TRANSACTIONS OF SOCIETY OF ACTUARIES
1993 VOL. 45

VALUING AMERICAN OPTIONS IN A PATH SIMULATION MODEL

JAMES A. TILLEY

ABSTRACT

The goal of this paper is to dispel the prevailing belief that American-style options cannot be valued efficiently in a simulation model, and thus remove what has been considered a major impediment to the use of simulation models for valuing financial instruments. We present a general algorithm for estimating the value of American options on an underlying instrument or index for which the arbitrage-free probability distribution of paths through time can be simulated. The general algorithm is tested by an example for which the exact option premium can be determined.

1. INTRODUCTION

Mathematicians seem to resort to simulation models to analyze a problem only when all other methods fail to yield a solution. In financial economics, evidence of this tendency to avoid simulation models is found in the proliferation of published binomial and multinomial lattice solutions (or their equivalent) to the problem of valuing instruments with cash flows or payoffs contingent on interest rates or stock prices [1], [2], [4], [5], [8], [9], [10], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], and [23]. The standard approach to valuing an American option is to utilize a one-factor (continuous) model of the stochastic price behavior of the option's underlying asset, then create a binomial or multinomial (discrete) connected lattice representation of that stochastic process, and finally solve the valuation problem by backward induction on the lattice. Market-makers who deal in today's complicated financial instruments and investors who buy and sell them are beginning to sense a need for more realistic multifactor models of the stochastic dynamics of interest rates, foreign exchange rates, stock prices, and commodity prices. These more complex models demand analysis by simulation, because constructing approximate solutions (whether by means of lattices or otherwise) to the nonlinear differential and integral equations to which they give rise is extremely difficult.

In general, the use of simulation models for valuing financial instruments has been restricted to assets that have path-dependent cash flows or payoffs, for example, mortgage-backed securities, including collateralized mortgage obligations (CMOs), and esoteric derivative instruments, such as "look-back" options [7], [11]. (An exception, at least in the academic literature if not in practice, is the paper by Boyle [3], which examines how Monte Carlo simulation can be used to value European-style options.) Indeed, it has been thought that simulation models could not be used to value American-style options efficiently, if at all ([7], [8], and Pedersen's discussion of [24]). Ideally, a broker-dealer would like to be able to use a *single* method to value its entire book, and a financial intermediary would like to be able to use a *single* method to analyze its entire asset-liability condition. I believe that simulation models offer that possibility.

Simulation models consume large amounts of computer processing time, and some problems have heretofore required too much execution time to be handled practically by simulation. However, the arrival of powerful workstations, servers, and parallel-processers has rendered simulation feasible in many situations where it previously was not, a condition that can only improve with time as the pace of major technological advances accelerates. In many situations, a *single sample* of paths can be generated and then *used repeatedly* to value many different instruments, for example, a dealer's entire book of interest rate swaps, caps, floors, and swaptions; a dealer's entire book of stock index derivatives or of currency swaps and options; or a financial intermediary's entire portfolio of fixed-income securities. Simulation may not be the best method when each financial instrument must be valued on the basis of its own random sample of paths, but this situation can often be avoided by designing the simulation properly.

The algorithm for valuing American options is described in Section 2 and tested by an example in Section 3. The issue of bias in the estimator of the option premium is examined in Section 4, after which the example is revisited in Section 5. Finally, Section 6 summarizes the paper.

2. THE VALUATION ALGORITHM

A textbook by Cox and Rubinstein [6] provides a comprehensive treatment of the subject of options. We assume that the reader is familiar with the general subject area, including various models for pricing options. For convenience, the option's underlying asset is referred to as a "stock," but the entire development in this section applies to any type of asset or index for which the arbitrage-free probability distribution of paths through time can be simulated. In my earlier paper [24] we discuss what is meant by "arbitrage-free" and show how arbitrage-free paths of interest rates can be sampled stochastically. The example in Section 3 of this paper utilizes paths of stock prices that are sampled from a probability distribution that is arbitrage-free because its mean has been adjusted appropriately.

We consider how to evaluate put and call options on a stock. The options are exercisable only at specified epochs t_1, t_2, ..., t_N, which are indexed 1, 2, ..., N for convenience. The origin of time is $t = 0$, which is indexed as epoch 0. The options can be considered to be first exercisable at epoch 0 or at epoch 1, as appropriate. A path of stock prices is a sequence $S(0)$, $S(1)$, $S(2)$, ..., $S(N)$, in which the arguments of S refer to the epoch indexes at which the stock prices occur. *All* paths of stock prices emanate from the initial stock price $S(0)$. The simulation procedure involves the random generation of a finite sample of R such paths and the estimation of option prices from that sample. The k-th path in the sample is represented by the sequence $S(0)$, $S(k,1)$, $S(k,2)$, ..., $S(k,N)$, in which the first index refers to the path and the second index refers to the epoch. Two paths of stock prices are represented in Figure 1. Let $d(k,t)$ be the present value at epoch t on path k of a \$1 payment occurring at epoch $t+1$ on path k. Let $D(k,t)$ be the present value at epoch 0 of a \$1 payment occurring at epoch t on path k, computed as the product of the discount factors $d(k,s)$ from $s=0$ to $s=t-1$.

Assume that the option has strike prices that can depend on the date of exercise but not on the stock price at the time of exercise. Let $X(1)$, $X(2)$, ..., $X(N)$ denote the sequence of exercise prices at epochs 1, 2, ..., N, respectively. Typical stock options have a constant strike price X independent of date of exercise, but typical call options in private placement bonds do not. The intrinsic value $I(k,t)$ of the option on path k at epoch t can now be defined as:

$$I(k,t) = \begin{cases} \text{maximum } [0, S(k,t) - X(t)] & \text{for a call option} \\ \text{maximum } [0, X(t) - S(k,t)] & \text{for a put option.} \end{cases}$$

FIGURE 1

TWO ILLUSTRATIVE PATHS OF STOCK PRICES SAMPLED FROM A DISCRETE-TIME CONTINUOUS-STATE MODEL OF STOCK PRICE MOVEMENTS

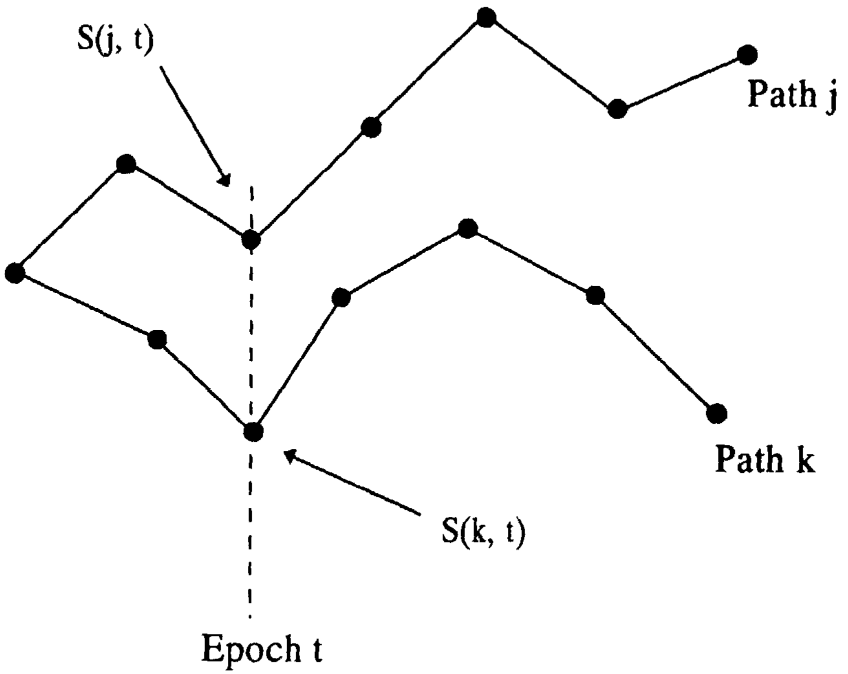

Finally, let $z(k,t)$ be the "exercise-or-hold" indicator variable, which takes the value 0 if the option is not exercised at epoch t on path k and which takes the value 1 if the option is exercised at epoch t on path k. Clearly, either $z(k,t)=0$ at all epochs t along path k or $z(k,t_*)=1$ at one and only one epoch t_* along path k. If such a t_* exists, it is the date at which the option is exercised on path k.

The price of any asset is known at epoch 0 if its cash flows are known at all epochs along all possible paths. That price is calculated in two steps: first, compute for each path k the present value at epoch 0 of the asset's cash flows along that path using the path-specific discount factors $D(k,t)$, and second, average across all paths the present values computed in the first step. The paths must be drawn from the appropriate arbitrage-free distribution. More details on this general valuation procedure can be found in my paper [24]. On a given stock price path, the "cash flow" for an option is 0 at every epoch other than the one at which the option is exercised. At exercise, the option's "cash flow" is equal to its payoff, which is its intrinsic value. Assuming the usual situation that all randomly sampled stock price paths are equally likely with probability weight R^{-1}, we can express the option premium estimator by the following equation:

$$\text{Premium Estimator} = R^{-1} \sum_{\substack{\text{all} \\ \text{paths} \\ k}} \sum_{\substack{\text{all} \\ \text{epochs} \\ t}} z(k,t)D(k,t)I(k,t).$$

Thus, to estimate the price of the option, we need to estimate the exercise-or-hold indicator function $z(k,t)$, given a finite sample of R paths drawn from an arbitrage-free distribution of paths. The algorithm presented in this section for estimating $z(k,t)$ mimics the standard backward induction algorithm implemented on a connected lattice for estimating the value of an American option. A discussion of this standard technique can be found in the textbook by Cox and Rubinstein [6].

The backward induction is begun at the latest epoch at which the option can be exercised, that is, at its expiration date. On that date, represented by epoch N, the option, if it is still "alive" on path k (that is, not previously exercised), will be exercised if and only if $I(k,N)>0$. The general step of the backward induction performed at an arbitrary epoch t involves determining whether it is optimal to hold the option for possible exercise beyond epoch t or to exercise the option immediately at epoch t. This decision is made by comparing the option's "holding value" to its "exercise value." The option's exercise value is equal to its intrinsic value and can be directly calculated for each path, because the price of the underlying stock is known at each epoch on each path. The option's holding value on any path is calculated as the present value of the expected one-period-ahead option value.

Many believe that utilizing the path structure illustrated in Figure 1 precludes estimation of an option's holding values, because the only point from which *many* paths emerge is epoch 0. On any particular path, at any epoch $t>0$, only a *single* path is simulated. One might think that many paths would need to be simulated from each such point to estimate closely the mathematical expectation of the one-period-ahead option value. Unfortunately, such an approach would lead to a multinomial "tree" in which the number of paths grows exponentially with the number of epochs—a computational infeasibility. Instead, computational feasibility can be achieved by utilizing the path structure illustrated in Figure 1 and then estimating the option's holding value by means of a distinct partitioning at each epoch of the R paths into Q bundles of P paths each. The hope is that the paths within a given bundle are sufficiently alike that they can be considered to have the same expected one-period-ahead option value; in other words, Q must not be too small. The mathematical expectation of the one-period-ahead option value is estimated as an average over all the paths in the bundle. Thus, the estimate of the option holding value will be good only if there are sufficiently many paths in the bundle; in other words, P must not be too small.

In general, there is at least one bundle in which the decision for some paths is to hold the option, while the decision for the rest of the paths in the bundle is to exercise the option. Such a bundle generally has a "transition zone" in stock price from a decision to hold the option to a decision to exercise the option. Specifically, for a call option, there exist stock prices $S_L(t)$ and $S_U(t)$ at epoch t, with $S_L(t)<S_U(t)$, such that the decision is to hold for $S \leq S_L(t)$ and to exercise for $S \geq S_U(t)$. However, for $S_L(t)<S<S_U(t)$, the decision is "inconsistent"; that is, there exist stock prices $S_l(t)$ and $S_u(t)$ such that $S_L(t)<S_l(t)<S_u(t)<S_U(t)$, yet the decision is to exercise at $S=S_l(t)$ and to hold at $S=S_u(t)$! The transition zone from an unambiguous hold decision to an unambiguous exercise decision often extends across several consecutive bundles. The algorithm can be refined to eliminate the transition zone by creating a "sharp boundary" at $S=S_*(t)$, such that the decision is to hold for $S<S_*(t)$ and to exercise for $S \geq S_*(t)$.

The general step that is performed at epoch t in the backward induction algorithm includes eight substeps, as follows:

1. Reorder the stock price paths by stock price, from lowest price to highest price for a call option or from highest price to lowest price for a put option. Reindex the paths from 1 to R according to the reordering.
2. For each path k, compute the intrinsic value $I(k,t)$ of the option.
3. Partition the set of R ordered paths into Q distinct bundles of P paths each. Assign the first P paths to the first bundle, the second P paths

to the second bundle, and so on, and finally the last P paths to the Q-th bundle. We assume that P and Q are integer factors of R.

