IMPROVING SIMULATION EFFICIENCY WITH QUASI CONTROL VARIATES

Markus Emsermann and Burton Simon

University of Colorado at Denver

ABSTRACT

In a simulation one can often identify a random variable, Y, that is likely to be highly correlated with a random variable of interest, X. If $\mu_Y = E(Y)$ is known then Y can be used as a control variate to estimate $\mu_X = E(X)$ more efficiently than by a direct simulation of X. We study the asymptotic properties of a method that uses Y to potentially speed up the simulation when μ_Y is not known. The method is effective when μ_Y can be efficiently estimated in an auxiliary simulation that does not involve X. We call Y a quasi control variate. For a simulation of length $t > 0$ time units, we invest pt units estimating μ_Y with the auxiliary simulation, yielding $\overline{Y}_{pt}$. The remaining $qt = (1-p)t$ units are spent on the main simulation yielding estimates $(\tilde{X}_{qt}, \tilde{Y}_{qt})$ for (μ_X, μ_Y). The two simulations can be interleaved so they are effectively done simultaneously. For each $p \in (0,1)$ and $\alpha \in \Re$ we have a quasi control variate estimator for μ_X,

$$\hat{X}_t(p,\alpha) = \tilde{X}_{qt} + \alpha(\tilde{Y}_{qt} - \overline{Y}_{pt}), \quad t > 0.$$

We find p and α that minimize the asymptotic variance parameter (AVP) of $\hat{X}_t(p,\alpha)$ in terms of statistics that are estimated during the simulations, and then describe an easily implemented adaptive procedure that achieves the minimum AVP. The adaptive procedure evolves into the optimal quasi control variate procedure if it is more efficient than a direct simulation, $\overline{X}_t \to \mu_X$; otherwise it evolves into the direct simulation. Applications in stochastic linear programming, stochastic partial differential equations (PDE's) and queueing theory are cited.

1 PRELIMINARIES

An experimentor wants to estimate a quantity μ_X by simulation. In order to speed up the simulation the experimentor considers using a certain control variate, but finds that the mean of the proposed control variate is unknown. It is possible to estimate the mean of the control variate in a separate simulation, but the result will not be exact, and the time spent estimating the control variate mean could have been spent on the direct simulation of μ_X. How is the experimentor to proceed?
In this paper we develop an adaptive procedure that is guaranteed to make the best of this situation, where "best" means smallest asymptotic variance for the estimate of μ_X. The adaptive procedure utilizes three simulation programs: simulation A, simulation B, and simulation C. Simulation A is a direct simulation of μ_X, simulation B is a simulation of μ_X and a control variate, and simulation C is a direct simulation of the control variate.
In the remainder of this section we describe the three simulations in more detail and define a *quasi control variate* (QCV) procedure. In section 2 we find the optimal QCV procedure in terms of statistics that can be estimated in simulation B and C. If simulation C is too slow, or if the correlation between the primary and control variates is too small, it is possible that simulation A is more efficient than even the optimal QCV procedure. In section 3 we describe an adaptive procedure that "evolves" into the optimal QCV procedure; unless simulation A is superior, in which case it evolves into simulation A. In section 4 we discuss some applications of QCV procedures. Section 5 summarizes our results.

Simulation A generates a pair of stochastic processes $(\overline{X}_t, \sigma_X^2(t)),\ t > 0$, that satisfy as $t \to \infty$,

$$\overline{X}_t \to \mu_X \quad a.s., \tag{1.1}$$

$$\sigma_X^2(t) \to \sigma_X^2 \quad a.s., \tag{1.2}$$

and

$$s\sqrt{t}\left(\overline{X}_{st} - \mu_X\right) \Rightarrow \sigma_X W_A(s), \quad s \geq 0, \tag{1.3}$$

where $W_A(s)$ is a standard Wiener process, and "$\Rightarrow$" denotes convergence in distribution. Setting $s = 1$ in (1.3) we see that for large t, $\overline{X}_t$ is approximately a Normal random variable with mean μ_X and variance σ_X^2/t. We therefore interpret σ_X^2 as a measure of the simulation's asymptotic efficiency. In general, we measure simulation efficiency as follows.

Definition 1.1: Suppose a simulation generates estimators $Z_t = (Z_t^1, Z_t^2, \ldots, Z_t^m),\ t > 0$ for $\mu = (\mu_1, \mu_2, \ldots, \mu_m)$, satisfying as $t \to \infty$,

$$Z_t \to \mu \quad a.s.,$$

and

$$s\sqrt{t}(Z_{st} - \mu) \Rightarrow \vec{W}(s)\Sigma^{1/2},$$

where $\Sigma = (\Sigma^{1/2})^2$ is an $m \times m$ symmetric positive definite matrix and

$$\vec{W}(s) = (W_1(s), W_2(s), \ldots, W_m(s))$$

is a vector of independent standard Wiener processes. We call Σ the *asymptotic covariance matrix*, and its diagonal elements $(\sigma_1^2, \sigma_2^2, \ldots, \sigma_m^2)$ are the *asymptotic variance parameters* (AVP's) for the estimators $(Z_t^1, Z_t^2, \ldots, Z_t^m)$.

The time variable in our analysis is "real" or "CPU" time, i.e., the elapsed time on the experimentor's watch; and all limits are as $t \to \infty$. The assumptions (1.1), (1.2) and (1.3) are fairly mild. For example, if the experiment is a Monte Carlo simulation then $\mu_X = E(\phi)$ is estimated by sample averages of an i.i.d. sequence $\{\phi_1, \phi_2, \ldots\}$ of replicates of ϕ. If we define N_t to be the number of samples generated by time t then

$$\overline{X}_t = N_t^{-1} \sum_{i=1}^{N_t} \phi_i$$

and

$$\sigma_X^2(t) = \frac{t}{N_t} \left(N_t^{-1} \sum_{i=1}^{N_t} \phi_i^2 - \overline{X}_t^2 \right)$$

satisfy (1.1), (1.2) and (1.3), as long as the variance of ϕ is finite and the times to generate the ϕ_i's form an i.i.d. sequence with finite means [1]. Note that $\sigma_X^2(t) \to \tau_\phi \sigma_\phi^2$, the product of the expected time to generate ϕ and the variance of ϕ. We can expect (1.1), (1.2) and (1.3) to hold for most discrete event simulations as well, e.g., finite horizon, and regenerative and other steady state simulations.

Simulation B also estimates μ_X, but it keeps track of some additional quantities in the hope that the overhead associated with collecting them pays off in improved simulation efficiency. It generates a vector of stochastic processes

$$\left(\tilde{X}_t, \tilde{Y}_t, \tilde{\sigma}_X^2(t), \tilde{\sigma}_Y^2(t), \tilde{\sigma}_{XY}(t) \right), \ t > 0,$$

satisfying

$$(\tilde{X}_t, \tilde{Y}_t) \to (\mu_X, \mu_Y) \quad a.s., \tag{1.4}$$

$$\begin{pmatrix} \tilde{\sigma}_X^2(t) & \tilde{\sigma}_{XY}(t) \\ \tilde{\sigma}_{XY}(t) & \tilde{\sigma}_Y^2(t) \end{pmatrix} \to \begin{pmatrix} \tilde{\sigma}_X^2 & \tilde{\sigma}_{XY} \\ \tilde{\sigma}_{XY} & \tilde{\sigma}_Y^2 \end{pmatrix} \quad a.s., \tag{1.5}$$

and

$$s\sqrt{t}[\tilde{X}_{st} - \mu_X, \tilde{Y}_{st} - \mu_Y] \Rightarrow [W_{B_1}(s), W_{B_2}(s)] \begin{pmatrix} \tilde{\sigma}_X^2 & \tilde{\sigma}_{XY} \\ \tilde{\sigma}_{XY} & \tilde{\sigma}_Y^2 \end{pmatrix}^{1/2}, \quad s \geq 0, \tag{1.6}$$

where $W_{B_1}(s)$ and $W_{B_2}(s)$ are independent standard Wiener processes and $\{\tilde{\sigma}_X^2, \tilde{\sigma}_{XY}^2, \tilde{\sigma}_Y^2\}$ are the asymptotic variance/covariance parameters as defined in definition 1.1.
In the classical method of *control variates*, e.g., [2], simulation B is used to construct

$$\tilde{X}_t^* = \tilde{X}_t - \frac{\tilde{\sigma}_{XY}(t)}{\tilde{\sigma}_Y^2(t)} (\tilde{Y}_t - \mu_Y). \tag{1.7}$$

