RESOURCE ALLOCATION AMONG SIMULATION TIME STEPS

PAUL GLASSERMAN

Graduate School of Business, Columbia University, New York, New York 10027, pg20@columbia.edu

JEREMY STAUM

Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208, staum@iems.nwu.edu

Motivated by the problem of efficient estimation of expected cumulative rewards or cashflows, this paper proposes and analyzes a variance reduction technique for estimating the expectation of the sum of sequentially simulated random variables. In some applications, simulation effort is of greater value when applied to early time steps rather than shared equally among all time steps; this occurs, for example, when discounting renders immediate rewards or cashflows more important than those in the future. This suggests that deliberately stopping some paths early may improve efficiency. We formulate and solve the problem of optimal allocation of resources to time horizons with the objective of minimizing variance subject to a cost constraint. The solution has a simple characterization in terms of the convex hull of points defined by the covariance matrix of the cashflows. We also develop two ways to enhance variance reduction through early stopping. One takes advantage of the statistical theory of missing data. The other redistributes the cumulative sum to make optimal use of early stopping.

(*Simulation, efficiency: variance reduction. Statistics, design of experiments: data missing by design. Finance, asset pricing: computational methods.*)

Received May 2001; revisions received April 2002, September 2002; accepted October 2002.
Area of review: Simulation.

1. INTRODUCTION

In a standard finite-horizon simulation, every path reaches every time step. This is an implicit allocation of computational resources among the time steps. This paper proposes and analyzes a variance reduction technique that sets the number of paths reaching each time step optimally according to the solution of a resource allocation problem.

This investigation is motivated by problems of estimating cumulative discounted rewards or cashflows. In various applications, cashflows resulting from business operations or a financial security are potentially generated at each time step; their timing and magnitude are determined by an underlying stochastic model of physical operations or market prices. Mortgage-backed securities provide a specific illustration. The cashflows of these securities result from payments from a pool of mortgages, and valuing such a security entails calculating the expected present values of these payments. Valuing a security backed by 30-year mortages making monthly payments nominally requires simulating 360 time steps; but most of the value of the security is typically determined over a far shorter horizon. Cashflows occurring late in the security's life are less important because of two effects: discounting, which reduces the present value of money received in the future, and the homeowners' right to prepay their mortgages, which reduces the nominal amount of mortgage payments. Discounting is favorable for our method, while prepayments may be favorable or unfavorable, depending on the effect they have on the covariances among cashflows, as we explain in §§6.1–6.2.

In §2 we pose the problem abstractly. To summarize briefly, our goal is efficient estimation of the expected value of a finite sum of sequentially generated random variables. The summands are correlated—for example, they may be functions of the state of a Markov chain. Simulating each summand entails simulating all previous summands. The length of a path is the number of summands simulated. We plan in advance to simulate a fixed number of paths of each length; by appropriately weighting paths of different lengths, we obtain an unbiased estimator. We find the optimal number of paths of each length to simulate by minimizing variance subject to a computational budget constraint and a monotonicity constraint requiring that the number of paths reaching each horizon decreases as the horizon increases. The optimal allocation has a simple characterization in terms of the convex hull of points that reflect the contributions of each step to total work and variance. Section 3 makes this precise.

This method may be interpreted as *stratification in time.* Indeed, our optimal allocation bears a superficial resemblance to the optimal allocation for stratified sampling. However, whereas in stratified sampling different strata can ordinarily be sampled independently, in our setting later time steps can be sampled only if earlier time steps are sampled, too.

The idea of varying the time horizon of paths in simulation recalls Fox and Glynn's (1989) investigation of estimation of expected reward over infinite time horizons. However, their work uses random stopping to remove bias from truncating an infinite integral, as does Asmussen

(1990), while we consider a strategy of deterministic stopping optimized to reduce estimator variance.

The variables not generated when we stop a path early may be viewed as missing data. This interpretation suggests the possibility of improved output analysis using missing data techniques, as in Little and Rubin (1987, §6.5). With these techniques, Hocking and Smith (1972) find the cost-minimizing experimental design for bivariate data subject to simultaneous constraints on the variance of estimators of six parameters: means, variances, and regression coefficients. In §4, we show how the theory of missing data can be applied to reduce variance in our more general setting and how to solve the resource allocation problem when estimation employs missing data techniques.

Estimators resulting from the application of missing data techniques may be interpreted as redistributing cashflows across time steps. This suggests consideration of other estimators that produce "fictitious cashflows" at intermediate time steps while keeping the sum constant. Section 5 analyzes this case and finds the optimal fictitious cashflows for early stopping.

A discussion of the effectiveness of the techniques developed in this paper appears in §6. Section 7 concludes the paper. All proofs are deferred to Appendix A, and Appendix B contains an algorithm for solving the resource allocation problem.

2. THE SIMPLE EARLY STOPPING PROBLEM

We seek to estimate, by simulation, the expectation of the sum X of m discounted cashflows $X_1, \ldots, X_m$ having finite, positive variance. Each X_k is a function of the history $\{S_1, \ldots, S_k\}$, with S_k denoting the state of the simulation at time step k. The cost of generating S_k and X_k given the process up to step $k-1$ is a constant c_k, and the total computational budget is C. We plan in advance to simulate X_k only on paths $1, \ldots, n_k$, where the n_k satisfy the monotonicity constraint

$$n_k \geqslant n_{k+1}, \qquad k = 1, \ldots, m-1, \tag{1}$$

and the budget constraint

$$\sum_{k=1}^{m} c_k n_k \leqslant C. \tag{2}$$

Define the feasible set $\mathcal{N}$ as the set of all vectors $\mathbf{n} = (n_1, \ldots, n_m)$ satisfying these two constraints.

There are $n = \max\{n_k\} = n_1$ paths. Let X_{ik} be the value of X_k on path i. The simulation generates this value if and only if $i \leqslant n_k$. This simulation structure also allows us to define the length of the ith path $m_i = \max\{k \mid n_k \geqslant i\}$, so that X_{ik} is observed (i.e., simulated) if and only if $k \leqslant m_i$, and $m = \max\{m_i\} = m_1$ by the monotonicity constraint (1). Let $\mathbf{m}$ be the vector $(m_1, \ldots, m_n)$. The simulation produces an $n \times m$ matrix of data $\mathbf{X}$ whose ith row (the path i) is observed up to column m_i and then unobserved, and whose kth column (the random variable X_k) is observed up to row n_k and then unobserved.

EXAMPLE. When the vector $\mathbf{n} = (4, 2, 1)$, then the vector $\mathbf{m} = (3, 2, 1, 1)$ and the matrix $\mathbf{X}$ is

$$\begin{bmatrix} X_{11} & X_{12} & X_{13} \\ X_{21} & X_{22} & \cdot \\ X_{31} & \cdot & \cdot \\ X_{41} & \cdot & \cdot \end{bmatrix}.$$

This structure for simulation is equivalent to stipulating that if X_k is generated on path i, so were $X_1, \ldots, X_{k-1}$ and hence necessarily $S_1, \ldots, S_{k-1}$. Then we say that for $k < l$, X_k is more observed than X_l. This is necessary in simulation when the process is specified in terms of transition probabilities, so that generating S_k requires knowing S_{k-1}.

Now define $\overline{X}_{kl} \triangleq (1/n_l)\sum_{i=1}^{n_l} X_{ik}$, the average of X_k on those paths where X_l is observed. For this to make sense, we must have $k \leqslant l$ so that X_k is at least as observed as X_l. The matrix $\overline{\mathbf{X}}$ is an upper-triangular $m \times m$ matrix.

EXAMPLE. Continuing the previous example, the matrix $\overline{\mathbf{X}}$ is

$$\begin{bmatrix} (X_{11}+X_{21}+X_{31}+X_{41})/4 & (X_{11}+X_{21})/2 & X_{11} \\ \cdot & (X_{12}+X_{22})/2 & X_{12} \\ \cdot & \cdot & X_{13} \end{bmatrix}.$$

In this simple setting, we estimate $\mu_k \triangleq \mathbf{E}[X_k]$ by $\hat{\mu}_k = \overline{X}_{kk}$ and $\mu \triangleq \sum_{k=1}^{m} \mu_k$ by $\hat{\mu} = \sum_{k=1}^{m} \hat{\mu}_k$. Finding the experimental design $\mathbf{n} = (n_1, \ldots, n_m)$ that minimizes variance given the fixed computational budget is a resource allocation problem with constraints (1) and (2). We ignore the issue that the n_i must actually be integral; because they will be large in reasonable applications, rounding errors are likely to be insignificant. We suggest rounding up for safety.

To solve the resource allocation problem, we first give an explicit expression for the objective.

LEMMA 1. *The variance objective is*

$$\mathbf{Var}[\hat{\mu}] = \sum_{k=1}^{m} \frac{v_k}{n_k}, \tag{3}$$

where $\sigma_{kl} = \mathbf{Cov}[X_k, X_l]$ form the covariance matrix $\boldsymbol{\Sigma}$ and

$$v_k = \sigma_{kk} + 2\sum_{l=k+1}^{m} \sigma_{kl} = \mathbf{Var}\left[\sum_{l=k}^{m} X_l\right] - \mathbf{Var}\left[\sum_{l=k+1}^{m} X_l\right]. \tag{4}$$

The resource allocation objective is then

$$\min_{n \in \mathcal{N}} \sum_{k=1}^{m} \frac{v_k}{n_k}. \tag{5}$$

3. SOLUTION OF THE RESOURCE ALLOCATION PROBLEM

The variance component v_k defined in Equation (4) is that part of the variance of $\sum_{k=1}^{m} X_k$ attributable to step k, much as c_k is the component of a complete path's cost attributable to simulating step k. Because of the monotonicity constraint (1), doing a unit of work at step k requires doing a unit of work at all steps $j \leqslant k$, so we are also interested in partial sums of the variance components v_k and cost components c_k. Define the partial sums

$$V_k \triangleq \sum_{j=1}^{k} v_j \quad \text{and} \quad C_k \triangleq \sum_{j=1}^{k} c_j$$

and the point set

$$\mathcal{V} \triangleq \{(C_k, V_k) \mid k = 0, \ldots, m\}.$$

The tail sum $V_m - V_k = \sum_{l=k+1}^{m} v_l = \mathbf{Var}[\sum_{l=k+1}^{m} X_l]$ is the variance of the last $m-k$ discounted cashflows. Likewise, $C_m - C_k$ is the cost of simulating the last $m-k$ steps.

