Discrete-Time Continuous-State Interest Rate Models

Michael A. Sullivan [*]

Risk Analysis Division, Comptroller of the Currency, 250 E St. S.W.,Washington, DC 20219.

Abstract

We show how to implement arbitrage-free models of the short-term interest rate in a discrete-time setting that allows a continuum of rates at any particular date. Discrete time allows approximate pricing of interest rate contingent claims that cannot be valued in continuous-time models. It is usually associated with discrete states, with possible interest rates restricted to a limited number of outcomes, as in the lattice model of Hull and White (1994). We develop a method for approximating the prices of contingent claims without that restriction. We use numerical integration to evaluate the risk-neutral expectations that define those prices, and function approximation to efficiently summarize the information. The procedure is simple and flexible. We illustrate its properties in the extended Vasicek model of Hull and White and show it to be an effective alternative to lattice methods.

JEL classification: C63; G13

Keywords : American options; Bond options; Lattice approach; Gaussian quadrature; Chebyshev approximation.

---

1. Introduction

Arbitrage-free models of the term structure are widely used to value interest rate derivative securities. They have the advantages of being able to fit the current term structure exactly and of delivering prices of a variety of claims whose payoffs are contingent on interest rate movements. One class of models specifies the dynamics of the short interest rate and builds up the term structure and pricing results based on arbitrage relations. Continuous-time models that have a continuum of possible short rates (continuous state) can price bonds of various maturities and associated European options. However, not all securities can be priced explicitly in a such a setting, which explains the need for effective approximation techniques.

Lattice methods are widely used to price interest rate contingent claims in arbitrage-free interest rate models. They assume that rates evolve randomly at a set of discrete dates and restrict rates at each date to a finite number of possible outcomes (discrete state). Prices calculated in that setting are taken as approximations to those in the corresponding continuous-time continuous-state model. The usual assumption of mean reversion of short rates complicates the application of the original binomial lattice model of Cox, Ross, and Rubinstein (1979). Hull and White (1994) develop an effective technique for preserving recombination of the lattice that allows application of the method to interest rate options. Alternatively, Black, Derman, and Toy (1990) begin with a recombining lattice specification and show how to derive the term structure and the prices of derivative securities from that process.

These approaches have several appealing features. They follow the logic of risk-neutral pricing in their recursive valuation, deliver prices that converge in the limit to their continuous-time counterparts, and can be adapted to a variety of security specifications. However, the efficiency of lattice methods as an approximation method is diminished by their uneven convergence. Derman, et al. (1995) demonstrate this shortcoming for barrier options by showing substantial fluctuations in approximate prices even for a large number of time steps. Persistent errors are caused by a misalignment of the discrete outcomes allowed on the lattice and critical points in the payoff function of the contingent claim. Techniques to improve the

behavior of lattice methods can reduce their other advantage, ease of construction.

This paper develops a method to approximate prices of interest rate contingent claims while allowing for a continuum of short rates on a set of discrete dates. Rather than using a discrete probability distribution, we assume that next period's short rate has a continuous distribution. Instead of approximating the prices of interest rate derivatives only at discrete points, we approximate the function which relates their value to the current period's short rate. In this discrete-time continuous-state setting, the goal is to approximate the pricing function for a security. We use known results in numerical integration and function approximation to reach that goal. Judd (1992) develops a similar method for application to a stochastic growth model. Sullivan (2000) applies the technique to stock options, and it is extended here to a situation in which the state variable is not an asset price.

The following section lays out the Vasicek model in a discrete-time setting and derives analytical solutions for bond prices and European options. This is used to demonstrate the approximation errors of imposing discrete states. Section 3 explains how to apply Gauss-Legendre quadrature and Chebyshev polynomial approximation to price American-style options. Section 4 demonstrates the convergence, accuracy, and efficiency of the method with particular comparison to the Hull and White lattice approach. Section 5 concludes with an outline for extensions of the method.

2. Analytical results for the Vasicek model in discrete time

Vasicek (1977) develops a continuous-time model of the interest rate that incorporates mean reversion and a risk premium. Hull and White (1994) adapt it as a specification of the risk-neutral process for the interest rate in the spirit of arbitrage-free models of the term structure. They also extend the model to allow for a time-varying mean reversion level to fit the current term structure exactly. We follow their interpretation to specify a discrete-time continuous-state model of the short-rate, assuming a time interval of length Δt for a mean reverting process subject to random shocks that are normally distributed:

$$r_{t+\Delta t} - r_t = a(h - r_t)\Delta t + \sigma z_{t+\Delta t}\sqrt{\Delta t} \tag{1}$$

where r_t is the short interest rate at time t, a determines the speed of mean reversion to the level h, σ is constant, and z is a standard normal shock to the interest rate. The prices of bonds of various maturities and of European-style options on those bonds can be derived from this specification of the short-rate process.

Zero coupon bonds are priced using the risk-neutral value relation:

$$B(t,\tau) = e^{-r_t \Delta t} E\left[B(t + \Delta t, \tau)|r_t\right] \tag{2}$$

where $B(t,\tau)$ is the time t value of a bond that pays \$1 at time τ with an associated yield $y(t,\tau) = -\ln B(t,\tau)/(\tau - t)$. The expectation is taken with respect to the risk-neutral distribution of $r_{t+\Delta t}$ conditional on r_t, so the current price of the bond depends on the current value of the short rate. To evaluate the time t bond price, begin at time $\tau - \Delta t$ by pricing a one-period bond that matures at time τ and finding its relation to the contemporaneous short rate. Use that relation to get the discounted expectation that defines the pricing function for a two-period bond. Continue by evaluating the expectation of the derived function to get the pricing function for a bond with an additional period to maturity until the current date is reached.

