

Pricing Options under Generalized GARCH and Stochastic Volatility Processes

PETER RITCHKEN and ROB TREVOR*

ABSTRACT

In this paper, we develop an efficient lattice algorithm to price European and American options under discrete time GARCH processes. We show that this algorithm is easily extended to price options under generalized GARCH processes, with many of the existing stochastic volatility bivariate diffusion models appearing as limiting cases. We establish one unifying algorithm that can price options under almost all existing GARCH specifications as well as under a large family of bivariate diffusions in which volatility follows its own, perhaps correlated, process.

THROUGH AN EQUILIBRIUM ARGUMENT Duan (1995) shows that options can be priced when the dynamics for the price of the underlying instrument follows a General Autoregressive Conditionally Heteroskedastic (GARCH) process. The theory of pricing under such processes is now well understood; however, the design of efficient numerical procedures for pricing them is lacking. Most applications resort to large sample simulation methods with a variety of variance reduction techniques. The complexity of pricing arises from the massive path dependence inherent in GARCH models. This path dependence causes typical lattice-based procedures to grow exponentially in the number of time increments. The lack of efficient numerical schemes hinders empirical tests among the wide array of competing GARCH models. Most tests of GARCH models limit themselves to the use of high frequency data on spot assets, with little attention placed on the information content of options.¹

*Ritchken is at the Weatherhead School of Management, Case Western Reserve University, Cleveland; Trevor is at the Centre for Studies in Money, Banking and Finance, Macquarie University, Sydney, Australia. The authors gratefully acknowledge support from the Australian Research Council, Macquarie Investment Management Limited, and Axiom Funds Management Corporation under Collaborative Research Grant C595301128. We have benefited from the comments of René Stulz, the referee, Jin-Chuan Duan, Ivilina Popova, L. Sankarasubramanian, and participants at the 4th Annual CAP Workshop on Mathematical Finance at Columbia University, the 7th Annual Derivatives Securities Conference at Queen's University, the 10th Annual Australian Finance and Banking Conference, the 1997 Quantitative Methods in Finance Conference in Australia, the 1997 INFORMS in Dallas, the 1998 American Finance Association Meetings in Chicago, and at seminars at Macquarie Bank, Macquarie University, Princeton University, the University of Auckland, the University of New South Wales, and the University of Western Australia.

¹ See Bollerslev, Chou, and Kroner (1992) and Bollerslev, Engle, and Nelson (1994) for surveys of the literature.

Indeed, to our knowledge, other than Duan (1996a), there have been no approaches that imply out GARCH parameters using a theoretically justified GARCH option pricing model.² Moreover, pricing American options under GARCH processes using simulation methods has not been well studied.

GARCH type processes can be linked to bivariate diffusion models and vice versa. Nelson (1990), for example, shows that certain GARCH processes can be used to approximate some bivariate diffusion models. More recently, Duan (1996b, 1997) generalizes these results and brings the largely separate GARCH and bivariate diffusion literatures together. Indeed, Duan (1997, 1996b) shows that most of the existing bivariate diffusion models that have been used to model asset returns and volatility can be represented as limits of a family of GARCH models. As a result, even if one prefers modeling prices and volatilities by a bivariate process, there may be advantages in estimating parameters of the model using GARCH techniques. Conversely, if one prefers the GARCH paradigm, there may be some advantages in implementing option models using the existing vast literature on numerical procedures for pricing under bivariate diffusions.

The purpose of this article is to do just the reverse. We establish an efficient lattice algorithm for pricing European and American options under discrete time GARCH processes. This algorithm is then extended to price options under generalized GARCH processes, which contain many of the existing bivariate diffusion models as limiting cases. By setting up an efficient computation scheme for the generalized GARCH model, we are able to solve option problems not only for GARCH processes, but also for many bivariate diffusions. For example, by suitably curtailing the parameters of generalized GARCH processes, we can obtain European and American option prices under the stochastic volatility models of Hull and White (1987), Scott (1987), Wiggins (1987), Stein and Stein (1991), and Heston (1993).

Providing simple computational schemes for pricing options under GARCH and stochastic volatility should be of direct interest to empiricists who are interested in comparing alternative GARCH structures for the asset return dynamics and would like to incorporate information implied by the existing set of option prices. These models should also be of interest to researchers who want to compare alternative stochastic volatility bivariate diffusion models. The lattice algorithm presented here should allow them to experiment with many alternative structures in a single modeling paradigm. Finally, this article contributes to the literature by providing an efficient way to compute values of American options under GARCH processes, as well as stochastic volatility bivariate diffusions.³

² Duan (1996a) follows Engle and Mustafa (1992) in using simulation methods to imply out GARCH parameters. However, Engle and Mustafa assume that agents are risk neutral.

³ Duan and Simonato (1997) concurrently develop a Markov chain approximation to a GARCH process which they implement using sparse matrix techniques to price American options under GARCH.

The paper proceeds as follows. In Section I we review the basic GARCH option pricing model of Duan (1995). In Section II we describe the problems caused by the massive path dependence inherent in GARCH, and provide the main result that permits the design of an efficient lattice-based approximation to the GARCH process. In Section III we discuss how option contracts can be priced using this lattice. In Section IV we illustrate the convergence properties of the algorithm and show how it can be readily modified to price claims under a variety of GARCH processes. Section V extends the lattice algorithm to approximate generalized GARCH models which converge to bivariate diffusions that include processes used by Hull and White (1987), Heston (1993), Stein and Stein (1991), Scott (1987), Wiggins (1987), and others. This section shows that option prices under many stochastic volatility bivariate diffusions can be efficiently computed using the approximating GARCH model. Finally, in Section VI we provide a summary and outline the types of empirical issues that now can be more readily addressed.

I. The GARCH Option Model

Let S_t be the asset price at date t , and h_t be the conditional variance, given information at date t , of the logarithmic return over the period $[t, t + 1]$, which (without loss of generality) we call a “day”. The dynamics of prices are assumed to follow the process

$$\ln\left(\frac{S_{t+1}}{S_t}\right) = r_f + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}v_{t+1}, \quad (1)$$

$$h_{t+1} = \beta_0 + \beta_1 h_t + \beta_2 h_t(v_{t+1} - c)^2, \quad (2)$$

where v_{t+1} , conditional on information at time t , is a standard normal random variable. The riskless rate of return over the period is r_f . The unit risk premium for the asset is λ .

The particular structure imposed in equation (2) is the nonlinear asymmetric GARCH (NGARCH) model, that has been studied by Engle and Ng (1993) and Duan (1995). The nonnegative parameter c captures the negative correlation between return and volatility innovations that is frequently observed in equity markets. (The model simplifies to the popular GARCH model of Bollerslev (1986) when this correlation is absent.) To ensure that the conditional volatility stays positive β_0 , β_1 , and β_2 should be nonnegative.

Duan (1995) derives a pricing mechanism for derivative securities when the price of the underlying security follows the above process. In particular, under suitable preference restrictions, he establishes a local risk-neutralized

probability measure under which option prices can be computed as simple discounted expected values.⁴ The process under the local risk-neutralized measure is

$$\ln\left(\frac{S_{t+1}}{S_t}\right) = \left(r_f - \frac{1}{2}h_t\right) + \sqrt{h_t}\epsilon_{t+1}, \quad (3)$$

$$h_{t+1} = \beta_0 + \beta_1 h_t + \beta_2 h_t (\epsilon_{t+1} - c^*)^2, \quad (4)$$

where ϵ_{t+1} , conditional on time t information, is a standard normal random variable with respect to the risk-neutralized measure, and $c^* = c + \lambda$. The above model has five unknown parameters, namely β_0 , β_1 , β_2 , c^* , and the initial variance h_0 .

Most implementations of GARCH option pricing models use Monte Carlo simulation to price European options.⁵ Typically, very large replications are required to obtain precise estimates, even when variance reduction techniques are used. Moreover, simulation procedures for American options are still quite tedious. In the next section we develop a simple lattice-based algorithm for pricing claims under the preceding GARCH process. The algorithm is extremely efficient and readily permits the pricing of American options.