Equations (1.4) and (1.5) imply that

$$\tilde{X}_t^* \to \mu_X \quad a.s..$$

Assuming $\tilde{\sigma}_Y^2 > 0$, equations (1.5), (1.6) and [3], problem 4.1, imply

$$\begin{aligned} s\sqrt{t}(\tilde{X}_{st}^* - \mu_X) &= s\sqrt{t}[\tilde{X}_{st} - \mu_X, \tilde{Y}_{st} - \mu_Y] \begin{bmatrix} 1 \\ -\frac{\tilde{\sigma}_{XY}(st)}{\tilde{\sigma}_Y^2(st)} \end{bmatrix} \\ &\Rightarrow \sqrt{(1-\tilde{\rho}^2)\tilde{\sigma}_X^2}\, W(s), \end{aligned}$$

where $W(s)$ is a standard Wiener process and

$$\tilde{\rho}^2 = \frac{\tilde{\sigma}_{XY}^2}{\tilde{\sigma}_X^2 \tilde{\sigma}_Y^2} \tag{1.8}$$

is the (squared) asymptotic correlation coefficient for $(\tilde{X}_t, \tilde{Y}_t)$. The AVP for $\tilde{X}_t^*$ is

$$\sigma_{cv}^2 = (1-\tilde{\rho}^2)\tilde{\sigma}_X^2, \tag{1.9}$$

which is the lowest possible AVP for an estimator of μ_X from simulation B. If $(1-\tilde{\rho}^2)\tilde{\sigma}_X^2 < \sigma_X^2$ then simulation B estimates μ_X (asymptotically) more efficiently than simulation A. Clearly $\sigma_X^2 < \tilde{\sigma}_X^2$ since simulation A is dedicated to estimating μ_X. In the control variate literature it is often (implicitly) assumed that $\sigma_X^2 \approx \tilde{\sigma}_X^2$, so there is no reason to consider simulation A. This assumption is valid when simulation A can be easily turned into simulation B by keeping track of quantities that are (essentially) already there, e.g., "internal" control variates. In that case, the expected squared errors of $\overline{X}_t$ and $\tilde{X}^*_{(1-\tilde{\rho}^2)t}$ are approximately the same for large t, which means that simulation B attains the same degree of accuracy as simulation A in a fraction $(1-\tilde{\rho}^2)$ of the time. If the additional overhead in simulation B is *not* insignificant then $\tilde{\sigma}_X^2 \gg \sigma_X^2$ and the speedup (if there is any) is proportionately less, e.g., "external" control variates.
The method of control variates requires that μ_Y is known exactly. There are many examples where μ_Y is unknown, but $\tilde{\rho}^2$ is (potentially) large, suggesting a potential for a significant improvement in efficiency by using simulation B instead of A. We propose to use simulation C in conjunction with simulation B to construct an estimator that is analogous to (1.7) when μ_Y is unknown. Simulation C generates $(\overline{Y}_t, \sigma_Y^2(t)),\ t > 0$, satisfying

$$\overline{Y}_t \to \mu_Y \quad a.s., \tag{1.10}$$

$$\sigma_Y^2(t) \to \sigma_Y^2 \quad a.s., \tag{1.11}$$

and

$$s\sqrt{t}\left(\overline{Y}_{st} - \mu_Y\right) \Rightarrow \sigma_Y W_C(s), \quad s \geq 0, \tag{1.12}$$

where $W_C(s)$ is a standard Wiener process. Clearly $\sigma_Y^2 < \tilde{\sigma}_Y^2$ since simulation C is dedicated to the task of estimating μ_Y, and in many cases

$$r^2 = \frac{\sigma_Y^2}{\tilde{\sigma}_Y^2} \tag{1.13}$$

can be made very small. We can divide t units of time between simulations B and C, and construct an estimator

$$\hat{X}_t(p,\alpha) = \tilde{X}_{qt} + \alpha(\tilde{Y}_{qt} - \overline{Y}_{pt}), \quad t > 0, \tag{1.14}$$

where p is the fraction of time spent executing simulation C, $q = 1 - p$ is the fraction of time spent executing simulation B, and $\alpha \in \Re$ is a constant. We call the resulting scheme a *quasi control variate* (QCV) procedure. If p and α are fixed in advance and do not change during the experiment then the QCV procedure is *static*.
Our goal is to find optimal experimental designs using simulations A, B and C, where the criteria for optimality is minimizing the AVP. Of course, in some cases the optimal design is not a QCV procedure. Even the best QCV procedure may be inferior to simulation B by itself (using $\tilde{X}_t$ and ignoring $\tilde{Y}_t$) if $\tilde{\rho}^2$ is too small, or if r^2 is too big. Since $\sigma_X^2 < \tilde{\sigma}_X^2$, simulation A is in fact optimal in those cases. In the next section we find the optimal QCV procedure of the form (1.14) using simulations B and C. Our recommended procedure, described in section 3, takes simulations A, B and C, and adds a simple "controlling program" that collects statistics from the three simulations as they run, and uses the statistics to assign CPU times to them in dynamically changing ratios. Except in the unlikely pathological case where σ_X^2 is exactly equal to the AVP of the optimal static QCV procedure, the resulting procedure is provably optimal in the sense that no other estimator based on data available from simulations A, B and C yields a lower AVP.

Remark 1.1. In order to avoid trivial or pathological situations we will assume that estimators constructed from our simulations have positive AVP's, and all correlations are less than unit magnitude.

Quasi control variates (by another name) were considered in [4]. The analysis of QCV procedures in [4] relies on "cost" measures associated with simulations B and C that are determined heuristically. The focus is on sensitivity to sample size (small sample analysis). In [5], a strong case is made for using a biased estimate $\hat{\mu}_Y$ for μ_Y, instead of simulation C, in a time constrained simulation experiment where the approximation error, $|\hat{\mu}_Y - \mu_Y|$, is sufficiently small in comparison with the simulation error, $|\tilde{X}_t - \mu_X|$.

2 THE OPTIMAL STATIC QCV PROCEDURE

For fixed $p \in [0, 1)$ and $\alpha \in \Re$ we construct the static QCV estimator

$$\hat{X}_t(p, \alpha) = \tilde{X}_{qt} + \alpha(\tilde{Y}_{qt} - \overline{Y}_{pt}),$$

where $q = 1 - p$ and $(\tilde{X}_s, \tilde{Y}_s),\ s > 0$ and $\overline{Y}_s,\ s > 0$ come from simulations B and C respectively. In this section we find p^* and α^* that minimize the AVP for $\hat{X}_t(p, \alpha)$. In our analysis we assume that $W_{B_1}(\cdot)$, $W_{B_2}(\cdot)$ and $W_C(\cdot)$ are independent Wiener processes.

Remark 2.1. In general it is beneficial to make the correlation between $\tilde{X}_{qt}$ and $\tilde{Y}_{qt} - \overline{Y}_{pt}$ large, so designs that induce a correlation between simulations B and C could be superior in some cases. However, there is no obvious way to induce correlations between simulations B and C in a systematic and predictable manner. Our analysis apparently extends without *theoretical* difficulties to the case where simulations B and C are correlated, but what is gained in generality is lost in clarity of exposition, so we will not work out the details.

From (1.6) and (1.12) the AVP for $\hat{X}_t(p,\alpha)$ is found to be

$$v^2(p,\alpha) = q^{-1}\tilde{\sigma}_X^2 + 2\alpha q^{-1}\tilde{\sigma}_{XY} + \alpha^2(q^{-1}\tilde{\sigma}_Y^2 + p^{-1}\sigma_Y^2) \tag{2.1}$$

as long as $0 < p < 1$. If $p = 0$ then $\hat{X}_t(p,\alpha)$ does not converge to μ_X unless $\alpha = 0$ (since presumably $\mu_Y \neq 0$). When $p = \alpha = 0$ the estimator reduces to $\tilde{X}_t$. If $p = 1$ then $\hat{X}_t(p,\alpha)$ does not estimate μ_X, so this case is not relevant.

Theorem 2.1 $v^2(p,\alpha),\ p \in [0,1),\ \alpha \in \Re$ has a unique global minimum at (p^*,α^*), where

$$p^* = \begin{cases} \frac{\sqrt{r^2\tilde{\rho}^2(1-r^2)(1-\tilde{\rho}^2)} - r^2(1-r^2)}{(1-r^2)(1-r^2-\tilde{\rho}^2)} & \text{if } r^2+\tilde{\rho}^2 \neq 1 \text{ and } \tilde{\rho}^2 > r^2, \\ 1 - \frac{1}{2\tilde{\rho}^2} & \text{if } r^2+\tilde{\rho}^2 = 1 \text{ and } \tilde{\rho}^2 > r^2, \\ 0 & \text{if } \tilde{\rho}^2 \leq r^2, \end{cases} \tag{2.2}$$

and $\alpha^* = -\frac{\tilde{\sigma}_{XY}}{\tilde{\sigma}_Y^2}\tilde{\alpha}^*$, where

$$\tilde{\alpha}^* = \begin{cases} \frac{\sqrt{\tilde{\rho}^2(1-r^2)(1-\tilde{\rho}^2)} + r(r^2-1)}{(1-r^2)(\sqrt{\tilde{\rho}^2(1-r^2)(1-\tilde{\rho}^2)} - r\tilde{\rho}^2)} & \text{if } r^2+\tilde{\rho}^2 \neq 1 \text{ and } \tilde{\rho}^2 > r^2, \\ \frac{2\tilde{\rho}^2-1}{\tilde{\rho}^2} & \text{if } r^2+\tilde{\rho}^2 = 1 \text{ and } \tilde{\rho}^2 > r^2, \\ 0 & \text{if } \tilde{\rho}^2 \leq r^2. \end{cases} \tag{2.3}$$

Furthermore,
(a) p^* and α^* are continuous in the variables $\tilde{\sigma}_X^2$, $\tilde{\sigma}_{XY}$, $\tilde{\sigma}_Y^2$ and σ_Y^2, and
(b) $p^* > 0$ if and only if $\tilde{\rho}^2 > r^2$.

