

STOPPING SIMULATED PATHS EARLY

Paul Glasserman

Graduate School of Business
Columbia University
New York, NY 10027, U.S.A.

Jeremy Staum

School of Operations Research and Industrial Engineering
Cornell University
Ithaca, NY 14853, U.S.A.

ABSTRACT

We provide results about stopping simulation paths early as a variance reduction technique, adding to our earlier work on this topic. The problem of pricing a financial instrument with cashflows at multiple times, such as a mortgage-backed security, motivates this approach, which is more broadly applicable to problems in which early steps are more informative than later steps of a path. We prove a limit theorem that demonstrates that this relative informativeness of simulation steps, not the number of steps, determines the effectiveness of the method. Next we consider an extension of the idea of stopping simulation paths early, showing how early stopping can be random and depend on the state a path has reached, yet still produce an unbiased estimator. We illustrate the potential effectiveness of such estimators, and describe directions for future research into their design.

1 INTRODUCTION

In a standard simulation for a purpose such as pricing a financial derivative, every simulated path reaches each of a fixed number of time steps. We consider simulations that estimate an expected total reward, for concreteness focusing on a sum of discounted cashflows that a derivative security pays. When discounted cashflows at early steps are more important than those at later steps, devoting equal computational resources to simulation at early steps and later steps seems to be a suboptimal allocation of resources.

An example is a mortgage-backed security, whose monthly cashflows are the sum of all payments made on a pool of mortgages. These payments are nominally constant, but the value of discounted cashflows tends to decline because of the time value of money and because some homeowners prepay their mortgages, thus exiting the pool.

In previous work, we showed how to find a policy of early stopping, with some paths scheduled to terminate before the final time step, that minimizes the variance of an estimator with a fixed computational budget (Glasserman

and Staum 2001b). In Section 2, we review the formalism and main results of the previous paper, including the use of the method for some problems to which it is seemingly inapplicable, such as those with only a terminal cashflow. Section 3 contains a new theorem, which proves that the variance reduction obtained by early stopping depends not on the number of time steps, but on the shape of the covariance matrix of discounted cashflows. An extension of early stopping occupies Section 4: here we introduce a method where the decision to stop is randomized and depends on the current state, rather than being planned in advance. We then suggest directions for future research into its effective use.

2 REVIEW OF EARLY STOPPING

This section presents without proof a result of Glasserman and Staum (2001b). Motivated by the problem of pricing a mortgage-backed security, we consider a simulation to estimate the expectation of a sum $X = \sum_{k=1}^m X_k$, interpreting the terms X_1, \dots, X_m as discounted cashflows. When some paths stop early, there are n_k paths that reach step k , producing the observations $X_{1k}, \dots, X_{n_k k}$.

The numbers of paths reaching each step must satisfy a monotonicity constraint: $n_k \geq n_{k+1}$. Let the cost of step k be c_k and the total computational budget be C . The budget constraint is $\sum_{k=1}^m c_k n_k \leq C$.

The obvious estimator of the mean $\mu = \mathbf{E}[X]$ is now

$$\hat{\mu} = \sum_{k=1}^m \frac{1}{n_k} \sum_{i=1}^{n_k} X_{ik} \quad (1)$$

We want to minimize over all feasible choices of n_1, \dots, n_m the variance of $\hat{\mu}$, which is $\mathbf{Var}[\hat{\mu}] = \sum_{k=1}^m v_k / n_k$ where

$$v_k = \mathbf{Var}[X_k] + 2\mathbf{Cov} \left[X_k, \sum_{j=k+1}^m X_j \right].$$

The “variance components” v_1, \dots, v_m on which the solution is based are typically unknown, but they can be estimated from a set of initial simulated paths, and the optimal resource allocation applied to the rest of the simulation.

The resource allocation problem is

$$\min \sum_{k=1}^m \frac{v_k}{n_k} \quad \text{s.t.} \quad \sum_{k=1}^m c_k n_k \leq C \quad (2)$$

$$\text{and } n_k \geq n_{k+1}, \quad k = 1, \dots, m - 1$$

The solution to problem (2) involves the convex hull of a graph determined by the costs c_k and the variance components v_k . Define the partial sums

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

and the graph

$$\mathcal{V} = \{(C_k, V_k) | k = 0, \dots, m\}.$$

Let the function V^* be the upper convex hull of this graph, and let V_k^* denote $V^*(C_k)$. The solution is based on the slopes $u_k^* = v_k^*/c_k$ where v_k^* is the increment $V_k^* - V_{k-1}^*$. Here we restate Theorem 1 of Glasserman and Staum (2001b).