II. Approximating the GARCH Process

Pricing options under GARCH processes using lattices has been difficult because of the inherent path dependence that leads to an exploding number of states. To illustrate this, assume that a simple Bernoulli sequence is used to proxy the standard normal random variable in equations (3) and (4). After one period there are two prices with two distinct updated variances. Because the variances in the up-node and down-node are different, there is no reason that the price resulting from an up and then down path equals that obtained from the down and then up path. As a result, after two periods there are four asset prices with four variances. Since paths do not reconnect, such an approximation leads to an exploding number of states for the two state variables.

The key to an efficient implementation is to design an algorithm that avoids an exponentially exploding number of states. Toward this goal, we begin by approximating the sequence of single period lognormal random variables in equation (3) by a sequence of discrete random variables. In par-

⁴ Kallsen and Taqqu (1998) derive the same result by a no-arbitrage condition in a continuous time version of the GARCH model.

⁵ See Duan (1995, 1996a), Amin and Ng (1993), and Engle and Rosenberg (1995), for example. The simulations in the latter study are based on the estimated dynamics of the spot process, assuming risk neutral agents.

ticular, assume the information set at date t is $\{S_t, h_t\}$ and let $y_t = \ln(S_t)$. Then, viewed from date t , y_{t+1} is a normal random variable with conditional moments

$$E_t[y_{t+1}] = y_t + r_f - \frac{1}{2}h_t, \quad (5)$$

$$\text{Var}_t[y_{t+1}] = h_t. \quad (6)$$

We establish a discrete state Markov chain approximation, $\{(y_t^a, h_t^a) | t = 0, 1, 2, \dots\}$, for the dynamics of the discrete time state variables that converges to the continuous state, discrete time, GARCH process, $\{(y_t, h_t) | t = 0, 1, 2, \dots\}$. In particular, we approximate the sequence of conditional normal random variables by a sequence of discrete random variables. Given this period's logarithmic price and conditional variance, the conditional normal distribution of the next period's logarithmic price is approximated by a discrete random variable that takes on $2n + 1$ values. There are n values larger than the current price, n values smaller than its current value, and a value that is unchanged. If, for example, $n = 1$, then the approximation consists of a trinomial random variable. The lattice we construct has the property that the conditional means and variances of one-period returns match the true means and variances given in equations (5) and (6), and the approximating sequence of discrete random variables converges to the true sequence of normal random variables as n increases.⁶

In the usual binomial approximation of a Wiener process the size of the local Bernoulli jumps equals the volatility over the time increment. Indeed, if the first two conditional moments on the lattice are to match the true moments, then the magnitude of the jumps is *uniquely* determined by the volatility. In GARCH processes the volatility changes through time and hence Bernoulli approximations require different jump sizes along the path. This leads to a nonrecombining tree in which adjacent logarithmic prices are not separated by a fixed constant. On the other hand, if a sequence of equally spaced trinomial (or, more generally, multinomial) variables is used, then the first two moments can be matched without uniquely locking in this space partition. This observation permits the same grid of points to be used to approximate a wide range of normal random variables that differ in means and variances. As we shall see, this fact is helpful in that it permits lattices to be built off a single grid of equidistant logarithmic prices. Hence, in what follows we always use a multinomial approximation to the normal with at least three points—that is, with $n \geq 1$. For the degenerate case where volatility is constant (i.e., $\beta_1 = \beta_2 = 0$) and $n = 1$, the lattice we construct

⁶ In what follows, variables that have superscripts of a denote variables on the lattice that approximate the same variables without superscripts. By convergence, we mean that over each period, as $n \rightarrow \infty$ the approximating conditional distributions converge to normal distributions.

collapses to the trinomial model of Kamrad and Ritchken (1991). Further, for constant volatility and $n > 1$ our approximating sequence consists of identical multinomial jumps over all periods.⁷

We first establish a grid of logarithmic prices to approximate the possible states of nature during the life of the option. Let γ be a fixed constant that determines the gap between adjacent logarithmic prices on this grid. Over each day we require a $2n + 1$ point approximation to a normal distribution with mean and variance given by equations (5) and (6). The gap between adjacent logarithmic prices is determined by γ and n . In particular, the gap is γ_n where

$$\gamma_n = \frac{\gamma}{\sqrt{n}}.$$

At date t assume that the logarithmic price and variance are given. The actual size of these $2n + 1$ jumps depends on the variance, but we restrict them to be integer multiples of γ_n . If the variance is “small” then discrete probability values can be found over the grid of $2n + 1$ prices surrounding the current price, such that the conditional mean and variance of these $2n + 1$ points match the values of the conditional normal distribution being approximated. However, if the variance is sufficiently large then it may not be possible to find valid probability values for the surrounding $2n + 1$ prices such that the first two moments of the approximating distribution match. In this case, we construct an approximating distribution that uses every second point above (below) the current value, as well as the unchanged value. Each of these points is separated by a gap of $2\gamma_n$. If these “double-sized” jumps do not permit the means and variances to be exactly matched while producing valid probabilities, then the approximating scheme uses points separated by larger γ_n multiples.

Let η be the smallest integer multiple that allows the mean and variance of next period’s logarithmic price to be matched to the true moments, at the same time ensuring that all the $2n + 1$ probability values are valid numbers in the interval $[0, 1]$. We refer to η as the jump parameter. The proposition below provides us with a method for determining η for any given volatility level. Clearly, η will depend on the space parameter, γ .

More formally, assume that at date t the logarithmic price is y_t^a and the conditional variance is h_t^a , where the superscript a denotes values derived from using $2n + 1$ point approximations to the conditional normal distributions along the lattice. Over the next period the logarithmic price moves to one of $2n + 1$ points and the variance is updated according to the magnitude of the move in the asset price,

$$y_{t+1}^a = y_t^a + j\eta\gamma_n, \quad (7)$$

$$h_{t+1}^a = \beta_0 + \beta_1 h_t^a + \beta_2 h_t^a [\epsilon_{t+1}^a - c^*]^2, \quad (8)$$

⁷ For related discussions on these constant volatility models see Boyle, Evnine, and Gibbs (1989).

where

$$\epsilon_{t+1}^a = \frac{j\eta\gamma_n - (r_f - h_t^a/2)}{\sqrt{h_t^a}} \quad (9)$$

and $j = 0, \pm 1, \pm 2, \dots, \pm n$. The jump parameter η is the smallest integer that ensures that the probabilities in equation (11) below are between 0 and 1. It is chosen such that

$$(\eta - 1) < \frac{\sqrt{h_t^a}}{\gamma} \leq \eta, \quad (10)$$

where γ is the fixed constant that in conjunction with n determines the logarithmic price grid. We choose $\gamma = \sqrt{h_0^a}$ for the calculations in this paper.⁸ In the following proposition we see that viewed from time t , ϵ_{t+1}^a is a discrete state random variable, with mean 0 and variance 1, that converges in distribution to a continuous state standard normal random variable as $n \rightarrow \infty$.

All that remains to complete the specification of the dynamics of the approximating process is the assignment of probabilities for the transitions. We do this by effectively splitting each period (day) into n subintervals of equal length, $1/n$. The variance is constant over each of these subintervals, but is updated at the end of the day. The price may be thought of as evolving over the course of the day as a sequence of n independent, constant volatility, normal processes that we approximate by a sequence of n trinomial distributions.⁹ Over each day, then, we have $2n + 1$ possible outcomes. The probability of each of the $2n + 1$ jumps of size $\eta\gamma_n$ for the whole day is the same as the probability for the corresponding paths along an n -step trinomial tree that lead to that node. The probability distribution for y_{t+1}^a conditional on y_t^a and h_t^a is then given by

$$\text{Prob}(y_{t+1}^a = y_t^a + j\eta\gamma_n) = P(j) \quad j = 0, \pm 1, \pm 2, \dots, \pm n,$$

where

$$P(j) = \sum_{j_u, j_m, j_d} \binom{n}{j_u j_m j_d} p_u^{j_u} p_m^{j_m} p_d^{j_d} \quad (11)$$

⁸ It is not advisable to have an extremely large value for γ . Kamrad and Ritchken (1991) explore the convergence rate of their trinomial model to the Black–Scholes equation as γ increases. Their recommendation is to choose $\gamma = \sqrt{3/2}\sigma$ where σ is the Black–Scholes volatility. At this setting, the probabilities of all three jumps are close to each other. If γ is far from this value (e.g., $\gamma = 10\sigma$) option prices still converge, although at a slower rate.