4. For each path k, the option's holding value $H(k,t)$ is computed as the following mathematical expectation taken over all paths in the bundle containing the path k:

$$H(k,t) = d(k,t)P^{-1} \sum_{\substack{\text{all } j \\ \text{in bundle} \\ \text{containing } k}} V(j, t+1).$$

The variable $V(k,t)$ is fully defined in substep 8 below. At epoch N, $V(k,N)=I(k,N)$ for all k.

5. For each path k, compare the holding value $H(k,t)$ to the intrinsic value $I(k,t)$, and decide "tentatively" whether to exercise the option or to hold it. Define an indicator variable $x(k,t)$ as follows:

$$x(k,t) = \begin{cases} 1 & \text{if} \quad I(k,t) > H(k,t) \quad \text{Exercise} \\ 0 & \text{if} \quad H(k,t) \geq I(k,t) \quad \text{Hold.} \end{cases}$$

6. Examine the sequence of 0's and 1's $\{x(k,t);\ k=1, 2, \ldots, R\}$. Determine a sharp boundary between the hold decision and the exercise decision as the start of the first string of 1's, the length of which exceeds the length of every subsequent string of 0's. Let $k_*(t)$ denote the path index (in the sample as *ordered* in substep 1 above) of the leading 1 in such a string. The transition zone between *hold* and *exercise* is defined as the sequence of 0's and 1's that begins with the first 1 and ends with the last 0. An example is given below:

$$\begin{array}{c} \text{Boundary} \\ \downarrow \\ 0\,0 \ldots 0\,1\,1\,0\,0\,0\,1\,1\,1\,1\,1\,0\,0\,1 \ldots 1\,1 \end{array}$$

7. Define a new exercise-or-hold indicator variable $y(k,t)$ that incorporates the sharp boundary as follows:

$$y(k,t) = \begin{cases} 1 & \text{for} \quad k \geq k_*(t) \\ 0 & \text{for} \quad k < k_*(t). \end{cases}$$

8. For each path k, define the current value $V(k,t)$ of the option as follows:

$$V(k,t) = \begin{cases} I(k,t) & \text{if} \quad y(k,t) = 1 \\ H(k,t) & \text{if} \quad y(k,t) = 0. \end{cases}$$

After the algorithm has been processed backward from epoch N to epoch 1 (or epoch 0 if immediate exercise is permitted), the indicator variable $z(k,t)$ for $t \leq N$ is estimated as follows:

$$z(k,t) = \begin{cases} 1 & \text{if } y(k,t) = 1 \text{ and } y(k,s) = 0 \text{ for all } s < t \\ 0 & \text{otherwise.} \end{cases}$$

This completes the description of the algorithm for valuing an American option.

The partition of R paths into Q bundles of P paths each can be characterized by defining a "bundling parameter" α by means of the equation $Q=R^{\alpha}$, and therefore, $P=R^{1-\alpha}$. It is clear that $0 \leq \alpha \leq 1$. The value $\alpha=0$ corresponds to the partition into a single bundle of R paths, and the value $\alpha=1$ corresponds to the partition into R bundles of one path each. A particular American option valuation algorithm can now be fully described by the sample size R, the technique used to sample paths, and the bundling parameter α. If α is restricted to rational numbers, we can fix α and take sensible limits as $R \rightarrow \infty$ to investigate the convergence properties of the option premium estimator. For example, with $\alpha=2/5$, we can examine sample sizes equal to 2^5, 3^5, 4^5, ... paths for which we can study the estimators associated with the partitions $Q=2^2$, 3^2, 4^2, ... bundles and $P=2^3$, 3^3, 4^3, ... paths per bundle, respectively.

For any exercise-hold decision algorithm with α fixed and $0<\alpha<1$, it can be proved that the option premium estimate must converge to the proper result as $R \rightarrow \infty$. This follows from the observation that the algorithm for determining the exercise-hold decision variable is based on the standard backward induction algorithm for valuing American options and that all sources of error arise from P, Q, and R being finite. For finite R, imprecision in the premium estimates arises because: (1) the continuous distribution of stock prices at each epoch is not sampled finely enough and (2) the mathematical expectation in substep 4 above is approximated by an average over a finite number of paths. Imprecision of the first type can be reduced by increasing Q, the number of bundles. Imprecision of the second type can be reduced by increasing P, the number of paths per bundle. For fixed R, increasing Q means decreasing P, and vice versa, implying a tradeoff between the first and second types of imprecision. However, if α is held constant at some value in the interval $(0,1)$, then both types of imprecision are eliminated simultaneously as $R \to \infty$, because then *both* $Q \to \infty$ *and* $P \to \infty$.

The distinction between the variables $y(k,t)$ and $x(k,t)$ disappears as $R \to \infty$ and α is held constant at a value other than 0 or 1. As $R \to \infty$, the boundary between a decision to exercise the option and a decision to hold the option becomes sharper and sharper; that is, at each epoch, the transition zone with alternating strings of 1's and 0's occurs over a smaller and smaller interval of stock prices. Defining a sharp boundary by means of substep 6 above generally improves the convergence of the algorithm considerably for any α in the interval $(0,1)$ and also generally broadens considerably the interval of α over which the option premium estimates are good. In general, the option premium estimate based on a given sample size, sampling technique, and bundling parameter is more accurate when a sharp exercise-hold boundary is determined than when it is not. However, the *ultimate* convergence of the exercise-hold decision algorithm to the exact option premium does *not* depend at all on whether substeps 6 and 7 above are implemented. If substeps 6 and 7 were omitted from the algorithm, $x(k,t)$ would be used in lieu of $y(k,t)$, both in substep 8 and in the calculation of $z(k,t)$.

3. AN EXAMPLE

To test the algorithm presented in the preceding section, we consider the situation of a non-dividend-paying stock. Let $S(t)$ denote the price of the stock at time t. We assume that the random variable $\ln[S(t)/S(0)]$ is normally distributed with mean μt and variance $\sigma^2 t$. We further assume that the yield curve is flat and that interest rates are constant over time at an annual effective rate r. For the distribution of stock price movements to be arbitrage-free over time, it must be true that $\mu = \ln[1+r] - \sigma^2/2$. Refer to the textbook by Cox and Rubinstein [6] for a proof of this statement.

When a non-dividend-paying stock is the underlying asset, the price of an American call option must be exactly the same as the price of an otherwise identical European call option [6]. The price of an American put option must be no less than the price of an otherwise identical European put option, but the former will in general exceed the latter [6]. Therefore, we test the valuation algorithm on a *put option* that is first exercisable in one quarter of a year and is exercisable every quarter of a year thereafter until its expiration in three years. The stock price has logarithmic volatility σ equal to 30 percent. The initial price of the stock $S(0)$ is 40; the strike price X of the option is 45 at all epochs; and the annual effective interest rate r is 7 percent. Paths of stock price movements are generated randomly by stratified sampling of the standard normal density as described in my paper [24]. Random samples of size $7!=5{,}040$ are used so that many different partitions can be examined. Table 1 lists the values of the bundling parameter α that correspond to each of the 60 different partitions of 5,040 paths into equal bundles.

The exact price of the three-year American put option with quarterly exercise intervals was determined to be 7.941 by using a binomial lattice with 1,200 time periods constructed according to the procedure described in Cox and Rubinstein [6]. This is approximately 1.61 higher than the price of the corresponding three-year European put option. Using a *single* sample of 5,040 paths, the exercise-decision algorithm described in the preceding section was tested for all partitions having at least 12 bundles but no more than 420 bundles. The results are displayed in Figure 2.

In Figure 2, the solid line connecting "diamonds" corresponds to application of the algorithm without substeps 6 and 7—that is, with a *transition zone* from hold to exercise, not a sharp boundary between hold and exercise. The broken line connecting "squares" corresponds to application of the algorithm with substeps 6 and 7 included—that is, with a *sharp boundary* between hold and exercise. The horizontal line across the graph at a vertical axis value of 7.941 marks the exact option premium. Figure 2 clearly demonstrates the importance of including substeps 6 and 7 in the algorithm. When a sharp boundary is determined, the option premium estimates are essentially flat across an interval from $\alpha=0.29$ to $\alpha=0.71$ and cover a range of only 12 cents. However, when only a transition zone is utilized, the option premium estimates rise more or less steadily as the bundling parameter is increased and cover a range

TABLE 1

BUNDLING PARAMETER ALPHA FOR VARIOUS PARTITIONS OF 5,040 PATHS

Partition		Bundling	Partition		Bundling
Number of Bundles	Paths per Bundle	Parameter Alpha	Number of Bundles	Paths per Bundle	Parameter Alpha
1	5040	0.00000	72	70	0.50165
2	2520	0.08131	80	63	0.51401
3	1680	0.12887	84	60	0.51973
4	1260	0.16261	90	56	0.52783
5	1008	0.18879	105	48	0.54591
6	840	0.21017	112	45	0.55348
7	720	0.22825	120	42	0.56157
8	630	0.24392	126	40	0.56730
9	560	0.25773	140	36	0.57965
10	504	0.27009	144	35	0.58296
12	420	0.29148	168	30	0.60104
14	360	0.30956	180	28	0.60913
15	336	0.31765	210	24	0.62721
16	315	0.32522	240	21	0.64288
18	280	0.33904	252	20	0.64860
20	252	0.35140	280	18	0.66096
21	240	0.35712	315	16	0.67478
24	210	0.37279	336	15	0.68235
28	180	0.39087	360	14	0.69044
30	168	0.39896	420	12	0.70052
35	144	0.41704	504	10	0.72991
36	140	0.42035	560	9	0.74227
40	126	0.43270	630	8	0.75608
42	120	0.43843	720	7	0.77175
45	112	0.44652	840	6	0.78983
48	105	0.45409	1008	5	0.81121
56	90	0.47217	1260	4	0.83739
60	84	0.48027	1680	3	0.87113
63	80	0.48599	2520	2	0.91869
70	72	0.49835	5040	1	1.00000

of approximately 63 cents, more than five times the range obtained when a sharp boundary is utilized!

To study the efficiency of estimation, the "70 bundles by 72 paths per bundle" partition was used on 1,000 independent samples of 5,040 paths ach. Each sample gives rise to an estimate of the put option premium. The frequency histogram of these 1,000 estimates is plotted in Figure 3. The mean of the estimates is 7.971 and the standard deviation of the estimates is 0.053. The solid line graph superimposed on the frequency histogram is that of a normal density function with the same mean and standard deviation as the option premium estimator. We can see that the algorithm produces premium estimates that are normally distributed. What

FIGURE 2

PREMIUM ESTIMATES FOR 3-YEAR AMERICAN PUT OPTION
(5,040 Paths Partitioned 40 Ways into Exercise-Decision Bundles)

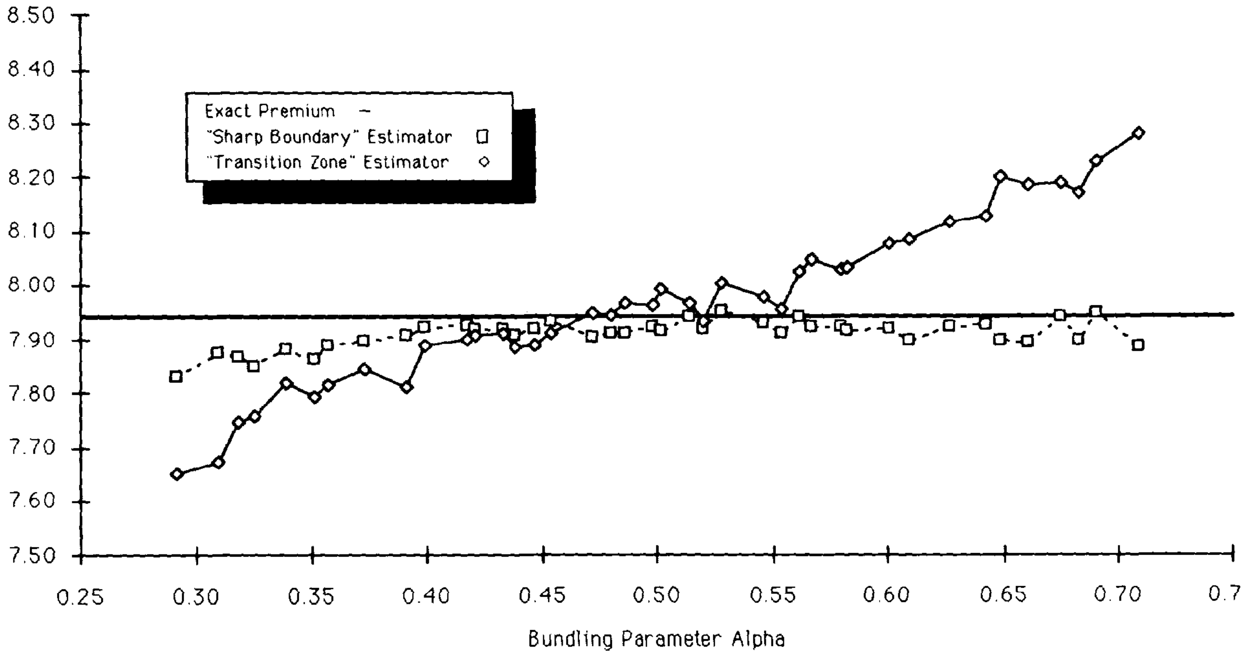

FIGURE 3

FREQUENCY HISTOGRAM FOR $\alpha=0.50$ PREMIUM ESTIMATOR
(Based on 1,000 Samples of 5,040 Paths Each)