Proof: If $\tilde{\rho}^2 > r^2$ we show that the Hessian matrix $\nabla^2 v^2$ is positive definite, and then solve $\nabla v^2 = 0$ to obtain the unique global minimum. To show that $\nabla^2 v^2$ is positive definite it suffices to verify

$$\frac{\partial^2 v^2}{\partial \alpha^2} > 0$$

and

$$\mathrm{Det}(\nabla^2 v^2) = \frac{\partial^2 v^2}{\partial \alpha^2}\frac{\partial^2 v^2}{\partial p^2} - \left(\frac{\partial^2 v^2}{\partial p \partial \alpha}\right)^2 > 0.$$

We find

$$\frac{\partial^2 v^2}{\partial \alpha^2} = 2\left(q^{-1}\tilde{\sigma}_Y^2 + p^{-1}\sigma_Y^2\right) > 0,$$

and

$$\mathrm{Det}(\nabla^2 v^2) = \frac{4[(\tilde{\sigma}_X^2\tilde{\sigma}_Y^2 - \tilde{\sigma}_{XY}^2)p^3 + q\sigma_Y^2(\tilde{\sigma}_X^2 p^2 + \tilde{\sigma}_Y^2\alpha^2 + 2\tilde{\sigma}_{XY}p\alpha)]}{p^3 q^4}. \tag{2.4}$$

It remains to show the numerator in (2.4) is positive. Since $\tilde{\rho}^2 < 1$ (see Remark 1.1) we have

$$(\tilde{\sigma}_X^2\tilde{\sigma}_Y^2 - \tilde{\sigma}_{XY}^2)p^3 + q\sigma_Y^2(\tilde{\sigma}_X^2 p^2 + \tilde{\sigma}_Y^2\alpha^2 + 2\tilde{\sigma}_{XY}p\alpha) \ \geq \ \tilde{\sigma}_X^2\tilde{\sigma}_Y^2(1-\tilde{\rho}^2)p^3 \tag{2.5}$$

which is positive, proving that $\nabla^2 v^2$ is positive definite. To find the global minimum we solve

$$\frac{\partial v^2}{\partial p} = q^{-2}\tilde{\sigma}_X^2 + 2q^{-2}\tilde{\sigma}_{XY}\alpha + \left(q^{-2}\tilde{\sigma}_Y^2 - p^{-2}\sigma_Y^2\right)\alpha^2 = 0 \tag{2.6}$$

and

$$\frac{\partial v^2}{\partial \alpha} = 2\alpha \left(q^{-1}\tilde{\sigma}_Y^2 + p^{-1}\sigma_Y^2 \right) + 2q^{-1}\tilde{\sigma}_{XY} = 0. \tag{2.7}$$

For fixed $p \in (0,1)$, solving (2.7) yields the corresponding optimal value of α,

$$\alpha_p^* = -\frac{\tilde{\sigma}_{XY}}{\tilde{\sigma}_Y^2} \left(\frac{p}{p+qr^2} \right), \tag{2.8}$$

where r^2 is given by (1.13). Substituting (2.8) for α in (2.6) and simplifying yields (2.2) and (2.3). One can verify that the expressions (2.2b) and (2.3b) are limits of the expressions (2.2a) and (2.3a). Since $\nabla^2 v^2 > 0$, the implicit function theorem applied to (2.6) and (2.7) implies that p^* and α^* are continuous functions of the remaining variables $\tilde{\sigma}_X^2$, $\tilde{\sigma}_{XY}$, $\tilde{\sigma}_Y^2$ and σ_Y^2 when $\tilde{\rho}^2 > r^2$. By taking the limit $r^2 \to \tilde{\rho}^2$ in (2.2a), (2.2b), (2.3a) and (2.3b) it follows that p^* and α^* are continuous on the boundary $r^2 = \tilde{\rho}^2$.
We now check that our expression (2.2) for p^* takes values in the interval $(0,1)$ when $\tilde{\rho}^2 > r^2$. Since $0 < r^2 < 1$ and $0 \le \tilde{\rho}^2 < 1$ we can consider the following four cases: (i) $\tilde{\rho}^2 > \max(r^2, 1-r^2)$, (ii) $r^2 < \min(\tilde{\rho}^2, 1-\tilde{\rho}^2)$, (iii) $r^2 + \tilde{\rho}^2 = 1$, and (iv) $\tilde{\rho}^2 \le r^2$. Case (iv) is trivial since (2.2c) gives us $p^* = 0$. Case (iii) is covered by case (iv) when $\tilde{\rho}^2 \le r^2$; if $\tilde{\rho}^2 > r^2$ then $1/2 < \tilde{\rho}^2 < 1$ so (2.2b) gives us $0 < p^* < 1$. There remains cases (i) and (ii). In case (i) it follows that the numerator and denominator in (2.2a) are negative, and the denominator is less than the numerator; in case (ii) the numerator and denominator are positive and the numerator is less than the denominator.

If $r^2 \ge \tilde{\rho}^2$ we use (2.1) and (2.8) to write

$$v^2(p, \alpha_p^*) = \frac{\tilde{\sigma}_X^2}{q} \left(1 - \frac{\tilde{\rho}^2 p}{p+qr^2} \right).$$

It follows that

$$v^2(p, \alpha_p^*) < \tilde{\sigma}_X^2 \iff 0 < p < \frac{\tilde{\rho}^2 - r^2}{1-r^2}. \tag{2.9}$$

Thus, if $r^2 \ge \tilde{\rho}^2$ then $v^2(p, \alpha_p^*) \ge \tilde{\sigma}_X^2$ for every $p \in (0,1)$, which implies that $p^* = 0$, and therefore $\alpha^* = 0$ as well. □

Corollary 2.1.1 The optimal QCV procedure has AVP,

$$v^2(p^*, \alpha^*) = \frac{\tilde{\sigma}_X^2 (1 - \hat{\rho}^2)}{q^*}, \tag{2.10}$$

where

$$\hat{\rho}^2 = \frac{p^*}{p^* + q^* r^2}\, \tilde{\rho}^2.$$

Proof: Define

$$\beta = \frac{p^*}{p^* + q^* r^2}.$$

From (2.1) and (2.8) we write

$$\begin{aligned} v^2(p^*,\alpha^*) &= \frac{\tilde{\sigma}_X^2}{q^*} + \tilde{\sigma}_X^2\tilde{\rho}^2\left(\frac{\beta^2-2\beta}{q^*}+\frac{\beta^2 r^2}{p^*}\right) \\ &= \frac{\tilde{\sigma}_X^2(1-\tilde{\rho}^2\beta)}{q^*} = \frac{\tilde{\sigma}_X^2(1-\hat{\rho}^2)}{q^*} \quad \square \end{aligned}$$

Remark 2.2. Since simulations B and C are statistically independent we can interpret $\hat{\rho}$ as the asymptotic correlation between $\tilde{X}_{q^*t}$ and $\tilde{Y}_{q^*t} - \overline{Y}_{p^*t}$. If $r^2 \approx 0$ then the QCV procedure is "almost" a control variate procedure, which would have AVP, $\tilde{\sigma}_X^2(1-\tilde{\rho}^2)$. This is consistent with (2.10) since $r^2 \approx 0$ implies that $q^* \approx 1$ and $\hat{\rho}^2 \approx \tilde{\rho}^2$. In general, we have $\hat{\rho}^2 < \tilde{\rho}^2$ and $q^* < 1$, so there is a price for not knowing μ_Y.

Corollary 2.1.2 The following statements are equivalent:

(a) $p^* > 0$, (It is optimal to devote some time to simulation C.)

(b) $v^2(p^*,\alpha^*) < \tilde{\sigma}_X^2$ (The optimal static QCV procedure has a lower AVP than simulation B alone.)

(c) $r^2 < \tilde{\rho}^2$.