⁹ A similar concept is used by Kallsen and Taqqu (1998) who also assume that volatility remains constant within each day of their continuous time GARCH model.

with $j_u, j_m, j_d \geq 0$ such that $n = j_u + j_m + j_d$ and $j = j_u - j_d$. The expression in parentheses denotes the trinomial coefficient $n!/(j_u!j_m!j_d!)$, and the trinomial probabilities are

$$p_u = \frac{h_t^a}{2\eta^2\gamma_n^2} + \frac{(r_f - h_t^a/2)\sqrt{1/n}}{2\eta\gamma_n}, \quad (12)$$

$$p_m = 1 - \frac{h_t^a}{\eta^2\gamma_n^2}, \quad (13)$$

$$p_d = \frac{h_t^a}{2\eta^2\gamma_n^2} - \frac{(r_f - h_t^a/2)\sqrt{1/n}}{2\eta\gamma_n}. \quad (14)$$

Proposition 1 investigates the behavior of this discrete state process.

PROPOSITION 1: *The discrete state approximating process given in equations (7) through (14) has the following properties:*

1. *The distribution of the discrete state random variable y_{t+1}^a , given the state variables (y_t^a, h_t^a) at date t , has mean $r_f - h_t^a/2$ and variance h_t^a .*
2. *The discrete state random variable ϵ_{t+1}^a , given information at date t , has zero mean, unit variance, and converges to a continuous state standard normal distribution as $n \rightarrow \infty$. Equivalently, y_{t+1}^a , given information at date t , converges to a normal distribution as $n \rightarrow \infty$.*
3. *The process $\{(y_t^a, h_t^a) \text{ for } t = 0, 1, \dots\}$ in equations (7) through (9) converges to the GARCH process given in equations (3) and (4) as $n \rightarrow \infty$.*

The proof is given in the Appendix.

Proposition 1 provides the basis for our discrete state approximation to the continuous state GARCH process described by equations (3) and (4). It is best understood by considering an example.

Assume that over each day the underlying price process is GARCH as in equations (1) and (2). We approximate the risk-neutralized process (equations (3) and (4)) by a lattice of prices for the case $n = 1$. Over each day the approximating logarithmic prices on the lattice then represent consecutive drawings from trinomial distributions with means and variances that match the true means and variances of the GARCH process. Suppose that the current underlying price is $S_0 = 1000$, the risk free rate is $r_f = 0$, the risk premium is $\lambda = 0$, $\beta_0 = 0.000006575$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, $c = 0$, and the initial daily variance of $h_0 = 0.0001096$ corresponds to an annual volatility of 20 percent. Choose a grid of approximating logarithmic prices at spacings of $\gamma_1 = \gamma = \sqrt{h_0} = 0.0105$ around the initial value of the logarithmic price $y_0^a = \ln S_0 = 6.9078$. Figure 1 shows the first three days of the evolution of the lattice defined by equations (7) through (14) on this grid. Since the parameter γ controls the spacing between the logarithmic returns, all the approximating logarithmic prices are separated by this value.




Figure 1. Lattice of state variables over three days. This figure shows the first three days of the first phase of the lattice for an NGARCH model with parameters $r_f = 0$, $\lambda = 0$, $\beta_0 = 6.575 \times 10^{-6}$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, and $c = 0$. The grid of values for the logarithmic price of the underlying, $y = \ln S$ is determined by taking intervals of size $\gamma = \sqrt{h_0} = 0.0105$ around the log of the initial price $S_0 = 1000$. In this example, $n = 1$, giving three possible paths from each node for a given variance. Each node is represented by a box containing two numbers. The top (bottom) number is the maximum (minimum) variance (multiplied by 10^5) of all paths to that node. These variances determine whether the successor nodes are one or more units of γ apart on the grid. At some nodes a variance value produces three paths that are two units apart.

To help describe the construction of the approximating process, let node (t, i) represent the node at day t when the logarithm of the stock price is at "level" i , where i counts the net number of γ up jumps since date 0. In particular, if $y_t^a(i)$, represents the logarithmic price at this node, then

$$y_t^a(i) = y(0) + i\gamma$$

where $i = -M_d(t), \dots, 0, \dots, M_u(t)$ and $M_d(t)$ ($M_u(t)$) is the maximum number of units of γ that the price can jump down (up) over t consecutive periods.¹⁰ To construct the lattice we begin by using equation (10) to compute η at the initial node. With γ chosen as $\sqrt{h_0}$, the value for the jump parameter at date 0 is $\eta = 1$. This implies that the jumps are "single γ " jumps. Equation (7) is used to find the three successor stock prices. For each of these three jumps, equation (9) is computed to establish the normalized innovation, and equation (8) is used to calculate the variance for the next period. Note that after one period there are three states, each with different prices and variances. The probabilities associated with the three jumps can be computed using equations (11) through (14). By construction, these probabilities are always between 0 and 1. At this stage, however, these probabilities are not required.

In Figure 1 we see that at nodes $(1, -1)$ and $(1, 0)$ the value of $\eta = 1$ ("single γ " jumps), and at node $(1, 1)$ the value is 2 ("double γ " jumps). So at node $(1, 1)$, the three successor nodes are separated from each other by 2γ . Figure 1 shows that the successor nodes at $(1, 1)$ are $(2, -1)$, $(2, 1)$, and $(2, 3)$. The volatility at node $(1, 1)$ is sufficiently large that if we restricted the jumps to be "single γ " jumps, then, in order to match the mean and variance of the conditional normal distribution from that node, some probabilities would fall outside the range $[0, 1]$. By allowing the jumps to be integer multiples of γ we can ensure that the probabilities stay in the interval $[0, 1]$. Notice also that once period 2 is reached, the variances at individual nodes may not be unique. For example, there are three paths to node $(2, -1)$, namely from $(1, 1)$, $(1, 0)$, and $(1, -1)$. In Figure 1 the maximum and minimum variances along all possible paths to each node are reported in the boxes at the nodes. (They are multiplied by 10^5 to ease reading.)

With $n = 1$ there are exactly three successor nodes for each price and variance combination on the lattice. For example, consider node $(2, -1)$. Since there are three possible paths to this node from the initial $(0, 0)$ node, there are three possible variances here. Associated with each of these variances is an η value that determines the successor nodes. At node $(2, -1)$ the three different variances produce two different η values. For the two lower variance values, $\eta = 1$ produces valid probabilities and the successor nodes are the three closest nodes. However, for the higher variance value, $\eta = 2$ and "double-sized" γ jumps are required. Without conditioning on the variance, there are five successor nodes from node $(2, -1)$. However, there are only three successor nodes for each possible variance.

¹⁰ The means for computing $M_d(t)$ and $M_u(t)$ are discussed shortly.

Once all the η values are established for day 1, then the lowest and highest values of the logarithmic prices for day 2 are known. Specifically, the maximum number of down jumps is $M_d(2) = 2$, and the maximum number of up jumps is $M_u(2) = 3$. In general, once the state variables are identified in period t , the η values can be computed and the extreme nodes in period $t + 1$ can be identified.

This example illustrates two properties of the lattice. First, in spite of the fact that variances are heavily path dependent, for a given n we are able to fix the topology so that the logarithms of stock prices on the lattice are separated by a constant, γ_n . As a result, the number of possible stock prices grows linearly in the number of periods, rather than exponentially as in a nonrecombining tree. This desirable property is made possible by Proposition 1, which limits the choice of successor prices to a particular grid of values and ensures that the probability values for the successor nodes can be chosen in such a way that the true means and variances are matched, without violating the necessary condition that the probabilities remain between 0 and 1.

Second, at any given node there are a number of different volatilities. In general, the number of distinct volatilities at a particular node depends on the number of different paths that can be taken to the node. An option price critically depends on volatility, so at each node there could be as many distinct option prices as paths to the node. Since the number of distinct paths to nodes may increase exponentially as the number of time periods increases, it is not feasible to track all distinct volatilities (and option prices) at each node.

