Stochastic volatility: option pricing using a multinomial recombining tree

Ionuţ Florescu[1,3] and Frederi G. Viens[2,4]

[1]Department of Mathematical Sciences, Stevens Institute of Technology,
Castle Point on the Hudson, Hoboken, NJ 07030
[2]Department of Statistics, Purdue University,
150 N. University St, West Lafayette, IN 47907-2067

[3]ifloresc@stevens.edu
[4]viens@purdue.edu

Abstract

In this article we treat the problem of option pricing when the volatility component of the underlying asset price is stochastic. The basic model we consider is commonly known as the Stochastic Volatility model: $dS_t = \mu S_t dt + \sigma(Y_t) S_t dW_t$, where Y_t is an exogenous mean-reverting-type process. We seek a discrete approximation of this model, one that will converge in distribution to the above continuous model, and that will allow us to calculate the option price numerically. First, we show how to estimate the distribution of the volatility component, using an interacting particle filtering algorithm due to Del Moral, Jacod and Protter (Del Moral et al., 2001). Then we use this distribution to construct two different models which converge to the solution of the original model. The first model is our main contribution to the research in Mathematical Finance, a recombining tree for the stock process, featuring four successors at every branch. The second model uses an Euler method to generate future stock prices and calculate option price.

We are in the incomplete market situation, and in order to price options on the stock, we use classical arbitrage-free valuation, combined with resampling and Monte-Carlo methods to generate versions of the stock price and compute expectations. Finally, we compare our methods with the classical Black-Scholes prices, using daily European Call Options on the SP500 index price in April 2004, and IBM in July 2005.

Key words and phrases: incomplete markets, Monte-Carlo method, option market, option pricing, particle method, random tree, stochastic filtering, stochastic volatility.

1 Introduction

We cannot start an article about option valuation without quoting the most celebrated articles in the domain (Black and Scholes, 1973) and (Merton, 1973). Despite significant development in the option pricing theory, the Black-Scholes formula for the European Call Option remains the most widely used application of stochastic analysis in Finance.

Nevertheless, the above quoted formula has significant biases, see for example (Rubinstein, 1985). The model's failure to describe the structure of reported option prices is thought to arise principally from its constant volatility assumption. Allowing the volatility to change over time means that it should generally be modeled as a stochastic process. However, accounting for such *stochastic volatility* within an option valuation formula is not an easy task. (Hull and White, 1987), (Chesney and Scott, 1989), (Stein and Stein, 1991), (Heston, 1993) all have constructed various specific stochastic volatility models for option pricing. A notable example of an attempt to find analytic formulas for option prices under stochastic volatility is (Fouque et al., 2000a). Even so, there are no simple formulas for the price of options on stochastic-volatility-driven stocks. When some means of implicit or explicit equations are found, the relations involved are cumbersome. Approximations have been constructed to these and other specific volatility models, e.g. (Hilliard and Schwartz, 1996), and (Ritchken and Trevor, 1999).

When the tree-based *Binomial* approximation was developed (Sharpe, 1978), the option pricing model became accessible to a wider audience. (Cox et al., 1979) constructed a binomial model that converged in distribution to the lognormal diffusion of Black-Scholes, and they also showed that the limit of the computed option value was the same as the one given by Black-Scholes valuation. Later, (Cox and Rubinstein, 1985) used the same approach to value the American style options on dividend paying stock, and they also relaxed some other assumptions of the original Black-Scholes model.

We believe that if one hopes to find any concrete results for stochastic volatility option pricing, one will have to revert to numerical methods. Inspired by the success of the binomial models, we seek a tree-based approximation. Our work shows how to achieve our goal using a *quadrinomial* tree model, one in which, at any time, a stock price's increment may take any one of four values.

For our underlying continuous-time stochastic process model, we assume that the price process S_t and the volatility driving process Y_t solve the equations:

$$\begin{cases} dS_t &= \mu S_t dt + \sigma(Y_t) S_t dW_t \\ dY_t &= \alpha(\nu - Y_t)dt + \psi(Y_t)dZ_t \end{cases} \tag{1.1}$$

This model spans all the stochastic volatility models considered previously for different specifications of the functions $\sigma(x)$ and $\psi(x)$. For simplicity, we assume W_t and Z_t are two *independent* Brownian motions. The case when they are correlated can be treated using ideas similar to those presented herein. We choose a mean-reverting type process to drive the volatility because this seems to be the most lucrative choice from the practical point of view, as noted for example in (Fouque et al., 2000a).

We shall note here that the model presented above implies that the market is incomplete, and as a consequence any derivative price calculated using this model will not be unique. We adopt a classical approach to this problem by fixing a derivative price as given apriori, thus in fact, completing the market. We present the details in Section 4

When trying to implement a tree approximation to price an option on a stock driven by this kind of model, one is faced with two problems: modeling the volatility component, and modeling the price itself. Modeling the volatility is a particularly difficult problem since the volatility cannot be observed directly from the market, only the stock price is directly observable.

This tree approximation approach has been tried in former articles, most notably (Leisen, 2000) who uses a binomial tree for the volatility and a so-called 8 successors tree for the price. His idea is similar with the one applied in (Nelson and Ramaswamy, 1990) for the case when the volatility is deterministic. However, that idea falls short of replicating the theoretical model when applied to Leisen's case, since the transformation used to eliminate the volatility does not work with stochastic volatility. Another interesting article is (Aingworth et al., 2003) where a Markov chain is used for the volatility process; unfortunately, the price tree therein is not recombining.

Section 2 contains the details of the conditions necessary for the convergence of our tree approximation to the solution of the model (1.1). The theory presented is not new, it is rather an adaptation of an older result about convergence of Markov Chains to the solution of diffusion equations. Notably, all previous articles on the subject only pointed to the referenced books (Stroock and Varadhan, 1979), and (Ethier and Kurtz, 1986) without checking if they were indeed in the specified case. Verifying this fact is not a trivial matter, thus we choose to include Section 2 to detail the specific application to our case.

In fact one of the main merits of this paper is the construction of a specific discrete Markov Chain process (our quadrinomial tree) that will converge in distribution to the solution of the given diffusion equation (1.1). The method we present herein is not restricted to the specified Ornstein-Uhlembeck model for the volatility process Y_t; indeed for any reasonable type of stochastic process the construction will proceed in a similar fashion. However, to be able to use the convergence result we need to construct a discrete distribution approximation of the volatility driving process Y_t in (1.1).

The method for estimating this volatility distribution uses a genetic-type algorithm introduced by Del Moral, Jacod, and Protter (Del Moral et al., 2001). This algorithm works with a fixed number of interacting particles, and gives a good approximation of the *optimal* estimation (in the sense of least squares) of the volatility given the observed stock prices. Hence it can be best described as an *optimal stochastic volatility particle filter*. Section 3 details this algorithm.

Based on this estimated volatility distribution we will construct two different models: a Static model in Section 4 (see page 9) and a Dynamic model in Section 5. Subsections 4.1 and 4.2 present the process of constructing a quadrinomial recombining tree which converges in distribution to the price process, and includes the stochastic volatility particle filter. Subsection 4.3 gives the proof of convergence of our quadrinomial tree to the solution of equation (1.1).

It is possible to prove – although we will leave a full proof out of our article for the sake of conciseness – that a tree with anything less than four basic successors could not possibly converge to the Markov process (1.1) it tries to approximate. This includes any binomial or trinomial tree construction. The proof includes the following idea: the main convergence theorem in Section 2 cannot be verified for smaller trees, because we simply do not have enough parameters in the model to verify all the equations involved. Another fact worth mentioning, which we also leave out in this paper, is that if the correlation between the processes W and Z in (1.1) is nonzero, our quadrinomial tree will not work anymore, but a pentanomial tree should be more than sufficient to handle it.

The Dynamic model is presented in Section 5 for the sake of comparison, it is what one would expect to have to construct in order to be consistent with the stochastic volatility model and the stochastic volatility particle filter estimation method. We present a standard Monte-Carlo method for pricing in this case, based on an Euler-type discretization of the governing system of stochastic differential equations (1.1). Our Static model differentiates itself from the Dynamic because, once the stochastic volatility has been estimated up to the present time[1], the volatility remains unchanged for the future construction of the model. Note however that this volatility is not a constant, but rather a random variable given as a function of all passed discrete stock observations.

Although the Static model may seem to be less consistent with the underlying stochastic processes driving our stock price in the future, we show in this article that it performs significantly better than the more mathematically natural Dynamic model in terms of option pricing, when performance is judged by comparison with prices actually observed on the option market. There are a number of ways of interpreting this mathematically counter-intuitive result. The easiest way is to note that any option pricing is done conditional on the filtration available today ($\mathcal{F}_0$). Since our constructed process (X_{t_i}, Y_{t_i}) is a Markov Chain the option price that we are calculating should only depend on the volatility distribution at time 0. Such an explanation is not a surprise when one knows how popular the straight constant-volatility Black-Scholes formula is among practitioners.

Section 6 contains numerical results obtained when applying our algorithms to SP500 and IBM option data. Section 7 contains the interpretation of the results, and directions of future study using our model.

2 The model and theoretical results

We work under an equivalent martingale measure, and instead of the stock price directly we work with the logarithm of the price (the return). We denote $X_t = \log S_t$. Under this measure the system of equations (1.1) becomes:

$$\begin{cases} dX_t &= \left(r - \frac{\sigma^2(Y_t)}{2}\right) dt + \sigma(Y_t) dW_t \\ dY_t &= \alpha(\nu - Y_t) dt + \psi(Y_t) dZ_t \end{cases} \tag{2.1}$$

Here r is the short-term risk-free rate of interest. We used the same notations W_t and Z_t for the corresponding Brownian Motions under the martingale measure obtained applying the Girsanov's theorem. We would like to obtain discrete versions of the processes (X_t, Y_t) which would converge in distribution to the continuous processes (2.1). Using the fact that e^x is a continuous function, and that the price of the European Option can be written

[1]the time at which the option pricing question is being asked

as a conditional expectation of a continuous function of the price, this is enough for the convergence in distribution of the option price found with our discrete approximation to the real price of the option.

To achieve this goal, we construct a discrete time Markov chain (our Static model) in Section 4, and using the theory in Section 11.3 in the book (Stroock and Varadhan, 1979), we show the convergence in distribution of this Markov chain to the solution of the diffusion equation (2.1). The same theory can also be found in (Ethier and Kurtz, 1986), though in a slightly less general form.

We choose to present details of the Theorems from the above cited books, since they require certain modifications in order to be applicable to our case. Once we present Theorem 2.2 in our specific context, we only need to verify it to obtain convergence of our Static model to the solution of the diffusion equation (2.1), a task that we accomplish in Subsection 4.3

As noted in Section 4, the market is incomplete. Thus, even though we can prove convergence in distribution under a fixed equivalent martingale measure, there is an infinite number of such equivalent martingale measures. However, we can cope with this problem using the standard approach of fixing one value of the option as given apriori. This fixed value will help us determine the "proper" martingale measure to use, and solve the pricing problem.

Let T be the maturity date of the option we are trying to price and N the number of steps in our tree. Let us denote the time increment by $\Delta t = \frac{T}{N} = h$.

We will assume that $S_1, S_2, \ldots S_K$ historical stock prices are known. We will use this history of prices to estimate the volatility process in the next Section 3. For now let us assume that we have constructed an approximating process Y_t^n that converges in distribution to the volatility process Y_t in (2.1) for each $t = t_i$ time $i = 1, 2, \ldots, t_K$. We are only going to be interested in the last discrete distribution $Y_{t_K}^n$ which will converge in distribution to Y_{t_K} when $n \to \infty$. To be precise, there are two periods of time that concern us. There is the past which contains the times t_i, $i = 1, 2, \ldots, K$, that we use to estimate the discrete distribution $Y_{t_K}^n$. Then there is the future time where we construct our Markov chain, dependent only on the present stock price $S_K = S$, and the distribution of the present day volatility $Y_{t_K}^n = Y_0^n$. For simplicity of notation we will drop the subscript in Y and use Y^n for the discrete distribution and Y for the continuous process both at time $t_K = 0$ (present time).

In the following we will prove a convergence result which applies to our Static model detailed in Section 4; the result implies that our quadrinomial tree convergences in distribution to the solution of the following equation:

$$dX_t = \left(r - \frac{\sigma^2(Y)}{2} \right) dt + \sigma(Y) dW_t, \tag{2.2}$$

where Y is a random variable with the same distribution as the volatility process at time $t = 0$ i.e., Y_0. In modeling terms, we assume that Y_0 is consistent with the past evolution of the volatility process given by the autonomous model for $\{Y_s : s \leq 0\}$ in (2.1). As alluded to in the introduction, there is a difference between the distributions of X under the Dynamic model (2.1) and the Static model (2.2). On the other hand, a simple Euler method will be used to converge to the distribution of X under the Dynamic model.

We assume that the coefficients of (2.1) are regular enough to guarantee that the martingale problem associated with the diffusion process X_t in (2.1) has a unique solution starting from $x = \log S_K$, the last data point available. This is equivalent to saying that equation (2.1) has a unique solution in the weak sense, i.e. that the law of the solution is unique. Standard texts, such as (Stroock and Varadhan, 1979) can be consulted to see that this hardly places any practical restriction on the coefficients.