Proof: Theorem 2.1 shows that (a) and (c) are equivalent. Since p^* and α^* are the unique optimal parameters we have $v^2(p^*,\alpha^*) < v^2(0,0) = \tilde{\sigma}_X^2$ when $p^* > 0$. If $p^* = 0$ then (2.8) implies that $\alpha^* = 0$, so (a) and (b) are equivalent. $\square$

Figure 1 shows p^* and $\tilde{\alpha}^*$ (defined in (2.3)) as functions of r^2 and $\tilde{\rho}^2$. Since $\tilde{\alpha}^* = 1$ in the "control variate limit" (i.e., as $r^2 \to 0$), $\tilde{\alpha}^*$ can be interpreted as the "correction factor" for QCV estimators. From (2.10) we see that the "speedup factor" of the optimal static QCV procedure compared to simulation B alone is $q^*(1-\hat{\rho}^2)^{-1}$ (the ratio of their AVP's), which is plotted as a function of r^2 and $\tilde{\rho}^2$ in figure 2.

3 AN OPTIMAL DYNAMIC QCV PROCEDURE

In the classical control variate environment, the experimentor has access to simulations A and B. It is well known that the optimal utilization of simulation B (minimum AVP) is to form the estimator $\tilde{X}_t^*$ given by (1.7). This tactic is easy to implement. However, an "honest" experimentor must determine that $\tilde{X}_t^*$ is more efficient than $\overline{X}_t$ from simulation A before going ahead with simulation B. Fortunately, a rough comparison can be done at very little cost. One simple heuristic has the experimentor observing $\overline{X}_t$ from simulation A and $\tilde{X}_t^*$ from simulation B "simultaneously" in an initial phase, until a statistical analysis reveals that one is more efficient than the other (at, say, a 99% confidence level). From that point on, the more efficient simulation runs. If one choice is much better than the other, then the initial phase is short. If the choices are comparable then little is lost by

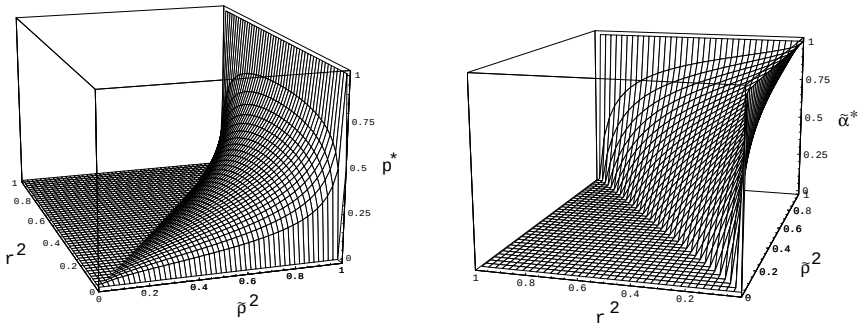

Figure 1: p^* and $\tilde{\alpha}^*$ as functions of r^2 and $\tilde{\rho}^2$

running the inferior one. Likewise, in the QCV environment the experimentor must determine that the optimal procedure based on simulations B and C is more efficient than simulation A before going ahead with a QCV procedure.
In this section we describe a dynamic procedure using simulations A, B and C that is no more difficult to use than the simple heuristic for control variate experiments, and is guaranteed to provide an estimator that has an AVP as close to optimal as the experimentor desires ("ϵ–optimal"). An optimal estimator has AVP

$$v_*^2 = \min\left(\sigma_X^2, v^2(p^*, \alpha^*)\right),$$

i.e., use the optimal static QCV procedure if it is superior to simulation A; otherwise use simulation A. The experimentor can in fact "gamble" and obtain an optimal estimator by using a procedure even simpler than the ϵ–optimal procedure. The simpler procedure is not guaranteed to work if $v^2(p^*, \alpha^*) = \sigma_X^2$, however one can reasonably argue that those two quantities being precisely equal is a somewhat pathological case.

At the beginning of the simulation experiment we do not know the values of any of the parameters that would allow us to run the optimal static QCV procedure derived in the previous section or decide whether or not it is more efficient than simulation A. We

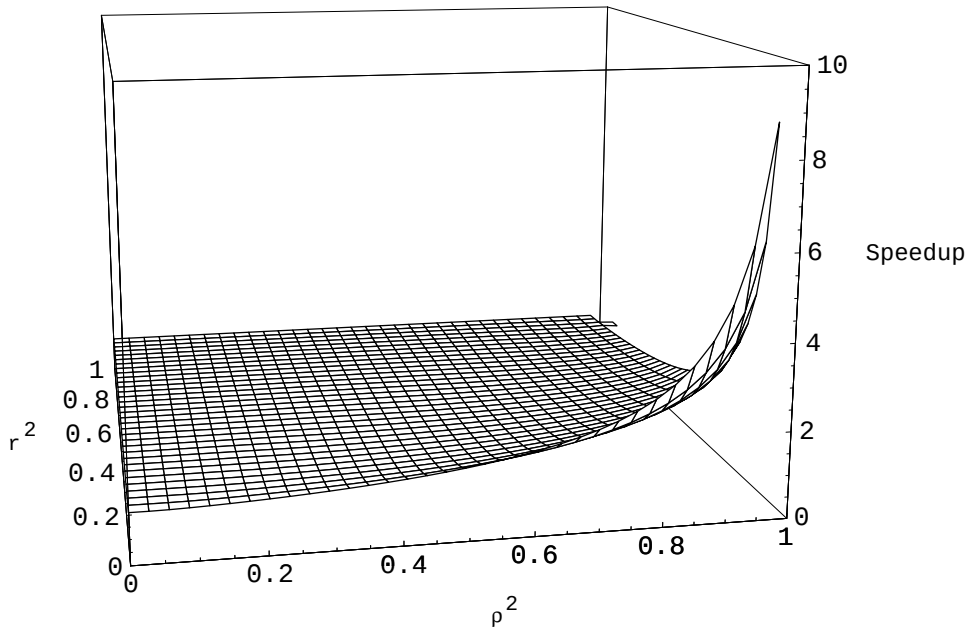

Figure 2: Speedup factor as a function of r^2 and $\tilde{\rho}^2$

therefore must periodically revise our strategy based on the information from the simulations as they run. The simulation experiment is divided into time segments, which (without loss of generality) are one unit in duration. The nth segment begins at time $t = n,\ n = 0, 1, 2, \ldots$. (We assume that one unit of time is enough so that any overhead associated with our procedure is negligible.) At the end of the $(n-1)$st segment we construct estimates of relevant parameters based on all the simulation data obtained up till that point, and devise a strategy for the nth segment. The strategy for the nth segment is simply the fractions of time that will be devoted to simulations A, B and C, which we will denote by $(\tau_n^A, \tau_n^B, \tau_n^C)$. Define a_t, b_t and c_t to be the CPU times devoted to simulations A, B and C up to time t, and note that

$$\sum_{n=0}^{\lfloor t \rfloor - 1} (\tau_n^A, \tau_n^B, \tau_n^C) \leq (a_t, b_t, c_t) \leq \sum_{n=0}^{\lfloor t \rfloor} (\tau_n^A, \tau_n^B, \tau_n^C). \tag{3.1}$$

Let

$$r_t^2 = \frac{\sigma_Y^2(c_t)}{\tilde{\sigma}_Y^2(b_t)} \qquad \text{and} \qquad \tilde{\rho}_t^2 = \frac{\tilde{\sigma}_{XY}^2(b_t)}{\tilde{\sigma}_X^2(b_t)\tilde{\sigma}_Y^2(b_t)}$$

be the estimates of r^2 and $\tilde{\rho}^2$ based on information from simulations B and C up to time t, and define

$$p_t = \begin{cases} \frac{\sqrt{r_t^2\tilde{\rho}_t^2(1-r_t^2)(1-\tilde{\rho}_t^2)} - r_t^2(1-r_t^2)}{(1-r_t^2)(1-r_t^2-\tilde{\rho}_t^2)} & \text{if } r_t^2 + \tilde{\rho}_t^2 \neq 1 \text{ and } \tilde{\rho}_t^2 > r_t^2, \\ 1 - \frac{1}{2\tilde{\rho}_t^2} & \text{if } r_t^2 + \tilde{\rho}_t^2 = 1 \text{ and } \tilde{\rho}_t^2 > r_t^2, \\ 0 & \text{if } \tilde{\rho}_t^2 \leq r_t^2, \end{cases}$$

and

$$\alpha_t = \begin{cases} -\frac{\tilde{\sigma}_{XY}(b_t)}{\tilde{\sigma}_Y^2(b_t)} \left(\frac{\sqrt{\tilde{\rho}_t^2(1-r_t^2)(1-\tilde{\rho}_t^2)} + r_t(r_t^2-1)}{(1-r_t^2)(\sqrt{\tilde{\rho}_t^2(1-r_t^2)(1-\tilde{\rho}_t^2)} - r_t\tilde{\rho}_t^2)} \right) & \text{if } r_t^2 + \tilde{\rho}_t^2 \neq 1 \text{ and } \tilde{\rho}_t^2 > r_t^2, \\ -\frac{\tilde{\sigma}_{XY}(b_t)}{\tilde{\sigma}_Y^2(b_t)} \left(\frac{2\tilde{\rho}_t^2 - 1}{\tilde{\rho}_t^2} \right) & \text{if } r_t^2 + \tilde{\rho}_t^2 = 1 \text{ and } \tilde{\rho}_t^2 > r_t^2, \\ 0 & \text{if } \tilde{\rho}_t^2 \leq r_t^2, \end{cases}$$

to be the estimates of p^* and α^* from (2.2) and (2.3). Let $q_t = 1 - p_t$, and from (2.10) construct the estimate of $v^2(p^*, \alpha^*)$,

$$v_t^2 = \frac{\tilde{\sigma}_X^2(b_t)}{q_t} \left(1 - \frac{p_t}{p_t + q_t r_t^2} \tilde{\rho}_t^2 \right).$$