To capture this path dependence we keep track of only the maximum and minimum variances, rather than of all the variances, that can occur at each node. We approximate the state space of variances at each node by K values selected to span the range between these maximums and minimums. Let $h_t^{\min}(i)$ and $h_t^{\max}(i)$ represent the minimum and maximum variances that can be attained at node (t, i) . Option prices at this node are computed for K levels of variance ranging from the lowest to the highest at equidistant intervals. In particular, let $h_t^a(i, k)$ represent the k th level of the variance at node (t, i) with $k = 1, \dots, K$, where

$$h_t^a(i, k) = h_t^{\min}(i) + \delta_t(i)(k - 1) \quad \text{for } k = 1, 2, 3, \dots, K \quad (15)$$

and

$$\delta_t(i) = \frac{h_t^{\max}(i) - h_t^{\min}(i)}{K - 1}.$$

Here $\delta_t(i)$ is the constant gap between adjacent variances at node (t, i) . The total number of stock price and variance combinations on the lattice at time period t equals the number of different y_t^a values multiplied by K . Since the number of different stock prices on the lattice grows linearly in the number of time periods, and K is a chosen constant, our lattice captures the path

dependency of the GARCH process without inducing an exploding tree. Convergence of this approximating process to the true GARCH process is ensured as $n \rightarrow \infty$ and $K \rightarrow \infty$.

It is simple to obtain the maximum and minimum values for the variance at any particular node. For the GARCH process in our example, the extreme values at node (t, i) are uniquely determined by the extreme values at the period $(t - 1)$ nodes that can reach node i in period t . These extremes can readily be established using a forward dynamic program.

Figure 1 illustrates this forward scan for the first three days of the lattice. At each day we determine the set of feasible logarithmic prices and at each of these nodes we determine the minimum and maximum variance of all paths to that node. This completes the first phase of the algorithm.

III. Pricing Options on the Lattice

Once the above information is established, option prices can be computed on the lattice using standard backward recursion procedures. The procedure we use follows that used by Ritchken, Sankarasubramanian, and Vijh (1993) in pricing American options on the average, by Hull and White (1993) in pricing a variety of exotics, and by Li, Ritchken, and Sankarasubramanian (1995) in providing efficient implementation of some Heath, Jarrow, and Morton (1992) interest rate models. At each node we evaluate option prices over a grid of K points covering the state space of the variances from the minimum to the maximum. We begin by setting up a vector of option prices of size K at each terminal node. Each entry in the vector corresponds to an option price, with the price in the first cell of the vector corresponding to the smallest conditional variance, and the price at the K th entry corresponding to the price when variance is a maximum. Let $C_t^a(i, k)$ correspond to the k th option price at the node (t, i) (for $k = 1, 2, \dots, K$) when the underlying asset price is $S_t^a(i) = e^{y_t^a(i)}$ and the variance is $h_t^a(i, k)$.

Note that in the final period, T , the payout of a standard claim is fully determined by that period's asset price alone. Hence, each K -vector of option prices at each terminal node consists of K equal entries. For example, the boundary condition for a standard call option with strike X which expires in period T is

$$C_T^a(i, 1) = C_T^a(i, 2) = \dots = C_T^a(i, K) = \text{Max}[0, S_T^a(i) - X]. \quad (16)$$

We apply backward recursion to establish the option price at date 0. Consider a node in period t , say node (t, i) , and assume that the option price $C_t^a(i, k)$, corresponding to variance $h_t^a(i, k)$ at that node, is to be computed. We use Proposition 1 to compute the $2n + 1$ (forward) successor nodes. In particular, given the variance $h_t^a(i, k)$, we compute the appropriate jump parameter, η , by equation (10). The successor nodes for this variance are $(t + 1, i + j\eta)$, where $j = 0, \pm 1, \dots, \pm n$. Equations (8) and (9) are used to compute the period $(t + 1)$

variance for each of these nodes. Specifically, for the transition from the k th variance element of node (t, i) to node $(t + 1, i + j\eta)$, the period $(t + 1)$ variance is given by

$$h^{next}(j) = \beta_0 + \beta_1 h_t^a(i, k) + \beta_2 h_t^a(i, k) [(j\eta\gamma_n - r_f + h_t^a(i, k)/2)/\sqrt{h_t^a(i, k)} - c^*]^2. \quad (17)$$

However, at node $(t + 1, i + j\eta)$, we have stored option prices for only K different variance levels. Though these variances span the space of all possible variances generated along all paths to that node, there may not be a variance entry corresponding exactly to $h^{next}(j)$. Hence, there may not be a corresponding option price stored. If this is the case, then linear interpolation of the two stored option prices corresponding to the two stored variance entries closest to $h^{next}(j)$ is used to obtain the option price corresponding to a variance of $h^{next}(j)$. Let L be an integer smaller than K defined such that

$$h_{t+1}^a(i + j\eta, L) < h^{next}(j) \leq h_{t+1}^a(i + j\eta, L + 1). \quad (18)$$

The interpolated option price is

$$c^{interp}(j) = q(j)C_{t+1}^a(i + j\eta, L) + (1 - q(j))C_{t+1}^a(i + j\eta, L + 1), \quad (19)$$

where

$$q(j) = \frac{h_{t+1}^a(i + j\eta, L + 1) - h^{next}(j)}{h_{t+1}^a(i + j\eta, L + 1) - h_{t+1}^a(i + j\eta, L)}. \quad (20)$$

In this way an option price is identified for each of the $2n + 1$ jumps from node (t, i) with variance $h_t^a(i, k)$. In each case, either node $(t + 1, i + j\eta)$ contains a variance entry (and hence option value) that matches $h^{next}(j)$, or the relevant information is interpolated from the closest two entries. We use equations (11) through (14) to compute the expectation of these option prices over all these $2n + 1$ successor nodes, and discount it at the riskless rate to obtain the unexercised option value $C_t^{a go}(i, k)$. That is,

$$C_t^{a go}(i, k) = e^{-r_f} \sum_{j=-n}^n P(j) c^{interp}(j). \quad (21)$$

Denote the exercised value of the claim by $C_t^{a stop}(i, k)$. For an American call option with strike X , this is

$$C_t^{a stop}(i, k) = \text{Max}[S_t^a(i) - X, 0]. \quad (22)$$

The value of the claim at the k th entry of node (t, i) is then

$$C_t^a(i, k) = \text{Max}[C_t^{a \text{ stop}}(i, k), C_t^{a \text{ go}}(i, k)]. \quad (23)$$

The final option price, obtained by backward recursion of this procedure, is given by $C_0^a(0, 1)$.¹¹

To illustrate this option pricing procedure, reconsider our three-period example where each day a trinomial approximation to the conditional normal distribution is used ($n = 1$). Let $K = 3$, so that the minimum, maximum, and midpoint variances are stored at each node, as are the corresponding option values. Figure 2 illustrates the calculations for the case of a three-period ($T = 3$) at-the-money European call option. In the box representing each node, the maximum and minimum variances over all possible paths to that node are shown in the left column. (These are the same as the values shown in Figure 1.) The right column shows the option values corresponding to these maximum and minimum variances, as well as the values corresponding to the midpoint variances. Notice that at the expiration date all $K = 3$ option values at each node are identical, due to the expiration condition for a European call. At the origin, all $K = 3$ option prices are also identical, since there is only one initial variance.

IV. Convergence Properties of the Algorithm

In this section we illustrate the rate of convergence of option prices produced on the lattice to their true values. Confidence intervals for the theoretical prices of European options are obtained using large sample sizes in Monte Carlo simulations of equations (3) and (4). The rate of convergence of the algorithm to these true option prices is governed by the choice of the parameters n (which governs the order of each day's multinomial approximation to the conditional normal distribution) and K (the number of variances used to span the variance space), with more precision being obtained as these values are increased (for a given choice of γ).

We first investigate the convergence behavior as n increases for a fixed value of $K = 20$. Table I shows the convergence of at-the-money European call option prices for different maturities. The rate of convergence to theoretical prices is quite rapid. For longer term contracts, a choice of $n = 1$ works satisfactorily. For short term European options, one might imagine that a larger n may be necessary. Table I indicates that as the maturity date declines, larger n values are needed. Surprisingly, though, small n values do appear to produce reasonable results. Indeed, using more than nine points (i.e., $n > 4$) to approximate the conditional normal distributions appears to provide little improvement in precision.

¹¹ Actually all K entries of $C_0^a(0, k)$ will be the same because all K entries of $h_0^a(0, k)$ are equal to the initial variance h_0^a .




