FINANCIAL OPTIONS RESEARCH CENTRE

University of Warwick

Simulating the Evolution of the Implied Distribution

George Skiadopoulos
and
Stewart Hodges

December 2000

Financial Options Research Centre
Warwick Business School
University of Warwick
Coventry
CV4 7AL
Phone: (0)24 76 524118

FORC Preprint: 2000/111

Simulating the Evolution of the Implied Distribution*

George Skiadopoulos† Stewart Hodges‡

Abstract

Motivated by the implied stochastic volatility literature (Britten-Jones and Neuberger (1998), Derman and Kani (1997), Ledoit and Santa-Clara (1998)) this paper proposes a new and general method for constructing smile-consistent stochastic volatility models. The method is developed by recognizing that option pricing and hedging can be accomplished via the simulation of the implied risk neutral distribution. We devise an algorithm for the simulation of the implied distribution, when the first two moments change over time. The algorithm can be implemented easily, and it is based on an economic interpretation of the concept of mixture of distributions. It can also be generalized to cases where more complicated forms for the mixture are assumed.

JEL classification: G13.

Keywords: Smile-Consistent stochastic volatility models, Implied Distribution, Mixture of Distributions, Simulation.

1 Introduction

There are two approaches that can be taken for the construction of an option pricing model. The first approach specifies the process for the underlying asset in advance, and then the model is developed. The simplest example of this approach is the Black-Scholes [6] model, which assumes a lognormal diffusion with constant volatility for the underlying. Richer processes which include either stochastic volatility (see among others Hull and White [22], Johnson and Shanno [24], Scott [32]), or jumps (see Bates [3] Merton [28]), or both (see Bates [4], [5], Scott [33]), can give rise to implied volatility patterns which are somewhat similar to the ones observed. However, none of these models fully explains the empirically observed implied volatilities (see Das and Sundaram [15], Taylor and Xu [34]).

*We would like to thank participants at the 2000 CEMAF Conference in Lisbon, the 2000 European Financial Management Association Conference in Athens, and the 2000 European Finance Association Conference in London, and especially the discussants Peter Carayanopoulos and Nicole Branger for their useful suggestions. We would also like to thank Abhay Abhyankar, Silio Aparicio, Anthony Neuberger, Alessandro Rossi, Mark Rubinstein and Nick Webber for helpful discussions and comments. Financial support from the Corporate Members of FORC is gratefully acknowledged. Any remaining errors are our responsibility alone.

†Associate Research Fellow at the Financial Options Research Centre, University of Warwick. Address: Chrisanthemon 76, N. Iraklio 14121, Athens, Greece; e-mail: gskiadopoulos@hotmail.com

‡Director of the Financial Options Research Centre, Warwick Business School, University of Warwick, Coventry CV4 7AL, UK; Tel: 01203-523606, e-mail: S.D.Hodges@warwick.ac.uk

The second approach reverses the option pricing problem. Rather than specifying the asset process exogenously, it starts from the observed European option prices and it implies the process from them (implied process). There are two streams of literature falling within this approach. The first stream, extracts a deterministic volatility implied process (see Derman and Kani [16], Dupire [20], Rubinstein [31], Jackwerth [23]). The second, derives a no-arbitrage stochastic volatility implied process (see Britten-Jones and Neuberger [10], Derman and Kani [17], Ledoit and Santa-Clara [26]).

The deterministic volatility implied process ensures an exact fit of the current observed smile, but is not able to account for its stochastic evolution (see Dumas, Fleming, and Whaley [19], and Buraschi and Jackwerth [11]). This drawback can be dealt with a stochastic volatility implied process; its evolution should be in a no-arbitrage fashion, so that option pricing to be feasible.

Derman and Kani combine an implied trinomial tree with Monte-Carlo simulation by starting from today's local volatilities. Ledoit and Santa-Clara, start from today's implied volatilities and they derive the risk-neutral implied volatility process; Britten-Jones and Neuberger construct a trinomial tree under stochastic volatility which is consistent with today's option prices. These models are very promising, but their implementation is subject to some theoretical and practical limitations. Derman and Kani's model is very computer intensive. Britten-Jones and Neuberger's is based on the assumption that the asset's transition probabilities can be decomposed, in a multiplicative way, into two other functions; in Ledoit's and Santa-Clara's, the implementation is very sensitive to the chosen interpolation method for the implied volatility surface.

In this paper, we propose a new and general method for constructing smile-consistent stochastic volatility models. Rather than simulating the evolution of implied, or local volatilities, *we simulate the evolution of the implied risk-neutral distribution* through time. The simulation starts from today's implied distribution, and proceeds in an arbitrage free way. This type of approach provides a natural tool for risk assessment and for applications in pricing and hedging. However, even though the way to extract the risk-neutral distribution from European option prices has been studied extensively (see for a survey Bahra [2], and Mayhew [27]), we are not aware of any research so far, on its simulation.

In principle, we can calculate the European option price by integrating the payoff of the option over the risk-neutral distribution. Hence, the simulation of the implied distribution enables us to generate price paths for the options having the same expiry, simultaneously. In addition, the technique can be used to examine the behavior of hedges. For example, a riskless hedge is constructed by using a deterministic volatility model. Then, its performance can be investigated under a stochastic volatility scenario generated by our model. The technique can also be used for economic policy purposes. Since the implied distributions reveal the expectations of risk-neutral agents, their evolution shows how these expectations change over time. Hence, provided that the subjective and the risk-neutral distributions do not differ a lot, issues like the credibility of economic policy announcements can be addressed (see Bahra [2]). On the other hand, the proposed algorithm can not be used for valuing simultaneously options with different expiries and American type products, as is the case with the conventional Monte Carlo simulation.

Our method uses an economic interpretation of the concept of mixture of distributions; the current implied distribution can be regarded as a conditional expectation of the future possible distributions. Using mixture of distributions has the advantage that the simulated probabilities evolve as martingales, i.e. the model does not allow any arbitrage opportunities. The method can be implemented easily. Once the form of the mixtures has been specified, the only input that the algorithm requires is the current implied distribution.

The algorithm is developed in two stages. First, the implied distribution changes due to shocks to its mean. Second, the implied distribution changes because of shocks to the variance. The two shocks are assumed to be orthogonal (in the spirit of Hull and White's model [22]), and therefore they can be studied separately. The simulation of the evolution of the implied distribution can be performed by applying them sequentially. Moreover, the algorithm can be extended to cases where more complicated assumptions for the mixture are considered.

The rest of the paper is structured as follows. In the second section, we explain why the simulation of the implied distribution is relevant for option pricing. In the third section, we describe briefly the idea of mixture of distributions, we provide an economic interpretation of the formal concept, and we present the main steps for the construction of the algorithm. In the fourth and fifth section, we present the algorithm for the simulation of the implied distribution when its mean and variance change, respectively. Numerical examples are presented. The sixth section concludes, and it addresses issues for further research.

2 Simulation of Implied Distributions and Option Pricing

There are two analogies between the conventional Monte Carlo simulation, and the simulation of the implied distribution. In the Monte Carlo simulation, for each simulation run we generate a path for the asset price S_t, until we reach the maturity T of the option. In the simulation of the implied distribution, a path for the evolution of the whole distribution is created, until we reach $t = T$, where the distribution degenerates to a single point S_T. The mean of the probability density function (PDF) delivers the path for S_t. In both cases, we require some thousands of simulation runs to obtain a credible estimate of the option price.

Denote by $\pi_t(S_T)$ the risk-neutral probability, as formed at current time t, that the asset price S_T will be reached at time T $(t < T)$. $\pi_t(S_T)$ can be extracted from the observed at time t market call option prices C maturing at time T. The way to do this comes from the well-known relationship

$$\frac{\partial^2 C}{\partial K^2} = e^{-r(T-t)} f_{t,T}(K) \tag{1}$$

established by Breeden and Litzenberger [9], where K is the strike price, r is the interest rate, and $f_{t,T}(K)$ is the risk neutral PDF conditional on the information at time t. Assuming that we observe option prices for a continuum of strikes, we can calculate the probabilities for reaching every S_T.

The simulation of the implied density, can reveal to us the path of the asset price. This is because the current asset price S_t is proportional to the expected value of S_T (martingale property of the implied distribution). The expectation is conditional on the information at time t, and it is formed with respect to the risk-neutral probability measure $\pi_t(S_T)$ (see

Dothan [18]). Hence,

$$S_t = e^{-r(T-t)} E_t^*(S_T) \tag{2}$$

If we knew what the implied risk neutral probability $\pi_{t+1}(S_T)$ would be in the next time step $t+1$ for $\forall S_T$, then we would be able to calculate the asset price S_{t+1}, as equation (2) shows. In this way we would have in front of us a path of the asset price for all times n for $t < n < T$. Moreover, the simulation of the implied density allows the simultaneous generation of price paths for European options with different strike prices, but having the same expiry T. This is because the price Y_t of any European option with payoff h_Y satisfies

$$Y_t = e^{-r(T-t)} \int h_Y(S_T) f_{t,T}(S_T) dS_T \tag{3}$$

In fact, the implied density at $t+1$ is unobserved at time t; Today's known implied density is perturbed due to a number of shocks which affect its moments; as a result, the next time step implied density emerges. Hence, the simulation of the implied distribution is a multivariate problem since it should take into account these shocks. In the next sections, we perturb the initial implied density by assuming that orthogonal shocks affect its mean and variance. The orthogonality assumption is necessary, so as to convert the multivariate problem to a simpler univariate one, and to examine the effect of the shocks, separately. The shocks which are orthogonal, can be applied sequentially, and the multivariate problem becomes a sequence of univariate ones. The shock on the mean is the shock on the asset price, and we expect to shift the PDF to the left, or to the right; the shock on the variance should shrink the PDF, as we approach the maturity T.