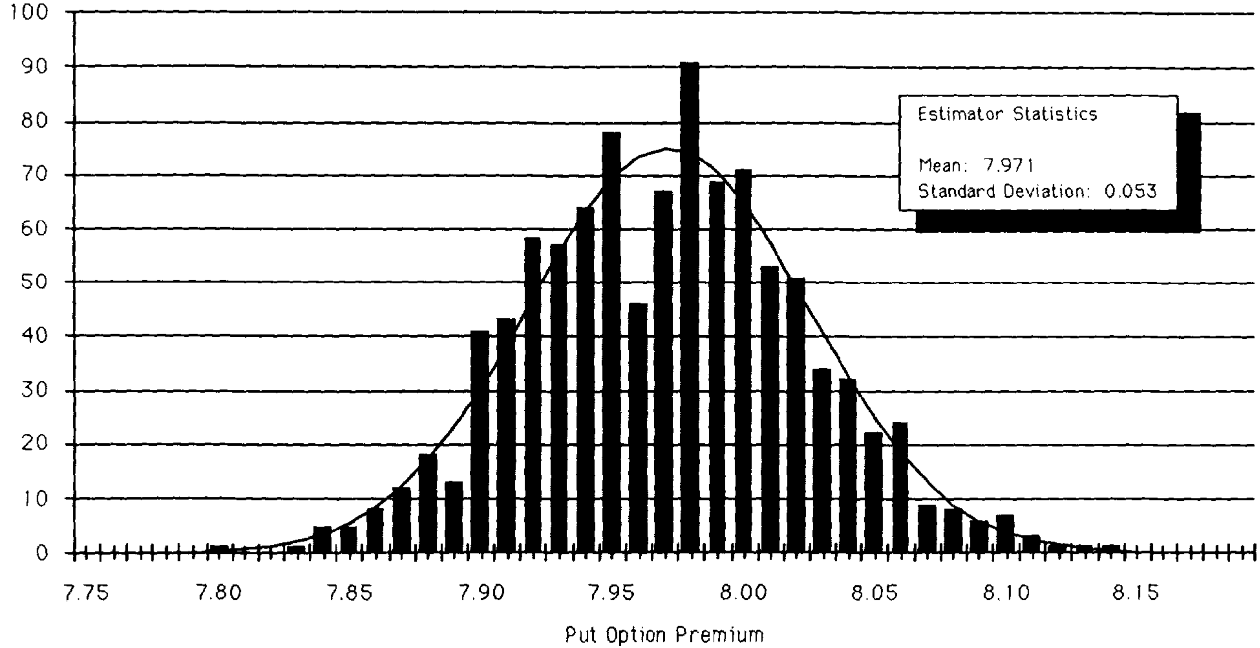

seems surprising is that the premium estimator is *biased*. The mean estimate of \$7.971 is 3 cents higher than the exact premium of \$7.941, which is about 17.9 times the standard deviation of $5.3/\sqrt{1000}$ cents. Despite the bias, the algorithm can estimate the option premium quite tightly.

4. ESTIMATOR BIAS

In this section, we investigate the source of the bias in the option premium estimates that was discovered by means of the example presented in the last section. It turns out that the bias arises because the "optimization" is done over a *finite* sample. The bias vanishes in the limit of infinite sample size. The description of the exercise-decision algorithm in Section 2 makes it evident that estimating the premium for an American option is equivalent to estimating the exercise-hold stock price boundary at each epoch at which the option can be exercised. Accordingly, we determined the "exact" boundary between holding and exercising the put option at each of the 12 exercise-decision epochs by using the Cox-Rubinstein binomial lattice that was described in the preceding section. With *full knowledge* of the exact exercise-hold boundaries, the American option premium was estimated again by simulation using the *same 1,000 samples* of 5,040 paths on which the results shown in Figure 3 were based. The resulting frequency histogram of the premium estimates is shown in Figure 4.

In Figure 4 the premium estimates are normally distributed. The standard deviation of the estimates is 5.3 cents, the same as in Figure 3. However, the mean of the estimates is \$7.943, only 0.2 cents higher than the exact premium. This deviation is not statistically significant at a 5 percent level of confidence, since it is only about 1.2 times the standard deviation of $5.3/\sqrt{1000}$ cents. Thus, with full knowledge of the exact exercise-decision boundaries, the American option premium estimator is unbiased, even for finite samples of paths. We must conclude that the process of estimating the exercise-hold boundaries from a finite sample of paths introduces the bias. The following analysis demonstrates the truth of this assertion.

The exact price of an American option is the value given by the premium estimator equation in Section 2 when the *infinite sample space* of

FIGURE 4

FREQUENCY HISTOGRAM FOR "BEST" PREMIUM ESTIMATOR
(Based on 1,000 Samples of 5,040 Paths Each)

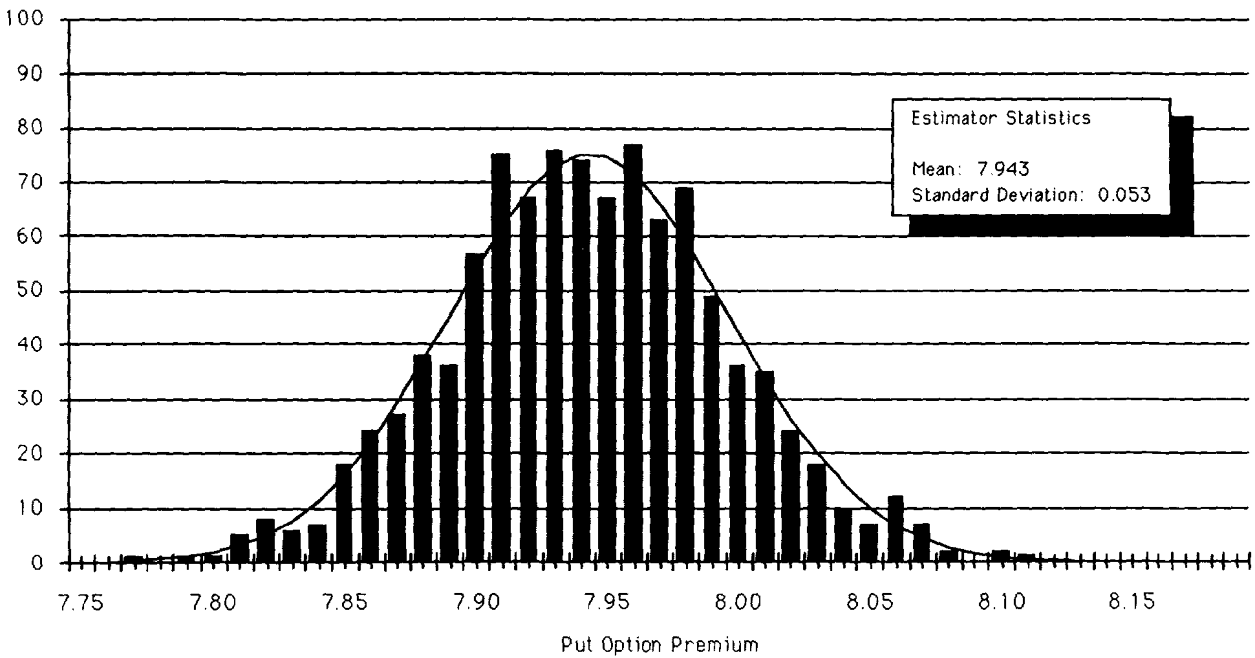

stock price paths and the *exact exercise-hold boundaries* are used. Determining the exact price of the option is equivalent to finding the exercise-hold boundaries at all exercise-decision epochs that maximize the value given by the premium estimator equation when the infinite sample space of stock price paths is used. An approximation to the exact price is obtained by finding the exercise-hold boundaries at all exercise-decision epochs that maximize the value given by the premium estimator equation when a *finite sample* of R stock price paths is used. A different approximation to the exact price is obtained by implementing the backward induction algorithm with eight substeps at each epoch that was described in Section 2. This latter estimate of the exact option price is itself an approximation to the former estimate of the exact option price, by reason of the construction of the backward induction algorithm as an optimization.

Let E_i denote the option premium estimate obtained when the i-th sample of R paths is used together with some premium estimation method. The dependence of the estimate on the estimation method used is denoted by an appropriate superscript. The superscript *∞-optimal* is used to represent the estimation method that utilizes the exact boundaries determined from the infinite sample space of stock price paths. The superscript *R-optimal* is used to represent the estimation method that utilizes the boundaries that optimize the value given by the premium estimator equation when the finite sample of R paths is used. Finally, the superscript *R-algorithm* is used to represent the estimation method that utilizes the boundaries determined from the eight-substep backward induction algorithm applied to the finite sample of R paths. As a consequence of the definitions of the various estimates and the construction of the different estimation methods, the following inequalities hold for *any* sample i consisting of R paths:

$$E_i^{\infty\text{-}optimal} \leq E_i^{R\text{-}optimal} \quad \text{and} \quad E_i^{R\text{-}algorithm} \leq E_i^{R\text{-}optimal}.$$

Thus, the means of the various estimators computed over any finite number of samples of R paths each also satisfy the same inequalities. In practice, the strict inequality will hold "almost surely." When the sample size is infinite, the inequalities become equalities. Because the *∞-optimal* estimator is unbiased, the first inequality demonstrates that the *R-optimal* estimator must always have positive bias. The bias tends to zero as $R \rightarrow \infty$. Furthermore, the second inequality demonstrates that the *R-optimal* estimator must be positively biased relative to the *R-algorithm* estimator.

The relative bias tends to zero as $R \to \infty$. It is indeterminable whether the *R-algorithm* estimator has positive or negative bias with respect to the *∞-optimal* estimator. The two inequalities also show that we should not try too hard to "perfect" the *R-algorithm* estimator in the sense of making it better approximate the *R-optimal* estimator, because the latter always has positive bias relative to the unbiased *∞-optimal* estimator.

5. EXAMPLE REVISITED

Now that we understand that the sign of the bias of the *R-algorithm* estimator is indeterminable, but is likely to be positive if the *R-algorithm* estimates the *R-optimal* exercise-hold decision boundaries closely, we should conduct further empirical studies of the bias. Table 2 presents results obtained by using the *R-algorithm* estimator of Section 2 on 100

TABLE 2

STATISTICS FOR $\alpha = 0.50$ ESTIMATORS OF PREMIUMS FOR AMERICAN PUT OPTIONS (STOCK PRICE VOLATILITY OF 30 PERCENT)

Stock Price: 40
Option Expiration: 3.00 Years — Stock Volatility: 30 Percent
Exercise Interval: 0.25 Years — Annual Interest Rate: 7 Percent

Strike Price	'Exact' Premium*	Estimator Mean†	Estimator Bias	Estimator Standard Deviation†
10	0.003	0.003	0.000	0.001
15	0.046	0.046	0.000	0.005
20	0.242	0.239	−0.003	0.012
25	0.744	0.744	0.000	0.020
30	1.689	1.694	0.005	0.027
35	3.172	3.185	0.013	0.038
40	5.247	5.268	0.021	0.044
45	7.941	7.968	0.027	0.055
50	11.255	11.289	0.034	0.063
55	15.136	15.161	0.025	0.059
60	19.469	19.485	0.016	0.054
65	24.100	24.109	0.009	0.044
70	28.894	28.899	0.005	0.034
75	33.764	33.763	−0.001	0.028
80	38.665	38.662	−0.003	0.024
85	43.576	43.574	−0.002	0.017
90	48.491	48.486	−0.005	0.015
95	53.407	53.400	−0.007	0.014
100	58.323	58.316	−0.007	0.012

*Calculated using the Cox-Rubinstein binomial model with 1,200 time intervals.
†Calculated using a simulation model with 100 samples of 5,040 paths and exercise boundary determined by first dominant string of 1's in the transition zone.

independent samples of 5,040 paths each by using a partition of 70 bundles by 72 paths per bundle. Results are shown for 3-year American put options with strike prices ranging from 10 to 100 in multiples of 5. All other assumptions are the same as in the earlier example. The "exact" premiums were calculated as before, by using the Cox-Rubinstein binomial lattice with 1,200 time intervals. The estimator bias ranges from a low of -0.7 cents to a high of $+3.4$ cents. The standard deviations of the estimates peak at 6.3 cents for a put option somewhat in the money. The premium estimates must be considered very accurate.