We assume that each tail variance $V_m - V_k$ is finite and strictly positive. If for some k it were zero, then it would be proper to end the simulation at step k, because $\sum_{l=k+1}^{m} X_l$ would be a constant. This is equivalent to setting $n_l = 0$ for $l > k$, which, as one may verify, is what the solution in Theorem 1 yields in this case.

The solution to the resource allocation problem (5) involves the upper convex hull of the graph $\mathcal{V}$ as follows. Let conv($\mathcal{V}$) denote the convex hull of $\mathcal{V}$ and let V^* be the function on $[0, C_m]$ defined by $V^*(x) = \sup\{y\colon (x, y) \in \text{conv}(\mathcal{V})\}$; this is the upper portion of the boundary of the convex hull, connecting $(C_0, V_0) = (0, 0)$ to (C_m, V_m). Also, V^* is a concave function and it is the least concave function that lies above $\mathcal{V}$. Write V_k^* for $V^*(C_k)$ and define v_k^* as the increment $V_k^* - V_{k-1}^*$. Finally, define the slopes

$$u_k \triangleq v_k / c_k \quad \text{and} \quad u_k^* \triangleq v_k^* / c_k.$$

The transformation from v to v^* is illustrated in Figure 1. The upper left panel illustrates a hypothetical v; the upper right panel plots its cumulative sum V; the lower left panel shows the upper convex hull V^* of V; the increments v^* of V^* appear in the lower right panel. We now have Theorem 1.

Theorem 1. *The solution to the resource allocation problem* (5) *is*

$$n_k = \sqrt{\frac{u_k^*}{\nu}}, \tag{6}$$

$$\nu = \left(\frac{1}{C} \sum_{k=1}^{m} c_k \sqrt{u_k^*}\right)^2, \tag{7}$$

and the ratio of optimal variance to standard variance is

$$R = \frac{\sum_{i=1}^{m} \sum_{j=1}^{m} v_i c_j \sqrt{u_j^*/u_i^*}}{\sum_{i=1}^{m} \sum_{j=1}^{m} v_i c_j}. \tag{8}$$

The result parallels the classical result for stratified sampling that the variance of an estimator of a mean is minimized when the number of samples from each stratum is proportional to the square root of the ratio of the variance in that stratum and the cost to sample from that stratum.

Figure 1. Illustration of convex hull solution.

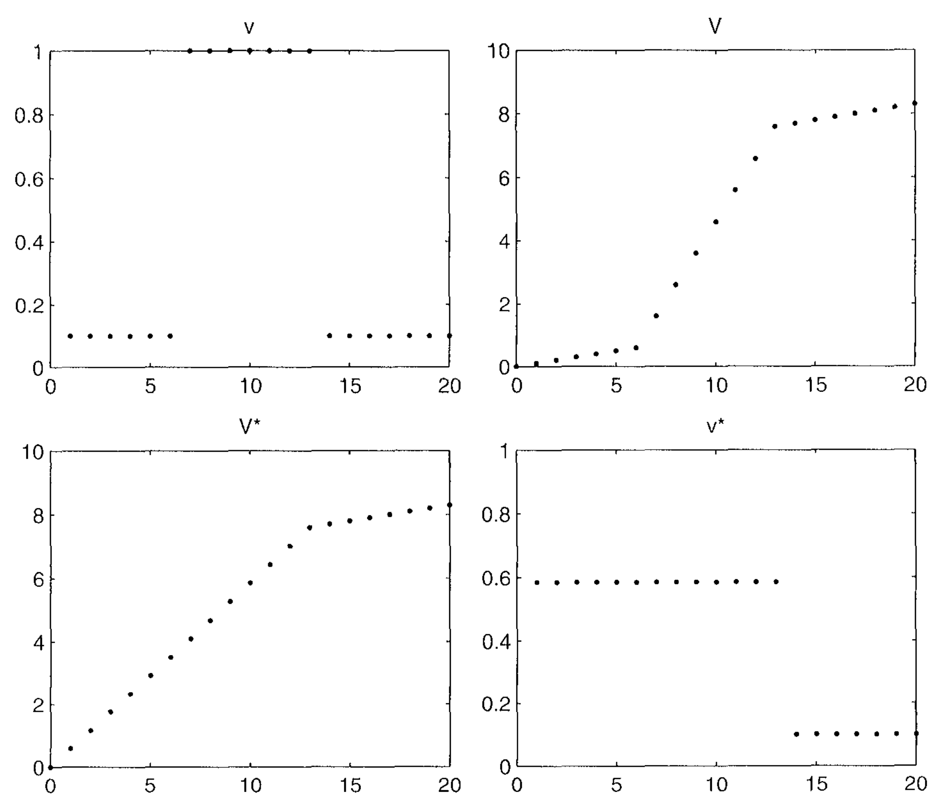

See, for example, Cochran (1953, Theorem 5.7). To make the analogy, interpret the time steps as strata of equal size. The v_k^* are not variances but they play a similar role.

This resource allocation problem is also nearly a special case of a problem that arises in planning production-distribution systems. Maxwell and Muckstadt (1985, p. 1323) pose "problem RP," which is to minimize $\sum_k (K_k/T_k + g_k T_k)$, where the sum is over all nodes k in a directed graph, subject to the constraints $T_j \geqslant T_k$ for each arc (j, k) in the graph. Federgruen and Zheng (1992, Theorem 1) describe the solution of problem RP in greater generality, but the case considered here has special features—in particular, the convex hull characterization—not shared by the general case. Roundy (1986, p. 720) notes connections between the production-distribution planning problem and statistical applications such as isotonic regression and multidimensional scaling and provides references.

The solution to the resource allocation problem depends on the covariance matrix of the variables X_k, which is presumably unknown in realistic situations. However, the covariance matrix can be estimated in pilot runs, as demonstrated numerically in §6. In either case, the design optimal for a good estimate of the covariance matrix may be better than either the standard design where all paths reach all time steps or a guess at a design. Also, we have assumed that the costs c_k are fixed. If the cost of generating random variables at step k is itself random, we could use its expected value as c_k, provided we interpret C as a constraint on the *expected* cost.

The savings due to implementing the scheme should be weighed against the overhead costs of finding the optimal n_k. As expressed in the resource allocation problem, the new scheme is a variance reduction technique, but this is equivalent to a reduction of cost to attain the same variance. If the technique reduces variance to a fraction R of the original variance, then by applying it, one could attain the original variance with a fraction R of the original budget. So such a reduction is equivalent to saving the fraction $1 - R$ of the cost, which is $\mathcal{O}(mn)$.

Finding a planar convex hull is ordinarily $\mathcal{O}(m \log m)$, because of the intimate connection to sorting, as remarked by Preparata and Shamos (1985, p. 94). In our case, the points of $\mathcal{V}$ are already sorted by abscissa C_i. Taking advantage of this and our lack of interest in the lower convex hull, we can use a stripped-down Graham scan which is only $\mathcal{O}(m)$. We present the algorithm in Appendix A; for a discussion of the Graham scan, see Preparata and Shamos (1985, §3.3.2). Estimating the variance components v_k as given in Equation (4) from a pilot run with a fixed number of paths also takes $\mathcal{O}(m)$ work, in computing partial sums of the m random variables $X_1, \dots, X_m$, sample variances of these m partial sums, and differences of these m sample variances. So when the number n of paths planned for a standard simulation is large, it is indeed efficient to use this algorithm.

4. MISSING DATA PERSPECTIVE

If some paths stop early, we can view the cashflows not generated as missing data and use missing data techniques to get an estimator with lower variance. However, the resource allocation problem we solved assumed a particular estimator $\hat{\mu}$. If we know in advance that we plan to use a different estimator, that should be reflected in the allocation problem.

In this section, we first formulate estimators that take advantage of missing data techniques. We then show, conveniently, that the optimal resource allocation problem for these estimators has the same form as the problem solved in §3, but with a different set of variance contributions v_k. Thus, the solution in §3 continues to apply in this more general setting.

As discussed in §2, stopping early on some paths results in a data matrix $\mathbf{X}$ in which X_k is more observed than X_l for $k < l$. Little and Rubin (1987, Example 6.7) consider data sets with this structure, but with the further assumption that the vector $(X_1, \dots, X_m)$ is multivariate normal. Then the maximum likelihood estimate of the mean μ_k is

$$\hat{\mu}_k = \overline{X}_{kk} + \sum_{i=1}^{k-1} \hat{\gamma}_{ik} (\hat{\mu}_i - \overline{X}_{ik}), \tag{9}$$

where $\hat{\gamma}_{ik}$ is the estimated coefficient of X_k on X_i in the regression on $X_1, \dots, X_{k-1}$, using the n_k observations where all of these variables are observed. The resulting estimator of $\mu = \sum_{k=1}^{m} \mu_k$ is the sum $\sum_{k=1}^{m} \hat{\mu}_k$. In our setting, the estimator (9) has the following interpretation: Values of $X_1, \dots, X_{k-1}$ are used to predict missing values of X_k by taking advantage of the relation between $X_1, \dots, X_{k-1}$ and X_k on paths where all are observed.

In the case of simulation, requiring the random variables X_k to be multivariate normal would be too restrictive, but their sample averages should be approximately normal. Define the constants $q_i = n_i/n$, so $q_1 = 1$. As the budget C goes to infinity, the number of paths n goes to infinity, and thus each n_i goes to infinity. Therefore, the joint distribution of the appropriately scaled averages $\overline{X}_{ik}$ converges to the multivariate normal, by the central limit theorem. This large-sample result frees us to use functions of the data $T_h = t_h(S_1, \dots, S_h)$ for $h < k$ as predictors of X_k (rather than just $X_1, \dots, X_{k-1}$), as long as they have finite variance. The vector $(\overline{T}_{1k}, \dots, \overline{T}_{k-1,k}, \overline{X}_{kk})$ of sample averages on the first n_k paths is also approximately multivariate normal, although in general $(T_1, \dots, T_{k-1}, X_k)$ is not. This approximate multivariate normality justifies the use of Little and Rubin's (1987) maximum likelihood estimator for use with missing data.

