
Simulation for American Options: Regression Now or Regression Later?

Paul Glasserman and Bin Yu

Columbia University, Graduate School of Business, New York, NY 10027
pg20@columbia.edu, by52@columbia.edu

Summary. Pricing American options requires solving an optimal stopping problem and therefore presents a challenge for simulation. This article investigates connections between a weighted Monte Carlo technique and regression-based methods for this problem. The weighted Monte Carlo technique is shown to be equivalent to a least-squares method in which option values are regressed at a later time than in other regression-based methods. This “regression later” technique is shown to have two attractive features: under appropriate conditions, (i) it results in less-dispersed estimates, and (ii) it provides a dual estimate (an upper bound) with modest additional effort. These features result, more generally, from using martingale regressors.

1 Introduction

At the MCQMC 2002 conference in Singapore we presented work on weighted Monte Carlo estimators reported in Glasserman and Yu [9]. That work was motivated by applications in finance, including the model calibration technique of Avellaneda et al. [2, 3] and a method for pricing American options proposed by Broadie, Glasserman, and Ha (BGH) [5]. The method of BGH [5] uses weights in approximating the dynamic programming problem involved in the calculation of American option prices. The weights are chosen to minimize a convex objective subject to linear constraints.

One of the implications of the general analysis of weighted Monte Carlo estimators presented at the conference and in [9] is an equivalence between such estimators and regression-based estimators when the convex objective used is quadratic. This, then, raises a question about the connection between the weighted method of BGH [5] and regression-based methods for pricing American options proposed in Carrière [6], Longstaff and Schwartz [11], and Tsitsiklis and Van Roy [13]. The purpose of this article is to develop this connection. Several presentations at the Singapore conference addressed other aspects of these methods.

When reformulated using least-squares regression, the weighted Monte Carlo method of BGH [5] differs from other methods in how it combines

regression with the backward induction required for American options. In particular, we contrast methods that, at time i , regress option values from time $i + 1$ against basis function values at time i (regression now) with methods that regress against basis function values at time $i + 1$ (regression later). We show that “regression later” requires some stronger conditions but has two benefits: it produces less-dispersed estimates, and it provides a dual estimate (an upper bound on the option price) with modest additional computational effort.

2 Optimal Stopping Problem

We consider the following class of problems. An \mathfrak{R}^d -valued Markov chain X_0, X_1, \dots, X_m (with X_0 fixed) records all relevant financial information, including the prices of underlying assets and any variables affecting the dynamics of the underlying assets. If exercised at time i , $i = 0, 1, \dots, m$, the option pays $h_i(X_i)$, for some known functions h_0, h_1, \dots, h_m mapping \mathfrak{R}^d into $[0, \infty)$. Let \mathcal{T}_i denote the set of randomized stopping times (as defined below) taking values in $\{i, i + 1, \dots, m\}$ and define

$$V_i^*(x) = \sup_{\tau \in \mathcal{T}_i} \mathbf{E}[h_\tau(X_\tau) | X_i = x], \quad x \in \mathfrak{R}^d, \quad (1)$$

for $i = 0, 1, \dots, m$. Then $V_i^*(x)$ is the value of the option at date i in state x , given that the option was not exercised at $0, 1, \dots, i - 1$. It is also the value of a new option issued at date i in state x . Our objective is to find $V_0^*(X_0)$.

Restricting τ to be an ordinary stopping time would mean requiring that each event $\{\tau = i\}$ be determined by X_1, \dots, X_i . In allowing randomized stopping times we are allowing such an event to depend also on other random variables independent of X_{i+1}, \dots, X_m . This extension is needed to accommodate stopping rules estimated by simulation.

The option values satisfy the dynamic programming equations

$$V_m^*(x) = h_m(x) \quad (2)$$

$$V_i^*(x) = \max\{h_i(x), \mathbf{E}[V_{i+1}^*(X_{i+1}) | X_i = x]\}, \quad (3)$$

$i = 0, 1, \dots, m - 1$. Most methods for estimating $V_0^*(X_0)$ by simulation rely on approximating this recursion in some way. We have not included discount factors in (1) and (2)–(3), but this formulation is sufficiently general to encompass discounted payoffs through appropriate definition of the X_i and h_i , as explained in Section 8.1 of Glasserman [8]. Also, the dynamic programming equations can be written in terms of continuation values

$$C_i^*(x) = \mathbf{E}[V_{i+1}^*(X_{i+1}) | X_i = x], \quad i = 0, 1, \dots, m - 1,$$

as

$$C_m^*(x) = 0 \tag{4}$$

$$C_i^*(x) = \mathbf{E}[\max\{h_{i+1}(X_{i+1}), C_{i+1}^*(X_{i+1})\} | X_i = x], \tag{5}$$

$i = 0, 1, \dots, m - 1$. The option values satisfy

$$V_i^*(x) = \max\{h_i(x), C_i^*(x)\},$$

so these can be calculated from the continuation values.

