn are independent of previously generated copies of U .

In order to optimize the performance of our estimator, we minimize the asymptotic product of the variance and expected running time, in the same spirit as stratified sampling (Glasserman 2004, Section 4.3.1), the splitting technique (Asmussen and Glynn 2007, Section V.5), MLMC (Giles 2008), and related methods (Rhee and Glynn 2015). This minimization is performed via a new geometric algorithm that solves in $O(d)$ time a d -dimensional optimisation problem. Our geometric algorithm is of independent interest and can be used to solve an optimization problem of the same type that was solved in (Rhee and Glynn 2015, Section 3) in $O(d^3)$ time. We are not aware of other previous algorithms that solve this problem. This work extends the research in (Kahalé 2016), where the variance properties of the randomized estimator presented in this paper were announced without proof.

Our method has the following features:

1. Under certain conditions, it estimates $E(f(U))$ with variance ϵ^2 in $O(d + \epsilon^{-2})$ expected time. We are not aware of any previous method that achieves, under the same condi-

tions, such a tradeoff between the expected running time and accuracy. In contrast with stratified sampling, which can be performed in practice only along a small number of dimensions (Glasserman 2004, Example 4.3.4), our method is targeted at high-dimensional problems. The efficiency of our method typically increases with d , even though it can be used in principle for any $d \geq 2$.

2. It is easy to implement, does not make any continuity assumptions on f , nor does it require a detailed knowledge of the structure of f or U . In contrast, importance sampling, multilevel splitting and multilevel Monte Carlo methods can achieve substantial variance reduction by exploiting the structure of the simulated model. The standard Quasi-Monte Carlo method does not necessitate a detailed knowledge of the model structure, but makes regularity assumptions on the function to be integrated, and its efficiency does not increase with d .

The rest of the paper is organized as follows. §2 presents our generic randomized dimension reduction algorithm and analyses its performance. §3 describes the aforementioned geometric algorithm and gives a numerical implementation of the randomized dimension reduction algorithm. §4 provides applications to Markov chains. §5 presents and analyses a deterministic version of our algorithm. §6 compares our algorithm to a class of MLMC algorithms. §7 gives numerical simulations. §8 contains concluding remarks. Most proofs are contained in the appendix. The connection between our approach and the ANOVA decomposition and truncation dimension is studied in the appendix. The appendix explores further the relation between our method, the splitting technique, and the conditional Monte Carlo method, and contains more numerical simulations.

2 The generic randomized dimension reduction algorithm

2.1 The algorithm description

We assume that all random variables in this paper are defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Our algorithm estimates $E(f(U))$ by performing n iterations, where n is an arbitrary positive integer. The algorithm samples more often the first arguments of f than the last ones. It implicitly assumes that, roughly speaking, the importance of the i -th argument of f decreases with i . In many Markov chain examples, the last random variables are more important than the first ones, but our algorithm can still be used efficiently after re-ordering the random variables, as described in detail in §4. A general approach to rank input variables according to their importance is described in (Sobol 2001), but we will not use such an approach in our examples.

Let

$$A = \{(q_0, \dots, q_{d-1}) \in \mathbb{R}^d : 1 = q_0 \geq q_1 \geq \dots \geq q_{d-1} > 0\}.$$

Throughout the paper, $q = (q_0, \dots, q_{d-1})$ denotes an element of A . Our generic algorithm takes such a vector q as parameter. Let (N_k) , $k \geq 1$, be a sequence of independent random integers in $[1, d]$ such that $\mathbb{P}(N_k > i) = q_i$ for $0 \leq i \leq d-1$ and $k \geq 1$. The algorithm simulates n copies $V^{(1)}, \dots, V^{(n)}$ of U and consists of the following steps:

1. First iteration. Simulate a vector $V^{(1)}$ that has the same distribution as U and calculate $f(V^{(1)})$.
2. Loop. In iteration $k+1$, where $1 \leq k \leq n-1$, let $V^{(k+1)}$ be the vector obtained from $V^{(k)}$ by redrawing the first N_k components of $V^{(k)}$, and keeping the remaining components unchanged. Calculate $f(V^{(k+1)})$.
3. Output the average of $f(V^{(1)}), \dots, f(V^{(n)})$.

More formally, consider a sequence $(U^{(k)})$, $k \geq 1$, of independent copies of U such that the two sequences (N_k) , $k \geq 1$, and $(U^{(k)})$, $k \geq 1$, are independent. Define the sequence $(V^{(k)})$, $k \geq 1$, in F^d as follows: $V^{(1)} = U^{(1)}$ and, for $k \geq 1$, the first N_k components of $V^{(k+1)}$ are the same as the corresponding components of $U^{(k+1)}$, and the remaining components of $V^{(k+1)}$ are the same as the corresponding components of $V^{(k)}$. The algorithm then outputs

$$f_n \triangleq \frac{f(V^{(1)}) + \cdots + f(V^{(n)})}{n}.$$

Note that f_n is an unbiased estimator of $E(f(U))$ since $V^{(k)} \stackrel{d}{=} U$ for $1 \leq k \leq n$.

2.2 Performance analysis

For ease of presentation, we ignore the time needed to generate N_k and the running time of the third step of the algorithm. For $1 \leq i \leq d$, let t_i be the expected time needed to generate $V^{(k+1)}$ and calculate $f(V^{(k+1)})$ when $N_k = i$. Equivalently, t_i is the expected time needed to perform Step 2 of the algorithm when $N_k = i$. Thus, t_i is the expected time needed to re-draw the first i components of U and recalculate $f(U)$, and t_d is the expected time needed to simulate $f(U)$. By convention, $t_0 = 0$. We assume for simplicity that t_i is a strictly increasing function of i . In many examples (see §2.4 and §4), it can be shown that $t_i = O(i)$. As $\mathbb{P}(N_k = i) = q_{i-1} - q_i$ for $1 \leq i \leq d$ and $k \geq 1$, where $q_d = 0$, the expected running time of a single iteration of our algorithm, excluding the first one, is equal to T , where

$$T \triangleq \sum_{i=1}^d (q_{i-1} - q_i) t_i = \sum_{i=0}^{d-1} q_i (t_{i+1} - t_i). \quad (2.1)$$

For $0 \leq i \leq d$, define

$$C(i) \triangleq \text{Var}(E(f(U)|U_{i+1}, \dots, U_d)).$$

Thus, $C(0) = \text{Var}(f(U))$, while $C(d) = 0$, and we can interpret $C(i)$ as the variance captured by the last $d - i$ components of U . Note that if f depends only on its first i arguments, then $f(U)$ is independent of (U_{i+1}, \dots, U_d) , and so

$$E(f(U)|U_{i+1}, \dots, U_d) = E(f(U)),$$

which implies that $C(i) = 0$. More generally, if the last $d - i$ arguments of f are not important, the conditional expectation $E(f(U)|U_{i+1}, \dots, U_d)$ is “almost” constant, and its variance $C(i)$ should be small. Thus $C(i)/C(0)$ can be used to measure the importance of the last $d - i$ components of U . As shown in the appendix, when U is uniformly distributed on the d -dimensional unit cube $[0, 1]^d$, this ratio coincides with a global sensitivity index for the subset $\{i + 1, \dots, d\}$, defined in (Sobol 2001, Definition 3) in terms of the ANOVA decomposition of f . Proposition 2.1 below shows that $(C(i))$, $0 \leq i \leq d$, is always a decreasing sequence, and gives an alternative expression for $C(i)$, which can be viewed as a variant of Theorem 2 of (Sobol 2001).

Proposition 2.1. *The sequence $(C(i))$, $0 \leq i \leq d$, is decreasing. If U'_1, \dots, U'_i are random variables such that $U'_j \stackrel{d}{=} U_j$ for $1 \leq j \leq i$, and U'_1, \dots, U'_i, U are independent, then*

$$C(i) = \text{Cov}(f(U), f(U'_1, \dots, U'_i, U_{i+1}, \dots, U_d)). \quad (2.2)$$

Theorem 2.1 below establishes a formal relationship between the variance of f_n and the $C(i)$'s. Let ν^* be the vector of \mathbb{R}^{d+1} with $\nu_0^* = C(0)$ and $\nu_i^* = 2C(i)$ for $1 \leq i \leq d$.

Theorem 2.1. For $n \geq 1$,

$$n\text{Var}(f_n) \leq \sum_{i=0}^{d-1} \frac{\nu_i^* - \nu_{i+1}^*}{q_i}. \quad (2.3)$$

Furthermore, the LHS of (2.3) converges to its RHS as n goes to infinity.

As, for $\nu = (\nu_0, \dots, \nu_d) \in \mathbb{R}^d \times \{0\}$,

$$\sum_{i=0}^{d-1} \frac{\nu_i - \nu_{i+1}}{q_i} = \nu_0 + \sum_{i=1}^{d-1} \nu_i \left(\frac{1}{q_i} - \frac{1}{q_{i-1}} \right), \quad (2.4)$$

the RHS of (2.3) is a weighted combination of the $C(i)$'s, with positive weights. Thus, the smaller the $C(i)$'s, the smaller the RHS of (2.3). Furthermore, as $C(i)$ is the variance of the conditional expectation $E(f(U)|U_{i+1}, \dots, U_d)$, which can be considered as a smoothed version of $f(U)$, we expect our algorithm to be resilient to discontinuities of f .

Denote by \mathbb{R}_+ the set of nonnegative real numbers. For $q \in A$, and $\vartheta = (\vartheta_0, \dots, \vartheta_d) \in \{0\} \times \mathbb{R}_+^d$, and $\nu = (\nu_0, \dots, \nu_d) \in \mathbb{R}_+^d \times \{0\}$, set

$$R(q; \vartheta, \nu) = \left(\sum_{i=0}^{d-1} q_i (\vartheta_{i+1} - \vartheta_i) \right) \left(\sum_{i=0}^{d-1} \frac{\nu_i - \nu_{i+1}}{q_i} \right). \quad (2.5)$$

The expected time needed to perform n iterations of the algorithm, including the first one, is $T_n = (n-1)T + t_d$. Theorem 2.1 and (2.1) imply that $T_n \text{Var}(f_n)$ converges to $R(q; t, \nu^*)$ as n goes to infinity, where $t = (t_0, \dots, t_d)$. By (2.4), $R(q; \vartheta, \nu)$ is an increasing function with respect to ν , i.e. $R(q; \vartheta, \nu) \leq R(q; \vartheta, \nu')$ for $\nu \leq \nu'$, where the symbol \leq between vectors represents componentwise inequality. Similarly, it is easy to see that $R(q; \vartheta, \nu)$ is increasing with respect to ϑ . Let $T^{\text{tot}}(q, \epsilon)$ be the total expected time it takes for our algorithm to guarantee that $\text{Std}(f_n) \leq \epsilon$. Corollary 2.1 below gives an upper bound on $T^{\text{tot}}(q, \epsilon)$ in terms of $R(q; t, \nu^*)$. It also implies that, if $R(q; t, \nu^*)$ is upper bounded by a constant independent of d , and $\text{Var}(f(U)) = \Theta(1)$, then our algorithm outperforms the standard Monte Carlo algorithm by a factor of order t_d . More precisely, running our algorithm for $n = \lceil t_d T^{-1} \rceil$ iterations has the same expected cost, up to a constant, as a single iteration of the standard Monte Carlo method, but produces an unbiased estimator of $E(f(U))$ with $O(1/t_d)$ variance.

Corollary 2.1. For $\epsilon > 0$,

$$T^{\text{tot}}(q, \epsilon) \leq t_d + R(q; t, \nu^*) \epsilon^{-2}. \quad (2.6)$$

Furthermore, if $n = \lceil t_d T^{-1} \rceil$, the expected running time of n iterations of the algorithm is at most $2t_d$, and

$$\text{Var}(f_n) \leq \frac{R(q; t, \nu^*)}{t_d}. \quad (2.7)$$

Proof. Theorem 2.1 and (2.1) imply that $n\text{Var}(f_n)T \leq R(q; t, \nu^*)$. Thus, $\text{Std}(f_n) \leq \epsilon$ for $n = \lceil R(q; t, \nu^*)T^{-1}\epsilon^{-2} \rceil$. The expected time needed to calculate f_n is T_n , which is upper-bounded by $t_d + R(q; t, \nu^*)\epsilon^{-2}$ since $n-1 \leq R(q; t, \nu^*)T^{-1}\epsilon^{-2}$. Hence (2.6). On the other hand, if $n = \lceil t_d T^{-1} \rceil$, then $T_n \leq 2t_d$ since $(n-1)T \leq t_d$, and (2.7) holds since $nT \geq t_d$. \square

Theorem 2.2 below establishes a central limit theorem on f_n . It also establishes a central limit theorem on the estimate of $E(f(U))$ that can be obtained with a computational budget c , using the framework described by Glynn and Whitt (1992). Denote by $\tilde{N}(c)$ the number of iterations generated by our algorithm in c units of computation time. In other words, $\tilde{N}(c)$ is the maximum integer n such that f_n is calculated within c time (with $f_0 \triangleq 0$). As the time to calculate f_n is random, $\tilde{N}(c)$ is a random integer. Let \Rightarrow denote weak convergence (see (Billingsley 1999)).

Theorem 2.2. As $n \rightarrow \infty$,

$$\sqrt{n}(f_n - E(f(U))) \Rightarrow N(0, \sigma^2), \quad (2.8)$$

where

$$\sigma^2 = \sum_{i=0}^{d-1} \frac{\nu_i^* - \nu_{i+1}^*}{q_i}.$$

Furthermore, as $c \rightarrow \infty$,

$$\sqrt{c}(f_{\tilde{N}(c)} - E(f(U))) \Rightarrow N(0, R(q; t, \nu^*)). \quad (2.9)$$

In light of above, we will use $R(q; t, \nu^*)$ to measure the performance of our algorithm. The smaller the $C(i)$'s and t_i 's, the smaller $R(q; t, \nu^*)$, and the better the performance of our algorithm. Proposition 2.2 below shows that $C(i)$ is small if f is well-approximated by a function of its first i arguments.

Proposition 2.2. For $1 \leq i \leq d$, if f_i is a measurable function from F^i to \mathbb{R} such that $f_i(U_1, \dots, U_i)$ is square-integrable, then

$$C(i) \leq \text{Var}(f(U) - f_i(U_1, \dots, U_i)).$$

2.3 Explicit and semi-explicit distributions

An optimal choice for q is a one that minimizes $R(q; t, \nu^*)$. A numerical algorithm that performs such minimization is presented in §3. This subsection gives explicit or semi-explicit choices for q , with corresponding upper-bounds on $R(q; t, \nu^*)$.

Proposition 2.3 below gives upper bounds on $R(q; t, \nu^*)$ if $t_i = O(i)$ and $C(i)$ decreases at a sufficiently high rate. It implies in particular that, if $t_i = O(i)$ and $C(i) = O((i+1)^\gamma)$ with $\gamma < -1$, then $T^{\text{tot}}(q, \epsilon) = O(d + \epsilon^{-2})$.

Proposition 2.3. Assume that $d \geq 2$ and there are constants c and c' and $\gamma < 0$ independent of d such that $t_i \leq ci$ and $C(i) \leq c'(i+1)^\gamma$ for $0 \leq i \leq d$. Then, for $q_i = (i+1)^{(\gamma-1)/2}$, $0 \leq i \leq d-1$, there is a constant c_1 independent of d such that

$$R(q; t, \nu^*) \leq \begin{cases} c_1, & \gamma < -1, \\ c_1 \ln^2(d), & \gamma = -1, \\ c_1 d^{\gamma+1}, & -1 < \gamma < 0. \end{cases} \quad (2.10)$$

Below is a simple example where $C(0) = 1$ and the $C(i)$'s do not meet the conditions of Proposition 2.3.

Example 2.1. Suppose that $F = \mathbb{R}$ and that U_1, \dots, U_d are square-integrable real-valued random variables with unit variance. Assume that $f(x_1, \dots, x_d) = d^{-1/2}(\sum_{j=1}^d x_j)$ for $(x_1, \dots, x_d) \in \mathbb{R}^d$. As f depends equally on its arguments, our algorithm does not improve upon the standard Monte Carlo method. Since

$$E(f(U)|U_{i+1}, \dots, U_d) = d^{-1/2}(E(U_1 + \dots + U_i) + U_{i+1} + \dots + U_d),$$

$C(i) = (d-i)/d$. The conditional variance $C(i)$ decreases very slowly i since $C(d/2)$ has the same order of magnitude as $C(0)$. Thus the $C(i)$'s do not meet the conditions of Proposition 2.3.

When upper-bounds on the $C(i)$'s and t_i 's satisfying a convexity condition are known, Proposition 2.4 below gives an explicit vector q together with an upper bound on $R(q; t, \nu^*)$.

Proposition 2.4. Assume that $t_i \leq \vartheta_i$ for $0 \leq i \leq d$, where $\vartheta_0, \dots, \vartheta_d$ is a strictly increasing sequence with $\vartheta_0 = 0$. Assume further that ν_0, \dots, ν_{d-1} are positive real numbers such that $C(i) \leq \nu_i$ for $0 \leq i \leq d-1$, and that the sequence

$$\theta_i = \frac{\nu_{i+1} - \nu_i}{\vartheta_{i+1} - \vartheta_i},$$

$0 \leq i \leq d-1$, is increasing (by convention, $\nu_d = 0$). Then, for $q_i = \sqrt{\theta_i/\theta_0}$, $0 \leq i \leq d-1$,

$$R(q; t, \nu^*) \leq 2 \left(\sum_{i=0}^{d-1} \sqrt{(\nu_i - \nu_{i+1})(\vartheta_{i+1} - \vartheta_i)} \right)^2.$$

Proof. We first observe that $\theta_i \leq \theta_{d-1} < 0$ for $0 \leq i \leq d-1$. Thus q is well-defined and belongs to A . Let $\vartheta = (\vartheta_0, \dots, \vartheta_d)$. As $t \leq \vartheta$ and $\nu^* \leq 2\nu$, and since $R(q; \cdot, \cdot)$ is increasing with respect to its second and third arguments, we have $R(q; t, \nu^*) \leq R(q; \vartheta, 2\nu)$. This concludes the proof. \square

Proposition 2.5 below yields an upper bound on $\sqrt{R(q; t, \nu^*)}$ in terms of a weighted sum of the square roots of the $C(i)$'s, for a semi-explicit vector q .

Proposition 2.5. Assume that $C(d-1) > 0$. If, for $0 \leq i \leq d-1$,

$$q_i = \sqrt{\frac{t_1 C(i)}{t_{i+1} C(0)}},$$

then

$$R(q; t, \nu^*) \leq 8 \left(\sum_{i=0}^{d-1} (\sqrt{t_{i+1}} - \sqrt{t_i}) \sqrt{C(i)} \right)^2. \quad (2.11)$$

Proposition 2.6 below gives an explicit distribution which is optimal up to a logarithmic factor, without requiring any prior knowledge on the $C(i)$'s.

Proposition 2.6. For any $q \in A$,

$$R(q; t, \nu^*) \geq \sum_{i=0}^{d-1} C(i)(t_{i+1} - t_i). \quad (2.12)$$

Furthermore, if $q_i = t_1/t_{i+1}$ for $0 \leq i \leq d-1$, then

$$R(q; t, \nu^*) \leq 2(1 + \ln(\frac{t_d}{t_1})) \sum_{i=0}^{d-1} C(i)(t_{i+1} - t_i). \quad (2.13)$$

2.4 A Lipschitz function example

Assume that $F = \mathbb{R}$ and that U_1, \dots, U_d are square-integrable real-valued random variables, with $\sigma_1 \geq \dots \geq \sigma_d > 0$, where σ_i is the standard deviation of U_i . Assume also that $f(x_1, \dots, x_d) = g(\sum_{j=1}^d x_j)$ for $(x_1, \dots, x_d) \in \mathbb{R}^d$, where g is a real-valued 1-Lipschitz function on \mathbb{R} that can be calculated in constant time. For instance, $f(x_1, \dots, x_d) = \max(\sum_{j=1}^d x_j - K, 0)$, where K is a constant, satisfies this condition. Assume further that each U_i can be simulated in constant time. For $1 \leq k \leq n$, let S_k be the sum of all components of $V^{(k)}$. Thus S_{k+1} can be calculated recursively in $O(N_k)$ time by adding to S_k the first N_k components of $V^{(k+1)}$ and subtracting the first N_k components of $V^{(k)}$. Hence $t_i \leq ci$, for some constant c .