Using these $T_1, \dots, T_{k-1}$ instead of $X_1, \dots, X_{k-1}$ as predictors for X_k, (9) is replaced with

$$\hat{\mu}_k = \overline{X}_{kk} + \sum_{h=1}^{k-1} \hat{\beta}_{hk} (\hat{\tau}_h - \overline{T}_{hk}), \tag{10}$$

where $\hat{\beta}_{hk}$ is the estimated multiple regression coefficient of X_k on T_h in the regression on $T_1, \dots, T_{k-1}$, while $\hat{\tau}_h$ is a missing-data estimate of the mean $\tau = \mathbf{E}[T_h]$. Expressed with a different subscript, this estimate is

$$\hat{\tau}_k = \overline{T}_{kk} + \sum_{h=1}^{k-1} \hat{b}_{hk}(\hat{\tau}_h - \overline{T}_{hk}), \tag{11}$$

where $\hat{b}_{hk}$ is the estimated coefficient of T_k on T_h in the regression on $T_1, \ldots, T_{k-1}$.

Our objective now is to show that the missing data estimator $\sum_{k=1}^{m} \hat{\mu}_k$ can be put in a form in which the resource allocation of §3 applies. To this end, we define an adjusted data matrix $\mathbf{X}' = \mathbf{X} + \mathbf{TW}$ where $\mathbf{W}$ is a matrix of estimated weights. The matrix $\mathbf{T}$ has typical element T_{ik}, the realization of the predictor T_k on the ith path. As Lemma 2 states, the missing data estimator takes the form of a sum of averages of random variables X'_k. This means that the minimization of this missing data estimator's variance is a resource allocation problem based on $\mathbf{X}'$ rather than $\mathbf{X}$. Let $\mathbf{1}$ and $\mathbf{0}$ be the vectors whose components are respectively 1 and 0. The condition $\mathbf{W1} = \mathbf{0}$ will make the new estimator unbiased, as shown in Theorem 2.

LEMMA 2. *Let $\hat{\mu}_k$ be as given in Equation* (10). *Then,*

$$\sum_{k=1}^{m} \hat{\mu}_k = \hat{\mu} \triangleq \sum_{k=1}^{m} \frac{1}{n_k} \sum_{i=1}^{n_k} X'_{ik}, \tag{12}$$

with $\mathbf{X}' \triangleq \mathbf{X} + \mathbf{TW}$ *for an upper-triangular $m \times m$ matrix* $\mathbf{W}$ *whose elements w_{hi} are given as follows. For $h \leqslant i \leqslant k$, define*

$$\omega_{hik} \triangleq \begin{cases} 1 & \text{if } h = i = k \\ -\hat{b}_{hi} & \text{if } h < i = k \\ \sum_{j=i}^{k-1} \hat{b}_{jk} \omega_{hij} & \text{otherwise} \end{cases} \tag{13}$$

$$w_{hik} \triangleq \begin{cases} 0 & \text{if } h = i = k \\ -\hat{\beta}_{hi} & \text{if } h < i = k \\ \sum_{j=i}^{k-1} \hat{\beta}_{jk} \omega_{hij} & \text{otherwise} \end{cases} \tag{14}$$

$$w_{hi} \triangleq \sum_{k=i}^{m} w_{hik}. \tag{15}$$

This matrix $\mathbf{W}$ *satisfies* $\mathbf{W1} = \mathbf{0}$.

In the definition of the elements of $\mathbf{W}$, the third subscript k indicates that the $\omega_{\cdot\cdot k}$ is a weight used in estimating $\hat{\tau}_k$ and the $w_{\cdot\cdot k}$ are weights in $\hat{\mu}_k$; the w with only two subscripts are weights in the estimate $\hat{\mu}$ of the sum.

Motivated by this case, we henceforth let $\mathbf{W}$ be any upper-triangular $m \times m$ random matrix that satisfies $\mathbf{W1} = \mathbf{0}$. Then

$$\mathbf{X}'\mathbf{1} = \mathbf{X1} + \mathbf{TW1} = \mathbf{X1} + \mathbf{T0} = \mathbf{X1},$$

and thus the adjusted X'_k differ from the discounted cash-flows X_k, but their sums on a complete path are the same.

THEOREM 2. *The estimator $\hat{\mu}$ is unbiased for μ.*

If $\mathbf{W}$ is constant, not random due to dependence on the simulated data, the covariance matrix of $(X'_1, \ldots, X'_m)$ is

$$\mathbf{\Sigma}' = \mathbf{\Sigma} + \mathbf{W}^{\top}\mathbf{\Xi} + \mathbf{\Xi}^{\top}\mathbf{W} + \mathbf{W}^{\top}\mathbf{\Theta}\mathbf{W}, \tag{16}$$

where $\mathbf{\Sigma}$, $\mathbf{\Xi}$, and $\mathbf{\Theta}$ are the matrices defined by having the typical element in row k and column l be respectively $\mathbf{Cov}[X_k, X_l]$, $\mathbf{Cov}[X_k, T_l]$, and $\mathbf{Cov}[T_k, T_l]$, all assumed finite. The following theorem implies that the optimal resource allocation for the estimator $\hat{\mu}$ is given by Theorem 1 with $\mathbf{\Sigma}'$ replacing $\mathbf{\Sigma}$. It also establishes asymptotic normality in order to facilitate the construction of a confidence interval for $\hat{\mu}$.

THEOREM 3. *If* $\mathbf{W}$ *is constant, $\hat{\mu}$ has variance ς^2 given by Equations* (3) *and* (4), *but with* $\mathbf{\Sigma}'$ *replacing* $\mathbf{\Sigma}$. *As the budget C goes to infinity, the distribution of $(\hat{\mu} - \mu)/\varsigma$ converges to standard normal.*

What if $\mathbf{W}$ is not constant? It could depend on the simulated data, in which case we think of it as an estimator of some constant matrix $\mathbf{W}^{\circ}$.

THEOREM 4. *If there is a constant matrix* $\mathbf{W}^{\circ}$ *such that each element w_{hi} of* $\mathbf{W}$ *is a consistent estimator of w°_{hi}, then as the budget C goes to infinity, the distribution of $(\hat{\mu} - \mu)/\varsigma$ converges to standard normal.*

The matrix $\mathbf{W}$ of Lemma 2, based on estimated least-squares regression coefficients, does satisfy this consistency condition with the matrix $\mathbf{W}^{\circ}$ of true regression coefficients. Thus, through this result, the solution to the resource allocation problem in §3 applies to the missing data estimator.

This general missing data procedure requires computation of $\mathcal{O}(m^3)$ coefficients w_{hik}. However, restricted models with fewer coefficients to estimate can save computation, reduce estimation error, and make implementation easier. The key to the effective use of this variance reduction technique is to use knowledge about the model to find predictors which are known as early as possible in the simulation and explain as much as possible of the variance in payoffs. Such knowledge may suggest that it is unnecessary to allow for a completely general relationship between all predictors and payoffs.

For instance, in the important special case when T_k is Markov, $b_{hk} = 0$ for $h < k - 1$ and should not be estimated. If, furthermore, X_k is a function of T_k only, likewise $\beta_{hk} = 0$ for $h < k - 1$. Plugging into Equations (13)–(15) and expanding the recursion, in this case the elements of $\mathbf{W}$ are

$$w_{hi} = \begin{cases} \sum_{k=i+1}^{m} \hat{\beta}_{k-1,k} \prod_{j=i+1}^{k-1} \hat{b}_{j-1,j} & \text{if } h = i, \\ -\hat{\beta}_{hi} - \hat{b}_{hi} \sum_{k=i+1}^{m} \hat{\beta}_{k-1k} \prod_{j=i+1}^{k-1} \hat{b}_{j-1,j} & \text{if } h < i. \end{cases} \tag{17}$$

To compute these takes only $\mathcal{O}(m^2)$ work because component sums and products can be reused.

This use of missing data techniques to produce improved estimates of the means μ_k bears a strong resemblance to control variates. In the above setting, the predictor mean $\tau_k = \mathbf{E}[T_k]$ is unknown. If each T_k is a control variate for X_k, τ_k is known, and $X'_k = X_k + \hat{\beta}_k(\tau_k - T_k)$, where $\hat{\beta}_k$ is the estimated simple regression coefficient of X_k on T_k. Again, the covariance matrix relevant to resource allocation is that of $(X'_1, \ldots, X'_m)$. A key distinction is that whereas the control variate method relies on $\mathbf{E}[T_k] - T_k$ having mean zero, the missing data estimator relies on the property (embodied in the condition $\mathbf{W1} = \mathbf{0}$) that each T_k gets a net weight of zero when we take the expectation of $\mathbf{X}'$. This is possible only because different paths stop at different times in simulation with early stopping, and would not be possible in standard simulation.

5. FICTITIOUS CASHFLOWS

The effect of using predictors of future cashflows is to replace the actual step-k cashflow X_k with X'_k, which we may interpret as a *fictitious* cashflow at that time. This suggests consideration of other estimators based on redistributing cumulative cashflows across time steps. More precisely, we consider sequences $\{X'_1, \ldots, X'_m\}$ having two properties. First, X'_k is measurable with respect to $\mathcal{F}_k$, the sigma-algebra generated by the state vectors $S_1, \ldots, S_k$: that is, it is actually known at step k of the simulation. Second, the sum on a complete path stays the same: $\sum_{k=1}^{m} X'_k = \sum_{k=1}^{m} X_k \triangleq X$. We now have $\mathbf{X}' = \mathbf{X} + \mathbf{TW}$ after the fashion of §4, with $\mathbf{W}$ equal to the identity matrix and the random variable $T_k = X'_k - X_k$. Even though $\mathbf{W1} \neq 0$, still $\mathbf{X}'\mathbf{1} = \mathbf{X1}$, and Theorem 3 applies. Let $\mathcal{X}$ be the set of random vectors $(X'_1, \ldots, X'_m)$ satisfying the two properties of summation to X and adaptation to the filtration $\{\mathcal{F}_k\}$. What is the optimal $(X'_1, \ldots, X'_m)$ in $\mathcal{X}$ for resource allocation?

The significance of fictitious cashflows is that the covariance matrices $\mathbf{\Sigma}$ of $(X_1, \ldots, X_m)$ and $\mathbf{\Sigma}'$ of $(X'_1, \ldots, X'_m)$ generally differ even though the X_k and X'_k have the same sum. Let v'_k and V'_k be respectively the variance components, derived from $\mathbf{\Sigma}'$ according to Equation (4), and their partial sums. Because the sums $\sum_{k=1}^{m} X_k \triangleq X$ and $\sum_{k=1}^{m} X'_k$ are equal, their respective variances V_m and V'_m are equal, but the intermediate partial sums V_k and V'_k may be unequal. As a consequence, the variance achieved after stopping early may be different. The next result identifies the best set of fictitious cashflows, as measured by the remaining variance after we apply optimal early stopping.

