

# Sequential Monte Carlo Pricing of American-Style Options under Stochastic Volatility Models

B. Ricky Rambharat\*and Anthony Brockwell†

April 24, 2006

## Abstract

We introduce a new method to price American-style options on underlying investments governed by stochastic volatility models. The method combines a standard gridding approach to solving the associated dynamic programming problem, with a sequential Monte Carlo scheme to estimate required posterior distributions of the latent volatility process. The method represents a refinement of previous algorithms since it does not require the volatility process to be directly observable. Instead, the sequential Monte Carlo scheme provides accurate estimates of the required conditional distributions. Furthermore, the method incorporates market price of volatility risk, and is generalizable to handle different kinds of stochastic volatility models. We also demonstrate that with historical data for two stocks, and appropriately chosen market price of volatility risk, the algorithm yields option prices which are highly consistent with market data.

**Keywords:** American option, pricing, stochastic volatility model, arbitrage, risk-neutral, dynamic programming, sequential, Monte Carlo, decision

---

\*Institute of Statistics and Decision Sciences, Duke University, Durham, NC 27708; e-mail: ricky@stat.duke.edu.

†Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213; e-mail: abrock@stat.cmu.edu.

# 1 Introduction

American-style option contracts are traded extensively over several exchanges. These early-exercise financial derivatives are typically written on equity stocks, foreign currency, and some indices, and include, among other examples, options on individual equities traded on The American Stock Exchange (AMEX), options on currency traded on the Philadelphia Stock Exchange (PHLX), and the OEX index options on the S&P 100 Index traded on the Chicago Board Options Exchange (CBOE). As with any other kind of option, methods for pricing are based on assumptions about the probabilistic model governing the evolution of the underlying asset. Arguably, stochastic volatility models are the most realistic models to date for underlying equities, but existing methods for pricing American-style options have mostly been developed using less realistic models, or else assuming that volatility is observable. In this paper we develop a new method for pricing American-style options, when the underlying process is governed by a stochastic volatility model, and the volatility is not directly observable. The method yields near-optimal solutions under the model assumptions, and can also formally take into account market price of volatility risk.

It follows from the fundamental theorem of arbitrage that an option price can be determined by computing the discounted expectation of the payoff of the option under a risk-neutral measure, assuming that the exercise decision is made so as to maximize the payoff. While this is simple to compute for European-style options, the pricing problem is enormously more difficult for American-style options, due to the possibility of early exercise. For American-style options, the price is in fact the supremum over a large range of possible stopping times of the discounted expected payoff under the risk-neutral measure. A range of methods has been developed to find this price, or equivalently, to solve a corresponding stochastic dynamic programming problem. Due to the difficulty of the problem, certain assumptions are usually made. For instance, a number of effective algorithms (including those developed in Brennan and Schwartz (1977), Geske and Johnson (1984), Carr et al. (1992), Broadie and Glasserman (1997), Longstaff and Schwartz (2001), and Rogers (2002)) are based on the assumption that the underlying asset price is governed by a univariate diffusion process with constant and/or directly observable volatility.

As recognized by Black and Scholes (1973) and others, the assumption of constant volatility is typically inconsistent with observed data. A variety of more realistic models has subsequently been developed for asset prices, with stochastic volatility models arguably representing the best models to-date. This has led researchers to develop pricing methods (Hull and White, 1987; Heston, 1993; Stein and Stein, 1991; Fouque et al., 2000) for European-style options when the underlying asset price is governed by stochastic volatility models. However, work on pricing of American-style options under stochastic volatility models is far less-developed. A number of authors (including Clarke and Parrott, 1999; Finucane and Tomas, 1997; Tzavalis and Wang, 2003; Guan and Xiaoqiang, 2000; Fouque et al., 2000) have made valuable inroads in addressing this problem, but most assume that volatility is observable. For example, Clarke and Parrott (1999) use a multi-grid approach where both the asset price and volatility are state variables. Guan and Xiaoqiang (2000) derive a lattice-based solution to pricing American options with stochastic volatility. They construct a lattice that depends directly on the asset price and volatility and they illustrate an empirical study where they back out the parameters from the stochastic volatility model using data on American-style S&P500 futures options. Fouque et al. (2000) also provide an approximation scheme based on the assumption of fast mean-reversion and they use a clever asymptotic expansion method to correct the constant volatility option price to account for stochastic volatility. The correction involves parameters estimated from the implied volatility surface and they derive a pricing equation that does not depend directly on the volatility process.

The approach we develop, in contrast with the aforementioned methods, combines the optimal decision-making problem with the volatility estimation problem. We assume that the asset price obeys a stochastic volatility model, that observations are made at discrete points in time  $t = 0, 1, 2, \dots$ , and that exercise decisions are made immediately after each observation. Our pricing scheme is based on two key observations. First, we can use a sequential Monte Carlo (also referred to as “particle filtering”) scheme to perform inference on the unobserved volatility process at any given point in time. Second, conditional distributions of the unobserved volatility at a given point in time, given current and past observations of the price process, which are necessary for finding an exact solution to the dynamic programming problem, can be well-approximated by

log-normal distributions. Since market price of volatility risk is reflected in the equation for the evolution of log-volatility in the stochastic volatility model, our method also provides a formal way to incorporate this quantity into the pricing scheme.

It is beyond the scope of this paper to address the issue of estimation and/or model selection for stochastic volatility models. A number of useful papers can be found covering this topic, including, Ghysels et al. (1996); Gallant et al. (1997); Jacquier et al. (1994); Kim et al. (1998); Chernov and Ghysels (2000); Pan (2002); Eraker (2004), and further references found within these papers. The purpose of this paper is to address the problem of pricing an American-style option, once a good model has already been found. It is worth noting, however, that since neither the sequential Monte Carlo scheme nor the gridding approach used in our pricing scheme are tied to a particular model, the method is generalizable in a straightforward manner to handle a fairly wide range of stochastic volatility models. Thus it could be used to perform option pricing under a range of variants of stochastic volatility models, such as those discussed in Chernov et al. (2003).

The paper is organized as follows. In Section 2, we formally state the class of stochastic volatility models we work with. In Section 3 we review the dynamic programming approach for pricing American-style options and demonstrate how it can be transformed into an equivalent form, and introduce (1) a sequential Monte Carlo scheme that yields certain conditional distributions, and (2) a gridding algorithm that makes use of the sequential Monte Carlo scheme to compute option prices. In Section 4 we apply the method to observed data and comment on the results. We then conclude with additional discussion in Section 5.

## 2 The Pricing Problem

### 2.1 The Stochastic Volatility Model

Let  $(\Omega, \mathcal{F}, P)$  be a probability space, and let  $\{S(t), t \geq 0\}$  be a stochastic process defined on  $(\Omega, \mathcal{F}, P)$ , describing the evolution of our asset price over time. Time  $t = 0$  will be referred to as the “current” time, and we will be interested in an option with expiry time  $(T\Delta) > 0$ , with  $T$  being some positive integer, and  $\Delta$  some positive real-valued constant. Assume that we observe the process only at the discrete time points  $t = 0, \Delta, 2\Delta, \dots, T\Delta$ , and that exercise decisions are made immediately after each observation. (We would typically take the time unit here to be one year, and  $\Delta$  to be  $1/252$ , representing one trading day. However, both  $\Delta$  and the time units can be chosen arbitrarily subject to the constraints mentioned above.)