Table 3 presents results similar to those in Table 2, but for a partition of 504 bundles by 10 paths per bundle. In this case, substep 6 of the exercise-decision algorithm was refined to account not only for the first dominant string of 1's in the transition zone but also the last dominant string of 0's in the transition zone. As in substep 6, a boundary index is determined as the start of the first string of 1's, the length of which

TABLE 3

STATISTICS FOR $\alpha=0.73$ ESTIMATORS OF PREMIUMS FOR AMERICAN PUT OPTIONS (STOCK PRICE VOLATILITY OF 30 PERCENT)

Stock Price: 40
Option Expiration: 3.00 Years — Stock Volatility: 30 Percent
Exercise Interval: 0.25 Years — Annual Interest Rate: 7 Percent

Strike Price	'Exact' Premium*	Estimator Mean†	Estimator Bias	Estimator Standard Deviation†
10	0.003	0.003	0.000	0.001
15	0.046	0.048	0.002	0.005
20	0.242	0.246	0.004	0.012
25	0.744	0.750	0.006	0.018
30	1.689	1.697	0.008	0.028
35	3.172	3.178	0.006	0.039
40	5.247	5.255	0.008	0.049
45	7.941	7.943	0.002	0.052
50	11.255	11.260	0.005	0.066
55	15.136	15.139	0.003	0.059
60	19.469	19.468	−0.001	0.058
65	24.100	24.094	−0.006	0.049
70	28.894	28.889	−0.005	0.037
75	33.764	33.752	−0.012	0.034
80	38.665	38.654	−0.011	0.032
85	43.576	43.566	−0.010	0.021
90	48.491	48.480	−0.011	0.021
95	53.407	53.398	−0.009	0.020
100	58.323	58.315	−0.008	0.016

*Calculated using the Cox-Rubinstein binomial model with 1,200 time intervals.
†Calculated using a simulation model with 100 samples of 5,040 paths and exercise boundary determined by dominant strings of both 0's and 1's in the transition zone.

exceeds the length of every subsequent string of 0's. Another boundary index is determined as the end of the last string of 0's, the length of which exceeds the length of every previous string of 1's. In many cases, the two boundaries are identical, but if not, the *dominant 0-string* boundary must occur before the *dominant 1-string* boundary. The boundary index actually used in the revised algorithm is the arithmetic mean of the two boundary indexes, rounded appropriately. The estimator bias shown in Table 3 ranges from a low of −1.2 cents to a high of +0.8 cents. The standard deviations of the estimates are generally a little higher than their counterparts in Table 2.

Table 4 presents results similar to those in Table 3, except that the stock price volatility has been doubled to 60 percent. Again, the estimator biases are small, ranging from −0.8 cents to +2.4 cents. The standard deviations of the estimates are much larger, but are still very small when expressed as a percentage of the exact premiums.

TABLE 4

STATISTICS FOR α=0.73 ESTIMATORS OF PREMIUMS FOR AMERICAN PUT OPTIONS (STOCK PRICE VOLATILITY OF 60 PERCENT)

Stock Price: 40
Option Expiration: 3.00 Years
Exercise Interval: 0.25 Years

Stock Volatility: 60 Percent
Annual Interest Rate: 7 Percent

Strike Price	'Exact' Premium*	Estimator Mean†	Estimator Bias	Estimator Standard Deviation†
10	0.486	0.489	0.003	0.013
15	1.409	1.414	0.005	0.022
20	2.810	2.815	0.005	0.034
25	4.636	4.643	0.007	0.044
30	6.834	6.840	0.006	0.054
35	9.357	9.367	0.010	0.064
40	12.162	12.180	0.018	0.075
45	15.220	15.240	0.020	0.088
50	18.504	18.526	0.022	0.102
55	21.986	22.009	0.023	0.120
60	25.650	25.667	0.017	0.118
65	29.475	29.493	0.018	0.131
70	33.453	33.477	0.024	0.129
75	37.566	37.584	0.018	0.148
80	41.798	41.809	0.011	0.145
85	46.137	46.139	0.002	0.149
90	50.571	50.567	−0.004	0.149
95	55.086	55.078	−0.008	0.144
100	59.670	59.662	−0.008	0.144

*Calculated using the Cox-Rubinstein binomial model with 1,200 time intervals.
†Calculated using a simulation model with 100 samples of 5,040 paths and exercise boundary determined by dominant strings of both 0's and 1's in the transition zone.

6. SUMMARY AND CONCLUSIONS

This paper has presented an algorithm for valuing American options in a path simulation model and has demonstrated its accuracy by an example involving a put option on a non-dividend-paying stock for which the exact premium could be determined. The demonstration of the existence of a useful algorithm for valuing American options in a path simulation model should remove what has been perceived as a major impediment to the use of simulation models in valuing a broker-dealer's derivatives book and in analyzing the asset-liability condition of financial intermediaries.

In many situations involving the use of multifactor models to describe realistic market price behavior, simulation is the only method that can handle the American option valuation problem satisfactorily. Furthermore, it is usually straightforward to apply a simulation technique, whereas solving complicated partial differential equations numerically generally requires great care as well as sophistication in applied mathematical methods. This paper has not dealt with some of the complexities that arise in determining exercise-hold decision boundaries when multifactor stochastic models of asset price behavior are utilized. Empirical studies that I have conducted suggest that some modification to the algorithm presented in this paper is required to handle those situations adequately. For example, the bundling must often be carried out in at least two dimensions rather than the single dimension presented in this paper. Boundary points become boundary lines or surfaces.

REFERENCES

1. Black, F., Derman, E., and Toy, W. "A One-Factor Model of Interest Rates and Its Application to Treasury Bond Options," *Financial Analysts Journal* 46, no. 1 (1990): 33–39.
2. Bookstaber, R., Jacob, D. P., and Langsam, J. A. "The Arbitrage-Free Pricing of Options on Interest-Sensitive Instruments." In *Advances in Futures and Options Research,* Vol. 1 Part A, edited by F. J. Fabozzi, 1–23. Greenwich, Conn.: JAI Press, 1986.
3. Boyle, P. P. "Options: A Monte Carlo Approach," *Journal of Financial Economics* 4 (1977): 323–38.
4. Clancy, R. P. "Options on Bonds and Applications to Product Pricing," *TSA* XXXVII (1985): 97–130.
5. Cox, J. C., Ross, S. A., and Rubinstein, M. "Option Pricing: A Simplified Approach," *Journal of Financial Economics* 7 (1979): 229–63.

6. COX, J. C., AND RUBINSTEIN, M. *Options Markets*. Englewood Cliffs, N.J.: Prentice-Hall, 1985.
7. DYER, L. J., AND JACOB, D. P. "An Overview of Fixed Income Option Models." In *The Handbook of Fixed Income Securities*, edited by F. J. Fabozzi, 742–73. Homewood, Ill.: Business One Irwin, 1991.
8. GESKE, R., AND SHASTRI, K. "Valuation by Approximation: A Comparison of Alternative Option Valuation Techniques," *Journal of Financial and Quantitative Analysis* 20, no. 1 (1985): 45–71.
9. HALL, A. R. "Valuing the Mortgage Borrower's Prepayment Option," *Journal of the American Real Estate and Urban Economics Association* 13 (1985): 229–47.
10. HO, T. S. Y., AND LEE, S.-B. "Term Structure Movements and Pricing Interest Rate Contingent Claims," *Journal of Finance* 41 (1986): 1011–29.
11. HULL, J. *Options, Futures, and Other Derivative Securities*. Englewood Cliffs, N.J.: Prentice-Hall, 1989.
12. HULL, J., AND WHITE, A. "Valuing Derivative Securities Using the Explicit Finite Difference Method," *Journal of Financial and Quantitative Analysis* 25 (1990): 87–100.
13. HULL, J., AND WHITE, A. "Pricing Interest-Rate-Derivative Securities," *Review of Financial Studies* 3, no. 4 (1990): 573–92.
14. HULL, J., AND WHITE, A. "Valuing Interest Rate Options: Trinomial Versus Binomial Trees," Working Paper, University of Toronto, 1990.
15. JACOB, D. P., LORD, G., AND TILLEY, J. A. "A Generalized Framework for Pricing Contingent Cash Flows," *Financial Management* 16, no. 3 (1987): 5–14.
16. MORGAN, I. G., AND NEAVE, E. H. "A Mean Reverting Process for Pricing Treasury Bills and Futures Contracts," *Proceedings of the 2nd AFIR Colloquium* 1 (1990): 237–66.
17. PEDERSEN, H. W., AND SHIU, E. S. W. "Pricing of Options on Bonds by Binomial Lattices and by Diffusion Processes," *ARCH* 1988.2: 115–39.
18. PEDERSEN, H. W., SHIU, E. S. W., AND THORLACIUS, A. E. "Arbitrage-Free Pricing of Interest-Rate Contingent Claims," *TSA* XLI (1989): 231–65.
19. PITTS, M. "The Pricing of Options on Debt Securities," *Journal of Portfolio Management* 11 (1985): 41–50.
20. RENDLEMAN, R. J., JR., AND BARTTER, B. J. "Two State Option Pricing," *Journal of Finance* 34 (1979): 1093–1110.
21. RENDLEMAN, R. J., JR., AND BARTTER, B. J. "The Pricing of Options on Debt Securities," *Journal of Financial and Quantitative Analysis* 15 (1980): 11–24.
22. RITCHKEN, P., AND BOENAWAN, K. "On Arbitrage-Free Pricing of Interest Rate Contingent Claims," *Journal of Finance* 45 (1990): 259–64.

23. STAPLETON, R. C., AND SUBRAHMANYAM, M. G. "The Valuation of Options When Asset Returns Are Generated by a Binomial Process," *Journal of Finance* 39 (1984): 1525–39.
24. TILLEY, J. A. "An Actuarial Layman's Guide to Building Stochastic Interest Rate Generators," *TSA* XLIV (1992): 509–564.

DISCUSSION OF PRECEDING PAPER

JACQUES F. CARRIERE:

First, I want to say that I greatly enjoyed Dr. Tilley's paper. Second, I want to embellish this paper by relating it to the theory of *stopping times* and by presenting a variation of Dr. Tilley's valuation algorithm. When I read the article, it struck me that the algorithm was very similar to one that Chow, Robbins and Siegmund [1] presented in their treatise about optimal stopping times. I believe that restating the *backward induction* result given there [1, p. 50] will be useful, even though this result is well-known. The ensuing discussion is an application of this backward induction technique.

Let $t=0, 1, \ldots, N$ and let I_t denote the intrinsic value of an American option at time t. Let D_t denote the price of a zero-coupon bond that is purchased at time 0 and matures at time $t \geq 0$ with a unit redemption value and let $d_t = D_{t+1}/D_t$. Let $P_t = D_t \times I_t$ denote the present value of the option and let $\mathbf{X}_t = (X_{1,t}, \ldots, X_{m,t})$ denote the state variables. Let

$$\mathcal{H}_t = \sigma\{\mathbf{X}_s | s = 0, 1, \ldots, t\}$$

denote the history (sub-σ-field) and suppose that I_t is a measurable function with respect to $\mathcal{H}_t$. Let V_t denote the discounted value of the American option and let H_t denote the discounted holding value. According to [1], $V_N = P_N$, and for $t = N-1, \ldots, 1$, we must have

$$V_t = \max\{P_t, H_t\},$$

where $H_t = \mathrm{E}(V_{t+1} | \mathcal{H}_t)$, and for $t=0$, we must have

$$V_0 = H_0 = \mathrm{E}(V_1 | \mathcal{H}_0).$$

Moreover, Chow et al. [1] suggest that the optimal exercise strategy is to exercise at the first time that $V_t = P_t$. This means that

$$z_t = I\{V_s = H_s \quad \text{for} \quad s < t \quad \text{and} \quad V_t = P_t\}$$

is the exercise indicator function. The result in Chow et al. [1] also states that

$$V_0 = \mathrm{E}\left\{\sum_{t=1}^{N} z_t D_t I_t | \mathcal{H}_0\right\}$$

and that

$$V_0 \geq \mathrm{E}\left\{\sum_{t=1}^{N} z_t^* D_t I_t | \mathcal{H}_0\right\}$$

for any other exercise indicator function z_t^* that is measurable with respect to $\mathcal{H}_t$. In other words, V_0 is the value of the American option.