3 Partitions, Mixtures, and the Simulation of the Implied Distribution

In this section, we provide an economic interpretation to the concept of the mixture of distributions (for more details on the theory of mixture of distributions, see Timmerington, Smith, and Makov [35]). Then, we present, in a general context, the algorithm with which we can perturb the implied distribution, when one of its moments changes. A mixture of distributions is assumed for the construction of the algorithm.

3.1 Partitions and Mixtures

Let $f_0(S_T)$ be the PDF of the asset price S_T, as viewed from time 0 (i.e. conditional on the information at time 0), and $f_t(S_T)$ be the PDF of S_T, as viewed from a later time t. The evolution of the initial PDF $f_0(S_T)$ to $f_t(S_T)$ can be regarded as a partition of $f_0(S_T)$. $f_0(S_T)$ consists of infinite many different slices; one of them is the realized $f_t(S_T)$[1]. This is a result from the law of iterated expectations. More specifically, $\Pr[S_T \leq s \mid I_k] = E[1_{S_T \leq s} \mid I_k]$, and $\Pr[S_T \leq s \mid I_t] = E[1_{S_T \leq s} \mid I_t]$, where $k < t$, and the information set $I_k \subset I_t$. Then, from

[1] The area of each slice does not integrate up to one, since it is a subset of the original distribution. Therefore, it has to be scaled so that a proper PDF to be defined.

the law of iterated expectations (see Oksendal [29]) we get that

$$E[E[1_{S_T \leq s} \mid I_t] \mid I_k] = E[1_{S_T \leq s} \mid I_k]$$

Hence, the original PDF at time k, evolves to a distribution on date t, conditional on the information at the subsequent date t. The set of possible conditional distributions has to be constrained so that *the original distribution is seen as a mixture (weighted average) over the possible subsequent conditional ones.* Figure 1 shows how the original PDF is partitioned in two possible slices. One of them is going to be the subsequent conditional PDF (once its area is normalized to one).

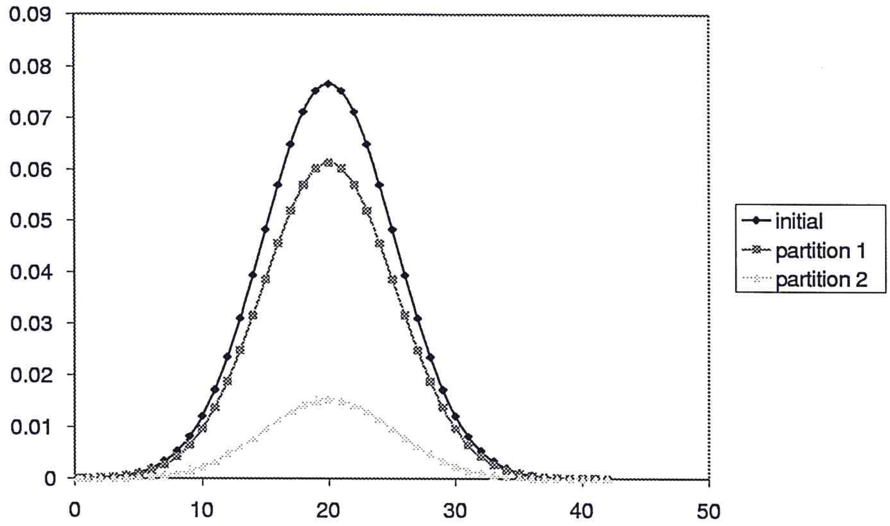

Figure 1: Partitioning of the Original Distribution in two possible realizations.

A mixed distribution is generated from two PDFs. One distribution, called the *structural distribution*, has a parameter which itself is distributed according to a second distribution, called the *mixing distribution*. The structural distribution of a variable X is assumed to have a density $g(X \mid \theta)$, where θ is a parameter which varies across subsets of X. In our context, θ will be one of the moments of the distribution. It is assumed that the distribution of the variable θ is specified by a second distribution, a mixing density $k(\theta)$. The resulting distribution for X is the mixture of the density functions $g(X \mid \theta)$ and $k(\theta)$. This observed, or mixed density is defined by

$$m(X) = \int g(X \mid \theta)k(\theta)d\theta \tag{4}$$

Equation (4) defines a continuous mixture. Similarly, a discrete mixture can be defined as

$$m(X) = \sum_{i=1}^{l} g(X \mid \theta_i)k(\theta_i) \tag{5}$$

Equations (4), and (5) are also satisfied by the cumulative density functions (CDF) $M(X)$, $G(. \mid \theta)$, and the PDF (rather than the CDF) $k(\theta)$ (we denote with capital cases a CDF, and with lower cases a PDF, hereafter), i.e.[2]

[2]We are grateful to Nicole Branger for indicating this point to us.

$$M(X) = \int G(X \mid \theta)k(\theta)d\theta \tag{6}$$

$$M(X) = \sum_{i=1}^{l} G(X \mid \theta_i)k(\theta_i) \tag{7}$$

3.2 Simulating the Implied Distribution

The General Approach We can relate now the concept of a mixture to the one of a partition by regarding $m(X)$ as a model for $f_0(S_T)$[3]. In order to apply this model so that we can simulate the evolution of $f_0(S_T)$ into $f_t(S_T)$, we need to establish a mapping between our model variable X, and the actual variable S_T. In the next section, we describe how we use the distributions of $m(X)$ and $f_0(S_T)$ to accomplish this.

The Details We assume a mixture of distributions, so as to construct the algorithm which simulates the evolution of the implied distribution. The mixture is used as a model for today's implied PDF[4]. The simulation is performed by assuming that one of the moments of the PDF changes, while keeping the other fixed. To simplify the analysis, we simulate the implied distribution over only two time steps.

Let F_0^{-1} be the inverse function of the cumulative density function (CDF) F_0, as viewed from time 0 (today's CDF), i.e.

$$F_0^{-1} : P \rightarrow \; S_T. \tag{8}$$

where P is the cumulative probability, i.e. $P \equiv \Pr[S_T \leq s \mid I_0] = \int_0^s f_0(S_T)dS_T$. Similarly, the inverse of the cumulative distribution as viewed from time t F_t^{-1}, is defined as

$$F_t^{-1} : Q \rightarrow \; S_T \tag{9}$$

where $Q \equiv \Pr[S_T \leq s \mid I_t] = \int_0^s f_t(S_T)dS_1$[5].

The aim is to obtain the next time step inverse function F_t^{-1} when one of the moments of today's PDF changes. We only know F_0^{-1}, as this is implied from the European option prices. However, we can derive F_t^{-1}, by establishing a mapping from Q to P ($Q \rightarrow P$).

[3] $k(\theta)$ is the area (less than one) of one of the possible slices of $m(X)$. It is chosen once the corresponding value of θ has been realized. The structural distribution $g(X \mid \theta)$ is $f_t(S_T)$; it is the slice given by $k(\theta)$, but its area has been normalized to one. The mixing density $k(\theta)$ provides the weights with which the $f_t(S_T)$ are averaged.

[4] Mixtures of distributions have long been used in the finance literature to explain the observed significant skewness, and kurtosis in the distribution of rates of returns (see among others Blattberg and Gonedes [7] and Kon [25]). There, the motivation for using a mixture of distributions comes from the empirical evidence that daily returns deviate from normality more than monthly ones (see among others Blattberg and Gonedes [7], and Fama [21]). The mixture of distributions gives rise to a fat-tailed unconditional distribution with a finite variance, and finite higher moments. Since all moments are finite, the Central Limit theorem applies, and long-horizon returns will tend to be closer to the normal distribution than short-horizon returns (see Campbell, Lo, and Mackinlay [12]). Our work can be viewed as a generalization of this literature to enable us to work with an arbitrary initial pricing kernel.

[5] We prefer using the inverse function of a CDF rather than that of a PDF. This enables us to work in the interval [0,1], rather than in the interval [-∞, +∞].

Then, since we know what F_0 looks like, we can calculate the corresponding asset value S_T from the cumulative probability P. Now we know $Q \to S_T$. The mapping ($Q \to S_T$) is established via four steps.

1. We assume a mixture distribution

$$M(x) = \Pr[X_T \leq x] = \int G(X_T \mid \theta)k(\theta)d\theta \tag{10}$$

with CDF $G(. \mid \theta)$, and PDF $K(\theta)$[6].

The mixing is done over θ which can be the mean, or the variance, or the skewness, etc. We use the mixture $M(X)$ as a model for today's implied distribution, i.e. $M(X) = F_0(S_T)$, and $G(X_T \mid \theta)$ as a model for the new distribution. We have a great deal of flexibility in the choice of the CDFs $G(. \mid \theta)$, and $k(\theta)$ to provide the kind of stochastic evolution we seek. The main requirement for computational purposes is that $G(. \mid \theta)$, $G^{-1}(. \mid \theta)$ and $k(\theta)$ can be calculated reasonably efficient.

2. For any given probability level $Q \in (0, 1)$, we calculate X_T by inverting the CDF $G(. \mid \theta)$ for every Q and for a given realisation of θ, i.e.

$$X_T = G^{-1}(Q \mid \theta) \tag{11}$$

3. From the calculated X_T, we calculate $M(X_T)$ from equation (10). Since $M(X_T) = P$, the mapping $Q \to P$ is established.