In order to bound the $C(i)$'s, we show that $f(U)$ can be approximated by $f_i(U_1, \dots, U_i)$, where $f_i(x_1, \dots, x_i) = g(\sum_{j=1}^i x_j + \sum_{j=i+1}^d E(U_j))$ for $(x_1, \dots, x_i) \in \mathbb{R}^i$. Let $\|Z\| = \sqrt{E(Z^2)}$ for a real-valued random variable Z . By Proposition 2.2,

$$\begin{aligned} C(i) &\leq \|f(U) - f_i(U_1, \dots, U_i)\|^2 \\ &\leq \left\| \sum_{j=i+1}^d (U_j - E(U_j)) \right\|^2 \\ &= \text{Var}\left(\sum_{j=i+1}^d U_j \right) \\ &= \sum_{j=i+1}^d \sigma_j^2. \end{aligned} \tag{2.14}$$

The second equation follows from the assumption that g is 1-Lipschitz. By applying Proposition 2.4, with $\vartheta_i = ci$ and $\nu_i = \sum_{j=i+1}^d \sigma_j^2$, and setting $q_i = \sigma_{i+1}/\sigma_1$, $0 \leq i \leq d-1$, we infer that

$$R(q; t, \nu^*) \leq 2c \left(\sum_{i=1}^d \sigma_i \right)^2.$$

Thus, if $\sigma_i = O(i^\gamma)$, with $\gamma < -1$, then $R(q; t, \nu^*) = O(1)$ and $T^{\text{tot}}(q, \epsilon) = O(d + \epsilon^{-2})$. Also, by (2.14) and a standard calculation, $C(i) = O((i+1)^{2\gamma+1})$ for $0 \leq i \leq d$. As $2\gamma + 1 < -1$, Proposition 2.3 is also applicable in this case.

3 The optimal distribution

We now seek to calculate a vector q that minimizes $R(q; t, \nu^*)$. Given a vector ν in $\mathbb{R}^d \times \{0\}$ whose first d components are positive, Theorem 3.1 below gives a geometric algorithm that finds in $O(d)$ time a vector q^* that minimizes $R(q; t, \nu)$ under the constraint that $q \in A$. Note that the vector q^* depends on ν . In (Rhee and Glynn 2015, Section 3), a dynamic programming algorithm that calculates such a vector q^* in $O(d^3)$ time has been described.

Let $\nu' = (\nu'_0, \dots, \nu'_d) \in \mathbb{R}^{d+1}$ be such that the set $\{(t_i, \nu'_i) : 0 \leq i \leq d\}$ forms the lower hull of the set $\{(t_i, \nu_i) : 0 \leq i \leq d\}$. In other words, ν' is the supremum of all sequences in \mathbb{R}^{d+1} such that $\nu' \leq \nu$ and the sequence (θ_i) is increasing, where

$$\theta_i = \frac{\nu'_{i+1} - \nu'_i}{t_{i+1} - t_i}, \tag{3.1}$$

$0 \leq i \leq d-1$. For instance, if $d = 6$, with $t_i = i$ and $\nu = (20, 21, 13, 8, 7, 2, 0)$, then $\nu' = (20, 16, 12, 8, 5, 2, 0)$, as illustrated in Fig. 1. §3.1 shows how to calculate ν' in $O(d)$ time.

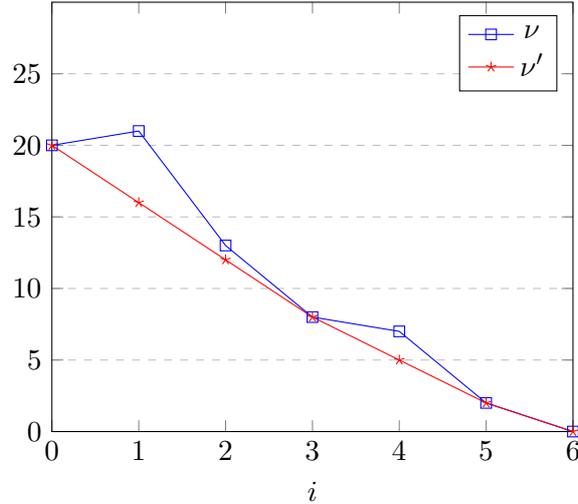
Theorem 3.1. *Let ν be a vector in $\mathbb{R}^d \times \{0\}$ whose first d components are positive. For $0 \leq i \leq d-1$, set $q_i^* = \sqrt{\theta_i/\theta_0}$, where θ_i is given by (3.1), and let $q^* = (q_0^*, \dots, q_{d-1}^*)$. Then $q^* = \arg \min_{q \in A} R(q; t, \nu)$, and*

$$R(q^*; t, \nu) = \left(\sum_{i=0}^{d-1} \sqrt{(\nu'_i - \nu'_{i+1})(t_{i+1} - t_i)} \right)^2. \tag{3.2}$$

3.1 Lower hull calculation

Given ν , the following algorithm, due to (Andrew 1979), first generates recursively a subset $B(j)$ of $\{1, \dots, d\}$, $2 \leq j \leq d$, then calculates ν' via $B(d)$. The algorithm runs in $O(d)$ time.

Figure 1: Lower hull



1. Set $B(2) = \{1, 2\}$.
2. For $j = 3$ to d , denote by $i_1 < \dots < i_m$ the elements of $B(j-1)$. Let k be the largest element of $\{2, \dots, m\}$ such that (t_{i_k}, ν_{i_k}) lies below the segment $[(t_{i_{k-1}}, \nu_{i_{k-1}}), (t_j, \nu_j)]$, if such k exists, otherwise let $k = 1$. Set $B(j) = \{i_1, \dots, i_k, j\}$.
3. For $i = 1$ to d , let i' and i'' be two elements of $B(d)$ with $i' \leq i \leq i''$. Set ν'_i so that (t_i, ν'_i) lies on the segment $[(t_{i'}, \nu_{i'}), (t_{i''}, \nu_{i''})]$.

3.2 Estimating the $C(i)$'s

The calculation of a vector $q^* \in A$ that minimizes $R(q; t, \nu^*)$ requires the knowledge of the $C(i)$'s. Proposition 3.1 below can be used to estimate $C(i)$ via Monte Carlo simulation. Assuming that $f(U)$ can be approximated by a function of its first i arguments, we expect that both components of the product in the RHS of (3.3) to be small, on average. Thus, (3.3) can be considered as a “control variate” version of (2.2), and should yield a more accurate estimate of $C(i)$ via Monte Carlo simulation for large values of i .

Proposition 3.1. *Assume that U'_1, \dots, U'_d , and U''_{i+1}, \dots, U''_d , are random variables such that $U'_j \stackrel{d}{=} U_j$ for $1 \leq j \leq d$, and $U''_j \stackrel{d}{=} U_j$ for $i+1 \leq j \leq d$, and U'_1, \dots, U'_d , U , U''_{i+1}, \dots, U''_d are independent. Then*

$$C(i) = E((f(U) - f(U_1, \dots, U_i, U'_{i+1}, \dots, U'_d)) \\ (f(U'_1, \dots, U'_i, U_{i+1}, \dots, U_d) - f(U'_1, \dots, U'_i, U''_{i+1}, \dots, U''_d))). \quad (3.3)$$

3.3 Numerical algorithm

Building upon the previously discussed elements, the algorithm that we have used for our numerical experiments is as follows. It constructs a vector (ν_0, \dots, ν_d) and uses it as a proxy for ν^* .

1. For $i = 0$ to $d-1$, if $i+1$ is a power of 2, estimate $C(i)$ by Monte Carlo simulation with 1000 samples via Proposition 3.1.

2. Set $\nu_0 = C(0)$, and $\nu_d = 0$. For $1 \leq i \leq d-1$, let $\nu_i = 2C(j)$, where j is the largest index in $[0, i]$ such that $j+1$ is a power of 2.
3. For $i = d-1$ down to 1, set $\nu_i \leftarrow \max(\nu_i, \nu_{i+1})$. Set $\nu_0 \leftarrow \max(\nu_0, \nu_1/2)$.
4. Let $(\nu'_0, \dots, \nu'_d) \in \mathbb{R}^{d+1}$ be such that the set $\{(t_i, \nu'_i) : 0 \leq i \leq d\}$ forms the lower hull of the set $\{(t_i, \nu_i) : 0 \leq i \leq d\}$. For $0 \leq i \leq d-1$, set $q_i = \sqrt{\theta_i/\theta_0}$, where θ_i is given by (3.1).
5. Calculate T via (2.1). For $0 \leq i \leq d-1$, set

$$q_i \leftarrow \min\left(1, \max\left(q_i, \frac{T}{t_{i+1} \ln(t_d/t_1)}\right)\right).$$

6. Run Steps 1 through 3 of the generic randomized dimension reduction algorithm of §2.1 using q .

The purpose of Steps 3 and 5 is to reduce the impact on q of statistical errors that arise in Step 1. Because of statistical errors, Step 1 may underestimate or overestimate the $C(i)$'s. Step 3 guarantees that the ν_i 's are non-negative, so that the q_i 's can be calculated in Step 4. Step 5 yields a cap on the RHS of (2.3) by ensuring that the q_i 's are not too small. Using (2.1) and the proof of Proposition 2.6, and assuming that $t_d \geq 2t_1$, it can be shown that Step 5 increases T by at most a constant multiplicative factor. An alternative way to implement our algorithm is to skip Steps 1 through 5 and run the generic algorithm with $q_i = t_1/t_{i+1}$ for $0 \leq i \leq d-1$. By Proposition 2.6, the resulting vector q is optimal up to a logarithmic factor.

4 Applications to Markov chains

In queueing systems, the performance metrics at a specific time instant are heavily dependent on the last busy cycle, i.e., the events that occurred after the queue was empty for the last time. Thus, the performance metrics depend a lot more on the last random variables driving the system than on the initial ones. Nevertheless, we can apply our algorithm to queueing systems by using a time-reversal transformation inspired from (Glynn and Rhee 2014). More generally, using such a time-reversal transformation, this section shows that our algorithm can efficiently estimate the expected value of a function of the state of a Markov chain at time-step d , for a class of Markov chains driven by independent random variables.

Let (X_m) , $0 \leq m \leq d$, be a Markov chain with state-space F' and deterministic initial value X_0 . Assume that there are independent random variables Y_i , $0 \leq i \leq d-1$, that take values in F , and measurable functions g_i from $F' \times F$ to F' such that $X_{i+1} = g_i(X_i, Y_i)$ for $0 \leq i \leq d-1$. We want to estimate $E(g(X_d))$ for a given positive integer d , where g is a deterministic real-valued measurable function on F' such that $g(X_d)$ is square-integrable. For $1 \leq i \leq d$, set $U_i = Y_{d-i}$. It can be shown by induction that $X_d = G_i(U_1, \dots, U_i, X_{d-i})$, where G_i , $0 \leq i \leq d$, is a measurable function from $F^i \times F'$ to F' , and so there is a real-valued measurable function f on F^d with $g(X_d) = f(U_1, \dots, U_d)$. We can thus use our randomized dimension reduction algorithm to estimate $E(g(X_d))$. Recall that, in iteration $k+1$ in Step 2 of the generic algorithm of §2.1, conditioning on $N_k = i$, the first i arguments of f are re-drawn, and the remaining arguments are unchanged. This is equivalent to re-drawing the last i random variables driving the Markov chain, and keeping the first $d-i$ variables unchanged. In light of above, the generic randomized dimension reduction algorithm for Markov chains estimation takes as parameter a vector $q \in A$ and consists of the following steps:

1. First iteration. Generate recursively X_0, \dots, X_d . Calculate $g(X_d)$.
2. Loop. In iteration $k+1$, where $1 \leq k \leq n-1$, keep X_0, \dots, X_{d-N_k} unchanged, and calculate recursively X_{d-N_k+1}, \dots, X_d by re-drawing $Y_{d-N_k}, \dots, Y_{d-1}$, where N_k is a random integer in $[1, d]$ such that $\mathbb{P}(N_k > i) = q_i$. Calculate $g(X_d)$.

3. Output the average of g over the n copies of X_d generated in the first two steps.

We assume that g and the g_i 's can be calculated in constant time, and that the expected time needed to simulate each Y_i is upper-bounded by a constant independent of d . Thus, given N_k , the expected time needed to perform iteration $k + 1$ is $O(N_k)$. Hence $t_i \leq ci$, for some constant c independent of d . Proposition 4.1 below shows that, roughly speaking, $C(i)$ is small if X_{d-i} and X_d are ‘‘almost’’ independent.

Proposition 4.1. *For $0 \leq i \leq d$, we have $C(i) = \text{Var}(E(g(X_d)|X_{d-i}))$.*

By Proposition 2.3, if there are constants $c' > 0$ and $\gamma < -1$ independent of d such that $C(i) \leq c'(i+1)^\gamma$ for $0 \leq i \leq d-1$, then $R(q; t, \nu^*)$ is upper-bounded by a constant independent of d , where $q_i = (i+1)^{(\gamma-1)/2}$ for $0 \leq i \leq d-1$. The analysis in (Asmussen and Glynn 2007, Section IV.1a), combined with Proposition 4.1, suggests that $C(i)$ decreases exponentially with i for a variety of Markov chains.

For $x \in \mathbb{F}'$, and $0 \leq i \leq d$, let

$$X_{i,x} = G_i(U_1, \dots, U_i, x).$$

In other words, $X_{i,x}$ is the state of the chain at time-step d if the chain is at state x at time-step $d-i$. Intuitively, we expect $X_{i,x}$ to be close to X_d for large i if X_d depends mainly on the last Y_j 's. By Proposition 2.2, if $g(X_{i,x})$ is square-integrable,

$$C(i) \leq \|g(X_d) - g(X_{i,x})\|^2. \quad (4.1)$$

In the following examples, we prove that under certain conditions, $R(q; t, \nu^*)$ is upper-bounded by a constant independent of d for an explicit vector $q \in A$, and so $T^{\text{tot}}(q, \epsilon) = O(d + \epsilon^{-2})$.

4.1 GARCH volatility model

In the GARCH(1,1) volatility model (see (Hull 2014, Ch. 23)), the variance X_i of an index return between day i and day $i+1$, as estimated at the end of day i , satisfies the following recursion:

$$X_{i+1} = w + \alpha X_i Y_i^2 + \beta X_i,$$

$i \geq 0$, where w , α and β are positive constants with $\alpha + \beta < 1$, and Y_i , $i \geq 0$, are independent standard Gaussian random variables. The variable Y_i is known at the end of day $i+1$. At the end of day 0, given $X_0 \geq 0$, a positive integer d and a real number z , we want to estimate $\mathbb{P}(X_d > z)$. In this example, $F = F' = \mathbb{R}$, and $g_i(x, y) = w + \alpha xy^2 + \beta x$, with $g(u) = \mathbf{1}\{u > z\}$ for $u \in \mathbb{R}$. Proposition 4.2 below shows that $C(i)$ decreases exponentially with i .

Proposition 4.2. *There is a constant κ independent of d such that $C(i) \leq \kappa(\alpha + \beta)^{i/2}$ for $0 \leq i \leq d-1$.*

By applying Proposition 2.4 with $\vartheta_i = ci$ and $\nu_i = \kappa(\alpha + \beta)^{i/2}$, and setting $q_i = (\alpha + \beta)^{i/4}$ for $0 \leq i \leq d-1$, we infer that $R(q; t, \nu^*)$ is upper-bounded by a constant independent of d .

4.2 $G_t/D/1$ queue

Consider a queue where customers arrive at time-step i , $1 \leq i \leq d$, and are served by a single server in order of arrival. Service times are all equal to 1. Assume the system starts empty at time-step 0, and that A_i customers arrive at time-step i , $0 \leq i \leq d$, where $A_0 = 0$ and the A_i 's are independent square-integrable random variables. Let X_i be the number of customers waiting in the queue at time-step i . Then $X_0 = 0$ and (X_i) satisfies the Lindley equation

$$X_{i+1} = (X_i + Y_i)^+,$$

for $0 \leq i \leq d-1$, with $Y_i = A_{i+1} - 1$. We want to estimate $E(X_d)$. In this example, g is the identity function, $F = F' = \mathbb{R}$, and $g_i(x, y) = (x + y)^+$. Proposition 4.3 below shows that $C(i)$ decreases exponentially with i under certain conditions on the service times.

Proposition 4.3. *If there are constants $\gamma > 0$ and $\kappa < 1$ independent of d such that*

$$E(e^{\gamma Y_i}) \leq \kappa \quad (4.2)$$

for $0 \leq i \leq d-1$, then $C(i) \leq \gamma' \kappa^i$ for $0 \leq i \leq d-1$, where γ' is a constant independent of d .

By applying Proposition 2.4 with $\vartheta_i = ci$ and $\nu_i = \gamma' \kappa^i$, and setting $q_i = \kappa^{i/2}$ for $0 \leq i \leq d-1$, we conclude that, under the assumption of Proposition 4.3, $R(q; t, \nu^*)$ is upper-bounded by a constant independent of d . The assumption in Proposition 4.3 can be justified as follows. Given $i \in [0, d-1]$, if $E(A_i) < 1$ and the function $h(\gamma) = E(e^{\gamma Y_i})$ is bounded on a neighborhood of 0, then $h'(0) = E(Y_i) < 0$. As $h(0) = 1$, there is $\gamma > 0$ such that $h(\gamma) < 1$, and (4.2) holds for $\kappa = h(\gamma)$. The assumption in Proposition 4.3 says that γ and κ can be chosen independently of i and of d .

4.3 $M_t/GI/1$ queue

Consider a $M_t/GI/1$ queue where customers are served by a single server in order of arrival. We assume that customers arrive according to a Poisson process with positive and continuous time-varying rate $\lambda_s \leq \lambda^*$, where λ^* is a fixed positive real number. The service times are assumed to be i.i.d. and independent of the arrival times. Assume that the system starts empty at time 0. For simplicity, we assume that the number of customers that arrive in any bounded time interval is finite (rather than finite with probability 1). Consider a customer present in the system at a given time s . If the customer has been served for a period of length τ , its remaining service time is equal to its service time minus τ , and if the customer is in the queue, its remaining service time is equal to its service time. The residual work W_s at time s is defined as the sum of remaining service times of customers present in the system at s . We want to estimate the expectation of W_θ , where θ is a fixed time. Let $d = \lceil \lambda^* \theta \rceil$, and assume that $d \geq 2$. For $0 \leq i \leq d$, let $X_i = W_{i\theta/d}$ be the residual work at time $i\theta/d$. For $0 \leq i \leq d-1$, let Y_i be the vector that consists of arrival and service times of customers that arrive during the interval $(i\theta/d, (i+1)\theta/d]$. In this example, g is the identity function, F is equal to the set of real-valued sequences with finite support, and $F' = \mathbb{R}$. Let $0 \leq s < s'$. If no costumers arrive in $(s, s']$ then $W_{s'} = (W_s - s' + s)^+$. On the other hand, if no costumers arrive in (s, s') and a customer with service time S arrives at s' , then $W_{s'} = S + (W_s - s' + s)^+$. Thus, given the set of arrival and service times of customers that arrive in $(s, s']$, we can calculate iteratively $W_{s'}$ from W_s . This implies that X_{i+1} is a deterministic measurable function of X_i and Y_i , for $0 \leq i \leq d-1$. Proposition 4.4 below shows that $C(i)$ decreases exponentially with i under certain conditions on the arrival and service times.