Theorem 1. *The solution to the resource allocation problem (2) is*

$$n_k = \sqrt{\frac{u_k^*}{v}} \quad \text{and} \quad v = \left(\frac{1}{C} \sum_{k=1}^m c_k \sqrt{u_k^*} \right)^2$$

and the ratio of optimal variance to standard variance is

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

We go on to describe two methods that can enhance the effectiveness of early stopping and make it more broadly applicable. The first involves the statistical perspective on missing data. Viewing the simulation as an experiment, the data table has many missing observations X_{ik} when path i stops before reaching step k . The theory of missing data leads to estimators which are superior to $\hat{\mu}$ because they take advantage of dependence between the variables X_1, \dots, X_m and use information about the more frequently observed variables to learn about those less frequently observed.

The other approach is to study new random variables X'_1, \dots, X'_m that satisfy the following two properties.

- At step k of the simulation, X'_k can be computed.
- The sum $\sum_{k=1}^m X'_k = X = \sum_{k=1}^m X_k$.

The estimator (1) based on the new random variables has the same expectation μ , but different variance components v_1, \dots, v_m , which may render early stopping more effective. The optimal choice of X'_1, \dots, X'_m would satisfy $\sum_{j=1}^k X'_j = \mathbf{E}[X | \mathcal{F}_k]$, but of course this conditional expectation is unknown. However, there are many problems for which we have a decent approximation to the conditional expectation $\mathbf{E}[X | \mathcal{F}_k]$, and this enables early stopping to provide more variance reduction. It is this perspective to which we will return in Section 4 as we develop methods of random early stopping.

3 CONTINUOUS LIMIT

This section shows that the “shape” of the covariance matrix of discounted cashflows determines the solution and effectiveness of the resource allocation problem (2). First define a sequence of resource allocation problems with convergent shape as follows. Let v be a positive, continuous function on $[0, 1]$, and let $V(t) = \int_0^t v(s) ds$. The sequence of problems is indexed by size m : in the m th problem, $C^{(m)}$ is the budget and $\mathcal{V}^{(m)}$ is the point set $\{(C_k^{(m)}, V_k^{(m)}) | k = 0, \dots, m\}$, where

$$\begin{aligned} C_k^{(m)} &= \frac{kC^{(m)}}{m} \\ V_k^{(m)} &= V^{(m)}(k/m) \\ V^{(m)}(t) &= \int_0^t v^{(m)}(s) ds \\ v^{(m)}(t) &= v([tm]/m) \end{aligned}$$

The m th problem has variance components $v_k^{(m)} = v((k - 1)/m)/m$ and equal costs $c_k^{(m)} = C^{(m)}/m$.

Let $V^{(m)*}$ be the upper convex hull of $V^{(m)}$ and V^* be the upper convex hull of V . We will prove in the lemmas in the Appendix that this sequence of problems has convergent shape, as described by $V^{(m)*}$, and that the upper convex hulls $V^{(m)*}$ and their slopes $v^{(m)*}$ also converge respectively to V^* and its derivative v^* .

Then Theorem 2 asserts that this sequence of problems has solution and effectiveness that also converge. To make this precise, we define

$$n^{(m)}(t) = \frac{\sqrt{v^{(m)*}(t)}}{\int_0^1 \sqrt{v^{(m)*}(s)} ds} \quad (3)$$

and

$$F^{(m)} = \left(\int_0^1 \sqrt{v^{(m)*}(s)} ds \right) \left(\int_0^1 \frac{v^{(m)}(s)}{\sqrt{v^{(m)*}(s)}} ds \right). \quad (4)$$

These are the quantities of interest.