4. Since $M(X_T)$ is only a model which may not represent accurately today's implied CDF, we make the pairs S and X equivalent through the relation $M(X_T) = F_0(S_T)$. So, we can scale S from X as $S = F_0^{-1}[M(X_T)]$ (or the less interesting case, X from S as $X = M^{-1}[F_0(S)]$). By doing this, we incorporate in our method the information from today's implied distribution. In order to invert $F_0(S_T)$ we need to approximate $F_0^{-1}(S_T)$ with a CDF whose inverse function has an analytic formula (see Section 3.3).

The whole procedure can be repeated, as described, for any probability values Q are needed.

Figure 2 shows how the mapping $Q \to S_T$ is established via the four steps; it shows, for a given value of θ, the transition from Q to X_T, then from X_T to $M(X_T)$, and eventually from X_T to the actual level S_T.

[6] Alternatively, the mixture can be discrete for the sake of tractability in the numerical implementation.

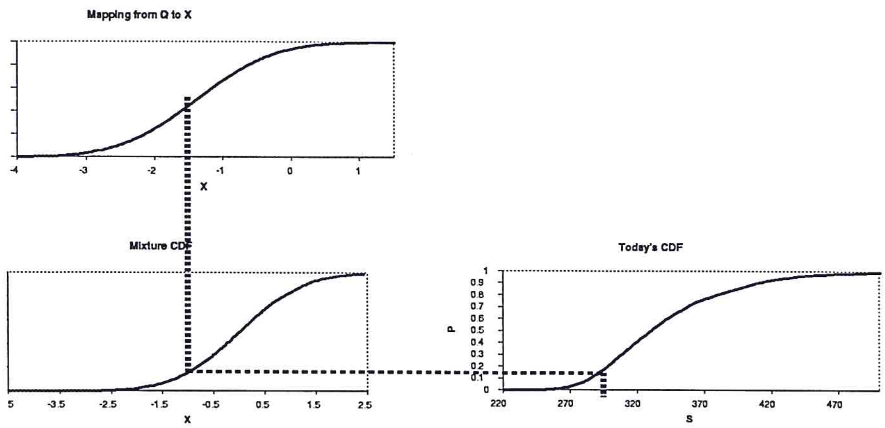

Figure 2: Establishing the Mapping from Q to S_T, using the mixture of distributions $M(X)$ for a given value of θ.

The process for establishing $Q \to P$ is a legitimate because it ensures that the date T risk-neutral distribution for S_T is a martingale, and hence so are all securities which are priced by it, e.g. options and the underlying asset. Figure 3 shows the martingale property of the implied distribution assuming that θ can take only two values, with respective probabilities k_1 and k_2 i.e.

$$P(S_T) = k_1 Q(S_T \mid \theta_1) + k_2 Q(S_T \mid \theta_2)$$

where S_T takes a discrete distribution.

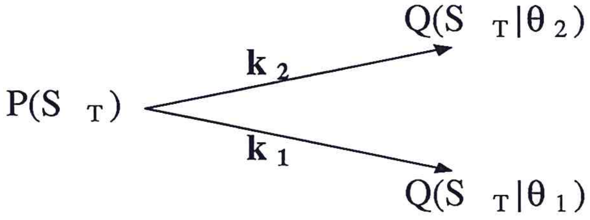

Figure 3: Martingale Property of the Implied Distribution.

The following theorem proves that the algorithm simulates the CDF so that it is a martingale.

Theorem 1 *The established mapping $Q \to S_T$ ensures that the implied distribution evolves as a martingale, i.e. $\int Q(S_T \mid \theta)k(\theta)d\theta = F_0(S_T)$ (martingale property).*

Proof. $\int Q(S_T \mid \theta)k(\theta)d\theta = \int G(X_T \mid \theta)k(\theta)d\theta = M(X_T) = F_0(S_T)$. The last equality follows because X_T is chosen so as $M(X_T) = F_0(S_T)$.

■

The established mapping $Q \to S_T$, simulates the CDF. The following proposition shows how we can simulate the PDF, i.e. perturb $f_0(S_T)$ to $f_t(S_T)$, by using $Q \to P$.

Proposition 1 *The PDF for the underlying asset at the maturity of the option $f_t(S_T)$, as viewed from time t, is given by*[7]

$$f_t(S_T) = f_0(S_T) \times \frac{dQ}{dP} \tag{12}$$

Proof. The mapping $Q \to P$, implies that $Q = Q(P)$. In addition, $P = P(S)$. Hence, $Q = Q(P(S))$. Differentiating Q with respect to S, we have

$$\frac{dQ}{dS} = \frac{dQ}{dP}\frac{dP}{dS} \tag{13}$$

which can be written as

$$f_t(S_T) = f_0(S_T)\frac{dQ}{dP}$$

■

Proposition 2 *The PDF $f_t(S_T)$ calculated from equation (12), integrates up to one, i.e.*

$$\int_{-\infty}^{+\infty} f_t(S_T)dS = 1$$

Proof. Using equations (12) and (13) we have that

$$\int_{-\infty}^{+\infty} f_t(S_T)dS = \int_{-\infty}^{+\infty} f_0(S_T)\frac{dQ}{dP}dS = \int_{-\infty}^{+\infty} f_0(S_T)\frac{1}{f_0(S_T)}\frac{dQ}{dS}dS = \int_{-\infty}^{+\infty} dQ = 1.$$

■

It is worth noting that the mapping $Q \to P$ depends on the postulated mixture given by equation (10) and on the realisation of θ, too. For a different mixture, the mappings $Q \to X$, $X \to M(X)$, and $M(X) \to S_T$, described by Steps 1-4 would have been different. The analytic (or numerical calculation) of the derivative $\frac{dQ}{dP}$ depends on the assumed mixture of distributions, as well. In general, the analytic expression for the derivative is

$$\frac{dQ}{dP} = \frac{\frac{dQ}{dX}}{\frac{dP}{dX}} = \frac{g(X_T \mid \theta)}{m(X_T)} \tag{14}$$

In sections 4 and 5, we study the simulation of the evolution of the implied distribution when either the mean, or the variance changes. Imposing shocks only on the first two moments should make the simulated evolution of the implied distribution to be close to the empirical one. This is because the observed skewness may be explained by shifts in the mean, and the observed fat tails are consistent with shifts in the variance (see Kon [25]). However, the algorithm can be easily generalized to other cases, by choosing more complicated assumptions for the mixture.

[7] The evolution of $f_0(S_T)$ to $f_t(S_T)$ can be regarded as a change of measure, and the ratio $\frac{dQ}{dP}$ can be interpreted as the Radon-Nikodym derivative.

3.3 Inverting the Initial Implied Distribution

In principle, we can use any suitably rich way of representing the implied distribution measured from option prices. However, in order to perform our Monte Carlo simulation, we need to have an analytic expression for $F_0^{-1}(S_T)$. We therefore approximate $F_0^{-1}(S_T)$ by using the Ramberg distribution [30]. This is because it is defined analytically in terms of the inverse of the CDF (percentile function), given that we have pre-specified its first four moments. Moreover, it is a four parameter distribution, which includes a wide variety of curve shapes. Hence, it is useful for the representation of data when the underlying model is unknown. All these characteristics make it ideal for Monte Carlo simulation purposes[8]. Its percentile function $R(P) : P \rightarrow S$ is

$$S = R(P) = \lambda_1 + [P^{\lambda_3} - (1-P)^{\lambda_4}]/\lambda_2 \tag{15}$$

where $0 \leq P \leq 1$, is the cumulative probability, λ_1 is a location parameter, λ_2 is a scale parameter and λ_3 and λ_4 are shape parameters. The values for the λ_1, and λ_2 parameters, are given in tables (see Ramberg et al. [30]) for a variable with a mean of zero and a variance of one. To calculate them for a variable with mean μ and variance σ, we use

$$\lambda_1(\mu, \sigma) = \lambda_1(0,1)\sigma + \mu$$

$$\lambda_2(\mu, \sigma) = \lambda_2(0,1)/\sigma$$

The values of λ_3, and λ_4 correspond to values of skewness and kurtosis, and they are given by tables, as well. The PDF corresponding to equation (15) is given by

$$f(x) = f(R(P)) = \lambda_2[\lambda_3 P^{\lambda_3 - 1} + \lambda_4(1-P)^{\lambda_4 - 1}]^{-1} \tag{16}$$

In order to plot the density for given λ_1, λ_2, λ_3, λ_4, we need to evaluate equations (15) and (16) simultaneously, for values of P ranging from zero to one. Then, $f(R(P))$ is plotted on the y-axis versus $R(P)$ on the x-axis.

4 Simulating Changes in the Asset Price

In this section, we present the algorithm for the simulation of the PDF when its mean changes. This corresponds to changes in the asset price. We provide a numerical example, and we check the numerical accuracy of our procedure.

4.1 The Algorithm

In order to study how the PDF changes when the mean changes (keeping all the other moments fixed), we assume that the model variable X follows a standard Brownian motion W_t, i.e. $X_t = W_t$. This implies that the X is distributed normally.

[8] Other rich classes of distributions defined in terms of the asset price such as the Generalized Beta-2 distribution (see Bookstaber and McDonald [8]) are less suitable for the purposes of our study.

Imagine that we stand at time 0, and the terminal time is $T = 1$. From the properties of the standard Brownian motion process, the model's variable as viewed from time 0, ${}_0X_1$, is normally distributed with mean zero, and variance one, i.e.

$$ {}_0X_1 \sim N(0,1) \tag{17} $$

The asset price as viewed from time t, is normally distributed with mean μ_t, and variance $1-t$, i.e.[9]

$$ {}_tX_1 \sim N(\mu_t, 1-t) \tag{18} $$

where μ_t is normally distributed with zero mean and variance t, i.e.

$$ \mu_t \sim N(0,t) \tag{19} $$

This is because under the equivalent martingale measure, the relative model asset price (using the bank account as numeraire) at time t is

$$ X_t = E_t^*[X_1] = \mu_t \tag{20} $$

and μ_t is one of the possible realizations of the Brownian Motion as viewed from time 0 for the time interval $[0,t]$. Then, $Q \to S_T$ is defined through the four steps of the generic procedure that we have already described.