Let us start with a discrete Markov chain $(x(ih), \mathcal{F}_{ih})$ with transition probabilities denoted p_x^z of jumping from the point x to the point z. These transition probabilities also depend on h, but for simplicity of notation we let that subscript out. For each h let P_x^h be the probability measure on $\mathbf{R}$ characterized by:

$$\begin{cases} (i) & \mathbf{P}_x^h\left(x(0) = x\right) = 1 \\ (ii) & \mathbf{P}_x^h\Bigg(x(t) = \frac{(i+1)h-t}{h}\, x(ih) + \frac{t-ih}{h}\, x((i+1)h) \\ & \qquad\qquad , \quad ih \le t < (i+1)h \Bigg) = 1, \quad \forall\, i \ge 0 \\ (iii) & \mathbf{P}_x^h\left(x((i+1)h) = z | \mathcal{F}_{ih}\right) = p_x^z, \quad \forall\, z \in \mathbf{R} \text{ and } \forall\, i \ge 0 \end{cases} \tag{2.3}$$

Remark 2.1. We can see the following:

1. Properties (i) and (iii) say that $(x(ih), \mathcal{F}_{ih})$, $i \ge 0$ is a time-homogeneous Markov Chain starting at x with transition probability p_x^z under the probability measure $\mathbf{P}_x^h$.

2. Condition (ii) assures us that the process $x(t)$ is linear between $x(ih)$ and $x((i+1)h)$. In turn this will later guarantee that the process $x(t)$ we construct is a tree.

3. We will show in Section 4 precisely how to construct this Markov chain $x(ih)$

Conditional on being at x and on the Y^n variable, we construct the following quantities as functions of $h > 0$:

$$b_h(x, Y^n) = \frac{1}{h} \sum_{\text{z successor of x}} p_x^z (z - x) = \frac{1}{h} E^Y\left[\Delta x(ih)\right]$$

$$a_h(x, Y^n) = \frac{1}{h} \sum_{\text{z successor of x}} p_x^z (z - x)^2 = \frac{1}{h} E^Y\left[\Delta^2 x(ih)\right],$$

where the notation $\Delta x(ih)$ is used for the increment over the interval $[ih, (i+1)h]$, and E^Y denotes conditional expectation with respect to the sigma algebra $\mathcal{F}_{t_K}^Y$ generated by the Y variable. Here the successor z is determined using both the predecessor x and the Y^n random variable. We will see exactly how z is defined in Section 4 when we construct our specific Markov chain. Similarly, we define the following quantities corresponding to the infinitesimal generator of the equation (2.1):

$$\begin{aligned} b(x, Y) &= r - \frac{\sigma^2(Y)}{2}, \\ a(x, Y) &= \sigma^2(Y). \end{aligned}$$

We make the following assumptions, where $\xrightarrow{\mathrm{D}}$ denotes convergence in distribution:

$$\lim_{h\searrow 0} b_h(x, Y^n) \xrightarrow{\mathrm{D}} b(x, Y)\,, \text{ when } n \to \infty \tag{2.4}$$

$$\lim_{h\searrow 0} a_h(x, Y^n) \xrightarrow{\mathrm{D}} a(x, Y)\,, \text{ when } n \to \infty \tag{2.5}$$

$$\lim_{h\searrow 0} \max_{z \text{ successor of } x} |z - x| = 0. \tag{2.6}$$

Theorem 2.2. *Assume that the martingale problem associated with the diffusion process X_t in* (2.1) *has a unique solution $\mathbf{P}_x$ starting from $x = \log S_K$ and that the functions $a(x, y)$ and $b(x, y)$ are continuous and bounded. Then conditions* (2.4), (2.5) *and* (2.6) *are sufficient to guarantee that $\mathbf{P}_x^h$ as defined in* (2.3) *converges to $\mathbf{P}_x$ as $h \searrow 0$ and $n \to \infty$. Or equivalently saying: $x(ih)$ converges in distribution to X_t the unique solution of the equation* (2.2)

Proof. The proof of the theorem consists in showing the convergence of the infinitesimal generators formed using the discretized coefficients $b_h(.,.)$ and $a_h(.,.)$ to the infinitesimal generator of the continuous version. Since it is very similar to the proof of Theorem 11.3.4 in (Stroock and Varadhan, 1979) we omit it here. □

3 Estimating the filtered stochastic volatility distribution

In this section we describe the method used to find the distribution of the volatility process given the discrete stock price observations, also known as the stochastic volatility particle filter.

The issue of estimating the coefficients of the volatility process Y_t is itself a very difficult problem, which has not led to many satisfactory answers. We plan to address this problem using a systematic statistical analysis in a subsequent article. For now, the reader is directed to current work in (Fouque et al., 2000b), (Nielsen and Vestergaard, 2000) or (Bollerslev and Zhou, 2002).

We assume that the coefficients ν, α and the functions $\sigma(y)$ and $\psi(y)$ are known or have already been estimated. We now use an algorithm due to Del Moral, Jacod, and Protter (Del Moral et al., 2001) adapted to our specific case, in order to estimate, in the optimal filtering sense, the actual volatility process given all past stock observations. Specifically, define the random probability measure for all $i = 1, \cdots, K$,

$$p_i\,(dy) = \mathbf{P}\left[Y_{t_i} \in dy | X_{t_1}, \cdots, X_{t_i}\right].$$

This is the filtered stochastic volatility process at time i given all discrete passed observations of the stock price. Note that if $X_{t_1}, \cdots, X_{t_i}$ are assumed to be known (observed), then p_i is non-random and depends explicitly on these observed values. (Del Moral et al., 2001), Section 5, provides an algorithm which constructs n time varying particles $\left\{Y_i^j : i = 1, \cdots, K; j = 1, \cdots, n\right\}$ together with corresponding probabilities $\{p_i^j :$

$i = 1, \cdots, K; j = 1, \cdots, n\}$ such that for each i, and for any sequence of observations $X_{t_1}, \cdots, X_{ti}$, the empirical distribution of these particles converges to the probability measure $p_i\,(dy)$. This algorithm is a two step genetic-type algorithm with a Mutation step - Selection step sequence. We refer to the above cited article (Theorem 5.1) for the proof of convergence. Here we present the algorithm in detail.

The data we work with is a sequence of returns: $\{x_0 = \log S_0, x_1 = \log S_1, \ldots, x_K = \log S_K\}$, observed from the market. We need an initial distribution for the volatility process Y_t. For practical purposes, we use $\delta_{\{\nu\}}$ for this distribution. Here $\delta_{\{x\}}$ is the Dirac point mass. The only condition we need is: the functions $\sigma(x)$ and $\psi(x)$ have to be twice differentiable with bounded derivatives of all orders up to 2.

Let us define the function:

$$\phi(x) = \begin{cases} 1 - |x| & \text{if } -1 < x < 1 \\ 0 & \text{otherwise} \end{cases}$$

Another function we could use is $\phi(x) = e^{-2|x|}$. In fact, the only requirement on ϕ is that it have finite L^1 norm. In order to obtain good results, we need ϕ to be concentrated near 0. For $n > 0$ we define the contraction corresponding to $\phi(x)$ as:

$$\phi_n(x) = \sqrt[3]{n}\,\phi(x\sqrt[3]{n}) = \begin{cases} \sqrt[3]{n}\,(1 - |x\sqrt[3]{n}|) & \text{if } -\frac{1}{\sqrt[3]{n}} < x < \frac{1}{\sqrt[3]{n}} \\ 0 & \text{otherwise} \end{cases}. \tag{3.1}$$

First we choose $m = m_n$ an integer.

Step 1: We start with $X_0 = x_0$ and $Y_0 = y_0 = \nu$.

Mutation step: This part calculates a random variable with approximately the same distribution as (X_1, Y_1) using the well known Euler scheme for the equation (2.1). More precisely we set:

$$\begin{aligned} Y(m, y_0)_{i+1} \ &: = Y_{i+1} = Y_i + \frac{1}{m}\alpha(\nu - Y_i) + \frac{1}{\sqrt{m}}\psi(Y_i)U_i \\ X(m, x_0)_{i+1} \ &: = X_{i+1} = X_i + \frac{1}{m}(r - \frac{\sigma^2(Y_i)}{2}) + \frac{1}{\sqrt{m}}\sigma(Y_i)U_i'. \end{aligned} \tag{3.2}$$

Here U_i and U_i' are iid Normal random variates with mean 0 and variance 1. At the end of this first evolution step we obtain:

$$\begin{aligned} X_1 &= X(m, x_0)_m, \\ Y_1 &= Y(m, y_0)_m. \end{aligned} \tag{3.3}$$

Selection step: We repeat the evolution step n times to obtain n pairs: $\{(X_1^j, Y_1^j)\}_{j=\overline{1,n}}$. Now we introduce the discrete probability measure:

$$\Phi_1^n = \begin{cases} \frac{1}{C}\sum_{j=1}^n \phi_n(X_1^j - x_1)\delta_{\{Y_1^j\}} & \text{if } C > 0 \\ \delta_{\{0\}} & \text{otherwise.} \end{cases} \tag{3.4}$$

Here the constant C is chosen to make Φ_1^n a probability measure, i.e., $C = \sum_{j=1}^n \phi_n(X_1^j - x_1)$. The idea is to "select" only the values of Y_1 which correspond to values of X_1 not far away from the realization x_1. We end the first Selection step by simulating n iid variables $\{Y'^j_1\}_{j=\overline{1,n}}$.

Steps 2 to K: For each step $i = 2, 3, \ldots, K$, we first apply the evolution step to each of the variables selected at the end of the previous step, that is, starting with $X_0 = x_{i-1}$ and $Y_0 = Y'^j_{i-1}$ for each $j = 1, 2, \ldots, n$ in (3.2). Thus, we obtain n pairs $\{(X_i^j, Y_i^j)\}_{j=1,2,\ldots,n}$. Then we apply the selection step to these pairs. That is, we use them in the distribution (3.4) instead of the $\{(X_1^j, Y_1^j)\}_{j=1,2,\ldots,n}$ pairs, and x_i instead of x_1.

At the end of each step i we obtain a discrete distribution Φ_i^n, and this is our estimate for the transition probability of the process Y_t at the respective time t_i. In our construction of the quadrinomial tree, we use only the latest estimated probability distribution, i.e., Φ_K^n. We will refer to this distribution by saying that it is the set of particles $\{\bar{Y}_1, \bar{Y}_{2,} \cdots, \bar{Y}_n\}$ together with their corresponding probabilities (or weights) $\{\bar{p}_1, \bar{p}_2, \ldots, \bar{p}_n\}$.

4 The Static Model: Constructing the tree

We will present two models that will approximate the price of the option. We called the first model presented the "Static" model since the volatility distribution is not changing at every step in the future. It remains static, equal with the distribution at the present time. As we shall see from Section 7 this is the better model when compared with what we will later call the "Dynamic" model.

We assume that we have an option with maturity T. Our purpose in this section is to construct a discrete tree which will assist us calculating an estimate of the given option's price. The data available is the value of the stock price today S, and a history of earlier stock prices. As described in the previous section, we use them to compute a set Y^n of particles $\{\bar{Y}_1, \bar{Y}_{2,} \cdots, \bar{Y}_n\}$ with weights $\{\bar{p}_1, \bar{p}_2, \ldots, \bar{p}_n\}$, whose empirical law approximates the best estimate of Y_0 at time 0 for the volatility process given all past observations.

Remark 4.1. The market is incomplete. Thus, the option price is not unique.

It is easy to see that the remark is true as is the case with all the stochastic volatility models since the number of sources of randomness (2) is bigger than the number of tradable assets (1). Remember that the volatility process is not a tradable asset, and cannot, in practice, be observed in *discrete* time. This means that the price of a specific derivative will not be completely determined by the specification (2.1) of the (X, Y) dynamics observed in discrete time, and by the requirement that the market is free of arbitrage. However, the requirement of no arbitrage does imply that the prices of various derivatives will have to satisfy certain internal consistency relationships, in order to avoid arbitrage possibilities on the derivative market.

To take advantage of this fact and to be able to use the classical pricing idea in incomplete markets, we need to make an assumption:

Assumption 4.2. There is a liquid market for every contingent claim.