Let us give an equivalent representation of the *backward algorithm* that conforms better with the paper's algorithm. Define

$$V_t^* = V_t/D_t \quad \text{and} \quad H_t^* = d_t \mathrm{E}(V_{t+1}^* | \mathcal{H}_t).$$

Then we must have $V_N^* = I_N$, and for $t = N-1, \ldots, 1$, we must have

$$V_t^* = \max\{I_t, H_t^*\}.$$

Moreover, we must have $V_0 = H_0^*$ because $d_0 = D_1$. The exercise indicator function can now be written as

$$z_t = I\{V_s^* = H_s^* \quad \text{for} \quad s < t \quad \text{and} \quad V_t^* = I_t\}.$$

The example in this paper is a Markov example with one state variable X_t. This means that

$$\mathrm{E}(V_{t+1}^* | \mathcal{H}_t) = \mathrm{E}(V_{t+1}^* | X_t).$$

The key to calculating V_0 is to calculate or approximate the conditional expectation

$$\mathrm{E}(V_{t+1}^* | X_t) \quad \text{for} \quad t = 0, 1, \ldots, N$$

or equivalently to approximate the exercise indicator function z_t. The valuation algorithm proposed by Dr. Tilley uses an approximation of

$$\mathrm{E}(V_{t+1}^* | X_t)$$

(substeps 1 to 4) to get an approximation of z_t (substeps 5 to 8), which is subsequently used in the empirical analog of the American option valuation formula

$$V_0 = \mathrm{E}\left\{\sum_{t=1}^{N} z_t D_t I_t | X_0\right\}.$$

The major contribution of this paper is the demonstration that V_0 can be estimated by a Monte Carlo simulation, a technique that many researchers thought would be computationally impossible. The advantage of the simulation technique is that it can readily be generalized to handle options that depend on a multivariate state variable $\mathbf{X}_t$.

We know that the value of the American option is equal to

$$V_0 = d_0 \mathrm{E}(V_1^* | X_0).$$

This means that if we can construct an unbiased estimator of

$$\mathrm{E}(V_{t+1}^* | X_t) \quad \text{for each} \quad t = 0, 1, \ldots, N,$$

then the backward algorithm will yield an unbiased estimate of V_0. With this approach, it is not necessary to approximate the exercise indicator function z_t, because we are not using the equivalent option valuation formula

$$V_0 = \mathrm{E}\left\{ \sum_{t=1}^{N} z_t D_t I_t | X_0 \right\}.$$

Therefore, I investigated a variation of the algorithm that refines substep 4 and excludes substeps 6 and 7.

From the Markov property we know that

$$V_{t+1}^* = \max\{I_{t+1}, H_{t+1}^*\}$$

is a function of the state variable X_{t+1}. Fix t and suppose that we have approximated

$$V_{t+1}^* \equiv f(X_{t+1}).$$

Our problem is to estimate the conditional expectation

$$\mathrm{E}(f(X_{t+1}) | X_t).$$

By simulation, we have R replications of the pair $\{X_t, f(X_{t+1})\}$. This means that we can use *nonlinear regression analysis* to estimate

$$\mathrm{E}(f(X_{t+1}) | X_t)$$

in substep 4. I claim that the bundling method proposed by Dr. Tilley is actually a regression method, albeit crude. One way of regressing $f(X_{t+1})$ on X_t is to use a *spline* method. We found that our estimates of V_0 were almost unbiased when a spline was used in substep 4. Using the paper's example, we found that after 1,000 samples of 5,040 paths, the mean of the estimates was 7.963 and the standard deviation of the mean was $0.0822 \div \sqrt{1{,}000} = 0.0026$. Obviously, the spline method that we used yielded a biased estimator of the true value, 7.941, but its performance is comparable to Dr. Tilley's method. Other regression techniques were also investigated with good success.

REFERENCE

1. CHOW, Y. S., ROBBINS, H., AND SIEGMUND, D. *Great Expectations: The Theory of Optimal Stopping.* New York: Houghton Mifflin Co., 1971.

HANS U. GERBER AND ELIAS S. W. SHIU:

Dr. Tilley is to be congratulated for this paper, shattering the myth that American options cannot be evaluated by simulations. Indeed, Hull and White's recent article ([7], p. 22) states that "Monte Carlo simulation . . . involves working forward, simulating paths for the asset price. It can handle options where the payoff is path-dependent, but it cannot handle American options, because there is no way of knowing whether early exercise is optimal when a particular stock price is reached at a particular time."

A main difficulty in pricing American options with a finite expiration date is the determination of the optimal exercise-hold boundary (as a function of stock price and time). Because it was thought that backward induction was not possible in a simulation setting, option exercises were determined by "myopic" decision rules. Finnerty and Rose [2] have shown that such evaluations of American options are systematically biased.

The optimal exercise-hold boundary of a *perpetual* American option with a constant exercise price is a fixed stock value independent of time [see (4.6) and (4.8) below]. This makes the valuation problem easier and closed-form formulas are available. Indeed, closed-form formulas for *deferred* perpetual American options can also be obtained. A purpose of this discussion is to present these formulas. Note that an upper bound for the price of an American option with a finite expiry date is the value of the corresponding perpetual option. On the other hand, a lower bound for the price of an American option with a finite expiry date T is the difference between the value of the perpetual option and that of the deferred perpetual option with deferral period T.

1. The Stock Price Model

Let $S(t)$ denote the price of a stock at time t, $t \geq 0$. We assume that $\{S(t);\ t \geq 0\}$ is a geometric Brownian motion. Hence, for $t \geq 0$, $\ln[S(t)/S(0)]$ is a normal random variable, and there exist constants μ and σ such that

$$E(\ln[S(t)/S(0)]) = \mu t$$

and

$$\mathrm{Var}(\ln[S(t)/S(0)]) = \sigma^2 t.$$

(Our μ and σ are the same as those in Section 3 of the paper.) The assumption, expressed as a stochastic differential equation, is

$$\frac{dS(t)}{S(t)} = \left(\mu + \frac{\sigma^2}{2}\right) dt + \sigma dW(t), \quad t \geq 0,$$

where $\{W(t);\ t \geq 0\}$ denotes the standardized Wiener process. We note that the process $\{\ln S(t);\ t \geq 0\}$ has stationary and independent increments, and many of the results below can be generalized to stock-price processes with this property.

We also assume that the risk-free force of interest is constant, and it is denoted by δ. The market is frictionless, and trading is continuous. There are no taxes, no transaction costs, and no restriction on borrowing or short sales. All securities are perfectly divisible.

Extending the classical Black-Scholes model, we assume that the stock pays a continuous stream of dividends, at a rate proportional to its price; that is, there is a positive constant ρ such that the amount of dividends paid between time t and $t+dt$ is

$$S(t)\rho dt.$$

Observe that, if all dividends are reinvested in the stock, each share of the stock at time 0 grows to $e^{\rho t}$ shares at time t.

2. Equivalent Martingale Measure

A fundamental insight in advancing the Black-Scholes theory of option-pricing is the concept of *risk-neutral valuation* introduced by Cox and Ross [1]. Further elaboration on this idea was given by Harrison and

Kreps [5] and by Harrison and Pliska [6]. To rule out arbitrage opportunities in the model, we are to seek the probability measure that is equivalent to the original measure and with respect to which the process

$$\{e^{-\delta t}S(t)e^{\rho t}; \quad t \geq 0\}$$

is a martingale. This probability measure is called the *equivalent martingale measure,* or *risk-neutral measure*. The price of each option on the stock is the supremum of its expected discounted payoffs, where the expectation is taken with respect to the risk-neutral measure.

To determine the risk-neutral measure, we observe that, for each constant h,

$$\{S(t)^h/\mathrm{E}[S(t)^h]; t \geq 0\}$$

is a positive martingale, with which we can define a change of measure. We ([3], [4]) call the new measure the *Esscher measure* of parameter h and write the expectation, variance, and probability with respect to the Esscher measure as $\mathrm{E}(\cdot; h)$, $\mathrm{Var}(\cdot; h)$, and $\Pr(\cdot; h)$, respectively. For each measurable function $\psi(\cdot)$,

$$\mathrm{E}[\psi(S(t)); h] = \frac{\mathrm{E}[\psi(S(t))S(t)^h]}{\mathrm{E}[S(t)^h]}. \tag{2.1}$$

Using the formula

$$\mathrm{E}[S(t)^k] = S(0)^k e^{(k\mu + k^2\sigma^2/2)t},$$

we [3] have shown that

$$\mathrm{E}(\ln[S(t)/S(0)]; h) = \mu(h)t$$

and

$$\mathrm{Var}(\ln[S(t)/S(0)]; h) = \sigma^2 t,$$

where

$$\mu(h) = \mu + h\sigma^2. \tag{2.2}$$

The risk-neutral measure is the Esscher measure of parameter $h=h^*$, where h^* is determined by

$$S(0) = \mathrm{E}[e^{-(\delta-\rho)t}S(t); h^*], \tag{2.3}$$

or

$$\delta - \rho = \mu(h^*) + \tfrac{1}{2}\sigma^2.$$

It follows from (2.2) that, for each constant k,

$$\mu(h^* + k) = \mu(h^*) + k\sigma^2 \tag{2.4}$$

$$= \delta - \rho + (k - \tfrac{1}{2})\sigma^2. \tag{2.5}$$

3. *European Option Formulas*

Let $1_A(\cdot)$ denote the indicator function,

$$1_A(x) = \begin{cases} 1 & x \in A \\ 0 & x \notin A \end{cases}.$$

The payoff, at time T, of a *call* option on the stock exercisable at time T with exercise price X is

$$[S(T) - X]1_{(X,\infty)}(S(T)).$$

Similarly, the payoff, at time T, of a *put* option on the stock exercisable at time T with exercise price X is

$$[X - S(T)]1_{(0,X)}(S(T)).$$

For positive constants T, X and Y, let us calculate

$$c(S(0), X, Y, T) = \mathrm{E}(e^{-\delta T}[S(T) - X]1_{(Y,\infty)}(S(T)); h^*), \tag{3.1}$$

which we need later.

$$\begin{aligned} c(S(0), X, Y, T) &= e^{-\delta T}\mathrm{E}[S(T)1_{(Y,\infty)}(S(T)); h^*] \\ &\quad - e^{-\delta T}X\mathrm{E}[1_{(Y,\infty)}(S(T)); h^*] \\ &= e^{-\delta T}\mathrm{E}[S(T)1_{(Y,\infty)}(S(T)); h^*] \\ &\quad - e^{-\delta T}X\Pr[S(T) > Y; h^*]. \end{aligned}$$

To evaluate the expectation

$$\mathrm{E}[S(T)1_{(Y,\infty)}(S(T)); h^*],$$

we apply the following lemma, which is an immediate consequence of (2.1).

Factorization Lemma

Let h and k be two real numbers and $\varphi(\cdot)$ a measurable function. Then

$$E[S(t)^k\varphi(S(t)); h] = E[S(t)^k; h]E[\varphi(S(t)); h + k]. \tag{3.2}$$

Applying the Factorization Lemma [with $k=1$, $\varphi(\cdot)=1_{(Y,\infty)}(\cdot)$ and $h=h^*$] and Formula (2.3) yields

$$\begin{aligned} E[S(T)1_{(Y,\infty)}(S(T)); h^*] &= E[S(T); h^*]E[1_{(Y,\infty)}(S(T)); h^* + 1] \\ &= E[S(T); h^*]\Pr[S(T) > Y; h^* + 1] \\ &= e^{(\delta-\rho)T}S(0)\Pr[S(T) > Y; h^* + 1]. \end{aligned}$$

Hence

$$\begin{aligned} c(S(0), X, Y, T) = e^{-\rho T}S(0)\Pr[S(T) > Y; h^* + 1] \\ - e^{-\delta T}X\Pr[S(T) > Y; h^*]. \end{aligned} \tag{3.3}$$

To evaluate the two probabilities in (3.3), note that

$$\begin{aligned} E(\ln[S(T)/S(0)]; h^* + k) &= \mu(h^* + k)T \\ &= [\delta - \rho + (k - 1/2)\sigma^2]T \end{aligned} \tag{3.4}$$

by (2.5). Thus

$$\begin{aligned} &\Pr[S(T) > Y; h^* + k] \\ &\quad = \Pr(\ln[S(T)/S(0)] > \ln[Y/S(0)]; h^* + k) \\ &\quad = 1 - \Phi\left(\frac{\ln[Y/S(0)] - [\delta - \rho + (k - 1/2)\sigma^2]T}{\sigma\sqrt{T}}\right) \\ &\quad = \Phi\left(\frac{\ln[S(0)/Y] + [\delta - \rho + (k - 1/2)\sigma^2]T}{\sigma\sqrt{T}}\right), \end{aligned} \tag{3.5}$$

where $\Phi(\cdot)$ denotes the standardized normal distribution function. Consequently,

$$c(S, X, Y, T) = e^{-\rho T} S \Phi\left(\frac{\ln(S/Y) + \left(\delta - \rho + \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}\right) - e^{-\delta T} X \Phi\left(\frac{\ln(S/Y) + \left(\delta - \rho - \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}\right), \quad (3.6)$$

which reduces to the celebrated Black-Scholes formula when $X=Y$ and $\rho=0$.