At time t we can estimate the AVP for $\overline{X}_t$ by $\sigma_X^2(a_t)$. Let

$$\Delta_t = \sigma_X^2(a_t) - v_t^2$$

be the estimate of the difference of the two AVP's, and define

$$\Delta = \sigma_X^2 - v^2(p^*, \alpha^*)$$

to be the true difference. If a_t, b_t and c_t each grow without bound, estimates of all the parameters needed to decide whether $\overline{X}_t$ or $\hat{X}_t(p^*, \alpha^*)$ is more efficient will converge almost surely to their exact values. On the other hand, we want to spend an asymptotically negligible fraction of time running a suboptimal experimental design.

We now describe our recommended QCV procedure. Let $\delta_i \in (0, \frac{1}{3}),\ i = 1, 2, \ldots$ be a sequence satisfying

$$\delta_i \to 0, \tag{3.2}$$

$$\sum_{i=1}^{\infty} \delta_i = \infty, \tag{3.3}$$

and

$$n^{-\frac{1}{2}} \sum_{i=1}^{n} \delta_i \to 0. \tag{3.4}$$

For example, $\delta_i = ci^{-\gamma}$ satisfies (3.2), (3.3) and (3.4) for any $0 < c < 1/3$ and $1/2 < \gamma \leq 1$. Assign $\tau_0^A > 0$, $\tau_0^B > 0$ and $\tau_0^C > 0$ time units to simulations A, B and C in the zeroth segment ($\tau_0^A + \tau_0^B + \tau_0^C = 1$, but otherwise arbitrary). Choose $\epsilon \geq 0$. For $n \geq 1$, the simulations are assigned

$$\tau_n^B = \begin{cases} (1 - 2\delta_n) \max(q_n, \delta_n) & \text{if } \Delta_n > \epsilon \\ \delta_n & \text{if } \Delta_n < -\epsilon \\ \delta_n + \frac{\Delta_n + \epsilon}{2\epsilon}[(1 - 2\delta_n) \max(q_n, \delta_n) - \delta_n] & \text{if } -\epsilon \leq \Delta_n \leq \epsilon \end{cases} \tag{3.5}$$

$$\tau_n^C = \begin{cases} (1 - 2\delta_n) \max(p_n, \delta_n) & \text{if } \Delta_n > \epsilon \\ \delta_n & \text{if } \Delta_n < -\epsilon \\ \delta_n + \frac{\Delta_n + \epsilon}{2\epsilon}[(1 - 2\delta_n) \max(p_n, \delta_n) - \delta_n] & \text{if } -\epsilon \leq \Delta_n \leq \epsilon \end{cases} \tag{3.6}$$

and

$$\tau_n^A = 1 - \tau_n^B - \tau_n^C \tag{3.7}$$

time units in the nth segment. At time t we can estimate μ_X by using any weighted average of $\overline{X}_{a_t}$ (from simulation A) and $\tilde{X}_{b_t} + \alpha_t(\tilde{Y}_{b_t} - \overline{Y}_{c_t})$ (from simulations B and C). For $t > 0$ let

$$X_t^\epsilon = \frac{a_t}{t}\overline{X}_{a_t} + \frac{b_t + c_t}{t}\left(\tilde{X}_{b_t} + \alpha_t(\tilde{Y}_{b_t} - \overline{Y}_{c_t})\right), \tag{3.8}$$

and let v_ϵ^2 be the AVP for X_t^ϵ.

Remark 3.1. Both (3.5) and (3.6) linearly interpolate between the values at $-\epsilon$ and ϵ. If $\epsilon = 0$ then the alternatives in (3.5) and (3.6) reduce to $\Delta_n > 0$ and $\Delta_n \leq 0$. The resulting procedure, X_t^0, is well defined, and in fact is arguably easier to implement than procedures $X_t^\epsilon,\ \epsilon > 0$. We recommend using the estimator X_t^0 unless it is possible that Δ is extremely small.

Remark 3.2. The dynamic QCV procedure can be simplified and applied to standard control variate procedures. When μ_Y is known we define $\Delta_t = \sigma_X^2(a_t) - \tilde{\sigma}_X^2(b_t)(1 - \tilde{\rho}_t^2)$, and set

$$\tau_n^B = \begin{cases} 1 - \delta_n & \text{if } \Delta_n > \epsilon \\ \delta_n & \text{if } \Delta_n < -\epsilon \\ \delta_n + \frac{\Delta_n + \epsilon}{2\epsilon}(1 - 2\delta_n) & \text{if } -\epsilon \leq \Delta_n \leq \epsilon \end{cases}$$

and $\tau_n^A = 1 - \tau_n^B$. (There is no simulation C, so $\tau_n^C = 0$.)

Theorem 3.1 If simulations A, B and C are statistically independent then for every $\epsilon > 0$,

$$v_\epsilon^2 \leq v_*^2 + \epsilon/8.$$

If $\Delta \neq 0$ then X_t^0 is an optimal estimator, i.e., $v_0^2 = v_*^2$.

Proof: We first show that a_t, b_t and c_t each grow without bound as $t \to \infty$. From (3.5) and (3.6) we have $\tau_n^B \geq \delta_n - 2\delta_n^2$ and $\tau_n^C \geq \delta_n - 2\delta_n^2$, so (3.1), (3.2) and (3.3) imply that $b_t \to \infty$ *a.s.* and $c_t \to \infty$ *a.s.*. From (3.7) we have

$$\begin{aligned} \tau_n^A &\geq 1 - (1 - 2\delta_n)\left(\max(q_n, \delta_n) + \max(p_n, \delta_n)\right) \\ &\geq 1 - (1 - 2\delta_n)(1 + \delta_n) = \delta_n + 2\delta_n^2, \end{aligned}$$

so $a_t \to \infty$. Assumptions (1.1), (1.2), (1.4), (1.5), (1.10) and (1.11) imply that $r_t \to r$, $\tilde{\rho}_t \to \tilde{\rho}$, $\sigma_X^2(a_t) \to \sigma_X^2$, $\overline{X}_t \to \mu_X$, $(\tilde{X}_t, \tilde{Y}_t) \to (\mu_X, \mu_Y)$ and $\overline{Y}_t \to \mu_Y$, almost surely. From Theorem 2.1 we conclude that $p_t \to p^*$ and $\alpha_t \to \alpha^*$, and since (2.1) is continuous in p and α, we have $v_t^2 \to v^2(p^*, \alpha^*)$, so $\Delta_t \to \Delta$. We consider the three cases $\Delta > \epsilon$, $\Delta < -\epsilon$ and $-\epsilon \leq \Delta \leq \epsilon$ separately.
If $\Delta > \epsilon$ then there is almost surely an $N < \infty$ so that $\Delta_n > \epsilon$ for every $n > N$. Thus, for $n > N$ we have $\tau_n^A \leq 2\delta_n$, so (3.4) implies

$$t^{-\frac{1}{2}} a_t \leq t^{-\frac{1}{2}} \left(N + 2 \sum_{n=N}^{\lfloor t \rfloor} \delta_n \right) \to 0. \tag{3.9}$$

Write

$$\begin{aligned} s\sqrt{t}(X_{st}^\epsilon - \mu_X) &= \frac{a_{st}}{st} s\sqrt{t}(\overline{X}_{a_{st}} - \mu_X) \\ &\quad + \frac{b_{st} + c_{st}}{st} s\sqrt{t}[\tilde{X}_{b_{st}} - \mu_X, \tilde{Y}_{b_{st}} - \mu_Y] \begin{bmatrix} 1 \\ \alpha_{st} \end{bmatrix} \\ &\quad - \frac{b_{st} + c_{st}}{st} \alpha_{st} s\sqrt{t}(\overline{Y}_{c_{st}} - \mu_Y) \end{aligned} \tag{3.10}$$