For example, the time $\tau - \Delta t$ yield on a bond that pays \$1 at time τ must equal $r_{\tau - \Delta t}$. Moving back one period using risk-neutral valuation,

$$B(\tau - 2\Delta t, \tau) = e^{-r_{\tau - 2\Delta t}\Delta t} E\left[e^{-r_{\tau - \Delta t}\Delta t}|r_{\tau - 2\Delta t}\right] \tag{3}$$

Since future short rates are assumed to be normally distributed, the expectation can be explicitly evaluated to give a two-period yield of

$$y(\tau - 2\Delta t, \tau) = \frac{1}{2}\left[r_{\tau - 2\Delta t}(1 + (1 - a\Delta t)) + ha\Delta t - (\sigma\Delta t)^2/2\right] \tag{4}$$

Because this yield is a linear function of the short rate, the procedure can be repeated to solve for yields and prices of bonds with various maturities. The general solution for the yield on an m-period zero coupon bond is

$$y(t, t + m\Delta t) = \frac{1}{m}\left[\theta_m r_t + \sum_{j=1}^{m-1}\gamma_j\right] \tag{5}$$

3

$$\text{where} \quad \theta_m = \sum_{j=0}^{m-1} (1 - a\Delta t)^j \quad \text{and} \quad \gamma_j = \theta_j ha\Delta t - (\theta_j \sigma \Delta t)^2/2 \tag{6}$$

Because the parameters of the interest rate process are assumed to be constant over time, the yield for a given maturity depends only on the short rate. If the parameters of the short rate process changed over time, then there would be a time dependence in the functions relating yield to the state variable.

European options can also be priced in this model. Consider a call option that matures at time $t + n\Delta t$ with a strike price K written on a zero coupon bond that matures at time $t + m\Delta t$, where $m > n$. Its current price is

$$C(t) = B(t, t + n\Delta t)E\left[(B(t + n\Delta t, t + m\Delta t) - K)^+ | r_t\right] \tag{7}$$

where the expectation is taken with respect to the forward neutral measure as in Jamshidian (1989). Using results in Rubinstein (1976), the option value can be stated as

$$C(t) = B(t, t + m\Delta t)N(b) - B(t, t + n\Delta t)KN(b - \sigma_{m-n}) \tag{8}$$

where

$$b = \ln\left(B(t, t + m\Delta t)/KB(t, t + n\Delta t)\right)\sigma_{m-n} + \sigma_{m-n}/2 \tag{9}$$

and the variance of the $m - n$ period yield is

$$\sigma_{m-n}^2 = (\theta_{m-n}\sigma)^2 \sum_{j=0}^{n-1} (1 - a\Delta t)^{2j}\Delta t \tag{10}$$

A discrete-time version of Jamshidian's continuous-time model, this model converges to his results as Δt goes to zero. Since both are single-factor models, the routine explained in Jamshidian (1989) can be used to price European options on coupon bonds. But neither allows an explicit solution for American-style options.

Hull and White (1994) put forward a method for approximating the continuous-time model using a sequence of trinomial distributions over discrete time intervals. They show how to match the mean reversion of the short rate while ensuring that the number of possible short rates does not grow too quickly as Δt is decreased. They also explain how to calibrate the model parameters to arbitrary term structures. In the resulting lattice, the procedure

for pricing interest rate contingent claims is formally the same as outlined above, but the analytical results for bond prices and European option prices are replaced by sums over probability-weighted values for discrete short rate outcomes. Hull and White demonstrate convergence of their routine to continuous-time continuous-state analytical results as Δt goes to zero.

We are concerned with measuring the error in the lattice approach for a given number of time steps. There are two sources of error: first, assuming discrete timing; second, imposing discrete states. To measure these errors separately, we extend the discrete-time continuous-state model above to allow for time-varying h to fit arbitrary term structures. The method for deriving analytical results is the same as followed above, but the notation becomes more complex to keep track of the dating of the mean reversion level. The difference between prices found using the two analytical results, one in discrete time and the other in continuous time, measures the error associated with discrete timing. The difference between prices found using the two discrete-time models, the analytical solution and the Hull and White approximation, measures the error from using discrete states.

In figure 1, a comparison across different strike prices shows that the use of a discrete-state process causes errors of significant size and variation. In contrast, the errors due to the use of discrete timing are nearly constant across different strike prices. While the trinomial model can result in smaller errors than the analytical solution, one cannot be sure of the size of the error that may arise for a given option. This suggests that eliminating the reliance on the lattice can improve the performance of approximation methods.

3. Approximating the value of American options

For an American option we need to accommodate the potential for early exercise. If the payoff to immediate exercise exceeds the value of holding on to the option, the option holder will exercise. If exercise is allowed only at discrete dates, the value of the option will be

$$\mathrm{C}(r_t, t) = \max\left[\mathrm{B}(r_t, t, t + m\Delta t) - K, \mathrm{B}(r_t, t, t + \Delta t)\mathrm{E}\left[\mathrm{C}(r_{t+\Delta t}, t + \Delta t)|r_t\right]\right] \quad (11)$$

where the dependence of bond prices on the contemporaneous short rate is noted explicitly. For a call option, the early exercise boundary is the largest short rate that satisfies $\mathrm{C}(r_t, t) =$

$\mathrm{B}(r_t, t, t + m\Delta t) - K$. If r^* denotes the exercise boundary at time t, then for short rates below r^* the option will be immediately exercised. The expectation term in (11) then has two components, one conditional on the short rate moving to a value that leads to early exercise at time $t + \Delta t$, and the other conditional on it staying above the exercise boundary.