THEOREM 5. *Given* $X = \sum_{k=1}^{m} X_k$, *the optimal sequence of random variables*

$$\arg \min_{\{X'_k\} \in \mathcal{X}} \min_{n \in \mathcal{N}} \sum_{k=1}^{m} \frac{v'_k}{n_k}$$

is given by $X'_k = \mathbf{E}[X \mid \mathcal{F}_k] - \mathbf{E}[X \mid \mathcal{F}_{k-1}]$, $k = 1, \ldots, m$; *i.e.,* $\sum_{j=1}^{k} X'_j = \mathbf{E}[X \mid \mathcal{F}_k]$.

The optimal fictitious cashflows are thus the martingale differences associated with the martingale $\mathbf{E}[X \mid \mathcal{F}_k]$, $k = 1, \ldots, m$. Of course, the conditional expectations $\mathbf{E}[X \mid \mathcal{F}_k]$ are presumably unknown. However, the proof of the theorem shows that early stopping is most effective with fictitious cashflows whose partial sums get as close as possible to these conditional expectations, because this reduces the variance $\mathbf{Var}[X - \sum_{j=1}^{k} X'_j]$ attributable to steps after k. Even if $\sum_{j=1}^{k} X'_j$ is biased for $\mathbf{E}[X \mid \mathcal{F}_k]$, Theorem 2 applies, and the estimator under consideration is unbiased for $\mathbf{E}[X] \triangleq \mu$. We benefit as long as $\sum_{j=1}^{k} X'_j$ is a better estimator of $\mathbf{E}[X \mid \mathcal{F}_k]$ than $\sum_{j=1}^{k} X_j$ is. Examples of fictitious cashflows designed to resemble unknown conditional expectations appear in §6.

We can re-interpret the approach of §4 in this light. We want to reduce $\mathbf{Var}[X - \sum_{j=1}^{k} X'_j]$, but we do not know $\mathbf{E}[X \mid \mathcal{F}_k] = \sum_{j=1}^{k} X_k + \mathbf{E}[\sum_{l=k+1}^{m} X_l \mid \mathcal{F}_k]$. We do know $\sum_{j=1}^{k} X_j$, and we find the projection of $\sum_{l=k+1}^{m} X_l$ onto the space of random variables linear in the predictors $T_1, \ldots, T_k$. The primary attraction of this approach is the possibility of estimating the regression coefficients involved in the linear projection; for success it requires only a good choice of linear predictors, not extensive knowledge of the conditional expectations. The other embellishment is the substitution of missing-data estimators $\hat{\tau}_k$ for sample averages $\overline{T}_k$ in the final estimator of the mean.

6. EFFECTIVENESS

Here we will offer some numerical examples that illustrate the dependence of variance reduction on the particular problem. The goal is to give some guidance about when the variance reduction technique is effective and what kind of missing data predictors and fictitious cashflows are practical. First, we discuss more abstractly the characteristics of a problem that make this method effective. Effectiveness consists in achieving a low ratio R of optimal variance to standard variance as given in Equation (8), possibly after introducing fictitious cashflows or missing data techniques. The infimum of this ratio R over all problems with m steps of equal cost is $1/m$, as follows. Consider a sequence of such problems for which the proportion of the total variance that is associated with the first step, $v_1 / \sum_{k=1}^{m} v_k$, goes to 0. Then $n_1 \to C$ and $n_k \to 0$ for $k > 1$, that is, almost all resources are optimally allocated to the first step. For instance, if we consider a sequence of problems with identical cashflows and a discount rate going to infinity, steps after the first become negligible compared to the first. With a sufficiently high discount rate, the ratio R can be made arbitrarily close to $1/m$ for any $X_1, \ldots, X_m$. Even better variance reduction is possible if the costs are not constant, as the first step becomes cheaper relative to the rest of the problem. In most of the examples we have tried, the variance reduction achieved is modest, but the method is not difficult to apply.

However, suppose that the number of steps m is not small, and the first step is not of overwhelming importance. Then m is not very important at all. The degree of variance reduction is determined primarily by the shape of the graph $\mathcal{V}$. A sequence of problems whose size m increases and whose cost-variance shape converges also has convergence of resource allocation shape and variance objective value. For a precise statement and proof of this assertion, see Glasserman and Staum (2001).

In the numerical examples that follow, for the sake of transparency, each step has unit computational cost $c_k = 1$ and the budget is m. Then standard simulation has resource allocation $n_k = 1$ for each step. (Of course, any reasonable budget C would be much larger, and then the standard resource allocation would be $n_k = C/m$. We are simply standardizing $C = m$.) The results use the optimal n_k computed from the true covariance matrix, ignoring deviation due to pilot estimation of the covariance matrix and rounding the n_k. All results are reported for standard simulations of $n = 10{,}000$ paths.

We also measure costs solely in terms of the numbers of steps n_k, disregarding any extra costs incurred in computing predictors or fictitious cashflows from the state vector and true cashflows. Assessment of such costs, while of importance in practice, would be dependent on implementation of numerical algorithms, which is unrelated to the variance reduction method under consideration. For instance, in the example of §6.3, simulating the state vector involves generating a normal random variable and taking its exponential, while computing the fictitious cashflows requires two evaluations of the standard normal cumulative distribution function. These costs depend on the desired accuracy of the approximate algorithms. The issue of the overhead in computing predictors is essentially the same as the overhead in computing control variates.

Another issue is the overhead involved in computing the convex hull solution. The total time taken by our algorithm is a function of p, the fraction of total work used in the pilot run to estimate variances. Let $R(p)$ be the variance ratio discussed above, but based on the random realization of sample variance components from the pilot run. It tends to decrease in p as better estimation leads to a resource allocation closer to optimal. Let T_0 be the standard simulation run time, T_1 be the time required to compute sample variance components using all paths, and T_2 be the time required to solve the convex hull problem to get a resource allocation. Our algorithm's time as a fraction of that required by standard simulation to achieve the same variance is

$$R(p) + \frac{pT_1 + T_2}{T_0}. \tag{18}$$

A larger pilot run requires more overhead for computing sample variance components but results in a resource allocation that probably yields more variance reduction.

Also, because the pilot run uses up some of the computing budget by simulating m-step paths, if it uses too many paths it may result in overcommitting resources to later time steps, preventing the optimal resource allocation from being attained at a fixed budget level. For instance, a pilot run of 500 paths might reveal that the optimal number of paths to reach the last step m should be only 100. Then step m has received five times more resources than is optimal.

We now investigate whether the theoretical variance ratio R gives an accurate picture of the actual variance reduction (18). We do this by comparing the two quantities for a particular example, that of §6.1. In that example, a pilot run of approximately 700 paths produced the minimal cost for our algorithm, as described in Equation (18). The difference between the reduced variance R and the actual cost was 0.10%. The solution time T_2 and the overcommitment of resources to later time steps turned out to be negligible in comparison to the other factors. These are the impact of imperfect estimation, which leads to a suboptimal resource allocation, leading to an attained variance ratio $R(p) > R$, and the cost pT_1 of computing sample variance components. For a version of the same problem with only 30 instead of 360 steps, the efficiency loss was only 0.03%, minimized with about 200 paths in the pilot run, and the solution time T_2 was no longer negligible. These results suggest that the theoretical variance ratio R gives a good indication of the effectiveness of our method; we therefore report estimates of R in the examples that follow.

6.1. Mortgage-Backed Security

In this application, the simulation values a mortgage-backed security (MBS), a financial security whose cashflows are the total payments made on a pool containing a large number of mortgages. The difficulty of pricing this security arises from the possibility of prepayment: The mortgages include an option for the homeowner to prepay at any date the balance of the principal and cease making payments thereafter. Prepayments increase when interest rates are low and homeowners have an incentive to refinance their mortgages at more favorable rates.

The prepayment model is based on Richard and Roll (1989), who provide further background on mortgage-backed securities. We assume that the pool is divided into individual mortgages of equal and negligible size. The refinancing ratio is the ratio of the current 30-year mortgage rate to the rate locked in for this pool of mortgages at its creation. A refinancing ratio less than 1 gives homeowners an incentive to refinance their mortgages and thus prepay the old ones. Figure 2 graphs the fraction of existing mortgages that prepay per year as a function of the refinancing ratio. Even when the refinancing ratio is very low, only approximately half the homeowners will refinance per year, perhaps because some have lost creditworthiness and have no opportunity to refinance, or because it takes them time to become aware of and respond to financial incentives. Even when the refinancing ratio is very high and unfavorable, some homeowners prepay, perhaps when moving house.

The instantaneous interest rate r_t obeys the Vasicek model, following the stochastic differential equation $dr_t = \phi(\bar{r} - r_t)\,dt + \sigma\,dW_t$ with volatility $\sigma = 1\%$, long-term average interest rate $\bar{r} = 8\%$, and mean reversion strength $\phi = 2\%$. The initial interest rate is $r_0 = 6.5\%$ and the mortgages are 30 years long. The refinancing rate is always 2% above the yield of a 30-year risk-free bond. The parameters

Figure 2. Annual prepayment fraction.

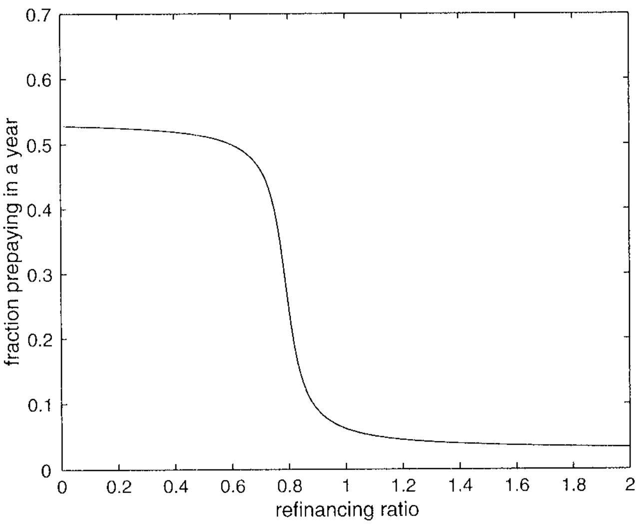

are such that the expected fraction of mortgages surviving to term is about 0.11 and the price of a 30-year zero-coupon bond is about 17 cents for a dollar of face value. These two quantities reflect the insignificance of the last step relative to the first: At the end of the MBS's life, there are fewer mortgages still in the pool, and payments have a lesser present value.