Lemma 1. *The m th problem of form (2) has solution given by $n_k^{(m)} = n^{(m)}(k/m)$ and its optimal variance objective is $F^{(m)}$.*

Proof. Because $V^{(m)}$ is continuous and piecewise linear, so is its upper convex hull $V^{(m)*}$, and $v^{(m)*}$ is piecewise constant. We have

$$v_k^{(m)*} = v^{(m)*}((k-1)/m)/m$$

just like $v_k^{(m)} = v^{(m)}((k-1)/m)/m$. By Theorem 1, the solution to the m th problem is

$$\begin{aligned} n_k^{(m)} &= \sqrt{\frac{v_k^{(m)*}}{vc_k^{(m)}}} \\ &= C^{(m)} \frac{\sqrt{v_k^{(m)*}/c_k^{(m)*}}}{\sum_{j=1}^m \sqrt{c_j^{(m)} v_j^{(m)*}}} \\ &= \sqrt{\frac{v^{(m)*}((k-1)/m)}{\frac{1}{m} \sum_{j=1}^m \sqrt{v^{(m)*}((j-1)/m)}}}. \end{aligned}$$

Because $v^{(m)*}$ is piecewise constant, $n_k^{(m)}$ is indeed equal to $n^{(m)}(t)$ as defined in equation (3) for $t \in [(k-1)/m, k/m)$.

The variance objective of the m th problem is

$$\begin{aligned} \sum_{k=1}^m \frac{v_k^{(m)}}{n_k^{(m)}} &= \left(\frac{1}{m} \sum_{j=1}^m \sqrt{v_j^{(m)*}} \right) \left(\sum_{k=1}^m \frac{v_k^{(m)}}{\sqrt{v_k^{(m)*}}} \right) \\ &= \left(\frac{1}{m} \sum_{j=1}^m \sqrt{\frac{v^{(m)*} \left(\frac{j-1}{m} \right)}{m}} \right) \left(\sum_{k=1}^m \frac{\frac{1}{m} v^{(m)} \left(\frac{k-1}{m} \right)}{\sqrt{\frac{1}{m} v^{(m)*} \left(\frac{k-1}{m} \right)}} \right) \\ &= \left(\frac{1}{m} \sum_{j=1}^m \sqrt{\frac{v^{(m)*} \left(\frac{j-1}{m} \right)}{m}} \right) \left(\frac{1}{m} \sum_{k=1}^m \frac{v^{(m)} \left(\frac{k-1}{m} \right)}{\sqrt{v^{(m)*} \left(\frac{k-1}{m} \right)}} \right) \end{aligned}$$

and again because $v^{(m)*}$ is piecewise constant, this equals $F^{(m)}$ as defined in equation (4). \square

Also define $n(t)$ and F as in equations (3) and (4), but without the superscript (m). Now we can state the main theorem, whose proof, involving many lemmas, is deferred to the Appendix.

Theorem 2. *For $t \in (0, 1)$, $\lim n^{(m)}(t) = n(t)$; also $\lim F^{(m)} = F$.*

With no early stopping, the variance objective of the m th problem would be $V_m^{(m)}$, which is the m -step Riemann approximation to the integral $\int_0^1 v(t) dt$; these converge because we assumed v is continuous on the compact set $[0, 1]$. Therefore the convergence of the optimal variance objectives $F^{(m)}$ to F is equivalent to convergence of the variance reduction ratios $F^{(m)}/V_m^{(m)}$.

The next theorem states that these limits of solutions to discrete optimization problems are also the solutions to a continuous optimization problem that is a limit of the discrete problems. Its proof is similar to that of the discrete result Theorem 1 of Glasserman and Staum (2001b), but simpler and more elegant.

Theorem 3. *If v^* is continuously differentiable, then the limits n and F are the solution and objective of the calculus of variations problem*

$$\min \int_0^1 \frac{v(t)}{n(t)} dt \quad \text{s.t.} \quad \int_0^1 n(t) dt \leq 1$$

and n nondecreasing and continuously differentiable on $(0, 1)$.

Proof. Because n is differentiable, the monotonicity constraint is $n' \leq 0$, so the Lagrangian integrand function is

$$v(t)/n(t) + \nu n(t) + \lambda(t)n'(t).$$

The Euler-Lagrange equation is then

$$-v(t)/n(t)^2 + \nu - \lambda'(t) = 0.$$

Complementary slackness is $\lambda(t)n'(t) = 0$.

We will show that

$$\begin{aligned} \nu &= \left(\int_0^1 \sqrt{v^*(t)} dt \right)^2 \\ n(t) &= \sqrt{\frac{v^*(t)}{\nu}} \\ \lambda(t) &= \frac{V^*(t) - V(t)}{n(t)^2} \end{aligned}$$

satisfy these conditions as well as the constraints. Monotonicity holds because the slope v^* of the upper convex

hull V^* must be nonincreasing. Because v is positive, so is v^* , and as it is continuously differentiable, n must also be continuously differentiable on $(0, 1)$. The budget constraint $\int_0^1 n(t) dt \leq 1$ is satisfied because the numerator of $n(t)$ integrates to \sqrt{v} .