The integral representation of the expectation term is

$$\mathrm{E}[\mathrm{C}(r_{t+\Delta t}, t + \Delta t)|r_t] = \int_{-\infty}^{r^*_{t+\Delta t}} \left(\mathrm{B}(r_{t+\Delta t}, t + \Delta t, t + m\Delta t) - K\right) f(r_{t+\Delta t}|r_t) dr_{t+\Delta t}$$

$$+ \int_{r^*_{t+\Delta t}}^{\infty} \mathrm{C}(r_{t+\Delta t}, t + \Delta t) f(r_{t+\Delta t}|r_t) dr_{t+\Delta t} \quad (12)$$

where $f(r_{t+\Delta t}|r_t)$ is the normal density function reflecting the conditional distribution of $r_{t+\Delta t}$ given r_t as specified in (1). The endogenous nature of the early exercise decision combined with the recursive structure of the pricing relation defeats attempts to find a closed-form solution. Our approximation will be done in two stages: first, using numerical integration to evaluate the time t expectation given the time $t + \Delta t$ pricing function, and then using function approximation to find a continuous function that closely fits those values.

Quadrature methods approximate an integral by evaluating the integrand at a set of points, called abscissas, and then constructing a weighted sum. A general statement of the method is

$$\int_a^b W(x)g(x)dx \approx \sum_{i=1}^{q} g(x_i)w_i \quad (13)$$

where $g(x)$ is the function to be integrated over the interval (a, b) with respect to the weighting function, $W(x)$. Gaussian quadrature rules choose the q abscissas, x_i, and the q weights, w_i, to give exact results when the integrand is $W(x)$ times a polynomial of order up to $2q - 1$ over the integration interval. If the integrand can be closely approximated by a polynomial times $W(x)$, then a small number of points are enough to achieve high accuracy. Various Gaussian quadrature formulas have been derived and tabulated for specific weighting functions and intervals.

Gauss-Hermite quadrature is for $W(x) = e^{-x^2}$ and $(a, b) = (-\infty, \infty)$. While the weighting function corresponds to the normal density in (12), the interval would cross the exercise boundary. In a discrete-time setting there is a discontinuity in the first derivative of option's payoff function at the exercise boundary, therefore, the function is not approximated

6

very well by a polynomial over the entire range of r_t and that substantially reduces the accuracy of Gauss-Hermite quadrature. Gauss-Laguerre quadrature is for $W(x) = x^\alpha e^{-x}$ and $(a, b) = (0, \infty)$. The interval corresponds to the half-finite intervals in (12), but the weighting function does not.

Testing the approximation on a one-period option showed good performance for Gauss-Laguerre quadrature for short rates near the exercise boundary. However, for larger r_t there is small chance of getting close to $r_{t+\Delta t}^*$ at the next date, yet the abscissas would still be concentrated there. For our problem, it turns out to be better to directly focus on a finite interval around the current short rate by using Gauss-Legendre quadrature, for which $W(x) = 1$ and $(a, b) = (-1, 1)$. In terms of the formal statement of the quadrature method, our integrand, $g(x)$, is the probability-density-weighted value of the option at the next date. While this choice worked well for our problem, in other situations a different numerical integration routine may be appropriate. Our procedure is not tied to a particular quadrature rule, but rather allows use of any rule that proves efficient. It also gives the means to test alternative rules.

Abscissas and weights for Gauss-Legendre quadrature are tabulated in Abramovitz and Stegun (1964) for the interval $[-1, 1]$ and can be applied to an arbitrary interval $[r_1, r_2]$ by scaling the abscissas to $z_i = (x_i(r_2 - r_1) + r_2 + r_1)/2$ and the weights to $y_i = w_i(r_2 - r_1)/2$.

We set the interval to cover the most probable short-rate outcomes while resticting it to the region in which the option value has a continuous first derivative. A range of six standard deviations around the mean short rate, truncated at the exercise boundary if necessary, proved to work well. Specifically, conditional on exercise at time $t + \Delta t$, the interval is given by $r_1 = \min(\mathrm{E}[r_{t+\Delta t}|r_t] - 6\sigma\sqrt{\Delta t}, r_{t+\Delta t}^*)$ and $r_2 = r_{t+\Delta t}^*$. Conditional on no exercise, the interval is $r_1 = \max(r_{t+\Delta t}^*, \mathrm{E}[r_{t+\Delta t}|r_t] - 6\sigma\sqrt{\Delta t})$ and $r_2 = \mathrm{E}[r_{t+\Delta t}|r_t] + 6\sigma\sqrt{\Delta t}$. The first interval captures the value associated with the possible exercise of the option at the next date; the second captures the value associated with holding on to the option for possible exercise at subsequent dates.

Gaussian quadrature can be used for multidimensional integration; however, its efficiency relies, in part, on using unequally spaced abscissas. Its recursive application to handle the multiple time steps in the option-pricing problem leads to exponential growth in the number

of evaluations. If a bond option has one period to maturity, then a single application of quadrature would suffice. The maturity date values at the abscissas, weighted by the normal density function, would be evaluated q times, multiplied by the corresponding q weights, and the discounted sum would approximate the option price. If there are two periods to maturity, then roughly q^2 calculations are made to get a single current price. The quadrature approximation of the two-period option value needs q one-period option prices, and each of those requires q evaluations of the maturity date value. In general, q^m evaluations are made to get the current price of an option with m periods to maturity.