For $st > N$ we have

$$\frac{c_{st}}{st} = \frac{1}{st} \left(c_N + \sum_{n=N}^{\lfloor st \rfloor} (1 - 2\delta_n) \max(p_n, \delta_n) + \mathcal{O}(1) \right) \to p^* \quad a.s.,$$

since $p_n \to p^*$ *a.s.* and $\delta_n \to 0$. From Corollary 2.1.2 we know that $p^* > 0$ when $\Delta > 0$. Using the time change theorem ([3], Theorem 17.1) along with (1.12) we conclude that

$$s\sqrt{t}(\overline{Y}_{c_{st}} - \mu_Y) \Rightarrow \frac{\sigma_Y}{\sqrt{p^*}} W_C(s).$$

Since $\alpha_{st} \to \alpha^*$ *a.s.* and $\frac{b_{st}+c_{st}}{st} \to 1$ *a.s.*, we apply [3], problem 4.1 and conclude

$$\frac{b_{st} + c_{st}}{st} \alpha_{st} s\sqrt{t}(\overline{Y}_{c_{st}} - \mu_Y) \Rightarrow \frac{\alpha^* \sigma_Y}{\sqrt{p^*}} W_C(s). \tag{3.11}$$

Likewise, from (1.6) we can show that

$$\frac{b_{st}+c_{st}}{st}s\sqrt{t}\left[\tilde{X}_{b_{st}}-\mu_X,\tilde{Y}_{b_{st}}-\mu_Y\right]\begin{bmatrix}1\\ \alpha_{st}\end{bmatrix}$$

$$\Rightarrow \frac{1}{\sqrt{q^*}}\left[W_{B_1}(s),W_{B_2}(s)\right]\begin{pmatrix}\tilde{\sigma}_X^2 & \tilde{\sigma}_{XY}\\ \tilde{\sigma}_{XY} & \tilde{\sigma}_Y^2\end{pmatrix}^{\frac{1}{2}}\begin{bmatrix}1\\ \alpha^*\end{bmatrix}. \tag{3.12}$$

Finally, since $a_t\to\infty$, (1.1) implies that $(\overline{X}_{a_{st}}-\mu_X)\to 0$ *a.s.*. Thus, (3.9) yields

$$\frac{a_{st}}{st}s\sqrt{t}(\overline{X}_{a_{st}}-\mu_X)\to 0 \quad a.s. \tag{3.13}$$

Equations (3.11), (3.12), (3.13) and (3.10) together imply

$$\begin{aligned} s\sqrt{t}(X^\epsilon_{st}-\mu_X) &\Rightarrow \frac{1}{\sqrt{q^*}}\left[W_{B_1}(s),W_{B_2}(s)\right]\begin{pmatrix}\tilde{\sigma}_X^2 & \tilde{\sigma}_{XY}\\ \tilde{\sigma}_{XY} & \tilde{\sigma}_Y^2\end{pmatrix}^{\frac{1}{2}}\begin{bmatrix}1\\ \alpha^*\end{bmatrix}-\frac{\alpha^*\sigma_Y}{\sqrt{p^*}}W_C(s)\\ &=_d v(p^*,\alpha^*)W(s),\end{aligned}$$

where $W(s)$ is a standard Wiener process. Thus, in this case $v_\epsilon^2=v_*^2$.

If $\Delta<-\epsilon$ then there is an $N<\infty$ so that $\Delta_n<-\epsilon$ for every $n>N$. This time (3.4) implies $t^{-\frac{1}{2}}(b_t+c_t)\to 0$, which along with (1.4) and (1.10) implies that the second and third term in (3.10) converge to zero almost surely. Also,

$$\frac{a_{st}}{st}=\frac{1}{st}\left(a_N+\sum_{i=N+1}^{\lfloor st\rfloor}(1-2\delta_i)+\mathcal{O}(1)\right)\to 1 \quad a.s.,$$

so Billingsley's time change theorem and problem 4.1 applied to the first term in (3.10) yields

$$\frac{a_{st}}{st}s\sqrt{t}(\overline{X}_{a_{st}}-\mu_X)\Rightarrow\sigma_X W_A(s),$$

so

$$s\sqrt{t}(X^\epsilon_{st}-\mu_X)\Rightarrow\sigma_X W_A(s),$$

and again, $v_\epsilon^2=v_*^2$.

If $\Delta\neq 0$ we have proven that $v_\epsilon^2=v_*^2$ for any $\epsilon\geq 0$. Thus, if $\Delta\neq 0$ then X_t^0 is optimal.

If $-\epsilon\leq\Delta\leq\epsilon$ then

$$\tau_n^A\to\frac{\epsilon-\Delta}{2\epsilon},\quad \tau_n^B\to\frac{\epsilon+\Delta}{2\epsilon}q^* \quad\text{and}\quad \tau_n^C\to\frac{\epsilon+\Delta}{2\epsilon}p^*,$$

so

$$\frac{a_{st}}{st}\to\frac{\epsilon-\Delta}{2\epsilon},\quad \frac{b_{st}}{st}\to\frac{\epsilon+\Delta}{2\epsilon}q^* \quad\text{and}\quad \frac{c_{st}}{st}\to\frac{\epsilon+\Delta}{2\epsilon}p^*.$$

Thus, the time change theorem applied to (3.10) yields

$$\begin{aligned} s\sqrt{t}(X^{\epsilon}_{st} - \mu_X) \quad \Rightarrow \quad & \sigma_X \sqrt{\frac{\epsilon - \Delta}{2\epsilon}} W_A(s) \\ & + \sqrt{\frac{\epsilon + \Delta}{2\epsilon q^*}} [W_{B_1}(s), W_{B_2}(s)] \begin{pmatrix} \tilde{\sigma}_X^2 & \tilde{\sigma}_{XY} \\ \tilde{\sigma}_{XY} & \tilde{\sigma}_Y^2 \end{pmatrix}^{\frac{1}{2}} \begin{bmatrix} 1 \\ \alpha^* \end{bmatrix} \\ & - \alpha^* \sigma_Y \sqrt{\frac{\epsilon + \Delta}{2\epsilon p^*}} W_C(s), \end{aligned}$$

which yields

$$v_\epsilon^2 = \frac{\epsilon - \Delta}{2\epsilon} \sigma_X^2 + \frac{\epsilon + \Delta}{2\epsilon} v^2(p^*, \alpha^*).$$

A simple calculation shows that $v_\epsilon^2 - v_*^2$ is maximized when $|\Delta| = \epsilon/2$, in which case its value is $\epsilon/8$. $\square$

If Z_1 and Z_2 are independent with $E(Z_1) = E(Z_2) = \mu_X$, $\mathrm{Var}(Z_1) = \sigma_1^2$ and $\mathrm{Var}(Z_2) = \sigma_2^2$ then the minimum variance weighted average of Z_1 and Z_2 is

$$Z = \frac{\sigma_1^{-2}}{\sigma_1^{-2} + \sigma_2^{-2}} Z_1 + \frac{\sigma_2^{-2}}{\sigma_1^{-2} + \sigma_2^{-2}} Z_2.$$

With this in mind, it may be preferable to use

$$\hat{X}_t^{\epsilon} = \beta_t \overline{X}_{a_t} + (1 - \beta_t) \left(\tilde{X}_{b_t} + \alpha_t (\tilde{Y}_{b_t} - \overline{Y}_{c_t}) \right), \tag{3.14}$$

where

$$\beta_t = \frac{\sigma_X^{-2}(a_t) a_t}{\sigma_X^{-2}(a_t) a_t + v_t^{-2}(t - a_t)}$$

instead of X_t^{ϵ} as a QCV estimator for μ_X. Let $\hat{v}_\epsilon^2$, $\epsilon \geq 0$ be the AVP for $\hat{X}_t^{\epsilon}$. The following result shows that asymptotically, $\hat{X}_t^{\epsilon}$ and X_t^{ϵ} are equivalent.

Corollary 3.1.1. If Simulations A, B and C are statistically independent then for every $\epsilon > 0$

$$\hat{v}_\epsilon^2 \leq v_*^2 + \epsilon/8.$$

If $\Delta \neq 0$ then $v_0^2 = v_*^2$.