Figure 3 graphs the variance components v of the MBS's discounted cashflows. More explicitly, at each step k, the height of the graph of v is the variance of the cashflow at step k plus twice the sum of the covariances between the cashflow at k and the cashflows at all subsequent steps. As expected, there is very little variance at later steps. The graph also reveals a characteristic feature of MBSs: negative covariance between cashflows at early and late steps. Consider a single mortgage. An unusually large cashflow (the principal balance) occurs when it prepays, and afterwards cashflows are unusually low (zero). This accounts for the presence of negative variance components in the first five steps. At these steps, the negative covariance of X_k with the sum of future discounted cashflows outweighs the positive variance of X_k in Equation (4).

This feature is unfavorable for variance reduction by resource allocation among time steps. Theorem 1 indicates that there is greater variance reduction when early variance components v_k are large and positive, not when they are negative. Figure 4 shows that, in this example, the optimal resource allocation $\mathbf{n}$ does not drop below that of standard simulation (which is 1) until approximately year 20 out of 30. With this optimal $\mathbf{n}$, variance is reduced to 84.0% of standard simulation.

It is possible to reduce the impact of negative covariance by using the missing data technique. At step k, we have both a discount factor and the number of remaining mortgages at step $k+1$. The discount factor for step $k+1$ is the product of the stochastic discount realized up to step k and a one-period bond price, known in closed form for this model. The number of mortgages left at $k+1$ is the product of the number left at k and the prepayment ratio based on the history up to k. We take the predictor T_k to be the product of this discount factor and number remaining at step $k+1$, which is known at step k. The intuition is that, ignoring the variability of the nominal cashflow per mortgage, the discounted cashflow is proportional to both the discount factor and the number of mortgages. For this reason, we estimate a reduced model with $\beta_{hk} = 0$ for $h < k-1$.

Applying the missing data technique yields variance components with a smaller initial negative dip. Their partial sums V' are in Figure 5, which shows that V' is larger than V after one year, often much larger. Then Theorem 5 suggests that the variance of the estimator using the missing data technique will be better than that without. Indeed the variance is now only 50.5% compared to standard estimation. The new optimal resource allocation $\mathbf{n}'$ in Figure 4 has

Figure 3. MBS example: Variance components.

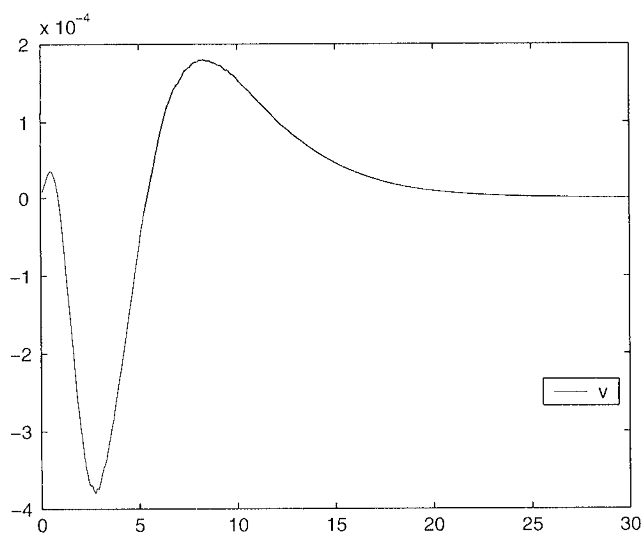

Figure 4. MBS example: Resource allocation.

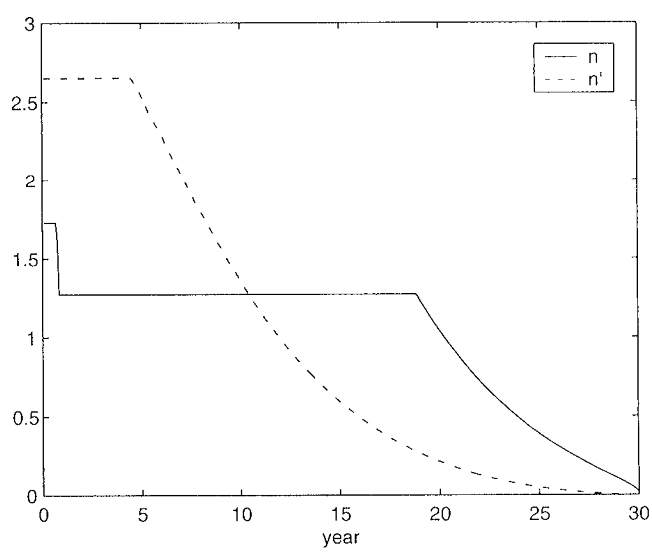

Figure 5. MBS example: Cumulative sums.

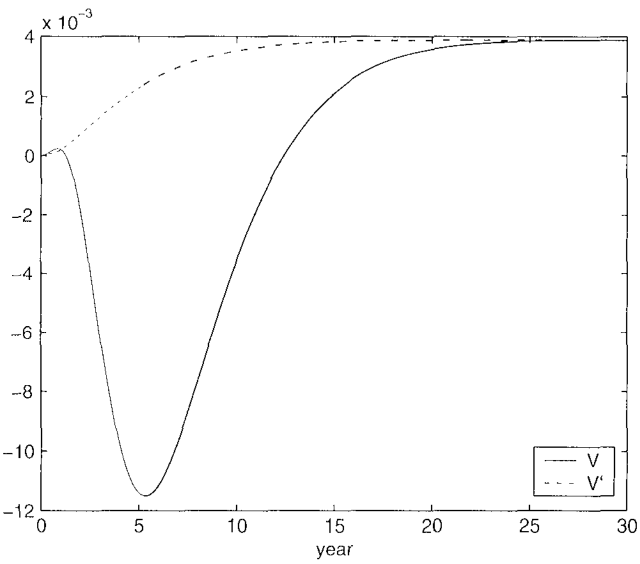

more resources devoted to the earliest steps and less at later steps. The optimal allocation stops early more aggressively when using the cashflows after adjustment using the missing data technique than for the actual cashflows, because the former bring more of the total variance to early steps in the simulation, which is the recipe for success.

6.2. Seasoned Interest-Only Security

The cashflows generated by a pool of mortgages are often divided unevenly among different "tranches" to create different types of securities. One example of this is an interest-only security (IO) that is backed solely by the interest portion of the monthly payments of the mortgage holders. As with an ordinary MBS, an IO is valued by calculating the expected present value of its cashflows, and these cashflows are sensitive to prepayments.

In this example, the parameters for the interest rate process are $\sigma = 1\%$, long-term average interest rate $\bar{r} = 5\%$, and mean reversion strength $\phi = 0.5\%$. The initial interest rate is $r_0 = 10\%$ and the mortgages are 25 years long. This is an example of a seasoned security, that is, one that was issued in the past, in this case, five years ago. The rate on the mortgages is 12%, or 2% above the 10% yield that a risk-free 30-year bond offered when the mortgages were issued. The idea behind this example is that since the mortgages were issued, a new expectation has developed that interest rates will start to decline over the long term.

This means that it is likely that many homeowners will prepay, which is favorable for this method. The variance reduction is to 22.3% without use of the missing data technique and to 13.7% with it.

That we find greater variance reduction with the IO than with the ordinary MBS is to be expected in light of the discussion in §6.1. For an ordinary MBS, prepayments introduce negative correlation between current and future cashflows because a larger payment now implies a smaller payment later. The effect on an IO is different: A prepayment has no impact on the current interest received (the prepayment is an additional payment of principal) and it reduces interest cashflows at all future dates.

6.3. Asian Option

An Asian derivative is a financial security whose payoff depends on the average of an underlying price over time. There are m averaging dates $t_1, \ldots, t_m$, and the arithmetic average of the underlying price S up to step k is $A_k = \sum_{j=1}^{k} S_j/k$, where S_j is the price at time t_j. This example is an Asian call, with payoff $(A_m - K)^+$, where K is a strike price. The underlying price obeys the Black-Scholes lognormal model, following the stochastic differential equation $dS_t = S_t(r\,dt + \sigma\,dW_t)$ under a risk-neutral probability measure. The difficulty of pricing the Asian derivative is that while the geometric average of jointly lognormal random variables is lognormal, there is no convenient expression for the distribution of the arithmetic average.

This example has $m = 5$ averaging dates, which are the last five days in the option's one-year life. The constant interest rate $r = 6.5\%$ and the volatility $\sigma = 20\%$. The strike price K and initial underlying price S_0 are both 100.

The simulation has m steps because we need to generate a price at each averaging date, but the only cashflow occurs at the terminal date. Therefore resource allocation applied directly produces no benefit. However, it is usable in combination with fictitious cashflows or the missing data approach. The intuition is that the first step in the simulation is most important because S_1 both appears directly in the average and has a great influence on later prices.

To design fictitious cashflows, we rely on our knowledge about the distribution of a geometric average of lognormal prices. There is a formula $f(k, S_k, G_k)$, which gives the value at time t_k of a call on the geometric average, given the price S_k and geometric average price to date G_k at step k. That is, $f(k, S_k, G_k) = e^{-r(t_m - t_k)}\mathbf{E}[(G_m - K)^+ \mid S_k, G_k]$, and this price is given by the Black-Scholes call pricing formula, but incorporating the parameters of the distribution not of S_m but of G_m, which Curran (1994, §2.2) provides. Our approximation for the step-k value of the arithmetic Asian call is $f(k, S_k, A_k)$. This approximation is exactly correct at step m. The fictitious cashflows are $X_k = e^{-rt_k} f(k, S_k, A_k) - e^{-rt_{k-1}} f(k-1, S_{k-1}, A_{k-1})$.

These produce the cumulative sums V of variance components in Figure 6. This curve coincides with its convex hull. The resulting optimal allocation n in Figure 7 reflects this absence of binding monotonicity constraint: n decreases at each step. However, the variance reduction is to 26.2%, close to the best possible reduction, which is 20%, as discussed at the beginning of this section.

The purpose of the Asian feature in this option might be to smooth the price used in computing the payoff, diluting the effect of possible large short-term deviations. We also considered an Asian option with $m = 5$ averaging dates spaced equally over a year. The averaging feature

Figure 6. Asian option example: Cumulative sums.