Figure 2. Illustrative lattice for three-period at-the-money call option. This figure shows the valuation of a three-period at-the-money European call option from the second phase of the lattice procedure. Each node is represented by a box containing five numbers. The top (bottom) number in the first column is the maximum (minimum) variance (multiplied by 10^5) of all paths to that node, as shown in Figure 1. In this example $K = 3$, so three option values are carried at each node. These are shown in the second column. The top number is the option value corresponding to the maximum variance, the bottom number is the value corresponding to the minimum variance, and the middle number corresponds to the midpoint variance.

Table II also compares the prices of different maturity at-the-money call options to prices produced by simulation. In this table $n = 5$ and the sensitivity to the volatility space parameter, K , is explored. For short term contracts a very small K leads to accurate prices. There are few paths to any

Table I
Convergence of At-The-Money Call Option Prices

This table shows the convergence of at-the-money European call option prices from the lattice as the order of the multinomial approximation to each day's conditional normal distribution increases. The lattice is for an NGARCH model with $r_f = 0$, $\lambda = 0$, $\beta_0 = 6.575 \times 10^{-6}$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, $c = 0$, $S_0 = 100$, and $h_0 = 0.0001096$, equivalent to an annualized volatility of 20 percent. The maturities of the contracts are in days. The prices are reported as the number of trinomial steps, n , increases, for fixed $K = 20$ and $\gamma = \sqrt{h_0}$. For example, when $n = 5$, then $(2n + 1) = 11$ points are used to approximate each daily conditional normal distribution. The bottom two rows, ∞_L and ∞_U , show the 95 percent confidence intervals for the true prices based on 500,000 simulations. The table shows that prices converge very rapidly.

Trinomial Steps n	Maturity of Option (days)							
	2	5	10	20	50	75	100	200
1	0.588	0.909	1.327	1.858	2.944	3.607	4.165	5.893
2	0.567	0.906	1.302	1.849	2.931	3.592	4.157	5.872
3	0.573	0.928	1.310	1.852	2.930	3.591	4.079	5.870
4	0.574	0.921	1.307	1.851	2.931	3.591	4.148	5.870
5	0.584	0.927	1.309	1.851	2.930	3.591	4.148	5.869
10	0.584	0.925	1.309	1.851	2.929	3.590	4.147	5.868
25	0.588	0.927	1.309	1.850	2.929	3.589	4.147	5.868
∞_L	0.587	0.923	1.306	1.846	2.918	3.573	4.142	5.862
∞_U	0.592	0.931	1.317	1.862	2.944	3.605	4.179	5.916

node, so there is little path dependence. For longer term contracts there is more path dependence and K needs to be increased. For all maturities, however, prices produced with $K = 20$ and $K = 40$ are almost identical.¹²

Table III compares the prices of a range of in- and out-the-money contracts from a lattice with $n = 5$ and $K = 20$ to their theoretical values. In all cases the algorithm produces results close to the theoretical values.

Since there is no available method for computing the theoretical value of an American option under a GARCH process, we examine the behavior of the early exercise premium with respect to the parameter n . Table IV shows prices from lattices with different n 's and $K = 20$ for a variety of at-the-money American puts. Also shown are the corresponding European prices (which converge well on the basis of the preceding analysis) and the early exercise premium expressed as a percentage of the American price. It is clear from this table that, as was the case for the European call options in Table I, there is little advantage to using multinomial approximations with $n > 4$.

The algorithm presented is not restricted to NGARCH models. It can be applied to most GARCH models where variances are updated according to their current values and recent asset innovations. By modifying the vari-

¹² We assume a constant K throughout the lattice. Though we can make the algorithm more efficient by allowing K to vary according to the range of variances at each node, our goal is to evaluate the convergence of as simple an algorithm as possible.

Table II
Convergence of At-The-Money Call Option Prices

This table shows the convergence of at-the-money European call option prices when the number of variances, K , carried at each node in the lattice increases. The other parameters ($r_f = 0$, $\lambda = 0$, $\beta_0 = 6.575 \times 10^{-6}$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, $c = 0$, $S_0 = 100$, $h_0 = 0.0001096$, and $\gamma = \sqrt{h_0}$) are the same as in Table I, with $n = 5$. The bottom two rows, ∞_L and ∞_U , show the 95 percent confidence intervals for the true prices based on 500,000 simulations. As K increases, the prices converge to their true values. The table shows that for very short term contracts, $K = 2$ may suffice. However, for contracts exceeding 10 days, more points are required.

Number of Variances K	Maturity of Option (days)								
	5	10	30	60	100	150	200	250	300
2	0.926	1.303	2.213	3.085	3.957	4.831	5.570	6.222	6.811
3	0.927	1.307	2.246	3.155	4.061	4.967	5.731	6.404	7.013
4	0.927	1.308	2.257	3.181	4.100	5.018	5.791	6.473	7.089
5	0.928	1.309	2.261	3.192	4.118	5.042	5.821	6.507	7.127
10	0.927	1.309	2.267	3.206	4.141	5.073	5.858	6.550	7.174
20	0.927	1.309	2.268	3.210	4.148	5.082	5.869	6.562	7.188
40	0.927	1.310	2.268	3.211	4.149	5.083	5.871	6.564	7.190
∞_L	0.923	1.306	2.257	3.197	4.142	5.074	5.862	6.556	7.164
∞_U	0.931	1.317	2.277	3.225	4.179	5.120	5.916	6.617	7.230

ance updating scheme in equation (8), European and American option prices can readily be computed.

V. Pricing under Generalized GARCH and Bivariate Diffusion Processes

Some researchers find the GARCH specification mechanical from the perspective of option pricing and prefer modeling volatility as a separate diffusion, perhaps correlated with the return generating diffusion. It turns out, however, that certain *univariate* GARCH processes can be used to approximate some *bivariate* diffusion models with stochastic volatility. Nelson (1990), for example, has shown that the linear and exponential GARCH models converge weakly to specific stochastic volatility bivariate diffusion models. In a very insightful paper, Duan (1997) generalizes these results. In particular, he shows that his augmented family of stationary generalized GARCH processes can be used to approximate a rich array of bivariate diffusion models.

The relationship between GARCH processes and bivariate diffusion models allows the rich set of statistical tools for GARCH to be used to estimate the parameters of bivariate diffusions. Duan (1997) also argues that the rich family of numerical option pricing techniques, already developed for diffusion systems, could be used to compute option prices even if the GARCH option pricing model is the preferred paradigm. We now show that the reverse is true. Specifically, pricing options under bivariate diffusions can be

Table III
Call Option Prices for Different Strikes

This table compares the prices of European call options for different strike prices with 95 percent confidence intervals from simulations with 500,000 replications. The other parameters ($r_f = 0$, $\lambda = 0$, $\beta_0 = 6.575 \times 10^{-6}$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, $c = 0$, $S_0 = 100$, $h_0 = 0.0001096$, $\gamma = \sqrt{h_0}$, and $K = 20$) are the same as for Table I, with $n = 5$. All but two of the lattice prices are within these confidence intervals.

Strike <i>X</i>	Maturity (days)				
	5	10	30	50	100
95.0	5.012 (5.008,5.021)	5.086 (5.082,5.100)	5.560 (5.543,5.571)	6.030 (6.009,6.044)	7.028 (7.010,7.056)
97.5	2.665 (2.660,2.671)	2.915 (2.909,2.925)	3.712 (3.696,3.720)	4.316 (4.299,4.329)	5.468 (5.456,5.497)
100.0	0.927 (0.923,0.931)	1.309 (1.306,1.317)	2.268 (2.257,2.277)	2.930 (2.918,2.944)	4.148 (4.142,4.179)
102.5	0.178 (0.177,0.181)	0.439 (0.438,0.444)	1.263 (1.256,1.271)	1.885 (1.876,1.897)	3.069 (3.067,3.100)
105.0	0.018 (0.019,0.020)	0.108 (0.109,0.112)	0.639 (0.635,0.646)	1.148 (1.143,1.159)	2.214 (2.214,2.242)

accomplished by first approximating the bivariate diffusions by generalized GARCH processes, and then efficiently pricing contracts under these processes using modifications to our simple lattice algorithm.