With similar ease, we can calculate

$$p(S(0), X, Y, T) = \mathrm{E}[e^{-\delta T}(X - S(T))1_{(0,Y)}(S(T)); h^*]. \quad (3.7)$$

The resulting formula is

$$p(S, X, Y, T) = e^{-\delta T} X \Phi\left(\frac{\ln(Y/S) - \left(\delta - \rho - \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}\right) - e^{-\rho T} S \Phi\left(\frac{\ln(Y/S) - \left(\delta - \rho + \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}\right). \quad (3.8)$$

4. Perpetual American Options

Under the risk-neutral measure, the process $\{e^{-(\delta-\rho)t}S(t); t\geq 0\}$ is a martingale. Also, there exist two numbers θ for which the process $\{e^{-\delta t}S(t)^\theta; t\geq 0\}$ is a martingale under the risk-neutral measure. From the condition

$$S(0)^\theta = \mathrm{E}[e^{-\delta t}S(t)^\theta; h^*], \quad (4.1)$$

we obtain

$$\delta = \mu(h^*)\theta + \tfrac{1}{2}\sigma^2\theta^2 \quad (4.2)$$

or

$$\sigma^2\theta^2 + (2\delta - 2\rho - \sigma^2)\theta - 2\delta = 0. \quad (4.3)$$

The roots of this quadratic equation are

$$\theta_0 = \frac{-(2\delta - 2\rho - \sigma^2) - \sqrt{(2\delta - 2\rho - \sigma^2)^2 + 8\sigma^2\delta}}{2\sigma^2} \tag{4.4}$$

and

$$\theta_1 = \frac{-(2\delta - 2\rho - \sigma^2) + \sqrt{(2\delta - 2\rho - \sigma^2)^2 + 8\sigma^2\delta}}{2\sigma^2}. \tag{4.5}$$

Note that the quadratic function on the left-hand side of (4.3) is negative for $\theta=0$ (because $\delta>0$) and for $\theta=1$ (because $\rho>0$). From this, it follows that $\theta_0<0$ and $\theta_1>1$.

Applying the *optional sampling theorem* to the martingale

$$\{e^{-\delta t}S(t)^{\theta_1};\quad t \geq 0\},$$

we [4] derive the value (at time 0) of the perpetual American call option on the stock with exercise price X, which we denote as $C(S(0), X)$. The optimal exercise-hold boundary turns out to be

$$X_1 = \frac{\theta_1}{\theta_1 - 1} X, \tag{4.6}$$

which is independent of the time variable, and the pricing formula is

$$C(S, X) = \begin{cases} (X_1 - X)\left(\dfrac{S}{X_1}\right)^{\theta_1} & S < X_1 \\ S - X & S \geq X_1 \end{cases} \tag{4.7}$$

Similarly, applying the optional sampling theorem to the martingale

$$\{e^{-\delta t}S(t)^{\theta_0};\quad t \geq 0\},$$

we [4] obtain the value (at time 0) of the perpetual American put option on the stock with exercise price X, $P(S(0), X)$. The time-independent optimal exercise-hold boundary is

$$X_0 = \frac{-\theta_0}{1 - \theta_0} X, \tag{4.8}$$

and

$$P(S, X) = \begin{cases} X - S & S < X_0 \\ (X - X_0)\left(\dfrac{S}{X_0}\right)^{\theta_0} & S \geq X_0 \end{cases}. \tag{4.9}$$

5. Deferred Perpetual American Options

We are now ready to determine the value of the perpetual American call option that cannot be exercised for the first n years,

$$_{n|}C(S(0), X) = \mathrm{E}[e^{-\delta n}C(S(n), X); h^*]. \tag{5.1}$$

Writing (4.7) as

$$C(S(n), X) = (X_1 - X)X_1^{-\theta_1}S(n)^{\theta_1}1_{(0,X_1)}(S(n)) + [S(n) - X]1_{(X_1,\infty)}(S(n)),$$

and applying (3.1), we obtain

$$_{n|}C(S(0), X) = (X_1 - X)X_1^{-\theta_1}\mathrm{E}[e^{-\delta n}S(n)^{\theta_1}1_{(0,X_1)}(S(n)); h^*] + c(S(0), X, X_1, n).$$

We can evaluate the expectation by applying the Factorization Lemma and Formula (4.1),

$$\mathrm{E}[e^{-\delta n}S(n)^{\theta_1}1_{(0,X_1)}(S(n)); h^*] = S(0)^{\theta_1}\Pr[S(n) < X_1; h^* + \theta_1]. \tag{5.2}$$

Now, by (2.4) and (4.2),

$$\begin{aligned} \mathrm{E}[\ln[S(n)/S(0)]; h^* + \theta_1] &= \mu(h^* + \theta_1)n \\ &= [\mu(h^*) + \theta_1\sigma^2]n \\ &= \left(\frac{\delta}{\theta_1} + \frac{\theta_1\sigma^2}{2}\right)n. \end{aligned} \tag{5.3}$$

Thus

$$_{n|}C(S, X) = (X_1 - X)\left(\frac{S}{X_1}\right)^{\theta_1} \Phi\left(\frac{\ln\left(\dfrac{X_1}{S}\right) - \left(\dfrac{\delta}{\theta_1} + \dfrac{\theta_1\sigma^2}{2}\right)n}{\sigma\sqrt{n}}\right) + c(S, X, X_1, n). \tag{5.4}$$

Similarly, we can determine the value (at time 0) of the n-year deferred perpetual American put option on the stock with exercise price X,

$$ {}_{n|}P(S(0), X) = \mathrm{E}[e^{-\delta n} P(S(n), X); h^*]. \tag{5.5}$$

It follows from (4.9) that

$$ {}_{n|}P(S, X) = p(S, X, X_0, n) $$

$$ + (X - X_0)\left(\frac{S}{X_0}\right)^{\theta_0} \Phi\left(\frac{\ln\left(\dfrac{S}{X_0}\right) + \left(\dfrac{\delta}{\theta_0} + \dfrac{\theta_0\sigma^2}{2}\right)n}{\sigma\sqrt{n}}\right) \tag{5.6}$$

6. *Non-Dividend-Paying Stocks*

As the dividend yield rate ρ tends to 0, the exponent θ_1 tends to 1 and the optimal exercise-hold boundary X_l tends to ∞. It follows from (5.4) that

$$\lim_{\rho \to 0} {}_{n|}C(S, X) = S.$$

The price of a deferred perpetual American call option on a non-dividend-paying stock is the current stock price. This result reaffirms the fact mentioned in Section 3 of the paper that an American call option on a non-dividend-paying stock is never optimally exercised before its maturity date.

On the other hand,

$$\lim_{\rho \to 0} \theta_0 = \frac{-2\delta}{\sigma^2}. \tag{6.1}$$

With ρ=0, (4.8) and (5.6) become

$$X_0 = \frac{2\delta}{2\delta + \sigma^2} X \tag{6.2}$$

and

$$
{}_{n|}P(S, X) = e^{-\delta n} X \Phi\left(\frac{\ln(X_0/S) - \left(\delta - \dfrac{\sigma^2}{2}\right)n}{\sigma\sqrt{n}} \right)
$$

$$
- S\Phi\left(\frac{\ln(X_0/S) - \left(\delta + \dfrac{\sigma^2}{2}\right)n}{\sigma\sqrt{n}} \right)
$$

$$
+ (X - X_0)(X_0/S)^{2\delta/\sigma^2}\, \Phi\left(\frac{\ln(S/X_0) - \left(\delta + \dfrac{\sigma^2}{2}\right)n}{\sigma\sqrt{n}} \right), \quad (6.3)
$$

respectively.

7. *Executive Stock Options*

We were motivated in studying the pricing of deferred perpetual options by the problem of valuing executive stock options [9]. These employee stock options are American call options with a very long expiry date, but there are vesting restrictions prohibiting the executives to exercise the options for several years.

Suppose that an executive is granted an American call option that will expire m years from now, and there is a vesting period of n years, $m>n$. If the stock pays no dividends ($\rho=0$), then the option value is the same as that of the m-year European call option, $c(S(0), X, X, m)$, which does not depend on n. If $\rho>0$, there is no closed-form formula for the option value; however, an upper bound is ${}_{n|}C(S(0), X)$, and a lower bound is

$$
{}_{n|}C(S(0), X) - {}_{m|}C(S(0), X).
$$

8. *Several Risky Assets*

Our method for obtaining the risk-neutral measure or equivalent martingale measure can be generalized to the case of several securities. For $j=1, 2, \ldots, n$, let $S_j(t)$ denote the price of the j-th stock at time t and ρ_j denote the constant instantaneous dividend-yield rate. Each vector

$$
\mathbf{h} = (h_1, h_2, \ldots, h_n)
$$

determines a positive martingale

$$\left\{\frac{S_1(t)^{h_1}\,S_2(t)^{h_2}\,\dots\,S_2(t)^{h_n}}{\mathrm{E}[S_1(t)^{h_1}\,S_2(t)^{h_2}\,\dots\,S_n(t)^{h_n}]};\, t \geq 0\right\}$$

which, in turn, determines a change of measure, the Esscher measure of parameter vector **h**. The risk-neutral measure is the Esscher measure with respect to which each process

$$\{e^{-(\delta-\rho_j)t}\,S_j(t);\, t \geq 0\}, \quad j = 1, 2, \dots, n,$$

is a martingale. For further elaboration, see the second half of our paper [3].

9. Concluding Remark

In his address to the 1989 Centennial Celebration of the Actuarial Profession in North America, Dr. Tilley [8, p. 535] remarked: "In case there are any doubters about the relevance of martingales to actuarial science, one should note the new street address of the headquarters office of the Society of Actuaries, the American Academy of Actuaries, and the Conference of Actuaries in Public Practice." We certainly concur with his sentiment, as we have just presented a story of martingales.

REFERENCES

1. Cox, J. C., and Ross, S. A. "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics* 3 (1976): 145–66.
2. Finnerty, J. D., and Rose, M. "Arbitrage-Free Spread: A Consistent Measure of Relative Value," *Journal of Portfolio Management* 17, no. 3 (Spring 1991): 65–77.
3. Gerber, H. U., and Shiu, E. S. W. "Option Pricing by Esscher Transforms," *TSA* XLVI (1994), in press.
4. Gerber, H. U., and Shiu, E. S. W. "Martingale Approach to Pricing Perpetual American Options," *Proceedings of the 4th AFIR International Colloquium, Orlando, April 20–22, 1994*, Vol. 2, pp. 659–89.
5. Harrison, J. M., and Kreps, D. M. "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory* 20 (1979): 381–408.
6. Harrison, J. M., and Pliska, S. "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and Their Applications* 11 (1981): 215–60.
7. Hull, J., and White, A. "Efficient Procedures for Valuing European and American Path-Dependent Options," *Journal of Derivatives* 1, no. 1 (1993): 21–31.

8. TILLEY, J. A. "Whither Actuarial Science? Borrowing Concepts and Mathematics from Other Disciplines," *1989 Centennial Celebration Proceedings of the Actuarial Profession in North America*, 523–41.
9. YOUNG, C. K. "What's the Right Black-Scholes Value?" *Financial Executive* 9, no. 5 (Sept./Oct., 1993): 57–9; Letters, 9, no. 6 (Nov./Dec. 1993): 5, 6, 57; and 10, no. 1 (Jan./Feb. 1994): 5–6.

BENJAMIN W. WURZBURGER*:

Dr. Tilley has provided a stimulating and novel approach to the very important problem of valuing American-style options. Section 1 of this discussion contrasts the Tilley approach with the mainstream and finds a fundamental appeal in the Tilley approach. It is most noteworthy that, despite the assertion in a leading textbook [5, p. 402] that "Unfortunately, American style options cannot be priced with Monte Carlo [i.e., stochastic] simulations," Tilley (page 499) does claim to "present a general algorithm for estimating the value of American options [via simulation]." Section 2, the key section in this discussion, raises some theoretical and practical concerns about the Tilley solution for his stock option example. In general, finance/economic theory deals with models in which the decision-makers act optimally, subject to the then-available (the "ex ante") information. In Tilley's model, however, the decision-makers (the decision whether to exercise the option or to hold) do take advantage of ex post information (information not yet revealed to the market), but do so with arbitrary rules of thumb that are clearly suboptimal. Section 3 follows with some brief comments on the tables.

The Tilley algorithm is intended to be applicable to the valuation of American options on any instrument. Section 4 notes some potential problems in extending the analysis from the stock option example to the more complex case of fixed-income/term structure dynamics. Tilley's example relies on his ability to order the stock prices, but the yield curves within many term structure models (for example, even the single-factor version of Heath, Jarrow and Morton [3]) do not unambiguously lend themselves to a simple ordering. Section 5 concludes.

*Dr. Wurzburger, not a member of the Society, is in the Investment Policy and Research Department at the John Hancock Mutual Life Insurance Company.