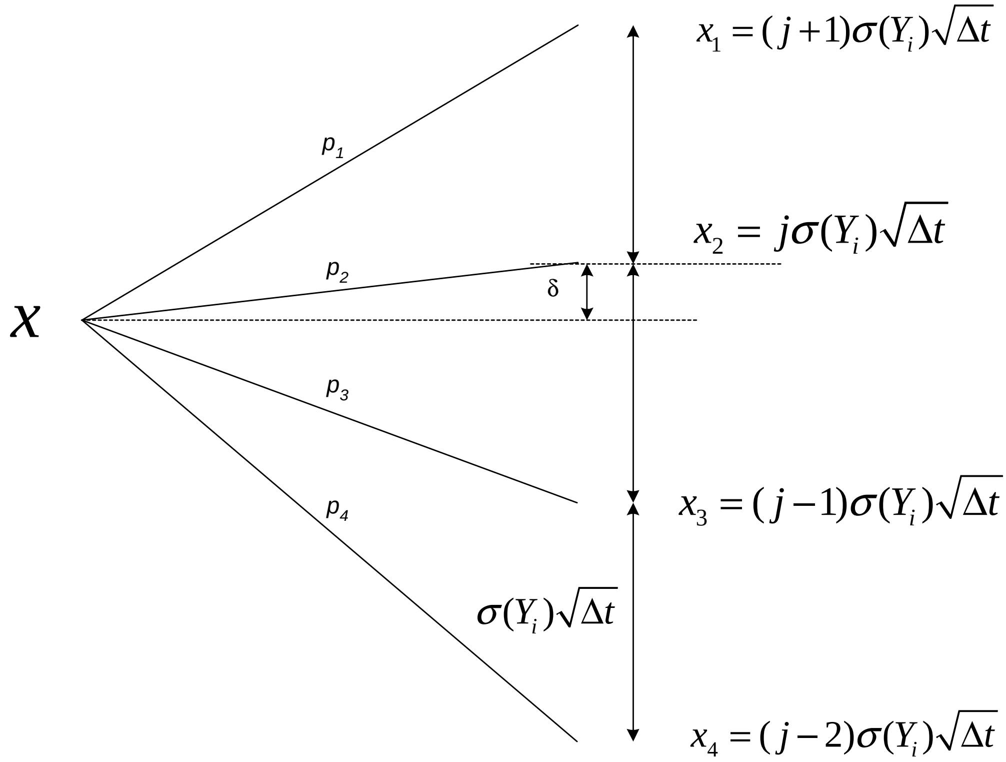

Figure 1: The basic successors for a given volatility value. Case 1.

This assumption assures us that the derivatives are tradable assets. Thus, taking the price of one particular option (called the "benchmark" option) as given a priori will allow us to find a unique price for all the other derivatives. Indeed, we would then have two sources of randomness and two tradable assets (the stock and the benchmark), and the price of any derivative would be uniquely determined.

Let us divide the interval $[0, T]$ into N subintervals each of length $\Delta t = \frac{T}{N} = h$. At each of the points $i\Delta t = ih$ the tree is branching. The nodes on the tree represent possible return values $X_t = \log S_t$.

4.1 Construction of the one period model

Now, assume that we are at a point x. What are the possible successors of x?

We sample a volatility value from the discrete approximating distribution Y^n at each time period ih, $i \in \{1, 2, \ldots, N\}$. Denote the value drawn at step i corresponding to time ih by Y_i. Corresponding to this volatility value Y_i we will construct the successors in the following way.

We consider a grid of points of the form $l\sigma(Y_i)\sqrt{\Delta t}$ with l taking integer values. No matter where the parent x is, it will fall at one such point or between two grid points. In this grid, let j be the integer that corresponds to the point above x. Mathematically, j is

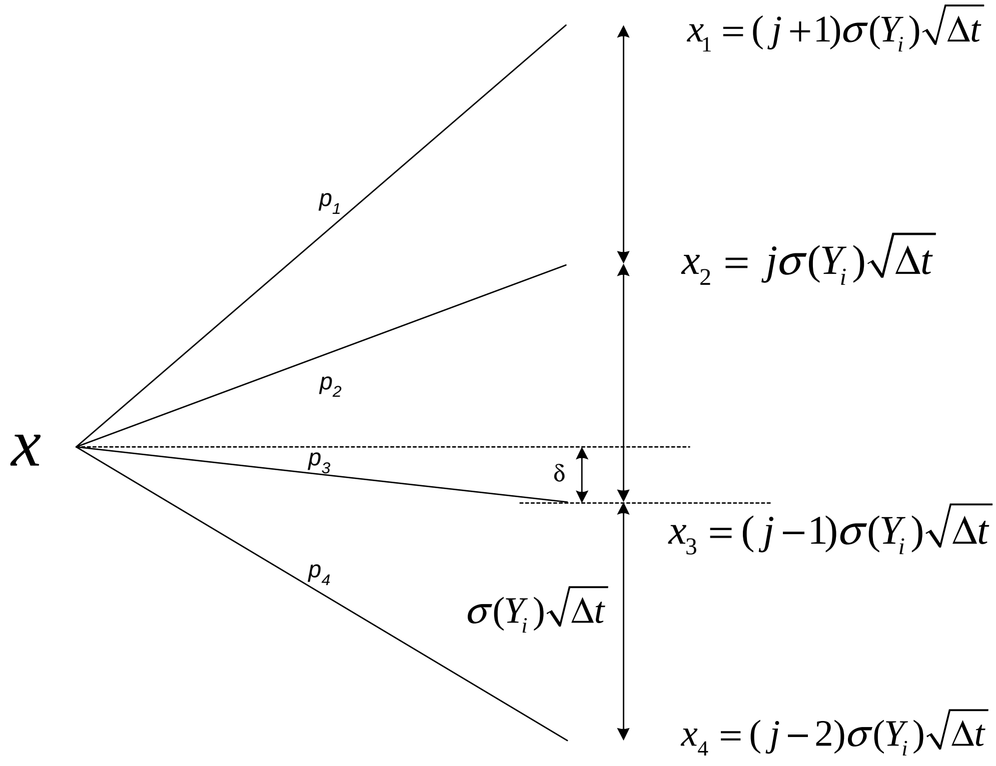

Figure 2: The basic successors for a given volatility value. Case 2.

the point that attains:

$$\inf\left\{l \in \mathbf{N} \quad | \quad l\,\sigma(Y_i)\sqrt{\Delta t} \geq x\right\}$$

We will have two possible cases: either the point $j\,\sigma(Y_i)\sqrt{\Delta t}$ on the grid corresponding to j is closer to x, or the point $(j-1)\sigma(Y_i)\sqrt{\Delta t}$ corresponding to $j-1$ is closer. We will treat the two cases separately.

Let us denote

$$\delta := x - j\,\sigma(Y_i)\sqrt{\Delta t},$$
$$q := \delta / \left(\sigma(Y_i)\sqrt{\Delta t}\right).$$

Remark 4.3. With the notation above we have $\delta \in \left[-\frac{\sigma(Y_i)\sqrt{\Delta t}}{2}, \frac{\sigma(Y_i)\sqrt{\Delta t}}{2}\right]$ or $q \in \left[-\frac{1}{2}, \frac{1}{2}\right]$

Case 1. *$j\,\sigma(Y_i)\sqrt{\Delta t}$ is the point on the grid closest to x, $\left(q \in \left[-\frac{1}{2}, 0\right]\right)$.*

Figure 1 on page 10 refers to this case.

One of the assumptions we need to verify is (2.4), which asks the mean of the increment to converge to the drift of the X_t process in (2.2). In order to simplify this requirement, we add the drift quantity to each of the successors. This trick will simplify the conditions (2.4) to ask now the convergence of the mean increment to zero. This idea has been previously used by many authors including Leisen as well as Nelson & Ramaswamy.

Explicitly, we take the 4 successors to be:

$$\begin{cases} x_1 & = (j+1)\sigma(Y_i)\sqrt{\Delta t} + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \\ x_2 & = j\sigma(Y_i)\sqrt{\Delta t} + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \\ x_3 & = (j-1)\sigma(Y_i)\sqrt{\Delta t} + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \\ x_4 & = (j-2)\sigma(Y_i)\sqrt{\Delta t} + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \end{cases} \tag{4.1}$$

First notice that condition (2.6) is trivially satisfied by this choice of successors. The plan is to set a system of equations consisting of the variance condition (2.5), and the mean condition (2.4), and to solve it for the joint probabilities p_1, p_2 , p_3 and p_4. Because the market is incomplete, we cannot expect to have a unique solution to the system. *However, each solution will give us an equivalent martingale measure.*

Algebraically, we write: $j\,\sigma(Y_i)\sqrt{\Delta t} = x - \delta$, and using this we infer that the increments over the period Δt are:

$$\begin{cases} x_1 - x & = \sigma(Y_i)\sqrt{\Delta t} - \delta + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \\ x_2 - x & = -\delta + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \\ x_3 - x & = -\sigma(Y_i)\sqrt{\Delta t} - \delta + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \\ x_4 - x & = -2\sigma(Y_i)\sqrt{\Delta t} - \delta + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \end{cases} \tag{4.2}$$

Conditions (2.4) and (2.5) translate here as:

$$\mathbf{E}[\Delta x|Y_i] = \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t$$

$$\mathbf{V}[\Delta x|Y_i] = \sigma^2(Y_i)\Delta t$$

where by Δx we denote the increment over the period Δt.

We will solve the following system of equations with respect to p_1, p_2 p_3 and p_4:

$$\begin{cases} \left(\sigma(Y_i)\sqrt{\Delta t} - \delta\right)p_1 & +(-\delta)p_2 + \left(-\sigma(Y_i)\sqrt{\Delta t} - \delta\right)p_3 + \left(-2\sigma(Y_i)\sqrt{\Delta t} - \delta\right)p_4 = 0 \\ \left(\sigma(Y_i)\sqrt{\Delta t} - \delta\right)^2 p_1 & +(-\delta)^2 p_2 + \left(-\sigma(Y_i)\sqrt{\Delta t} - \delta\right)^2 p_3 + \left(-2\sigma(Y_i)\sqrt{\Delta t} - \delta\right)^2 p_4 \\ & -\mathbf{E}[\Delta x|Y_i]^2 = \sigma^2(Y_i)\Delta t \\ & p_1 + p_2 + p_3 + p_4 = 1 \end{cases}. \tag{4.3}$$

Eliminating the terms in the first equation of the system we get:

$$\sigma(Y_i)\sqrt{\Delta t}\ (p_1 - p_3 - 2p_4) - \delta = 0$$

or

$$p_1 - p_3 - 2p_4 = \frac{\delta}{\sigma(Y_i)\sqrt{\Delta t}}. \tag{4.4}$$

Neglecting the terms of the form $\left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t$ when using (4.4) in the second equation in (4.3) we obtain the following:

$$\sigma^2(Y_i)\Delta t = \sigma^2(Y_i)\Delta t\,(p_1 + p_3 + 4p_4) + 2\delta\sigma(Y_i)\sqrt{\Delta t}\,(p_3 - p_1 + 2p_4) + \delta^2 - \left(\sigma(Y_i)\sqrt{\Delta t}\;(p_1 - p_3 - 2p_4) - \delta\right)^2.$$

After simplifications, we obtain the equation:

$$(p_1 + p_3 + 4p_4) - (p_1 - p_3 - 2p_4)^2 = 1.$$

So now the system of equations to be solved looks like:

$$\begin{cases} p_1 + p_3 + 4p_4 = 1 + \frac{\delta^2}{\sigma^2(Y_i)\Delta t} \\ p_1 - p_3 - 2p_4 = \frac{\delta}{\sigma(Y_i)\sqrt{\Delta t}} \\ p_1 + p_2 + p_3 + p_4 = 1 \end{cases} \tag{4.5}$$

Note that this is a system with 4 unknowns and 3 equations. Thus, there exists an infinite number of solutions to the above system. Since we are interested in the solutions in the interval $[0, 1]$, we are able to reduce somewhat the range of the solutions. Let us denote by p the probability of the branch furthest away from x. In this case $p := p_4$. Expressing the other probabilities in term of p and the q defined in Remark 4.3, we obtain:

$$\begin{cases} p_1 = \frac{1}{2}(1 + q + q^2) - p \\ p_2 = 3p - q^2 \\ p_3 = \frac{1}{2}(1 - q + q^2) - 3p \end{cases} \tag{4.6}$$

Now using the condition that every probability needs to be between 0 and 1, we solve the following three inequalities:

$$\begin{cases} \frac{1}{2}(-1 + q + q^2) \le p \le \frac{1}{2}(1 + q + q^2) \\ \frac{q^2}{3} \le p \le \frac{1+q^2}{3} \\ \frac{1}{6}(-1 - q + q^2) \le p \le \frac{1}{6}(1 - q + q^2) \end{cases} \tag{4.7}$$

It is not difficult to see that the solution of the inequalities (4.7) is $p \in [\frac{1}{12}, \frac{1}{6}]$. Respectivelly, we will obtain an equivalent martingale measure for every $p \in [\frac{1}{12}, \frac{1}{6}]$ thanks to the first equation in (4.3).

We postpone the statement of this result until after Case 2 bellow.

Case 2. $(j-1)\sigma(Y_i)\sqrt{\Delta t}$ *is the point on the grid closest to* x, $\left(q \in \left[0, \frac{1}{2}\right]\right)$.

Figure 2 on page 11 refers to this case.

Remark 4.4. This second case is just the mirror image of the first case with respect to x.

The 4 successors are the same as in Case 1; the increments are going to be calculated similarly with (4.2) and using the Remark 4.4 together with the equations (2.4) and (2.5) will give the following solution:

$$\begin{cases} p_2 = \frac{1}{2}(1+q+q^2) - 3p \\ p_3 = 3p - q^2 \\ p_4 = \frac{1}{2}(1-q+q^2) - p \end{cases}, \tag{4.8}$$

where p is the probability of the successor furthest away, in this case p_1. This is just the solution given in (4.6) with $p_1 \rightleftarrows p_4$ and $p_2 \rightleftarrows p_3$ taking into account the interval for δ. Thus, we are able to state the following result.