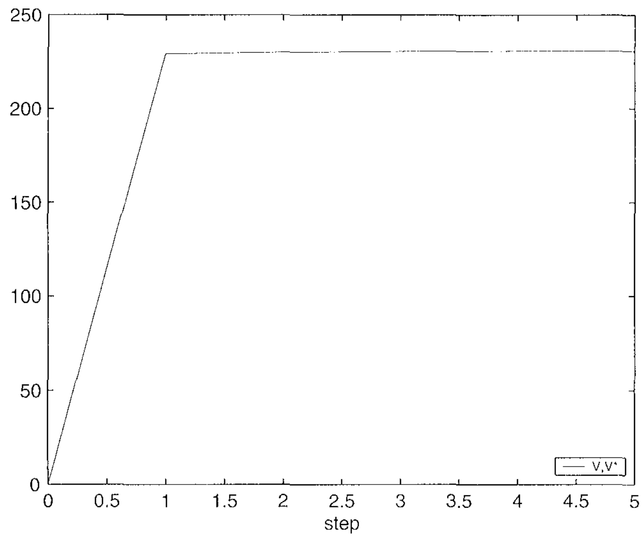

Figure 7. Asian option example: Resource allocation.

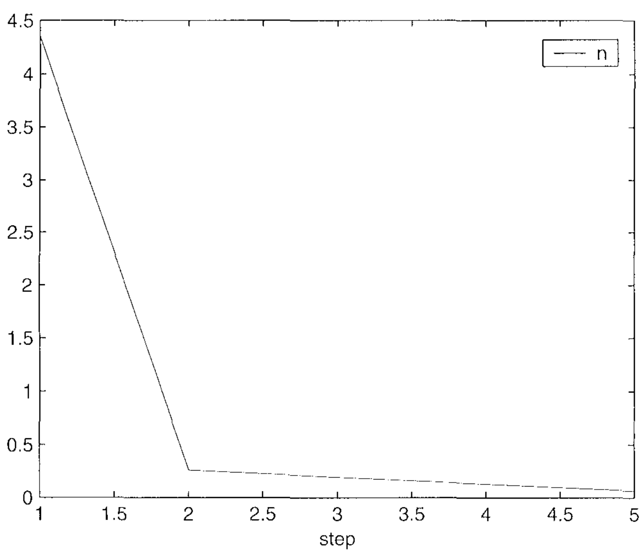

makes such an option useful, for instance, for limiting the risk of a company which plans to make regular purchases of a commodity whose price underlies the option payoff. For this Asian option with equal spacing, the results are not as good: variance reduction only to 87.6% using fictitious cashflows. The first Asian option was more favorable because almost all the information about the payoff is contained in S_1, because it has nearly a year of variability in it, while the later steps take only days.

7. CONCLUSION

We have proposed and analyzed a variance reduction technique based on allocating greater computational effort to earlier steps along a simulated path. The method is motivated by the application of simulation to estimating the expected present value of a finite stream of cashflows, where discounting or other features render later times steps less important than earlier time steps. We pose and solve a resource allocation problem to determine the optimal number of paths to simulate of each length and develop connections with the statistical theory of missing data. Further variance reduction can be achieved through a decomposition into fictitious cashflows; we find the optimal such decomposition.

Examples confirm that this method is most effective when early steps are indeed more important than later steps, as would often be the case at higher discount rates. Somewhat less obviously, the covariances between cashflows at different dates also have a significant influence on the effectiveness of the method: The method works best when early cashflows have positive covariance with later ones. Resource allocation yields variance reduction only if the problem has this sort of structure. The computational overhead incurred in solving the allocation problem is justified when the number of paths to be simulated is large relative to the number of steps per path.

APPENDIX A: PROOFS

PROOF OF LEMMA 1.

$$\begin{aligned}\mathbf{Var}[\hat{\mu}] &= \mathbf{Var}\left[\sum_{i=1}^{n}\sum_{k=1}^{m_i}\frac{X_{ik}}{n_k}\right] = \sum_{i=1}^{n}\mathbf{Var}\left[\sum_{k=1}^{m_i}\frac{X_{ik}}{n_k}\right] \\ &= \sum_{i=1}^{n}\sum_{k=1}^{m_i}\sum_{l=1}^{m_i}\frac{\sigma_{kl}}{n_k n_l} \\ &= \sum_{k=1}^{m}\sum_{l=1}^{m}\sum_{i=1}^{\min\{n_k,n_l\}}\frac{\sigma_{kl}}{n_k n_l} = \sum_{k=1}^{m}\sum_{l=1}^{m}\frac{\sigma_{kl}}{\max\{n_k,n_l\}} \\ &= \sum_{k=1}^{m}\frac{1}{n_k}\left(\sigma_{kk}+2\sum_{l=k+1}^{m}\sigma_{kl}\right).\end{aligned}$$

PROOF OF THEOREM 1. First, see that the stated solution (6) is primal-feasible. It clearly satisfies the budget constraint (2). The upper convex hull V^* is a concave function, so the slopes u_k^* of its segments are nonincreasing. These segments' endpoints are extreme points of the original graph $\mathscr{V}$, so in particular there is a $k < m$ such that $u_m^* = (V_m - V_k)/(C_m - C_k)$. Assuming that costs are positive, $C_m - C_k > 0$. Also $V_m - V_k = \sum_{l=k+1}^{m} v_l = \mathbf{Var}[\sum_{l=k+1}^{m} X_l] > 0$. Therefore $u_m^* > 0$, and since these are nonincreasing, all are positive. Consequently, all n_k are positive and nonincreasing, satisfying the monotonicity constraint (1).

Next, show that the solution is dual-feasible. The Lagrangian is

$$\mathscr{L} = \sum_{k=1}^{m}\frac{v_k}{n_k} - \nu\left(C - \sum_{k=1}^{m} c_k n_k\right) - \sum_{k=1}^{m-1}\lambda_k(n_k - n_{k+1}).$$

Differentiating with respect to n_j, the first-order conditions are, for $j = 1, \ldots, m$,

$$\lambda_{j-1} = \lambda_j + v_j/n_j^2 - \nu c_j,$$

where λ_0 denotes zero.

Dual-feasibility is the existence of nonnegative Lagrange multipliers λ_j and ν which satisfy these first-order equations and the complementary slackness conditions $\lambda_j(n_j - n_{j+1}) = 0$. Let ν be as given in Equation (7), which is positive, and let

$$\lambda_j = (V_j^* - V_j)/n_j^2,$$

which is nonnegative by definition of the upper convex hull. So j is the index of an extreme point if and only if $\lambda_j = 0$, and otherwise $\lambda_j > 0$. Therefore, $\lambda_j \neq 0$ implies that j is not the index of an extreme point, and because V^* is linear between extreme points, the slopes $u_{j+1}^* = u_j^*$, hence, $n_j - n_{j+1} = 0$. Thus complementary slackness is satisfied.

Define $l(j)$ as the index of the first extreme point to the right of (C_j, V_j):

$$l(j) \triangleq \min\{k > j \mid V_k^* = V_k\}.$$

As mentioned above, $V_m^* = V_m$ so this is well-defined for $j = 1, \ldots, m-1$. Then

$$\begin{aligned} \lambda_j n_j^2 = V_j^* - V_j &= (V_{l(j)} - V_j) - (V_{l(j)}^* - V_j^*) \\ &= \sum_{k=j+1}^{l(j)} (v_k - v_k^*). \end{aligned}$$

Again, because V^* is linear between extreme points, n_k is constant for $k = j+1, \ldots, l(j)$. Also, from Equation (6), $n_k^2 = v_k^*/\nu c_k$, so

$$\lambda_j = \sum_{k=j+1}^{l(j)} \frac{v_k - v_k^*}{n_k^2} = \sum_{k=j+1}^{l(j)} \left(\frac{v_k}{n_k^2} - \nu c_k\right).$$

Suppose j is not the index of an extreme point. Then $l(j-1) = l(j)$ and we see directly from this that the first-order condition is satisfied. If j is the index of an extreme point, then $l(j-1) = j$ and $\lambda_j = 0$, so $\lambda_{j-1} = 0 + v_j/n_j^2 - \nu c_j$ and the first-order condition is satisfied.

Finally, there is the matter of the variance reduction ratio. Under standard simulation,

$$n_k = \frac{C}{\sum_{i=1}^m c_i} \quad \text{and} \quad \mathbf{Var} = \sum_{k=1}^m \frac{v_k}{n_k} = \left(\sum_{i=1}^m c_i\right)\left(\sum_{k=1}^m v_k\right)\Big/ C.$$

Using the optimal solution,

$$n_k = \frac{C\sqrt{u_k^*}}{\sum_{i=1}^m c_i\sqrt{u_i^*}} \quad \text{and} \quad \mathbf{Var} = \left(\sum_{i=1}^m c_i\sqrt{u_i^*}\right)\left(\sum_{k=1}^m \frac{c_k u_k}{\sqrt{u_k^*}}\right)\Big/ C,$$

so changing indices and substituting for slopes $u = v/c$, the ratio of variances is

$$\frac{(\sum_{j=1}^m c_j\sqrt{u_j^*})(\sum_{i=1}^m v_i/\sqrt{u_i^*})}{(\sum_{j=1}^m c_j)(\sum_{i=1}^m v_i)}.$$

which equals the formula given in Equation (8).

PROOF OF LEMMA 2. First we show that $\sum_{k=1}^m \hat{\mu}_k$ can indeed be written in the form given by Equations (12)–(15). By substituting for $\hat{\tau}_h$ in the recursive definition (11), we see that $\hat{\tau}_k$ is a linear combination of averages $\overline{T}$, where the weights in the linear combination involve estimated regression coefficients $\hat{b}$. As remarked in §2, the average indexed as $\overline{T}_{hi}$ must have $h \leqslant i$ because it is the average of T_h on paths $i, \ldots, n_i$ so we must have $n_h \geqslant n_i$ for T_h to be observed on all these paths. Also, if $\overline{T}_{hi}$ is to feature in $\hat{\tau}_k$, then we must have $i \leqslant k$ so that $\overline{T}_{hi}$ is based on at least as many paths as $\overline{T}_{kk}$, the obvious estimate of τ_k, and thus can be of use in correcting it. Having established that the random variable T_{ik} does not appear in $\hat{\tau}_h$ for $h < k$, by inspection, $\omega_{kkk} = 1$ is the coefficient of the average $\overline{T}_{kk}$ in $\hat{\tau}_k$ and $\omega_{hkk} = -\hat{b}_{hk}$ is the coefficient of the average $\overline{T}_{hk}$ in $\hat{\tau}_k$. For $i < k$, the random variable T_{hi} does not appear directly in $\hat{\tau}_k$, only through its appearance in $\hat{\tau}_j$ for $j = i, \ldots, k-1$. Therefore we can write

$$\hat{\tau}_k = \sum_{i=1}^k \sum_{h=1}^i \omega_{hik}\overline{T}_{hi}$$

following the definition of Equation (13).