Assume that under the physical measure  $P$ , the asset price  $S(t)$  evolves according to the Itô stochastic differential equations

$$dS(t) = \mu S(t)dt + \sigma(Y(t))S(t)dW_1(t) \quad (1)$$

$$\sigma(Y(t)) = \exp(Y(t)) \quad (2)$$

$$dY(t) = \alpha[\beta - Y(t)]dt + \gamma dW_2(t), \quad (3)$$

where  $\mu$  is a constant referred to as the “drift”,  $\sigma(Y(t))$  is referred to as “volatility”, and  $\alpha, \beta$ , and  $\gamma$  are constants with  $\alpha > 0$ .  $\{W_1(t)\}$  and  $\{W_2(t)\}$  are assumed to be two independent standard Brownian motions. (The American option pricing problem and the valuation algorithm that we propose do not depend on the specific form of the stochastic volatility model. Indeed, the procedure we introduce later is easily adaptable to other types of stochastic volatility models as well.)

Under a risk-neutral measure, the asset price evolves according to the (similar) Itô stochastic differential equations

$$dS(t) = rS(t)dt + \sigma(Y(t))S(t)dW_3(t) \quad (4)$$

$$dY(t) = [\alpha(\beta - Y(t)) - \lambda\gamma]dt + \gamma dW_4(t), \quad (5)$$

where  $r$  represents the risk-free interest rate (measured in appropriate time units),  $\lambda$  is a constant referred to as the “market price of volatility risk” (Musielka and Rutkowski, 1998; Melenberg and Werker, 2001; Bakshi and Kapadia, 2003), and  $\{W_3(t)\}$  and  $\{W_4(t)\}$  are independent standard Brownian motions. (Note that  $\lambda$  is not uniquely determined in this system of SDEs. There is a range of possible risk-neutral measures for the model, each one having a different value of  $\lambda$ . In fact, there are also risk-neutral measures under which  $\lambda$  varies over time. However, for the sake of simplicity, we will assume that  $\lambda$  is a constant. Later in this paper, we discuss one way to determine this constant.)

Since our observations occur at discrete time-points  $0, \Delta, 2\Delta, \dots$ , we will make extensive use of the discrete-time approximation to the solution to the risk-neutral stochastic differential equations (4, 5) given by

$$S_{t+1} = S_t \cdot \exp \left\{ (r - \sigma_t^2/2) \Delta + \sigma_t \sqrt{\Delta} Z_{t+1} \right\} \quad (6)$$

$$\sigma_t = \exp(Y_t) \quad (7)$$

$$Y_{t+1} = \beta^* + e^{-\alpha\Delta}(Y_t - \beta^*) + \gamma\xi_{t+1}, \quad (8)$$

where

$$\beta^* = \beta - \alpha^{-1}\lambda\gamma, \quad (9)$$

$\{Z_t\}$  is an independent and identically distributed (iid) sequence of random variables with standard normal ( $N(0,1)$ ) distributions,  $\{\xi_t\}$  is an iid sequence of random variables with normal distributions,  $\xi_t \sim N\left(0, \frac{1-e^{-2\alpha\Delta}}{2\alpha}\right)$ , and all other parameters are as defined previously. Thus  $S_t$  and  $Y_t$  represent approximations, respectively, to  $S(t\Delta)$  and  $Y(t\Delta)$ . (The expression for  $Y_t$  is obtained directly from the exact solution to (5), while the expression for  $S_t$  is the solution to (4) that one would obtain by regarding  $\sigma_t$  to be constant on successive intervals of length  $\Delta$ .) It will sometimes be convenient to express (6) in terms of the log-returns  $R_t = \log(S_{t+1}/S_t)$ , as

$$R_t = (r - \sigma_t^2/2)\Delta + \sigma_t \sqrt{\Delta} Z_{t+1}. \quad (10)$$

To complete the specification of the model, we can assign  $Y_0$  a normal distribution,

$$Y_0 \sim N(\beta^*, \gamma^2/(2\alpha)). \quad (11)$$

This is simply the stationary (limiting) distribution of the first-order autoregressive process  $\{Y_t\}$ . However, for practical purposes, it will usually be preferable to replace this distribution by the conditional distribution of  $Y_0$ , given some historical observed price data  $S_{-1}, S_{-2}, \dots$ .

## 2.2 The Price of an American-Style Option

The arbitrage-free price for an American-style option is

$$\sup_{\tau \in \mathcal{T}} E_{RN}[\exp(-r\tau)g(S_\tau)], \quad (12)$$

where  $\tau$  is a random stopping time at which an exercise decision is made,  $\mathcal{T}$  is the set of all possible stopping times with respect to the filtration  $\{\mathcal{F}_t, t = 0, 1, \dots\}$  defined by

$$\mathcal{F}_t = \sigma(S_0, \dots, S_t),$$

$E_{RN}(\cdot)$  represents the expectation under a risk-neutral probability measure, of its argument, and  $g(s)$  denotes the payoff from exercise of the option when the underlying asset price is equal to  $s$ . For example, a call option with strike price  $K$  has payoff function  $g(s) = \max(s - K, 0)$ , and a put option with strike price  $K$  has payoff function  $g(s) = \max(K - s, 0)$ . Since  $\tau$  is a stopping time, the event  $\{\tau \leq t\}$  must be  $\mathcal{F}_t$ -measurable, or equivalently, the decision to exercise or hold at a given time must be made only based on observations of the previous and current values of the underlying price process. To allow for the possibility that the option is never exercised, we adopt the convention that  $\tau = \infty$  if the option is not exercised at or before expiry, along with the convention that  $\exp(-r\infty)g(S_\infty) = 0$ .

## 3 Dynamic Programming

The pricing problem for the American option is to find the stopping time  $\tau$  at which the supremum in (12) is achieved. While it is not immediately obvious how one might search through the space

of all possible stopping times, this problem is equivalent to a stochastic control problem, which can be solved (in theory) using the dynamic programming algorithm.

### 3.1 General Method

The dynamic programming algorithm (originally developed by Bellman (1953) and discussed, for example, in the context of financial analysis by Glasserman (2004)) constructs the exact optimal decision functions recursively, working its way from the terminal decision point (at time  $T\Delta$ ) back to the first possible decision point (at time 0). In addition, the procedure yields the expectation in the expression (12), which is our desired option price. The algorithm works as follows.