1. Equations (17), (18), and (19) make us to choose the continuous mixture

$$ F_0(S_T) = M(X_T) = \int_{-\infty}^{\infty} N(X_T \mid \mu_t, 1-t) n(\mu_t \mid 0,t) d\mu \tag{21} $$

as a model for the implied distribution at time 0[10].

2. Since $\mu_t \sim N(0,t)$, we define the mapping

$$ N^{-1}(U \mid 0,t) : U \to \mu_t \tag{22} $$

where $U \equiv \Pr[\mu_t \leq \mu]$, and $N^{-1}(U \mid 0,t)$ is the inverse of the normal CDF when a value for the probability mass U accumulated under μ, is drawn.

For a given drawing from the random variable U (i.e. for a given value of μ_t), we get by inverting the normal CDF in relationship (18) that

$$ X_T = N^{-1}(Q \mid \mu_t, 1-t) \tag{23} $$

Equation (23) is applied for every Q, with $Q \in [0,1]$.

---

[9] Equation (18) shows that the statement that we allow only for a change in the mean of the distribution, keeping all the other moments fixed, is not strictly true. There is a change in the variance, as well, due to the passage of time. However, we can control for it by choosing t infinitesimally small.

[10] This is because if a variable $X \sim N(0,1)$, and $\mu \sim N(0,t)$, then

$$ M(X) = \int N(X \mid \mu, 1-t) N(\mu \mid 0,t) d\mu $$

This follows from the properties of the convolution of the PDFs of two normally distributed variables.

3. We calculate $M(X_T)$ as

$$M(X_T) = N(X_T \mid 0, 1) \tag{24}$$

Since $M(X_T) = F_0(S_T)$, this establishes the mapping $Q \to P$, i.e.

$$P = N[N^{-1}(Q \mid \mu_t, 1-t) \mid 0, 1] \tag{25}$$

4. We scale S_T from X, as $S_T = F_0^{-1}[M(X_T)]$.

$f_t(S_T)$ is calculated from equation (12). The derivative $\frac{dQ}{dP}$ is calculated analytically by using equation (14), and our assumed mixture, i.e.[11]

$$\frac{dQ}{dP} = \frac{n(X_T \mid \mu_t, 1-t)}{n(X_T \mid 0, 1)}$$

4.2 A Numerical Example

In this section, we provide a numerical example of the simulation of an implied distribution, when the mean changes. The simulation is performed by using the mapping Q to P obtained from our assumed mixture, given by equation (21). The mapping has been established for $t = 0.2$. Then, we calculate $f_0(S_T)$ via the Ramberg PDF, by making an assumption about the values of its first four moments. Finally, we perturb $f_0(S_T)$, by using equation (12), in order to see what it looks like 0.2 units of time later.

Figure 4 shows the mapping from Q to P, for random drawings U= 0.001, 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.95, 0.99, and 0.999, i.e. for several values of μ_t, when $t = 0.2$.

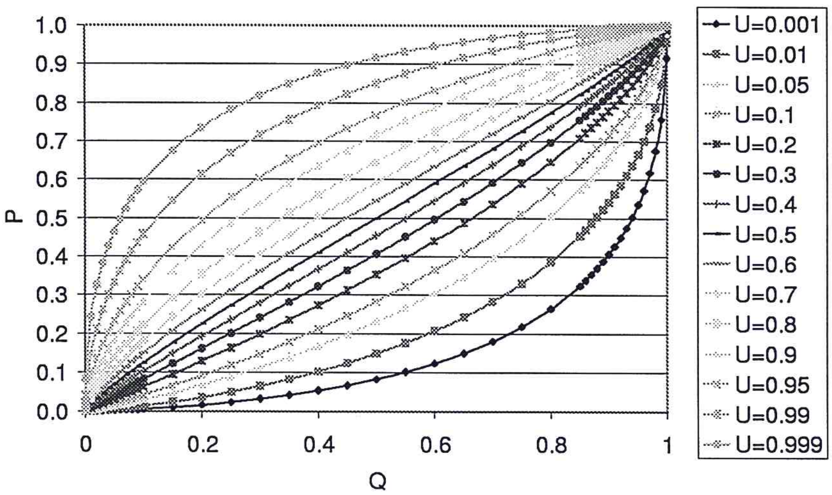

Figure 4: Mapping from Q to P when the Mean changes for t=0.2

[11] We have also tried calculating the derivative numerically, by using several numerical schemes (see Chapra and Canale [13]). However, the simulated PDFs were not smooth enough, because of the numerical errors.

The mapping shows that for the small values of U, $f_t(S_T)$ is shifted to the left of the initial $f_0(S_T)$, while for high values of U, $f_t(S_T)$ is shifted to the right. This is because for a given Q, U increases with P. An increase in U, increases μ_t. This shifts $F_t(S_T)$ to the right, since μ_t is its mean. Therefore, the values that S_T can take are higher, and consequently the probability mass P accumulated under them is larger, i.e.

$$U_1 < U_2 \Rightarrow \mu_{U_1} < \mu_{U_2} \Rightarrow S_{U_1} < S_{U_2} \Rightarrow P_{U_1} < P_{U_2}$$

Furthermore, the increase in P as U increases will be bigger, the bigger t is. This is because the variance of the distribution from which μ_t is drawn, increases when t increases. Hence, bigger shifts in the PDF are expected to happen. Therefore, for a given Q, an increase in U will produce a bigger increase in P, than with smaller t.

Another feature of the figures is that for a given U, P is increasing in Q. This is because increasing Q, increases S_T, and therefore P increases, as well.

For our numerical example, we assume that the mean of the initial implied PDF is 340, its variance is 2500, its skewness is 1, and its kurtosis is 4.2. From the tables of the Ramberg distribution, these values correspond to $\lambda_1 = -0.787$, $\lambda_2 = 0.1142$, $\lambda_3 = 0.0212$, and $\lambda_4 = 0.1244$. Then, we adjust λ_1, and λ_2 for the mean and variance.

Figures 5, 6, and 7 show $f_0(S_T)$, and $f_t(S_T)$. They are produced for $U = 0.01$, 0.5, and 0.95, respectively. The figures show that as U increases, the new density shifts to the right. This reflects the property of the mapping that we have already explained, i.e. that P increases with U, for a given Q.

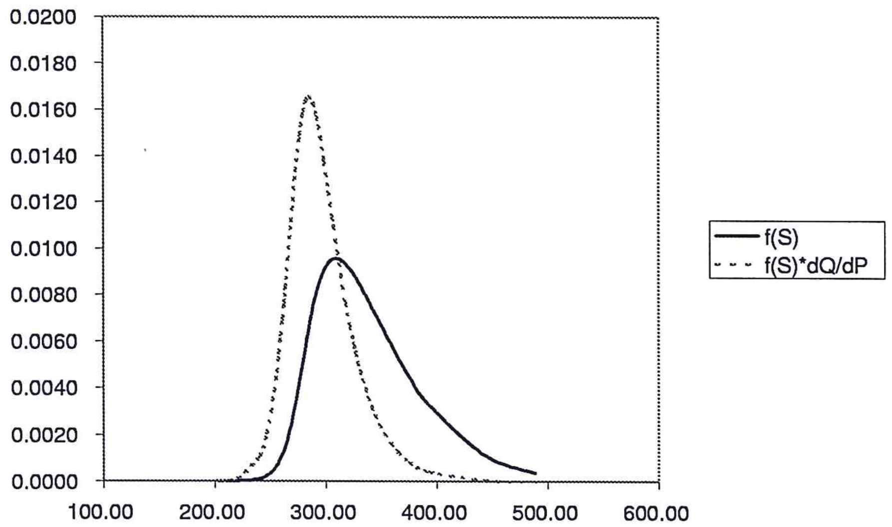

Figure 5: Initial and New PDF for U=0.01.

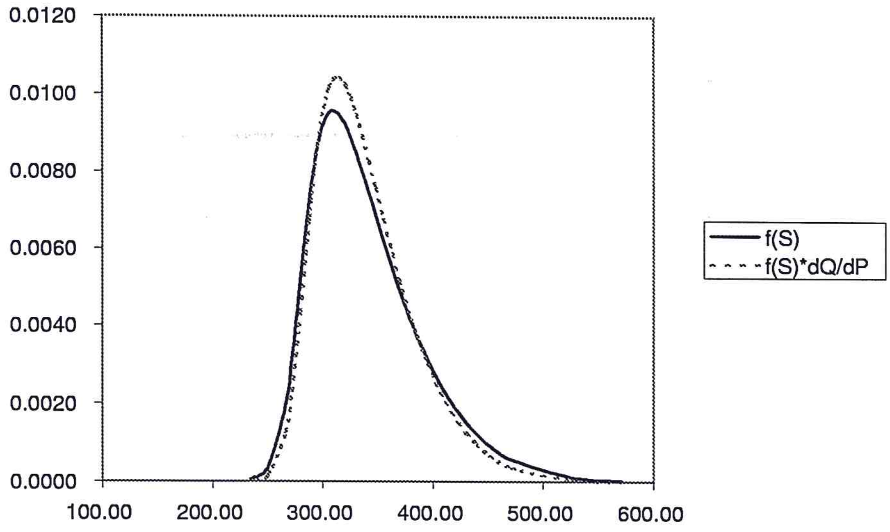

Figure 6: Initial and New PDF for U=0.5.