The exponential increase in computational effort can be avoided by using equally spaced points and a small q while keeping the recursive structure. This describes the lattice method and partly explains its popularity. It may be inefficient for simple problems, but it is feasible for more complicated ones. Another way to conserve on effort is to avoid the nested application of quadrature. An assumption of Gaussian quadrature is that the pricing function can be approximated by a polynomial function. This suggests fitting the function with a polynomial approximation. The information about the pricing function gained by a single application of Gaussian quadrature would be summarized in the coefficents of the polynomial. The recursion described above, in which quadrature relies on itself to produce evaluations of the integrand at abscissas, can be replaced by iteration, a sequence of alternating one-dimensional quadrature and function approximation.

To avoid recursion in pricing a two-period option, we use a two step procedure. First, we estimate a p-term polynomial to fit one-period option prices, based on q-point quadrature at a set of p short rates; next, we apply q-point quadrature to get the two-period price, but use the polynomial to get the q prices for one-period options. By doing it p times, we could estimate the function relating two-period option prices to the current short rate and approximate the price for any r_t.

In general, to produce an approximating function for an m-period option there are mqp^2 evaluations: at each of the m time steps we need to get p prices that each use q-point quadrature built on evaluations of a p-term polynomial. The linear increase in the number of operations with m would be an important feature for longer term options or for smaller Δt. Whether the method works well in practice depends on having efficient routines for

one-dimensional numerical integration and for function approximation.

An effective procedure for polynomial approximation uses as basis functions the set of Chebyshev polynomials, $\mathbf{T}_j(x) = \cos(j \arccos x)$, where $j = 0, 1, \ldots, p - 1$ when a p-term polynomial is used to approximate a given function, and $x \in [-1, 1]$. Coefficients in the approximating polynomial, c_j, are chosen so that the target function is matched exactly at the p zeros of $\mathbf{T}_p(x)$, which occur at the points $x_j = \cos(\pi(j - 1/2)/p), j = 1, 2, \ldots, p$. The basis functions are orthogonal at these points, so a simple recursion produces the necessary coefficients. Press, et al. (1992) give details of the routines, and Fox and Parker (1968) provide an extensive discussion. Chebyshev approximation has an advantage over other function approximation methods when the target function is smooth because it produces a close approximation with moderate p. This increases its efficiency by requiring fewer evaluations of the target function.

We must set a finite interval to determine coefficients in a polynomial that approximates the option value as a function of the short interest rate. Below the exercise boundary that value is known and given by the immediate exercise value. Above the exercise boundary the value in theory is always positive, but practically declines to zero for short rates that are large enough to push the option far out of the money. We will fit a polynomial in the finite interval $[r^*, r^{**}]$, where r^{**} denotes a short rate high enough that the value of the option is very small, say 0.01. Given the interval and the degree of the approximating polynomial, the zeros of $\mathbf{T}_p(x)$ are scaled from $[-1, 1]$ to give a set of short rates at which the option is valued. Coefficients are determined so that the polynomial matches the option value at those points. This polynomial will be used in place of quadrature to deliver prices for the option at other short rates. To reduce percent approximation errors and ensure positive values at short rates above r^{**}, the approximation is done in logs, fitting log option values to log short rates.

Using these two components the approximation starts at the date just before the option expires and then moves back in steps to the current date. Given a previous solution, q-point quadrature is used to price options at specific short rates. At each step, the exercise boundary is found by searching for the short rate at which the option holder is indifferent to immediate exercise, using quadrature to evaluate option prices. The upper boundary

for the approximation interval is found by searching for the short rate at which the option has a suitably small value. These set the interval for polynomial approximation, and the p coefficients are found from p prices in that interval, again evaluated using quadrature. With those coefficients representing the option's value function, another step back can be taken. The product of the routine is a sequence of approximations, one for each maturity date, that can be used to value options for any short rate. Hedge ratios – specifically, the sensitivity of the option price to changes in the short rate – can be evaluated directly by taking the derivative of the approximating polynomial.

4. Performance of the method

To illustrate the properties of the approximation method, we computed prices for an American option on a coupon bond for various specifications of the option and the underlying interest rate process along with a variety of parameterizations of the method. There are three basic parameters in the routine: Δt, the length of time steps in the approximating discrete time model; q, the number of quadrature points; and p, the number of terms in the approximating polynomial. Preliminary testing revealed that $p > 8$ led to substantial extrapolation errors for some option and interest rate combinations. This could be handled by using a second low-order polynomial to approximate the very low option values associated with high interest rates. We decided instead to use $p \leq 8$ since there were no apparent effects of extrapolation errors and no substantial impact of adopting the more complicated approximation. We also found that as Δt decreases and the number of time steps increases it is necessary to increase q and p to avoid propagation of approximation errors and variation in the prices.

An effective method should be capable of using small time steps to closely match the underlying continuous-time model. In table 1, approximate prices are shown as the number of time steps, denoted $n = 1/\Delta t$, ranges from 16 per year to 1,024 per year, with $p = 8$, and $q = 20$ for lower n and 40 for $n = 1,024$. The method displays smooth convergence at a variety of current short rates. Doubling n from 512 to 1,024 results in price changes only in the fourth decimal place. In the following, we will take the prices with $n = 1,024$ as

benchmark values.

For efficient approximation, a method should deliver adequate accuracy in a reasonable amount of time. The benchmark prices are almost matched by solutions using $n = 64$, with discrepancies below 0.01. The smooth convergence of the approximate prices also allows use of Richardson extrapolation. This was introduced for American stock option pricing by Geske and Johnson (1984) and uses information from solutions with small n to estimate the limiting values. Here we found a simple variant to be effective:

$$C_{\mathrm{RX}} = 2C_{2n} - C_n \tag{14}$$

where C_{RX} denotes the extrapolated price and C_j is a price found with j time steps. For $n = 32$, the extrapolated prices nearly duplicate the benchmark prices, matching to the third decimal place.