Proposition 4.4. *For $0 \leq s \leq \theta$, let $Z_\theta(s)$ be the cumulative service time of costumers that arrive in $[s, \theta]$. Assume there are constants $\gamma > 0$ and $\kappa < 1$ independent of d such that, for $0 \leq s \leq s' \leq \theta$ and $s' - s \leq 1/\lambda^*$,*

$$E(e^{\gamma(Z_\theta(s') - Z_\theta(s) - 1/\lambda^*)}) \leq \kappa. \quad (4.3)$$

Then $C(i) \leq \gamma' \kappa^{i/2}$ for $0 \leq i \leq d-1$, where γ' is a constant independent of d .

By applying Proposition 2.4 with $\vartheta_i = ci$ and $\nu_i = \gamma' \kappa^{i/2}$, and setting $q_i = \kappa^{i/4}$ for $0 \leq i \leq d-1$, we conclude that, under the assumption of Proposition 4.4, $R(q; t, \nu^*)$ is upper-bounded by a constant independent of d . The assumption in Proposition 4.4 can be justified as follows. For $0 \leq s \leq s' \leq \theta$ and $s' - s \leq 1/\lambda^*$, the cumulative service times of customers that arrive

in $[s, s']$ is $Z_\theta(s') - Z_\theta(s)$. If $E(Z_\theta(s') - Z_\theta(s)) < s' - s$ and $h(\gamma) = E(e^{\gamma(Z_\theta(s') - Z_\theta(s) - 1/\lambda^*)})$ is bounded on a neighborhood of 0, then $h'(0) < 0$. Thus $h(\gamma) < 1$ for some $\gamma > 0$ and (4.3) holds for $\kappa = h(\gamma)$. The assumption in Proposition 4.4 says that γ and κ can be chosen independently of d, s and s' .

5 Deterministic dimension reduction

This section studies a deterministic dimension reduction algorithm that performs the same steps as the generic randomized dimension reduction algorithm of §2.1, but uses a deterministic integral sequence (N_k) , $k \geq 1$, taking values in $[1, d]$, to estimate $E(f(U))$. As for the randomized algorithm, denote by f_n the output of the deterministic dimension reduction algorithm, and by T_n its expected running time, where n is the number of iterations, including the first one. The sequence (N_k) , $k \geq 1$, may depend on n . As the algorithm generates n copies of U , the random variable f_n is an unbiased estimator of $E(f(U))$.

Assume that $C(d-1) > 0$ and let $\hat{q} = \arg \min_{q \in A} R(q; t, C)$, where C denotes the vector $(C(0), \dots, C(d))$. The existence of \hat{q} follows from Theorem 3.1. Define the integers μ_0, \dots, μ_{d-1} recursively as follows. Let $\mu_0 = 1$ and, for $1 \leq i \leq d-1$, let μ_i be the largest multiple of μ_{i-1} in the interval $[0, 1/\hat{q}_i]$, i.e.

$$\mu_i = \mu_{i-1} \lfloor \frac{1}{\mu_{i-1} \hat{q}_i} \rfloor.$$

It can be shown by induction that μ_i is well-defined and positive. Let \bar{q} be the vector in A defined by $\bar{q}_i = 1/\mu_i$, for $0 \leq i \leq d-1$. As $\lfloor x \rfloor \leq x < 2\lfloor x \rfloor$ for $x \geq 1$, we have $\mu_i \leq 1/\hat{q}_i < 2\mu_i$. It follows that, for $0 \leq i \leq d-1$,

$$\hat{q}_i \leq \bar{q}_i < 2\hat{q}_i. \quad (5.1)$$

Define the sequence (\bar{N}_k) , $k \geq 1$, as follows:

$$\bar{N}_k = \max\{i \in [1, d] : k \text{ is a multiple of } \mu_{i-1}\}.$$

As $\mu_0 = 1$, such i always exists. For $0 \leq i \leq d-1$ and $k \geq 1$, if $\bar{N}_k = j$ with $j > i$, then k is a multiple of μ_{j-1} , and so k is a multiple of μ_i , since μ_{j-1}/μ_i is an integer. Conversely, if k is a multiple of μ_i then, by construction, $\bar{N}_k > i$. Hence

$$\bar{N}_k > i \Leftrightarrow k \equiv 0 \pmod{\mu_i}. \quad (5.2)$$

Given $i \in [0, d-1]$, the inequality $\bar{N}_k > i$ occurs once as k ranges in a set of μ_i consecutive positive integers. The sequence (\bar{N}_k) can thus be considered as a deterministic counterpart to the random sequence (N_k) generated by the randomized dimension reduction algorithm when $q = \bar{q}$.

Theorem 5.1 below gives a lower bound on the performance of the deterministic dimension reduction algorithm for any sequence (N_k) , $k \geq 1$, and analyses the algorithm when $N_k = \bar{N}_k$ for $k \geq 1$.

Theorem 5.1. *For $n \geq 1$ and any deterministic sequence (N_k) , $k \geq 1$,*

$$T_n \text{Var}(f_n) \geq R(\hat{q}; t, C). \quad (5.3)$$

If $N_k = \bar{N}_k$ for $k \geq 1$ then, for $n \geq 1$,

$$T_n = t_d + \sum_{i=0}^{d-1} \lfloor (n-1)\bar{q}_i \rfloor (t_{i+1} - t_i), \quad (5.4)$$

and

$$n \text{Var}(f_n) \leq \sum_{i=0}^{d-1} \frac{C(i) - C(i+1)}{\bar{q}_i}. \quad (5.5)$$

Furthermore, the LHS of (5.5) converges to its RHS as n goes to infinity.

Using again the framework of (Glynn and Whitt 1992), we measure the performance of an estimator via the work-normalized variance, i.e. the product of the variance and expected running time. If $N_k = \bar{N}_k$ for $k \geq 1$ then, by Theorem 5.1,

$$T_n \text{Var}(f_n) \rightarrow R(\bar{q}; t, C)$$

as n goes to infinity. Furthermore, it follows from (2.5) and (5.1) that $R(\bar{q}; t, C) \leq 2R(\hat{q}; t, C)$. Thus, up to a factor of 2, the sequence (\bar{N}_k) asymptotically minimizes the work-normalized variance of the deterministic dimension reduction algorithm. Moreover, by definition of \hat{q} , and since $C \leq \nu^*$,

$$R(\hat{q}; t, C) \leq R(q^*; t, C) \leq R(q^*; t, \nu^*),$$

where $q^* = \arg \min_{q \in A} R(q; t, \nu^*)$. Hence $R(\bar{q}; t, C) \leq 2R(q^*; t, \nu^*)$. Similarly, as $\nu^* \leq 2C$,

$$R(q^*; t, \nu^*) \leq R(\bar{q}; t, \nu^*) \leq 2R(\bar{q}; t, C).$$

Thus, the asymptotic work-normalized variances of the randomized dimension reduction algorithm, with $q = q^*$, and of the deterministic dimension reduction algorithm, with $N_k = \bar{N}_k$ for $k \geq 1$, are within a factor of 2 from each other.

Proposition 5.1 below shows that if, after generating the first copy of U , we generate the next $n(q_0 - q_1)$ samples by only changing the first component of the U in the previous iteration, and the next $n(q_1 - q_2)$ samples by only changing the first two components of the U in the previous iteration, and so on, the resulting algorithm is asymptotically less efficient than standard Monte Carlo.

Proposition 5.1. *Assume that $C(d-1) > 0$. Let $q \in A$ with $q_{d-1} < 1$. If $N_k = i$ for $1 \leq i \leq d$ and integer $k \in (n(1 - q_{i-1}), n(1 - q_i)]$, then $n \text{Var}(f_n) \rightarrow \infty$ as n goes to infinity.*

6 Comparison with a class of multilevel algorithms

We compare our method to a class of MLMC algorithms, adapted from (Giles 2008), that efficiently estimate $E(f(U))$ under the assumption that f is approximated, in the L^2 sense, by functions of its first arguments. Under conditions described in §6.1, we prove that, up to a constant, the randomized dimension reduction algorithm is at least as efficient as this class of MLMC algorithms. It should be stressed, however, that there may exist other MLMC algorithms that estimate $E(f(U))$ more efficiently than the class of MLMC algorithms described below.

6.1 The MLMC algorithms description and analysis

Let L be a positive integer and let (m_l) , $0 \leq l \leq L$, be a strictly increasing integral sequence, with $m_0 = 0$ and $m_L = d$. For $1 \leq l \leq L$, let ϕ_l be a square-integrable random variable equal to a deterministic measurable function of U_1, \dots, U_{m_l} , with $\phi_L = f(U)$. The ϕ_l 's are chosen so that, as l increases, ϕ_l gets closer to $f(U)$, in the L^2 sense. For instance, L can be proportional to $\ln(d)$, the m_l 's can increase exponentially with l , and ϕ_l could equal $f(U_1, \dots, U_{m_l}, \underbrace{x, \dots, x}_{d-m_l})$,

for some $x \in F$. For $1 \leq l \leq L$, let $\hat{\phi}_l$ be the average of n_l independent copies of $\phi_l - \phi_{l-1}$ (with $\phi_0 \triangleq 0$), where n_l is a positive integer to be specified later. Assume that the estimators $\hat{\phi}_1, \dots, \hat{\phi}_L$ are independent. As

$$E(f(U)) = \sum_{l=1}^L E(\phi_l - \phi_{l-1}),$$

$\hat{\phi} = \sum_{l=1}^L \hat{\phi}_l$ is an unbiased estimator of $E(f(U))$. Following the analysis in (Giles 2008),

$$\text{Var}(\hat{\phi}) = \sum_{l=1}^L \frac{V_l}{n_l},$$

where $V_l \triangleq \text{Var}(\phi_l - \phi_{l-1})$ for $1 \leq l \leq L$. The expected time needed to simulate $\hat{\phi}$ is $T_{\text{ML}} \triangleq \sum_{l=1}^L n_l \hat{t}_l$, where \hat{t}_l is the expected time needed to simulate $\phi_l - \phi_{l-1}$. As the variance of the average of n i.i.d. square-integrable random variables is proportional to $1/n$, for $\epsilon > 0$, we need $\lceil \text{Var}(\hat{\phi})\epsilon^{-2} \rceil$ independent samples of $\hat{\phi}$ to achieve an estimator variance at most ϵ^2 . Thus the total expected time $T^{\text{MLMC}}(\epsilon)$ needed for the MLMC algorithm to estimate $E(f(U))$ with variance at most ϵ^2 satisfies the relation

$$T^{\text{MLMC}}(\epsilon) = \Theta(T_{\text{ML}} + T_{\text{ML}} \text{Var}(\hat{\phi})\epsilon^{-2}). \quad (6.1)$$

The first term in the RHS of (6.1) accounts for the fact that $\hat{\phi}$ is simulated at least once. As shown in (Giles 2008), the work-normalized variance $T_{\text{ML}} \text{Var}(\hat{\phi})$ is minimized when the n_l 's are proportional to $\sqrt{V_l/\hat{t}_l}$ (ignoring the integrality constraints on the n_l 's), in which case

$$T_{\text{ML}} \text{Var}(\hat{\phi}) = \left(\sum_{l=1}^L \sqrt{V_l \hat{t}_l} \right)^2. \quad (6.2)$$

In line with (Giles 2008, Theorem 3.1), if $\hat{t}_l = O(2^l)$ and $\|\phi_l - \phi_L\|^2 = O(2^{\beta l})$, with $\beta < -1$, where the constants behind the O -notation do not depend on d , then $T_{\text{ML}} \text{Var}(\hat{\phi})$ is upper-bounded by a constant independent of d . This can be shown by observing that

$$V_l \leq \|\phi_l - \phi_{l-1}\|^2 \leq (\|\phi_l - \phi_L\| + \|\phi_{l-1} - \phi_L\|)^2.$$

Theorem 6.1 below shows that, under certain conditions, the randomized dimension reduction method is, up to a multiplicative constant, at least as efficient as the class of MLMC methods described above. Indeed, under the assumptions of Theorem 6.1, by (2.6), $T^{\text{tot}}(q, \epsilon) = O(d + T_{\text{ML}} \text{Var}(\hat{\phi})\epsilon^{-2})$. On the other hand, $T_{\text{ML}} \geq \hat{t}_L \geq \hat{c}d$ since $m_L = d$. Thus (6.1) implies that $T^{\text{MLMC}}(\epsilon) \geq c'(d + T_{\text{ML}} \text{Var}(\hat{\phi})\epsilon^{-2})$, for some constant c' .

Theorem 6.1. *Assume that there are constants c and \hat{c} independent of d such that $t_i \leq ci$ for $1 \leq i \leq d$, and $\hat{t}_l \geq \hat{c}m_l$ for $1 \leq l \leq L$, and that $C(d-1) > 0$. Then $R(q; t, \nu^*) \leq (32c/\hat{c})T_{\text{ML}} \text{Var}(\hat{\phi})$ if, for $0 \leq i \leq d-1$,*

$$q_i = \sqrt{\frac{C(i)}{(i+1)C(0)}}.$$

7 Numerical experiments

Our simulation experiments, using the examples in §4, were implemented in the C++ programming language. The randomized dimension reduction algorithm (RDR) was implemented as described in §3.3. The deterministic dimension reduction algorithm (DDR) was implemented similarly with $N_k = \bar{N}_k$. In both algorithms, n was chosen so that the expected total number of simulations of the U_i 's in iterations 2 through n is approximately $10d$. The actual total number of simulations of the U_i 's, denoted by ‘‘Cost’’ in our computer experiments, is about $11d$ because it includes the d simulations of the first iteration.

We have implemented the multilevel algorithm (MLMC) described in §6.1, with $L = \lfloor \log_2(d) \rfloor + 1$, and $m_l = \lfloor 2^{l-L}d \rfloor$ for $1 \leq l \leq L$, and $\phi_l = f(U_1, \dots, U_{m_l}, \underbrace{X_0, \dots, X_0}_{d-m_l})$. The V_l 's were estimated by Monte Carlo simulation with 1000 samples, and the n_l 's were scaled up so that the

Table 1: $\mathbb{P}(X_d > z)$ estimation in GARCH model, with $z = 4.4 \times 10^{-5}$, using 1000 samples, where X_d is the daily variance at time-step d .

		n	90% confidence interval	Std	Cost	Cost \times Std ²	VRF
$d = 1250$	RDR	277	$0.3918 \pm 2.1 \times 10^{-3}$	4.0×10^{-2}	$1.367 \times 10^4 \pm 8.9 \times 10^1$	21	14
	DDR	596	$0.3935 \pm 1.8 \times 10^{-3}$	3.4×10^{-2}	1.355×10^4	15	19
	MLMC	134	$0.3946 \pm 6.5 \times 10^{-3}$	1.2×10^{-1}	1.375×10^4	215	1.4
	QMC	4096	$0.39365 \pm 2.4 \times 10^{-4}$	4.6×10^{-3}	5.120×10^6	109	2.7
$d = 2500$	RDR	529	$0.3933 \pm 1.4 \times 10^{-3}$	2.8×10^{-2}	$2.745 \times 10^4 \pm 1.7 \times 10^2$	21	28
	DDR	1167	$0.3939 \pm 1.3 \times 10^{-3}$	2.4×10^{-2}	2.734×10^4	16	37
	MLMC	266	$0.3919 \pm 4.6 \times 10^{-3}$	8.9×10^{-2}	2.847×10^4	226	2.6
	QMC	4096	$0.39348 \pm 2.5 \times 10^{-4}$	4.8×10^{-3}	1.024×10^7	238	2.5
$d = 5000$	RDR	970	$0.3923 \pm 1.0 \times 10^{-3}$	2.0×10^{-2}	$5.490 \times 10^4 \pm 3.4 \times 10^2$	21	56
	DDR	1899	$0.39285 \pm 9.4 \times 10^{-4}$	1.8×10^{-2}	5.207×10^4	17	71
	MLMC	524	$0.3931 \pm 3.4 \times 10^{-3}$	6.5×10^{-2}	5.339×10^4	227	5.3
	QMC	4096	$0.39372 \pm 2.4 \times 10^{-4}$	4.6×10^{-3}	2.048×10^7	435	2.8

actual total number of simulations of the U_i 's is about $11d$. We have also implemented a randomized quasi-Monte Carlo method (QMC) with a random shift (Glasserman 2004, Section 5.4). Our implementation uses the C++ program available at <http://web.maths.unsw.edu.au/~fkuo/sobol> to generate d -dimensional Sobol sequences of length $n = 4096$. For practical reasons linked to computing time and storage cost, the QMC algorithm was tested for d up to 10^4 .

In Tables 1 through 6 and in the appendix, the variable Std refers to the standard deviation of f_n for the RDR and DDR algorithms, to the standard deviation of $\hat{\phi}$ for the MLMC algorithm, and to the standard deviation of the Quasi-Monte Carlo estimator for the QMC algorithm. The variable Std and a 90% confidence interval for $E(f(U))$ were estimated using 1000 independent runs of these two algorithms. For the RDR algorithm, a 90% confidence interval for the variable Cost was reported as well. The variance reduction factor VRF is defined as

$$\text{VRF} = \frac{d\text{Var}(f(U))}{\text{Cost} \times \text{Std}^2}.$$

We estimated $\text{Var}(f(U))$ by using 10000 independent samples of U .

7.1 GARCH volatility model

Table 1 shows results of our simulations of the GARCH volatility model for estimating $\mathbb{P}(X_d > z)$, with $z = 4.4 \times 10^{-5}$, $X_0 = 10^{-4}$, $\alpha = 0.06$, $\beta = 0.9$, and $w = 1.76 \times 10^{-6}$. As expected, the variable Cost is about $11d$ for the RDR, DDR, and MLMC algorithms. For these algorithms, the variable Cost \times Std² is roughly independent of d , and the variance reduction factors are roughly proportional to d . In contrast, for the QMC algorithm, the variable Cost \times Std² is roughly proportional to d , and the variance reduction factors are roughly constant. The 90% confidence interval of the RDR algorithm running time has a negligible length in comparison to the running time. The RDR algorithm outperforms the MLMC algorithm by about a factor of 10, and the QMC algorithm by a factor ranging from 5 to 20. In all our numerical experiments, the DDR algorithm outperforms the RDR algorithm by a factor between 1 and 2.