Let  $d_t \in \{E, H\}$  denote the decision made immediately after observation of  $S_t$ , either to exercise ( $E$ ) or hold ( $H$ ) the option. While either decision could be made, only one is optimal, given available information up to time  $t$ . We denote the optimal decision, as a function of the available observations, by

$$d_t^*(s_0, \dots, s_t) \in \{E, H\}.$$

Here and in the remainder of the paper, we adopt the usual convention of using  $S_j$  (upper case) to denote the random variable representing the equity price at time  $j$ , and  $s_j$  (lower case) to denote a particular possible realization of the random variable. Next, let

$$u_T(s_0, \dots, s_T, d_T) = \begin{cases} \exp(-r\Delta T)g(s_T), & d_T = E, \\ 0, & d_T = H, \end{cases} \quad (13)$$

and for  $t = 0, 1, \dots, T - 1$ , let  $u_t(s_0, \dots, s_t, d_t)$  denote the discounted expected payoff of the option at time  $(t\Delta)$ , assuming that decision  $d_t$  is made, and *also assuming* that optimal decisions are made at times  $(t + 1)\Delta, (t + 2)\Delta, \dots, T\Delta$ . It is obvious that at the expiration time,

$$d_T^*(s_0, \dots, s_T) = \arg \max_{d_T \in \{E, H\}} u_T(s_0, \dots, s_T, d_T)$$

The optimal decision functions  $d_{T-1}^*, \dots, d_0^*$  can then be obtained by defining

$$u_t^*(s_0, \dots, s_t) = u_t(s_0, \dots, s_t, d_t^*(s_0, \dots, s_t)), \quad t = 0, \dots, T, \quad (14)$$

and using the recursions

$$u_t(s_0, \dots, s_t, d_t) = \begin{cases} \exp(-r\Delta t)g(s_t), & d_t = E, \\ E_{RN}(u_{t+1}^*(S_0, \dots, S_{t+1})|S_0 = s_0, \dots, S_t = s_t), & d_t = H, \end{cases} \quad (15)$$

$$d_t^*(s_0, \dots, s_t) = \arg \max_{d_t \in \{E, H\}} u_t(s_0, \dots, s_t, d_t). \quad (16)$$

These recursions are used sequentially, for  $t = T - 1, T - 2, \dots, 0$ , and yield the (exact) optimal decision functions  $d_t^*$ ,  $t = 0, \dots, T$ . (Each  $d_t$  is optimal in the space of all possible functions of historical data  $s_0, \dots, s_t$ .) The corresponding stopping time  $\tau$  is simply

$$\tau = \min(\{t \in \{0, \dots, T\} | d_T = E\} \cup \{\infty\}),$$

Furthermore, the procedure also gives the risk-neutral option price, since

$$u_0^*(s_0) = \sup_{\tau \in T} E_{RN}[\exp(-r\tau)g(S_\tau)].$$

In practice, it is generally not possible to compute the optimal decision functions, since each  $d_t^*$  needs to be computed and stored for all (infinitely many) possible combinations of values of its arguments  $s_0, \dots, s_t$ . However, in what follows we will develop an approach which gives high-quality approximations to the exact solution. The approach relies on exploiting two key features of the American-style option pricing problem.

### 3.1.1 Equivalent Formulation of the Dynamic Programming Problem

First, it follows from the the Markov property of the bivariate process  $\{(S_t, Y_t)\}$  that we can transform the arguments of the decision functions so that they do not increase in number as  $t$  increases.

Let us define

$$\pi_t(y_t)dt = P(Y_t \in dy_t | S_0 = s_0, \dots, S_t = s_t), \quad t = 0, \dots, T, \quad (17)$$

where  $s_j$  denotes the observed value of  $S_j$ , so that  $\pi_t(\cdot)$  denotes the conditional density of the distribution of  $Y_t$  (with respect to Lebesgue measure), given historical information  $S_0 = s_0, \dots, S_t = s_t$ .

Then we have the following result.

**Lemma 3.1.** For each  $t = 0, \dots, T$ ,  $u_t(s_0, \dots, s_t, d_t)$  can be expressed as a functional of only  $s_t$ ,  $\pi_t$ , and  $d_t$ , that is,

$$u_t(s_0, \dots, s_t, d_t) = \tilde{u}_t(s_t, \pi_t, d_t).$$

Consequently for each  $t = 0, \dots, T$ ,  $d_t^*(s_0, \dots, s_t)$  and  $u_t^*(s_0, \dots, s_t)$  can also be expressed, respectively, as functionals

$$\begin{aligned} d_t^*(s_0, \dots, s_t) &= \tilde{d}_t^*(s_t, \pi_t) \\ u_t^*(s_0, \dots, s_t) &= \tilde{u}_t^*(s_t, \pi_t). \end{aligned}$$

A proof of this result is given in the appendix. It is important to note here that the argument  $\pi_t$  to the functions  $\tilde{u}_t$ ,  $\tilde{d}_t^*$ , and  $\tilde{u}_t^*$  is a function itself.

Lemma 3.1 states that each optimal decision function can be expressed as functional depending only on the current price  $s_t$  and the conditional distribution of the latent log-volatility  $Y_t$ , given observations of prices  $s_0, \dots, s_t$ . In other words, we can write the exact equivalent form of the dynamic programming recursions,

$$\tilde{u}_T(s_T, \pi_T, d_T) = \begin{cases} \exp(-r\Delta T)g(s_T), & d_T = E, \\ 0, & d_T = H, \end{cases} \quad (18)$$

$$\tilde{u}_t(s_t, \pi_t, d_t) = \begin{cases} \exp(-r\Delta t)g(s_t), & d_t = E, \\ E_{RN}[\tilde{u}_{t+1}^*(S_{t+1}, \pi_{t+1}) | S_t = s_t, \pi_t], & d_t = H, \end{cases} \quad (19)$$

where

$$\tilde{d}_t^*(s_t, \pi_t) = \arg \max_{d_t \in \{E, H\}} \tilde{u}_t(s_t, \pi_t, d_t).$$

and

$$\tilde{u}_t^*(s_t, \pi_t) = \tilde{u}_t(s_t, \pi_t, \tilde{d}_t^*(s_t, \pi_t)).$$

### 3.1.2 Sequential Monte Carlo and Log-Normal Approximations

The second feature we exploit is the empirical observation that the distributions  $\pi_t$  (recall the definition (17)) can be well-approximated by normal distributions. This means that we can use just

two parameters to represent a good approximation to each density functions  $\pi_t$ .

To see this, we used a recently-developed method known as “sequential Monte Carlo” simulation (also known as “particle filtering”, see, e.g. Doucet et al., 2001, for discussion, analysis and examples of these algorithms) to obtain approximate samples from the distributions  $\pi_t$ . A full treatment of these methods is beyond the scope of this paper, but the most basic form of the algorithm is as follows.

**Algorithm 3.1: Sequential Monte Carlo Estimation of  $\pi_0, \dots, \pi_T$**

**Initialization.** Choose a number of “particles”  $m > 0$ . (Typically  $m$  would be around 1000.) Set  $t = 0$ . Draw a sample  $\{y_0^{(1)}, \dots, y_0^{(m)}\}$  from the distribution of  $Y_0$  (see equation (11)).