Proof: It suffices to show

$$\frac{\sigma_X^{-2}(a_t) a_t}{\sigma_X^{-2}(a_t) a_t + v_t^{-2}(t - a_t)} - \frac{a_t}{t} \to 0 \quad a.s. \tag{3.15}$$

We write

$$\frac{\sigma_X^{-2}(a_t) a_t}{\sigma_X^{-2}(a_t) a_t + v_t^{-2}(t - a_t)} - \frac{a_t}{t} = \frac{\sigma_X^{-2}(a_t) \frac{a_t}{t}(1 - \frac{a_t}{t}) - v_t^{-2} \frac{a_t}{t}(1 - \frac{a_t}{t})}{\sigma_X^{-2}(a_t) \frac{a_t}{t} + v_t^{-2}(1 - \frac{a_t}{t})}.$$

From the proof of Theorem 3.1 we know that if $\Delta \neq 0$ then either $\frac{a_t}{t} \to 0$ or $\frac{a_t}{t} \to 1$ almost surely. Thus, if $\Delta \neq 0$ then (3.15) holds. There remains the case $\Delta = 0$. Choose $0 < \eta < \sigma_X^{-2}$ and let T be large enough so that if $t > T$ then $\max(|\sigma_X^{-2}(a_t) - \sigma_X^{-2}|, |v_t^{-2} - \sigma_X^{-2}|) < \eta$. Then for $t > T$,

$$\begin{aligned} \left| \frac{\sigma_X^{-2}(a_t)\frac{a_t}{t}(1-\frac{a_t}{t}) - v_t^{-2}\frac{a_t}{t}(1-\frac{a_t}{t})}{\sigma_X^{-2}(a_t)\frac{a_t}{t} + v_t^{-2}(1-\frac{a_t}{t})} \right| &\leq \frac{2\eta\frac{a_t}{t}(1-\frac{a_t}{t})}{\sigma_X^{-2} - \eta} \\ &\leq \frac{\eta/2}{\sigma_X^{-2} - \eta}. \end{aligned}$$

Since η is arbitrary we have shown that (3.15) holds when $\Delta = 0$. □

4 APPLICATIONS

We begin this section by describing a "generic" example of a simulation where a particular type of quasi control variate improves the efficiency of the experiment. We then describe three applications, two of which are variations on the generic example.

Generic example: Let $\phi = (\phi^1, \phi^2, \ldots, \phi^m)$ be a random vector, and let $f : \Re^m \to \Re$ be a "well behaved" function (e.g. uniformly continuous) that is nevertheless time consuming to evaluate. We wish to estimate $\mu = E(f(\phi))$. Choose $\psi_j \in \Re^m,\ j = 1, 2, \ldots, M$ and compute and store the values of $f(\psi_j),\ j = 1, 2, \ldots, M$. Let $\phi_1, \phi_2, \ldots$ be i.i.d. replicates of ϕ and define $\pi_i = \text{argmin}_{j\in\{1,2,\ldots,M\}}\|\phi_i - \psi_j\|$ to be the index of the element of $\Psi \equiv \{\psi_1, \psi_2, \ldots, \psi_M\}$ closest to ϕ_i, with ties going to the choice with the smallest index. Simulation A estimates μ by straight Monte Carlo. Let N_t be the number of replicates of $f(\phi_i)$ the simulation generates up to time t. Then simulation A provides

$$\overline{X}_t = \frac{1}{N_t}\sum_{i=1}^{N_t} f(\phi_i)$$

and

$$\sigma_X^2(t) = \frac{t}{N_t}\left(\frac{1}{N_t}\sum_{i=1}^{N_t} f(\phi_i)^2 - \overline{X}_t^2\right).$$

Simulation B estimates μ and $\mu_Y = E(f(\psi_{\pi_i}))$ along with the asymptotic covariance matrix by

$$(\tilde{X}_t, \tilde{Y}_t) = \frac{1}{M_t}\sum_{i=1}^{M_t}(f(\phi_i), f(\psi_{\pi_i})), \tag{4.1}$$

and

$$\begin{pmatrix} \tilde{\sigma}_X^2(t) & \tilde{\sigma}_{XY}(t) \\ \tilde{\sigma}_{XY}(t) & \tilde{\sigma}_Y^2(t) \end{pmatrix} = \frac{t}{M_t}\left(\frac{1}{M_t}\sum_{i=1}^{M_t}\begin{bmatrix} f(\phi_i) \\ f(\psi_{\pi_i}) \end{bmatrix}[f(\phi_i), f(\psi_{\pi_i})] - \begin{pmatrix} \tilde{X}_t^2 & \tilde{X}_t\tilde{Y}_t \\ \tilde{X}_t\tilde{Y}_t & \tilde{Y}_t^2 \end{pmatrix}\right),$$

where M_t is the number of pairs $(f(\phi_i), f(\psi_{\pi_i}))$ simulation B generates up to time t.

Simulation C estimates μ_Y and its AVP by

$$\overline{Y}_t = \frac{1}{K_t} \sum_{i=1}^{K_t} f(\psi_{\pi_i}) \tag{4.2}$$

and

$$\sigma_Y^2(t) = \frac{t}{K_t} \left(\frac{1}{K_t} \sum_{i=1}^{K_t} f(\psi_{\pi_i}) - \overline{Y}_t^2 \right),$$

where K_t is the number of replicates of $f(\psi_{\pi_i})$ that are generated up to time t. Since $\{f(\psi_j)\},\ j = 1, 2, \ldots, M$ is evaluated before the main experiment starts, there is no need to evaluate $f(\cdot)$ in simulation C. Evaluating π_i involves finding the closest element of Ψ to ϕ_i, which has complexity at most $\mathcal{O}(M)$, and is therefore fast (unless M is too big). Thus, r^2 is small. Also, if f is fairly smooth and Ψ is "representative" of ϕ then $f(\phi_i)$ and $f(\psi_{\pi_i})$ will be highly correlated. (A reasonable way to choose Ψ is to generate M independent replicates of ϕ.) We therefore expect the QCV estimator (3.8) or (3.14) to be far superior to $\overline{X}_t$.

Stochastic Linear Programs: Let A be an $m \times n$ matrix and let $c \in \Re^n$. For $b \in \Re^m$, let

$$f(b) = \min_{x \in \Re^n} \{cx \ : \ Ax \geq b,\ x \geq 0\} \tag{4.3}$$

be the solution of the linear program specified by A, c, and "right hand side", b. We will consider the *stochastic* linear program seeking

$$\mu = E(f(\phi)),$$

where $\phi = (\phi^1, \phi^2, \ldots, \phi^m)'$ is a random (right hand side) vector. (These techniques apply to more general stochastic linear programs with random constraint matrices and cost coefficients [6].) Since evaluating the linear program is time consuming, we have a variant of our generic example. Choose a set of vectors $\Psi = \{\psi_1, \psi_2, \ldots, \psi_M\}$ as a set representative of ϕ, and evaluate $f(\psi_j),\ j = 1, 2, \ldots, M$. Let $\{\phi_1, \phi_2, \ldots\}$ be i.i.d. replicates of ϕ, and let $\pi_i \in \{1, 2, \ldots, M\}$ be the index of the element of Ψ closest to ϕ_i. Then $(\tilde{X}_t, \tilde{Y}_t)$ given by (4.1) should be a highly correlated pair (if Ψ is well chosen). Simulation C, which generates $\overline{Y}_t$ given by (4.2) is very efficient since (4.3) does not have to be solved. We can consider the QCV procedure just described as the "primal" approach to the stochastic linear programming problem. In [6], Emsermann develops a "dual" QCV procedure for this problem that has some clear advantages over the primal approach. Both the primal and dual approaches reduce the AVP significantly in the cases studied because the correlation $\tilde{\rho}$ between $f(\phi_i)$ and $f(\psi_{\pi_i})$ is large when Ψ is chosen to be $M \approx 20$ independent replicates of ϕ, and $M \approx 20$ is small enough so that r^2 is small.
For example, the method can be applied to problems in power system reliability evaluation [7]. Here the linear program models the power network equations and constraints; the stochastic right hand side represents possible failures in the system. The objective is to minimize the interruption of power supply (load curtailment). The most popular reliability indices are *expected power not supplied* (EPNS) and *loss of load probability* (LOLP). In the context of (4.3), ϕ would represent random generator and circuit outages causing overloads

by disrupting the flow of power. These overloads can often be eliminated by rescheduling the system generators and $f(\phi)$ would be the resulting total minimized load curtailed. Thus, in this case μ represents the resulting EPNS. A related power performance measure, LOLP, is the probability of energy not being supplied and is defined as $P(f(\phi) > 0)$, so $\mu = E(1_{f(\phi)>0})$. In general the LOLP is close to zero, making it difficult to estimate accurately, so increasing simulation efficiency is crucial in this application.