As for complementary slackness, where $V^*(t) = V(t)$, $\lambda(t) = 0$. It remains to consider intervals $[a, b]$ where V^* and V are equal at the endpoints but not on the interior. Here the upper convex hull V^* must be linear, i.e. its second derivative $(v^*)'$ is zero, so n' is also zero, and complementary slackness is satisfied.

Turning to the Euler-Lagrange equation, we differentiate λ to get

$$\begin{aligned} \lambda'(t) &= \frac{v^*(t) - v(t)}{n(t)^2} - \frac{2(V^*(t) - V(t))n'(t)}{n(t)^3} \\ &= v - \frac{v(t)}{n(t)^2} - \frac{2\lambda(t)n'(t)}{n(t)} \end{aligned}$$

By complementary slackness, the third term is zero, and the Euler-Lagrange equation is verified. \square

The assumption that v^* be continuously differentiable is technical: it simplifies the use of the calculus of variations. Likewise the constraint that n be continuously differentiable does not seem to be central to the result.

4 STATE-DEPENDENT STOPPING

The early stopping methodology of Section 2 succeeds when the unconditional variance components v are decreasing. It allocates more computational resources to early steps at which their expenditure produces more variance reduction. Next we investigate how it might be possible to allocate resources on the basis of conditional variance. This is potentially much more effective because the conditional variance of future cashflows often varies greatly across the possible states at one time.

For instance, if a barrier option gets knocked out, then it is certain that all future cashflows are zero, and it is obvious that a simulated path should stop at this step. More generally, when a path reaches a state with low conditional variance, it might be advantageous to do less work in simulating the rest of the path, which will provide little information about the expectation being estimated. But how can a lesser but positive amount of work be done conditional on the current simulated state, when the only choice can be to simulate the next step of this path or not? One possibility is to randomize the decision whether or not to continue simulating this path, with the probability of continuation positively related to the benefit of continuation.

We formalize such a scheme by introducing a new cemetery state Δ to the state space Ω . From the probability measure \mathbf{P} for the state vector path (S_1, \dots, S_m) , we define

a new a probability measure \mathbf{P}^* by

$$\mathbf{P}^*[S_{k+1} \in Q \mid S_k = s] = p(s)\mathbf{P}[S_{k+1} \in Q \mid S_k = s]$$

for any event $Q \subseteq \Omega$ and state $s \in \Omega$,

$$\mathbf{P}^*[S_{k+1} = \Delta \mid S_k = s] = 1 - p(s)$$

for any state $s \in \Omega$, and

$$\mathbf{P}^*[S_{k+1} = \Delta \mid S_k = \Delta] = 1.$$

That is, at each step k of the simulation there is state-dependent probability $1 - p(S_k)$ of entering an absorbing cemetery state, but if that does not happen, the transition probabilities are the same. Define the indicator variable

$$A_k = \mathbf{1}\{S_k \neq \Delta\}$$

which has the value 1 when a path is still ‘‘alive.’’

This is the opposite of the approach of Glasserman and Staum (2001a), where we considered conditioning on the event that the state vector not enter a real cemetery state, for instance, that a barrier option not knock out. There we introduced a scheme to reduce estimator variance, despite doing more work per path, by replacing the indicator variable A_k with a likelihood ratio

$$L_k = \prod_{j=0}^{k-1} p(S_j)$$

which expresses the probability of surviving up to step k given the path (S_1, \dots, S_{k-1}) . Here the goal is to achieve variance reduction, despite increased variance per path, by doing less work per path and thus running more paths with a fixed computational budget.

Let \mathcal{F}_k be the sigma-algebra generated by (S_1, \dots, S_k) and let \mathcal{A}_k be the event $\{A_k = 1\}$.

Lemma 2. *The Radon-Nikodym derivative of the probability measure \mathbf{P} with respect to the restriction of \mathbf{P}^* to \mathcal{A}_k (which is merely a measure) relative to \mathcal{F}_k is $1/L_k$, i.e.*

$$\mathbf{E}^*[XA_k/L_k] = \mathbf{E}[X]$$

for any \mathcal{F}_k -measurable random variable X such that $E[X]$ exists and is finite.