7.2 $G_t/D/1$ queue

Assume that A_i has a Poisson distribution with time-varying rate $\lambda_i = 0.75 + 0.5 \cos(\pi i/50)$, for $1 \leq i \leq d$, (recall that $A_0 = 0$). These parameters are taken from (Whitt and You 2016). Table 2 estimates $E(X_d)$, and Table 3 gives VRFs in the estimation of $\mathbb{P}(X_d > z)$, for selected values of z . Once again, for the RDR, DDR, and MLMC algorithms, the variable Cost \times Std² is roughly independent of d , and the variance reduction factors are roughly proportional to d . The VRFs of the RDR and DDR algorithms in Table 3 are greater than or equal to the corresponding VRFs in Table 2, which confirms the resiliency of these algorithms to discontinuities of g . In

Table 2: $E(X_d)$ estimation in $G_t/D/1$ queue, 1000 samples, where X_d is the number of customers in the queue at time-step d .

		n	90% confidence interval	Std	Cost	Cost \times Std ²	VRF
$d = 10^4$	RDR	2.3×10^3	$5.5243 \pm 4.6 \times 10^{-3}$	8.8×10^{-2}	$1.106 \times 10^5 \pm 6.3 \times 10^2$	8.5×10^2	1.8×10^2
	DDR	2.8×10^3	$5.5221 \pm 4.5 \times 10^{-3}$	8.6×10^{-2}	1.032×10^5	7.7×10^2	2.0×10^2
	MLMC	7.1×10^3	$5.524 \pm 5.9 \times 10^{-3}$	1.1×10^{-1}	1.076×10^5	1.4×10^3	1.1×10^2
$d = 10^5$	RDR	2.3×10^4	$5.5232 \pm 1.5 \times 10^{-3}$	2.8×10^{-2}	$1.101 \times 10^6 \pm 5.3 \times 10^3$	8.6×10^2	1.8×10^3
	DDR	3.6×10^4	$5.5229 \pm 1.5 \times 10^{-3}$	2.8×10^{-2}	1.046×10^6	8.3×10^2	1.8×10^3
	MLMC	6.3×10^4	$5.5241 \pm 2.0 \times 10^{-3}$	3.8×10^{-2}	1.096×10^6	1.6×10^3	9.5×10^2
$d = 10^6$	RDR	2.3×10^5	$5.52325 \pm 4.9 \times 10^{-4}$	9.4×10^{-3}	$1.103 \times 10^7 \pm 4.8 \times 10^4$	9.8×10^2	1.5×10^4
	DDR	2.6×10^5	$5.52363 \pm 4.5 \times 10^{-4}$	8.7×10^{-3}	1.047×10^7	7.9×10^2	1.9×10^4
	MLMC	6.6×10^5	$5.5231 \pm 5.7 \times 10^{-4}$	1.1×10^{-2}	1.119×10^7	1.4×10^3	1.1×10^4

Table 3: VRFs for $\mathbb{P}(X_d > z)$ estimation in $G_t/D/1$ queue.

z		0	2	4	6	8	10
$d = 10^4$	RDR	2.8×10^2	2.3×10^2	2.1×10^2	2.1×10^2	1.9×10^2	1.8×10^2
	DDR	4.7×10^2	2.8×10^2	2.3×10^2	2.4×10^2	2.2×10^2	1.9×10^2
	MLMC	2.1×10^1	3.7×10^1	5.1×10^1	6.5×10^1	7.7×10^1	7.2×10^1
$d = 10^5$	RDR	3.1×10^3	2.2×10^3	2.1×10^3	1.9×10^3	1.7×10^3	1.8×10^3
	DDR	4.3×10^3	3.0×10^3	2.4×10^3	2.5×10^3	2.2×10^3	2.0×10^3
	MLMC	2.1×10^2	3.6×10^2	4.5×10^2	5.9×10^2	5.9×10^2	5.8×10^2
$d = 10^6$	RDR	3.3×10^4	2.2×10^4	1.9×10^4	1.8×10^4	1.7×10^4	1.9×10^4
	DDR	4.1×10^4	2.9×10^4	2.2×10^4	2.2×10^4	1.9×10^4	2.2×10^4
	MLMC	2.0×10^3	3.3×10^3	4.4×10^3	5.3×10^3	5.7×10^3	5.8×10^3

contrast, the VRFs of the MLMC algorithm in Table 3 are lower than the corresponding VRFs in Table 2. The RDR algorithm outperforms the MLMC algorithm by a factor ranging from 1 to 2 in Table 2, and a factor ranging from 2 to 17 in Table 3. Table 4 estimates $E(X_d)$ for shifted values of d . The results in Table 4 are similar to those of Table 2, but the values of $E(X_d)$ in Table 4 are significantly smaller than those in Table 2. This can be explained by observing that λ_d is maximized (resp. minimized) at the values of d listed in Table 2 (resp. Table 4).

7.3 $M_t/GI/1$ queue

Assume that $\lambda_s = 0.75 + 0.5 \cos(\pi s/50)$ for $s \geq 0$. These parameters are taken from (Whitt and You 2016). Assume further that, for $j \geq 1$, the service time S_j for the j -th customer has a Pareto distribution with $\mathbb{P}(S_j \geq z) = (1 + z/\alpha)^{-3}$ for $z \geq 0$, for some constant $\alpha > 0$. A simple calculation shows that $E(S_j) = \alpha/2$. In our simulations, we have set $d = \lceil \theta \rceil$. Table 5 gives our simulation results for estimating $\mathbb{P}(W_\theta > 1)$ when $\alpha = 2$, and Table 6 lists VRFs for estimating $\mathbb{P}(W_\theta > 1)$ for selected values of α . Here again, for the RDR, DDR, and MLMC algorithms, the variable Cost \times Std² is roughly independent of d , and the VRFs are roughly proportional to d . The RDR algorithm outperforms the MLMC algorithm by a factor ranging from 1 to 10, depending on the value of α . In Table 6, the RDR, DDR, and MLMC algorithms become less efficient as α increases. This can be explained by noting that, as α increases, the length of the

Table 4: $E(X_d)$ estimation in $G_t/D/1$ queue, 1000 samples, with shifted dimensions.

		n	90% confidence interval	Std	Cost	Cost \times Std ²	VRF
$d = 10050$	RDR	8.5×10^2	$0.6599 \pm 4.1 \times 10^{-3}$	7.8×10^{-2}	$1.100 \times 10^5 \pm 6.1 \times 10^2$	6.7×10^2	6.2×10^1
	DDR	2.4×10^3	$0.6599 \pm 4.1 \times 10^{-3}$	7.9×10^{-2}	1.065×10^5	6.6×10^2	6.3×10^1
	MLMC	1.5×10^3	$0.659 \pm 4.5 \times 10^{-3}$	8.6×10^{-2}	1.106×10^5	8.2×10^2	5.0×10^1
$d = 100050$	RDR	7.8×10^3	$0.6594 \pm 1.3 \times 10^{-3}$	2.5×10^{-2}	$1.101 \times 10^6 \pm 5.3 \times 10^3$	7.2×10^2	5.6×10^2
	DDR	7.0×10^3	$0.6593 \pm 1.4 \times 10^{-3}$	2.6×10^{-2}	1.025×10^6	7.2×10^2	5.6×10^2
	MLMC	1.1×10^4	$0.6587 \pm 1.6 \times 10^{-3}$	3.1×10^{-2}	1.115×10^6	1.1×10^3	3.6×10^2
$d = 1000050$	RDR	6.7×10^4	$0.66007 \pm 4.5 \times 10^{-4}$	8.6×10^{-3}	$1.101 \times 10^7 \pm 4.8 \times 10^4$	8.2×10^2	5.3×10^3
	DDR	6.2×10^4	$0.66019 \pm 4.4 \times 10^{-4}$	8.4×10^{-3}	1.034×10^7	7.4×10^2	5.9×10^3
	MLMC	1.1×10^5	$0.66033 \pm 5.6 \times 10^{-4}$	1.1×10^{-2}	1.135×10^7	1.3×10^3	3.3×10^3

Table 5: $\mathbb{P}(W_\theta > 1)$ estimation in $M_t/GI/1$ queue, $\alpha = 2$, with 1000 samples, where W_θ is the residual work at time θ .

		n	90% confidence interval	Std	Cost	Cost \times Std ²
$\theta = 10^4$	RDR	3.9×10^3	$0.85389 \pm 4.9 \times 10^{-4}$	9.4×10^{-3}	$1.108 \times 10^5 \pm 6.2 \times 10^2$	10
	DDR	7.6×10^3	$0.85333 \pm 4.0 \times 10^{-4}$	7.7×10^{-3}	1.037×10^5	6
	MLMC	1.0×10^4	$0.8541 \pm 1.6 \times 10^{-3}$	3.1×10^{-2}	1.142×10^5	106
$\theta = 10^5$	RDR	3.0×10^4	$0.85385 \pm 1.6 \times 10^{-4}$	3.0×10^{-3}	$1.100 \times 10^6 \pm 5.3 \times 10^3$	10
	DDR	4.8×10^4	$0.85373 \pm 1.3 \times 10^{-4}$	2.5×10^{-3}	1.031×10^6	6
	MLMC	8.3×10^4	$0.85322 \pm 4.9 \times 10^{-4}$	9.5×10^{-3}	1.092×10^6	98
$\theta = 10^6$	RDR	2.5×10^5	$0.853762 \pm 5.1 \times 10^{-5}$	9.8×10^{-4}	$1.103 \times 10^7 \pm 4.9 \times 10^4$	11
	DDR	4.5×10^5	$0.853779 \pm 4.4 \times 10^{-5}$	8.4×10^{-4}	1.040×10^7	7
	MLMC	8.5×10^5	$0.8539 \pm 1.6 \times 10^{-4}$	3.2×10^{-3}	1.114×10^7	111

Table 6: VRF for $\mathbb{P}(W_\theta > 1)$ estimation in $M_t/GI/1$ queue.

α		0.5	1	1.5	2
$\theta = 10^4$	RDR	4.3×10^2	2.6×10^2	2.3×10^2	1.3×10^2
	DDR	5.6×10^2	4.0×10^2	3.0×10^2	2.0×10^2
	MLMC	3.2×10^2	1.4×10^2	5.2×10^1	1.2×10^1
$\theta = 10^5$	RDR	4.1×10^3	2.6×10^3	1.8×10^3	1.2×10^3
	DDR	5.2×10^3	3.8×10^3	2.4×10^3	1.9×10^3
	MLMC	2.3×10^3	1.1×10^3	5.3×10^2	1.2×10^2
$\theta = 10^6$	RDR	3.8×10^4	2.7×10^4	1.7×10^4	1.2×10^4
	DDR	4.7×10^4	3.5×10^4	2.5×10^4	1.7×10^4
	MLMC	2.6×10^4	1.1×10^4	4.4×10^3	1.1×10^3

last busy cycle increases as well, which renders W_θ more dependent on the first Y_i 's.

8 Conclusion

We have described a randomized dimension reduction algorithm that estimates $E(f(U))$ via Monte Carlo simulation, assuming that f does not depend equally on all its arguments. We formally prove that under some conditions, in order to achieve an estimator variance ϵ^2 , our algorithm requires $O(d + \epsilon^{-2})$ computations as opposed to $O(d\epsilon^{-2})$ under the standard Monte Carlo method. Our algorithm can be used to efficiently estimate the expected value of a function of the state of a Markov chain at time-step d , for a class of Markov chains driven by random variables. The numerical implementation of our algorithm uses a new geometric procedure of independent interest that solves in $O(d)$ time a d -dimensional optimisation problem that was previously solved in $O(d^3)$ time. We have argued intuitively that our method is resilient to discontinuities of f , and have described and analysed a deterministic version of our algorithm. Our numerical experiments confirm that our approach highly outperforms the standard Monte Carlo method for large values of d , and show its high resilience to discontinuities. Whether our approach can be combined with the Quasi-Monte Carlo method to produce a provably efficient estimator is left for future work.

Acknowledgments

This research has been presented at the 9th NIPS Workshop on Optimization for Machine Learning, Barcelona, December 2016, the Stochastic Methods in Finance seminar at École des Ponts Paris-Tech, March 2018, and the 23rd International Symposium on Mathematical Programming, Bordeaux, July 2018. The author thanks Bernard Lapeyre, seminar and conference participants for helpful conversations. He is grateful to three anonymous referees, an anonymous associate editor, and Baris Ata (department editor), for insightful comments and suggestions. This work was achieved through the Laboratory of Excellence on Financial Regulation (Labex ReFi) under the reference ANR-10-LABX-0095. It benefitted from a French government support managed by the National Research Agency (ANR).

A Proof of Proposition 2.1

For $0 \leq i \leq d$, let

$$f^{(i)} = E(f(U)|U_{i+1}, \dots, U_d).$$

Thus, $C(i) = \text{Var}(f^{(i)})$. By the tower law, for $0 \leq i \leq d-1$,

$$E(f^{(i)}|U_{i+2}, \dots, U_d) = f^{(i+1)},$$

and so

$$C(i+1) = \text{Var}(E(f^{(i)}|U_{i+2}, \dots, U_d)).$$

As the variance decreases by taking the conditional expectation, it follows that $C(i+1) \leq C(i)$, as desired.

We now prove (2.2). Let $W = (U'_1, \dots, U'_i, U_{i+1}, \dots, U_d)$. Since U and W are conditionally independent given U_{i+1}, \dots, U_d , and $E(f(W)|U_{i+1}, \dots, U_d) = f^{(i)}$,

$$E(f(U)f(W)|U_{i+1}, \dots, U_d) = (f^{(i)})^2.$$

Hence, by the tower law,

$$E(f(U)f(W)) = E((f^{(i)})^2).$$

On the other hand, using the tower law once again,

$$E(f(U)) = E(f(W)) = E(f^{(i)}),$$

and so the RHS of (2.2) is equal to $\text{Var}(f^{(i)})$, as required. \square

B Proof of Theorem 2.1

We first prove Lemma B.1 below, which follows by classical calculations (see e.g. (Asmussen and Glynn 2007, Section IV.6a)).

Lemma B.1. *Let (Z_k) , $k \geq 1$, be an homogeneous stationary Markov chain in \mathbb{R}^d , and let g be a real-valued Borel-measurable function on \mathbb{R}^d such that $g(Z_1)$ is square-integrable, and $a_j = \text{Cov}(g(Z_1), g(Z_{1+j}))$ is non-negative for $j \geq 0$. Assume that $\sum_{j=1}^{\infty} a_j$ is finite. Then*

$$n^{-1} \text{Var}\left(\sum_{m=1}^n g(Z_m)\right) \leq a_0 + 2 \sum_{j=1}^{\infty} a_j. \quad (\text{B.1})$$

Furthermore, the LHS of (B.1) converges to its RHS as n goes to infinity.

Proof. Since (Z_k) , $k \geq 1$, is homogeneous and stationary, $\text{Cov}(g(Z_m)g(Z_{m+j})) = a_j$ for $m \geq 1$ and $j \geq 0$. Thus,

$$\begin{aligned} \text{Var}\left(\sum_{m=1}^n g(Z_m)\right) &= \sum_{m=1}^n \text{Var}(g(Z_m)) + 2 \sum_{1 \leq m < m+j \leq n} \text{Cov}(g(Z_m)g(Z_{m+j})) \\ &= na_0 + 2 \sum_{j=1}^n (n-j)a_j. \end{aligned}$$

Hence

$$n^{-1} \text{Var}\left(\sum_{m=1}^n g(Z_m)\right) = a_0 + 2 \sum_{j=1}^n \frac{n-j}{n} a_j,$$

which implies (B.1). Using Lebesgue's dominated convergence theorem concludes the proof. \square

We now prove Theorem 2.1. Since the $C(i)$'s and the variance of f_n remain unchanged if we add a constant to f , we can assume without loss of generality that $E(f(U)) = 0$. For $0 \leq i \leq d-1$, set $p_i = 1 - q_i$. Let $m < k$ be two integers in $[1, n]$. For $0 \leq i \leq d-1$,

$$\mathbb{P}(\max_{m \leq j < k} N_j \leq i) = p_i^{k-m}$$

and so, for $1 \leq i \leq d-1$,

$$\mathbb{P}(\max_{m \leq j < k} N_j = i) = p_i^{k-m} - p_{i-1}^{k-m}.$$

For $i \in [1, d]$, conditional on the event $\max_{m \leq j < k} N_j = i$, the first i components of $V^{(m)}$ and $V^{(k)}$ are independent, and the last $d-i$ components of $V^{(m)}$ and $V^{(k)}$ are the same. This is because, if $N_j = i$, with $m \leq j < k$, the vector $V^{(m)}$ and the first i components of $V^{(j+1)}$ are independent. Thus, by Proposition 2.1,

$$E(f(V^{(m)})f(V^{(k)}) | \max_{m \leq j < k} N_j = i) = C(i). \quad (\text{B.2})$$

As $C(d) = 0$, it follows from Bayes' formula that

$$\begin{aligned} E(f(V^{(m)})f(V^{(k)})) &= \sum_{i=1}^{d-1} \mathbb{P}(\max_{m \leq j < k} N_j = i) C(i) \\ &= \sum_{i=1}^{d-1} (p_i^{k-m} - p_{i-1}^{k-m}) C(i). \end{aligned}$$

Let $a_j = \text{Cov}(f(V^{(1)}), f(V^{(1+j)}))$. Thus, for $j > 0$,

$$a_j = \sum_{i=1}^{d-1} (p_i^j - p_{i-1}^j) C(i)$$

is non-negative, and

$$\begin{aligned} \sum_{j=1}^{\infty} a_j &= \sum_{i=1}^{d-1} \left(\frac{p_i}{1-p_i} - \frac{p_{i-1}}{1-p_{i-1}} \right) C(i) \\ &= \sum_{i=1}^{d-1} \left(\frac{1}{1-p_i} - \frac{1}{1-p_{i-1}} \right) C(i) \\ &= -C(1) + \sum_{i=1}^{d-1} \frac{C(i) - C(i+1)}{q_i} \end{aligned}$$

is finite. Since $a_0 = C(0)$, it follows that

$$a_0 + 2 \sum_{j=1}^{\infty} a_j = C(0) - 2C(1) + 2 \sum_{i=1}^{d-1} \frac{C(i) - C(i+1)}{q_i}. \quad (\text{B.3})$$

We conclude the proof using Lemma B.1. □

C Proof of Proposition 2.2

Let U'_1, \dots, U'_i be random variables satisfying the conditions of Proposition 2.1, and $W = (U'_1, \dots, U'_i, U_{i+1}, \dots, U_d)$. Since (U_1, \dots, U_i) and W are independent,

$$\text{Cov}(f_i(U_1, \dots, U_i), f(W)) = 0.$$

Similarly,

$$\text{Cov}(f_i(U_1, \dots, U_i), f_i(U'_1, \dots, U'_i)) = \text{Cov}(f(U), f_i(U'_1, \dots, U'_i)) = 0.$$

Thus, by Proposition 2.1 and bilinearity of the covariance,

$$\begin{aligned} C(i) &= \text{Cov}(f(U) - f_i(U_1, \dots, U_i), f(W) - f_i(U'_1, \dots, U'_i)) \\ &\leq \text{Std}(f(U) - f_i(U_1, \dots, U_i)) \text{Std}(f(W) - f_i(U'_1, \dots, U'_i)) \\ &= \text{Var}(f(U) - f_i(U_1, \dots, U_i)). \end{aligned}$$

The last equation follows by observing that $f(W) - f_i(U'_1, \dots, U'_i) \stackrel{d}{=} f(U) - f_i(U_1, \dots, U_i)$. \square

D Proofs of Propositions 2.3 and 2.5

We first prove the following.