**Step 1.** Compute the weights

$$w_t^{(i)} = p(r_t | Y_t = y_t^{(i)}), \quad (20)$$

where the term on the right denotes the conditional density of the log-return  $R_t$ , given  $Y_t = y_t^{(i)}$ , evaluated at the observed value  $r_t$ . (These weights are trivially obtained from (10).)

**Step 2.** Draw a new sample  $\{\tilde{y}_t^{(1)}, \dots, \tilde{y}_t^{(m)}\}$  by sampling with replacement from  $\{y_t^{(1)}, \dots, y_t^{(m)}\}$ , with probabilities proportional to  $w_t^{(1)}, \dots, w_t^{(m)}$ .

**Step 3.** For  $i = 1, 2, \dots, m$ , draw  $y_{t+1}^{(i)}$  from the distribution  $p(y_{t+1} | Y_t = \tilde{y}_t^{(i)})$ . (This distribution is Gaussian, specified by (8).)

**Repetition.** Replace  $t$  by  $t + 1$ . If  $t < T$  go back to Step 1, otherwise terminate the procedure.

This algorithm yields  $T + 1$  collections of particles,  $\{y_t^{(1)}, \dots, y_t^{(m)}\}$ ,  $t = 0, 1, \dots, T$ , with the property that for each  $t$ ,  $\{y_t^{(1)}, \dots, y_t^{(m)}\}$  can be regarded as an approximate sample of size  $m$  from the distribution  $\pi_t$ . The algorithm has the convenient property that as  $m$  increases, the empirical

distributions of the particle collections converge to the desired distributions  $\pi_t$ . (For more details and various refinements of this algorithm, refer to, among others, Kitagawa and Sato, 2001; Liu and West, 2001; Pitt and Shephard, 1999).

To investigate the form of the distributions  $\pi_t$  for real data under the stochastic volatility model (6, 8), we obtained daily closing share prices for Xerox from January 1st, 2002 to December 31st, 2003, from the Wharton Research Data Services website<sup>1</sup>.



Figure 1: Daily closing prices of Xerox stock, for all trading days during the years 2002 and 2003.

The prices over this time period are plotted in Figure 1. We then fit a stochastic volatility model of the form (1, 2, 3). We fit the model parameters under the objective measure by maximum likelihood estimation<sup>2</sup>, obtaining

$$\mu = 0.296, \quad \alpha = 17.723, \quad \beta = -0.828, \quad \text{and} \quad \gamma = 2.932. \quad (21)$$

Then we implemented the sequential Monte Carlo algorithm described above with  $m = 10,000$  particles, to estimate the posterior distributions  $\pi_t = p(Y_t | S_1 = s_1, \dots, S_t = s_t)$ , for each possible value of  $t$ , under a risk-neutral measure, with  $\lambda = 0$  and  $r = 0.01$ .

<sup>1</sup>Access to the Wharton Research Data Services website was obtained through the Tepper School of Business at Carnegie Mellon University.

<sup>2</sup>Maximum likelihood parameter estimation was performed using the software package Cronos ([www.stat.cmu.edu/~abrock/cronos](http://www.stat.cmu.edu/~abrock/cronos)).

Figure 2 shows four of the resulting distributional approximations, at time points  $t$  corresponding to Jul. 1st, 2002, Dec. 27th, 2002, Jul. 27th, 2003, and Dec. 24th, 2003. Each plot shows a kernel-smoothed density estimate based on the appropriate collection of particles  $\{y_t^{(1)}, \dots, y_t^{(10000)}\}$ , along with a superimposed line showing the density of the Gaussian approximation obtained by matching the sample mean and variance of the collection of particles. Clearly, the distribution  $\pi_t$  is very close to Gaussian in each of the four cases. These results are typical; we observed similar near-Gaussian distributions for all other values of  $t$ , using the Xerox data, and also using several other share price time series. Note that the choice of  $\lambda$  will not affect these results since  $\lambda$  only affects the drift of the log-volatility process.

In light of these empirical results, we will represent each  $\pi_t$  using only two parameters: the mean and standard deviation, respectively, of a Gaussian distributional approximation,

$$\mu_t = \int x\pi_t(x)dx, \quad \zeta_t = \sqrt{\int x^2\pi_t(x)dx - \mu_t^2}. \quad (22)$$

### 3.2 The Pricing Algorithm

Now we are in a position to use the observations of Subsections 3.1.1 and 3.1.2 to construct a formal pricing algorithm. Combining the equivalent form of the dynamic programming recursions (18,19) with the property that  $\pi_t$  can be well-approximated by a Gaussian distribution with mean  $\mu_t$  and variance  $\zeta_t^2$ , we can approximate the dynamic programming recursions by

$$\hat{u}_T(s_T, \mu_T, \zeta_T, d_T) = \begin{cases} \exp(-r\Delta T)g(s_T), & d_T = E, \\ 0, & d_T = H, \end{cases} \quad (23)$$

$$\hat{u}_t(s_t, \mu_t, \zeta_t, d_t) = \begin{cases} \exp(-r\Delta t)g(s_t), & d_t = E, \\ E_{RN}[\hat{u}_{t+1}^*(S_{t+1}, \mu_{t+1}, \zeta_{t+1})|S_t = s_t, \mu_t, \zeta_t], & d_t = H, \end{cases} \quad (24)$$

with

$$\hat{d}_t^*(s_t, \mu_t, \zeta_t) = \arg \max_{d_t \in \{E, H\}} \hat{u}_t(s_t, \mu_t, \zeta_t, d_t). \quad (25)$$

and

$$\hat{u}_t^*(s_t, \mu_t, \zeta_t) = \hat{u}_t(s_t, \mu_t, \zeta_t, \hat{d}_t^*(s_t, \mu_t, \zeta_t)), \quad (26)$$

where  $\mu_t$  and  $\zeta_t$  are as defined in (22). The recursion (24) is particularly convenient because the required expectation can be computed using the core of the sequential Monte Carlo update algorithm. Since the steps of the sequential Monte Carlo algorithm are designed to make the transition from a specified  $\pi_t$  to the corresponding distribution  $\pi_{t+1}$ , we can combine a standard Monte Carlo simulation approach with the use of Steps 1, 2, and 3 of Algorithm 3.1 to compute the expectation. To be more specific, the following procedure computes the expectation on the right side of equation (24).

**Algorithm 3.2: Estimation of Conditional Expectation**

For fixed  $s_t, \mu_t$  and  $\zeta_t$ ,

$$\begin{aligned} E_{RN}[\hat{u}_{t+1}^*(S_{t+1}, \mu_{t+1}, \zeta_{t+1}) | S_t = s_t, \mu_t, \zeta_t] \\ \simeq \frac{1}{n} \sum_{j=1}^n \hat{u}_{t+1}^*(s_{t+1}^{(j)}, m^{(j)}, z^{(j)}), \end{aligned} \quad (27)$$

where to get  $\{(s_{t+1}^{(j)}, m^{(j)}, z^{(j)}), i = 1, \dots, n\}$ , we use the following procedure.