This lemma strongly resembles well-known results of importance sampling, and the technicalities of its proof are analogous to those of Lemma 1 of Glasserman and Staum (2001a).

The state-dependent probability of continuation $p(S_k)$ controls how much expected work will be done on the rest

of a path which has reached state S_k . We can construct an unbiased estimator, using the likelihood ratio L_k , even though the simulation is less likely to reach some states than others. States which are unlikely tend to be reached with a low value of L_k , and receive a correspondingly large weight.

Let $\tau = \max\{k \mid A_k = 1\}$ be the random step at which the simulated path stops. If it is less than m , the number of steps, there is early stopping, and we have not generated the total sum of discounted cashflows $X = \sum_{k=1}^m X_k$, but we have learned something about it. Let Y_τ represent a guess at the expected value of X given this information; this guess could be a sophisticated approximation based on knowledge about the simulation problem, or as simple as $\sum_{k=1}^\tau X_k$, or even just zero. That is, for each value of k , $Y_k = Y_k(S_1, \dots, S_k)$ is a deterministic function of the path to date, which is generally not equal to the unknown conditional expectation $\mathbf{E}[X \mid \mathcal{F}_k]$. However, we require $Y_m = X$.

An estimator of $\mathbf{E}[X]$ under \mathbf{P}^* is

$$X^* = \frac{Y_\tau}{L_\tau} - \sum_{j=1}^{\tau-1} \left(1 - \frac{1}{p(S_j)}\right) \frac{Y_j}{L_j}.$$

The intuition is that Y_τ is our best guess at what X would be if we were to continue the simulation rather than stopping at step τ . This guess deserves a weight $1/L_\tau$ to make up for the hazard of early stopping before reaching this step. To avoid bias, we must subtract something to make up for the guesses that we would have made had we stopped at earlier steps.

Although this way of writing the estimator is most intuitive, another expression is more useful for theory. Let Y_0 denote 0.

Lemma 3. *The estimator X^* also equals*

$$\sum_{j=1}^m \frac{A_j}{L_j} (Y_j - Y_{j-1}).$$

Proof. The indicator function for the event that the path stops at step j is $A_j - A_{j+1}$. That is, $A_\tau - A_{\tau+1} = 1$ while $j \neq \tau$ implies $A_j - A_{j+1} = 0$. Therefore

$$\begin{aligned} X^* &= \frac{Y_\tau}{L_\tau} - \sum_{j=1}^{\tau-1} \left(1 - \frac{1}{p(S_j)}\right) \frac{Y_j}{L_j} \\ &= \frac{Y_\tau}{L_\tau} + \sum_{j=1}^{\tau-1} Y_j \left(\frac{1}{L_j} - \frac{1}{L_{j+1}}\right) \\ &= \sum_{j=1}^m \frac{A_j}{L_j} (Y_j - Y_{j-1}). \quad \square \end{aligned}$$

Using this, we can prove the key theorem.

Theorem 4. *The estimator X^* is unbiased under the \mathbf{P}^* simulation scheme, i.e. $\mathbf{E}^*[X^*] = \mathbf{E}[X]$.*

Proof. Using the formulation of Lemma 3,

$$\begin{aligned} \mathbf{E}^*[X^*] &= \mathbf{E}^* \left[\sum_{j=1}^m \frac{A_j}{L_j} (Y_j - Y_{j-1}) \right] \\ &= \sum_{j=1}^m \mathbf{E}^*[Y_j A_j / L_j - Y_{j-1} A_j / L_j] \\ &= \sum_{j=1}^m \mathbf{E}[Y_j] - \mathbf{E}[Y_{j-1}] \\ &= \mathbf{E}[Y_m] - \mathbf{E}[Y_0] = \mathbf{E}[X] - 0. \end{aligned}$$

The third equality is justified by Lemma 2 since Y_j and Y_{j-1} are both \mathcal{F}_j -measurable. \square

This approach is closely related to the methodology of introducing new variables X'_1, \dots, X'_m , described at the end of Section 2. Just as there, the optimal choice of Y_j is the unknown conditional expectation $\mathbf{E}[X \mid \mathcal{F}_j]$, and in practice one must attempt to use knowledge about the simulation problem to construct a good approximation. Here we also face two more challenges: finding a good randomized stopping policy \mathbf{P}^* , and designing a computationally efficient algorithm.