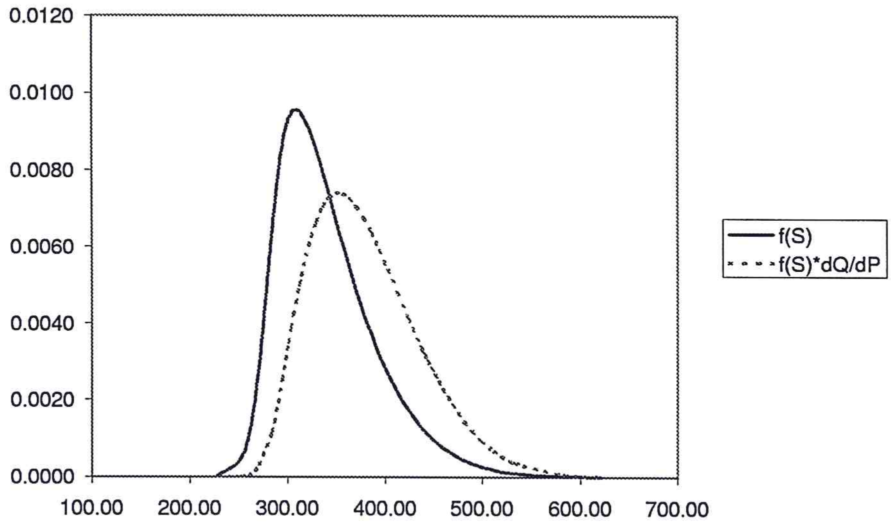

Figure 7: Initial and New PDF for U=0.95

4.3 Checking the Accuracy of the Transformation

The mapping $Q \to P$ has been established for discrete values of Q and P, while the assumed mixture given by equation (21) is a continuous one. This may incur numerical errors in our simulation procedure. In this section we check for this by assuming that $S_T \sim N(0,1)$. Using $Q \to P$, we calculate $f_0(S_T)$ by means of the Ramberg distribution. Then, we calculate $f_t(S_T)$ by means of equation (12). Finally, we compare $f_t(S_T)$ with the normal PDF when $S_T \sim N(\mu_t, 1-t)$. The comparison is shown in Figures 8, 9, and 10, for the cases that $U = 0.01$, 0.5, and 0.95, respectively.

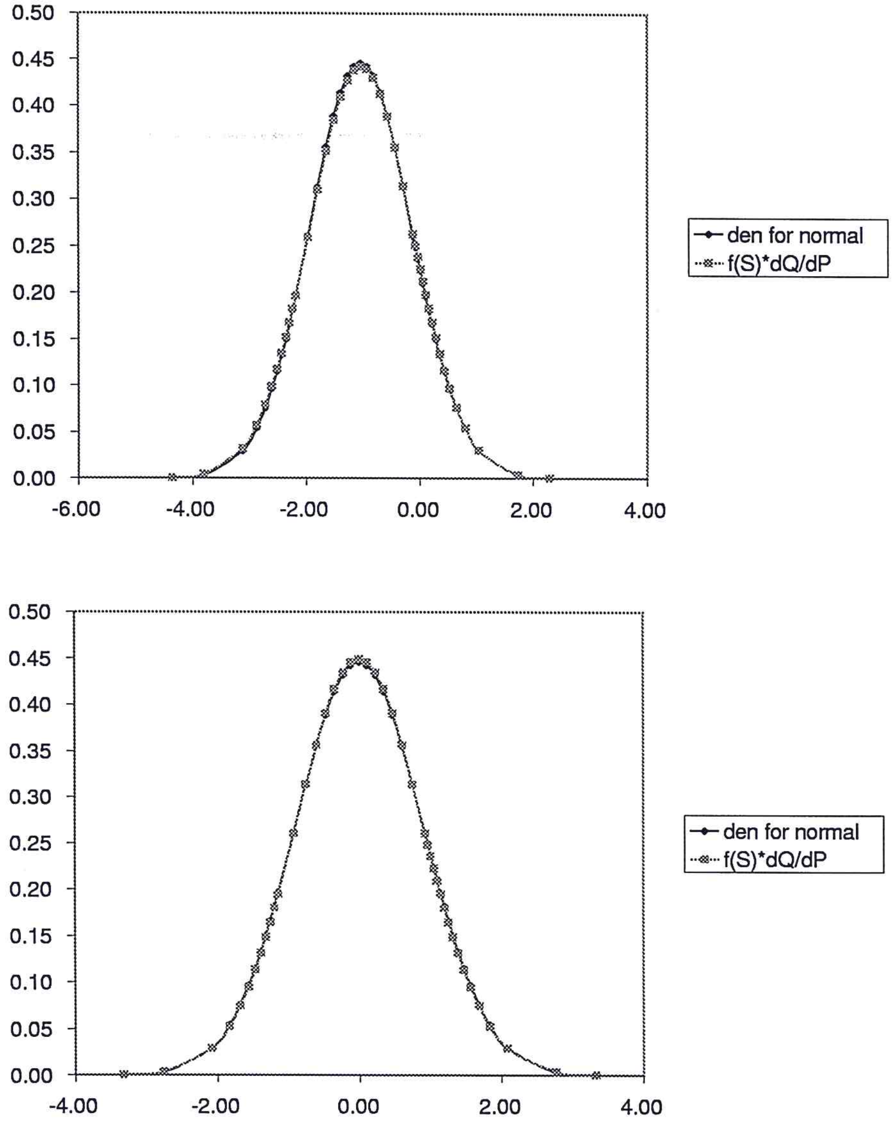

Figure 9: Comparison between the Ramberg-Normal and the Normal PDFs for U=0.5.

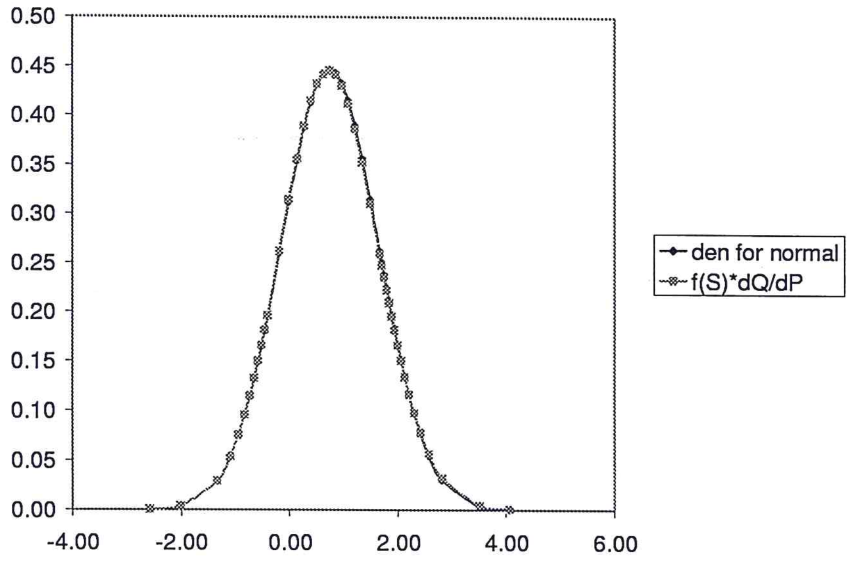

Figure 10: Comparison between the Ramberg-Normal and the Normal PDFs for U=0.95.

It is obvious that the Ramberg-Normal and the Normal PDFs coincide, showing that there is no numerical bias in our procedure.

5 Simulation of the Implied Distribution when the Volatility Changes

In this section, we develop the algorithm for the simulation of the PDF due to a change in its variance σ^2. We show how we determine the evolution of the variance, and we present some numerical examples.

5.1 The Algorithm

We need to establish the mapping $Q \rightarrow S_T$, in order to find F_t^{-1}, just as when the mean changes. However, now the mapping has to take into account the change in the variance; it is established through, the familiar by now, four steps.

1. We introduce the effect of the variance by assuming that the initial implied CDF can be represented by a discrete mixture $M(X_T)$, i.e.

$$M(x) = \Pr[X_T \leq x] = \sum_{i=1}^{l} N(X_T \mid 0, \theta_i) k(\theta_i) \tag{26}$$

where θ_i is the ith realization for the variance variable, $k(\theta_i)$ is the PDF for the variance, and $N(0, \theta_i)$ is the normal CDF of X_T conditional on θ_i occurring. In order to simplify the notation, we set $k_i = k(\theta_i)$.

2. We assume a PDF for the variance for some discrete values of θ, so as to calculate k_i, and θ_i. Then, a θ is drawn from the distribution $k(\theta)$, and we get

$$X_T = N^{-1}(Q \mid 0, \theta)$$

for every Q.

3. For the calculated from the second step X_T, we compute $M(X_T)$.

4. We calculate S_T, since $M(X_T) = P(S_T)$, as $S_T = F_0^{-1}[M(X_T)]$.

Having established the mapping $Q \to S_T$, $f_t(S_T)$ is calculated from equation (12), with the derivative being calculated as

$$\frac{dQ}{dP} = \frac{n(X_T \mid 0, \theta)}{\sum_{i=1}^{l} n(X_T \mid 0, \theta_i) K(\theta_i)}$$

5.2 Calculating the Mean and the Variance of the Variance

In a mixture of distributions with different variances, the kurtosis (Kurt) of the mixture depends directly on the variance of these variances. The choice of our mixture as a model for today's implied distribution, is constrained by the following equations (see Appendix 1 for the proof)

$$E(X) = 0 \tag{27}$$

$$Var(X) = E[\sigma^2] \tag{28}$$

$$Skewness(X) = 0 \tag{29}$$

$$Kurt(X) = 3(1 + (\frac{S.D.(\sigma^2)}{E(\sigma^2)})^2) \tag{30}$$

The factor 3 is the kurtosis of the next time step model distribution, as represented by our normal model. The abbreviation S.D. stands for the standard deviation[12].