Proposition D.1. *Let $\nu = (\nu_0, \dots, \nu_d)$ be an element of $\mathbb{R}^d \times \{0\}$. Assume that ν_0, \dots, ν_{d-1} are positive, and that the sequence (ν_i/t_{i+1}) , $0 \leq i \leq d-1$, is decreasing. For $0 \leq i \leq d-1$, set $q_i = \sqrt{(t_1 \nu_i)/(\nu_0 t_{i+1})}$. Then*

$$R(q; t, \nu) \leq 4 \left(\sum_{i=0}^{d-1} \sqrt{\nu_i} (\sqrt{t_{i+1}} - \sqrt{t_i}) \right)^2.$$

Proof. Applying the inequality $x - y \leq 2\sqrt{x}(\sqrt{x} - \sqrt{y})$, which holds for $x \geq 0$ and $y \geq 0$, to $x = \nu_i$ and $y = \nu_{i+1}$ yields

$$\begin{aligned} \sum_{i=0}^{d-1} \frac{\nu_i - \nu_{i+1}}{q_i} &\leq 2\sqrt{\frac{\nu_0}{t_1}} \sum_{i=0}^{d-1} (\sqrt{\nu_i} - \sqrt{\nu_{i+1}}) \sqrt{t_{i+1}} \\ &= 2\sqrt{\frac{\nu_0}{t_1}} \sum_{i=0}^{d-1} \sqrt{\nu_i} (\sqrt{t_{i+1}} - \sqrt{t_i}). \end{aligned}$$

Similarly, since $t_{i+1} - t_i \leq 2\sqrt{t_{i+1}}(\sqrt{t_{i+1}} - \sqrt{t_i})$,

$$\sum_{i=0}^{d-1} q_i (t_{i+1} - t_i) \leq 2\sqrt{\frac{t_1}{\nu_0}} \sum_{i=0}^{d-1} \sqrt{\nu_i} (\sqrt{t_{i+1}} - \sqrt{t_i}).$$

Taking the product completes the proof. \square

We now prove Proposition 2.3. Set $t' = (t'_0, \dots, t'_d)$, with $t'_i = ci$, $0 \leq i \leq d$, and let $\nu = (\nu_0, \dots, \nu_{d-1}, 0)$, with $\nu_i = c'(i+1)^\gamma$, $0 \leq i \leq d-1$. Thus $R(q; t, \nu^*) \leq R(q; t', 2\nu)$ since $t \leq t'$ and $\nu^* \leq 2\nu$. By Proposition D.1,

$$R(q; t', 2\nu) \leq 8cc' \left(\sum_{i=0}^{d-1} (i+1)^{\gamma/2} (\sqrt{i+1} - \sqrt{i}) \right)^2.$$

As $\sqrt{i+1} - \sqrt{i} \leq (i+1)^{-1/2}$ for $i \geq 0$, it follows that

$$R(q; t, \nu^*) \leq 8cc' \left(\sum_{i=1}^d i^{(\gamma-1)/2} \right)^2.$$

The inequality

$$\sum_{i=1}^d i^{(\gamma-1)/2} \leq 1 + \int_1^d x^{(\gamma-1)/2} dx$$

implies that

$$\sum_{i=0}^{d-1} i^{(\gamma-1)/2} \leq \begin{cases} 1 + 2/(\gamma + 1), & \gamma < -1, \\ 1 + \ln(d), & \gamma = -1, \\ 2d^{(\gamma+1)/2}/(\gamma + 1), & -1 < \gamma < 0. \end{cases}$$

We conclude that there is a constant c_1 such that (2.10) holds. This concludes the proof of Proposition 2.3. \square

We now prove Proposition 2.5. By applying Proposition D.1 to $\nu = (2C(0), \dots, 2C(d-1), 0)$, it follows that $R(q; t, \nu)$ is upper-bounded by the RHS of (2.11). Since $R(q; t, \nu^*) \leq R(q; t, \nu)$, this implies (2.11). \square

E Proof of Proposition 2.6

Since $R(q; t, \nu)$ is increasing with respect to ν ,

$$\begin{aligned} R(q; t, \nu^*) &\geq R(q; t, C(0), \dots, C(d)) \\ &= \left(\sum_{i=0}^{d-1} \frac{C(i) - C(i+1)}{q_i} \right) \left(\sum_{i=0}^{d-1} q_i (t_{i+1} - t_i) \right). \end{aligned}$$

However, as $\sum_{j=0}^{d-1} q_j (t_{j+1} - t_j) \geq q_i t_{i+1}$ for $0 \leq i \leq d-1$, and since $(C(i))$ is a decreasing sequence,

$$\left(\sum_{i=0}^{d-1} \frac{C(i) - C(i+1)}{q_i} \right) \left(\sum_{j=0}^{d-1} q_j (t_{j+1} - t_j) \right) \geq \sum_{i=0}^{d-1} (C(i) - C(i+1)) t_{i+1}.$$

This implies (2.12) since

$$\sum_{i=0}^{d-1} (C(i) - C(i+1)) t_{i+1} = \sum_{i=0}^{d-1} C(i) (t_{i+1} - t_i). \quad (\text{E.1})$$

Assume now that $q_i = t_1/t_{i+1}$ for $0 \leq i \leq d-1$. Since $R(q; t, \nu)$ is increasing with respect to ν ,

$$\begin{aligned} R(q; t, \nu^*) &\leq R(q; t, 2C(0), \dots, 2C(d)) \\ &= 2 \left(\sum_{i=0}^{d-1} \frac{C(i) - C(i+1)}{q_i} \right) \left(\sum_{i=0}^{d-1} q_i (t_{i+1} - t_i) \right) \\ &= 2 \left(\sum_{i=0}^{d-1} (C(i) - C(i+1)) t_{i+1} \right) \left(\sum_{i=0}^{d-1} \frac{t_{i+1} - t_i}{t_{i+1}} \right). \end{aligned}$$

Furthermore, as $(t_{i+1} - t_i)/t_{i+1} \leq \ln(t_{i+1}/t_i)$ for $1 \leq i \leq d-1$,

$$\sum_{i=0}^{d-1} \frac{t_{i+1} - t_i}{t_{i+1}} \leq 1 + \ln\left(\frac{t_d}{t_1}\right).$$

Using (E.1) once again yields (2.13). \square

F Relation with the ANOVA decomposition and the truncation dimension

This section assumes that f is a square-integrable function on $[0, 1]^d$, and that each U_i is uniformly distributed on $[0, 1]$, with $\text{Var}(f(U)) > 0$. Consider a decomposition of f in the following form:

$$f = \sum_{Y \subseteq \{1, \dots, d\}} f_Y, \quad (\text{F.1})$$

where f_Y is a measurable function on $[0, 1]^d$ and $f_Y(u)$ depends on u only through $(u_j)_{j \in Y}$, for $u = (u_1, \dots, u_d) \in [0, 1]^d$. For instance, f_\emptyset is a constant, $f_{\{j\}}(u)$ is a function of u_j , and $f_{\{j, k\}}(u)$ is a function of (u_j, u_k) . The relation (F.1) is called ANOVA representation of f if, for $Y \subseteq \{1, \dots, d\}$, any vector $u \in [0, 1]^d$, and any $j \in Y$,

$$\int_0^1 f_Y(u_1, \dots, u_{j-1}, x, u_{j+1}, \dots, u_d) dx = 0.$$

It can be shown (Sobol 2001, p. 272) that there is a unique ANOVA representation of f , that the f_Y 's are square-integrable, and that

$$\text{Var}(f(U)) = \sum_{Y \subseteq \{1, \dots, d\}} \sigma_Y^2,$$

where σ_Y is the standard deviation of $f_Y(U)$. Furthermore, for $0 \leq i \leq d - 1$,

$$E(f(U)|U_{i+1}, \dots, U_d) = \sum_{Y \subseteq \{i+1, \dots, d\}} f_Y(U),$$

and the covariance between $f_Y(U)$ and $f_{Y'}(U)$ is null if $Y \neq Y'$. Hence, for $0 \leq i \leq d - 1$,

$$C(i) = \sum_{Y \subseteq \{i+1, \dots, d\}} \sigma_Y^2. \quad (\text{F.2})$$

It follows that $C(i)$ is equal to the variance corresponding to the subset $\{i + 1, \dots, d\}$, as defined in (Sobol 2001, Eq. 4). Thus, $C(i)/C(0)$ is equal to the global sensitivity index $S_{\{i+1, \dots, d\}}$ for the subset $\{i + 1, \dots, d\}$ (see (Sobol 2001, Definition 3)).

Proposition F.1 below relates the performance of our algorithm to the truncation dimension d_t of f , defined in (Owen 2003) as

$$d_t \triangleq \frac{\sum_{Y \subseteq \{1, \dots, d\}, Y \neq \emptyset} \max(Y) \sigma_Y^2}{\text{Var}(f(U))}.$$

Under the conditions in Proposition F.1, if $t_d = \Theta(d)$ and d_t is upper bounded by a constant independent of d , the asymptotic (as n goes to infinity) work-normalized variance of our algorithm is $O(\ln(d)\text{Var}(f(U)))$, whereas the work-normalized variance of the standard Monte Carlo algorithm is $\Theta(d\text{Var}(f(U)))$.

Proposition F.1. *Assume that there is a real number c such that $t_i \leq ci$ for $1 \leq i \leq d$, and that $q_i = 1/(i + 1)$ for $0 \leq i \leq d - 1$. Then*

$$R(q; t, \nu^*) \leq 2c(1 + \ln(d))d_t \text{Var}(f(U)).$$

Proof. For $0 \leq i \leq d - 1$, we can rewrite (F.2) as

$$C(i) = \sum_{Y \subseteq \{1, \dots, d\}, Y \neq \emptyset} \mathbf{1}\{i < \min(Y)\} \sigma_Y^2.$$

Hence,

$$\begin{aligned}
\sum_{i=0}^{d-1} C(i) &= \sum_{Y \subseteq \{1, \dots, d\}, Y \neq \emptyset} \sum_{i=0}^{d-1} \mathbf{1}\{i < \min(Y)\} \sigma_Y^2 \\
&= \sum_{Y \subseteq \{1, \dots, d\}, Y \neq \emptyset} \min(Y) \sigma_Y^2 \\
&\leq d_t \text{Var}(f(U)).
\end{aligned}$$

Let $t'_i = ci$ for $0 \leq i \leq d$, and $t' = (t'_0, \dots, t'_d)$. The proof of (2.13) shows that this relation still holds if t is replaced by any increasing sequence of length $d + 1$ starting at 0. Replacing t by t' implies that

$$\begin{aligned}
R(q; t', \nu^*) &\leq 2c(1 + \ln(d)) \sum_{i=0}^{d-1} C(i) \\
&\leq 2c(1 + \ln(d)) d_t \text{Var}(f(U)).
\end{aligned}$$

Moreover, $R(q; t, \nu^*) \leq R(q; t', \nu^*)$ since $t \leq t'$. This completes the proof. \square

G Proof of Theorem 3.1

We use the following proposition, whose proof follows immediately from (2.4).

Proposition G.1. *If $\nu \in \mathbb{R}^d \times \{0\}$ and $\nu' \in \mathbb{R}^d \times \{0\}$ are such that $\nu' \leq \nu$, and $q \in A$, then $R(q; t, \nu') \leq R(q; t, \nu)$, with equality if $\nu_0 = \nu'_0$ and, for $1 \leq i \leq d - 1$, $(\nu_i - \nu'_i)(q_{i-1} - q_i) = 0$.*

By definition of the lower hull, (θ_i) , $0 \leq i \leq d - 1$, is an increasing sequence. Furthermore, $\theta_{d-1} < 0$ since it is equal to the slope of a segment joining (t_i, ν_i) to $(t_d, 0)$, for some $i \in [0, d - 1]$. Hence $\theta_i < 0$ for $0 \leq i \leq d - 1$, and so q^* is well defined and belongs to A . Furthermore, (ν'_i) , $0 \leq i \leq d$, is a decreasing sequence, and $\nu'_d = \nu_d = 0$. On the other hand, by (2.5),

$$R(q^*; t, \nu') = \left(\sum_{i=0}^{d-1} \sqrt{(\nu'_i - \nu'_{i+1})(t_{i+1} - t_i)} \right)^2.$$

Since, by the Cauchy-Schwartz inequality, for all non-negative sequences (x_i) and (y_i) ,

$$\left(\sum_{i=0}^{d-1} \sqrt{x_i y_i} \right)^2 \leq \left(\sum_{i=0}^{d-1} x_i \right) \left(\sum_{i=0}^{d-1} y_i \right),$$

it follows that $R(q^*; t, \nu') \leq R(q; t, \nu')$ for $q \in A$. Furthermore, by Proposition G.1, $R(q; t, \nu') \leq R(q; t, \nu)$, and so $R(q^*; t, \nu') \leq R(q; t, \nu)$. On the other hand, $(\nu_i - \nu'_i)(q_{i-1}^* - q_i^*) = 0$ for $1 \leq i \leq d - 1$. This is because, if $\nu_i \neq \nu'_i$, then the point (t_i, ν_i) does not belong to the lower hull of the set $\{(t_j, \nu_j) : 0 \leq j \leq d\}$. Hence (t_i, ν'_i) belongs to the segment $((t_{i-1}, \nu'_{i-1}), (t_{i+1}, \nu'_{i+1}))$, which implies that $\theta_{i-1} = \theta_i$ and $q_{i-1}^* = q_i^*$. Thus, as $\nu_0 = \nu'_0$, Proposition G.1 shows that $R(q^*; t, \nu) = R(q^*; t, \nu')$. This implies (3.2) and that $R(q^*; t, \nu) \leq R(q; t, \nu)$ for $q \in A$, as desired. \square

H Proof of Proposition 3.1

The random vectors $W^{(1)} = (U_1, \dots, U_i, U'_{i+1}, \dots, U'_d)$, $W^{(2)} = (U'_1, \dots, U'_i, U_{i+1}, \dots, U_d)$, and $W^{(3)} = (U'_1, \dots, U'_i, U''_{i+1}, \dots, U''_d)$ have the same distribution as U . Hence

$$E((f(U) - f(W^{(1)}))(f(W^{(2)}) - f(W^{(3)}))) = \text{Cov}(f(U) - f(W^{(1)}), f(W^{(2)}) - f(W^{(3)})).$$

As $W^{(1)}$ and $W^{(2)}$ are independent,

$$\text{Cov}(f(W^{(1)}), f(W^{(2)})) = 0.$$

Similarly,

$$\text{Cov}(f(W^{(1)}), f(W^{(3)})) = \text{Cov}(f(U), f(W^{(3)})) = 0.$$

By bilinearity of the covariance, it follows that

$$E((f(U) - f(W^{(1)}))(f(W^{(2)}) - f(W^{(3)}))) = \text{Cov}(f(U), f(W^{(2)})).$$

We conclude the proof using Proposition 2.1. \square

I Proof of Proposition 4.1

We first prove the following Markov property.

Proposition I.1. *Let $i \in [0, d]$. If H is a bounded random variable which is measurable with respect to the σ -algebra generated by X_i, Y_i, \dots, Y_{d-1} , then*

$$E(H|Y_0, \dots, Y_{i-1}) = E(H|X_i). \quad (\text{I.1})$$

Proof. Let \mathcal{H} be the vector space of bounded real-valued random variables H satisfying (I.1). Clearly, the constant random variables belong to \mathcal{H} . Let (H_m) , $m \geq 0$, be an increasing sequence of positive elements of \mathcal{H} such that $H = \sup_{m \geq 0} H_m$ is bounded. For $m \geq 0$,

$$E(H_m|Y_0, \dots, Y_{i-1}) = E(H_m|X_i). \quad (\text{I.2})$$

By the conditional Lebesgue dominated convergence theorem (Shiryaev 1996, Theorem 2, p. 218), the LHS (resp. RHS) of (I.2) converges to $E(H|Y_0, \dots, Y_{i-1})$ (resp. $E(H|X_i)$) as m goes to infinity, and so $H \in \mathcal{H}$.

Let \mathcal{G} (resp. \mathcal{G}') be the set of bounded real-valued random variables which are measurable with respect to the σ -algebra generated by X_i (resp. (Y_i, \dots, Y_{d-1})), and let \mathcal{C} be the set of random variables of the form GG' , with $G \in \mathcal{G}$ and $G' \in \mathcal{G}'$. For $G \in \mathcal{G}$ and $G' \in \mathcal{G}'$,

$$\begin{aligned} E(GG'|Y_0, \dots, Y_{i-1}) &= GE(G'|Y_0, \dots, Y_{i-1}) \\ &= GE(G'). \end{aligned}$$

The first equation holds since X_i is a measurable function of Y_0, \dots, Y_{i-1} , which implies that G is measurable with respect to the σ -algebra generated by Y_0, \dots, Y_{i-1} . The second equation follows from the independence of (Y_i, \dots, Y_{d-1}) and (Y_0, \dots, Y_{i-1}) . Similarly, since (Y_i, \dots, Y_{d-1}) and X_i are independent,

$$E(GG'|X_i) = GE(G'|X_i) = GE(G').$$

Thus, $GG' \in \mathcal{H}$, and so $\mathcal{C} \subseteq \mathcal{H}$. As \mathcal{C} is closed under pointwise multiplication, by the monotone class theorem (Revuz and Yor 1999, Theorem 2.2, p. 3), \mathcal{H} contains all bounded random variables which are measurable with respect to the σ -algebra generated by the elements of \mathcal{C} . Since Y_i, \dots, Y_{d-1} and X_i belong to \mathcal{C} , it follows that \mathcal{H} contains all bounded random variables which are measurable with respect to the σ -algebra generated by X_i, Y_i, \dots, Y_{d-1} . This completes the proof. \square

As $g(X_d)$ is a measurable function of $(X_i, Y_i, \dots, Y_{d-1})$, for any integer m , the random variable $\min(m, g^+(X_d))$ belongs to \mathcal{H} . By the conditional Lebesgue dominated convergence theorem, taking the limit as m goes to infinity implies that

$$E(g^+(X_d)|Y_0, \dots, Y_{i-1}) = E(g^+(X_d)|X_i).$$

A similarly equation holds for $g^-(X_d)$. Thus,

$$E(g(X_d)|Y_0, \dots, Y_{i-1}) = E(g(X_d)|X_i).$$

Replacing i with $d - i$ implies that

$$E(g(X_d)|U_{i+1}, \dots, U_d) = E(g(X_d)|X_{d-i}).$$

Taking the variance of both sides concludes the proof. \square

J Proof of Proposition 4.2

For $0 \leq i \leq d - 1$, set $Z_i = \alpha Y_i^2 + \beta$, so that $X_{i+1} = w + Z_i X_i$. It can be shown by induction on i that $X_d = w + wZ_{d-1}Z' + Z''$ for $2 \leq i \leq d$, where

$$Z' = \sum_{j=1}^{i-1} \prod_{k=2}^j Z_{d-k}$$

and

$$Z'' = Z_{d-i} \cdots Z_{d-1} X_{d-i}.$$

By convention, the product over an empty set is equal to 1. Since $X_{i,0}$ is equal to the state of the Markov chain at step d if $X_{d-i} = 0$, it follows that $X_{i,0} = w + wZ_{d-1}Z'$. Note that $E(Z'') = (\alpha + \beta)^i E(X_{d-i})$ as $E(Y_i^2) = 1$. By (Hull 2014, Eq. 23.13), for $0 \leq m \leq d$,

$$E(X_m) \leq \max(X_0, \frac{w}{1 - \alpha - \beta}),$$

and so

$$E(Z'') \leq \max(X_0, \frac{w}{1 - \alpha - \beta})(\alpha + \beta)^i.$$

Since the density of Y_{d-1} is upper-bounded by $1/2$, for $\gamma \leq \gamma'$,

$$\begin{aligned} \mathbb{P}(\gamma \leq Z_{d-1} \leq \gamma') &= 2\mathbb{P}\left(\sqrt{\frac{(\gamma - \beta)^+}{\alpha}} \leq Y_{d-1} \leq \sqrt{\frac{(\gamma' - \beta)^+}{\alpha}}\right) \\ &\leq \sqrt{\frac{(\gamma' - \beta)^+}{\alpha}} - \sqrt{\frac{(\gamma - \beta)^+}{\alpha}} \\ &\leq \sqrt{\frac{\gamma' - \gamma}{\alpha}}. \end{aligned} \tag{J.1}$$

The last equation follows from the inequality $\sqrt{y'} - \sqrt{y} \leq \sqrt{y' - y}$, which holds for $0 \leq y \leq y'$. On the other hand, as $X_{i,0} \leq X_d$,

$$g(X_d) - g(X_{i,0}) = \begin{cases} 1 & \text{if } X_{i,0} \leq z < X_d, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, by (4.1),

$$\begin{aligned} C(i) &\leq \|g(X_d) - g(X_{i,0})\|^2 \\ &= \mathbb{P}(X_{i,0} \leq z < X_d). \end{aligned}$$

But

$$\begin{aligned}
\mathbb{P}(X_{i,0} \leq z < X_d) &= \mathbb{P}\left(\frac{z - Z'' - w}{wZ'} < Z_{d-1} \leq \frac{z - w}{wZ'}\right) \\
&= E\left(\mathbb{P}\left(\frac{z - Z'' - w}{wZ'} < Z_{d-1} \leq \frac{z - w}{wZ'} \mid Z', Z''\right)\right) \\
&\leq E\left(\sqrt{\frac{Z''}{\alpha w Z'}}\right) \\
&\leq \kappa(\alpha + \beta)^{i/2},
\end{aligned}$$

where κ is a constant that depends only on w , X_0 , α and β . The third equation follows from (J.1) and the independence of Z_{d-1} and (Z', Z'') . The last equation follows from Jensen's inequality and the inequality $Z' \geq 1$. Hence $C(i) \leq \kappa(\alpha + \beta)^{i/2}$ for $2 \leq i \leq d$. This inequality also holds for $0 \leq i \leq 1$ by replacing κ with $\max(\kappa, 1/(\alpha + \beta))$. \square