1. The Lattice Approach and the Fundamental Appeal of the Tilley Stochastic Approach

Many researchers advocate valuing American options via backward recursion on a connected lattice, say, a binomial lattice [4] or a trinomial lattice [6]. These researchers seem to be relatively unconcerned that strong restrictions must be imposed on the underlying term structure dynamics in order that the lattice structure be applicable (that is, in order that the paths recombine). Ho and Lee [4] represent an extreme version of this lack of concern. Although the Ho and Lee algorithm implicitly requires (in the continuous-time limit) that yield curves shift in a parallel fashion, Ho and Lee do not even mention that they are implicitly imposing this strong restriction on the term structure dynamics.

Tilley represents an important counterweight to the prevailing school. His stated goal is to value an American option under any (arbitrage-free) model of the asset dynamics. In my view, Tilley has his priorities right. We should not adopt term structure models just because they conform to our solution algorithms; we should try to create algorithms that can handle our views about the term structure dynamics.

Having mentioned that the Tilley algorithm is intended to be applicable to any term structure model, we should also cite the Amin-Morton [1] exponential tree (all 2^N possible paths), a method that can handle any term structure dynamics. The Amin-Morton technique converges onto the true value as the step size shrinks; the problem is that the tree grows exponentially, and computer limitations therefore restrict us to a small number of periods. (Tilley—not in reference to Amin-Morton—characterizes [page 504] the exponential tree as a "computational infeasibility," but Amin and Morton achieve feasibility by restricting themselves to 7 to 10 periods.) The brute-force Amin-Morton technique does lack Tilley's ingenuity, but it also avoids the arbitrary rules, random scenarios, and computer programming complexity associated with the Tilley algorithm. It would be interesting if Tilley would report the computer run time for his technique, so we could determine which of the two (Amin-Morton or Tilley) provides the better accuracy for a given computer run time.

2. The Tilley Rule: A Suboptimal Processor of Ex Post Information

The "tentative" decision (page 505, step 5), whether to exercise the option or to hold, does rely on ex post information. The tentative decision is in turn transformed into a "new" (that is, final) decision; the transforming rule (steps 6 and 7) is asserted rather than derived via optimization.

a. Incorporation of Ex Post Information

The tentative exercise decision (step 5) relies on information about what set of paths actually emanates from the bundle, that is, information about which subset of the feasible scenarios was in fact randomly drawn by the computer. This represents information that is not available to the option holder at the exercise decision time.

Let us illustrate with an artificial numerical example. Suppose we are dealing with a put option at a strike price of 100, and consider a randomly generated bundle of 50 paths that emanate from nodes where the stock price is currently 90. Suppose the process that generates stock prices has an underlying drift of 5—the theoretical expected value of the future stock price is 95. In this finite random sample of size 50, however, the expected value of 95 will not be exactly realized—suppose, say, the average price of these particular 50 paths ends up at 85.[1] This information about the future, that stock prices will in fact be falling on average, will induce the put-option owners to hold instead of exercising early. (They know something that the market, which expects a mean of 95, does not know.) This reliance on "ex post" information allows the option owners to make excellent exercise/hold decisions and contributes an upward bias to the option valuation.

[1]The technique of antithetic variables (which Tilley [1] recommends in conjunction with stratified sampling) can usually ensure that the sample mean is equal to the theoretical expectation. (In his present paper, Tilley recommends stratified sampling but does not refer to antithetic variables.) It is not, however, sufficient to just ensure that the sample mean (of the 50 paths) is equal to the theoretical mean. The higher moments are also relevant. For example, a bundle in which the sample variance exceeds the theoretical variance is likely to induce the decision-maker not to exercise the option—and thus take advantage of ex post information. It is possible that the recommended stratified sampling may mitigate this sort of problem, but it would be helpful for the author to provide more details on this.

b. The Imperfect Rule

Having ordered the states (back in step 1) and having made a set of tentative exercise decisions, the author proceeds to construct a sequence of 0's and 1's, where the 0's indicate a "tentative" hold decision, and the 1's a "tentative" exercise decision. The 1's tend to the right of the sequence, at the nodes with the lower stock prices. He then proceeds to identify that point in the sequence at which the 1's begin to dominate; the transition point is denoted as a "sharp boundary."

The Tilley rule for identifying the "sharp boundary" (the first rule is on page 505, while page 506 offers an alternative rule) is asserted arbitrarily, rather than derived from any maximizing rule or general principle of statistical inference. His first rule, for example, identifies the "sharp boundary" as "the start of the first string of 0's, the length of which exceeds the length of every subsequent string of 0's."

While his rule does seem reasonable for his artificial example, namely, the sequence of 0's and 1's on page 505, I thought of the following example in which his rule looks less plausible.

00...0001101101101101101101000 ↓ 11...11

The arrow indicates where his rule (actually both his rules) would indicate that the 1's begin to dominate; I submit, however, that most observers would identify the general switch to 1's as occurring much further to the left. An imperfect rule is a factor tending to downwardly bias the estimate of the option premium.

In his primary numerical example [page 512], Tilley overestimates the option value: an estimate of $7.97 versus a true (via Cox-Ross-Rubinstein) value of $7.94. In this example, the upward bias from the ex post information more than offsets the downward bias from the imperfect rule. Tilley seems to welcome the partially offsetting downward bias from the imperfect rule and advises [page 515] not to try too hard to perfect the rule.

c. Comments

I regard this reliance on a deliberately imperfect rule as a way of partially offsetting another source of upward bias as a very dicey procedure. How did the author come to advocate his recommended rule? Did he experiment with better rules and with worse rules, and reject the better (less imperfect) rules because they contributed insufficient downward bias to offset the upward bias from the ex post information, and reject the worse rules because they tended to provide too much downward bias? If that is the case, we should be wary of extending the rule to another context (for example, the fixed income) without first somehow verifying that the biases also do roughly offset in that context.

The author does refer [pages 514 and 515] to the presence of biases. I hope that our discussion about ex ante and ex post information has helped clarify the conceptual underpinnings of the author's approach.

From a methodological perspective, it is probably easier to accept the reliance on an imperfect rule (for translating the "tentative" decision into a "sharp boundary") than to accept the utilization of the ex post information (in the construction of the "tentative" decision.) It is not that unusual to find models in financial economics that incorporate an imperfect suboptimal rule, as it is often very difficult to model optimal behavior. (This in turn suggests that economic decision-makers may not in fact actually optimize.) What is really unusual is that we are dealing with a model in which the agents can take advantage of ex post information.

3. The Tables

Tables 2 through 4 all display the interesting property that the bias becomes negative as the put option becomes more in-the-money. I'd be interested in Dr. Tilley's explanation of this phenomenon.

The transition from Table 2 to Table 3 involves two changes: a change in the Tilley "bundling parameter alpha," as well as a change in the rule for selecting the "sharp boundary." It would be interesting to see the separate impacts of these two changes in procedure.

4. Extending the Tilley Procedure beyond the Stock Price Context: The Ordering Problem

The Tilley algorithm for the stock option example relies on the ability to order the states in a special way: when he divides and bundles the states in accord with this ordering, the states within a bundle are similar (see page 504, "reorder the stock price paths by stock price . . ."). While it is clear how to order stock prices, it is by no means obvious how the user ought to order and bundle yield curves. The author should offer the reader detailed guidance on how to order yield curves.[2]

To better appreciate the problem, consider the Heath-Jarrow-Morton model (even the single-factor version) of the term structure dynamics. (This model has attracted much attention, both from academics as well as from some investment consulting firms.) Heath, Jarrow and Morton do not require that the short rate be Markovian, and the relevant history of the realized dynamics cannot be captured by a single yield. Since the task of ordering the states basically involves summarizing each state by a single number, the ordering step appears problematic in the Heath-Jarrow-Morton context.

The need to order the states would not cause any problems for the less general Hull and White or Ho and Lee term structure models. Hull and White restrict themselves to Markovian models, so the short rate captures all the relevant history, and it would therefore be reasonable to order states according to the short rate. The same ordering is also applicable to Ho and Lee, inasmuch as Ho and Lee can be regarded as a special case of Hull and White [see Hull and White].

5. Concluding Remarks

At the beginning of this discussion, we quoted the Hull textbook that "American options cannot be priced with Monte Carlo methods." Tilley's algorithm does price American options with Monte Carlo methods. The algorithm is, however, subject to some serious shortcomings: it can be viewed as a suboptimal processor of ex post information, and it is

[2]Set theory—the "Well Ordering Theorem"—does assure us that every set (for example, the set of yield curves) can be ordered. (The proof involves Zorn's Lemma/Axiom of Choice.) The question is how to order the states in a meaningful way, so that the states in the same bundle are similar.

Tilley's stock example readily lends itself to a meaningful ordering because he assumes that stock prices follow the Black-Scholes lognormal and are therefore Markovian: in order to determine the future distribution, we need only know the current price—there is no relevant incremental information in the path history. The empirical evidence indicates that stock prices are in fact non-Markovian, as a period of high volatility is likely to be followed by more high volatility. (For a nice survey on this topic, see Engle [2].) Nevertheless, the assumption that the stock price is Markovian would probably be regarded by most people as less problematical than the assumption that the short rate is Markovian.

Tilley (page 518) does refer to complexities arising in the "multifactor case," but we have noted that even the single-factor Heath-Jarrow-Morton model of the term structure does raise serious complexities for ordering. My point about "single-factor/multifactor" might just be reflective of a discrepancy in terminology; it is, however, standard in the literature to refer to the non-Markovian Heath-Jarrow-Morton model as a single-factor model.

not readily generalizable to the non-Markovian case. (For the Markovian case, we do not need to rely on Monte Carlo methods, but can instead use Hull and White or other similar methods.) Despite these problems, the author deserves credit and recognition for doing—albeit imperfectly—something that textbooks said could not be done.

ACKNOWLEDGMENT

I wish to thank Professor David Heath for a very helpful conversation.

REFERENCES

1. AMIN, K., AND MORTON, A. "Implied Volatility Functions in Arbitrage Free Term Structure Models," *Journal of Financial Economics* 35 (1994): in press.
2. ENGLE, R. F. "Statistical Models for Financial Volatility," *Financial Analysts Journal* 49 (January/February 1993): 72–78.
3. HEATH, D., JARROW, R., AND MORTON, A. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Evaluation," *Econometrica* 60 (1992): 77–105.
4. HO, T. S. Y., AND LEE, S.-B. "Term Structure Movements and Pricing Interest Rate Contingent Claims," *Journal of Finance* 41 (1986): 1011–29.
5. HULL, J. *Options, Futures, and Other Derivative Securities*. 2nd ed. Englewood Cliffs, N.J.: Prentice-Hall, 1992.
6. HULL, J., AND WHITE, A. "One-Factor Interest-Rate Models and the Valuation of Interest-Rate Derivative Securities," *Journal of Financial and Quantitative Analysis* 28 (1993): 235–54.
7. TILLEY, J. A. "An Actuarial Layman's Guide to Building Stochastic Interest Rate Generators," *TSA* XLIV (1992): 509–64.

(AUTHOR'S REVIEW OF DISCUSSIONS)

JAMES A. TILLEY:

I would like to thank Dr. Carriere, Drs. Gerber and Shiu, and Dr. Wurzburger for their discussions of my paper.

Drs. Gerber and Shiu present results on the exact analytical evaluation of certain perpetual American options utilizing an equivalent-martingale-measure technique they have developed. Their work provides a mathematical upper bound on the value of the corresponding finite-expiration American option. They also derive a lower bound for any such option by applying their technique to the valuation of a perpetual American option with a deferred start. Thus, the results of Drs. Gerber and Shiu offer a "reasonableness test" of a numerical procedure for evaluating the prices of American options in the class to which their work applies.

Many readers of my paper have done work of their own to reproduce the results shown in the paper and to investigate extensions of the techniques that I have developed. Dr. Carriere is the only one to have submitted a formal written discussion. I consider his discussion valuable in two important respects. First, he correctly states that the valuation of American options is a "stopping time" problem and that the exercise-hold indicator function gives the stopping times. His mathematical description of the problem is elegant and concise, and his characterization of my approach is accurate.

Second, Dr. Carriere has investigated variations of the algorithm presented in my paper, concentrating on nonlinear regression analysis to evaluate the conditional expectations needed to price the American option. He indicates that after 1,000 samples of 5,040 paths each (the same experiment reported in my paper), he finds a positive bias of +2.2 cents in his estimator. By multiplying his reported standard deviation of the mean by the square root of 1,000, a standard error of the sample-size-5,040 premium estimator equal to about 8.2 cents is found, 2.9 cents larger than the 5.3-cent standard error reported in my paper. I believe that a smaller standard error is achievable when the backward induction method is used to compute only the exercise-hold indicator function than when it is used to determine directly the option premium. The intrinsic value of the option is known with complete accuracy at each epoch on each of the simulated paths of stock prices. That information should be used in the most efficient way possible. Pushing all the estimation error, the source of which is the *finite* size of the sample of paths, into the determination of the exercise-hold indicator function achieves that efficiency.