Lemma 4.5. *If we construct a one step qudrinomial tree with the successors given by* (4.1), *and we denote with p the probability of the furthest away successor from x, then for every $p \in [\frac{1}{12}, \frac{1}{6}]$:*

(i) in **Case 1** ($q \in [-\frac{1}{2}, 0]$) *the relations* (4.6) *define an Equivalent Martingale Measure.*

(ii) in **Case 2** ($q \in [0, \frac{1}{2}]$) *the relations* (4.8) *define an Equivalent Martingale Measure.*

Remark 4.6. As we observed above, constructing the Equivalent Martingale Measure involves solving a system with 3 equations and 4 unknowns. It is natural to ask then why not try a tree with three successors, which will in turn imply solving a system with 3 equations and 3 variables. However, it turs out that such a system has no solution for any possible choice of successors[2].

4.2 Construction of the multi-period model. Option valuation.

Suppose now that we have to compute an option value. For illustrative purposes we will use an European type option, but the method should work with any kind of path dependent option e.g., American, Asian, Barrier etc.

Assume that the payoff function is: $\Phi(X_T)$. The maturity date of the option is T, and the purpose is to compute the value of this option at time $t = 0$ (for simplicity) using our model (2.1). We divide the interval $[0, T]$ into N smaller ones of length $\Delta t := \frac{T}{N}$. At each of the points $i\Delta t$ with $i \in \{1, 2, \ldots, N\}$ we then construct the successors in our tree as in the previous section. This tree converges in distribution to the solution of the stochastic model (2.2). A proof of this fact using Theorem 2.2 is found in the next subsection.

In order to calculate an estimate for the option price we will employ a resampling method based on the particles defining the approximate discrete distribution for the initial volatility Y. Suppose that we have the discrete probability distribution of Y, i.e. we know the stochastic volatility particle filter values $\{\bar{Y}_1, \bar{Y}_2, \ldots, \bar{Y}_n\}$, each with probability $\{\bar{p}_1, \bar{p}_2, \ldots, \bar{p}_n\}$. We sample N values from this distribution, and use them like the realization of volatility process Y along the N levels of the tree, into the future. In other words, call these sampled values

[2] This is to be expected, since the existence of such a solution will contradict the incompleteness of the market

$Y_1, \cdots, Y_N$, we start with the initial value x_0. We then compute the 4 successors of x_0 as in the previous section for the first sampled value, Y_1. After this, for each one of the 4 successors we compute their respective successors for the second sampled volatility value Y_2, and so on.

The tree we construct this way allows us to compute one instance of the option price by using the standard pricing technique that is consistent with a no-arbitrage condition: after creating the tree based on our sampled values, we compute the value of the payoff function Φ at the terminal nodes of the tree. Then, working backward in the path tree, we compute the value of the option at time $t = 0$ as the discounted expectation of the final node values. Because the tree is recombining by construction, the level of computation implied is manageable, typically of a polynomial order in N. The complexity of the filtering algorithm leading to the original particle values $\{\bar{Y}_1, \bar{Y}_2, \ldots, \bar{Y}_n\}$ is no greater.

If we are to iterate this procedure by using repeated samples $\{Y_1, \cdots, Y_N\}$, we can take the average of all prices obtained for each tree generated using each separate sample. This Monte Carlo method converges, as the number of particles n increases, to the true option price for the quadrinomial tree in which the original distribution of the volatility is the true law of Y_0 given past observations of the stock price. We leave a full proof of this convergence out of our article. It uses the following fact proved by Pierre del Moral, known as a *propagation of chaos* result: as n increases, for a fixed number N of particles $\{Y_1, \cdots, Y_N\}$ sampled from the distribution of particles $\{\bar{Y}_1, \bar{Y}_2, \ldots, \bar{Y}_n\}$ with probabilities $\{\bar{p}_1, \bar{p}_2, \ldots, \bar{p}_n\}$, the Y_i's are asymptotically independent and all identically distributed according to the law of Y_0 given all past stock price observations. Chapter 8, and in particular Theorem 8.3.3 in (Del Moral, 2004), can be consulted for this fact. The convergence proof in the next section is also based on del Moral's propagation of chaos.

In fact, a further convergence result can be established here. According to del Moral's propagation of chaos Theorem 8.3.3 in (Del Moral, 2004), if the number of samples $N = N(n)$ is taken to be a function of the number of particles n, and if the number of time steps $K = K(n)$, used before time 0 to simulate the particle approximation of Y_0, is also a function of n, then the speed at which the samples $\{Y_1, \cdots, Y_N\}$ converge to independent copies of the filtered value of Y_0 given all past continuously observed stock prices is given by

$$\frac{N^2(n)\, K(n)}{n}.$$

Thus if N and K are chosen so that this quantity tends to 0 as n tends to $+\infty$, we indeed have the announced convergence. But because of the very special nature of stochastic volatility filtering, at any time t, the squared volatility in the original model (2.1) is actually equal to the differential of the quadratic variation of the martingale X, which means that when the number of time observations N tends to infinity, $\sigma^2(Y_t)$ does not need to be filtered: it is actually known given the entire past of the path of X in continuous time. For simplicity, assume that σ^2 is a bijective function on the space where Y lives, or change the dynamics of Y so that $\sigma^2(Y_t)$ is Y_t itself. Alternately, we see that the filtered value of Y_0 tends to the actual objective value of Y_0 when N tends to $+\infty$. Hence in the above situation where we can apply the propagation of chaos, if $N(n)$ also tends to $+\infty$, we can guarantee that our sample $\{Y_1, \cdots, Y_N\}$ converges in distribution to independent copies of Y_0. Then, also invoking the convergence theorem of the next section, we conclude that the option valuation method described in this section converges to the true option price under the Static model (2.2) where the fixed random variable Y is the true objective volatility Y_0.

Figures 3 and 4, on pages 26 and 27, contain a number of simulated trees for various values of the parameter p. We can also visualize from Figures 3 and 4 that the trees recombine, and accordingly that the level of computation is not very high.

4.3 Convergence result for the quadrinomial tree

As noted in Section 2 we shall prove that our constructed tree converges to the solution of the process

$$dX_t = \left(r - \frac{\sigma^2(Y)}{2} \right) dt + \sigma(Y) dW_t,$$

where Y is a random variable with the same distribution as the actual volatility process at time $t = 0$ i.e., Y_0.

From the empirical results in Section 6 we shall see that this model actually performs better than the Dynamic model presented in Section 5, as announced in the introduction.

Theorem 4.7. *If we denote by p the probability of the successor furthest away from the parent point x, then the relations* (4.6) *and* (4.8) *respectively define an Equivalent Martingale Measure for every $p \in [\frac{1}{12}, \frac{1}{6}]$ depending on which successor is further away from x: x_4, respectively x_1.*

Furthermore, under any such measure the tree defined by the relations (4.1) *converges in distribution to the continuous process* (2.2) *as the time interval $\Delta t \to 0$ and the number of particles $n \to \infty$ in the discrete distribution Y^n*

Proof. The first equation in the systems we solved above guarantees that the discounted process has expected increments zero, thus assuring us that the resulting measure we find is a martingale measure.

It remains to show the convergence result, and to this end we are using Theorem 2.2. More specifically, we are going to prove that the two critical assumptions (2.4) and (2.5) are satisfied. Assume that at step i the tree is constructed using the volatility value Y_i sampled from the distribution Y^n.

Let us define the variable $\mathcal{X}$ in **Case 1**:

$$\mathcal{X} = \begin{cases} \sigma(Y_i)\sqrt{\Delta t} - \delta & \text{with probability } \frac{p_1}{\bar{p}_i} \\ -\delta & \text{with probability } \frac{p_2}{\bar{p}_i} \\ -\sigma(Y_i)\sqrt{\Delta t} - \delta & \text{with probability } \frac{p_3}{\bar{p}_i} \\ -2\sigma(Y_i)\sqrt{\Delta t} - \delta & \text{with probability } \frac{p_4}{\bar{p}_i} \end{cases}$$

and in **Case 2** as:

$$\mathcal{X} = \begin{cases} 2\sigma(Y_i)\sqrt{\Delta t} - \delta & \text{with probability } \frac{p_1}{\bar{p}_i} \\ \sigma(Y_i)\sqrt{\Delta t} - \delta & \text{with probability } \frac{p_2}{\bar{p}_i} \\ -\delta & \text{with probability } \frac{p_3}{\bar{p}_i} \\ -\sigma(Y_i)\sqrt{\Delta t} - \delta & \text{with probability } \frac{p_4}{\bar{p}_i} \end{cases}$$

where $\frac{\delta}{\sigma(Y_i)\sqrt{\Delta t}}$ is in the interval $[-\frac{1}{2}, 0]$ in case 1 and in $[0, \frac{1}{2}]$ in the case 2.

Let us note here that in either of the two cases $\Delta x|Y_i = \mathcal{X} + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t$. Also, remark that the system (4.3) gives the mean and variance of $\mathcal{X}$ as 0 and $\sigma^2(Y_i)\Delta t$, respectively.

Thus, we have:

$$\begin{aligned} \mathbf{E}[\Delta x|Y_i] &= \mathbf{E}[\mathcal{X}] + \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \\ &= \left(r - \frac{\sigma^2(Y_i)}{2}\right)\Delta t \end{aligned}$$

Now from the definition of $b_h(x)$ we have:

$$b_h(x) = \frac{\mathbf{E}[\Delta x|Y_i]}{\Delta t} = r - \frac{\sigma^2(Y_i)}{2}$$

At this point we need to invoke Theorem 8.3.3 in (Del Moral, 2004). When applied to our particular case it implies that each of the variables Y_1, Y_2, ... drawn from the discrete distribution Y^n converge in distribution as $n \to \infty$ to a random variable $Y^{\Delta t}$ whose law is that of the filtered value of Y_0 given the observed stock prices before time 0. Moreover, using the facts given in the last paragraph of the previous subsection, when $\Delta t \to 0$, the law of $Y^{\Delta t}$ tends to the law of the actual initial volatility Y_0. Thus for a n large enough and Δt small enough, using the continuity of $\sigma(y)$, we obtain that the Assumption (2.4) is satisfied.

Since $\mathbf{V}[\Delta x|Y_i] = \mathbf{V}[\mathcal{X}]$ we have:

$$\mathbf{E}[(\Delta x)^2|Y_i] = \mathbf{V}[\mathcal{X}] + \mathbf{E}[\Delta x|Y_i]^2 = \sigma^2(Y_i)\Delta t + \left(r - \frac{\sigma^2(Y_i)}{2}\right)^2 \Delta t^2$$

Thus:

$$a_h(x) = \frac{\mathbf{E}[\Delta x^2|Y_i]}{\Delta t} = \sigma^2(Y_i) + \left(r - \frac{\sigma^2(Y_i)}{2}\right)^2 \Delta t$$

Using the fact that r is a constant and that the function σ is locally bounded, the second term in $a_h(x)$ converges to 0 as $\Delta t \to 0$. Using again the Theorem 8.3.3 in (Del Moral, 2004) and the continuity of the function $\sigma^2(y)$, for a large enough n, we obtain the Assumption (2.5). □

To construct our tree we need to know the value of the parameter p described in the previous Subsection. In order to do that we use the price of a suitably chosen option from the market to calibrate for the parameter p. The option we chose to use in our numerics (see Section 6) is the previous day (April 21, 2004, in our case) at-the-money option, but theoretically, it could be any option from any moment in the past. For fixed values of p on a dense grid on the interval $[\frac{1}{12}, \frac{1}{6}]$ we generate trees and compute option prices corresponding to each p in the grid. Then we compare the results obtained with the price of the option from the market and we choose the value p that gave the closest value to the option on the market.

We use this parameter p to compute values for all the options at time 0 (April 22, 2004). A graphical illustration of this process applied to a specific example is presented in Figure 7. It turns out that these option prices are quite insensitive to the actual choice of p, for p in a wide range within the interval $[\frac{1}{12}, \frac{1}{6}]$. This robustness is a highly desirable property when one is faced with deciding in a rather arbitrary way how to choose a martingale measure.

Using our tree we can also approximate the sensitivities of the option price (the Greeks). If we denote by C the value of the option obtained using our tree method, we can compute:

$$\begin{aligned}
delta =& \frac{\partial C}{\partial S} = \frac{C(S+\Delta S) - C(S-\Delta S)}{2\Delta S} \\
gamma =& \frac{\partial^2 C}{\partial S^2} = \frac{C(S+\Delta S) - 2C(S) + C(S-\Delta S)}{\Delta S^2} \\
theta =& \frac{\partial C}{\partial t} = \frac{C(t+\Delta t) - C(t)}{\Delta t} \\
rho =& \frac{\partial C}{\partial r} = \frac{C(r+\Delta r) - C(r)}{\Delta r}
\end{aligned}$$

Here the value of the option is calculated using various initial conditions. For example, $C(S+\Delta S)$ will be estimated using an initial asset price of $S+\Delta S$ with ΔS small. $\Delta S = 0.001S$ would be a good choice here. Every price in a difference or a sum should be computed using the same set of volatility values to eliminate variability due to randomness.

Remark 4.8. Notice that we do not compute *vega* which is the derivative with respect to the volatility, because in our case it does not make much sense. We could also compute a nonstandard derivative with respect to the above described parameter p.

Remark 4.9. From the present moment when we start to construct our quadrinomial tree and until expiration the discrete distribution of the volatility remains unchanged in time. The next section details a simple model when the discrete distribution of the volatility evolves in time: the dynamic model.

5 The Dynamic Model: using stochastic volatility filtering in a Monte-Carlo method

The idea of this model is to start with the estimated volatility distribution at the present time $Y^n = Y_0^n$ and with the logarithm of the price of the stock today $x_0 = \log S_0$ in equation

(2.1), and then to take advantage of the same Euler scheme we used in the Section 3 to generate future stock prices and volatility values.