By substituting for $\hat{\tau}_h$ in definition (10) and repeating the reasoning of the previous paragraph,

$$\hat{\mu}_k = \overline{X}_{kk} + \sum_{i=1}^k \sum_{h=1}^i w_{hik}\overline{T}_{hi},$$

following the definition of Equation (14). Whereas ω_{kkk} is 1, w_{kkk} is 0, because X_{kk} appears in $\hat{\mu}_k$ where T_{kk} appears in $\hat{\tau}_k$. Then

$$\begin{aligned} \sum_{k=1}^m \hat{\mu}_k &= \sum_{k=1}^m \left(\overline{X}_{kk} + \sum_{i=1}^k \sum_{h=1}^i w_{hik}\overline{T}_{hi}\right) \\ &= \sum_{k=1}^m \overline{X}_{kk} + \sum_{i=1}^m \sum_{h=1}^i \left(\overline{T}_{hi} \sum_{k=i}^m w_{hik}\right) \\ &= \sum_{k=1}^m \overline{X}_{kk} + \sum_{i=1}^m \sum_{h=1}^i \overline{T}_{hi} w_{hi} = \sum_{k=1}^m \left(\overline{X}_{kk} + \sum_{h=1}^m \overline{T}_{hk} w_{hk}\right) \end{aligned}$$

using Equation (15), which says that $w_{hi} = \sum_{k=i}^m w_{hik}$ and is zero if $h > i$. Continuing,

$$\sum_{k=1}^m \hat{\mu}_k = \sum_{k=1}^m \frac{1}{n_k} \sum_{i=1}^{n_k} \left(X_{ik} + \sum_{h=1}^m T_{ih} w_{hk}\right) = \sum_{k=1}^m \frac{1}{n_k} \sum_{i=1}^{n_k} X'_{ik}.$$

Next, we show that $\mathbf{W1} = \mathbf{0}$ by virtue of the recursive definitions (13)–(15). Recall that ω_{hik} and w_{hik} are zero unless $h \leqslant i \leqslant k$. We begin with a proof by induction on k that for any $k > h$, $\sum_{i=h}^k \omega_{hik} = 0$.

$$\begin{aligned} \sum_{i=h}^k \omega_{hik} &= \omega_{hkk} + \sum_{i=h}^{k-1} \sum_{j=i}^{k-1} \hat{b}_{jk}\omega_{hij} = -\hat{b}_{hk} + \sum_{j=h}^{k-1} \hat{b}_{jk} \sum_{i=h}^j \omega_{hij} \\ &= -\hat{b}_{hk} + \hat{b}_{hk}\omega_{hhh} + \sum_{j=h+1}^{k-1} \hat{b}_{jk} \sum_{i=h}^j \omega_{hij} \\ &= -\hat{b}_{hk} + \hat{b}_{hk} + \sum_{j=h+1}^{k-1} \hat{b}_{jk} 0 = 0, \end{aligned}$$

where $\sum_{i=h}^{j} \omega_{hij} = 0$ is justified by inductive hypothesis because $h < j < k$.

To complete the proof that $\mathbf{W1} = \mathbf{0}$, we must establish that for any $k \geqslant h$, $\sum_{i=h}^{k} w_{hik} = 0$. If $k = h$, this is $w_{hhh} = 0$. Otherwise,

$$\sum_{i=h}^{k} w_{hik} = w_{hkk} + \sum_{i=h}^{k-1}\sum_{j=i}^{k-1} \hat{\beta}_{jk} w_{hij} = -\hat{\beta}_{hk} + \sum_{j=h}^{k-1} \hat{\beta}_{jk} \sum_{i=h}^{j} \omega_{hij}$$

$$= -\hat{\beta}_{hk} + \hat{\beta}_{hk}\omega_{hhh} + \sum_{j=h+1}^{k-1} \hat{\beta}_{jk} \sum_{i=h}^{j} \omega_{hij}$$

$$= -\hat{\beta}_{hk} + \hat{\beta}_{hk} + \sum_{j=h+1}^{k-1} \hat{\beta}_{jk} 0 = 0.$$

Consequently, for any h,

$$\sum_{i=h}^{m} w_{hi} = \sum_{i=h}^{m}\sum_{k=i}^{m} w_{hik} = \sum_{k=h}^{m}\sum_{i=h}^{k} w_{hik} = \sum_{k=h}^{m} 0 = 0.$$

PROOF OF THEOREM 2. The expectation is

$$\mathbf{E}\left[\sum_{k=1}^{m}\sum_{i=1}^{n_k} \frac{X'_{ik}}{n_k}\right] = \mathbf{E}\left[\sum_{k=1}^{m} X'_k\right].$$

But

$$\sum_{k=1}^{m} X'_k = \sum_{k=1}^{m}\left(X_k + \sum_{h=1}^{m} T_h w_{hk}\right)$$

$$= \sum_{k=1}^{m} X_k + \sum_{h=1}^{m} T_h \sum_{k=1}^{m} w_{hk}$$

$$= \sum_{k=1}^{m} X_k + \sum_{h=1}^{m} T_h 0 = \sum_{k=1}^{m} X_k,$$

because $\mathbf{W1} = 0$ so for each h, $\sum_{k=1}^{m} w_{hk} = 0$. By definition, $\mathbf{E}[\sum_{k=1}^{m} X_k] = \mu$.

PROOF OF THEOREM 3. We have $X'_{ik} = X_{ik} + \sum_{h=1}^{m} T_{ih} w_{hk} = X_{ik} + \sum_{h=1}^{k} T_{ih} w_{hk}$. When i and j index distinct paths, X'_{ik} and X'_{jk} are independent. Also, X'_{ik} does not involve the decision variables $\mathbf{n}$. Therefore, from $\sum_{k=1}^{m}\sum_{i=1}^{n_k} X'_{ik}/n_k = \sum_{i=1}^{n}\sum_{k=1}^{m_i} X'_{ik}/n_k$ the proof of Lemma 1 applies, with X' replacing X. Similarly, using the v'_k instead of v_k, Theorem 1 gives the optimal n_k for this estimator.

Rewrite from Equation (12)

$$\hat{\mu} = \sum_{i=1}^{n_1}\sum_{k=1}^{m_i} \frac{1}{n_k} X'_{ik} = \sum_{k=1}^{m} \hat{\nu}_k \quad \text{where}$$

$$\hat{\nu}_k \triangleq \frac{1}{n_k - n_{k+1}} \sum_{i=n_{k+1}}^{n_k} \sum_{j=1}^{k} X'_{ij}\left(\frac{n_k - n_{k+1}}{n_j}\right).$$

As the computational budget C goes to infinity, each $n_k - n_{k+1}$ goes to infinity, and $(n_k - n_{k+1})/n_j$ approaches some finite limit p_{jk} such that $\sum_{k=j}^{m} p_{jk} = 1$ for each j. Because $\sum_{j=1}^{k} p_{jk} X'_j$ has finite variance, the Lindeberg central limit theorem, for which see, e.g., Billingsley (1995, p. 359), implies that for each k, the distribution of

$$\frac{\hat{\nu}_k - \mathbf{E}[\sum_{j=1}^{k} p_{jk} X'_j]}{\sqrt{\mathbf{Var}[\sum_{j=1}^{k} p_{jk} X'_j]/(n_k - n_{k+1})}} \tag{19}$$

converges to standard normal as $C \to \infty$. The $\hat{\nu}_k$ are independent because they are taken over different sample paths, and they sum to $\hat{\mu}$. The sum of the expectations in (19) is

$$\sum_{k=1}^{m}\sum_{j=1}^{k} p_{jk}\mathbf{E}[X'_j] = \sum_{j=1}^{m} \mathbf{E}[X'_j] \sum_{k=j}^{m} p_{jk} = \sum_{j=1}^{m} \mathbf{E}[X'_j] = \mu,$$

and the sum of the variances in (19) is

$$\sum_{k=1}^{m} \frac{1}{n_k - n_{k+1}} \mathbf{Var}\left[\sum_{j=1}^{k} p_{jk} X'_j\right]$$

$$= \sum_{k=1}^{m} \frac{1}{n_k - n_{k+1}} \sum_{i=1}^{k}\sum_{j=1}^{k} p_{ik} p_{jk} \sigma'_{ij}$$

$$= \sum_{i=1}^{m}\sum_{j=1}^{m} \sigma'_{ij} \sum_{k=\max\{i,j\}}^{m} \left(\frac{p_{ik} p_{jk}}{n_k - n_{k+1}}\right). \tag{20}$$

From the proof of Lemma 1, with a change of indices, we have

$$\varsigma^2 = \sum_{i=1}^{m}\sum_{j=1}^{m}\left(\frac{\sigma'_{ij}}{\max\{n_i, n_j\}}\right) = \sum_{i=1}^{m}\sum_{j=1}^{m} \sigma'_{ij} \sum_{k=\max\{i,j\}}^{m} \left(\frac{n_k - n_{k+1}}{n_i n_j}\right).$$

The ratios

$$\left(\frac{p_{ik} p_{jk}}{n_k - n_{k+1}}\right) \Big/ \left(\frac{n_k - n_{k+1}}{n_i n_j}\right)$$

all converge to 1 as $C \to \infty$, so the ratio of (20) to ς^2 also converges to 1. Therefore, the distribution of $(\hat{\mu} - \mu)/\varsigma$ converges to standard normal as $C \to \infty$.

PROOF OF THEOREM 4. While ς is the standard deviation given by Equation (16), let ς° be that produced by substituting $\mathbf{W}^\circ$ in Equation (16). The assumption of consistency says that $w_{hi} - w^\circ_{hi}$ converges in probability to 0 for each h, i. Because of this convergence in probability, the random variable ς converges in probability to ς°. Put differently, $\varsigma/\varsigma^\circ$ converges in probability to one.