1. **Step 1.** Draw values  $\{y_t^{(i)}, i = 1, \dots, m\}$  independently from a  $N(\mu_t, \zeta_t^2)$  distribution.
2. **Step 2.** For each  $j = 1, \dots, n$ ,
  - (a) Draw  $s_{t+1}^{(j)}$  from the conditional distribution of  $S_{t+1}$ , given  $S_t = s_t$  and  $Y_t \sim N(\mu_t, \zeta_t^2)$ .
  - (b) Go through Steps 1 and 2 of Algorithm 3.1, but in computing weights  $\{w_t^{(i)}, i = 1, \dots, m\}$ , replace the actual log-return  $r_t$  by (the simulated log-return)  $r_t^{(i)} = \log(s_{t+1}^{(j)}/s_t)$ .
  - (c) Go through Step 3 of Algorithm 3.1, to obtain  $\{y_{t+1}^{(1)}, \dots, y_{t+1}^{(m)}\}$ .
  - (d) Compute  $m^{(j)} = \frac{1}{m} \sum_{i=1}^m y_{t+1}^{(i)}$  and  $z^{(j)} = \sqrt{\frac{1}{m} \sum_{i=1}^m y_{t+1}^{(i)2} - m^{(j)2}}$ .

Algorithm 3.2 works by drawing triples  $(s_{t+1}^{(i)}, m^{(i)}, z^{(i)})$  from the conditional distribution of  $(S_{t+1}, \mu_{t+1}, \zeta_{t+1})$ , given  $S_t = s_t, \mu_t$ , and  $\zeta_t$ , and using these to compute a Monte Carlo estima-

tor of the required conditional expectation. This allows us to evaluate  $\hat{u}_t$  at various points, given knowledge of  $\hat{u}_{t+1}^*$ , and will thus form a key component of the backward induction step in our final algorithm.

Since we will be interested in storing the functions  $\hat{u}_t^*(\cdot, \cdot, \cdot)$  and  $\hat{d}_t^*(\cdot, \cdot, \cdot)$ , we next introduce some additional notation. Let

$$\mathcal{G} = \{g_i \in \mathbb{R}^3, i = 1, 2, \dots, G\} \quad (28)$$

denote a collection of grid points in  $\mathbb{R}^3$ . These are points at which we will evaluate and store the functions  $\hat{u}_t^*$  and  $\hat{d}_t^*$ . We will typically take

$$\mathcal{G} = \mathcal{G}_1 \times \mathcal{G}_2 \times \mathcal{G}_3, \quad (29)$$

where

$$\begin{aligned} \mathcal{G}_1 &= \{\underline{s}, \underline{s} + \delta_s, \underline{s} + 2\delta_s, \dots, \bar{s}\}, \\ \mathcal{G}_2 &= \{\underline{m}, \underline{m} + \delta_m, \dots, \bar{m}\}, \text{ and} \\ \mathcal{G}_3 &= \{\underline{z}, \underline{z} + \delta_z, \dots, \bar{z}\}. \end{aligned}$$

Here  $\underline{s}$ ,  $\underline{m}$ , and  $\underline{z}$  are chosen to be lower bounds for likely possible values of  $S_t$ ,  $\mu_t$  and  $\zeta_t$ , respectively, while  $\bar{s}$ ,  $\bar{m}$ , and  $\bar{z}$  are corresponding upper bounds. The terms  $\delta_s$ ,  $\delta_m$ , and  $\delta_z$  determine grid spacing, with smaller values leading to a finer grid over the specified range.

We are now in a position to state the pricing algorithm itself. (This is a standard gridding approach to solving the dynamic programming equations, as described for instance in Brockwell and Kadane, 2003).

### Algorithm 3.3: American-Style Option Pricing Algorithm

**Initialization.** For each  $g \in \mathcal{G}$ , evaluate

$$\hat{u}_T(g, d_T), \quad \hat{d}_T^*(g), \quad \text{and} \quad \hat{u}_T^*(g),$$

using equations (23), (25), and (26). Store the results.

Set  $t = T - 1$ .

**Step 1.** For each  $g \in \mathcal{G}$ , evaluate

$$\hat{u}_t(g, d_t), \quad \hat{d}_t^*(g), \quad \text{and} \quad \hat{u}_t^*(g),$$

using equations (24), (25), and (26). To evaluate the expectations in equation (24), use Algorithm 3.2. Store the results.

**Step 2.** Replace  $t$  by  $t - 1$ . If  $t \geq 0$ , go back to Step 1.

**Step 3.** Evaluate the option price

$$\text{price} = \hat{u}_0^*(s_0, \beta^*, \gamma^2/(2\alpha)), \quad (30)$$

where  $\beta^*, \gamma$  and  $\alpha$  are the constants in the risk-neutral discrete-time stochastic volatility model (6, 8, 9).

*Remark:* This scheme gives a price which assumes that no information about volatility is available at time  $t = 0$ . In the absence of such information, we just assume that the initial log-volatility  $Y_0$  can be modeled as coming from the limiting distribution of the autoregressive process  $\{Y_t\}$ . However, in most cases, it is possible to estimate log-volatility at time  $t = 0$  using previous observations of the price process  $\{S_t\}$ . In such cases, we would replace equation (30) by

$$\text{price} = \hat{u}_0^*(s_0, e_1, e_2), \quad (31)$$

where  $e_1$  and  $e_2$  denote, respectively, the conditional mean and conditional variance of  $Y_0$ , given “previous” observations  $S_{-1}, S_{-2}, \dots, S_{-h}$  for some  $h > 0$ . These could be obtained in a straight-

forward manner by making use of the sequential Monte Carlo estimation procedure described in Algorithm 3.1 (appropriately modified so that time  $-h$  becomes time 0).

*Remark:* As with any quadrature-type approach, grid ranges must be chosen with some care. In order to preserve quality of approximations to the required expectations in Algorithm 3.3, it is necessary for the intervals  $[\underline{s}, \bar{s}]$ ,  $[\underline{m}, \bar{m}]$ ,  $[\underline{z}, \bar{z}]$  to contain observed values of  $S_t$ ,  $\mu_t$ , and  $\zeta_t$  with probability close to one, over the lifetime of the option. Once the stochastic volatility model has been fit, this is easily done. For instance, one could choose

$$\begin{aligned}(\underline{s}, \bar{s}) &= (0, S_0 \exp[k e^{\beta^* + \gamma^2 / (4\alpha)} \cdot T \Delta]), \\(\underline{m}, \bar{m}) &= \beta^* \pm k\gamma / \sqrt{2\alpha}, \text{ and} \\(\underline{z}, \bar{z}) &= (0, k\gamma / \sqrt{2\alpha}),\end{aligned}$$

where  $k$  is some constant determining how far we want to allow for various parameters to drift away from their respective means, which could be chosen to be, say, around 5.0 or more. Grid spacing parameters  $\delta_s$ ,  $\delta_m$  and  $\delta_z$  should be chosen to be as small as possible, subject to computing time constraints.