More precisely we start with x_0 and the Y_0^n distribution. We sample a volatility value y_0 from the distribution Y_0^n. We divide the time to expiration T into N intervals of length $\Delta t = T/N$ and generate a path (X, Y) recursively:

$$\begin{aligned} Y(y_0)_{i+1} \quad &: \quad = Y_{i+1} = Y_i + \alpha(\nu - Y_i)\Delta t + \psi(Y_i)U_i\sqrt{\Delta t}, \\ X(x_0)_{i+1} \quad &: \quad = X_{i+1} = X_i + (r - \frac{\sigma^2(Y_i)}{2})\Delta t + \sigma(Y_i)U_i'\sqrt{\Delta t}. \end{aligned} \tag{5.1}$$

where the variables U_i and U_i' are iid standard normal and $i \in \{0, 1, 2, \ldots, N-1\}$. In other words, we use the actual "Stochastic Volatility" dynamics for simulating future values of (X, Y), but started from the initial distribution $\delta_{\{x_0\}} \otimes Y_0^n$.

Once we find the value at the expiration $X(x_0)_N = X_T$ we can compute the value of the option at the expiration, and then we discount back to the present value using the risk-free rate. This represents one replication of a Monte Carlo method: to compute an estimate of the option price we generate many replications (typically of the order 10^6) then compute the average of the values obtained. This average is our estimate for the price of the option today.

Any Monte-Carlo method is notoriously inefficient, and one may try to improve the convergence of the method by such techniques as reduction of variance, and the like. However, since our dynamic model yields option prices which are not as close to those given in the market as the static model's prices, there seems little reason to improve the efficiency of our Monte-Carlo method for the dynamic model.

At this stage it is worth noting that this article's second-named author, in (Viens, 2002), provides a Monte Carlo method that solves a related stochastic portfolio optimization problem using elements of stochastic control, dynamically in time, based on the dynamic evolution of the stochastic volatility particle filter. Because of the non-linear nature of portfolio optimization, as opposed to the linear nature of option pricing, the numerics proposed in (Viens, 2002) are difficult to implement in general (see however the special case of power utility, for which a successful implementation can be found in (Batalova and Viens, 2005)). We hope that the successful option-pricing implementation in this article will be an invitation for researchers to apply the same models and methods to the optimization problem in (Viens, 2002).

6 Using real data: European Call options on S&P500 and IBM

We have chosen to illustrate our method with two sets of data. The first set is S&P500 data gathered on April 21-22 2004. We are using daily data from January 1st, 1999 to April 21, 2004 to compute the discrete volatility distribution according to the method described in Section 3. Figure 5 represents the evolution of the S&P500 index price over the time period mentioned above. The second set used is IBM data gathered on July 18-19, 2005. We present a more detailed explanation about this dataset on page 20

We are working with the model presented in (2.1) with $\phi(y) = \beta$ and $\sigma(y) = e^{-|y|}$, using the following parameters for the volatility equation: $\alpha = 50$, $m = -4.38$ $\beta = 1$ and $r = .01$ for the price. The parameters have been estimated from the data, and the short term interest rate is the value published for April 21, 2004.

We estimate the discrete volatility distribution using the Del Moral, Jacod, Protter method presented in Section 3. Figure 6(a) presents a plot of this distribution.

To compare our method, we also estimate the implied volatility on April 21 for a range of strike prices from the option data available that day. To do so, we use a simple bisection method. Figure 6(b) shows the implied volatility's behavior for various strike prices.

We should note that we used the option and stock (index) data available on April 21st to estimate these two plots. The implied volatility plot 6(b) corresponds to the 29-day-maturity options but it is representative for the other maturities as well. We notice very high implied volatility values for options deep in the money, and for the most others the implied volatility is around $0.12 - 0.135$.

Using the data available a day earlier we estimate option prices for that day for many values of the parameter p in the interval $[\frac{1}{12}, \frac{1}{6}]$. Then, we compare the estimated prices with the price of the "benchmark option" which we chose to be the option at the money. We can see computed values corresponding to a grid for the p parameter in Figure 7.

Beginning with $p = 0.135$ and ending with $p = 0.16$, the option values obtained are close to each other. In fact, this is a feature we have observed for all the option values calculated for the entire range of strike prices. The values of the 29-day-maturity options obtained for $p = 0.135$ are presented in Table 1 in the Appendix. For better illustration of the performance of the various methods we present in Figures 8, 9, and 10 on pages 30, 31, and 32, the values of the options separated in groups depending on the range of the strike prices (in the money, at the money, and out of the money, respectively).

All the above option-price graphs include, for comparison, the bid-ask spread of the actual prices seen on the option market, the price given via the standard Black-Scholes formula with constant (non-random) volatility. Figure 11 presents the estimated derivatives (sensitivities) with respect to various variables in the model.

One of the big advantages of our method is that it allows the computation of option prices even when there exists no formula for that option type. Since we lacked data for nonstandard options we have tried something else. In the third Friday of each month the European option with maturity that month expires. Thus the option with expiration 2 months becomes a one month option and so on. Also, during the following Monday and Tuesday an option with a new, intermediate, maturity date is starting to be traded. Since there is no implied volatility in the previous day for this option we suspected that our method will perform well. For this simulation we used IBM stock data from July 18-19, 2005. The new options with expiration in September did not appear until Tuesday July 19, 2005 so we are using the Monday July 18 for volatility calibration and finding the optimal p according to the method described above. The coefficients used in this case were: $\alpha = 11.85566$, $m = 0.9345938$ $\beta = 4.13415$ and $r = .0343$. We have estimated these coefficients from the historical data, the method we used is to be the subject of another article. We present the values obtained in the Table 2 in the Appendix, we also plot these values in Figure 12 on page 34.

7 Conclusions

The approximation methods presented in this work are computationally intensive, but are based on simple algorithms featuring low, manageable, complexity. By performing volatility estimation using a distribution corresponding to the best estimate of the present volatility level given all past stock price observations, we are able to implement a quadrinomial tree method yielding option prices that are very close to the same prices observed on the option market.

The two applied cases presented here are fundamentally different. In the first application we used an Index (S&P 500) which is not very sensitive to small movements in the market, therefore preferred by the majority of the theoretical papers in Mathematics of Finance. For the second application we choose a popular technology stock (IBM) but we should mention that the same basic conclusions were observed and calculated for many other stocks (Microsoft, Intel, Yahoo, Ford and GM), not included here for obvious space considerations.

In the first (S&P 500) case we observe that pricing using our Static model (Section 4) is clearly better than the pricing using the Dynamic model of Section 5, since the former typically falls within the option market's bid-ask spread, while the latter does not. This is despite the fact that the Dynamic model approximates pricing under the true Stochastic Volatility model, while in the Static model, volatility is assumed to be constantly distributed in the future according to its best present estimated distribution given all past stock price observations. We should also mention that we used 10^6 simulations for the Dynamic model and that the run time is about eight times larger than the run time for the Static model.

If we look at the three regions depending on how close the strike price is relatively to the stock price in that respective day, we see that the best performance is obtained for the options at-the-money – conveniently so, since those are the majority of options traded. We suspect that we have obtained better results for that region because we used an option at the money to calibrate the parameter p. This fact would suggest to use different optimal values of p for each of the three regions. Recalibrating p for each maturity date should provide even better estimates, although this might go against the idea of arbitrage-free consistency within a single option market.

On the other hand, since we are using European call options, for which there exists a formula that every market participant can look up at any moment, our results are naturally not far from the values obtained using the Black-Scholes formula. The strength of our method is that it works for any type of option including those that do not have a valuation formula, and may be path dependent.

This fact is illustrated in the second analysis (the IBM case). We see that the Static model once again performs better for at the money options (in the Strike price range 75-85). In this case we also wanted to investigate the reason why the Dynamic method performed worse in the previous simulation so we increased the number of simulations by 100 (thus now using 10^8 runs). We see that practically there is no difference now between the Static and the Dynamic models, with the exception of the runtime, which naturally is huge now for the Dynamic model when compared with the Static one.

Thus, in conclusion, this article's merit is in showing that in a well-known and mature market, our quadrinomial tree method outperforms others techniques, including the classical Black-Scholes method, or the Monte-Carlo method for the Stochastic Volatility model, even when the latter is based on an excellent particle approximation of the best possible stochastic volatility estimation technique.

Lastly, we mention the question of hedging. Since we can estimate the sensitivity of the estimated option price with respect to the various factors in the model (the Greeks: delta, gamma, theta and rho) we can do a heuristic *delta hedging*: starting with the option value at time zero, we can devise a dynamic trading strategy in stock and in a risk-free account, based on the dynamically observed values of the option's delta (for example), that approximately replicates the payoff of the option at maturity. We will investigate this topic in separate publication.

8 Appendix

Table 1: Results for 29 day SP500 Call Option on April 22

Strike Price	Bid-Ask Spread		Implied Volatility	Black-Scholes with const. vol = 0.13	Black-Scholes with vol = prev. day	Static Tree Method	Dynamic Method
700	435.9	437.9	0.99999994	440.4859435	440.4859435	440.8361291	441.6392171
750	386	388	0.99999994	390.5256538	390.5256538	390.8360651	389.4179381
800	336	338	0.99999994	340.565364	340.565364	340.8361125	343.1698244
825	311.1	313.1	0.99999994	315.5852191	315.5852191	315.8360728	317.7046472
850	286.1	288.1	0.99999994	290.6050743	290.6050743	290.8359979	288.6177766
875	261.1	263.1	0.99999994	265.6249294	265.6249294	265.8360605	264.3557368
900	236.2	238.2	0.99999994	240.6447845	240.6447845	240.8361814	242.4258073
925	211.3	213.3	0.99999994	215.6646397	215.6646397	215.8361042	216.070995
950	186.4	188.4	0.99999994	190.6844968	190.6844968	190.8361121	191.9580826
975	161.5	163.5	0.99999994	165.7044244	165.7044244	165.8361814	165.3575168
995	141.7	143.7	0.99999994	145.7211165	145.7211165	145.8365447	144.0436024
1005	131.9	133.9	0.99999994	135.7308655	135.7308655	135.8375831	136.3773022
1025	112.2	114.2	0.99999994	115.7639408	115.7639408	115.8489649	115.8511543
1035	102.5	104.5	0.99999994	105.8005096	105.8005096	105.1462257	104.9907379
1040	97.6	99.6	0.107155383	100.8298139	100.8298139	100.6728878	101.2075435
1050	88	90	0.135827243	90.92671721	90.92671721	91.05013255	90.85553399
1060	78.5	80.5	0.144771874	81.10857475	81.10857475	81.09046311	79.87881581

continued on next page

continued from previous page							
Strike Price	Bid-Ask Spread		Implied Volatility	Black Scholes formula constant vol = 0.13	Black Scholes formula vol = prev. day	Static Tree Method	Dynamic Method
1070	69.1	71.1	0.145104229	71.4350287	71.4350287	71.32386	72.79518986
1075	64.5	66.5	0.14579457	66.67796494	66.67796494	66.42264828	64.8763511
1080	59.9	61.9	0.144722283	61.99094358	61.99094358	61.05169626	64.4851441
1090	51	53	0.142792165	52.88699128	52.88699128	52.64029526	54.52657424
1100	42.4	44.4	0.1389274	44.25484143	44.25484143	43.86203883	44.96215623
1110	34.3	36.3	0.138166845	36.23617352	36.23617352	35.8019852	38.1543003
1115	30.4	32.4	0.133875668	32.50000074	32.50000074	32.01309525	33.06596175
1120	26.7	28.7	0.131962717	28.96654938	28.96654938	28.47139516	31.31540725
1125	24	24.7	0.131792247	25.64879123	25.64879123	25.12064397	26.1069527
1130	20.5	22	0.127446949	22.55710072	22.55710072	21.93006175	24.08772635
1135	17.1	18.6	0.126670897	19.69885882	19.69885882	19.10765002	21.01885359
1140	14.3	15.8	0.123505414	17.07818041	17.07818041	16.43635034	17.63322328
1145	12.2	13.3	0.121538699	14.6957816	14.6957816	14.03991679	15.35886164
1150	9.8	10.8	0.118664801	12.54899328	12.54899328	11.92985116	13.46731373
1155	7.8	8.8	0.118339717	10.63191686	10.63191686	8.807748226	10.92497799
1160	6	7	0.115783632	8.935708498	8.935708498	7.628036313	10.06979747
1165	4.7	5.4	0.116955578	7.448969646	7.448969646	6.944491652	9.177123201
1170	3.5	4	0.113949478	6.158216369	6.158216369	5.63323987	6.099003098
1175	2.7	3	0.113022745	5.048396347	5.048396347	4.606569754	6.228565193
1180	1.9	2.4	0.112161219	4.10342223	4.10342223	3.695063503	4.245166183
1185	1.3	1.8	0.112685382	3.306691969	3.306691969	2.962457681	4.203963659
1190	1	1.35	0.116848767	2.641570893	2.641570893	2.208950922	2.545651178
1200	0.5	0.8	0.117960393	1.641927513	1.641927513	1.417832905	2.068167719
1210	0.4	0.5	0.123467147	0.985039802	0.985039802	0.834163001	1.189949817
1215	0.25	0.4	0.130912602	0.752833946	0.752833946	0.62279235	1.145398155
1225	0.15	0.2	0.130754888	0.428105579	0.428105579	0.339846938	0.642447332

Table 2: Results for IBM Call Option with September Expiration Date on July 19

Strike Price	Bid-Ask Spread		Black-Scholes with const. vol = 0.234	Static Tree Method	Dynamic Method
60	23.8	24	24.0377906947886	24.1653535881753	24.0399236162121
70	13.9	14.1	14.161570849552	14.1653762415814	14.0904229344650
75	9	9.2	9.51642872827289	9.16997303037814	9.11994784069207
80	4.6	4.8	5.57383978902307	4.37994074833957	4.23440233858245
85	1.6	1.65	2.76178935046618	1.05711159482768	0.662309959257055
90	0.35	0.4	1.14035152327993	0.112329477448275	0.0454043560956927
95	0.1	0.15	0.391956502635727	0.00401547052393596	0.0050792747562195

References

Aingworth, D. D., Das, S. R., and Motwani, R. (2003). A simple approach for pricing equity options with markov switching state variables. unpubished.