K Proof of Proposition 4.3

By classical calculations (see, e.g., (Asmussen and Glynn 2007, §I, Eq. (1.4))), it can be shown by induction on d that $X_d = \max_{0 \leq j \leq d} S_j$, where $S_j = \sum_{k=d-j}^{d-1} Y_k$ for $1 \leq j \leq d$, with $S_0 = 0$. For $0 \leq i \leq d-1$, since $X_{i,0}$ is the number of customers in the queue at time-step d if there are no costumers in the queue at time-step $d-i$, it can be shown by induction on d that $X_{i,0} = \max_{0 \leq j \leq i} S_j$. Hence

$$X_d = \max(X_{i,0}, \max_{i+1 \leq j \leq d} S_j).$$

Thus

$$X_d - X_{i,0} \leq \max_{i+1 \leq j \leq d} S_j^+,$$

and so

$$\begin{aligned}
C(i) &\leq \|X_d - X_{i,0}\|^2 \\
&\leq \sum_{j=i+1}^d \|S_j^+\|^2.
\end{aligned}$$

On the other hand, since the Y_i 's are independent, $E(e^{\gamma S_j}) \leq \kappa^j$ for $0 \leq j \leq d$. Furthermore, as $(x^+)^2 \leq 2e^x$ for $x \in \mathbb{R}$, $\gamma^2 (S_j^+)^2 \leq 2e^{\gamma S_j}$. Taking expectations implies that $\gamma^2 \|S_j^+\|^2 \leq 2\kappa^j$, and so $C(i) \leq \gamma' \kappa^i$, where $\gamma' = 2\gamma^{-2}/(1 - \kappa)$. \square

L Proof of Proposition 4.4

It can be shown by induction on the number of arrivals in $(\theta', \theta]$ (see also (Ma and Whitt 2017)) that, if the system is empty at time $\theta' \in [0, \theta]$,

$$W_\theta = \sup_{s \in [\theta', \theta]} \{Z_\theta(s) - (\theta - s)\}.$$

Hence

$$X_d = \sup_{s \in [0, \theta]} \{Z_\theta(s) - (\theta - s)\}.$$

Also, for $0 \leq i \leq d$, since $X_{i,0}$ is the residual work at time θ if the residual work at time $(d-i)\theta/d$ is 0,

$$X_{i,0} = \sup_{s \in [(d-i)\theta/d, \theta]} \{Z_\theta(s) - (\theta - s)\}.$$

Thus, by a calculation similar to the proof of Proposition 4.3,

$$X_d - X_{i,0} \leq \sup_{0 \leq s \leq (d-i)\theta/d} \{Z_\theta(s) - (\theta - s)\}^+.$$

For non-negative integer j , set

$$S_j = Z_\theta((\theta - \frac{j}{\lambda^*})^+) - \frac{j-1}{\lambda^*}.$$

Fix $s \in [0, (d-i)\theta/d]$, and let $j = \lceil (\theta - s)\lambda^* \rceil$. As $\theta - j/\lambda^* \leq s$ and Z_θ is a decreasing function,

$$Z_\theta(s) \leq Z_\theta((\theta - \frac{j}{\lambda^*})^+).$$

Since $\theta - s \geq (j-1)/\lambda^*$, this implies that $Z_\theta(s) - (\theta - s) \leq S_j$. But $2j \geq i$ as $j \geq i\lambda^*\theta/d$ and $\lambda^*\theta \geq d/2$. Hence

$$X_d - X_{i,0} \leq \sup_{j \geq i/2} S_j^+.$$

By calculations similar to the proof of Proposition 4.3, it follows that

$$C(i) \leq \sum_{j \geq i/2} \|S_j^+\|^2.$$

On the other hand, (4.3) implies that $E(e^{\gamma(S_{j+1}-S_j)}) \leq \kappa$ for $j \geq 0$, and so $E(e^{\gamma S_j}) \leq e^{\gamma/\lambda^*} \kappa^j$. We conclude the proof in a way similar to the proof of Proposition 4.3. \square

M Proof of Theorem 2.2

If \mathcal{C} is a collection of random variables, denote by $\sigma[\mathcal{C}]$ the σ -algebra generated by the elements of \mathcal{C} . Let (ξ_k) , $k \geq 1$, be a stationary sequence of real-valued random variables. For $k \geq 1$, define the σ -algebras $\mathcal{F}_k = \sigma[\xi_n : 1 \leq n \leq k]$ and $\mathcal{F}'_k = \sigma[\xi_n : n \geq k]$. Following (Billingsley 1999, p. 203), for $n \geq 1$, let

$$\varphi_n = \sup_{k \geq 1} \sup \{ \mathbb{P}(B'|B) - \mathbb{P}(B') : B \in \mathcal{F}_k, \mathbb{P}(B) > 0, B' \in \mathcal{F}'_{k+n} \}.$$

The sequence (φ_n) is used to study the mixing proprieties of the sequence (ξ_k) . Theorem M.1 below establishes a functional central limit theorem on the partial sums of (ξ_k) under certain conditions on (φ_n) . Theorem M.1 follows immediately from Theorem 19.2 and the discussion on p. 203 of (Billingsley 1999). Let D_∞ denote the set of real-valued functions on the interval $[0, \infty)$ that are right-continuous and have left-hand limits, endowed with the Skorokhod topology (see (Billingsley 1999, Section 16)).

Theorem M.1 (Billingsley (1999)). *Assume that (ξ_k) , $k \geq 1$, is stationary, that ξ_1 is square-integrable, and that*

$$\sum_{n=0}^{\infty} \sqrt{\varphi_n} < \infty. \tag{M.1}$$

Then the series

$$\bar{\sigma}^2 = \text{Var}(\xi_1) + 2 \sum_{k=2}^{\infty} \text{Cov}(\xi_1, \xi_k) \tag{M.2}$$

converges absolutely. If $\bar{\sigma} > 0$ and $S_n = \sum_{k=1}^n (\xi_k - E(\xi_k))$ then, as $n \rightarrow \infty$,

$$\frac{S_{\lfloor ns \rfloor}}{\sqrt{n}} \Rightarrow \bar{\sigma} B_s \tag{M.3}$$

in the sense of D_∞ , where B_s is a standard Brownian motion.

We show that the conditions of Theorem M.1 hold when $\xi_k = f(V^{(k)})$ for $k \geq 1$. As $(V^{(k)})$, $k \geq 1$, is a stationary Markov chain, the sequence $(f(V^{(k)}))$, $k \geq 1$, is stationary. We first prove the following lemma, which follows intuitively from the fact that if all components of the copy of U are redrawn at a certain iteration, future copies of U are independent of past ones.

Lemma M.1. *For positive integers n and k , if $B \in \mathcal{F}_k$ and $B' \in \mathcal{F}'_{k+n}$, then B and $B' \cap I$ are independent, where*

$$I = I_{k,n} = \{\omega \in \Omega : \exists j \in [k, k+n-1], N_j(\omega) = d\}.$$

Proof. Let $\mathcal{G}_k = \sigma[N_j, U^{(j+1)} : j \geq k]$. Fix $l \geq k+n$. By construction, if $\omega \in I$, there is $j \in [k, k+n-1]$ is such that $V^{(j+1)}(\omega) = U^{(j+1)}(\omega)$. Hence, if $\omega \in I$, the vector $V^{(l)}$ does not depend $V^{(k)}$. More precisely, it can be shown by induction that there is a measurable function $H : \mathbb{F}^{d(l-k)} \times \mathbb{R}^{(l-k)} \rightarrow \mathbb{F}^d$ such that, for $\omega \in I$,

$$V^{(l)}(\omega) = H(U^{(k+1)}(\omega), \dots, U^{(l)}(\omega), N_k(\omega), \dots, N_{l-1}(\omega)).$$

For $\omega \in \Omega$, let

$$G(\omega) = f(H(U^{(k+1)}(\omega), \dots, U^{(l)}(\omega), N_k(\omega), \dots, N_{l-1}(\omega))).$$

Thus, $f(V^{(l)}(\omega)) = G(\omega)$ for $\omega \in I$. As the maps f and H are measurable, and since the random variables $U^{(k+1)}, \dots, U^{(l)}$ and N_k, \dots, N_{l-1} are measurable with respect to \mathcal{G}_k , the random variable G is measurable with respect to \mathcal{G}_k as well. For $z \in \mathbb{R}$, let $B_z = \{\omega \in \Omega : f(V^{(l)}(\omega)) < z\}$. Since $I \in \mathcal{G}_k$ and

$$B_z \cap I = \{\omega \in \Omega : G(\omega) < z\} \cap I,$$

we have $B_z \cap I \in \mathcal{G}_k$. Thus $B_z \in \mathcal{G}'_k$ for any real number z , where

$$\mathcal{G}'_k = \{B \in \mathcal{F} : B \cap I \in \mathcal{G}_k\}.$$

It is easy to check that \mathcal{G}'_k is a σ -algebra. We conclude that $f(V^{(l)})$ is measurable with respect to \mathcal{G}'_k , for any $l \geq k+n$, and so $\mathcal{F}'_{k+n} \subseteq \mathcal{G}'_k$. Consequently, if $B \in \mathcal{F}_k$ and $B' \in \mathcal{F}'_{k+n}$, then $B' \in \mathcal{G}'_k$, and so $B' \cap I \in \mathcal{G}_k$. Since \mathcal{F}_k and \mathcal{G}_k are independent, so are B and $B' \cap I$. \square

We now prove the theorem. For positive integers n and k , let $B \in \mathcal{F}_k$ with $\mathbb{P}(B) > 0$ and $B' \in \mathcal{F}'_{k+n}$, and define I as in Lemma M.1. Thus,

$$\mathbb{P}(B' \cap I | B) = \mathbb{P}(B' \cap I).$$

Since

$$\mathbb{P}(B') \leq \mathbb{P}(B' \cap I) + \mathbb{P}(I^c),$$

it follows that

$$\mathbb{P}(B') \leq \mathbb{P}(B' | B) + \mathbb{P}(I^c).$$

Replacing B' with its complement and noting that $\mathbb{P}(I^c) = (1 - q_{d-1})^n$, we conclude that

$$|\mathbb{P}(B' | B) - \mathbb{P}(B')| \leq (1 - q_{d-1})^n.$$

Hence $\varphi_n \leq (1 - q_{d-1})^n$, and so (M.1) holds. Thus the conclusions of Theorem M.1 hold for the sequence $(f(V^{(k)}))$, $k \geq 1$. It follows from (B.3) and (M.2) that $\bar{\sigma} = \sigma$. By (M.3), as $n \rightarrow \infty$,

$$\frac{[ns](f_{[ns]} - E(f(U)))}{\sqrt{n}} \Rightarrow \bar{\sigma} B_s \tag{M.4}$$

in D_∞ . Setting $s = 1$, which can be justified by applying the continuous mapping theorem (Billingsley 1999, Theorem 2.7) with the projection map (Billingsley 1999, Theorem 16.6), implies (2.8). Let τ_k denote the random amount time to generate $V^{(k)}$ and calculate $f(V^{(k)})$. Thus $E(\tau_k) = T$. By the strong law of large numbers, with probability 1,

$$\frac{1}{n} \sum_{k=1}^n \tau_k \rightarrow T \quad (\text{M.5})$$

as n goes to infinity. By (M.4), (M.5) and (Glynn and Whitt 1992, Theorem 1), it follows that

$$\sqrt{c}(f_{\tilde{N}(c)} - E(f(U))) \Rightarrow N(0, T\sigma^2),$$

as $c \rightarrow \infty$. Since $T\sigma^2 = R(q; t, \nu^*)$, this implies (2.9).

N Proof of Theorem 5.1 and of Proposition 5.1

For convenience, set $N_0 = \bar{N}_0 = d$. For $0 \leq i \leq d-1$, let

$$\mathcal{S}_i = \mathcal{S}_i(n) = \{k \in [0, n-1] : N_k > i\}. \quad (\text{N.1})$$

Let $u_i(0) = 0, u_i(1), \dots, u_i(|\mathcal{S}_i| - 1)$ be the elements of \mathcal{S}_i sorted in increasing order. Set $u_i(|\mathcal{S}_i|) = n$ and let $Q(i) = Q(i, n) = \sum_{l=1}^{|\mathcal{S}_i|} (u_i(l) - u_i(l-1))^2$. Note that $\mathcal{S}_0 = \{0, \dots, n-1\}$ and $Q(0) = n$. The following lemma gives an alternative characterization of $Q(i)$.

Lemma N.1. For $0 \leq i \leq d-1$,

$$\sum_{1 \leq m < k \leq n} \mathbf{1}\{\max_{m \leq j < k} N_j \leq i\} = \frac{1}{2}(Q(i) - n).$$

Proof. Given integers m and k with $1 \leq m < k \leq n$, the condition $\max_{m \leq j < k} N_j \leq i$ holds if and only if $[m, k) \cap \mathcal{S}_i = \emptyset$. It is thus equivalent to the existence of an integer $l \in [1, |\mathcal{S}_i|]$ such that $u_i(l-1) < m < k \leq u_i(l)$. Hence

$$\sum_{1 \leq m < k \leq n} \mathbf{1}\{\max_{m \leq j < k} N_j \leq i\} = \frac{1}{2} \sum_{l=1}^{|\mathcal{S}_i|} (u_i(l) - u_i(l-1))(u_i(l) - u_i(l-1) - 1).$$

We conclude the proof by noting that

$$\sum_{l=1}^{|\mathcal{S}_i|} (u_i(l) - u_i(l-1)) = n.$$

□

Lemma N.2 below relates the variance of f_n to the $Q(i)$'s and $C(i)$'s.

Lemma N.2. For $n \geq 1$,

$$\text{Var}(f_n) = n^{-2} \sum_{i=0}^{d-1} Q(i)(C(i) - C(i+1)).$$

Proof. As in the proof of Theorem 2.1, we assume without loss of generality that $E(f(U)) = 0$. Let $m < k$ be two integers in $[1, n]$. By arguments similar to those that lead to (B.2),

$$E(f(V^{(m)})f(V^{(k)})) = C(\max_{m \leq j < k} N_j).$$

Since, for $0 \leq l \leq d$,

$$\begin{aligned} C(l) &= \sum_{i=l}^{d-1} (C(i) - C(i+1)) \\ &= \sum_{i=0}^{d-1} \mathbf{1}\{l \leq i\} (C(i) - C(i+1)), \end{aligned}$$

it follows that

$$E(f(V^{(m)})f(V^{(k)})) = \sum_{i=0}^{d-1} \mathbf{1}\{\max_{m \leq j < k} N_j \leq i\} (C(i) - C(i+1)).$$

As

$$\text{Var}(f_n) = n^{-2} \left(\sum_{m=1}^n E((f(V^{(m)}))^2) + 2 \sum_{1 \leq m < k \leq n} E(f(V^{(m)})f(V^{(k)})) \right),$$

and $E((f(V^{(m)}))^2) = C(0)$ for $1 \leq m \leq n$, we conclude that

$$\text{Var}(f_n) = n^{-2} (nC(0) + 2 \sum_{i=0}^{d-1} \sum_{1 \leq m < k \leq n} \mathbf{1}\{\max_{m \leq j < k} N_j \leq i\} (C(i) - C(i+1))).$$

Thus, by Lemma N.1,

$$\text{Var}(f_n) = n^{-2} (nC(0) + \sum_{i=0}^{d-1} (Q(i) - n)(C(i) - C(i+1))),$$

which, after some simplifications, completes the proof. \square

We now prove Theorem 5.1. By the Cauchy-Schwartz inequality, for $0 \leq i \leq d-1$ and any sequence (x_l) , $1 \leq l \leq \mathcal{S}_i$,

$$\left(\sum_{l=1}^{|\mathcal{S}_i|} x_l \right)^2 \leq |\mathcal{S}_i| \left(\sum_{l=1}^{|\mathcal{S}_i|} x_l^2 \right).$$

Replacing x_l with $u_i(l) - u_i(l-1)$ shows that $n^2 \leq |\mathcal{S}_i|Q(i)$. By Lemma N.2, it follows that

$$\text{Var}(f_n) \geq \sum_{i=0}^{d-1} \frac{C(i) - C(i+1)}{|\mathcal{S}_i|}.$$

For $0 \leq k \leq n-1$, the expected running time of iteration $k+1$ is t_i if $N_k = i$. Furthermore, for $i \in [1, d]$, there are $|\mathcal{S}_{i-1}| - |\mathcal{S}_i|$ integers k in $[0, n-1]$ with $N_k = i$, where $\mathcal{S}_d \triangleq \emptyset$, and so

$$T_n = \sum_{i=1}^d (|\mathcal{S}_{i-1}| - |\mathcal{S}_i|) t_i = \sum_{i=0}^{d-1} |\mathcal{S}_i| (t_{i+1} - t_i). \quad (\text{N.2})$$

Thus,

$$T_n \text{Var}(f_n) \geq R(q; t, C),$$

where $q_i = |\mathcal{S}_i|/n$ for $0 \leq i \leq d-1$. This implies (5.3).

Assume now that $N_k = \bar{N}_k$ for $k \geq 1$. Then, by (5.2), for $0 \leq i \leq d-1$,

$$\mathcal{S}_i = \{k \in [0, n-1] : k \text{ is a multiple of } \mu_i\}. \quad (\text{N.3})$$

Thus $|\mathcal{S}_i| = 1 + \lfloor (n-1)\bar{q}_i \rfloor$ which, by (N.2), implies (5.4). Since $u_i(l) = l\mu_i$ for $0 \leq l \leq |\mathcal{S}_i| - 1$,

$$\begin{aligned} Q(i) &= \mu_i \sum_{l=1}^{|\mathcal{S}_i|-1} (u_i(l) - u_i(l-1)) + (n - \max(\mathcal{S}_i))^2 \\ &= \mu_i \max \mathcal{S}_i + (n - \max(\mathcal{S}_i))^2. \end{aligned}$$

As $0 \leq n - \max(\mathcal{S}_i) \leq \mu_i$, it follows that $\mu_i(n - \mu_i) \leq Q(i) \leq \mu_i n$. By Lemma N.2, this implies (5.5) and shows that, as n goes to infinity, the LHS of (5.5) converges to its RHS. \square

We now prove Proposition 5.1. Assume that (N_k) satisfies the conditions in the proposition, and that $nq_{d-1} > 1$. Then \mathcal{S}_{d-1} consists of the integers belonging to $\{0\} \cup (n - nq_{d-1}, n)$, and so $Q(d-1) \geq n^2(1 - q_{d-1})^2$. By Lemma N.2, it follows that $\text{Var}(f_n) \geq (1 - q_{d-1})^2 C(d-1)$, which completes the proof. \square

O Proof of Theorem 6.1

We first prove the following lemma.