Since, the chosen mixture is a model for the observed distribution, we choose $E(X^2) = E[\sigma^2] = \sigma_*^2$, where σ_*^2 is the variance of the observed implied distribution[13]. In addition, we choose the kurtosis of the mixture for X_T (model) to be less than the kurtosis of the observed implied distribution for S_T, $Kurt_*$ i.e.

$$Kurt_t[X] < Kurt_{t,*} \tag{31}$$

Otherwise, we would be likely to generate conditional distributions for S_T which have negative excess kurtosis. This would contradict the empirical evidence which shows that implied volatility smiles are more pronounced for shorter maturity options (Taylor and Xu [34]), i.e. the kurtosis of the implied distribution increases as we approach the maturity of the option.

The inequality in equation (31) implies that

$$3(1 + (\frac{S.D.(\sigma^2)}{E(\sigma^2)})^2) < Kurt_{t,*}$$

[12] Equation (30) shows that the chosen mixture generates a distribution with excess kurtosis. It also shows that the bigger the volatility of volatility, the higher is the kurtosis of the mixture.

[13] Since X is only a model for S, it would not have made a difference, if we had assumed that $E[X^2] = 1$. In that case, we would have constrained the evolution of the implied distribution via the choise of the standard deviation of the variance.

which can be written as[14]

$$S.D.(\sigma^2) < \sqrt{(\frac{Kurt_{t,*}}{3} - 1)\sigma_*^2} \tag{32}$$

5.3 Choosing a PDF for the Variance

We assume that the variance is distributed lognormally. This ensures that it can not take negative values. The variance is distributed with mean α, and variance β^2, i.e. $\sigma^2 \sim \Lambda(\alpha, \beta^2)$. Hence, $\ln \sigma^2 \sim N(\psi, v^2)$. Given that the lognormal distribution is the limit of the binomial distribution (see Cox, Ross, Rubnstein [14]), as the time interval $\delta t \to 0$, we construct a binomial tree for the evolution of the variance with N time steps, which are $\delta t = \frac{T}{N}$ spaced apart. The terminal time step delivers to us the values of θ_i with their associated probabilities k_i.

In order to construct the tree, we choose the upward and downward movement factors for the evolution of σ^2, u, and d ($u > 1$, $d < 1$), in the same way as Cox, Ross and Rubinstein [14] do, i.e. $u = \frac{1}{d}$, with $u = e^{v\sqrt{\delta t}}$, and $d = e^{-v\sqrt{\delta t}}$. Appendix 2 shows that

$$v\sqrt{\delta t} = \sqrt{\frac{1}{N} \ln[1 + \frac{S.D.(\sigma^2)^2}{E(\sigma^2)^2}]} \tag{33}$$

Notice that the construction of the tree for the evolution of the variance, takes into account the information from the current implied PDF.

The transition probabilities p, are calculated by requiring that $E[\sigma^2] = \sigma_*^2$, i.e.

$$pu\sigma_*^2 + (1-p)d\sigma_*^2 = \sigma_*^2$$

which simplifies to

$$p = \frac{1-d}{u-d} \tag{34}$$

Once the tree has been constructed, we use the formula for the binomial distribution to calculate the probabilities k_i for reaching the values θ_i at the terminal level of the tree. This is

$$k_i = \Pr(\sigma^2 = \theta_i) = \frac{N!}{i!(N-i)!} p^i (1-p)^{N-i} \tag{35}$$

for $i = 0, \ldots N$.

5.4 A Numerical Example

We give a numerical example for the construction of the binomial variance tree. We assume that the first four moments of the initial implied distribution are $E(S) = 340$, $Var(S) = 2500$, Skewness=1, and Kurtosis=8. Hence, $E[\sigma^2] = 2500$. Using equation (32), we find that

[14] Hence, we express the standard deviation of the variance in terms of the kurtosis of the actual distribution. We have to restrict ourselves to implied distributions with positive kurtosis $Kurt_*$, so as the square root to be defined.

θ_i	k_i
9266.95	0.02
5487.04	0.1
3248.92	0.26
1923.71	0.34
1139.05	0.22
674.44	0.06

Table 1: Values for the Variance with their associated Probabilities.

$S.D.(Var) < 3227.49$. We choose $S.D.(Var) = 1600$. The tree is constructed with five time steps ($N = 5$), and spacing $\delta t = 0.2$. Hence, the time horizon is $T = 1$. From equation (33), we find $\upsilon\sqrt{\delta t} = 0.26$. Then, $u = 1.3$, $d = 0.769$, and $p = 0.435$ (equation (34)). Figure 11 shows the constructed tree for the evolution of the variance. Table 1 shows the values θ_i with their associated probabilities k_i.

					9266.95
				7130.79	
			5487.04		5487.04
		4222.20		4222.20	
	3248.92		3248.92		3248.92
2500		2500		2500	
	1923.71		1923.71		1923.71
		1480.27		1480.27	
			1139.05		1139.05
				876.48	
					674.44

Figure 11: Constructed Tree for the Evolution of the Variance.

Figure 12 shows the mapping $Q \rightarrow P$ when the variance changes, for these different values of θ_i. We can see that for very high θ_i the initial PDF accumulates less mass than the new simulated one, for the low values of the asset price, while it accumulates more mass for the high asset prices. On the other hand, for the low θ_i the reverse happens.

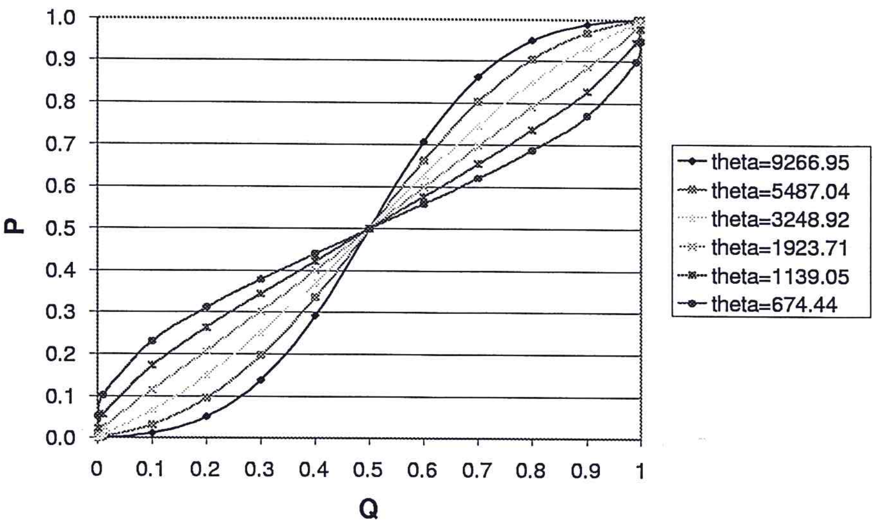

Figure 12: Mapping from Q to P for the different values of θ_i when the Variance changes.

Figures 13-18 show the original and the simulated $f_t(S_T)$ across θ_i. The pictures are consistent with the already discussed implications of the mapping $Q \rightarrow P$.

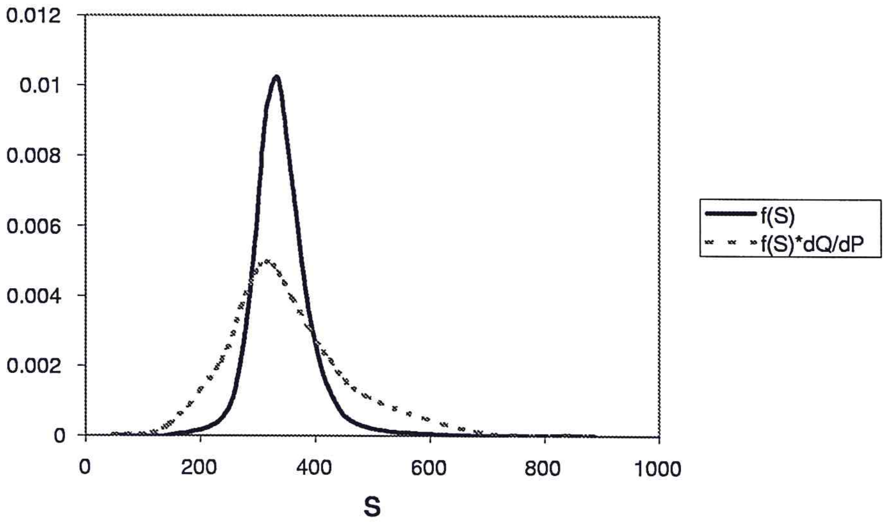

Figure 13: Initial and New PDF for $\theta = 9266.95$.

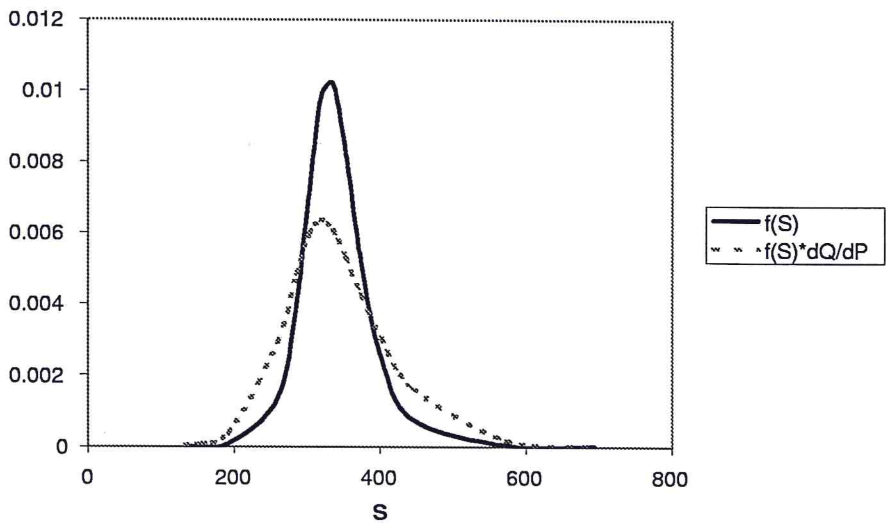

Figure 14: Initial and New PDF for $\theta = 5487.04$.