The problem of finding the variance-minimizing \mathbf{P}^* given a policy of guesses Y_1, \dots, Y_m appears difficult. However, preliminary investigations suggest that the optimal continuation probabilities $p(S_k)$ are large when proceeding to step $k + 1$ from state S_k yields a large reduction in the sum of the conditional variance of remaining discounted payoffs and the squared bias of the guess for the value of remaining discounted payoffs. As before, these quantities are unknown, but a decent guess at them produces an unbiased estimator with reduced variance that is merely suboptimal in the sense of not providing the maximum possible variance reduction.

5 CONCLUSIONS

In this paper we provided further analysis and an extension of the idea of stopping simulated paths early as a variance reduction technique. We proved that when paths stop early independent of the state reached, the extent of variance reduction depends on the shape of the covariance matrix of discounted cashflows, not on its size. The method is appropriate for problems where simulation produces more valuable information at early time steps than later ones. We also showed how to construct an unbiased estimator when the decision to stop paths early is random and depends on

the state reached. This opens up the possibility of greater variance reduction, if one can avoid doing work simulating the tails of paths when they are unlikely to produce much valuable information. These advances increase our ability to use our knowledge about the structure of simulation problems to apply computational resources efficiently.

ACKNOWLEDGMENTS

This research was supported in part by NSF grant DMS0074637, an IBM University Faculty Award, and an NSF graduate research fellowship.

APPENDIX: PROOF OF THEOREM 2

Lemma 4. *The sequence of functions $V^{(m)}$ converges uniformly to V .*

Proof. Each $V^{(m)}(t)$ is a Riemann approximation to the integral $V(t)$. Since v is continuous on the compact set $[0, 1]$, it is uniformly continuous there, and it is Riemann-integrable. Uniform continuity says

$$\forall \epsilon > 0, \exists \delta \ni \forall s, t \in [0, 1], |t-s| < \delta \Rightarrow |v(t)-v(s)| < \epsilon$$

Therefore by choosing $s = [tm]/m$, it follows that $|t-s| < 1/m$, and

$$\forall \epsilon > 0, \exists \delta \ni \forall t \in [0, 1], m > 1/\delta \Rightarrow$$

$$|v^{(m)}(t) - v(t)| = |v([tm]/m) - v(t)| < \epsilon$$

This is uniform convergence of $v^{(m)}$ to v . Next,

$$\forall \epsilon > 0, \exists \delta \ni \forall t \in [0, 1], m > 1/\delta \Rightarrow$$

$$|V^{(m)}(t) - V(t)| \leq \int_0^t |v^{(m)}(s) - v(s)| ds < t\epsilon \leq \epsilon$$

That is, $V^{(m)}$ converges uniformly to V . \square

Lemma 5. *The sequence of functions $V^{(m)*}$ converges uniformly to V^* .*

Proof. To prove the convergence of $V^{(m)*}$ to V^* , define $\epsilon^{(m)} = \sup_{t \in [0,1]} |V^{(m)}(t) - V(t)|$. By uniform convergence of $V^{(m)}$ to V , $\epsilon^{(m)}$ converges to zero. Because $V^{(m)*}(t)$ is concave, $V^{(m)*}(t) + \epsilon^{(m)}$ is concave. Further, using $V^{(m)*}(t) \geq V^{(m)}(t)$,

$$V^{(m)*}(t) + \epsilon^{(m)} \geq V^{(m)*}(t) - V^{(m)}(t) + V(t) \geq V(t)$$

Therefore, by definition of $V^*(t)$ as the least concave curve above $V(t)$, $V^*(t) \leq V^{(m)*}(t) + \epsilon^{(m)}$. The exact same

reasoning, switching the role of V^* and $V^{(m)*}$, shows that $V^{(m)*}(t) \leq V^*(t) + \epsilon^{(m)}$. Combining these,

$$\sup_{t \in [0,1]} |V^*(t) - V^{(m)*}(t)| \leq \epsilon^{(m)}$$

and thus $V^{(m)*}$ converges uniformly to V^* . \square

Lemma 6. *There exist nonincreasing functions $v^{(m)*}$ and v^* such that $V^{(m)*}(t) = \int_0^t v^{(m)*}(s) ds$ and $V^*(t) = \int_0^t v^*(s) ds$. Moreover, for $t \in (0, 1)$, $\lim v^{(m)*}(t) = v^*(t)$.*