3 Approximate Dynamic Programming

The methods we consider apply approximate versions of the dynamic programming recursions (2)–(3) or (4)–(5). These methods approximate the option values V_i^* or the continuation values C_i^* as linear combinations of basis functions. For each $i = 1, \dots, m$, let ψ_{ik} , $k = 0, \dots, K$, be functions from \mathfrak{R}^d to \mathfrak{R} and consider approximations of the form

$$V_i^*(x) \approx \sum_{k=0}^K \beta_{ik} \psi_{ik}(x)$$

and

$$C_i^*(x) \approx \sum_{k=0}^K \gamma_{ik} \psi_{ik}(x),$$

for some constants β_{ik} and γ_{ik} . Working with approximations of this type reduces the problem of finding the functions V_i^* or C_i^* to one of finding the coefficients β_{ik} or γ_{ik} .

These approximations can be made precise through the least-squares projection onto the span of $\psi_{ik}(X_i)$, $k = 0, 1, \dots, K$. Set $\psi_i = (\psi_{i0}, \dots, \psi_{iK})^\top$. For any square-integrable random variable Y define the projection

$$\Pi_i Y = \mathbf{E}[Y \psi_i(X_i)^\top] (\mathbf{E}[\psi_i(X_i) \psi_i(X_i)^\top])^{-1} \psi_i(X_i).$$

Thus,

$$\Pi_i Y = \sum_{k=0}^K a_k \psi_{ik}(X_i) \tag{6}$$

with

$$(a_0, \dots, a_K) = \mathbf{E}[Y \psi_i(X_i)^\top] (\mathbf{E}[\psi_i(X_i) \psi_i(X_i)^\top])^{-1} \tag{7}$$

and the residual $Y - \Pi_i Y$ is uncorrelated with $\psi_{i0}(X_i), \dots, \psi_{iK}(X_i)$. In a slight abuse of notation, we also write

$$(\Pi_i Y)(x) = \sum_{k=0}^K a_k \psi_{ik}(x)$$

for the function defined by the coefficients (7). These definitions require that the matrix $\mathbb{E}[\psi_i(X_i)\psi_i(X_i)^\top]$ be finite and nonsingular, which we assume throughout. In fact, we impose the following condition:

(C1). For each $i = 1, \dots, m$, $\psi_{i0} \equiv 1$, $\mathbb{E}[\psi_{ik}(X_i)] = 0$, $k = 1, \dots, K$, and

$$\mathbb{E}[\psi_i(X_i)\psi_i(X_i)^\top] = \begin{pmatrix} 1 & & & & \\ & \sigma_{i1}^2 & & & \\ & & \sigma_{i2}^2 & & \\ & & & \ddots & \\ & & & & \sigma_{iK}^2 \end{pmatrix},$$

with $0 < \sigma_{ik}^2 < \infty$ for all i, k .

The important point is that the basis variables have finite variance and are linearly independent. The further requirement that they be uncorrelated can then always be arranged through a linear transformation.

3.1 Regression Now

Define an approximation to (4)–(5) as follows:

$$C_m(x) = 0 \tag{8}$$

$$C_i(x) = (\Pi_i \max\{h_{i+1}(X_{i+1}), C_{i+1}(X_{i+1})\})(x). \tag{9}$$

As in (6), the application of the projection Π_i results in a linear combination of the basis functions, so

$$C_i(x) = \sum_{k=0}^K \beta_{ik} \psi_{ik}(x) \tag{10}$$

with $\beta_i^\top = (\beta_{i0}, \dots, \beta_{iK})$ defined as in (7) with Y replaced by

$$V_{i+1}(X_{i+1}) \equiv \max\{h_{i+1}(X_{i+1}), C_{i+1}(X_{i+1})\}.$$

Write

$$V_{i+1}(X_{i+1}) = \sum_{k=0}^K \beta_{ik} \psi_{ik}(X_i) + \epsilon_{i+1} \tag{11}$$

by defining the residual ϵ_{i+1} so that this holds. A sufficient condition for the approximation (8)–(9) to be exact is

(C2). For all $i = 0, \dots, m-1$, $\mathbb{E}[\epsilon_{i+1}|X_i] = 0$.