Batalova, N. V. and Viens, F. G. (2005). Selection of an optimal portfolio with stochastic volatility and discrete observations. *Preprint*.

Black, F. and Scholes, M. (1973). The valuation of options and corporate liability. *Journal of Political Economy*, 81:637–654.

Bollerslev, T. and Zhou, H. (2002). Estimating stochastic volatility diffusion using conditional moments of integrated volatility. *Journal of Econometrics*, 109:33–65.

Chesney, M. and Scott, L. (1989). Pricing european currency options: A comparison of the modified black-scholes model and a random variance model. *Journal of Financial and Quantitative Analysis*, 24:267–284.

Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Options pricing: A simplified approach. *Journal of Financial Economics*, 7:229–263.

Cox, J. C. and Rubinstein, M. (1985). *Options Markets*. Prentice-Hall, Englewood Cliffs, NJ.

Del Moral, P. (2004). *Feynman-Kac Formulae. Genealogical and Interacting Particle Systems with Applications*. Probability and its Applications (New York). Springer-Verlag.

Del Moral, P., Jacod, J., and Protter, P. (2001). The monte-carlo method for filtering with discrete time observations. *Probability Theory and Related Fields*, 120:346–368.

Ethier, S. N. and Kurtz, T. (1986). *Markov Processes: Characterization and Convergence*. John Wiley & Sons, New York.

Fouque, J.-P., Papanicolaou, G., and Sircar, K. R. (2000a). *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge University Press.

Fouque, J.-P., Papanicolaou, G., and Sircar, R. K. (2000b). Mean-reverting stochastic volatility. *International Journal of Theoretical & Applied Finance*, 3(1):101–142.

Heston, S. (1993). A closed form solution for options with stochastic volatilities with applications to bond and currency options. *The Review of Financial Studies*, 6:329–343.

Hilliard, J. E. and Schwartz, A. (1996). Binomial option pricing under stochastic volatility and correlated state variables. *The Journal of Derivatives*, 4:23–39.

Hull, J. and White, A. (1987). The pricing of options on assets with stochastic volatility. *Journal of Finance*, 42:281–300.

Leisen, D. (2000). Stock evolution under stochastic volatility: A discrete approach. *Journal of Derivatives*, 8:9–27.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 7:867–887.

Nelson, D. B. and Ramaswamy, K. (1990). Simple binomial processes as diffusion approximations in financial models. *The Review of Financial Studies*, 3(3):393–430.

Nielsen, J. N. and Vestergaard, M. (2000). Estimation in continuous-time stochastic volatility models using nonlinear filters. *Journal of Theoretical and Applied Finance*, 3(2):279–308.

Ritchken, P. and Trevor, R. (1999). Pricing options under generalized garch and stochastic volatility process. *Journal of Finance*, 54:377–402.

Rubinstein, M. (1985). Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active cboe option classes from august 23, 1976 through august 31, 1978. *Journal of Finance*, 40:455–480.

Sharpe, W. F. (1978). *Investments*. Prentice-Hall, Englewood Cliffs, NJ, 2 edition.

Stein, E. and Stein, J. (1991). Stock price distribution with stochastic volatility: An analytic approach. *The Review of Financial Studies*, 4:727–752.

Stroock, D. W. and Varadhan, S. R. S. (1979). *Multidimensional Diffusion Processes*. Springer-Verlag, Berlin.

Viens, F. G. (2002). Portfolio optimization under partially observed stochastic volatility. In Wells, W., editor, *The 8th International Conference on Advances in Communication and Control*, volume 1-12. Optim. Soft. Inc.