Lemma O.1. *Let (ν_i) , $0 \leq i \leq d$, be a decreasing sequence such that $\nu_{m_l} \leq \text{Var}(\phi_L - \phi_l)$ for $0 \leq l \leq L$, with $\nu_d = 0$. Then*

$$\sum_{i=0}^{d-1} \sqrt{\frac{\nu_i}{i+1}} \leq 2 \sum_{l=1}^L \sqrt{m_l V_l}.$$

Proof. Since the sequence (ν_i) is decreasing and $(i+1)^{-1/2} \leq 2(\sqrt{i+1} - \sqrt{i})$ for $i \geq 0$,

$$\sum_{i=j}^{k-1} \sqrt{\frac{\nu_i}{i+1}} \leq 2(\sqrt{k} - \sqrt{j})\sqrt{\nu_j},$$

for $0 \leq j \leq k \leq d$. Hence,

$$\begin{aligned} \sum_{i=0}^{d-1} \sqrt{\frac{\nu_i}{i+1}} &= \sum_{l=0}^{L-1} \sum_{i=m_l}^{m_{l+1}-1} \sqrt{\frac{\nu_i}{i+1}} \\ &\leq 2 \sum_{l=0}^{L-1} (\sqrt{m_{l+1}} - \sqrt{m_l}) \sqrt{\nu_{m_l}} \\ &\leq 2 \sum_{l=0}^{L-1} (\sqrt{m_{l+1}} - \sqrt{m_l}) \text{Std}(\phi_L - \phi_l) \\ &= 2 \sum_{l=1}^L \sqrt{m_l} (\text{Std}(\phi_L - \phi_{l-1}) - \text{Std}(\phi_L - \phi_l)) \\ &\leq 2 \sum_{l=1}^L \sqrt{m_l V_l}, \end{aligned}$$

where the last equation follows by sub-linearity of the standard deviation. \square

We now prove Theorem 6.1. Since ϕ_l is square-integrable and is a measurable function of U_1, \dots, U_{m_l} , by Proposition 2.2, $C(m_l) \leq \text{Var}(\phi_L - \phi_l)$ for $0 \leq l \leq L$. By Proposition 2.1, the sequence $(C(i))$, $0 \leq i \leq d$, is decreasing, and so it satisfies the conditions of Lemma O.1. Thus,

$$\sum_{i=0}^{d-1} \sqrt{\frac{C(i)}{i+1}} \leq 2 \sum_{l=1}^L \sqrt{m_l V_l}.$$

Furthermore, by Proposition 2.5,

$$R(q; t, \nu^*) \leq 8c \left(\sum_{i=0}^{d-1} (\sqrt{i+1} - \sqrt{i}) \sqrt{C(i)} \right)^2.$$

Since $\sqrt{i+1} - \sqrt{i} \leq (i+1)^{-1/2}$, it follows that

$$R(q; t, \nu^*) \leq 32c \left(\sum_{l=1}^L \sqrt{m_l V_l} \right)^2.$$

Using (6.2) concludes the proof. \square

P Relation with splitting and conditional Monte Carlo

Like the splitting algorithm described in (Asmussen and Glynn 2007, Section V.5) when $d = 2$, our method samples more often important random variables. This section explores further the relation between our method and the splitting and conditional Monte Carlo methods. Fix $n \geq 1$ and assume for simplicity that the sequence (N_k) is deterministic. We first analyse the relation between our method and the conditional Monte Carlo method in the general case, then show that the generic dimension reduction algorithm for Markov chains estimation can be efficiently cast as a splitting algorithm.

P.1 Relation with conditional Monte Carlo

Define \mathcal{S}_i via (N.1) for $0 \leq i \leq d-1$, and set $\mathcal{S}_d = \{0\}$. For $0 \leq i \leq d$ and $m \in \mathcal{S}_i$, let

$$\mathcal{S}_{i,m} = \{k \in [m, n-1] : (m, k] \cap \mathcal{S}_i = \emptyset\}, \quad (\text{P.1})$$

and

$$f_{i,m} = \frac{1}{|\mathcal{S}_{i,m}|} \sum_{k \in \mathcal{S}_{i,m}} f(V^{(k+1)}).$$

Thus, if (m, m') are consecutive elements of $\mathcal{S}_i \cup \{n\}$, then $\mathcal{S}_{i,m} = \{m, \dots, m'-1\}$. In particular, for $0 \leq m \leq n-1$, we have $\mathcal{S}_{0,m} = \{m\}$, and so $f_{m,0} = f(V^{(m+1)})$. Similarly, $\mathcal{S}_{d,0} = \{0, \dots, n-1\}$ and $f_{d,0} = f_n$. For $0 \leq i \leq d$, $m \in \mathcal{S}_i$, and $k \in \mathcal{S}_{i,m}$, we have $N_l \leq i$ for $m < l \leq k$, and so the last $d-i$ components of $V^{(k+1)}$ and $V^{(m+1)}$ are the same. In other words,

$$V_j^{(k+1)} = V_j^{(m+1)}, \quad \text{for } 0 \leq i < j \leq d \text{ and } k \in \mathcal{S}_{i,m}. \quad (\text{P.2})$$

We can thus view $f_{i,m}$ as a discrete analog to $E(f(V^{(m+1)}) | V_{i+1}^{(m+1)}, \dots, V_d^{(m+1)})$. The following proposition shows that the random variables $f_{i,m}$, for $0 \leq i \leq d$ and $m \in \mathcal{S}_i$, can be calculated inductively via (P.3). This is reminiscent of the conditional Monte Carlo method, where an expectation is estimated via an average of conditional expectations.

Proposition P.1. For $1 \leq i \leq d$ and $m \in \mathcal{S}_i$,

$$f_{i,m} = \sum_{k \in \mathcal{S}_{i-1} \cap \mathcal{S}_{i,m}} \frac{|\mathcal{S}_{i-1,k}|}{|\mathcal{S}_{i,m}|} f_{i-1,k}, \quad (\text{P.3})$$

and

$$|\mathcal{S}_{i,m}| = \sum_{k \in \mathcal{S}_{i-1} \cap \mathcal{S}_{i,m}} |\mathcal{S}_{i-1,k}|. \quad (\text{P.4})$$

Proof. We first show that

$$\mathcal{S}_{i,m} = \bigcup_{k \in \mathcal{S}_{i-1} \cap \mathcal{S}_{i,m}} \mathcal{S}_{i-1,k}. \quad (\text{P.5})$$

If $k \in \mathcal{S}_{i-1} \cap \mathcal{S}_{i,m}$ and $k' \in \mathcal{S}_{i-1,k}$, we have $(m, k] \cap \mathcal{S}_i = \emptyset$ and $(k, k'] \cap \mathcal{S}_{i-1} = \emptyset$. As $\mathcal{S}_i \subseteq \mathcal{S}_{i-1}$ and $m \leq k \leq k'$, this implies that $(m, k'] \cap \mathcal{S}_i = \emptyset$, and so $k' \in \mathcal{S}_{i,m}$. Thus

$$\bigcup_{k \in \mathcal{S}_{i-1} \cap \mathcal{S}_{i,m}} \mathcal{S}_{i-1,k} \subseteq \mathcal{S}_{i,m}.$$

Conversely, given $k' \in \mathcal{S}_{i,m}$, let

$$k = \max([0, k'] \cap \mathcal{S}_{i-1}). \quad (\text{P.6})$$

Since $m \in [0, k'] \cap \mathcal{S}_{i-1}$, the integer k is well-defined and $m \leq k$. Since $k \leq k'$ and $(m, k'] \cap \mathcal{S}_i = \emptyset$, we have $(m, k] \cap \mathcal{S}_i = \emptyset$. Hence $k \in \mathcal{S}_{i,m}$, and so $k \in \mathcal{S}_{i-1} \cap \mathcal{S}_{i,m}$. Furthermore, $(k, k'] \cap \mathcal{S}_{i-1} = \emptyset$ by (P.6). Thus $k' \in \mathcal{S}_{i-1,k}$, and so

$$\mathcal{S}_{i,m} \subseteq \bigcup_{k \in \mathcal{S}_{i-1} \cap \mathcal{S}_{i,m}} \mathcal{S}_{i-1,k}.$$

This implies (P.5). Moreover, if $k \in \mathcal{S}_{i-1}$ and $j \in \mathcal{S}_{i-1,k}$, then $(k, j] \cap \mathcal{S}_{i-1} = \emptyset$, and so $k = \max([0, j] \cap \mathcal{S}_{i-1})$. Thus, if k and k' are distinct elements of \mathcal{S}_{i-1} , the sets $\mathcal{S}_{i-1,k}$ and $\mathcal{S}_{i-1,k'}$ are disjoint. Together with (P.5), this immediately implies (P.3) and (P.4). \square

P.2 The Markov chains case

Using the same notation and assumptions as in §4, we show how to cast the generic dimension reduction algorithm for Markov chains estimation as a splitting algorithm. For each integer k in $[1, n]$, define the Markov chain $(X_i^{(k)})$, $0 \leq i \leq d$, by induction on i as follows: $X_0^{(k)} = X_0$ and $X_{i+1}^{(k)} = g_i(X_i^{(k)}, V_{d-i}^{(k)})$ for $0 \leq i \leq d-1$. Then it can be shown by induction that $X_d^{(k)} = G_d(V^{(k)}, X_0)$, and that $g(X_d^{(k)}) = f(V^{(k)})$. The generic dimension reduction algorithm for Markov chains estimation described in §4 thus outputs the average of $g(X_d^{(1)}), \dots, g(X_d^{(n)})$. Furthermore, it can be shown by induction that, for $0 \leq i \leq d$ and $1 \leq k \leq n$, the random variable $X_i^{(k)}$ is a deterministic function of $(V_j^{(k)})$, $d-i < j \leq d$. On the other hand, by (P.1), for $0 \leq i \leq d$ and $0 \leq k \leq n-1$, if $m = \max([0, k] \cap \mathcal{S}_{d-i})$ then $k \in \mathcal{S}_{d-i,m}$. The integer m is well-defined since $0 \in \mathcal{S}_{d-i}$. By (P.2), the last i components of $V^{(k+1)}$ and $V^{(m+1)}$ are the same, and so $X_i^{(k+1)} = X_i^{(m+1)}$. However, for $0 \leq i \leq d-1$ and $k \in \mathcal{S}_{d-i-1}$, we have $V_{d-i}^{(k+1)} = U_{d-i}^{(k+1)}$ since $N_k \geq d-i$, and so

$$X_{i+1}^{(k+1)} = g_i(X_i^{(m+1)}, U_{d-i}^{(k+1)}), \text{ with } m = \max([0, k] \cap \mathcal{S}_{d-i}). \quad (\text{P.7})$$

We can calculate $X_i^{(k+1)}$ by induction on i via (P.7) for $0 \leq i \leq d$ and $k \in \mathcal{S}_{d-i}$. Recalling that $\mathcal{S}_0 = \{0, \dots, n-1\}$, this allows us to simulate $X_d^{(1)}, \dots, X_d^{(n)}$. The generic dimension reduction algorithm for Markov chains estimation described in §4 can thus be rewritten as follows:

Table 7: Comparison with QMC in $E(X_d)$ estimation in $G_t/D/1$ queue, 1000 samples, where X_d is the number of customers in the queue at time-step d .

		n	90% confidence interval	Std	Cost	Cost \times Std ²	VRF
$d = 2500$	RDR	649	$5.525 \pm 8.9 \times 10^{-3}$	1.7×10^{-1}	$2.736 \times 10^4 \pm 1.5 \times 10^2$	8.0×10^2	48
	DDR	782	$5.526 \pm 8.6 \times 10^{-3}$	1.7×10^{-1}	2.608×10^4	7.2×10^2	53
	QMC	4096	$5.5235 \pm 1.0 \times 10^{-3}$	2.0×10^{-2}	1.024×10^7	3.9×10^3	10
$d = 5000$	RDR	1225	$5.521 \pm 6.5 \times 10^{-3}$	1.2×10^{-1}	$5.462 \times 10^4 \pm 3.2 \times 10^2$	8.4×10^2	92
	DDR	1400	$5.518 \pm 6.2 \times 10^{-3}$	1.2×10^{-1}	5.174×10^4	7.3×10^2	105
	QMC	4096	$5.5225 \pm 8.2 \times 10^{-4}$	1.6×10^{-2}	2.048×10^7	5.1×10^3	15
$d = 10000$	RDR	2295	$5.5243 \pm 4.6 \times 10^{-3}$	8.8×10^{-2}	$1.106 \times 10^5 \pm 6.3 \times 10^2$	8.5×10^2	182
	DDR	2822	$5.5221 \pm 4.5 \times 10^{-3}$	8.6×10^{-2}	1.032×10^5	7.7×10^2	202
	QMC	4096	$5.5234 \pm 1.1 \times 10^{-3}$	2.0×10^{-2}	4.096×10^7	1.7×10^4	9

1. Generate N_k for $1 \leq k \leq n - 1$ and calculate the sets \mathcal{S}_i , $0 \leq i \leq d$.
2. Set $X_0^{(1)} = X_0$.
3. For $i = 0, \dots, d - 1$ and $k \in \mathcal{S}_{d-i-1}$, sample a copy $U_{d-i}^{(k+1)}$ of Y_i and calculate $X_{i+1}^{(k+1)}$ via (P.7).
4. Output the average of $g(X_d^{(1)}), \dots, g(X_d^{(n)})$.

Step 3 generates one or several copies of X_{i+1} for each copy of X_i , and so the above algorithm may be viewed as a splitting algorithm. Assume now that $N_k = \bar{N}_k$ for $k \geq 1$. Step 1 can then be implemented via (N.3). In (P.7), $m = 0$ if $i = 0$ and, by (N.3), $m = \lfloor k/\mu_{d-i} \rfloor \mu_{d-i}$ if $1 \leq i \leq d - 1$. The expected running times of this algorithm and of the algorithm in §4 are within a constant from each other as they are both proportional to the number of sampled copies of the U_i 's.

Q Further numerical experiments

Q.1 Comparison with Quasi-Monte Carlo

Consider the $G_t/D/1$ queue with the same parameters as in §7.2. Table 7 estimates $E(X_d)$, and Table 8 gives VRFs in the estimation of $\mathbb{P}(X_d > z)$, for selected values of z . Our numerical results for the RDR and DDR algorithms are similar to those of §7.2. Once again, for the RDR and DDR algorithms, the variable Cost \times Std² is roughly independent of d , and the variance reduction factors are roughly proportional to d . In contrast, for the QMC algorithm, the variable Cost \times Std² is roughly proportional to d , and the variance reduction factors are roughly constant. The VRFs of the RDR and DDR algorithms in Table 8 are, in general, greater than or equal to the corresponding VRFs in Table 7, which confirms the resiliency of these algorithms to discontinuities of g . In contrast, the VRFs of the QMC algorithm in Table 8 are lower than the corresponding VRFs in Table 7. The RDR and DDR algorithms outperform the QMC algorithm. The VRFs of the QMC algorithm in Tables 7 and 8 are of the same order of magnitude as those obtained by L'Ecuyer and Lemieux (2000), who have reduced the variance in the simulation of a $M/M/1$ queue by a factor ranging between 5 and 10 via lattice rules.

Q.2 $G_t/D/1$ queue with time-varying amplitude

Consider a $G_t/D/1$ queue where A_i has a Poisson distribution with time-varying rate

$$\lambda_i = \left(1 - \frac{1}{\ln(i+2)}\right) \left(0.75 + 0.5 \cos\left(\frac{\pi i}{50}\right)\right),$$

Table 8: VRFs for $\mathbb{P}(X_d > z)$ estimation in $G_t/D/1$ queue.

z		0	2	4	6	8	10
$d = 2500$	RDR	83	60	59	49	50	51
	DDR	125	76	67	71	63	61
	QMC	1.8	2.2	2.5	2.7	2.1	1.7
$d = 5000$	RDR	149	120	109	101	101	107
	DDR	215	165	132	122	116	110
	QMC	1.1	1.8	2.8	2.4	2.2	1.8
$d = 10000$	RDR	280	227	206	208	188	179
	DDR	470	280	228	237	220	189
	QMC	1.2	2	2.5	2.2	2.3	1.9

Table 9: $E(X_d)$ estimation in $G_t/D/1$ queue with time-varying amplitude, 1000 samples, where X_d is the number of customers in the queue at time-step d .

		n	90% confidence interval	Std	Cost	Cost \times Std ²	VRF
$d = 10^4$	RDR	2.7×10^3	$3.693 \pm 3.5 \times 10^{-3}$	6.7×10^{-2}	$1.103 \times 10^5 \pm 6.1 \times 10^2$	4.9×10^2	2.1×10^2
	DDR	3.8×10^3	$3.6955 \pm 3.5 \times 10^{-3}$	6.6×10^{-2}	1.029×10^5	4.5×10^2	2.3×10^2
	MLMC	8.7×10^3	$3.696 \pm 4.1 \times 10^{-3}$	7.9×10^{-2}	1.084×10^5	6.8×10^2	1.5×10^2
$d = 10^5$	RDR	2.6×10^4	$4.0264 \pm 1.2 \times 10^{-3}$	2.3×10^{-2}	$1.099 \times 10^6 \pm 5.4 \times 10^3$	5.8×10^2	1.9×10^3
	DDR	2.9×10^4	$4.0255 \pm 1.2 \times 10^{-3}$	2.3×10^{-2}	1.040×10^6	5.7×10^2	1.9×10^3
	MLMC	7.3×10^4	$4.0271 \pm 1.4 \times 10^{-3}$	2.8×10^{-2}	1.103×10^6	8.5×10^2	1.3×10^3
$d = 10^6$	RDR	2.5×10^5	$4.2573 \pm 3.7 \times 10^{-4}$	7.2×10^{-3}	$1.098 \times 10^7 \pm 4.9 \times 10^4$	5.6×10^2	2.1×10^4
	DDR	2.7×10^5	$4.25687 \pm 3.8 \times 10^{-4}$	7.3×10^{-3}	1.051×10^7	5.6×10^2	2.1×10^4
	MLMC	7.3×10^5	$4.2566 \pm 4.6 \times 10^{-4}$	8.8×10^{-3}	1.127×10^7	8.7×10^2	1.3×10^4

for $1 \leq i \leq d$ (recall that $A_0 = 0$). Thus, up to a time-varying multiplicative factor, λ_i has the same expression as in §7.2. Table 9 estimates $E(X_d)$, and Table 10 gives VRFs in the estimation of $\mathbb{P}(X_d > z)$, for selected values of z . Our numerical results are similar to those of §7.2, except that there is a big difference in the means at different times. Once again, for the RDR, DDR, and MLMC algorithms, the variable Cost \times Std² is roughly independent of d , and the variance reduction factors are roughly proportional to d . The VRFs of the RDR and DDR algorithms in Table 10 are, in most cases, greater than or equal to the corresponding VRFs in Table 9, which confirms the resiliency of these algorithms to discontinuities of g . In contrast, the VRFs of the MLMC algorithm in Table 10 are lower than the corresponding VRFs in Table 9. The RDR algorithm outperforms the MLMC algorithm by a factor ranging from 1 to 2 in Table 9, and a factor ranging from 2 to 14 in Table 10.