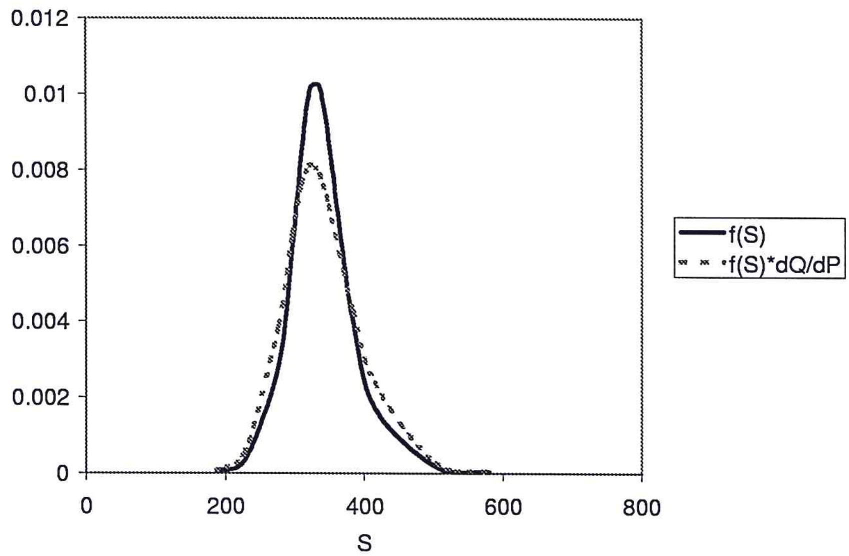

Figure 15: Initial and New PDF for $\theta = 3248.92$.

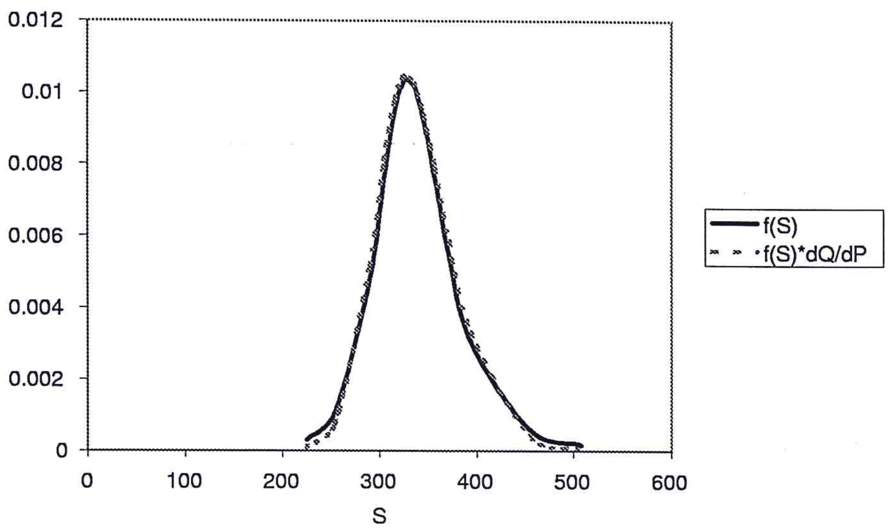

Figure 16: Initial and New PDF for $\theta = 1923.71$.

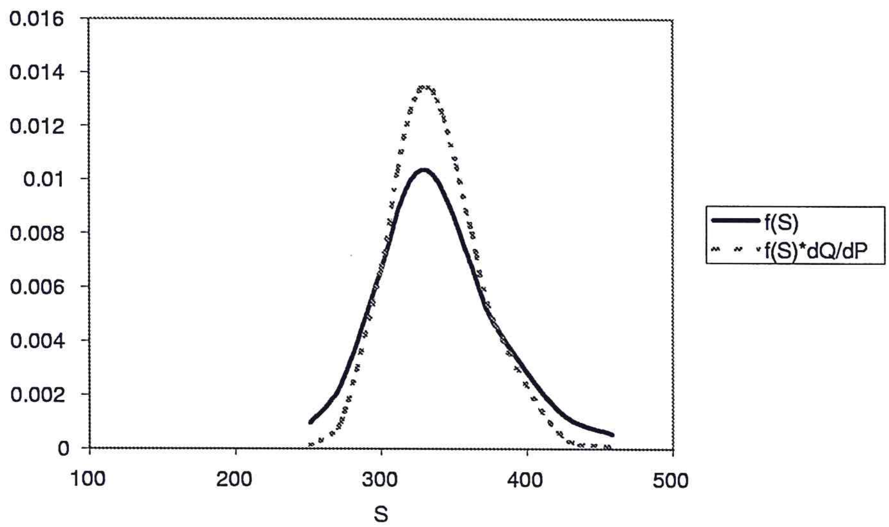

Figure 17: Initial and New PDF for $\theta = 1139.05$.

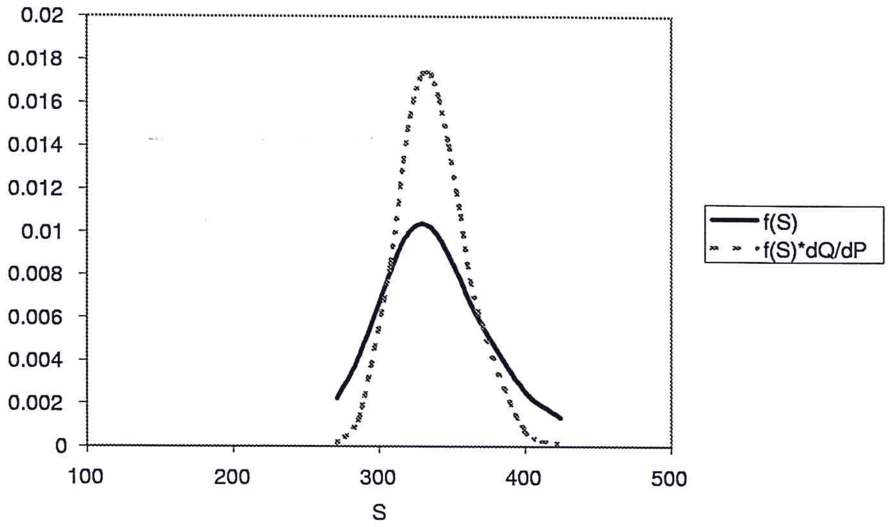

Figure 18: Initial and New PDF for $\theta = 674.44$

6 Conclusions and Issues for Further Research

We have presented a new method for smile consistent option pricing under stochastic volatility: the simulation of the implied risk-neutral probability density function (PDF). We have developed an algorithm for simulating the evolution of the implied PDF, by allowing for orthogonal shocks on its moments. The algorithm relies upon an assumed mixture of distributions. The mixture is used as a model for today's implied distribution. Then, a mapping between today's, and tomorrow's cumulative density function (CDF), for a given value of the asset price, is easily established.

We have dealt with the cases that the mean, or the variance changes, by assuming tractable, for the numerical implementation, mixtures. Some numerical examples are presented. Three are the main advantages of our algorithm. First, the mapping can be established through some very simple calculations. Second, once the form of the mixture has been specified, the algorithm takes as only input today's implied PDF, which we can easily extract from the European option prices. Third, the algorithm can be extended to cases where more complicated specifications of the mixture are assumed.

These advantages make the simulation of the implied PDF a very promising tool for option pricing. Moreover, issues like risk assessing and risk management and the credibility of economic policy announcements can be addressed. However, there is a number of issues that needs to be investigated by future research, in the context of our algorithm.

First, in order to assess the performance of the algorithm, we need to test it empirically. Second, numerical issues like convergence and speed of the algorithm should be considered. Third, the performance of various specifications for the employed mixture should be examined, so that to develop the method into a fully operational tool. This would also help clarifying the relation between stochastic volatility models and mixtures; a way of doing this is by considering the implied by different mixtures underlying's stochastic process and the reformulation of stochastic volatility models along the lines of the algorithm. The dynamics of implied distributions should be explored, as well. This will facilitate the search for the mixture of distributions which will simulate the evolution of the distribution, so as to mimic the evolution of the empirical distribution. Future research should also generalize the method to allow for correlation between the shocks of the asset price and the volatility, and for a jump process for the underlying to be incorporated. Furthrmore, the performance of hedges (constructed by a deterministic volatility model) should be investigated via our algorithm. Finally, the method should be extended in order to price simultaneously options with different expiries. This requires the modeling of the whole probability surface, rather than the modeling of a single-expiry terminal probability distribution.