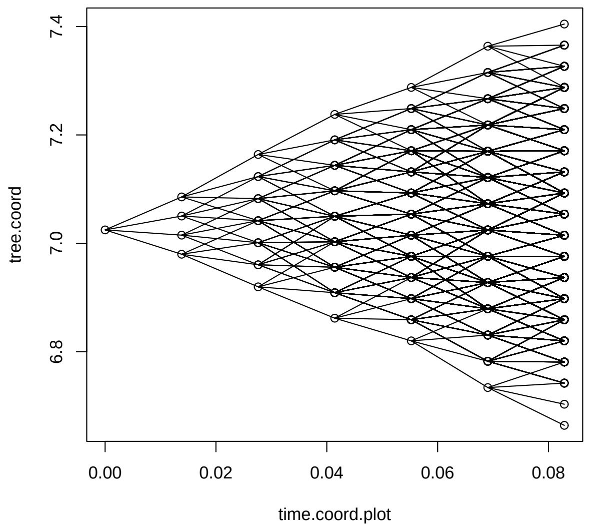

(a) Tree for $p = 0.09$

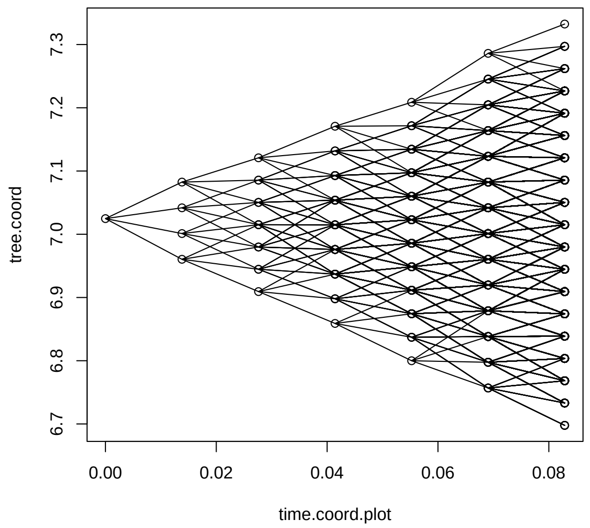

(b) Tree for $p = 0.1$

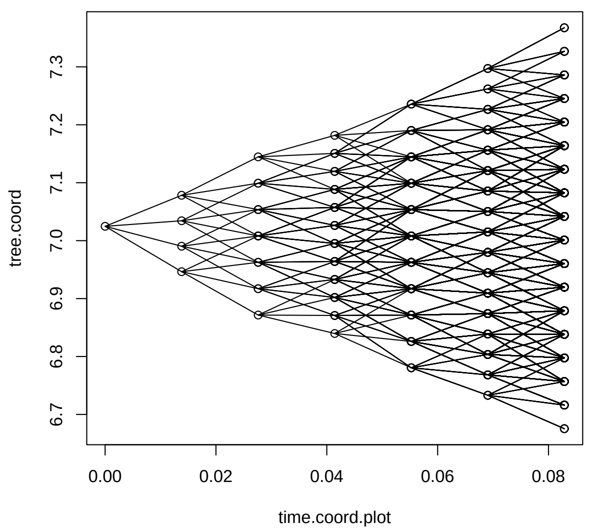

(c) Tree for $p = 0.11$

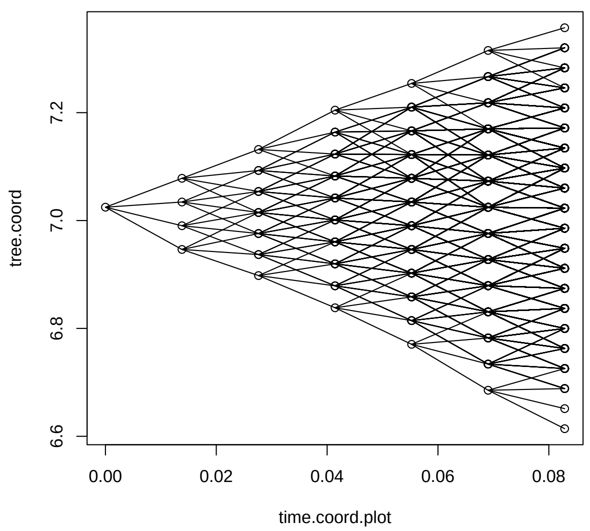

(d) Tree for $p = 0.12$

Figure 3: Example of reduced trees for values of p between 0.09 and 0.12

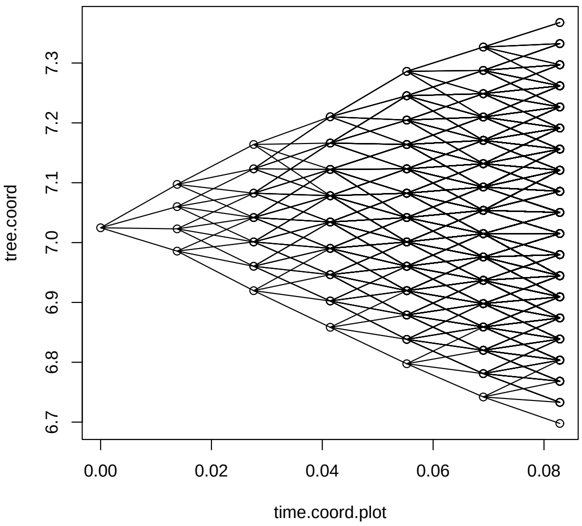

(a) Tree for $p = 0.13$

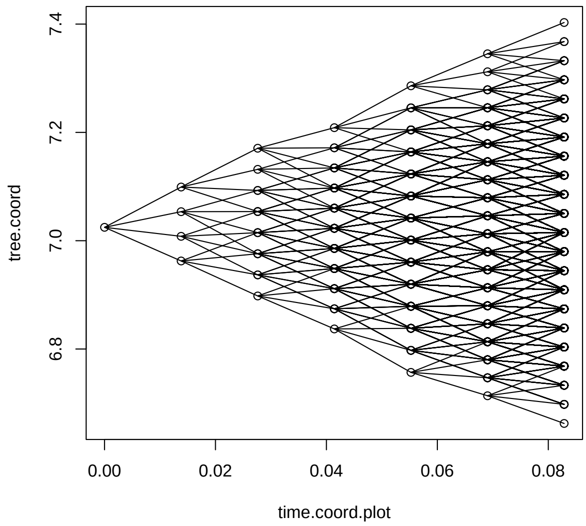

(b) Tree for $p = 0.14$

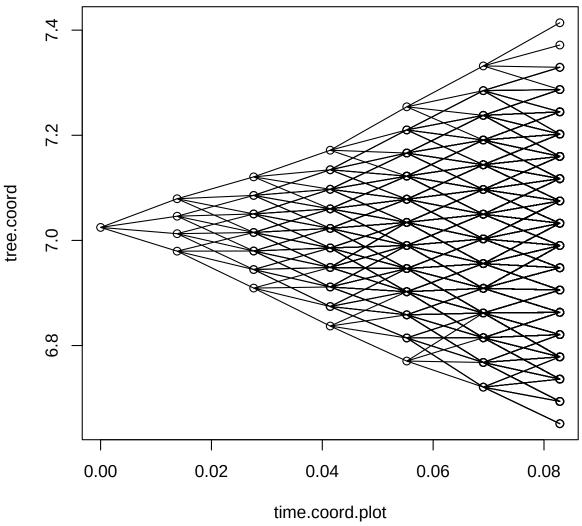

(c) Tree for $p = 0.15$

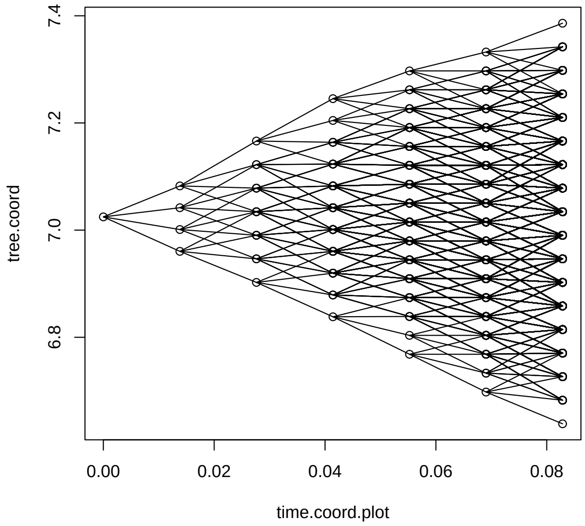

(d) Tree for $p = 0.16$

Figure 4: Example of reduced trees for values of p between 0.13 and 0.16

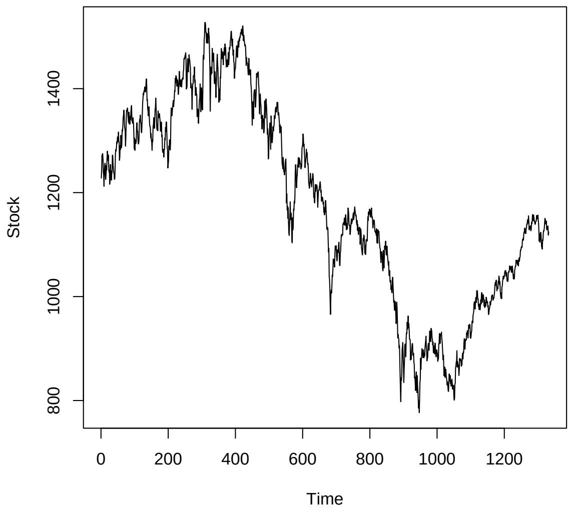

Figure 5: The S&P500 stock price over time

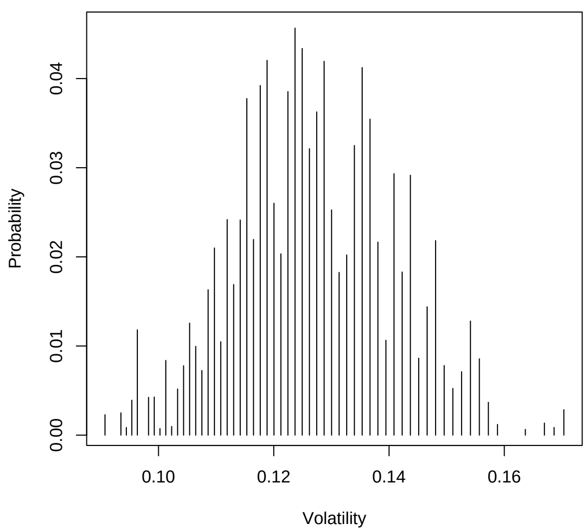

(a) Estimated discrete Volatility Distribution

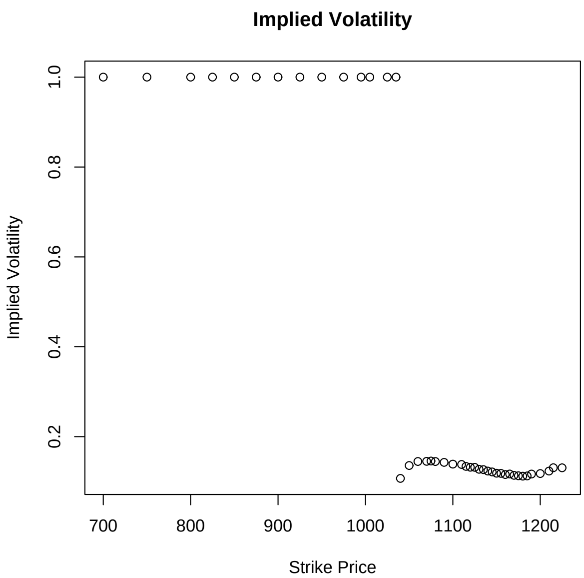

(b) Implied Volatility

Figure 6: Estimates from historical data