Either the lower n or the extrapolation would serve as an adequate substitute for the $n = 1,024$ solutions, and they are much easier to calculate. Table 1 includes computation times as measured on a 233Mhz Pentium computer. To get a solution using $n = 1,024$ takes about 30 seconds, which would be too long for some purposes. Computation time is linear in the number of time steps, and an approximation with "penny accuracy" would take about 2 seconds ($n = 64$). Extrapolation based on $n = 32$ and 64 improves the accuracy substantially while only adding another second to the calculations.

Even the shorter times are long compared with the time required to calculate European option prices from their analytical solution. However, the effort is necessary because of the important effects of the early exercise feature. These are illustrated by direct comparison of European and American option prices in figure 2. The difference in price is noticeable even at low option values. As time to maturity increases, the gap widens substantially. The early exercise provision anchors the American option value at the immediate exercise value. In contrast, a European option falls much below that value and is much less responsive to a change in the current short rate. And the European option can decline in value as its maturity lengthens. The differences in price and sensitivity are large enough to warrant the additional effort required to approximate American option prices.

The approximation technique developed here has important advantages over lattice meth-

ods. The number of time steps can be assigned separately from the number of option valuations required at each step. This results in more precise evaluations of option prices for a given time interval, gives a close match to the discrete time solution, and explains its smooth convergence. In turn, this allows effective use of extrapolation to conserve computation time. And, with computation time linear in the number of time steps, it is feasible to use a large number of time steps to improve accuracy.

In contrast, lattice methods require more time steps to get the greater number of function evaluations needed for increased accuracy. The computational effort grows exponentially with the number of time steps, and there is no way to separately reduce the errors in approximating the option pricing function from those which arise from using discrete time steps. The changes in the set of short rates at which the option is evaluated as Δt decreases also cause uneven convergence of approximate prices. This renders extrapolation methods useless. These shortcomings can overwhelm the strength of the lattice method, repetition of simple calculations, which otherwise give it an advantage over other techniques.

The results of a direct comparison to the Hull and White model, for given number of time steps, are shown in table 2. We measured deviations from benchmark prices (continuous-state with $n = 1,024$ time steps) for both routines. For a low number of time steps, $n = 16$, the lattice method has similar accuracy, while for $n = 32$, it produces a less exact match for our benchmark prices. Doubling n to 64 illustrates the uneven convergence of the discrete-state approximation. Extrapolation using continuous-state prices nearly duplicates the benchmark prices.

The tabulated prices do not completely reflect the relative accuracy of the methods. Figure 3 shows that there is substantial variation in errors for the lattice method across a range of interest rates. This variation should not be surprising given the previous results shown for European options. Because of the jagged profile of approximation errors in the lattice method, measures of interest rate sensitivity are less precise. The benefits of extrapolation for our method are also shown: errors are noticeable only near the exercise boundary.

While the lattice approach converges more slowly as the number of time steps increases, it also uses very simple operations. The greater precision of the continuous-state method comes at the cost of more extensive calculations at each time step. To control for these differences,

we compared solutions that use roughly the same number of computations. Denote the time to maturity of the option in years by M. The order of the number of operations for the discrete-state method is $(M/\Delta t)^2$, the number of nodes in the lattice. For the continuous-state solution, each time step has p nodes. From each node, q possible future values are found from a p term polynomial giving a total number of calculations on the order of $(M/\Delta t)qp^2$.

Table 3 contains information on the efficiency of each method, with measures of accuracy for comparable orders of number of calculations. For each option, we used a lattice method with 64 and 128 times steps per year to give two levels of precision and picked parameters for continuous-state solutions of matching order. For the six-month option, the discrete-state errors are less than a penny on average for $n = 64$. The continuous-state solutions do poorly by comparison. To accommodate the more involved computations per time step, a much smaller number of time steps must be used. The lower accuracy is especially evident in the maximum error. Increasing to $n = 128$, the two are closer. Errors for the lattice method are halved with the quadrupling of effort. They fall by much more for the continuous-state method, so that in terms of percent errors it has higher accuracy. Notice that the number of time steps for the continuous-state solution could quadruple if q and p were fixed. But it is better to also increase q and p as n goes up so that errors do not accumulate with the larger number of iterations.

The results for longer term options illustrate the general conclusion. The discrete-state method has an advantage when the number of time steps is lower: for shorter term options or for lower precision ($n = 64$). If the option has a longer time to maturity, or if more accurate pricing is needed, then the continuous-state approach generally does better. Comparing the two on the basis of differences from the benchmark price, the lattice method has lower root-mean-squared error and maximum absolute error for most of the cases presented. The extrapolated continuous-state solutions do poorly by the latter metric because of errors near the exercise boundary. Solutions with a smaller number of time steps overestimate the exercise boundary. This interrupts the smooth convergence of prices as n increases and reduces the accuracy of the extrapolation. In percent terms, this shortcoming has less of an impact and the continuous-state method has the same or lower errors compared to the discrete-state approach, except for the six-month option estimated with lower precision.

Overall, the continuous-state method is an effective substitute for a discrete-state lattice.

5. Extensions

This paper demonstrates an efficient alternative to lattice methods for valuing interest rate contingent claims. It has a flexible structure that allows it to be applied in a variety of situations. For example, Black, Derman, and Toy (1990) build an interest rate model based on the assumption of a lognormal distribution for short rates. The model is convenient because negative interest rates are ruled out. Unfortunately, there are no analytical results available for even the basic prices of zero coupon bonds. The method described here can be used to approximate the function relating yields to the short rate and from that build up valuation for more complicated instruments.