Q.3 A multi-frequency $G_t/D/1$ queue

Consider a $G_t/D/1$ queue where A_i has a Poisson distribution with time-varying rate

$$\lambda_i = 0.75 + 0.2 \cos\left(\frac{\pi i}{50}\right) + 0.1 \cos\left(\frac{\pi i}{5000}\right) + 0.05 \cos\left(\frac{\pi i}{500000}\right),$$

for $1 \leq i \leq d$, with $A_0 = 0$. Table 11 estimates $E(X_d)$. Once again, for the RDR, DDR, and MLMC algorithms, the variable Cost \times Std² is roughly independent of d , and the variance

Table 10: VRFs for $\mathbb{P}(X_d > z)$ estimation in $G_t/D/1$ queue with time-varying amplitude.

z		0	2	4	6	8	10
$d = 10^4$	RDR	3.1×10^2	2.4×10^2	2.2×10^2	2.1×10^2	2.3×10^2	2.2×10^2
	DDR	4.4×10^2	2.7×10^2	2.4×10^2	2.6×10^2	3.0×10^2	2.8×10^2
	MLMC	3.8×10^1	6.1×10^1	7.4×10^1	8.7×10^1	7.8×10^1	7.9×10^1
$d = 10^5$	RDR	3.3×10^3	2.3×10^3	2.0×10^3	2.0×10^3	1.9×10^3	2.1×10^3
	DDR	3.9×10^3	2.8×10^3	2.6×10^3	2.4×10^3	2.0×10^3	2.4×10^3
	MLMC	3.3×10^2	5.3×10^2	6.3×10^2	7.5×10^2	6.8×10^2	7.6×10^2
$d = 10^6$	RDR	3.5×10^4	2.3×10^4	2.0×10^4	1.9×10^4	1.8×10^4	1.8×10^4
	DDR	4.5×10^4	3.0×10^4	2.6×10^4	2.3×10^4	2.3×10^4	2.4×10^4
	MLMC	3.2×10^3	4.7×10^3	6.2×10^3	7.8×10^3	7.3×10^3	6.5×10^3

Table 11: $E(X_d)$ estimation in a multi-frequency $G_t/D/1$ queue, 1000 samples, where X_d is the number of customers in the queue at time-step d .

		n	90% confidence interval	Std	Cost	Cost \times Std ²	VRF
$d = 10^4$	RDR	4.2×10^2	$5.6 \pm 1.6 \times 10^{-2}$	3.0×10^{-1}	$1.104 \times 10^5 \pm 6.6 \times 10^2$	1.0×10^4	2.5×10^1
	DDR	9.8×10^2	$5.616 \pm 1.4 \times 10^{-2}$	2.6×10^{-1}	1.028×10^5	7.2×10^3	3.4×10^1
	MLMC	2.1×10^3	$5.607 \pm 2.0 \times 10^{-2}$	3.8×10^{-1}	1.065×10^5	1.6×10^4	1.6×10^1
$d = 10^5$	RDR	6.0×10^3	$5.1944 \pm 4.5 \times 10^{-3}$	8.7×10^{-2}	$1.101 \times 10^6 \pm 5.6 \times 10^3$	8.4×10^3	2.7×10^2
	DDR	7.8×10^3	$5.1961 \pm 4.0 \times 10^{-3}$	7.7×10^{-2}	1.030×10^6	6.1×10^3	3.6×10^2
	MLMC	2.0×10^4	$5.187 \pm 5.9 \times 10^{-3}$	1.1×10^{-1}	1.108×10^6	1.4×10^4	1.6×10^2
$d = 10^6$	RDR	4.8×10^4	$5.6118 \pm 1.7 \times 10^{-3}$	3.3×10^{-2}	$1.102 \times 10^7 \pm 5.0 \times 10^4$	1.2×10^4	2.1×10^3
	DDR	1.0×10^5	$5.6111 \pm 1.5 \times 10^{-3}$	3.0×10^{-2}	1.042×10^7	9.1×10^3	2.8×10^3
	MLMC	1.8×10^5	$5.6113 \pm 2.2 \times 10^{-3}$	4.2×10^{-2}	1.124×10^7	2.0×10^4	1.3×10^3

Table 12: $\mathbb{P}(X_d > z)$ estimation in GARCH model, with $z = 4.4 \times 10^{-5}$, where X_d is the daily variance at time-step d .

		n	90% confidence interval	Std	Cost	Cost \times Std ²	Bias
$d = 1250$	RDR	277	$0.393471 \pm 7.3 \times 10^{-5}$	3.9×10^{-2}	$1.368 \times 10^4 \pm 3.2$	21	0
	DDR	596	$0.393483 \pm 6.2 \times 10^{-5}$	3.4×10^{-2}	1.355×10^4	16	0
	Long-run	—	$0.415353 \pm 4.4 \times 10^{-5}$	7.9×10^{-2}	1.250×10^3	7.7	0.0219
$d = 2500$	RDR	529	$0.39339 \pm 7.2 \times 10^{-5}$	2.8×10^{-2}	$2.742 \times 10^4 \pm 8.5$	21	0
	DDR	1167	$0.393489 \pm 6.3 \times 10^{-5}$	2.4×10^{-2}	2.734×10^4	16	0
	Long-run	—	$0.404431 \pm 4.4 \times 10^{-5}$	5.6×10^{-2}	2.500×10^3	7.9	0.0109
$d = 5000$	RDR	970	$0.39346 \pm 7.4 \times 10^{-5}$	2.0×10^{-2}	$5.494 \times 10^4 \pm 23$	22	0
	DDR	1899	$0.393506 \pm 6.5 \times 10^{-5}$	1.8×10^{-2}	5.207×10^4	16	0
	Long-run	—	$0.398968 \pm 4.4 \times 10^{-5}$	4.0×10^{-2}	5.000×10^3	7.9	0.0055

reduction factors are roughly proportional to d . The RDR algorithm outperforms the MLMC algorithm by about a factor of 1.5.

Q.4 Comparison with a long-run average estimator

In several examples, when d is large, the Markov chain (X_m) , $0 \leq m \leq d$, has some notion of stationarity, and so $E(g(X_d))$ can be estimated via a suitable long-run average. A drawback of such an estimator is that it is biased. We compare below a long-run average estimator with the RDR and DDR algorithms. In Tables 12 through 15, the RDR and DDR algorithms use $10^9/d$ samples, while the long-run algorithm uses $11 \times 10^9/d$ samples. Thus, for each of the three algorithms and each d , the total number of simulations of the U_i s throughout the independent samples is roughly 11×10^9 . The bias of the long-run average estimator is calculated by taking the difference with the DDR estimator.

For the GARCH volatility example of §7.1, a natural long-run average estimator for $\Pr(X_d > z)$ is

$$\frac{1}{d} \sum_{i=1}^d \mathbf{1}\{X_i > z\}.$$

Table 12 compares this estimator with the RDR and DDR estimators. The work-normalized variance of each of the three algorithms is roughly independent of d . The work-normalized variance of the long-run average estimator is smaller than those of the RDR and DDR algorithms by about a factor of 3 and 2, respectively. The bias of the long-run average estimator decreases as d increases, but is much larger than the standard deviations of the long-run average and DDR estimators.

For the $G_t/D/1$ queue example of §7.2 (resp. §Q.2), the arrival rate is periodic (resp. almost periodic) with period 100, and so a long-run average estimator for $E(X_d)$ is

$$\frac{1}{\lceil d/100 \rceil} \sum_{i=0}^{\lceil d/100 \rceil - 1} X_{d-i*100}.$$

Table 13: $E(X_d)$ estimation in $G_t/D/1$ queue, where X_d is the number of customers in the queue at time-step d , with $A_i \sim \text{Poisson}(0.75 + 0.5 \cos(\pi i/50))$ for $1 \leq i \leq d$.

		n	90% confidence interval	Std	Cost	Cost \times Std ²
$d = 10^4$	RDR	2.3×10^3	$5.52351 \pm 4.8 \times 10^{-4}$	9.2×10^{-2}	$1.100 \times 10^5 \pm 6.2 \times 10^1$	9.2×10^2
	DDR	2.8×10^3	$5.52333 \pm 4.4 \times 10^{-4}$	8.5×10^{-2}	1.032×10^5	7.4×10^2
	Long-run	—	$5.5238 \pm 6.2 \times 10^{-4}$	3.9×10^{-1}	1.000×10^4	1.5×10^3
$d = 10^5$	RDR	2.3×10^4	$5.52324 \pm 4.8 \times 10^{-4}$	2.9×10^{-2}	$1.100 \times 10^6 \pm 1.7 \times 10^3$	9.4×10^2
	DDR	3.6×10^4	$5.52339 \pm 4.7 \times 10^{-4}$	2.8×10^{-2}	1.046×10^6	8.4×10^2
	Long-run	—	$5.5238 \pm 6.2 \times 10^{-4}$	1.2×10^{-1}	1.000×10^5	1.5×10^3
$d = 10^6$	RDR	2.3×10^5	$5.52325 \pm 4.9 \times 10^{-4}$	9.4×10^{-3}	$1.103 \times 10^7 \pm 4.8 \times 10^4$	9.8×10^2
	DDR	2.6×10^5	$5.52363 \pm 4.5 \times 10^{-4}$	8.7×10^{-3}	1.047×10^7	7.9×10^2
	Long-run	—	$5.5238 \pm 6.1 \times 10^{-4}$	3.9×10^{-2}	1.000×10^6	1.5×10^3

Table 14: $E(X_d)$ estimation in $G_t/D/1$ queue, where X_d is the number of customers in the queue at time-step d , with $A_i \sim \text{Poisson}((1 - 1/\ln(i + 2))(0.75 + 0.5 \cos(\pi i/50)))$

		n	90% confidence interval	Std	Cost	Cost \times Std ²	Bias
$d = 10^4$	RDR	2.7×10^3	$3.69517 \pm 3.5 \times 10^{-4}$	6.7×10^{-2}	$1.100 \times 10^5 \pm 6.0 \times 10^1$	5.0×10^2	0
	DDR	3.8×10^3	$3.69535 \pm 3.4 \times 10^{-4}$	6.5×10^{-2}	1.029×10^5	4.4×10^2	0
	Long-run	—	$3.4862 \pm 4.8 \times 10^{-4}$	3.1×10^{-1}	1.000×10^4	9.5×10^2	-0.209
$d = 10^5$	RDR	2.6×10^4	$4.02661 \pm 3.7 \times 10^{-4}$	2.3×10^{-2}	$1.100 \times 10^6 \pm 1.7 \times 10^3$	5.7×10^2	0
	DDR	2.9×10^4	$4.02676 \pm 3.7 \times 10^{-4}$	2.2×10^{-2}	1.040×10^6	5.3×10^2	0
	Long-run	—	$3.8854 \pm 5.1 \times 10^{-4}$	1.0×10^{-1}	1.000×10^5	1.1×10^3	-0.141
$d = 10^6$	RDR	2.5×10^5	$4.2573 \pm 3.7 \times 10^{-4}$	7.2×10^{-3}	$1.098 \times 10^7 \pm 4.9 \times 10^4$	5.6×10^2	0
	DDR	2.7×10^5	$4.25687 \pm 3.8 \times 10^{-4}$	7.3×10^{-3}	1.051×10^7	5.6×10^2	0
	Long-run	—	$4.1582 \pm 5.3 \times 10^{-4}$	3.4×10^{-2}	1.000×10^6	1.1×10^3	-0.099

Tables 13 and 14 compares this estimator with the RDR and DDR estimators. The work-normalized variance of the long-run average estimator is larger than those of the RDR and DDR estimators by a factor ranging from 1.5 and 2.2. The bias of the long-run algorithm is not reported in Table 13 because it is not statistically significant. In Table 14, however, the bias of the long-run average estimator is much larger than the standard deviations of the long-run average and DDR estimators.

Consider now the $M_t/GI/1$ queue example of §7.3, and assume that θ is an integer. In our experiments, we have set $d = \theta$ and $X_i = W_i$ for $0 \leq i \leq d$, and so a long-run average estimator for $\mathbb{P}(W_\theta > 1)$ is

$$\frac{1}{\lceil d/100 \rceil} \sum_{i=0}^{\lceil d/100 \rceil - 1} \mathbf{1}\{X_{d-i*100} > 1\}.$$

Table 15 compares this estimator with the RDR and DDR estimators. The work-normalized variance of the long-run average estimator is larger than those of the RDR and DDR estimators by about a factor of 1.3 and 2, respectively. Here again, the bias of the long-run algorithm is not reported because it is not statistically significant.

In summary, for large d and Markov chains with periodic features, $E(g(X_d))$ can be estimated via a suitable biased long-run average estimator. The order of magnitude of the bias depends on

Table 15: $\mathbb{P}(W_\theta > 1)$ estimation in $M_t/GI/1$ queue, $\alpha = 2$, where W_θ is the residual work at time θ .

		n	90% confidence interval	Std	Cost	Cost \times Std ²
$d = 10^4$	RDR	3.9×10^3	$0.853775 \pm 4.9 \times 10^{-5}$	9.3×10^{-3}	$1.100 \times 10^5 \pm 6.1 \times 10^1$	9.6
	DDR	7.6×10^3	$0.853788 \pm 3.9 \times 10^{-5}$	7.6×10^{-3}	1.037×10^5	5.9
	Long-run	—	$0.853734 \pm 5.6 \times 10^{-5}$	3.6×10^{-2}	1.000×10^4	13
$d = 10^5$	RDR	3.0×10^4	$0.85385 \pm 5.0 \times 10^{-5}$	3.0×10^{-3}	$1.100 \times 10^6 \pm 1.7 \times 10^3$	10
	DDR	4.8×10^4	$0.853799 \pm 4.1 \times 10^{-5}$	2.5×10^{-3}	1.031×10^6	6.5
	Long-run	—	$0.853837 \pm 5.6 \times 10^{-5}$	1.1×10^{-2}	1.000×10^5	13
$d = 10^6$	RDR	2.5×10^5	$0.853762 \pm 5.1 \times 10^{-5}$	9.8×10^{-4}	$1.103 \times 10^7 \pm 4.9 \times 10^4$	11
	DDR	4.5×10^5	$0.853779 \pm 4.4 \times 10^{-5}$	8.4×10^{-4}	1.040×10^7	7.3
	Long-run	—	$0.853847 \pm 5.6 \times 10^{-5}$	3.6×10^{-3}	1.000×10^6	13

the application and on the value of d . The bias of the long-run average estimator is difficult to evaluate without using an alternative estimator, though. In the examples in this subsection, the work-normalized variances of the long-run average, RDR and DDR estimators have the same order of magnitude. In the $G_t/D/1$ queue example of §Q.3, however, the arrival rate is periodic with period 10^6 . This example does not seem to admit a suitable long-run average estimator for the values of d listed in Table 11.

References

- Acworth, P. A., Broadie, M. and Glasserman, P. (1998). A comparison of some Monte Carlo and quasi Monte Carlo techniques for option pricing, *in* H. Niederreiter, P. Hellekalek, G. Larcher and P. Zinterhof (eds), *Monte Carlo and Quasi-Monte Carlo Methods 1996*, Vol. 127 of *Lecture Notes in Statistics*, Springer New York, pp. 1–18.
- Åkesson, F. and Lehoczky, J. P. (2000). Path generation for quasi-Monte Carlo simulation of mortgage-backed securities, *Management Science* **46**(9): 1171–1187.
- Andrew, A. M. (1979). Another efficient algorithm for convex hulls in two dimensions, *Inf. Process. Lett.* **9**(5): 216–219.
- Asmussen, S. and Glynn, P. W. (2007). *Stochastic simulation: algorithms and analysis*, Vol. 57, Springer Science & Business Media.
- Billingsley, P. (1999). *Convergence of probability measures*, second edn, Wiley, New York.
- Blanchet, J., Leder, K. and Shi, Y. (2011). Analysis of a splitting estimator for rare event probabilities in jackson networks, *Stochastic Systems* **1**(2): 306–339.
- Botev, Z. I., L’Ecuyer, P., Rubino, G., Simard, R. and Tuffin, B. (2013). Static network reliability estimation via generalized splitting, *INFORMS Journal on Computing* **25**(1): 56–71.
- Caffisch, R. E. (1998). Monte Carlo and quasi-Monte Carlo methods, *Acta Numerica* **7**: 1–49.
- Caffisch, R. E., Morokoff, W. J. and Owen, A. B. (1997). Valuation of mortgage backed securities using Brownian bridges to reduce effective dimension, *Journal of Computational Finance* **1**: 27–46.
- Ermakov, S. and Melas, V. (1995). *Design and analysis of simulation experiments*, Vol. 339, Springer Science & Business Media, Dordrecht, The Netherlands.
- Feldman, Z., Mandelbaum, A., Massey, W. A. and Whitt, W. (2008). Staffing of time-varying queues to achieve time-stable performance, *Management Science* **54**(2): 324–338.
- Giles, M. B. (2008). Multilevel Monte Carlo path simulation, *Operations Research* **56**(3): 607–617.
- Glasserman, P. (2004). *Monte Carlo methods in financial engineering*, Vol. 53, Springer, New York.
- Glasserman, P., Heidelberger, P. and Shahabuddin, P. (1999). Asymptotically optimal importance sampling and stratification for pricing path-dependent options, *Mathematical finance* **9**(2): 117–152.
- Glasserman, P., Heidelberger, P., Shahabuddin, P. and Zajic, T. (1999). Multilevel splitting for estimating rare event probabilities, *Operations Research* **47**(4): 585–600.

- Glynn, P. W. and Rhee, C.-h. (2014). Exact estimation for Markov chain equilibrium expectations, *Journal of Applied Probability* **51**(A): 377–389.
- Glynn, P. W. and Whitt, W. (1992). The asymptotic efficiency of simulation estimators, *Operations research* **40**(3): 505–520.
- Hull, J. (2014). *Options, Futures and Other Derivatives*, ninth edn, Prentice-Hall, Incorporated.
- Imai, J. and Tan, K. S. (2006). A general dimension reduction technique for derivative pricing, *Journal of Computational Finance* **10**(2): 129.
- Jiang, G. and Fu, M. C. (2017). Importance splitting for finite-time rare event simulation, *IEEE Transactions on Automatic Control*.
- Kahalé, N. (2016). Optimized sampling for Monte Carlo simulations via dimension reduction, *9th NIPS Workshop on Optimization for Machine Learning*, Barcelona, Spain. <http://opt-ml.org/index.html>.
- L’Ecuyer, P., Lécot, C. and Tuffin, B. (2008). A randomized quasi-Monte Carlo simulation method for Markov chains, *Operations Research* **56**(4): 958–975.
- L’Ecuyer, P. and Lemieux, C. (2000). Variance reduction via lattice rules, *Management Science* **46**(9): 1214–1235.
- Liu, R. and Owen, A. B. (2006). Estimating mean dimensionality of analysis of variance decompositions, *Journal of the American Statistical Association* **101**(474): 712–721.
- Ma, N. and Whitt, W. (2017). A rare-event simulation algorithm for periodic single-server queues, *INFORMS Journal on Computing* **30**(1): 71–89.
- Nagel, K., Wagner, P. and Woesler, R. (2003). Still flowing: Approaches to traffic flow and traffic jam modeling, *Operations Research* **51**(5): 681–710.
- Owen, A. B. (2003). The dimension distribution and quadrature test functions, *Statistica Sinica* **13**(1): 1–18.
- Paxson, V. (1994). Empirically derived analytic models of wide-area TCP connections, *IEEE/ACM Transactions on Networking (TON)* **2**(4): 316–336.
- Revuz, D. and Yor, M. (1999). *Continuous martingales and Brownian motion*, third edn, Springer-Verlag, Berlin.
- Rhee, C.-h. and Glynn, P. W. (2015). Unbiased estimation with square root convergence for SDE models, *Operations Research* **63**(5): 1026–1043.
- Rosenbaum, I. and Staum, J. (2017). Multilevel Monte Carlo Metamodeling, *Operations Research* **65**(4): 1062–1077.
- Rubinstein, R. Y. and Kroese, D. P. (2016). *Simulation and the Monte Carlo method*, Vol. 10, John Wiley & Sons, Hoboken, New Jersey.
- Shiryayev, A. N. (1996). *Probability*, Vol. 95 of *Graduate texts in mathematics*, Springer-Verlag, New York.
- Sloan, I. H. and Woniakowski, H. (1998). When are quasi-Monte Carlo algorithms efficient for high dimensional integrals?, *Journal of Complexity* **14**(1): 1 – 33.
- Sobol, I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, *Mathematics and computers in simulation* **55**(1-3): 271–280.

- Thompson, K., Miller, G. J. and Wilder, R. (1997). Wide-area internet traffic patterns and characteristics, *IEEE network* **11**(6): 10–23.
- Wang, X. (2006). On the effects of dimension reduction techniques on some high-dimensional problems in finance, *Operations Research* **54**(6): 1063–1078.
- Wang, X. and Fang, K.-T. (2003). The effective dimension and quasi-Monte Carlo integration, *Journal of Complexity* **19**(2): 101 – 124.
- Wang, X. and Sloan, I. H. (2005). Why are high-dimensional finance problems often of low effective dimension?, *SIAM Journal on Scientific Computing* **27**(1): 159–183.
- Wang, X. and Sloan, I. H. (2011). Quasi-Monte Carlo methods in financial engineering: An equivalence principle and dimension reduction, *Operations Research* **59**(1): 80–95.
- Wang, X. and Tan, K. S. (2013). Pricing and hedging with discontinuous functions: quasi-Monte Carlo methods and dimension reduction, *Management Science* **59**(2): 376–389.
- Whitt, W. (2017). Time-varying queues. Columbia University, New York, NY, <http://www.columbia.edu/~ww2040/allpapers.html>.
- Whitt, W. and You, W. (2016). Time-varying robust queueing. Columbia University, New York, NY, <http://www.columbia.edu/~ww2040/allpapers.html>.