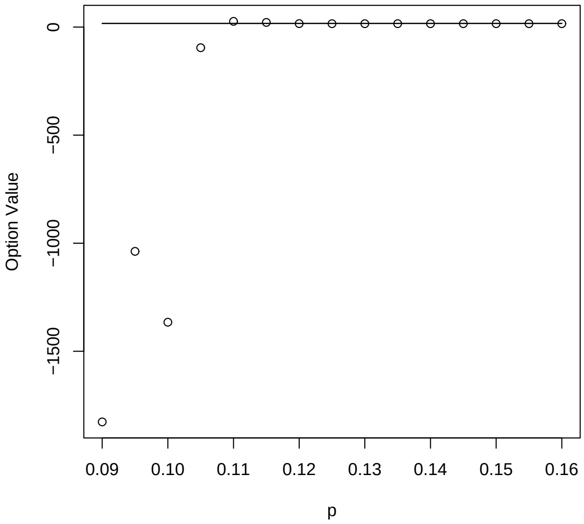

(a) All the values

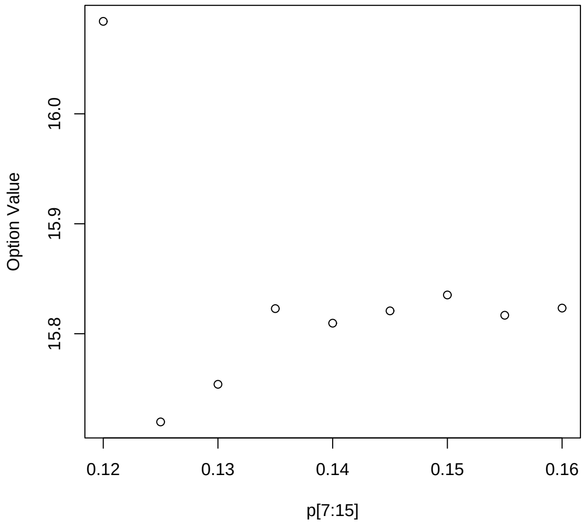

(b) Zoomed in figure

Figure 7: Determining optimal p parameter value

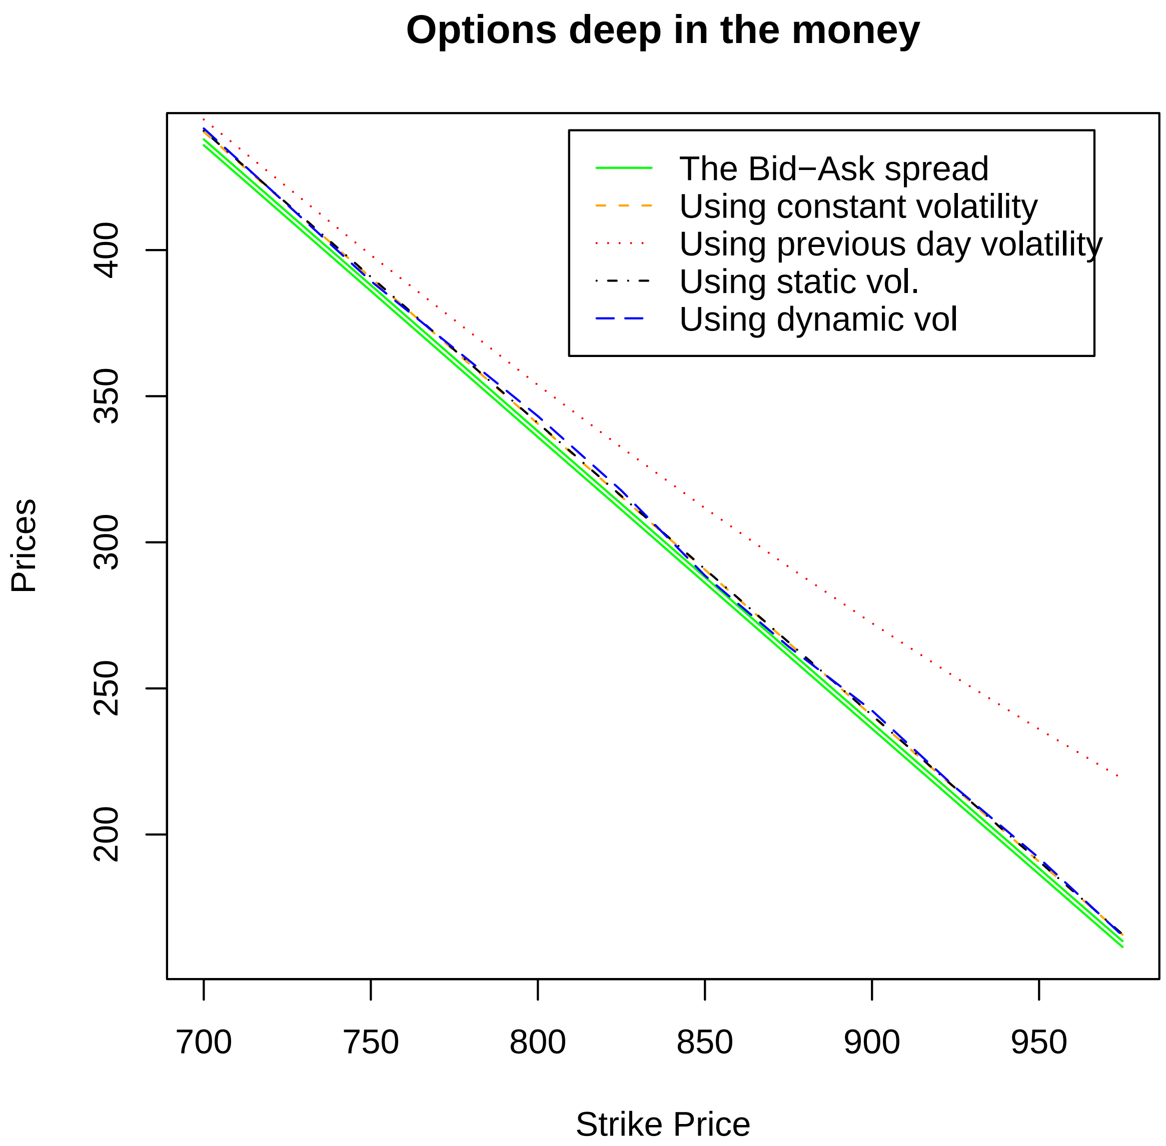

Figure 8: Estimated 29 day option prices: Deep in the money

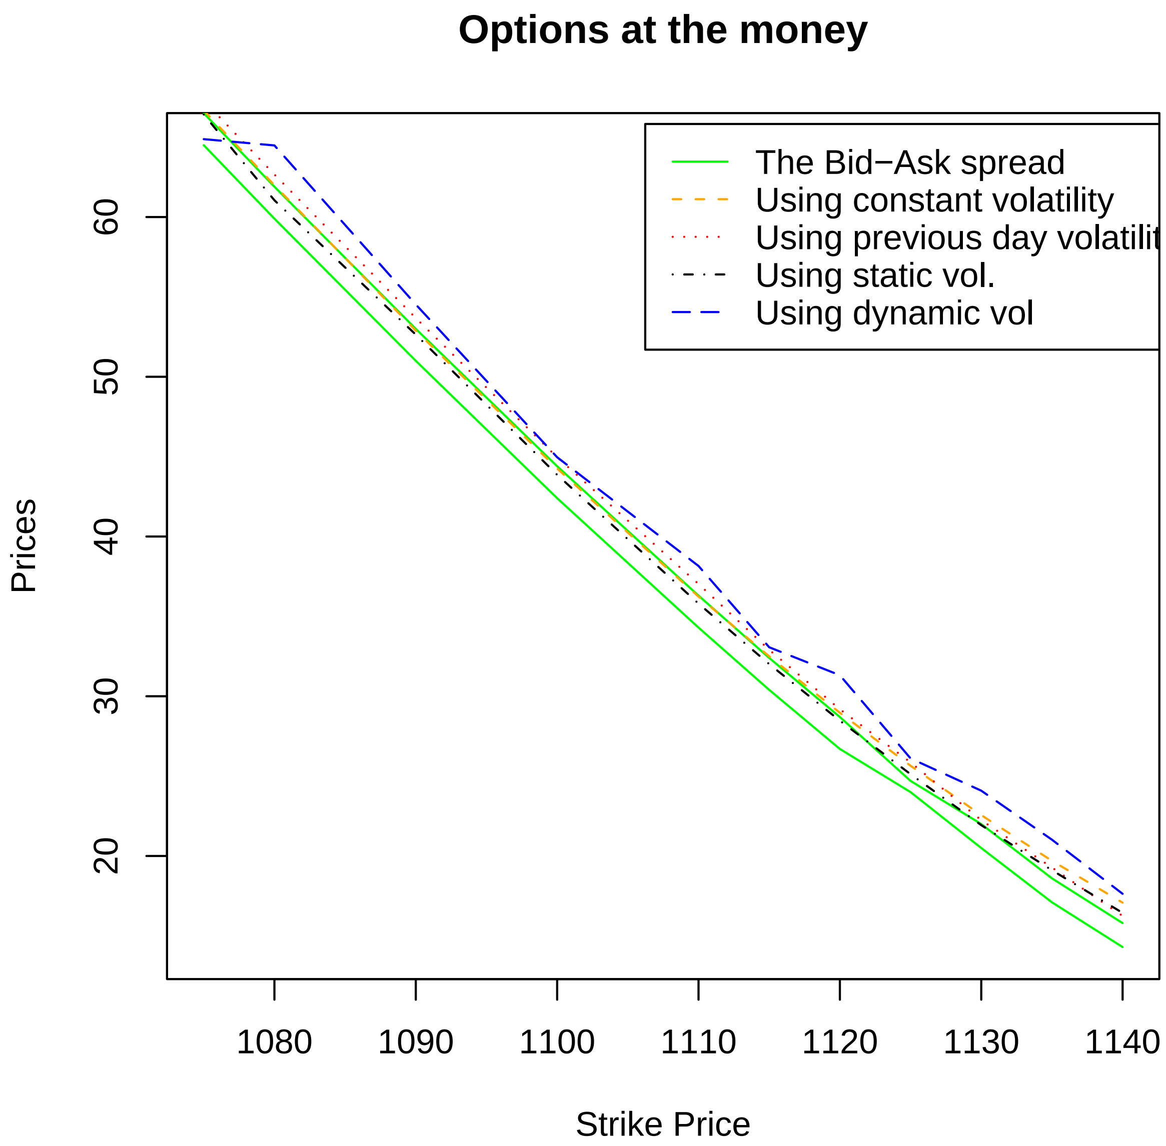

Figure 9: Estimated 29 day option prices: At the money

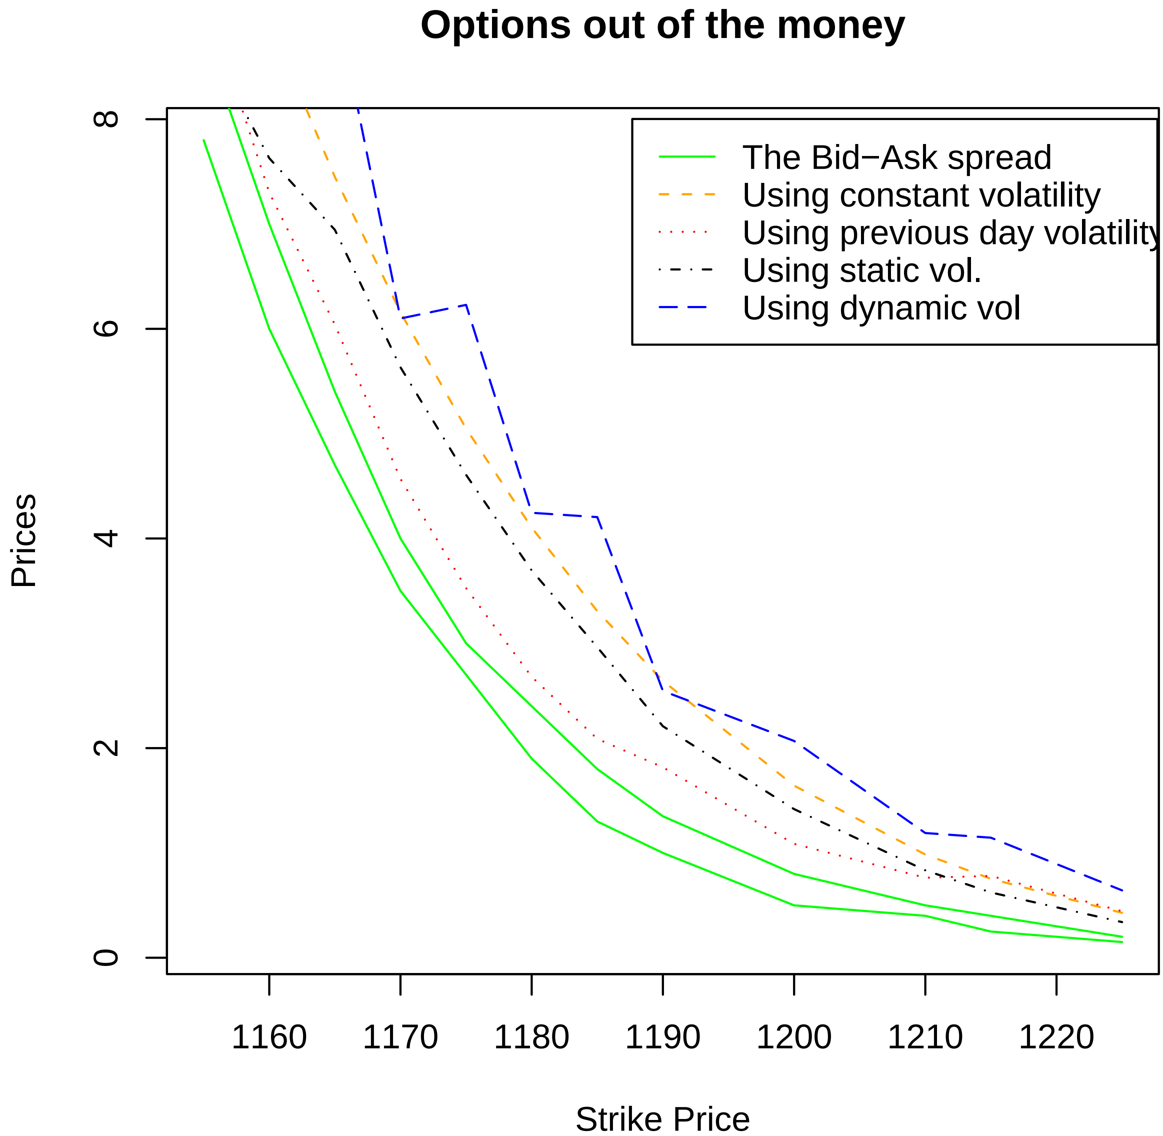

Figure 10: Estimated 29 day option prices: Deep out of the money

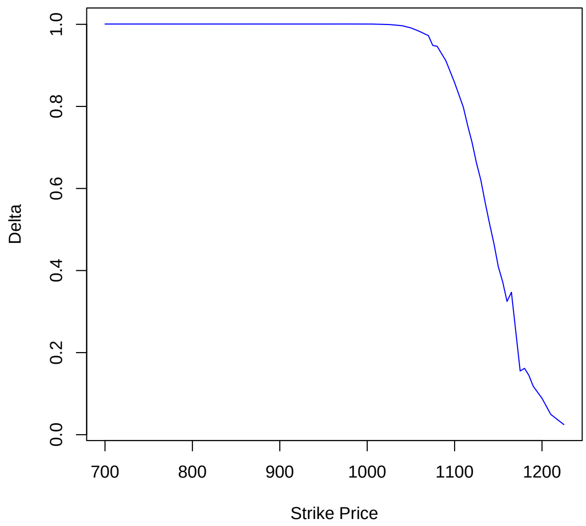

(a) Derivative with respect to stock

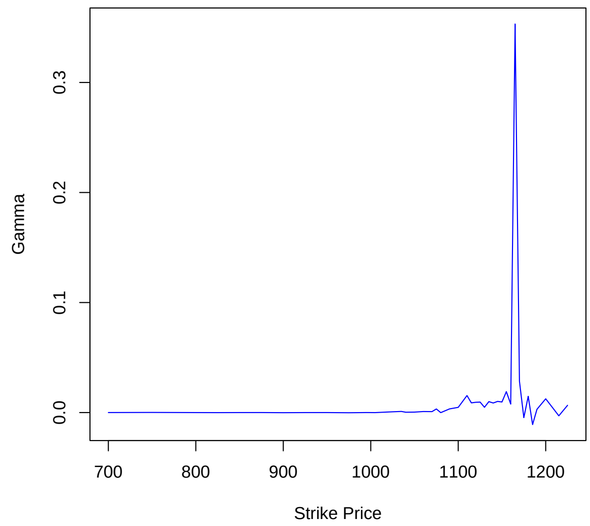

(b) Second derivative with respect to stock

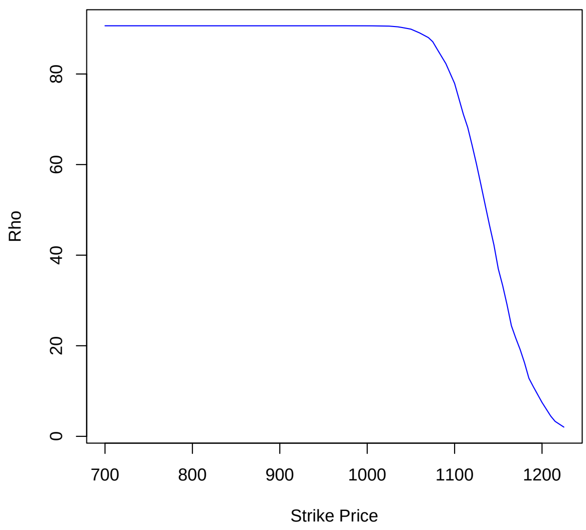

(c) Derivative with respect to interest rate

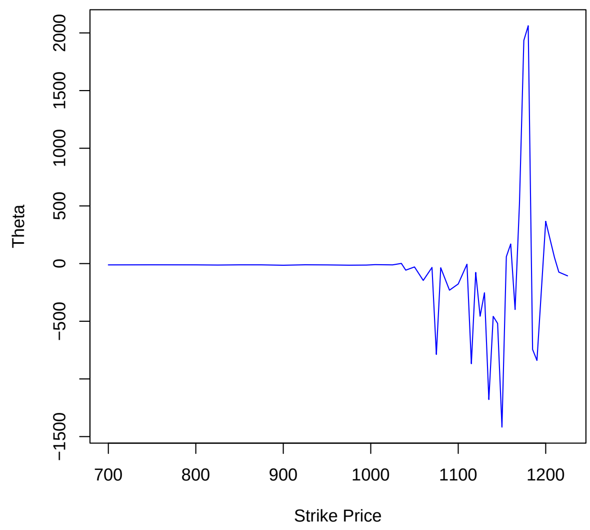

(d) Derivative with respect to time to maturity

Figure 11: The estimated sensitivities

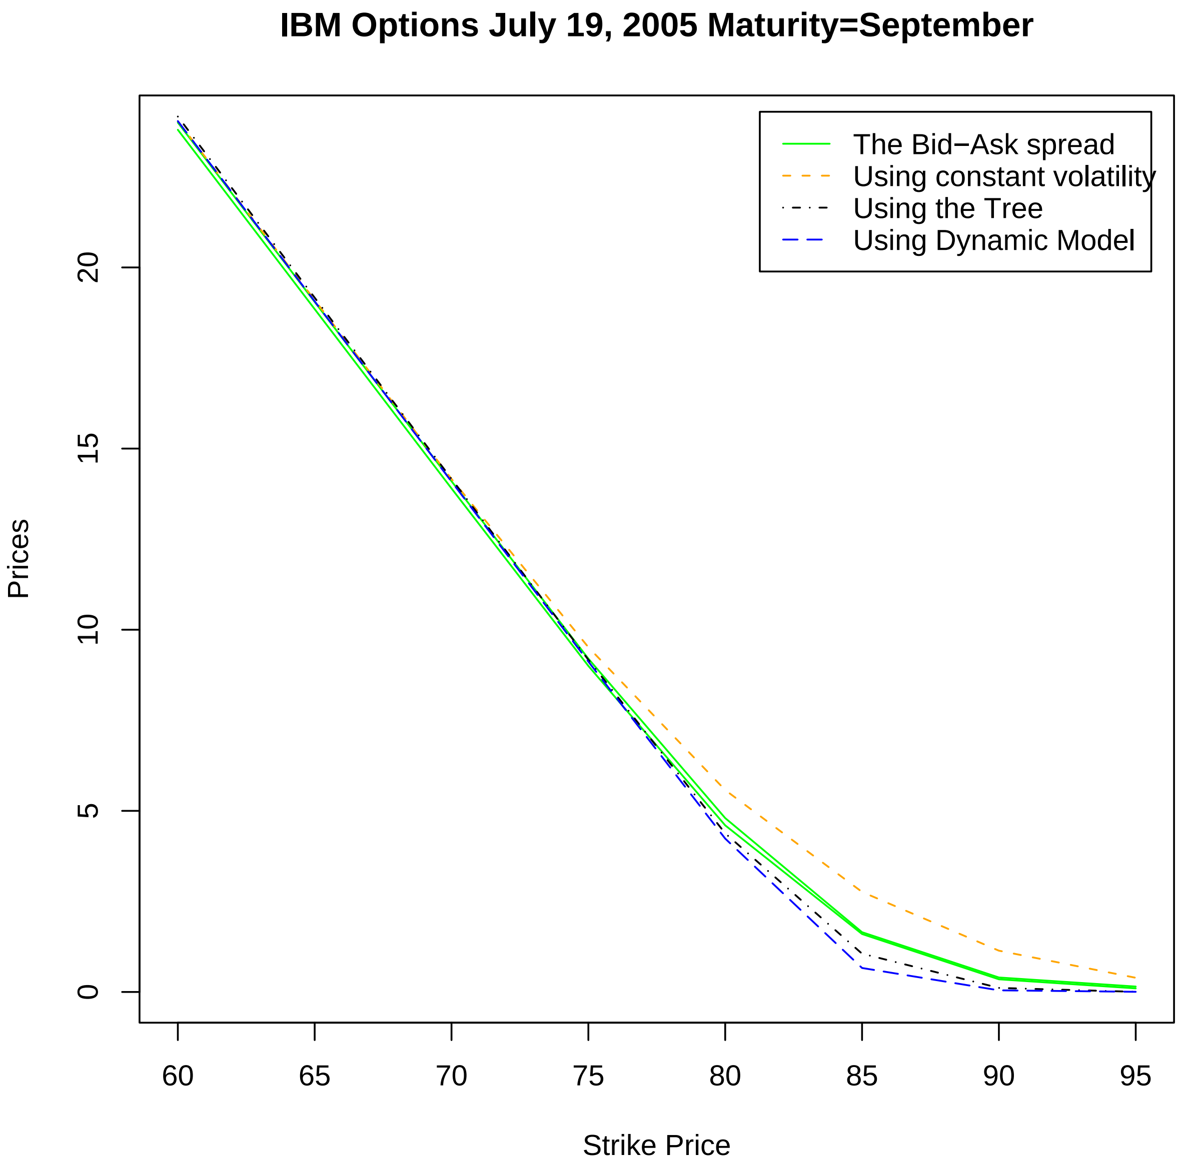

Figure 12: Estimated IBM call option values