Rather than describing the algorithm from the most general perspective of Duan’s augmented GARCH family, we make matters more concrete by illustrating it using the NGARCH specification used in the previous sections. The conditional variance process of the limiting diffusion for this model is a process which includes that used by Hull and White (1987) in their stochastic volatility option pricing model. Furthermore, Duan (1997) shows that the limiting diffusion is the same as the limiting form of the Glosten, Jagannathan, and Runkle (1993) GJR-GARCH model, which is commonly used to model the volatility in equity prices.

Partition each time period (“day”) into m “trading periods” of width $\Delta t = 1/m$. Label these trading periods by consecutive integers starting from the beginning of the current period, period 0. Let y_t be the logarithmic price at the beginning of the t th trading period, and let h_t be the variance for this trading period. Assume y_0 and h_0 are given. The dynamics of the return generating process are

$$y_{t+1} = y_t + (r_f + \lambda\sqrt{h_t} - h_t/2)\Delta t + \sqrt{h_t}\sqrt{\Delta t}v_{t+1}, \tag{24}$$

$$h_{t+1} = h_t + \beta_0\Delta t + h_t[\beta_1 + \beta_2q - 1]\Delta t + h_t\beta_2\sqrt{\Delta t}[(v_{t+1} - c)^2 - q], \tag{25}$$

Table IV
Convergence of American Put Option Prices

This table compares the prices of American at-the-money puts to European prices, as the order, n , of the multinomial approximation to each day's conditional normal distribution increases. For example, when $n = 2$, there are $2n + 1 = 5$ points used to approximate the daily conditional normal distribution. The other parameters ($\lambda = 0$, $\beta_0 = 6.575 \times 10^{-6}$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, $c = 0$, $S_0 = 100$, $h_0 = 0.0001096$, $\gamma = \sqrt{h_0}$, and $K = 20$) are the same as for Table I, except that $r_f = 0.1$. The early exercise premium is expressed as a percentage of the American option price. The table shows that American prices stabilize very quickly, especially when the number of days is greater than 10.

Maturity (days)	n	American Price	European Price	Early Exercise Premium (% of American)
2	1	0.563	0.563	0
	2	0.540	0.540	0
	3	0.546	0.546	0
	4	0.548	0.548	0
	5	0.556	0.556	0
10	1	1.216	1.194	1.83
	2	1.187	1.168	1.58
	3	1.193	1.176	1.43
	4	1.190	1.173	1.43
	5	1.192	1.175	1.40
50	1	2.419	2.294	5.19
	2	2.400	2.281	4.97
	3	2.399	2.281	4.94
	4	2.398	2.281	4.92
	5	2.398	2.281	4.92
100	1	3.168	2.899	8.50
	2	3.146	2.884	8.35
	3	3.144	2.882	8.33
	4	3.143	2.882	8.32
	5	3.143	2.882	8.32

where $q = (1 + c^2)$ and the sequence $\{v_t | t = 0, 1, 2, \dots\}$ is a sequence of independent standard normal random variables. Notice that when $\Delta t = m = 1$ this equation reduces to the GARCH model specified in equations (1) and (2).

Duan (1996b) shows through an equilibrium argument that a local risk-neutralized valuation relationship holds, in which options can be priced as if the economy were risk neutral under the pricing measure

$$y_{t+1} = y_t + (r_f - h_t/2)\Delta t + \sqrt{h_t}\sqrt{\Delta t}\epsilon_{t+1}, \quad (26)$$

$$h_{t+1} = h_t + \beta_0 \Delta t + h_t[\beta_1 + \beta_2 q - 1]\Delta t + h_t \beta_2 \sqrt{\Delta t}[(\epsilon_{t+1} - c - \lambda\sqrt{\Delta t})^2 - q], \quad (27)$$

where $\{\epsilon_t | t = 0, 1, 2, \dots\}$ is a sequence of independent standard normal random variables. When $\Delta t = m = 1$, these equations reduce to equations (3) and (4).

Duan shows that the limiting diffusion of the process for the price of the underlying security in equations (24) and (25) is

$$dy_t = (r_f + \lambda\sqrt{h_t} - h_t/2)dt + \sqrt{h_t}dW_1(t), \quad (28)$$

$$dh_t = [\beta_0 + (\beta_1 + \beta_2 q - 1)h_t]dt - 2c\beta_2 h_t dW_1(t) + \sqrt{2}\beta_2 h_t dW_2(t), \quad (29)$$

where $dW_1(t)$ and $dW_2(t)$ are independent Wiener processes. The limiting diffusion under the locally risk-neutralized measure is shown to be

$$dy_t = (r_f - h_t/2)dt + \sqrt{h_t}dZ_1(t), \quad (30)$$

$$dh_t = [\beta_0 + (\beta_1 + \beta_2 q - 1 + 2\lambda\beta_2 c)h_t]dt - 2c\beta_2 h_t dZ_1(t) + \sqrt{2}\beta_2 h_t dZ_2(t), \quad (31)$$

where $dZ_1(t)$ and $dZ_2(t)$ are independent Wiener processes.

With minor adjustments our lattice algorithm can price options under the bivariate diffusion process given in equations (30) and (31). This is accomplished by pricing the contracts under the approximating generalized GARCH process in equations (26) and (27). The adjustment involves partitioning each period (day) of the procedures, given in Sections II and III, into m trading periods, each of length $\Delta t = 1/m$, and modifying equations (7) through (9) to reflect equations (26) and (27) rather than equations (3) and (4). In particular, in Section II we approximate the “daily” GARCH model on our lattice by effectively dividing each day into n equal increments, where each increment has the same volatility. In the generalized GARCH model here, each of the m trading periods each day has a different volatility. We approximate this model on our lattice by dividing each of these m trading periods into n increments of equal volatility, approximated by trinomial distributions.

Thus, for each trading period of each day, we have an approximation for the conditional normal distribution of the logarithmic price that consists of a discrete distribution over $(2n + 1)$ points. The approximating distribution has the property that its first two moments match the true conditional moments. As m increases, $\Delta t \rightarrow 0$ and the generalized GARCH process converges to the required bivariate diffusion. Hence our algorithm for pricing options under *discrete time* GARCH is easily modified to price options under generalized GARCH processes and hence under *continuous time* bivariate diffusions.

Table V illustrates the prices, for a range of maturity at-the-money options, obtained using the lattice (with $n = 1$ and $K = 20$) as the number of trading periods per day, m , is increased. Also shown are the confidence intervals for the theoretical prices generated by simulation of the bivariate

Table V
Convergence of Generalized GARCH Option Prices

This table shows the convergence of at-the-money European call option prices as the number of trading periods per day, m , increases. The other parameters ($r_f = 0$, $\lambda = 0$, $\beta_0 = 6.575 \times 10^{-6}$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, $c = 0$, $S_0 = 100$, $h_0 = 0.0001096$, $\gamma = \sqrt{h_0}$, and $K = 20$) are the same as for Table I, with $n = 1$. Over each trading period, the price can move to one of three values. The last column, Diffusion Limit, shows 95 percent confidence intervals for the true prices based on 100,000 simulations of the limiting bivariate diffusion model with $m = 48$. The table clearly shows that for contracts with maturities greater than 20 days, $m = 1$ will suffice for the lattice.

Maturity (days)	Trading Periods per Day (M)					Diffusion Limit
	1	2	3	4	5	
2	0.589	0.617	0.603	0.598	0.595	0.580, 0.591
5	0.909	0.939	0.932	0.933	0.931	0.922, 0.939
10	1.312	1.318	1.315	1.315	1.315	1.313, 1.338
20	1.857	1.859	1.860	1.860	1.860	1.854, 1.890
50	2.942	2.943	2.943	2.942	2.942	2.932, 2.989
100	4.165	4.165	4.165	4.164	4.163	4.149, 4.231
200	5.893	5.893	5.893	5.893	5.893	5.804, 5.922

diffusion. Notice that the lattice prices converge very rapidly. Indeed, allowing the number of trading periods per day to be just one or two ($m = 1$ or 2) appears to suffice when the maturity of the contract exceeds ten days.

Table VI compares option prices, for a range of strikes and maturities, from the lattice (with $n = 1$ and $K = 20$) to prices obtained by simulation of the limiting bivariate diffusion process. For contracts with maturities shorter than 20 days, each day is broken up into four trading periods ($m = 4$); for longer term contracts only one trading period per day ($m = 1$) is used. The lattice produces results that are consistently within the confidence intervals produced by the simulation.