Stochastic PDE's: Consider a partial differential equation where one or more of the coefficient functions are not known exactly, so only statistical statements can be made about their values. It is often reasonable to model the set of unknown coefficient functions as a (single) random function, ϕ, that matches or approximates the statistical properties of the coefficient functions. We call the resulting equation a stochastic PDE. The solution, u, is random too, so the goal becomes estimating $\mu = E(\Phi(u))$, where $\Phi(u)$ is a scalar quantity of interest associated with the solution. One can estimate μ by generating independent replicates of the random coefficient functions, $\{\phi_1, \phi_2, \ldots\}$, and then solving (by deterministic numerical methods) the resulting sequence of PDE's. Thus, simulation A forms the estimator

$$\overline{X}_t = \frac{1}{N_t} \sum_{i=1}^{N_t} \Phi(u_i), \tag{4.4}$$

where u_i is the solution of the ith equation. Since the numerical solution of a PDE can be time consuming, (4.4) may not be an efficient estimator for μ. Let $(\mathcal{H}, d(\cdot,\cdot))$ be a metric space for the set of unknown coefficient functions and let $\Psi = \{\psi_1, \psi_2 \ldots, \psi_M\}$ be a set of representative elements of $\mathcal{H}$. Again, we can choose Ψ to be M independent replicates of the random coefficient functions. We evaluate $\Phi(\hat{u}_j),\ j = 1, 2, \ldots, M$, where $\hat{u}_j$ is the solution of the jth representative equation. Let $\pi_i \in \{1, 2, \ldots, M\}$ be the index that minimizes $d(\phi_i, \psi_j)$, and define

$$\mu_Y = E(\Phi(\hat{u}_{\pi_i})).$$

If $d(\cdot,\cdot)$ is easy to compute relative to solving the PDE then r^2 will be small. Thus, if $\Phi(u_i)$ and $\Phi(\hat{u}_{\pi_i})$ are well correlated then a QCV procedure will improve the efficiency of the Monte Carlo experiment. In this application it is crucial to find a metric that will induce a significant correlation; i.e., the map $\Phi(u) : \mathcal{H} \to \Re$ is smooth in that metric. We point out that this is another variant of the generic example; the random vector is now a random function.

As an example, the equations describing the flow of a solute in a porous medium (e.g., some liquid contaminant in the ground) are

$$\frac{\partial c}{\partial t} + \frac{\partial}{\partial x_i}(v_i c) - \frac{\partial}{\partial x_i}\left[\Gamma_{ij}\frac{\partial c}{\partial x_j}\right] = 0 \quad x \in D$$

$$c = c_b \quad x \in \partial D$$

$$c = c_0 \quad x \in D,\ t = 0$$

where $c \equiv c(x,t)$ is the concentration of the solute at location x at time t, D is the domain of interest, and summation over repeated indices is implied [8]. The coefficient functions $v_i(x)$ ("pore velocities") are unknown functions of the medium that are modeled as time

invariant nonstationary random fields. The initial concentration $c_0(x)$ may also be considered random. In this context, one might be interested in the expected center of mass of the contaminant as a function of time, or the expected time until a significant fraction of the contaminant reaches some sensitive region.

Queueing Theory: Here, we describe a quasi (external) control variate procedure for simulating complex queueing models, analogous to the applications described in [4] and [5]. The goal is to estimate some quantity μ associated with the complex model. An external control variate in this case typically takes the form of the analogous quantity $\hat{\mu}$ in an analogous but simplified model, where the quantity can be computed exactly. The two models are "synchronized", and then "common random numbers" are used for the two models in the hope of inducing a correlation between them, e.g., [2]. It is usually difficult to implement external control procedures successfully in practice. There are two reasons: First, simulating the second system adds a significant overhead to simulation B, i.e., we could have $\tilde{\sigma}_X^2 \approx 2\sigma_X^2$. Also, when drastic simplification of the original model is necessary in order to compute $\hat{\mu}$ it becomes very difficult to keep the two simulations synchronized, and the correlation suffers as a consequence. (Designing simulation B becomes very difficult too.) There is very little that can be done about the first problem, but relaxing the requirement that $\hat{\mu}$ must be known exactly can help alleviate the second problem. By allowing a wider choice of external control models this way, one can hope to find one where the synchronization problem is less serious, and therefore a larger correlation between the two systems is possible. Of course since an exact expression for $\hat{\mu}$ is not available the control system becomes a *quasi* control system. In order for a QCV procedure to work well in this setting, the simulation of the control model must be much more efficient than the simulation of the primary model. If the control model is Markovian, we can use the *approximating Markov process procedure* [9], to significantly speed up its simulation, and therefore successfully use a QCV procedure, if a good approximation for $\hat{\mu}$ is available. The approximation could come from (say) heavy traffic asymptotics, which are available for a much wider class of queueing systems than those with complete solutions, or any other method.
To illustrate, one might choose to model a wireless phone system as a network of flow controlled infinite server stations with Poisson arrivals and general service time distributions. Drastic simplification of this model is necessary to obtain an exact solution (e.g., some product form network). However, it is possible to construct heavy traffic limits for networks of flow controlled infinite server stations with Poisson arrivals and phase–type service times [10]. Such a system can be made very "close" to the desired model and may therefore serve as a good control. The heavy traffic limits yield approximations that allow the approximating Markov process procedure to be employed, making a QCV procedure feasible.

5 SUMMARY

We assume that three simulation programs are available to an experimentor who wishes to estimate an unknown scalar quantity of interest:

Simulation A: A direct simulation of the quantity of interest,

Simulation B: A simulation of the quantity of interest and a "control variate", and

Simulation C: A direct simulation of the control variate.

We measure the efficiency of a simulation by its asymptotic variance parameter (AVP), described in Definition 1.1. The goal is to find an estimator of the quantity of interest using simulations A, B and C that has the lowest possible AVP. The AVP for simulation A is the benchmark for comparisons with more sophisticated estimators using simulations B and C. The classical control variate procedure uses simulation B to estimate the quantity of interest via equation (1.7), although it is possible that simulation A has a lower AVP and is therefore (asymptotically) preferable. A quasi control variate (QCV) procedure also utilizes simulation B, but requires simulation C as well since the mean of the (quasi) control variate is unknown. The (static) QCV estimator has the form (1.14). Again, simulation A may have a lower AVP than even the optimal QCV procedure. The crucial variables in determining how good a QCV procedure can be are

r^2, the ratio of the AVP's for the control variate estimators in simulations B and C, given by (1.13), and

$\tilde{\rho}^2$, the asymptotic (squared) correlation coefficient between the primary and control variates in simulation B, given by (1.8).

The optimal static QCV procedure is specified by

p^*, the fraction of time devoted to simulation C, and

α^*, the weighting factor for the control variate term of the QCV estimator.

The values of p^* and α^* are given in Theorem 2.1. The AVP for the optimal QCV procedure is given in Corollary 2.1.1.
In practice one must estimate p^* and α^*, and consider the possibility that simulation A has a lower AVP than the optimal static QCV procedure. A sufficient condition that simulation A has a lower AVP is $\tilde{\rho}^2 \leq r^2$. Our recommended QCV procedure is a dynamic procedure that "evolves" into simulation A if the AVP for simulation A is lower than the AVP for the optimal static QCV procedure, and otherwise evolves into the optimal static QCV procedure utilizing simulations B and C. The situation is more complicated if the alternatives are precisely equal, i.e., $\Delta = 0$, but in practice this is unlikely. The recommended procedure is specified by (3.2) - (3.8) with $\epsilon = 0$. Theorem 3.1 shows that the recommended procedure is optimal in the sense that it achieves the lowest possible AVP for any simulation experiment utilizing simulations A, B and C as long as $\Delta \neq 0$. If $\Delta \approx 0$, Theorem 3.1 shows that for any $\epsilon > 0$ we can implement a dynamic QCV procedure that is "ϵ–optimal".

REFERENCES

[**1**] Glynn, P.W.; Whitt, W. The Asymptotic Efficiency of Simulation Estimators. Operations Research **1992**, 40, 505-520.

[**2**] Law, A.M.; Kelton, W.D. Variance Reduction–Techniques. *Simulation Modeling and Analysis*, 2nd Ed.; McGraw Hill: New York, **1991**.

[**3**] Billingsley, P. *Convergence of Probability Measures*; John Wiley & Sons: New York, **1968**.

[**4**] Schmeiser, B.W.; Taaffe, M.R.; Wang, J. Control–Variate Estimation Using Estimated Control Means, **2000**, to appear in IIE Transactions.

[**5**] Schmeiser, B.W.; Taaffe, M.R.; Wang, J. Biased Control–Variate Estimation. IIE Transactions **2001**, 33(3), 219-229.

[**6**] Emsermann, M. *Variance Reduction with Quasi Control Variates*, Ph.D. Thesis, Department of Mathematics, University of Colorado at Denver, Denver, CO., **2000**.

[**7**] Pereira, M.V.F.; Pinto, L.M.V.G. A New Computational Tool for Composite Reliability Evaluation. IEEE Transactions on Power Systems **1992**, 7 (1), 258-264.

[**8**] Graham, W.; McLaughlin, D. Stochastic Analysis of Nonstationary Subsurface Solute Transport, 1. Unconditional Moments. Water Resources Research **1989**, 25 (2), 215-232.

[**9**] Henderson, S.G.; Glynn, P.W. Approximating Martingales for Variance Reduction **2001**, to appear in Mathematics of Operations Research.

[**10**] Fleming, P.J.; Simon, B. Diffusion Approximations of an Infinite Server Network with Load Balancing. Probability in the Engineering and Informational Sciences **1999**, 13, 251-273.