*Remark:* In Step 1 of the algorithm, evaluation of  $\hat{u}_t(g, d_t)$  is performed making use of the Monte Carlo approximation given by (27). Obviously the expression relies on knowledge of  $\hat{u}_{t+1}^*(\cdot, \cdot, \cdot)$ , but since we have only evaluated  $\hat{u}_{t+1}^*$  at grid points  $g \in \mathcal{G}$ , it is necessary to interpolate in some manner. Strictly speaking, one could simply choose the nearest grid point, and rely on sufficient grid density (i.e. small values of  $\delta_s$ ,  $\delta_m$ , and  $\delta_z$ ) to control error. However, inspection of the surface suggests that local linear approximations are more appropriate. Therefore in our implementations, we use linear interpolation between grid points.

## 4 Evaluation of Observed Market Option Prices

To investigate the performance of Algorithm 3.3, we implemented it in an analysis of both Xerox and Dell put option prices.

In addition to the daily closing Xerox share prices during 2002 and 2003, to which we fit the stochastic volatility model (1,2,3) with parameters specified in (21), we obtained prices for American-style put options on these Xerox shares, during the first 20 trading days of 2004. We used a total of 60 options in this trading period. (We obtained these data from the website of the American Stock Exchange (AMEX).) The options all have a strike price equal to \$14.00 and maturities that range from 15 to 135 days. We also obtained daily data on the Federal Fund short rate to use in our pricing algorithm. For the period of the option valuation, this rate was approximately 1.0%, leading us to choose  $r = 0.01$  in our pricing algorithm.

Using a similar kind of approach to that of Eraker (2004), we determined the market price of volatility risk  $\lambda$  by pricing with a range of different values, and choosing the value that minimized the absolute value of the sum of the differences between our resulting 60 option prices and the corresponding market prices. For Xerox, using this method, we chose  $\lambda = 2.9$ . Figure 3 plots prices obtained using Algorithm 3.3 with  $\lambda = 2.9$ , against observed market prices. For the sake of comparison, it also shows the prices obtained using Algorithm 3.3 with  $\lambda = 0$ . The data is strongly suggestive of a non-zero market price of volatility risk in the market.

We also compared Xerox option prices using Algorithm 3.3 with  $\lambda = 2.9$  to those obtained using the least-squares Monte Carlo (LSM) Algorithm of Longstaff and Schwartz (2001). For the LSM Algorithm, we used a constant volatility model, with volatility obtained from the sample variance of log-returns during the years 2002-2003 (that is, the same data we used in estimating stochastic volatility model parameters). The results, shown in Figure 4, illustrate that our model generated option prices are more consistent with observed market data than those obtained using the LSM Algorithm. This is of course in part due to the ability to explicitly incorporate  $\lambda$  into the pricing procedure in Algorithm 3.3, but even beyond this, we appear to obtain slightly reduced variation of option prices around the approximate straight lines apparent in Figures 3 and 4.

We repeated the same analysis with Dell share prices and option prices from the first 20 trading days of 2004. There were a total of 120 options in this data set, which was also obtained from the AMEX website. The put options had strike prices of \$32.50, \$35.00, and \$37.50. The maturities

of the options ranged from 15 to 98 days. The results are shown in Figure 5 and 6. Interestingly, for Dell, the value of  $\lambda$  most consistent with observed prices was 2.6, which is close to the value obtained for Xerox puts. Qualitatively, the results are essentially the same as those for Xerox; it appears that market price of volatility risk is non-zero, and we are able to obtain prices somewhat more consistent with market prices than those obtained using the LSM Algorithm.

## 5 Discussion

We have introduced an algorithm for pricing American-style options under stochastic volatility models. The approach is based on (1) the empirical observation that conditional distributions  $\pi_t$  are approximately normal, (2) the use of a sequential Monte Carlo step to obtain and update the distributions  $\pi_t$ , and (3) a gridding (quadrature)-type approach to solving the associated dynamic programming problem. We have also shown that with appropriately chosen market-price of volatility risk (i.e. appropriately chosen risk-neutral measure), we obtain option prices that are consistent with market data.

The method is computationally intensive, but can be implemented using parallel processing. In particular, the computations for grid points in Step 1 can be divided between multiple CPUs, and evaluated simultaneously in several different blocks. This leads to a speed increase approximately proportional to the number of processors used. In the examples given in this paper, using a grid size ( $\#(\mathcal{G})$ ) of 300,000, along with parameters  $m = 200$  and  $n = 200$  in Algorithm 3.2, and  $T=135$ , for a single option, the algorithm took approximately 2 hours to complete on a Beowulf cluster of 32 dual-CPU Opteron 250 nodes.

Furthermore, in light of the analyses performed in Section 4, a potentially interesting line of future work could be to analyze market price of volatility risk more carefully. In the two shares we considered, we obtained values of  $\lambda$  equal to 2.6 and 2.9, but one could easily ask if these values apply universally across different shares and market sectors, and also whether or not they vary over

time. Investigation of these questions could be carried out by posing a joint probability model for a time and/or share-dependent  $\lambda$ , as well as appropriate share prices (in the spirit of the approach used by Eraker, 2004), and model likelihood could then be evaluated using the algorithm introduced in this paper within the computation.

## 6 Acknowledgements

The authors are grateful to John Lehoczky, Daniel Peris, Mark Schervish, and Duane Seppi for a number of useful conversations related to the work in this paper. This work was supported in part by equipment purchased with NSF SCREMS Grants DMS-0532083 and DMS-0422400. Ricky Rambharat gratefully acknowledges the assistance of Robert Wolpert and Merlise Clyde for granting access to the Duke Shared Cluster Resource at the Center for Computational Science, Engineering and Medicine.

## A Proofs

*Proof of Lemma 3.1.* We use an inductive argument. To begin with,  $u_T(s_0, \dots, s_T, d_T)$  is by its definition (13) obviously a function of  $s_T$  and  $d_T$ , and thus is trivially a functional of  $s_T, \pi_T$  and  $d_T$ , which we can denote by  $\tilde{u}_T(s_T, \pi_T, d_T)$ .