Our algorithm is easily modified to price options under a wide variety of stochastic volatility bivariate diffusions. In particular, Nelson (1990) and Duan (1997) identify the diffusion limits of a large family of generalized GARCH processes and show that they include almost all the stochastic volatility models that have been developed in the literature to price options. The risk-neutralized process, and its diffusion limit, is also derived in Duan (1995, 1996b) allowing options to be priced as expected terminal values discounted at the riskless rate. Table VII provides examples from this family. It shows a number of stochastic volatility bivariate diffusion models that have been used in the literature and the generalized GARCH processes that converge to these diffusion limits.

Our algorithm is easily modified to handle all these specifications.¹³ Indeed, our lattice procedure is viable as long as the distribution of the logarithm of the asset price is conditionally normal and the variance updating

¹³ Duan (1995, 1996b) shows that the risk-neutralized process for each generalized GARCH model shown in Table VII is obtained by replacing v_{t+1} by $\epsilon_{t+1} - \lambda\sqrt{\Delta t}$.

Table VI
Pricing Options Under Bivariate Diffusions

This table compares the prices of European call options for different strike prices with 95 percent confidence intervals from 100,000 simulations of the limiting bivariate diffusion model. The parameters ($r_f = 0$, $\lambda = 0$, $\beta_0 = 6.575 \times 10^{-6}$, $\beta_1 = 0.90$, $\beta_2 = 0.04$, $c = 0$, $S_0 = 100$, $h_0 = 0.0001096$, $\gamma = \sqrt{h_0}$, and $K = 20$) are the same as for Table I, with $n = 1$ and $m = 4$ ($m = 1$ for maturities of 20 days or longer). That is over each quarter-day (day for maturities of 20 days or longer), a three-point approximation is made to the conditional normal distribution. The lattice algorithm works well, even in this case, since every price falls within its corresponding confidence interval.

Maturity (days)	Strike				
	95.0	97.5	100.0	102.5	105.0
2	5.000	2.523	0.598	0.028	0.000
	(5.000, 5.000)	(2.507, 2.525)	(0.580, 0.591)	(0.028, 0.031)	(0.000, 0.000)
5	5.011	2.665	0.933	0.178	0.016
	(5.000, 5.013)	(2.642, 2.668)	(0.922, 0.939)	(0.174, 0.182)	(0.016, 0.018)
10	5.085	2.917	1.315	0.443	0.107
	(5.061, 5.101)	(2.902, 2.936)	(1.313, 1.338)	(0.440, 0.455)	(0.104, 0.111)
20	5.314	3.361	1.857	0.900	0.371
	(5.294, 5.347)	(3.339, 3.385)	(1.854, 1.890)	(0.890, 0.915)	(0.363, 0.379)
50	6.039	4.332	2.942	1.901	1.160
	(6.025, 6.102)	(4.316, 4.383)	(2.932, 2.989)	(1.888, 1.935)	(1.150, 1.186)
100	7.043	5.489	4.165	3.090	2.232
	(7.019, 7.121)	(5.465, 5.557)	(4.149, 4.231)	(3.069, 3.141)	(2.213, 2.274)
200	8.589	7.158	5.893	4.803	3.870
	(8.478, 8.617)	(7.055, 7.184)	(5.804, 5.922)	(4.718, 4.826)	(3.790, 3.888)

process is a predictable function of its current level and the current innovation. For pure GARCH models we only need to replace the variance updating scheme in equation (8) (which reflects equation (4)) with the appropriate scheme. Similarly, for the bivariate diffusion model we only need to replace equation (27) with the appropriate variance updating scheme that ensures convergence to the required bivariate diffusion.

VI. Conclusions

This paper develops an algorithm that permits American option prices to be computed when the underlying security price is driven by a wide variety of GARCH and stochastic volatility processes. The algorithm is described in detail for the NGARCH model, but readily extends to handle almost all GARCH processes where the variance updating mechanism is a predictable function of the current level of the variance and the current innovation. The algorithm differs from the usual lattice implementation of option pricing models in that a vector of option prices is carried along at each asset price in the tree. Nonetheless, computations are efficient. Indeed, the algorithm is useful not only for pricing American contracts but also for European contracts, since fewer computations are required than for simulation procedures.

Table VII
Some Generalized GARCH Models and their Diffusion Limits

This table shows some of the generalized GARCH models, and their diffusion limits, for which the lattice can compute option prices under the appropriate risk neutralized measure. For the GARCH models, ν_{t+1} is the univariate innovation driving both the current period's asset price and the conditional variance for next period's price. The corresponding Wiener process driving the asset price in the diffusion limit is $dW_1(t)$, which is independent of the other Wiener process, $dW_2(t)$, driving volatility. (See Duan (1997, 1996b) for further details.) References are given where a model nests, as a special case, a specification used in the literature.

Generalized GARCH Model	Stochastic Volatility Diffusion Limit
Linear: Bollerslev (1986) $h_{t+1} = h_t + \beta_0 \Delta t + (\beta_1 + \beta_2 - 1)h_t \Delta t + \beta_2 h_t (\nu_{t+1}^2 - 1)\sqrt{\Delta t}$ Glosten, Jagannathan, and Runkle (1993) $h_{t+1} = h_t + \beta_0 \Delta t + (\beta_1 + \beta_2 + \beta_3/2 - 1)h_t \Delta t$ $+ \beta_2 h_t (\nu_{t+1}^2 - 1)\sqrt{\Delta t} + \beta_3 h_t (\max(0, -\nu_{t+1})^2 - \frac{1}{2})\sqrt{\Delta t}$	$dh_t = \beta_0 dt + (\beta_1 + \beta_2 - 1)h_t dt + \sqrt{2}\beta_2 h_t dW_2(t)$ Hull and White (1987) $dh_t = \beta_0 dt + \beta_1 + \beta_2 + \beta_3/2 - 1)h_t dt - \beta_3 \sqrt{2/\pi} h_t dW_1(t)$ $+ \sqrt{2}\beta_2^2 + (5\pi - 8)/(4\pi)\beta_3^2 + 2\beta_2\beta_3 h_t dW_2(t)$
Nonlinear Asymmetric: Engle and Ng (1993) $h_{t+1} = h_t + \beta_0 \Delta t + (\beta_1 + \beta_2(1 + c^2) - 1)h_t \Delta t$ $+ \beta_2 h_t ((\nu_{t+1} - c)^2 - (1 + c^2))\sqrt{\Delta t}$	Hull and White (1987) $dh_t = \beta_0 dt + (\beta_1 + \beta_2(1 + c^2) - 1)h_t dt - 2c\beta_2 h_t dW_1(t)$ $+ \sqrt{2}\beta_2 h_t dW_2(t)$
Exponential: Nelson (1991) $\ln h_{t+1} = \ln h_t + (\beta_0 + \beta_1 - 1 + \beta_4(2/\sqrt{2\pi} - 1) + \beta_5(1/\sqrt{2\pi} - 1))\Delta t$ $+ (\beta_1 - 1)\ln h_t \Delta t$ $+ \beta_4(\nu_{t+1} - 2/\sqrt{2\pi})\sqrt{\Delta t} + \beta_5(\max(0, -\nu_{t+1}) - 1/\sqrt{2\pi})\sqrt{\Delta t}$	Wiggins (1987) $d \ln h_t = (\beta_0 + \beta_1 - 1 + \beta_4(2/\sqrt{2\pi} - 1) + \beta_5(1/\sqrt{2\pi} - 1))dt$ $+ (\beta_1 - 1)\ln h_t dt - \beta_5/2 dW_1(t)$ $+ \beta_4 + \beta_5/2 \sqrt{(\pi - 2)/\pi} dW_2(t)$
$\sqrt{h_{t+1}} = \sqrt{h_t} + (\beta_0 - \beta_1 + 1 + \beta_4(2/\sqrt{2\pi} - 1) + \beta_5(1/\sqrt{2\pi} - 1))/2\Delta t$ $+ (\beta_1 - 1)\sqrt{h_t} \Delta t + \beta_4(\nu_{t+1} - 2/\sqrt{2\pi})/2\sqrt{\Delta t}$ $+ \beta_5(\max(0, -\nu_{t+1}) - 1/\sqrt{2\pi})/2\sqrt{\Delta t}$	Scott (1987), Stein and Stein (1991), and Heston (1993) $d\sqrt{h_t} = (\beta_0 - \beta_1 + 1 + \beta_4(2/\sqrt{2\pi} - 1) + \beta_5(1/\sqrt{2\pi} - 1))/2dt$ $+ (\beta_1 - 1)\sqrt{h_t} dt - \beta_5/4 dW_1(t)$ $+ (\beta_4 + \beta_5/2 /2)\sqrt{(\pi - 2)/\pi} dW_2(t)$