Even in the model with normally distributed shocks a wide variety of contingent claims must be priced using approximation. Barrier options and Bermuda options are two examples. The lattice method can be adapted to handle these securities. But its efficiency is further diminished by the discrete nature of the payoffs. In contrast, those options are easier to value than American options using the above technique. The discrete knock out/in of a barrier option places an exogenous limit on the integration and estimation intervals that allows faster and more accurate approximation. The limited number of possible exercise dates in a Bermuda option would also simplify application. In principle, the new model can also be applied to multifactor term structure models.

Acknowledgments

References

Abramowitz, M., Stegun, I. A., eds., 1964. *Handbook of Mathematical Functions*, 2nd Ed. National Bureau of Standards Applied Mathematics Series 53, U. S. Government Printing Office.

Black, F., Derman, E., Toy, W., 1990. "A one-factor model of interest rates and its application to Treasury bond options." *Financial Analysts Journal* 46, 33-39.

Cox, J., Ross, S., Rubinstein, M., 1979. "Option pricing: a simplified approach." *Journal of Financial Economics* 7, 229-264.

Derman, E., Kani, I., Ergener, D., Bardhan, I., 1995. "Enhanced numerical methods for options with barriers." *Financial Analysts Journal* 51, 65-74.

Fox, L., Parker, I. B., 1968. *Chebyshev Polynomials in Numerical Analysis.* Oxford University Press, London.

Geske, R., Johnson, H. E., 1984. "The American put option valued analytically." *Journal of Finance* 34, 1511-1524.

Hull, J., White, A., 1993. "One-factor interest rate models and the valuation of interest rate derivative securities." *Journal of Financial and Quantitative Analysis* 28, 235-254.

Hull, J., White, A., 1994. "Numerical procedures for implementing term structure models I: single-factor models." *The Journal of Derivatives* 2, 7-16.

Jamshidian, F., 1989. "An exact bond option formula." *Journal of Finance* 44, 205-209.

Judd, K., 1992. "Projection methods for solving aggregate growth models." *Journal of Economic Theory* 58, 410-452.

Press, W., Teukolsky, S., Vetterling, W., Flannery, B., 1992. *Numerical Recipes in C*, 2nd ed. Cambridge University Press, New York.

Rubinstein, M., 1976. "The valuation of uncertain income streams and the pricing of options." *Bell Journal of Economics* 7, 407-425.

Sullivan, M. A., 2000. "Valuing American put options using Gaussian quadrature." *Review of Financial Studies* 13, 75-94.

Vasicek, O. A., 1977. "An equilibrium characterization of the term structure." *Journal of Financial Economics* 5, 177-188.

Table 1: Convergence and computation time.

Table entries are discrete-time continuous-state approximations of American option prices for a one-year option struck at par on a five-year bond paying annual coupons of 10 on a par value of 100. Extrapolations combine solutions according to $2C_{2n} - C_n$. Time measures the seconds to compute a solution with the indicated number of time steps per year, n, using programs written in the Gauss programming language and run on a 233Mhz Pentium computer. The interest rate parameters are $a = 0.05, h = 0.10$, and $\sigma = 0.02$.

			Short Interest Rate			
n	0.08	0.09	0.10	0.11	0.12	Time
16	6.3960	3.6779	1.9091	0.8711	0.3418	0.44
32	6.4107	3.6843	1.9117	0.8717	0.3417	0.88
64	6.4165	3.6875	1.9128	0.8720	0.3415	1.70
128	6.4191	3.6888	1.9133	0.8720	0.3414	3.52
256	6.4211	3.6896	1.9136	0.8720	0.3414	7.03
512	6.4221	3.6901	1.9137	0.8721	0.3414	14.28
1,024	6.4225	3.6903	1.9138	0.8721	0.3414	28.89
extrapolated						
16 and 32	6.4254	3.6907	1.9142	0.8724	0.3417	1.32
32 and 64	6.4224	3.6907	1.9140	0.8722	0.3413	2.58

Table 2: Approximation errors for equal number of time steps.
Table entries are percent differences from benchmark American option prices (continuous-state model with $n = 1,024$) for various number of time steps per year, n. Extrapolations combine continuous-state solutions as $2C_{2n} - C_n$. The option has one year to maturity and is struck at par on a five-year bond paying annual coupons of 10 on a par value of 100. The interest rate parameters are $a = 0.05, h = 0.10$, and $\sigma = 0.02$.

		Short Interest Rate			
n	0.08	0.09	0.10	0.11	0.12
discrete-state					
16	-0.514	-0.065	0.580	1.413	2.088
32	0.033	-0.460	0.199	-0.046	0.208
64	-0.117	-0.140	-0.421	-0.182	0.467
continuous-state					
16	-0.413	-0.335	-0.248	-0.120	0.106
32	-0.184	-0.162	-0.114	-0.042	0.096
64	-0.093	-0.075	-0.052	-0.019	0.035
extrapolated					
32 and 64	-0.002	0.012	0.009	0.005	-0.027