Proposition 1. *If (C2) holds, then $V_i = V_i^*$ for all $i = 0, 1, \dots, m$.*

Proof. Observe that $C_m = C_m^*$ so $V_m = V_m^*$. Now suppose that $C_{i+1} = C_{i+1}^*$ for some i . Then $V_{i+1} = V_{i+1}^*$ and (C2) implies that

$$\begin{aligned} C_i^*(x) &= \mathbb{E}[V_{i+1}^*(X_{i+1})|X_i = x] = \mathbb{E}[V_{i+1}(X_{i+1})|X_i = x] \\ &= \sum_{k=0}^K \beta_{ik} \psi_{ik}(x) = C_i(x). \end{aligned}$$

The result now follows by induction. \square

If (C2) fails to hold, (9) may still provide a useful approximation. Computing (9) is difficult but lends itself to further approximation through simulation. For $j = 1, \dots, b$ let (X_{1j}, \dots, X_{mj}) be independent replications of the underlying Markov chain. From these paths define a sample version of the projection Π_i through ordinary least-squares regression. In particular, set $\hat{C}_m = 0$,

$$\hat{C}_i(x) = \sum_{k=0}^K \hat{\beta}_{ik} \psi_{ik}(x) \tag{12}$$

with $\hat{\beta}_i^\top = (\hat{\beta}_{i0}, \dots, \hat{\beta}_{iK})$ the vector of regression coefficients

$$(\hat{\beta}_{i0}, \dots, \hat{\beta}_{iK}) = \left(\sum_{j=1}^b \hat{V}_{i+1}(X_{i+1,j}) \psi_i(X_{ij})^\top \right) \left(\sum_{j=1}^b \psi_i(X_{ij}) \psi_i(X_{ij})^\top \right)^{-1} \tag{13}$$

and

$$\hat{V}_{i+1} = \max\{h_{i+1}, \hat{C}_{i+1}\},$$

$i = 0, 1, \dots, m - 1$. Because the initial state X_0 is fixed, we set

$$\hat{C}_0(X_0) = \frac{1}{b} \sum_{j=1}^b \hat{V}_1(X_{1j}) \tag{14}$$

and $\hat{V}_0(X_0) = \max\{h_0(X_0), \hat{C}_0(X_0)\}$. Tsitisklis and Van Roy [13] prove convergence of the \hat{C}_i to the C_i as the number of paths b increases. Glasserman [8, §8.6.2] shows that this method corresponds to using a particular set of weights in the stochastic mesh method of Broadie and Glasserman [4]. Clément et al. [7] prove convergence of a related method in Longstaff and Schwartz [11].

3.2 Regression Later

Broadie, Glasserman, and Ha [5] develop a method for pricing American options by simulation in which the conditional expectations in the dynamic programming recursion are approximated using weighted averages of simulated downstream option values. The weights are chosen to minimize a separable

convex objective function subject to constraints. The constraints ensure that the weights correctly compute certain known conditional expectations.

To translate that method to the notation of this article, define

$$\bar{\psi}_i(x) = \mathbb{E}[\psi_{i+1}(X_{i+1})|X_i = x].$$

As before, let (X_{1j}, \dots, X_{mj}) , $j = 1, \dots, b$, denote independent replications of the underlying Markov chain. For each $j = 1, \dots, b$ consider the optimization problem

$$\min_{w_{j1}, \dots, w_{jb}} \sum_{\ell=1}^b w_{j\ell}^2 \quad (15)$$

$$\text{subject to } \sum_{\ell=1}^b w_{j\ell} \psi_{i+1}(X_{i+1,\ell}) = \bar{\psi}_i(X_{ij}). \quad (16)$$

The constraint (16) ensures that the weighted average of the basis function values one step ahead equals their conditional expectation evaluated at X_{ij} . (Because $\psi_{0,i+1} \equiv 1$, it also implies that the weights sum to 1.) The objective in (15) may be viewed as choosing a maximally uniform set of weights from the feasible set. BGH [5] also consider a maximum entropy objective for choosing the weights. Given optimal weights w_{j1}, \dots, w_{jb} , the continuation value at X_{ij} is estimated as

$$\hat{C}_i^+(X_{ij}) = \sum_{\ell=1}^b w_{j\ell} \hat{V}_{i+1}^+(X_{i+1,\ell}), \quad i = 1, \dots, m-1, \quad (17)$$

with $\hat{C}_m^+ \equiv 0$,

$$\hat{V}_{i+1}^+ = \max\{h_{i+1}, \hat{C}_{i+1}^+\}, \quad i = 0, 1, \dots, m-1,$$

and $\hat{C}_0^+(X_0)$ computed as in (14).

It follows from a general analysis of weighted Monte Carlo estimators in Glasserman and Yu [9] that (17) has a regression interpretation. Let $\hat{\gamma}_i = (\hat{\gamma}_{i0}, \dots, \hat{\gamma}_{iK})$ denote the vector of regression coefficients

$$\begin{aligned} &(\hat{\gamma}_{i0}, \dots, \hat{\gamma}_{iK}) = \\ &\left(\sum_{j=1}^b \hat{V}_{i+1}^+(X_{i+1,j}) \psi_{i+1}(X_{i+1,j})^\top \right) \left(\sum_{j=1}^b \psi_{i+1}(X_{i+1,j}) \psi_{i+1}(X_{i+1,j})^\top \right)^{-1} \end{aligned} \quad (18)$$

and note that these are defined by regression against $\psi_{i+1}(X_{i+1,j})$, $j = 1, \dots, b$, whereas (13) uses $\psi_i(X_{ij})$, $j = 1, \dots, b$. Thus, (13) uses current basis functions and (18) uses later basis functions.