This paper also shows how options can be priced using our lattice under generalized GARCH processes. These processes are important because they converge to a very rich family of stochastic volatility bivariate diffusions. We illustrate this by constructing a lattice procedure for the generalized NGARCH process which has a limiting bivariate diffusion that includes the system studied by Hull and White (1987). The convergence of the lattice prices to prices generated by the stochastic volatility process is explored in detail and found to be quite rapid. In general, our lattice algorithm is easily modified to price claims where the dynamics of the underlying are driven by a wide variety of stochastic volatility processes. For example, option prices under dynamics postulated by Heston (1993), Hull and White (1987), Scott (1987), Stein and Stein (1991), Wiggins (1987), and others, can readily be generated.

The research of Nelson (1990) and Duan (1997) has resulted in the merging of the very rich GARCH and stochastic volatility literatures. It is possible to use the statistical procedures for GARCH models to estimate the parameters of many stochastic volatility bivariate diffusions. Our research strengthens the connections between these two literatures by providing a single lattice algorithm that prices options under a very wide range of GARCH and bivariate diffusion processes.

Extensive empirical studies of option prices generated by GARCH models, and some bivariate diffusions, have been lacking in the literature, primarily due to the lack of efficient algorithms for pricing American, and even European, options. Duan (1996a), for example, illustrates how GARCH option models might be used to explain the volatility smile. Engle and Mustafa (1992) suggest a strong link between parameters implied from option prices and those estimated from the time series data for the price of the underlying security. In both studies, however, only very limited data sets are utilized since very costly recursive simulations are used to imply out the parameters of a GARCH model. With our efficient pricing algorithm we now can conduct more complete empirical tests over longer time horizons. As a result, the stability of GARCH parameters implied by option prices, and option pricing biases from such models, can be more thoroughly explored. Moreover, our algorithm permits alternative volatility evolution processes to be compared using the rich option data sets now available. It remains for future research to conduct these tests to obtain a clearer understanding of the processes governing the evolution of volatility in different markets.

Appendix

Proof of Proposition 1: We begin this proof by considering the approximation over a single period, which we call a “day”. Assume that the state variables at the beginning of this period are (y_t^a, h_t^a) . We construct a discrete approximation over the period such that the logarithmic price can move to n equidistant up and down locations or stay the same, where the gap between any two logarithmic prices is an integer multiple of $\gamma_n = \gamma/\sqrt{n}$, a constant.

Probabilities are set for these $2n + 1$ points such that the mean and variance of the logarithmic return match the mean and variance of the true process. We split this single period into n subintervals of equal width $1/n$. The variance stays unchanged over each of these n subintervals. At the end of the day the variance is updated according to equation (4). We approximate the day's conditional normal distribution by a sequence of n independent identical trinomial distributions over the n subintervals. Given h_t^a , η can be computed. Over each subinterval of the day allow the logarithm of the asset price to increase by $\eta\gamma_n$ (with probability p_u), stay unchanged (with probability p_m), or decrease by $\eta\gamma_n$ (with probability p_d). Notice that the determination of η ensures that these probability values are between 0 and 1. The expected change and the variance of the change over each subinterval are readily computed to be $(r_f - h_t^a/2)/n$ and h_t^a/n . Since the n trinomial trials in the day are independent and identically distributed, the mean and variance over the full day (period) are obtained as the required result.

The end of period distribution converges to a normal distribution, by the usual central limit theorem, as the number of subintervals in the period (n) increases. In particular, the limiting distribution of ϵ_{t+1}^a , viewed from date t , is a standard normal random variable.

It follows that, as n increases, y_{t+1}^a (conditional on y_t^a and h_t^a) converges to a normal random variable with mean and variance equal to the mean and variance of the true process. It then follows that the discrete time, discrete state process in equations (7) to (9) converges to the discrete time, continuous state process in equations (3) and (4).

REFERENCES

- Amin, Kaushik, and Victor Ng, 1993, ARCH processes and option valuation, Unpublished manuscript, University of Michigan.
- Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–327.
- Bollerslev, Tim, Ray Chou, and Kenneth Kroner, 1992, ARCH modeling in finance: A review of the theory and empirical evidence, *Journal of Econometrics* 52, 1–59.
- Bollerslev, Tim, Robert Engle, and Daniel Nelson, 1994, ARCH models; in Robert Engle and Daniel McFadden, ed.: *Handbook of Econometrics* (Elsevier Science B.V., Amsterdam).
- Boyle, Phelim, Jeremy Evnine, and Stephen Gibbs, 1989, Numerical evaluation of multivariate contingent claims, *Review of Financial Studies* 2, 241–250.
- Duan, Jin-Chuan, 1995, The GARCH option pricing model, *Mathematical Finance* 5, 13–32.
- Duan, Jin-Chuan, 1996a, Cracking the smile, *Risk* 9, 55–59.
- Duan, Jin-Chuan, 1996b, A unified theory of option pricing under stochastic volatility—from GARCH to diffusion, Unpublished manuscript, Hong Kong University of Science and Technology.
- Duan, Jin-Chuan, 1997, Augmented GARCH(p,q) process and its diffusion limit, *Journal of Econometrics* 79, 97–127.
- Duan, Jin-Chuan, and Jean-Guy Simonato, 1997, American option pricing under GARCH by a Markov chain approximation, Unpublished manuscript, Hong Kong University of Science and Technology.
- Engle, Robert, and Chowdhury Mustafa, 1992, Implied ARCH models from options prices, *Journal of Econometrics* 52, 289–311.

- Engle, Robert, and Victor Ng, 1993, Measuring and testing the impact of news on volatility, *Journal of Finance* 48, 1749–1779.
- Engle, Robert, and Joshua Rosenberg, 1995, GARCH gamma, *Journal of Derivatives* 4, 47–59.
- Glosten, Lawrence, Ravi Jagannathan, and David Runkle, 1993, Relationship between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.
- Heath, David, Robert Jarrow, and Andrew Morton, 1992, Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation, *Econometrica* 60, 77–105.
- Heston, Steven, 1993, A closed-form solution for options with stochastic volatility, *Review of Financial Studies* 6, 327–344.
- Hull, John, and Allan White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281–300.
- Hull, John, and Allan White, 1993, Efficient procedures for valuing European and American path dependent options, *Journal of Derivatives* 1, 21–32.
- Kallsen, Jan, and Murad Taqqu, 1998, Option pricing in ARCH-type models, *Mathematical Finance* 8, 13–26.
- Kamrad, Bardia, and Peter Ritchken, 1991, Multinomial approximating models for options with k state variables, *Management Science* 37, 1640–1653.
- Li, Anlong, Peter Ritchken, and L. Sankarasubramanian, 1995, Lattice models for pricing American interest rate claims, *Journal of Finance* 50, 719–737.
- Nelson, Daniel, 1990, ARCH models as diffusion approximations, *Journal of Econometrics* 45, 7–38.
- Nelson, Daniel, 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.
- Ritchken, Peter, L. Sankarasubramanian, and Anand Vih, 1993, The valuation of path dependent contracts on the average, *Management Science* 39, 1202–1213.
- Scott, Lois, 1987, Option pricing when the variance changes randomly: Theory, estimation and an application, *Journal of Financial and Quantitative Analysis* 22, 419–438.
- Stein, Elias, and Jeremy Stein, 1991, Stock price distributions with stochastic volatility: An analytic approach, *Review of Financial Studies* 4, 727–752.
- Wiggins, James, 1987, Option values under stochastic volatility: Theory and empirical evidence, *Journal of Financial Economics* 19, 351–372.