Also let $\hat{\mu}^\circ$ be like $\hat{\mu}$, but with the constant matrix $\mathbf{W}^\circ$ substituted for the random matrix $\mathbf{W}$. Then

$$\hat{\mu} - \hat{\mu}^\circ = \sum_{k=1}^{m} \frac{1}{n_k} \sum_{i=1}^{n_k}\sum_{h=1}^{k} T_{ih}(w_{hk} - w^\circ_{hk}),$$

so to prove that $\hat{\mu} - \hat{\mu}^\circ$ converges in probability to zero, it suffices to show that for all h and k, $T_h(w_{hk} - w^\circ_{hk})$ converges in probability to zero.

We already know that for any positive δ and ϵ, there is some $C(\delta, \epsilon)$ such that if the budget $C > C(\delta, \epsilon)$, then $\mathbf{P}[|w_{hk} - w^\circ_{hk}| > \epsilon] < \delta$ for each h, k. We must next show this is also true for $|T_h(w_{hk} - w^\circ_{hk})|$. Let F be the cumulative distribution function of $\max_h |T_h|$, and $\bar{F} = 1 - F$ be its tail probability. The event

$$\{|T_h(w_{hk} - w^\circ_{hk})| > \epsilon\} \subset \{|T_h| > \bar{F}^{-1}(\delta/2)\} \cup \{|w_{hk} - w^\circ_{hk}| > \epsilon/\bar{F}^{-1}(\delta/2)\}.$$

Then for $C > C(\delta/2, \epsilon/\overline{F}^{-1}(\delta/2))$,

$$\begin{aligned}\mathbf{P}[|T_h(w_{hk} - w^{\circ}_{hk})| > \epsilon] &< \mathbf{P}[|T_h| > \overline{F}^{-1}(\delta/2)] \\ &\quad + \mathbf{P}[|w_{hk} - w^{\circ}_{hk}| > \epsilon/\overline{F}^{-1}(\delta/2)] \\ &< \delta/2 + \delta/2 = \delta,\end{aligned}$$

i.e., $|T_h(w_{hk} - w^{\circ}_{hk})|$ converges in probability to zero. This proof works because there are a finite number of T_h, each of which is finite with probability one.

It is established that $\hat{\mu} - \hat{\mu}^{\circ}$ converges in probability to zero, and because ς° is positive, $(\hat{\mu} - \mu)/\varsigma^{\circ} - (\hat{\mu}^{\circ} - \mu)/\varsigma^{\circ}$ also converges in probability to zero. Using the convergence in probability of $\varsigma/\varsigma^{\circ}$ to one, finally we see $(\hat{\mu} - \mu)/\varsigma - (\hat{\mu}^{\circ} - \mu)/\varsigma^{\circ}$ converges in probability to zero. Because $\mathbf{W}^{\circ}$ is constant, by Theorem 3, $(\hat{\mu}^{\circ} - \mu)/\varsigma^{\circ}$ converges in distribution to standard normal, so $(\hat{\mu} - \mu)/\varsigma$ does as well.

Proof of Theorem 5. Suppose that sequences $X'_1, \ldots, X'_m$ and $X''_1, \ldots, X''_m$ in $\mathscr{X}$ satisfy $V'_k \geqslant V''_k$ for all $k = 0, \ldots, m$. One is tempted to say that this condition on the partial sums means that v' majorizes v. However, this is not quite so, for the definition of majorization involves partial sums of terms placed in decreasing order. Majorization is a relation between sets, not vectors; see, for instance, Marshall and Olkin (1979, p. 7, 12). Here the sequences retain their original order.

Let $\mathbf{n}'$ and $\mathbf{n}''$ be the optimal resource allocations for the problems using respectively X' and X''. Define the function

$$f(x_1, \ldots, x_m) \triangleq \sum_{k=1}^{m} \frac{x_k}{n''_k}.$$

The derivative of f with respect to its kth argument is $1/n''_k$. This is nondecreasing in k because the n''_k are nonincreasing, due to the monotonicity constraint (1). Because we can move from v' to v'' by subtracting from early components and adding equal amounts to later components, $f(v'_1, \ldots, v'_m) \leqslant f(v''_1, \ldots, v''_m)$. This informal argument parallels Theorem A.3 of Marshall and Olkin (1979), but without the restriction that components be decreasing. Then

$$\sum_{k=1}^{m} \frac{v'_k}{n'_k} \leqslant \sum_{k=1}^{m} \frac{v'_k}{n''_k} \leqslant \sum_{k=1}^{m} \frac{v''_k}{n''_k},$$

where the first inequality is true because n' minimizes the objective based on v' and the second inequality is the result about f just established. Consequently, the sequence $X'_1, \ldots, X'_m$ produces the resource allocation problem whose optimal objective is least if for any fictitious cashflow $X''_1, \ldots, X''_m$, $V'_k \geqslant V''_k$ for each k.

Next,

$$V'_k = V'_m - \sum_{l=k+1}^{m} v'_l = V'_m - \mathbf{Var}\left[\sum_{l=k+1}^{m} X'_l\right].$$

Also, because the sequence $X'_1, \ldots, X'_m$ is in $\mathscr{X}$, $\sum_{k=1}^{m} X'_k = X$, so $\sum_{l=k+1}^{m} X'_l = X - \sum_{j=1}^{k} X'_j$, and $\sum_{j=1}^{k} X'_j$ must be $\mathscr{F}_k$-measurable. Thus each V'_k is maximized by minimizing $\mathbf{Var}[X - Y]$ over $\mathscr{F}_k$-measurable random variables Y. As is well known, it is $\mathbf{E}[X \mid \mathscr{F}_k]$ that minimizes this residual variance; see, for instance, Williams (1991, §9.4).

APPENDIX B: ALGORITHM

Step 1. Construct a doubly linked list of points that are candidate extreme points of the upper convex hull in that they are above the line connecting the endpoints:

(a) Let $s = V_m/C_m$, $k = m$, and CARRY = START, a pointer to a newly created blank node.
(b) If $V_k \geqslant sC_k$, create a node with COST = C_k, VALUE = V_k, PREV = CARRY, NEXT = NULL, and let both PREV → NEXT and CARRY be pointers to this node.
(c) Decrease k by 1. If $k \geqslant 0$, go to Step 1b.

Step 2. Scan the list and eliminate points that are not extreme:

(a) Let the pointers K = START → NEXT, J = K → NEXT, and I = J → NEXT.
(b) Compute $\Delta = V_J(C_K - C_I) - (V_I(C_K - C_J) + V_K(C_J - C_I))$.
(c) If $\Delta \geqslant 0$, advance the scan: K = J, J = I, I = I → NEXT.
Otherwise, delete node J and back up: J = K, K = K → PREV.
(d) If I $\neq$ NULL, go to Step 2b.

Step 3. Produce n from the extreme points:

(a) Initialize $k = m$, K = START → NEXT, J = K → NEXT, and $u^*_k = (V_K - V_J)/(C_K - C_J)$.
(b) If $k \leqslant J$, let K = J and J = J → NEXT, and recompute u_k.
(c) Assign $v^*_k = u^*_k c_k$ and decrease k by 1. If $k > 0$, repeat Step 3b.
(d) Compute n from v^* as in Equation (6).

Preparata and Shamos (1985, §3.3.2) provide a proof that the Graham scan does produce the convex hull. Our algorithm departs from the standard Graham scan in three ways:

(1) It does not include a sort, because the points of the graph $\mathscr{V}$ are already sorted.

(2) In Step 1b, it discards points below the line connecting $(0, 0)$ and (C_m, V_m), which cannot form part of the upper convex hull. In Step 2d, the algorithm stops loop 2 when it reaches $(0, 0)$ and has completed the upper convex hull instead of going on to find the lower convex hull, too.

(3) This means that $(0, 0)$ must always be included in the output of loop 2. We already know $(0, 0)$ and (C_m, V_m) are the leftmost and rightmost points and must be in the upper convex hull. The algorithm also ensures that (C_m, V_m) is in the output. In the first iteration of Step 1b, the algorithm sets START → PREV = PREV. Then in Step 2b, if K = START and $\Delta < 0$, the result of backing up is to leave K = START and J = START. The next time Step 2b executes, $\Delta = 0$ and the scan will advance. The result is that the START node representing (C_m, V_m) can never be deleted.

ACKNOWLEDGMENTS

This research is supported in part by NSF grant DMS-0074637, an IBM University Partnership Award, and an NSF graduate research fellowship. The authors thank Mark Broadie and Robin Roundy for providing references, Erik Gilje for computational assistance, and the anonymous reviewers for many constructive suggestions.

REFERENCES

Asmussen, S. 1990. Exponential families and regression in the Monte Carlo study of queues and random walks. *Ann. Statist.* **18** 1851–1867.

Billingsley, P. 1995. *Probability and Measure*, 3rd ed. Wiley-Interscience, New York.

Cochran, W. G. 1953. *Sampling Techniques.* John Wiley and Sons, New York.

Curran, M. 1994. Valuing Asian and portfolio options by conditioning on the geometric mean price. *Management Sci.* **40** 1705–1711.

Federgruen, A., Y. S. Zheng. 1992. The joint replenishment problem with general joint cost structures. *Oper. Res.* **40** 384–403.

Fox, B. L., P. W. Glynn. 1989. Simulating discounted costs. *Management Sci.* **35** 1297–1315.

Glasserman, P., J. Staum. 2001. Stopping simulated paths early. B. A. Peters, J. S. Smith, D. J. Medeiros, M. W. Rohrer, eds. *Proc. 2001 Winter Simulation Conf.* IEEE Press, Piscataway, NJ, 318–324.

Hocking, R. R., W. B. Smith. 1972. Optimum incomplete multinormal samples. *Technometrics* **14** 299–307.

Little, R. J. A., D. B. Rubin. 1987. *Statistical Analysis with Missing Data.* John Wiley and Sons, New York.

Marshall, A. W., I. Olkin. 1979. *Inequalities: Theory of Majorization and Its Applications.* Academic Press, New York.

Maxwell, W. L., J. A. Muckstadt. 1985. Establishing consistent and realistic reorder intervals in production-distribution systems. *Oper. Res.* **33** 1316–1341.

Preparata, F. P., M. I. Shamos. 1985. *Computational Geometry.* Springer-Verlag, New York.

Richard, S. F., R. Roll. 1989. Prepayments on fixed-rate mortgage-backed securities. *J. Portfolio Management* **15** 73–82.

Roundy, R. 1986. A 98-percent-effective lot-sizing rule for a multiproduct, multistage production inventory system. *Math. Oper. Res.* **11** 699–727.

Williams, D. 1991. *Probability with Martingales.* Cambridge University Press, New York.