Dr. Wurzburger makes comments, raises questions, and notes concerns in four areas. I address each of these broad areas. To begin with, I would like to note that a demonstration that American options can be valued properly by means of a Monte Carlo simulation, *no matter how simple the example* (say, a one-factor Markovian process), does fully "shatter the myth that American options cannot be evaluated by simulations," in the words of Drs. Gerber and Shiu. However, Dr. Wurzburger notes correctly that it is important to determine whether the simulation method is efficient and extendible to more complicated situations, such as Heath-Jarrow-Morton (HJM) term structure processes.

I disagree with Dr. Wurzburger's characterization of the simulation technique in terms of "arbitrary rules, random scenarios, and computer programming complexity." I comment later on the use of rules to determine a sharp boundary between the decision to exercise the option and the decision to hold the option. As to randomness, everyone prefers to avoid simulation approaches, but it is well-known that the computational intractability of "high-dimensional" integration problems, of which the valuation of American options for realistic term structure processes is an example, is broken via randomization. In my opinion, Dr. Wurzburger is right to favor realistic, and necessarily complicated, models to unrealistically simplistic models that have the apparent advantage of computational ease. Realistic models almost invariably require the utilization of simulation.

I have personally implemented the Amin-Morton method to which Dr. Wurzburger refers. The computer program for the Amin-Morton method cannot be characterized as either shorter or less complex than the program that implements the method described in my paper—in my view, neither is particularly long or complex. The Amin-Morton paper is valuable in demonstrating that it may not always be necessary to use a very large number of periods to gain acceptable accuracy, but even Amin and Morton prefer to be able to use a much, much larger number of periods than ten. Many problems encountered on a derivatives trading desk demand that a considerably larger number of periods be analyzed. The computational infeasibility of an exponential tree is inescapable. Alternative approaches must be developed.

Dr. Wurzburger's references to "incorporation of ex post information" and "suboptimal processor of ex post information" merit commentary. For even the most general problem, computing the conditional expectation described in substep 4 relies only on historical and current information and on assumptions made about the stochastic process governing the evolution of stock prices in the future—in other words, it relies on ex ante information only. The technique is essentially equivalent to calculating the discounted expectation of one-period-ahead prices in any multinomial lattice. In a lattice, *two or more paths* emerge from each node. In the typical simulation model, only a *single path* emerges from any epoch on any path. In general, more than a single path must be used to calculate an option's holding value accurately. To overcome this problem in the simulation model, paths that are considered "nearly" the same are bundled together at each epoch, and then all the paths in a bundle are used to estimate the required mathematical expectation. By this procedure, in effect, a complicated lattice is created.

The results from the simulation approximate the exact answer. The simulation methodology has validity if the error tends to zero as the number of paths tends to infinity. For the example in my paper, as the total number of paths in the simulation is increased to infinity, one can both shrink the "radius" of the bundle to zero *and* increase the number of paths within the bundle to infinity. In that limit, it can be said that the agent makes his or her exercise decision by simulating a number of paths that originate from the current observable stock price. In other words, only ex ante information is used. While the description of this procedure has the right limiting behavior, one might argue that it remains to be shown that the actual approximations also have the right limiting behavior. I leave the proofs to mathematicians who have the powerful tools required to prove such limit theorems. In the interim, we should be highly encouraged by the empirical results that I have presented. They provide a strong inducement to undertake more empirical testing and more theoretical analysis of the simulation methodology.

Dr. Wurzburger's numerical example serves a useful purpose, despite the deliberate hyperbole. With any reasonable choice of stock price volatility, a 95 versus 85 discrepancy in the expected one-period-ahead stock price would be a highly unlikely statistical event, even for a sample size as small as 50 paths. However, for sake of example, suppose the event does occur. The resulting error in the option premium estimate, an upward bias as Dr. Wurzburger correctly concludes, arises because mathematical optimization techniques are efficient at exploiting biases inherent in finite-sample-size statistical fluctuations from the mean. The inaccuracy originates in the finiteness of the sample size, the affliction of any Monte Carlo method no matter how carefully or cleverly variance is controlled, but is made more prominent in American option valuation because it is an *optimization* problem. Dr. Carriere discovered the same phenomenon in his independent exploration of several variations on the simulation methodology presented in my paper.

Dr. Wurzburger expresses a concern that I have imposed a rule that requires the agent to look forward and that this invalidates, or severely weakens, my analysis. As most recently pointed out by Dr. Carriere, the time at which an agent optimally exercises his or her American option is a "stopping time." This means that the determination whether to exercise or to hold the option can depend only upon information that the agent has up to that date. Although the agent clearly cannot look into the future, he or she is permitted to model the future and to investigate what the model states about the future. There should be no objection to a procedure in which the agent at any point in time, starting from the then-existing yield curve or stock price, simulates future yield curves or stock prices and uses that simulation to determine optimal exercise, and hence option value. In fact, this procedure is at the heart of much of today's well-accepted and heavily used option-pricing methodology.

Dr. Wurzburger's comments about the rules for creating a sharp exercise-hold boundary are similar to those I have received from many people who have reproduced the results of my paper. I did indeed experiment with different reasonable rules, but I did not find any that gave bad results. Nor does it seem that anybody else did. The results reported by Dr. Carriere are quite close to mine. Why are option premium estimates at least somewhat sensitive to whether or not a "sharp" boundary is determined, but relatively insensitive to the particular rules employed for determining a sharp boundary? The reasons are supplied throughout my paper. I collect them here, and express them a little differently, to respond to Dr. Wurzburger's inquiry.

Solving the option valuation problem is equivalent to determining the boundary point, line, or surface (depending on the dimensionality of the problem) between exercise and hold that optimizes the value of the premium estimator equation shown in my paper. (As noted above, boundary determination algorithms must be based on ex ante information only.) In fact, option pricing is a problem in the calculus of variations. Because the true solution is "optimal," one knows that small deviations in the boundary away from the optimal boundary will generally result in very small deviations in the option premium away from its true value. In the usual language of the ordinary calculus, the first derivative of a function is zero at a local extremum. Unless the second derivative at the extremum is large in magnitude, small deviations away from the extremum result only in small changes in the value of the function away from its value at the extremum. Thus, different sensible rules for determining a sharp boundary between exercise and hold are very likely to produce essentially the same option premium. If that is so, why use any rules at all to establish a sharp boundary? Why not rely on the "transition zone" estimator?

The answer lies in the desire to accelerate the convergence of the option premium estimation. In my paper, I note that substeps 6 and 7 (the sharp-boundary rules) are completely unnecessary in the limit of an infinite number of paths, assuming that the limit is taken in a proper manner. Substeps 6 and 7 are very helpful in accelerating the convergence. In practice, one utilizes only "modest" sample sizes between 1,000 and 100,000, and thus wants to enhance convergence. My choice of sharp-boundary rules based on dominant strings of 0's and 1's certainly is arbitrary, but hopefully is reasonable. Several people have pointed out a rule that executes faster on a computer and that, incidentally, cures the difficulty described in Dr. Wurzburger's example. The rule is easily stated as moving all the 0's in the string to the left and all the 1's to the right. The boundary is then "obvious" and gives balance to all the 0's and 1's computed in substep 5. For the reason described above, use of this improved rule does not materially alter the estimates of the option premiums.

I regard the discussion in my paper of the source of bias as essentially complete. Ignoring the specific algorithm in the paper for the moment, it can be stated that the exact solution to the calculus-of-variations problem for any finite sample of paths produces an upwardly biased estimator of the true option premium. The upward bias tends to zero as the sample size increases and is equal to zero in the limit of infinite sample size. However, the calculus-of-variations problem for the finite sample of paths seems to be much too hard to solve. One must resort to a good approximation of the type described in my paper—namely, bundling paths to implement the traditional backward induction procedure, and then utilizing some sharp-boundary rule to accelerate convergence. The specific algorithm adopted may itself introduce some bias but, if constructed properly and if the sample size is large enough, should produce an option premium estimate that is close to the finite-sample-size optimal calculus-of-variations solution. However, the direction of the bias relative to the true option premium is indeterminate.

I accept Dr. Wurzburger's criticism of my comment about not trying too hard to perfect the approximation to the finite-sample-size optimal calculus-of-variations solution. One approach would be to try to perfect the approximation, and then use as large a sample of paths as is computationally feasible in order to reduce the positive bias to an acceptable level. Another approach would be to try to develop a different algorithm that has zero bias, even for finite sample size. Positive bias in the premium estimator can probably be reduced by creating a *distinct bundle for each path at each epoch* and ensuring that the paths included in that bundle produce very nearly the proper one-period-ahead expected stock price. Utilizing the suggested sharp-boundary rule described above might eliminate the source of negative bias. I do not know whether it is possible to construct a finite-sample estimator that has zero bias.

Dr. Wurzburger correctly points out that more research needs to be undertaken in determining the bundling algorithm. The bundling algorithm must address the multidimensional nature of the stochastic process describing the dynamics of the yield curve or of the option. The radius of the bundles must tend to zero as the number of paths in the simulation increases. As Dr. Wurzburger's comments might suggest, it is also desirable that the metric used to compute these radii fully captures the relevant dynamic process. For the simple stock case presented in my paper, the appropriate metric is the usual metric on one-dimensional Euclidean space; for path-average (Asian) options on stocks, a metric that incorporates two dimensions is necessary; for yield-curve-based options, it would appear that a metric measuring the distance between entire yield curves is necessary, although I have obtained good results using a one-dimensional metric on the option's payoff formula. I would like to comment in a little more detail on some of these more complicated situations in which the stochastic processes are non-Markovian.

First, I have done extensive work during the last year with Professor Neave of Queen's University, who has developed numerical algorithms to evaluate path-average stock options exactly. The valuation of such options provides a good test of the simulation/bundling methods described in my paper because the underlying process is non-Markovian. Yet, the process can be made Markovian by considering both the current stock price and the current value of the path-average stock price to define the current state of the world. The American option valuation problem for path-average stock options thus requires bundling in two dimensions: stock price in one dimension and path-average stock price in the other dimension. The necessary ordering and sharp-boundary algorithms then become straightforward extensions of the one-dimensional versions presented in my paper. Acceptable commercial accuracy, equal to a fraction of the option's bid-offer spread in the market, is achievable in many situations by using simulations based on 10,000 paths of stock prices. My research on this problem has highlighted the importance of estimating a sharp boundary line between exercise and hold for some classes of path-average options.

Second, whether the underlying stochastic process is Markovian or not or whether it is single-factor or multifactor, a lower bound on the premium for *any* American option, ignoring the upward bias due to the finiteness of the sample size, can be found by a straightforward application of the method described in my paper. Suppose the payoff formula for the option is f. Hence, the option's intrinsic value is given by $\max[0, f]$. The paths are ordered by the value of f and then arranged into appropriate bundles on the basis of the ordered values of f. If the process for f is non-Markovian and/or multifactor, the option premium obtained by this application of the techniques described in my paper must, in the limit of infinite sample size, be a lower bound to the true option premium, because the determination of the boundary at each epoch as a *point*, instead of a *line* or a *surface*, is suboptimal. The lower bound can be improved by bundling in two dimensions: the payoff function f in one dimension and the short-term interest rate, for example, in the other dimension.

Third, a fully general Heath-Jarrow-Morton model of the term structure of interest rates is non-Markovian, even for a one-factor version of the model! The algorithm presented in my paper can be modified by creating a distinct bundle at each epoch on each path. As always, the paths included in a bundle must be "close" to each other. In this situation, the appropriate definition of closeness, that is, the choice of metric, should relate to the *entire* yield curve. Instead of the entire yield curve, one might take "key benchmark" maturities along the yield curve and include paths in the bundle only if the yields at all the key benchmark maturities are sufficiently close. Unfortunately, even using only benchmark maturities, the bundling algorithm becomes a high-dimensional "nearest-neighbor" problem. It is computationally intensive, even with the massively parallel processing equipment available today. However, I have found that the lower bound obtained by using the one-dimensional metric based on the option payoff formula suffices for many yield-curve-based options.

I have focused on simulation models because many realistic features of the underlying stochastic processes can generally be incorporated quite easily. For most sufficiently realistic models, simulation is the natural method for solving security valuation problems. Sometimes, it is the only method available. Variance-reduction techniques can render simulation models quite efficient for securities with path-dependent payoffs and for non-path-dependent European options. Efficient algorithms for valuing

American options accurately in these realistic models need to be developed. A principal purpose of my paper was to induce many different researchers to bring their brainpower to this important problem. I am gratified by the number and quality of responses I have received, both the formal written discussions presented in these *Transactions* and the informal discussions I have had with many academics and practitioners.