Table 3: Approximation errors for equal computational effort.
The number of time steps per year (n_{DS}) for discrete-state solutions (DS) and the number of time steps per year, quadrature points, and polynomial terms (n_{CS}, q, p) for continuous-state solutions (CS) roughly equate the number of operations for each method. Extrapolations (RX) combine continuous-state solutions as $2C_{2n_{CS}/3} - C_{n_{CS}/3}$. Errors are measured against a continuous-state solution with $n_{CS} = 1,024, q = 40$, and $p = 8$ for 200 interest rates equally spaced between the exercise boundary and a rate at which the option value is 0.50. The option has M years to maturity and is struck at par on a five-year bond paying annual coupons of 10 on a par value of 100. The interest rate parameters are $a = 0.05, h = 0.10$, and $\sigma = 0.02$.

| M	n_{DS}	n_{CS}, q, p	Error			Percent Error		
			DS	CS	RX	DS	CS	RX
root mean square								
0.5	64	$18, 8, 4$	0.0063	0.0304	0.0299	0.37	2.01	1.59
	128	$30, 12, 5$	0.0034	0.0083	0.0059	0.21	0.19	0.18
1.0	64	$15, 12, 5$	0.0057	0.0161	0.0126	0.23	0.31	0.22
	128	$36, 13, 6$	0.0028	0.0070	0.0041	0.11	0.13	0.06
2.0	64	$18, 13, 6$	0.0051	0.0138	0.0093	0.14	0.24	0.12
	128	$60, 15, 6$	0.0025	0.0047	0.0023	0.07	0.07	0.04
maximum absolute value								
0.5	64	$18, 8, 4$	0.0178	0.0927	0.0624	1.47	4.66	2.92
	128	$30, 12, 5$	0.0094	0.0210	0.0329	0.92	0.33	0.52
1.0	64	$15, 12, 5$	0.0234	0.0466	0.0696	0.88	0.62	0.96
	128	$36, 13, 6$	0.0116	0.0190	0.0283	0.49	0.24	0.37
2.0	64	$18, 13, 6$	0.0249	0.0407	0.0628	0.41	0.48	0.76
	128	$60, 15, 6$	0.0128	0.0118	0.0171	0.23	0.15	0.19
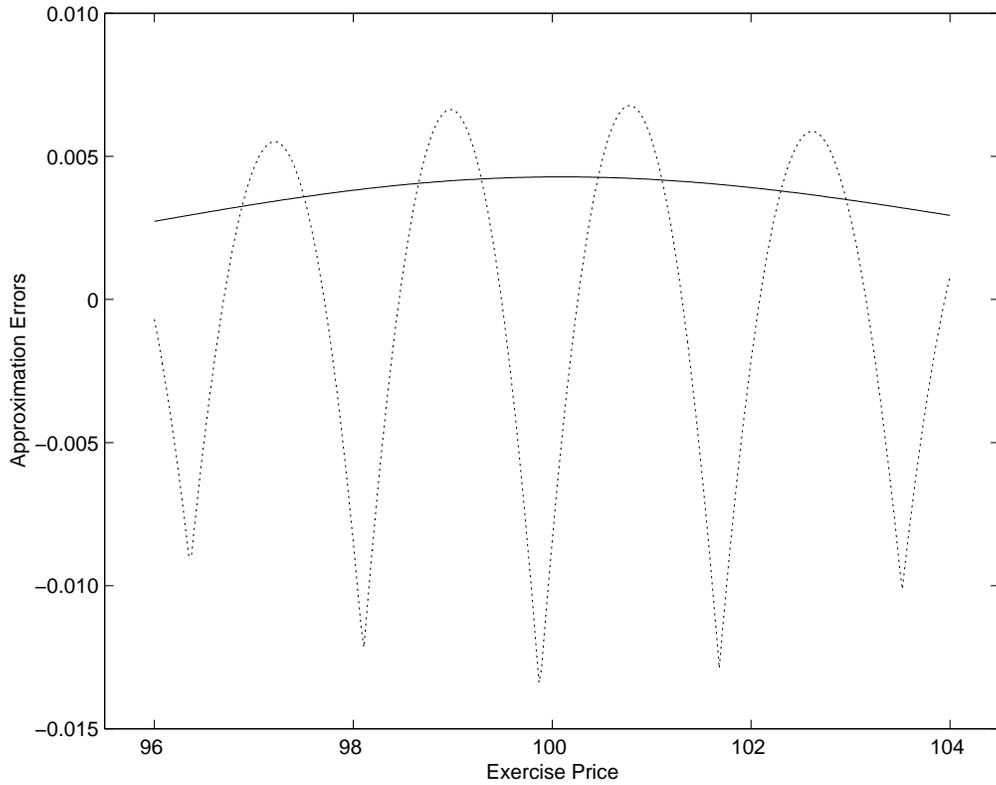
Figure 1: Approximation errors for European option prices.
The dotted line shows errors caused by using discrete states: the difference