Next, suppose that for some  $t$ , we can write  $u_{t+1}(s_0, \dots, s_{t+1}, d_{t+1}) = \tilde{u}_{t+1}(s_{t+1}, \pi_{t+1}, d_{t+1})$ . Then from (15),

$$u_t(s_0, \dots, s_t, E) = \exp(-r\Delta t)g(s_t), \tag{32}$$

and

$$\begin{aligned} u_t(s_0, \dots, s_t, H) &= E_{RN}(u_{t+1}^*(s_0, \dots, s_t, S_{t+1}) | S_0 = s_0, \dots, S_t = s_t) \\ &= E_{RN}(\tilde{u}_{t+1}^*(S_{t+1}, \pi_{t+1}) | S_0 = s_0, \dots, S_t = s_t) \end{aligned} \quad (33)$$

$$= \int E_{RN}(\tilde{u}_{t+1}^*(S_{t+1}, \pi_{t+1}) | S_0 = s_0, \dots, S_t = s_t, Y_t = y_t) \pi_t(y_t) dt \quad (34)$$

$$= \int E_{RN}(\tilde{u}_{t+1}^*(S_{t+1}, \pi_{t+1}) | S_t = s_t, Y_t = y_t) \pi_t(y_t) dt. \quad (35)$$

Equation (34) is obtained from (33) using a simple conditioning argument, and (35) then follows since  $\{(S_t, Y_t), t = 0, 1, \dots\}$  is a (bivariate) Markov process. The expression in (32) is obviously a function of  $s_t$ , and since  $s_t$  and  $\pi_t$  completely determine the distribution of the arguments  $S_{t+1}$  and  $\pi_{t+1}$  to the function  $\tilde{u}_{t+1}^*(\cdot, \cdot)$  in (35), it is also clear that the expression in (35) is a functional of  $s_t$  and  $\pi_t$ . Thus  $u_t(s_0, \dots, s_t, d_t)$  is a functional of  $s_t$ ,  $\pi_t$  and  $d_t$ , which we denote by  $\tilde{u}_t(s_t, \pi_t, d_t)$ .

Invoking this inductive step for  $t = T - 1, T - 2, \dots, 0$  gives the first part of the desired result. The second part of the result follows directly from the first part, along with the definitions (14) and (16).  $\square$

## References

- Bakshi, G. and N. Kapadia (2003). Delta-Hedged Gains and the Negative Market Volatility Risk Premium. *The Review of Financial Studies* 16(2), 527–566.
- Bellman, R. (1953). *An Introduction to the Theory of Dynamic Programming*. Rand Corporation.
- Black, F. and M. Scholes (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81, 637–659.
- Brennan, M. J. and E. S. Schwartz (1977). The Valuation of American Put Options. *The Journal of Finance* 32(2), 449–462.
- Broadie, M. and P. Glasserman (1997). Pricing American-Style Securities using Simulation. *Journal of Economic Dynamics and Control* 21, 1323–1352.

- Brockwell, A. and J. Kadane (2003). A gridding method for Bayesian sequential decision problems. *Journal of Computational and Graphical Statistics* 12(3), 566–584.
- Carr, P., R. Jarrow, and R. Myneni (1992). Alternative Characterizations of American Put Options. *Mathematical Finance* 2(2), 87–106.
- Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative models for stock price dynamics. *Journal of Econometrics* 116, 225–257.
- Chernov, M. and E. Ghysels (2000). A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation. *Journal of Financial Economics* 56, 407–458.
- Clarke, N. and K. Parrott (1999). Multigrid for American Option Pricing with Stochastic Volatility. *Applied Mathematical Finance* 6, 177–195.
- Doucet, A., N. de Freitas, and N. Gordon (2001). *Sequential Monte Carlo Methods in Practice* (first ed.). New York: Springer.
- Eraker, B. (2004). Do stock prices and volatility jump? reconciling evidence from spot and option prices. *The Journal of Finance* LIX(3), 1367–1404.
- Finucane, T. J. and M. J. Tomas (1997). American Stochastic Volatility Call Option Pricing: A Lattice Based Approach. *Review of Derivatives Research* 1, 183–201.
- Fouque, J. P., G. Papanicolaou, and K. R. Sircar (2000). *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge: Cambridge University Press.
- Gallant, A., D. Hsieh, and G. Tauchen (1997). Estimation of stochastic volatility models with diagnostics. *Journal of Econometrics* 81, 159–192.
- Geske, R. and H. Johnson (1984). The American Put Option Valued Analytically. *The Journal of Finance* 39(5), 1511–1524.

- Ghysels, E., A. Harvey, and E. Renault (1996). Stochastic volatility. In *Handbook of Statistics*, Volume 14. Elsevier.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer.
- Guan, L. K. and G. Xiaoqiang (2000). Pricing American Options with Stochastic Volatility: Evidence from S&P 500 Futures Options. *The Journal of Futures Markets* 20(7), 625–659.
- Heston, S. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies* 6(2), 327–343.
- Hull, J. and A. White (1987). The Pricing of Options on Assets with Stochastic Volatility. *The Journal of Finance* 42(2), 281–300.
- Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics* 12(4), 371–389.
- Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies* 65, 361–393.
- Kitagawa, G. and S. Sato (2001). Monte carlo smoothing and self-organizing state-space model. In *Sequential Monte Carlo Methods in Practice*, pp. 177–196. Springer.
- Liu, J. and M. West (2001). Combined parameter and state estimation in simulation-based filtering. In *Sequential Monte Carlo Methods in Practice*, pp. 197–217. Springer.
- Longstaff, F. A. and E. S. Schwartz (2001). Valuing American Options by Simulation: A Simple Least-Squares Approach. *The Review of Financial Studies* 14(1), 113–147.
- Melenberg, B. and B. J. Werker (2001). The Implicit Price of Volatility Risk: An Empirical Analysis. Technical report, Tilburg University, The Netherlands.
- Musiela, M. and M. Rutkowski (1998). *Martingale Methods in Financial Modeling*. New York: Springer.

- Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics* 63, 3–50.
- Pitt, M. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* 94(46), 590–.
- Rogers, L. (2002). Monte Carlo Valuation of American Options. *Mathematical Finance* 12(3), 271–286.
- Stein, E. and C. Stein (1991). Stock Price Distributions with Stochastic Volatility: An Analytic Approach. *The Review of Financial Studies* 4(4), 727–752.
- Tzavalis, E. and S. Wang (2003). Pricing American Options under Stochastic Volatility: A New Method using Chebyshev Polynomials to Approximate the Early Exercise Boundary. Working paper, Queen Mary, University of London, London.