Proof. As concave functions, $V^{(m)*}$ and V^* have both left and right derivatives at every point in $(0, 1)$. If $s < t < u$, focusing on V^* ,

$$\begin{aligned} \lim_{x \rightarrow s+} \frac{V^*(x) - V^*(s)}{x - s} &\geq \lim_{x \rightarrow t-} \frac{V^*(t) - V^*(x)}{t - x} \\ &\geq \lim_{x \rightarrow t+} \frac{V^*(x) - V^*(t)}{x - t} \\ &\geq \lim_{x \rightarrow u-} \frac{V^*(u) - V^*(x)}{u - x} \end{aligned}$$

The left and right derivatives are unequal only on a set of measure zero. For these assertions, see, for instance, Sundaram (1996, §7.2.2). It follows that any function which at each point equals either one-sided derivative is nonincreasing and integrates to V^* . The same holds for $V^{(m)*}$, which indeed has at most m points where the one-sided derivatives are unequal.

Next, for the convergence of $v^{(m)*}$ to v^* . By Lemma 5, we know that for any open interval $(a, b) \subset (0, 1)$,

$$\begin{aligned} \int_a^b (v^*(t) - v^{(m)*}(t)) dt &= \quad (A-1) \\ (V^*(b) - V^{(m)*}(b)) - (V^{(m)*}(a) - V^*(a)) &\rightarrow 0 \end{aligned}$$

The strategy is to show that

$$\exists t \in (0, 1) \ni \left\{ \begin{aligned} \liminf v^{(m)*}(t) &< v^*(t) \\ \text{or } \limsup v^{(m)*}(t) &> v^*(t) \end{aligned} \right\} \quad (A-2)$$

implies the contrary of (A-1), and thus (A-2) is false, i.e. for every $t \in (0, 1)$,

$$\liminf v^{(m)*}(t) \geq v^*(t) \geq \limsup v^{(m)*}(t)$$

and thus $\lim v^{(m)*}(t)$ exists and equals $v^*(t)$.

Suppose then that there is an $a \in (0, 1)$ such that $\liminf v^{(m)*}(a) < v^*(a)$. Define $\epsilon = v^*(a) - \liminf v^{(m)*}(a) > 0$. By definition of \liminf , there is a sequence $\{m_k\}$ such that for all k , $v^{(m_k)*}(a) < v^*(a) - \epsilon/2$. Because v^* is continuous and nonincreasing and $a < 1$,

there is a $b \in (a, 1)$ such that $0 < v^*(a) - v^*(b) < \epsilon/4$. Both v^* and $v^{(m_k)^*}$ are nonincreasing, so for $t \in (a, b)$,

$$\begin{aligned} v^*(t) - v^{(m_k)^*}(t) &\geq v^*(b) - v^{(m_k)^*}(a) \\ &= v^*(b) - v^*(a) + v^*(a) - v^{(m_k)^*}(a) \\ &> -\epsilon/4 + \epsilon/2 = \epsilon/4 > 0 \end{aligned}$$

Therefore

$$\int_a^b (v^*(t) - v^{(m_k)^*}(t)) dt > (b - a)\epsilon/4 > 0$$

for all m_k , thus $\liminf \int_a^b (v^*(t) - v^{(m)^*}(t)) dt > 0$, contradicting (A-1).

On the other hand, suppose there is a $b \in (0, 1)$ such that $\limsup v^{(m)^*}(b) > v^*(b)$. Now let ϵ be $\limsup v^{(m)^*}(b) - v^*(b) > 0$, and there is a sequence $\{m_k\}$ such that for all k , $v^{(m_k)^*}(b) > v^*(b) + \epsilon/2$. This time there is an $a \in (0, b)$ such that $0 < v^*(a) - v^*(b) < \epsilon/4$, and we get for $t \in (a, b)$

$$\begin{aligned} v^*(t) - v^{(m_k)^*}(t) &\leq v^*(a) - v^{(m_k)^*}(b) \\ &= v^*(a) - v^*(b) + v^*(b) - v^{(m_k)^*}(b) \\ &< \epsilon/4 - \epsilon/2 = -\epsilon/4 < 0 \end{aligned}$$

for all m_k , thus $\limsup \int_a^b (v^*(t) - v^{(m)^*}(t)) dt < 0$, contradicting (A-1). \square

Lemma 7. *The functions $v(t)$, $v^{(m)}(t)$, $v^*(t)$, and $v^{(m)^*}(t)$ are bounded above and bounded away from zero in m and t .*

Proof. Because v is positive and continuous, it is both bounded above and bounded away from zero on $[0, 1]$. By the construction of $v^{(m)}$, its values are values of v , so it too is both bounded above and bounded away from zero. Since v^* and $v^{(m)^*}$ are nonincreasing and positive, we can focus on $t = 0$ for proving their boundedness above and on $t = 1$ for proving their boundedness away from zero.