between prices from the Hull and White (1994) lattice routine with 100 time
steps and the analytical solution for that time spacing. The solid line shows
errors caused by using discrete time: the difference between an analytical
solution for a model with 100 time steps and the continuous-time solution.
Example parameters are from Hull and White (1993): $a = 0.1, \sigma = 0.014$,
and h_t chosen to fit a term structure with yields on discount bonds that rise
linearly from 9.5 percent to 11 percent at three years, and from 11 percent to
11.5 percent out to five years. The one-year option is on a five-year discount
bond. The exercise price is given as a percent of the forward bond price.
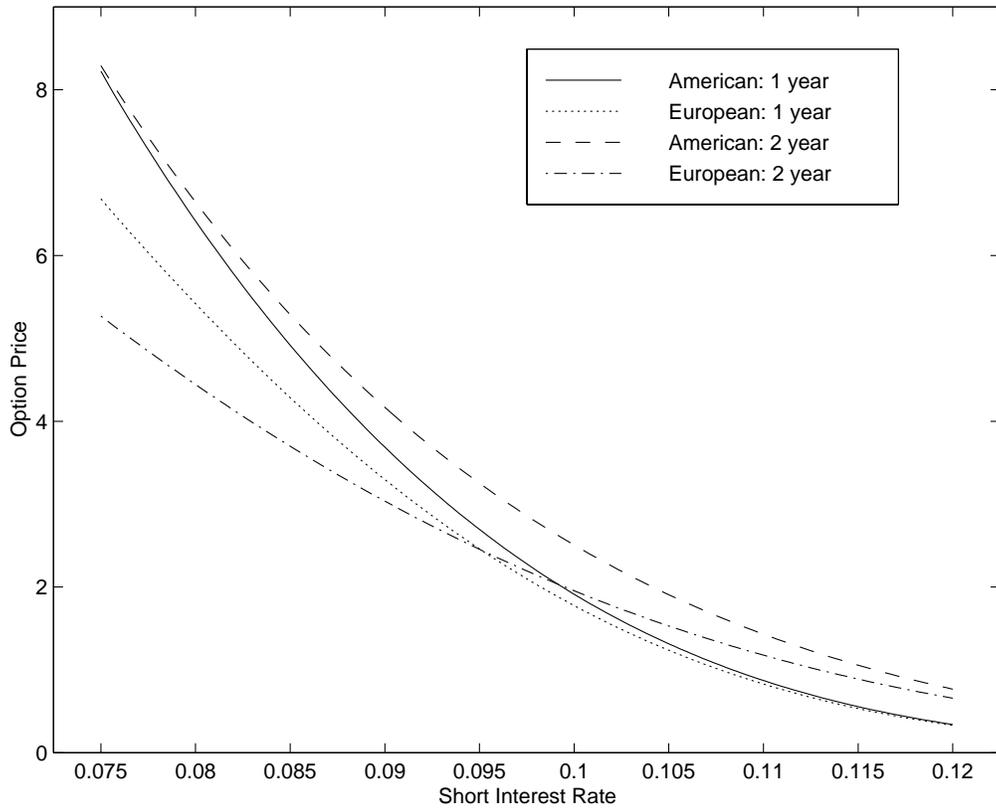
Figure 2: American vs. European option values.
The prices are for options struck at par on a five-year bond paying annual
coupons of 10 on a par value of 100. European option prices are calculated
using a continuous-time continuous-state model. American option prices are
calculated using a discrete-time continuous-state model with $n = 64$. The
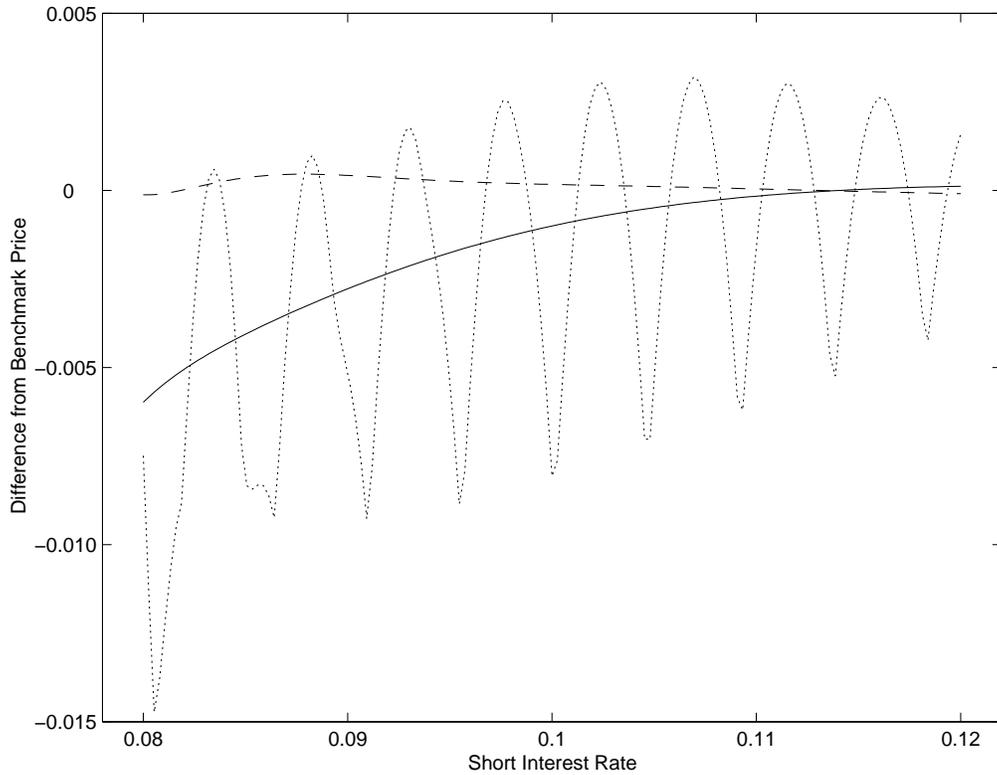interest rate parameters are $a = 0.05, h = 0.10$, and $\sigma = 0.02$.

Figure 3: Approximation errors for American option prices.
Differences from benchmark prices (discrete-time continuous-state with $n = 1,024$) are shown by the solid line for a continuous-state approximation and by the dotted line for a discrete-state approximation, both using $n = 64$. The dashed line shows differences for estimates by Richardson extrapolation based on continuous-state solutions with $n = 32$ and 64 time steps. The prices are for a one-year option struck at par on a five-year bond paying annual coupons of 10 on a par value of 100. The interest rate parameters are $a = 0.05, h = 0.10$, and $\sigma = 0.02$.