Proposition 2. *The BGH [5] estimator (17) admits the representation*

$$\hat{C}_i^+(X_{ij}) = \sum_{k=0}^K \hat{\gamma}_{ik} \bar{\psi}_{ik}(X_{ij}), \quad (19)$$

$i = 1, \dots, m-1$.

Proof. Proposition 1 and Theorem 2 of Glasserman and Yu [9] apply to pairs (Z_j, Y_j) , $j = 1, \dots, b$, with (row vectors) $Z_j \in \mathfrak{R}^n$ and $Y_j \in \mathfrak{R}$. They show that if w_1, \dots, w_b are chosen to minimize $w_1^2 + \dots + w_b^2$ subject to constraints

$$\sum_{j=1}^b w_j = 1, \quad \sum_{j=1}^b w_j Z_j = z,$$

for some (row vector) $z \in \mathfrak{R}^n$, then

$$\sum_{j=1}^b w_j Y_j = (1, z) \hat{\alpha},$$

with $\hat{\alpha}$ the (column) vector of coefficients obtained by least-squares regression of Y_1, \dots, Y_b against $(1, Z_1), \dots, (1, Z_b)$. Equation (19) follows once we identify Z_j with $(\psi_{i+1,1}(X_{i+1,j}), \dots, \psi_{i+1,K}(X_{i+1,j}))$, Y_j with $\hat{V}_{i+1}^+(X_{i+1,j})$, and z with $(\psi_{i1}(X_{ij}), \dots, \psi_{iK}(X_{ij}))$. \square

To further develop the connection between the weighted estimator (17) and regression, we impose the following stronger condition on the basis functions:

(C3). Martingale property: $\mathbb{E}[\psi_{i+1}(X_{i+1})|X_i] = \psi_i(X_i)$, $i = 0, 1, \dots, m-1$.

By the Markov property, $\mathbb{E}[\psi_{i+1}(X_{i+1})|X_i] = \mathbb{E}[\psi_{i+1}(X_{i+1})|X_1, \dots, X_i]$. Condition (C3) implies that $\bar{\psi}_i = \psi_i$ and comparison of (12) and (19) then shows that \hat{C}_i and \hat{C}_i^+ are linear combinations of the same basis functions. They differ only in the estimates of the coefficients they use.

To clarify what is being estimated by \hat{C}_i^+ , define $C_m^+ = 0$ and

$$C_i^+(x) = \sum_{k=0}^K \gamma_{ik} \psi_{ik}(x) = \mathbb{E} \left[\sum_{k=0}^K \gamma_{ik} \psi_{i+1,k}(X_{i+1}) | X_i = x \right] \quad (20)$$

$$= \mathbb{E} \left[(\Pi_{i+1} V_{i+1}^+(X_{i+1})) (X_{i+1}) | X_i = x \right] \quad (21)$$

$$= (\Pi_i \Pi_{i+1} V_{i+1}^+(X_{i+1})) (x) \quad (22)$$

with $V_{i+1}^+ = \max\{h_{i+1}, C_{i+1}^+\}$, $i = 0, \dots, m-1$. That the projection Π_i in (22) has the same effect as the conditional expectation in (21) is a consequence of (C3). Also, write

$$V_{i+1}^+(X_{i+1}) = \sum_{k=0}^K \gamma_{ik} \psi_{i+1,k}(X_{i+1}) + \epsilon_{i+1}^+ \quad (23)$$

with $\epsilon_{i+1}^+ = V_{i+1}^+(X_{i+1}) - \Pi_{i+1} V_{i+1}^+(X_{i+1})$ uncorrelated with the components of $\psi_{i+1}(X_{i+1})$. The usual regression estimate of the coefficient vector γ_i in (20), computed from b simulated paths, is exactly the $\hat{\gamma}_i$ in (18), so \hat{C}_i^+ is indeed a simulation estimate of C_i^+ ; the two stand in the same relation to each other as \hat{C}_i and C_i . We use the superscript “+” to emphasize that the coefficients in C_i^+ are obtained by regression against $\psi_{i+1}(X_{i+1})$ rather than $\psi_i(X_i)$. The step back to i is taken by the conditional expectation, via (C3).

We consider two conditions on the residuals ϵ_{i+1}^+ :

(C4). For all $i = 0, \dots, m-1$ and $k = 0, \dots, K$, $\mathbb{E}[\epsilon_{i+1}^+(\psi_{i+1,k}(X_{i+1}) - \psi_{ik}(X_i))] = 0$.

(C4’). For all $i = 0, \dots, m-1$, $\mathbb{E}[\epsilon_{i+1}^+ | X_i] = 0$.