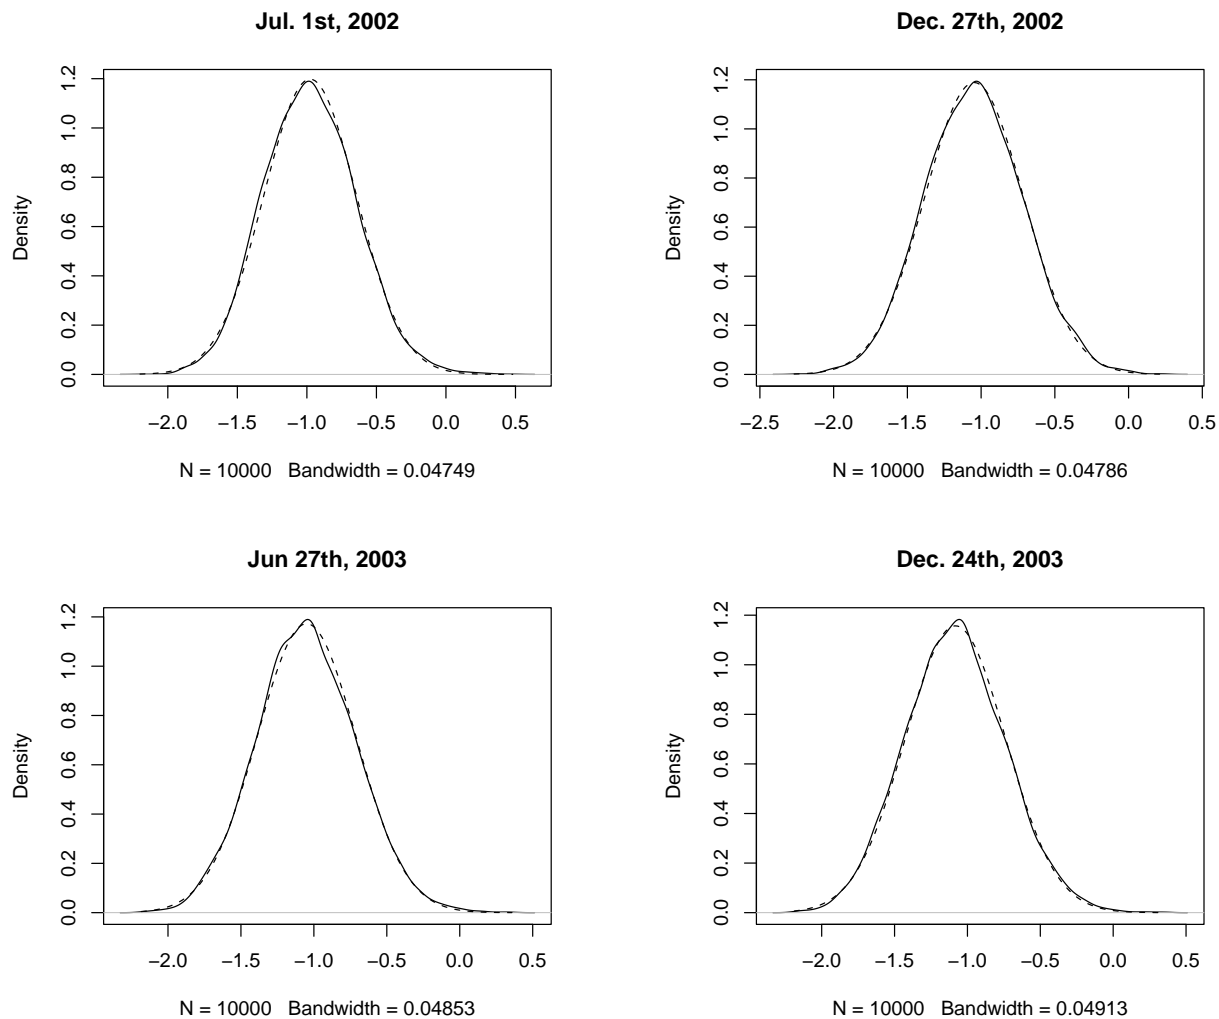


Figure 2: Comparison of conditional distributions  $\pi_t$  at four time points obtained using the particle filtering algorithm, and approximating Gaussian distributions, for the Xerox stock. The solid lines show kernel density estimates of  $\pi_t$  obtained using the sequential Monte Carlo algorithm with 10,000 particles. The dashed lines show approximating Gaussian distribution.




Figure 3: Comparison of market prices of put option prices for Xerox with those obtained using (a) Algorithm 3.3 with  $\lambda = 0$ , and (b) Algorithm 3.3 with  $\lambda = 2.9$ .

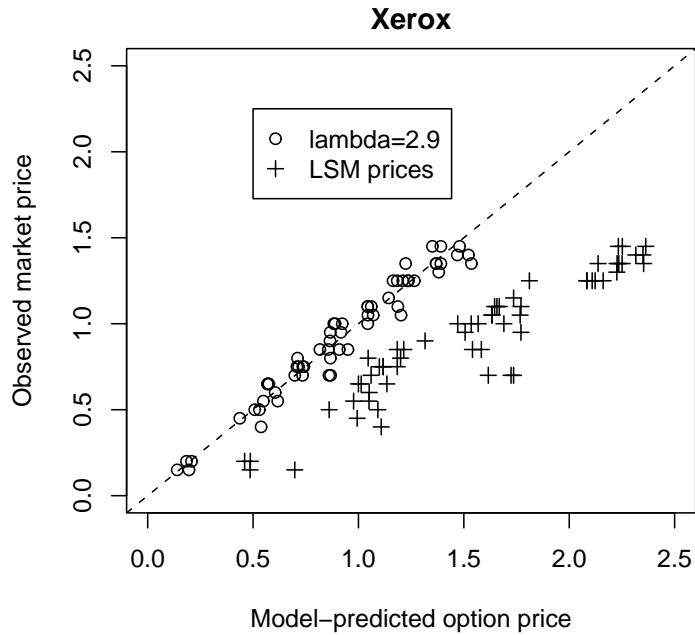


Figure 4: Comparison of market put option prices for Xerox and those obtained using (a) Algorithm 3.3 with  $\lambda = 2.9$ , and (b) Longstaff & Schwartz's Algorithm.

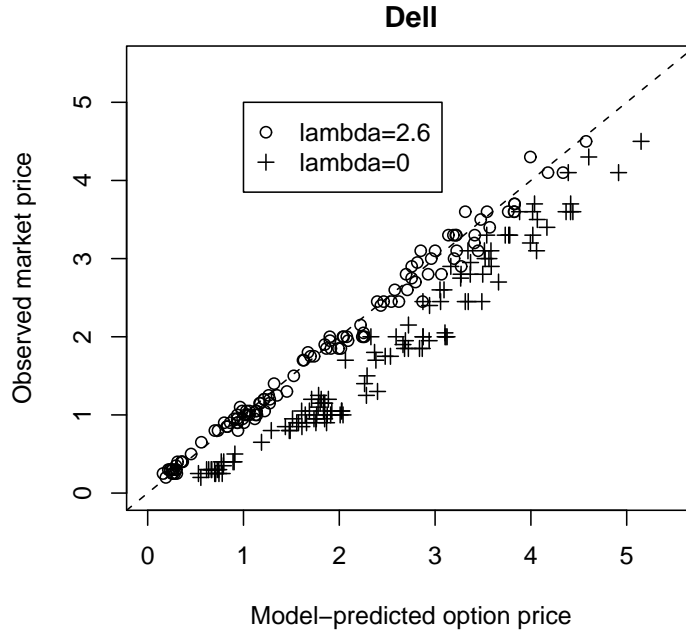


Figure 5: Comparison of market prices of put option prices for Dell with those obtained using (a) Algorithm 3.3 with  $\lambda = 0$ , and (b) Algorithm 3.3 with  $\lambda = 2.6$ .




Figure 6: Comparison of market put option prices for Dell and those obtained using (a) Algorithm 3.3 with  $\lambda = 2.6$ , and (b) Longstaff & Schwartz's Algorithm.