1 The Moments of a Mixture of (Normal) Distributions, Mixed across the Variance

We calculate the first four moments of the asset's price distribution for a mixture of a general zero mean distribution, mixed across the variance, i.e.

$$m(X) = \sum_{i=1}^{l} k_i f(X \mid 0, \theta_i)$$

where $\sigma_i^2 = \theta_i$. The mean is

$$\begin{aligned} E(X) &= \int Sm(X)dS = \int X \sum_{i=1}^{l} k_i f(X \mid 0, \theta_i)dX = \\ \sum_{i=1}^{l} k_i \int Xf(X &\mid 0, \theta_i)dS = \sum_{i=1}^{l} k_i E^C(X) = 0 \end{aligned}$$

where the superscript C, denotes the expectation of the conditional distribution. The variance is

$$Var(X) = E(X^2) = \int X^2 m(X)dX =$$

$$\int X^2 \sum_{i=1}^{l} k_i f(X \mid 0, \theta_i)dX = \sum_{i=1}^{l} k_i \int X^2 f(X \mid 0, \theta_i)dX$$

$$= \sum_{i=1}^{l} k_i E^C(X^2) = \sum_{i=1}^{l} k_i \theta_i = E[\sigma^2] \tag{36}$$

The skewness is

$$Skewness(X) \equiv \frac{E(X^3)}{[Var(X)]^{\frac{3}{2}}} = \frac{\int X^3 m(X)dX}{[Var(X)]^{\frac{3}{2}}} = \frac{\int X^3 \sum_{i=1}^{l} k_i f(X \mid 0, \theta_i)dX}{[Var(X)]^{\frac{3}{2}}} =$$

$$\frac{\sum_{i=1}^{l} k_i \int X^3 f(X \mid 0, \theta_i)dX}{[Var(X)]^{\frac{3}{2}}} = \frac{\sum_{i=1}^{l} k_i E^C(X^3)}{[Var(X)]^{\frac{3}{2}}} = 0 \tag{37}$$

The last step follows if $E^C(X^3) = 0$ (which is the case for a normal distribution).

The Kurtosis under equation (36) is

$$\begin{aligned} Kurt(X) &= \frac{E(X^4)}{[Var(X)]^2} = \frac{\int X^4 m(X)dX}{(E[\sigma^2])^2} = \\ \frac{\int X^4 \sum_{i=1}^{l} k_i f(X \mid 0, \theta_i)dX}{(E[\sigma^2])^2} &= \frac{\sum_{i=1}^{l} k_i \int X^4 f(X \mid 0, \theta_i)dX}{(E[\sigma^2])^2} = \frac{\sum_{i=1}^{l} k_i E^C(X^4)}{(E[\sigma^2])^2} \end{aligned}$$

Let the kurtosis of the conditional distribution be

$$Kurt^C(X) = Kurt^C = \frac{E^C(X^4)}{[Var^C(X)]^2}$$

and rearranging gives

$$E^C(X^4) = Kurt^C \sigma^4$$

We also have that

$$Var(\sigma^2) = E[\sigma^4] - [E(\sigma^2)]^2$$

Hence,

$$Kurt^M(X) = \frac{Kurt^C \sum k_i \sigma_i^4}{(E[\sigma^2])^2} = Kurt^C \frac{E[\sigma^4]}{(E[\sigma^2])^2} = Kurt^C \frac{Var(\sigma^2) + [E(\sigma^2)]^2}{(E[\sigma^2])^2} = Kurt^C(1+[\frac{\sqrt{Var}}{E[\sigma}$$

If the conditional distribution is normal, then $Kurt^C = 3$.

2 Constructing the Binomial Tree for the Evolution of the Variance

We calculate $\upsilon\sqrt{\delta t}$ by using the properties of the lognormal distribution . The mean and the variance of the lognormal distribution are given by (see Aitchison and Brown [1])

$$\alpha = e^{\psi + \frac{1}{2}\upsilon^2} \tag{38}$$

$$\beta^2 = e^{2\psi + \upsilon^2}(e^{\upsilon^2} - 1) \tag{39}$$

Rearranging expressions (38), and (39) we can calculate υ^2, i.e.

$$\upsilon^2 = \log(\frac{\beta^2}{\alpha^2} + 1) \tag{40}$$

Since σ^2 is distributed lognormally, Var($\ln \sigma_T^2 \mid \sigma_0^2) = \upsilon^2 T$. Therefore,

$$\upsilon^2 T = \ln[1 + \frac{S.D.(\sigma^2)^2}{E(\sigma^2)^2}] \tag{41}$$

The standard deviation of $\ln \sigma_T^2$ per step is

$$v\sqrt{\delta t} = \sqrt{\frac{v^2 T}{N}} \tag{42}$$

From equations (41), and (42) we get that

$$v\sqrt{\delta t} = \sqrt{\frac{1}{N}\ln[1 + \frac{S.D.(\sigma^2)^2}{E(\sigma^2)^2}]} \tag{43}$$

References

[1] Aitchison, J., and Brown, J. (1981): *The Lognormal Distribution with Special Reference to its Uses in Economics*, Cambridge University Press.

[2] Bahra, B. (1997):"Implied Risk-Neutral Probability Density Functions From Option Prices: Theory and Application", Bank of England Working Paper Series.

[3] Bates, D. (1988): "Pricing Options under Jump-Diffusion Processes", Working Paper, Rodney L. White Center.

[4] Bates, D. (1996): "Dollar Jump Fears, 1984-1992: Distributional Abnormalities Implicit in Currency Futures Options", *Journal of International Money and Finance* 15, pp. 65-93.

[5] Bates, D. (1996). "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options", *Review of Financial Studies* 9, pp. 69-107.

[6] Black, F., and Scholes, M. (1973): "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy* 81, pp. 637-654.

[7] Blattberg, R., and Gonedes, N. (1974): "A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices", *Journal of Business* 47, pp. 274-80.

[8] Bookstaber, R., and McDonald, J. (1987): "A General Distribution for Describing Security Price Returns", *Journal of Business* 60, pp. 401-424.

[9] Breeden, D., and Litzenberger, R. (1978): "Prices of State-Contingent Claims Implicit in Option prices", *Journal of Business*, 51 pp. 621-651.

[10] Britten-Jones, M., and Neuberger, A. (1999): "Option Prices, Implied Price Processes, and Stochastic Volatility", *Journal of Finance*, forthcoming.

[11] Buraschi, A., and Jackwerth, J. (1998): "Explaining Option Prices: Deterministic vs. Stochastic Models", Working Paper, London Business School.

[12] Cambell, J., Lo, A., and Mackinlay, A. (1997): *The Econometrics of Financial Markets*, Princeton University Press.

[13] Chapra, S., and Canale, A. (1998): *Numerical Methods for Engineers: With Programming and Software Applications*, Mc Graw Hill, Third Edition.

[14] Cox, J., Ross, S., and Rubinstein, M. (1979): "Option Pricing: A Simplified Approach", *Journal of Financial Economics* 7, pp. 229-263.

[15] Das, S., and Sundaram, R. (1999): "Of Smiles and Smirks: A Term-Structure Perspective", *Journal of Financial and Quantitative Analysis* 34, pp. 211-239.

[16] Derman, E., and Kani, I (1994): "Riding on a Smile", *Risk* 7, pp. 32-39.

[17] Derman, E., M., Kani, I. (1997): "Stochastic Implied Trees: Arbitrage Pricing With Stochastic Term and Strike Structure of Volatility", *International Journal of Theoretical and Applied Finance* 1, pp. 61-110.

[18] Dothan, M. (1990): *Prices in Financial Markets,* Oxford University Press.

[19] Dumas, B., Fleming, J. and Whaley, R. (1998): "Implied Volatility functions: Empirical Tests", *Journal of Finance* 53, pp. 2059-2106.

[20] Dupire, B. (1994): "Pricing with a Smile", *Risk*, 7, pp. 18-20.

[21] Fama, E. (1965): "The Behavior of Stock Market Prices", *Journal of Business* 38, pp. 34-107.

[22] Hull, J. and White,A (1987): "The Pricing of Options on Assets with Stochastic Volatilities", *Journal of Finance* 3, pp. 281-300.

[23] Jackwerth, J. (1997): "Generalized Binomial Trees", *Journal of Derivatives* 5, pp. 7-17.

[24] Johnson, H and Shanno, D. (1987): "Option Pricing when the Variance is Changing", *Journal of Financial and Quantitative Analysis* 22, pp. 143-151.

[25] Kon, S. (1984): "Models of Stock Returns-A Comparison", *Journal of Finance* 39, pp. 147-165.

[26] Ledoit, O., and Santa-Clara, P. (1998): "Relative Pricing of Options with Stochastic Volatility", Working Paper, University of California, Los Angeles.

[27] Mayhew, S. (1995): "On Estimating The Risk-Neutral Probability Distribution Implied by Option Prices", Working Paper, University of California, Berkeley.

[28] Merton, R. (1976): "Option Pricing when Underlying Stock Returns are Discontinuous", *Journal of Financial Economics* 3, pp. 125-144.

[29] Oksendal, B. (1992): *Stochastic Differential Equations: An Introduction with Applications*, Springer-Verlag.

[30] Ramberg, J., Dudewicz, E., Tadikamalla, P., and Mykytka, E. (1979): "A Probability Distribution and its Uses in Fitting Data", *Technometrics* 21, pp. 201-209.

[31] Rubinstein, M. (1994): "Implied Binomial Trees", *Journal of Finance* 69, pp. 771-818..

[32] Scott, L. (1987): "Option Pricing when the Variance changes Randomly: Theory, Estimation and an Application", *Journal of Financial and Quantitative Analysis* 22, pp. 419-438.

[33] Scott, L. (1997): "Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Applications of Fourier Inversion Methods", *Mathematical Finance* 7, pp. 413-426.

[34] Taylor, S., and Xu, X. (1994): "The Magnitude of Implied Volatility Smiles: Theory and Empirical Evidence for Exchange Rates", *Review of Futures Markets* 13, pp. 355-380.

[35] Timmerington, D., Smith, A., and Makov, U. (1985): *Statistical Analysis of Finite Mixture Distributions*, John Wiley & Sons.