Under (C3), the first of these states that the residuals are uncorrelated with the martingale differences. The second of these parallels (C2). Because ϵ_{i+1}^+ in (23) is uncorrelated with $\psi_{i+1,k}(X_{i+1})$, (C4’) implies (C4) if (C3) holds.

Proposition 3. *If (C3) and (C4) hold, then $C_i^+ = C_i$ for all i . If (C3) and (C4’) hold then in addition $C_i^+ = C_i^*$ for all i .*

Proof. Geometrically, the first part says that $\Pi_i \Pi_{i+1} V_{i+1}^+ = \Pi_i V_{i+1}^+$ when $V_{i+1}^+ - \Pi_{i+1} V_{i+1}^+$ is orthogonal to the space onto which Π_i projects. More explicitly, observe that $(C_m^+, V_m^+) = (C_m, V_m)$ and to argue by induction suppose that $V_{i+1}^+ = V_{i+1}$. From (23) we get

$$V_{i+1}^+(X_{i+1}) = \sum_{k=0}^K \gamma_{ik} \psi_{ik}(X_i) + \left\{ \sum_{k=0}^K \gamma_{ik} [\psi_{i+1,k}(X_{i+1}) - \psi_{ik}(X_i)] + \epsilon_{i+1}^+ \right\}$$

and if (C3)–(C4) hold then this decomposes $V_{i+1}^+(X_{i+1})$ as a linear combination of $\psi_i(X_i)$ and a term uncorrelated with $\psi_i(X_i)$. But since (11) does the same thing, we must have $\gamma_i = \beta_i$ and

$$\epsilon_{i+1} = \left\{ \sum_{k=0}^K \gamma_{ik} [\psi_{i+1,k}(X_{i+1}) - \psi_{ik}(X_i)] + \epsilon_{i+1}^+ \right\}. \quad (24)$$

Equality of the coefficients implies that $C_i^+ = C_i$ and then $V_i^+ = V_i$, concluding the induction. For the second assertion, apply (C3) and (C4’) to (24) to see that (C2) holds so Proposition 1 applies. \square

4 Comparison

We now turn to a comparison of the methods in Sections 3.1 and 3.2. Proposition 3 gives conditions under which the approximations C_i^+ and C_i are the same. But even when these conditions hold the simulation estimates \hat{C}_i^+ and \hat{C}_i are different, and it is natural to compare properties of these estimates.

The key difference is that \hat{C}_i^+ uses the estimates $\hat{\gamma}_i$ in (18) obtained by regressing against $\psi_{i+1}(X_{i+1})$ (regression later) whereas \hat{C}_i uses the estimates $\hat{\beta}_i$ in (13) obtained by regressing against $\psi_i(X_i)$ (regression now).

Intuitively, we expect “regression later” to give better results than “regression now” because the option values at time $i + 1$ should be more highly correlated with the basis functions at time $i + 1$ than with the basis functions at time i . Also, \hat{C}_i^+ takes advantage of the martingale property (C3) to compute the conditional expectation in (20) *exactly*, using simulation only to approximate Π_{i+1} . In contrast, with \hat{C}_i simulation is implicitly used for both steps when it is used to approximate Π_i . We now formulate a precise result.

Observe that (13) and (18) involve regressing different estimates of the option values at time $i + 1$: in (13) we have \hat{V}_{i+1} whereas in (18) we have \hat{V}_{i+1}^+ . In order to compare the two different ways of estimating coefficients, for the rest of this section we will suppose that the two methods regress the same values, and to be concrete we take these to be values of V_{i+1} . In effect, we are comparing two algorithms that proceed identically and exactly (as in (8)–(9)) backwards from time m to $i + 1$, and then use two different simulation estimates from $i + 1$ to i . We continue to use the same notation as before, despite this modification.

We use stronger conditions on the residuals:

$$(C5a). \quad \mathbb{E}[\epsilon_{i+1}^+ | \psi_{i+1}(X_{i+1})] = 0 \text{ and } \mathbb{E}[(\epsilon_{i+1}^+)^2 | \psi_{i+1}(X_{i+1})] = \text{Var}[\epsilon_{i+1}^+].$$

$$(C5b). \quad \mathbb{E}[\epsilon_{i+1} | \psi_i(X_i)] = 0 \text{ and } \mathbb{E}[(\epsilon_{i+1})^2 | \psi_i(X_i)] = \text{Var}[\epsilon_{i+1}].$$

As measures of regression precision, define the coefficients of determination

$$R_\beta^2 = \text{Var}[\beta_i^\top \psi_i(X_i)] / \text{Var}[V_{i+1}(X_{i+1})]$$

$$R_\gamma^2 = \text{Var}[\gamma_i^\top \psi_{i+1}(X_{i+1})] / \text{Var}[V_{i+1}(X_{i+1})].$$