If there is a (one-sided) neighborhood of $t = 0$ on which $V^* = V$, then $v^*(0) = v(0)$. If not, then there exists a point $t > 0$ such that $V^*(t) = V(t) = tv^*(0)$. Either way, $v^*(0)$ is finite and the function v^* is bounded above. Likewise, if there is a neighborhood of $t = 0$ on which $V^{(m)^*} = V^{(m)}$, then $v^{(m)^*}(0) = v^{(m)}(0) = v(0)$. If not, there exists a point $i/m > 0$ such that

$$\begin{aligned} (i/m)v^{(m)^*}(0) &= V^{(m)^*}(i/m) = V^{(m)}(i/m) \\ &= \int_0^{i/m} v^{(m)}(s) ds = \frac{1}{m} \sum_{j=0}^{i-1} v(j/m) \end{aligned}$$

So $v^{(m)^*}(0)$ is a Riemann average of v on the interval $[0, i/m]$, and this is bounded above in m because v is

bounded above on the whole interval $[0, 1]$. Either way, $v^{(m)^*}(0)$ is bounded above, and so is the function $v^{(m)^*}$.

The same reasoning applies to boundedness away from zero. If there is a neighborhood of $t = 1$ on which $V^* = V$, then $v^*(1) = v(1)$, and if not, there exists $t < 1$ such that $V^*(1) - V^*(t) = V(1) - V(t) = (1 - t)v^*(1)$. If there is a neighborhood of $t = 1$ on which $V^{(m)^*} = V^{(m)}$, then $v^{(m)^*}(1) = v^{(m)}(1) = v(1)$, and if not, for some i we have $v^{(m)^*}(1) = (\sum_{j=i}^{m-1} v(j/m))/(m - i)$. \square

Proof of Theorem 2. This proof frequently relies on Lemma 7. The functions $v^{(m)^*}$ and v^* are bounded away from zero, so we can take square roots, which are themselves bounded away from zero. Where $v^{(m)^*} \rightarrow v^*$, trivially $\sqrt{v^{(m)^*}} \rightarrow \sqrt{v^*}$. Since these are also positive and bounded above, by the Bounded Convergence Theorem, $\int_0^t \sqrt{v^{(m)^*}(s)} ds \rightarrow \int_0^t \sqrt{v^*(s)} ds$. Consequently, $n^{(m)}$ converges to n on $(0, 1)$. Also $v^{(m)}/\sqrt{v^{(m)^*}} \rightarrow v/\sqrt{v^*}$ because the denominators are bounded away from zero, and the numerators and denominators converge separately. With numerators positive and bounded above, and denominators bounded away from zero, these fractions are positive and bounded above. So again we can apply the Bounded Convergence Theorem, and $F^{(m)}$ converges to F . \square

REFERENCES

- Glasserman, P., and J. Staum. 2001a. Conditioning on one-step survival for barrier option simulations. *Operations Research*, forthcoming.
- Glasserman, P., and J. Staum. 2001b. Resource allocation among simulation time steps. Working paper, Graduate School of Business, Columbia University, New York, New York.
- Sundaram, R. K. 1996. *A First Course in Optimization Theory*. New York: Cambridge University Press.

AUTHOR BIOGRAPHIES

PAUL GLASSERMAN is the Jack R. Anderson Professor and chairman in the division of Decision, Risk, and Operations at Columbia Business School, and holds a secondary appointment in the Department of Industrial Engineering and Operations Research. He currently serves as departmental editor of *Management Science* for Stochastic Models and Simulation. His e-mail address is <pg20@columbia.edu> and web address is <www.paulglasserman.com>.

JEREMY STAUM is a Visiting Assistant Professor in the School of Operations Research and Industrial Engineering at Cornell University. He received his Ph.D. from Columbia University in 2001. His research interests include variance reduction techniques and financial engineering.