Write $\text{Cov}[\hat{\beta}]$ for the covariance matrix of $\hat{\beta}$ and let $\Sigma_\beta = \lim_{b \rightarrow \infty} b \text{Cov}[\hat{\beta}]$ whenever the limit exists. Let Σ_γ similarly denote the limiting covariance matrix of $\hat{\gamma}$. The existence of these limits is implied by the following uniform integrability conditions on the reciprocal sums of squares of the basis functions:

(C6). As $b \rightarrow \infty$,

$$b \mathbb{E} \left(\sum_{j=1}^b \psi_i(X_{i_j}) \psi_i(X_{i_j})^\top \right)^{-1} \rightarrow (\mathbb{E}[\psi_i(X_i) \psi_i(X_i)^\top])^{-1}$$

and

$$b\mathbf{E}\left(\sum_{j=1}^b\psi_{i+1}(X_{i+1,j})\psi_{i+1}(X_{i+1,j})^\top\right)^{-1}\rightarrow\left(\mathbf{E}[\psi_{i+1}(X_{i+1})\psi_{i+1}(X_{i+1})^\top]\right)^{-1}.$$

Theorem 1. *If (C1) and (C3)–(C4) hold, then $R_\beta^2 \leq R_\gamma^2$. If also (C5)–(C6) hold then $\Sigma_\gamma \leq \Sigma_\beta$.*

This says that, in a single-period problem, “regression later” yields a better fit (as measured by the coefficient of determination) and less variable estimates of coefficients than “regression now.” The matrix inequality is in the sense that $A \leq B$ if $B - A$ is positive semidefinite. It should be noted that once we impose (C3), the diagonalization condition in (C1) may be difficult to satisfy. We do not know if the comparison in this theorem continues to hold without it.

Proof. If (C3)–(C4) hold, we know from Proposition 3 that $\beta_i = \gamma_i$ so it suffices to show $\mathbf{Var}[\epsilon_{i+1}^+] \leq \mathbf{Var}[\epsilon_{i+1}]$. From (24) we see that ϵ_{i+1} is the sum of ϵ_{i+1}^+ and a term uncorrelated (in view of (C4)) with ϵ_{i+1}^+ . It follows that the variance of ϵ_{i+1} is at least as large as that of ϵ_{i+1}^+ .

From standard properties of least-squares regression (or by direct calculation) using (C5b), we know that

$$\mathbf{E}[\hat{\beta}|\psi_i(X_{i1}), \dots, \psi_i(X_{ib})] = \beta$$

and

$$\mathbf{Cov}[\hat{\beta}|\psi_i(X_{i1}), \dots, \psi_i(X_{ib})] = \left(\sum_{j=1}^b\psi_i(X_{ij})\psi_i(X_{ij})^\top\right)^{-1}\mathbf{Var}[\epsilon_{i+1}].$$

Because the conditional expectation is constant, the unconditional covariance matrix is obtained by taking the expectation of the conditional covariance matrix, which gives (using (C6))

$$\begin{aligned} b\mathbf{Cov}[\hat{\beta}] &= b\mathbf{E}\left(\sum_{j=1}^b\psi_i(X_{ij})\psi_i(X_{ij})^\top\right)^{-1}\mathbf{Var}[\epsilon_{i+1}] \\ &\rightarrow\left(\mathbf{E}[\psi_i(X_i)\psi_i(X_i)^\top]\right)^{-1}\mathbf{Var}[\epsilon_{i+1}] \equiv \Sigma_\beta. \end{aligned}$$

Similarly, using (C5a),

$$\mathbf{E}[\hat{\gamma}|\psi_{i+1}(X_{i+1,1}), \dots, \psi_{i+1}(X_{i+1,b})] = \beta$$

and

$$\begin{aligned} \text{Cov}[\hat{\gamma} | \psi_{i+1}(X_{i+1,1}), \dots, \psi_{i+1}(X_{i+1,b})] &= \\ & \left(\sum_{j=1}^b \psi_{i+1}(X_{i+1,j}) \psi_{i+1}(X_{i+1,j})^\top \right)^{-1} \text{Var}[\epsilon_{i+1}^+] \\ & \rightarrow (\mathbb{E}[\psi_{i+1}(X_{i+1}) \psi_{i+1}(X_{i+1})^\top])^{-1} \text{Var}[\epsilon_{i+1}^+] \equiv \Sigma_\gamma. \end{aligned}$$

Using (C1), we get

$$\Sigma_\beta = \begin{pmatrix} 1 & & & & \\ & \sigma_{i1}^{-2} & & & \\ & & \sigma_{i2}^{-2} & & \\ & & & \ddots & \\ & & & & \sigma_{iK}^{-2} \end{pmatrix} \text{Var}[\epsilon_{i+1}]$$

and

$$\Sigma_\gamma = \begin{pmatrix} 1 & & & & \\ & \sigma_{i+1,1}^{-2} & & & \\ & & \sigma_{i+1,2}^{-2} & & \\ & & & \ddots & \\ & & & & \sigma_{i+1,K}^{-2} \end{pmatrix} \text{Var}[\epsilon_{i+1}^+].$$

The proof of the first part of the theorem shows that $\text{Var}[\epsilon_{i+1}^+] \leq \text{Var}[\epsilon_{i+1}]$. The martingale property (C3) and Jensen’s inequality together ensure that $\sigma_{ik}^2 \leq \sigma_{i+1,k}^2$, for all $k = 1, \dots, K$, so this establishes the second part of the theorem. \square

It should be stressed that, as formulated, this result holds only over a single period because the comparison in the theorem assumes $\hat{V}_{i+1} = \hat{V}_{i+1}^+$ and this property would not be preserved by backward induction using the two sets of coefficients. Also, the practical scope of conditions (C4) and (C5) is unclear. We expect, however, that the comparison in the theorem will often hold even if the conditions are not met precisely. For example, the first inequality in the theorem would continue to hold if the two terms on the right side of (24) were positively correlated rather than uncorrelated.

5 Duality

Recent results of Haugh and Kogan [10] and Rogers [12] show that dual formulations of the dynamic programming equations (4)–(5) can be combined with simulation to produce upper bounds on American option prices. Upper bounds can be combined with lower bounds to produce interval estimates for prices. We now show that with “regression later” a dual estimate can be computed with minimal additional effort. A different approach to computing dual values by simulation is developed in Andersen and Broadie [1].

Fix the original b paths (X_{1j}, \dots, X_{mj}) , $j = 1, \dots, b$, used to estimate regression coefficients $\hat{\gamma}_i$, $i = 1, \dots, m-1$, and simulate a new path X_1, \dots, X_m independent of the other paths. Think of the coefficients $\hat{\gamma}_1, \dots, \hat{\gamma}_{m-1}$ as fixed, meaning that we now proceed conditional on the original set of paths. Set $\hat{\gamma}_m \equiv 0$. Conditional on the coefficients, we view

$$\hat{C}_i^+(\cdot) = \sum_{k=0}^K \hat{\gamma}_{ik} \psi_{ik}(\cdot), \quad i = 1, 2, \dots, m-1,$$

and

$$\tilde{V}_{i+1}^+(\cdot) \triangleq \sum_{k=0}^K \hat{\gamma}_{ik} \psi_{i+1,k}(\cdot), \quad i = 0, 1, \dots, m-1,$$

as deterministic functions on \mathfrak{R}^d .

Define

$$\hat{\tau} = \min\{i = 0, 1, \dots, m : h_i(X_i) \geq \hat{C}_i^+(X_i)\}; \quad (25)$$

this is the first time i at which the payoff from exercise $h_i(X_i)$ exceeds the continuation value estimated by regression. Further define $M_0 = 0$ and

$$M_n = \sum_{i=0}^{n-1} [\tilde{V}_{i+1}^+(X_{i+1}) - \hat{C}_i^+(X_i)], \quad n = 1, \dots, m. \quad (26)$$

Each summand is simply

$$\tilde{V}_{i+1}^+(X_{i+1}) - \hat{C}_i^+(X_i) = \sum_{k=0}^K \hat{\gamma}_{ik} [\psi_{i+1,k}(X_{i+1}) - \psi_{ik}(X_i)]. \quad (27)$$

Theorem 2. *If (C3) holds then*

$$\mathbb{E}[h_{\hat{\tau}}(X_{\hat{\tau}})] \leq V_0^*(X_0) \leq \mathbb{E}[\max_{n=0,1,\dots,m} (h_n(X_n) - M_n)].$$

Thus, the true value $V_0^*(X_0)$ is bounded above and below by terms that can be estimated through simulation. The lower bound can be estimated by simulating independent paths each stopping according to the rule defining $\hat{\tau}$. The upper bound can be estimated from these same independent paths by computing the differences (27) at each step, summing them to get M_n , and then taking the maximum of $h_n(X_n) - M_n$ along the path.

Proof. The lower bound follows from the fact that $V_0^*(X_0)$ is defined as a supremum in (1). The stopping rule $\hat{\tau}$ in (25) is not a stopping time with respect to the history of X_1, \dots, X_i because it depends on the estimated coefficients $\hat{\gamma}_i$. But it is a randomized stopping time because the event $\{\tau = i\}$ is contained in the sigma-algebra generated by X_1, \dots, X_i and $\hat{\gamma} = \{\hat{\gamma}_1, \dots, \hat{\gamma}_{m-1}\}$.

For the upper bound, the key observation is that $0 = M_0, M_1, \dots, M_m$ is a martingale, conditional on $\hat{\gamma}$; i.e.,

$$\mathbb{E}[M_{i+1}|X_1, \dots, X_i, \hat{\gamma}] = M_i.$$

This is evident from the fact that (27) has conditional expectation 0. The bound then follows along the lines in Haugh and Kogan [10] and Rogers [12]. In more detail, the conditional martingale property implies that for any $\tau \in \mathcal{T}_0$,

$$\begin{aligned} \mathbb{E}[h_\tau(X_\tau)|\hat{\gamma}] &= \mathbb{E}[h_\tau(X_\tau) - M_\tau|\hat{\gamma}] \\ &\leq \mathbb{E}[\max_{n=0,1,\dots,m} (h_n(X_n) - M_n)|\hat{\gamma}], \end{aligned}$$

and then

$$\mathbb{E}[h_\tau(X_\tau)] \leq \mathbb{E}[\max_{n=0,1,\dots,m} (h_n(X_n) - M_n)].$$

Because this inequality holds for all $\tau \in \mathcal{T}_0$, it also holds for $V_0^*(X_0)$. \square

The key point is that using “regression later” under (C3), the martingale terms (27) are available at almost no cost. In contrast, as explained in Glasserman [8] the superficially similar expression

$$\sum_{i=0}^{n-1} [\hat{V}_{i+1}(X_{i+1}) - \hat{C}_i(X_i)]$$

resulting from “regression now” is not in general a martingale (even conditional on the coefficients) and therefore does not result in a valid upper bound. With \hat{V}_{i+1} , one needs to use an expression of the form

$$\sum_{i=0}^{n-1} [\hat{V}_{i+1}(X_{i+1}) - \mathbb{E}[\hat{V}_{i+1}(X_{i+1})|X_i]].$$

Computing the i th conditional expectation in this sum typically requires simulating a large number of subpaths, each starting at X_i and advancing one time step. This use of subpaths to estimate conditional expectations involves a heavier computational burden than (27), which merely involves evaluating linear combinations of basis functions at each step. The savings results from taking advantage of the known conditional expectations provided by (C3).

While any martingale would provide an upper bound in Theorem 2, the one in (26) is close to optimal. The martingale

$$M_n^* = \sum_{i=0}^{n-1} [V_{i+1}^*(X_{i+1}) - C_i^*(X_i)] \tag{28}$$

constructed from the true value and continuation functions turns the upper bound in the theorem into an equality; see the derivation in Section 8.7 of Glasserman [8]. Thus, the martingale in (26) is in a sense a best approximation to the optimal martingale (28), given the choice of basis functions. The drawback to (27) is that it places more restrictive conditions on the available basis functions through (C3). The quality of the upper bound provided

by (26) compared with other bounds based on duality remains a topic for investigation.

Acknowledgements. We thank Mark Broadie for discussions related to this work. This research is supported by NSF Grant DMS0074637.

References

1. Andersen, L., and Broadie, M.: A primal-dual simulation algorithm for pricing multi-dimensional American options. Working paper, Columbia Business School, New York (2001)
2. Avellaneda, M.: Minimum-relative-entropy calibration of asset-pricing models. *International Journal of Theoretical and Applied Finance* **1**, 447–472 (1998)
3. Avellaneda, M., Buff, R., Friedman, C., Grandchamp, N., Kruk, L., and Newman, J.: Weighted Monte Carlo: a new technique for calibrating asset-pricing models. *International Journal of Theoretical and Applied Finance* **4**, 1–29 (2001)
4. Broadie, M., and Glasserman, P.: A stochastic mesh method for pricing high-dimensional American options. PaineWebber Papers in Money, Economics and Finance #PW9804, Columbia Business School, New York (1997)
5. Broadie, M., Glasserman, P., and Ha, Z.: Pricing American options by simulation using a stochastic mesh with optimized weights. In: Uryasev, S. (ed) *Probabilistic Constrained Optimization: Methodology and Applications*, Kluwer Publishers, Norwell, Mass. (2000)
6. Carrière, J.: Valuation of early-exercise price of options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* **19**, 19–30 (1996)
7. Clément, E., Lamberton, D., and Protter, P.: An analysis of a least squares regression algorithm for American option pricing. *Finance and Stochastics* **6**, 449–471 (2002)
8. Glasserman, P.: *Monte Carlo Methods in Financial Engineering*, Springer-Verlag, New York (2004)
9. Glasserman, P., and Yu, B.: Large sample properties of weighted Monte Carlo. Working Paper DRO-2002-07, Columbia Business School, New York (2002)
10. Haugh, M., and Kogan, L.: Pricing American options: a duality approach. *Operations Research*, to appear.
11. Longstaff, F.A., and Schwartz, E.S.: Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies* **14**, 113–147 (2001)
12. Rogers, L.C.G.: Monte Carlo valuation of American options. *Mathematical Finance* **12**, 271–286 (2002)
13. Tsitsiklis, J., and Van Roy, B.: Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* **12**, 694–703 (2001)