

Estimating Option Implied Risk-Neutral Densities using Spline and Hypergeometric Functions*

Ruijun Bu[†]

The University of Liverpool

Kaddour Hadri

Durham University

*The authors would like to thank Gael Martin from Monash University for very constructive and insightful comments. We also would like to thank the editor and an anonymous referee for helpful comments and suggestions that greatly improved the paper. All remaining errors are ours.

[†]Corresponding author: The School of Management, The University of Liverpool, Chatham Street, Liverpool L69 7ZH, UK, Tel: +44-151-795-3122, Fax: +44-151-795-3005, Email: RuijunBu@liv.ac.uk (R. Bu).

Estimating Option Implied Risk-Neutral Densities using Spline and Hypergeometric Functions

Abstract

We examine the ability of two recent methods — the smoothed implied volatility smile method (SML) and the density functionals based on confluent hypergeometric functions (DFCH) — for estimating implied risk-neutral densities (RNDs) from European-style options. Two complementary Monte Carlo experiments are conducted and the performance of the two RND estimators is evaluated by the Root Mean Integrated Squared Error (RMISE) criterion. Results from both experiments show that the DFCH method outperforms the SML method for the overall quality of the estimated RNDs concerning both accuracy and stability. An application of the two methods to the OTC currency options market is also presented.

JEL classification: C14; C15; C52; E58

Keywords: Risk-neutral density; Natural spline; Hypergeometric functions; Root Mean Integrated Squared Error

1. Introduction

Cross sections of observed option prices have long been used to estimate the implied risk-neutral probability density functions (RNDs). Since these RNDs represent forward-looking forecasts of the distributions of the prices of the underlying asset, they prove to be particularly useful for various applications. They are used for pricing complex derivatives; estimating parameters of the underlying stochastic processes — Bates (1996); testing market rationality — Bondarenko (1997); estimating risk preferences — Ait-Sahalia and Lo (2000), Jackwerth (2000), Rosenberg and Engle (2002), Bliss and Panigirtzoglou (2004); and so on. In particular, option implied RNDs have found an extensive use for monetary policy purposes by an increasing number of Central Banks. Examples include Bahra (1997), Söderlind and Svensson (1997), Melick and Thomas (1997), McManus (1999), Jondeau and Rockinger (2000), Söderlind (2000), etc.

Numerous econometric methods have been developed for recovering the implied RNDs. Parametric methods essentially rely on specific assumptions on the data generating process, which depends on some unknown parameters. These parameters can be estimated by optimizing certain objective functions. Examples of parametric methods include: generalized distribution methods of Aparicio and Hodges (1998), Rosenberg (1998) and Lim et al. (2005); expansion methods of Jarrow and Rudd (1982) and Rubinstein (1998); lognormal mixture models of Bahra (1997) and Melick and Thomas (1997); and models for stochastic processes of Heston (1993), Bates (1996) and Wu and Huang (2004). Nonparametric methods, in contrast, are flexible data-driven methods. Examples of nonparametric methods for estimating RNDs include: implied trees of Rubinstein (1994); kernel estimation methods of Ait-Sahalia and Lo (1998, 2000) and Ait-Sahalia et al. (2001); smoothing techniques of Shimko (1993), Malz (1997a,b), Campa et al. (1998) and Bliss and Panigirtzoglou (2002); maximum entropy methods of Buchen and Kelly (1996) and Stutzer (1996); and neural network approaches of Garcia and Gencay (2000) and Gottschling et al. (2000). Jackwerth (1999), Jondeau and Rockinger (2000) and Bliss and Panigirtzoglou (2002) provide excellent surveys of existing methods.

While many papers have estimated and interpreted the option implied RNDs, relatively few have considered the reliability of these methods for estimating implied RNDs. Söderlind and Svensson (1997) and Melick and Thomas (1998) both worked with the parameter variance-covariance matrix. Based on the assumption that the distribution of the estimated parameters are multivariate normal,

the confidence intervals of the estimated RNDs were obtained for examination of stability. While Söderlind and Svensson (1997) derived the confidence intervals analytically using the delta method applied to the parameter estimator, Melick and Thomas (1998) used a Monte Carlo simulation. One disadvantage of such parameter randomization is that it only allows us to examine the stability of the estimated RNDs. It does not enable us to assess how close the estimated RNDs resemble the true RNDs since the true RNDs are unknown. Other authors, including Cooper (1999), Söderlind (2000) and Bondarenko (2003), therefore chose to use the pseudo-prices method. The pseudo-prices method begins with known RNDs which are used to generate fitted prices. These fitted prices are then randomly perturbed to generate pseudo-prices. These pseudo-prices are used to estimate the implied RNDs, and the properties of the estimated RNDs are finally assessed against certain criteria regarding accuracy and stability. Söderlind (2000) and Bondarenko (2003) both estimated implied RNDs from observed option prices and then applied Monte Carlo methods to the fitted option prices based on presumed error distributions. The study of particular relevance to the current work is from Cooper (1999), who generated pseudo-prices from Heston (1993) stochastic volatility model and using Monte Carlo simulations he examined the relative performance of the two most commonly used methods — the double lognormal approximating function method (DLN) and the smoothed implied volatility smile method (SML) — for estimating implied RNDs. By examining the distribution of a set of summary statistics calculated from the estimated RNDs, he concluded that the SML method dominates the DLN method as a technique for estimating option implied RNDs. In addition to Cooper (1999), Bliss and Panigirtzoglou (2002) also compared these two methods. But unlike Cooper (1999) they focused only on the stability of the estimated implied RNDs. They therefore chose to perturb observed option prices. Also by examining a set of summary statistics from estimated RNDs, they reached the same conclusion that the SML method outperforms the DLN method. Many authors have since used the SML method in various studies. Most recent applications are given by Bliss and Panigirtzoglou (2004) and Panigirtzoglou and Skiadopoulos (2004).

Recently, Abadir and Rockinger (2003) proposed an alternative method for estimating option implied RNDs. We call it the density functionals based on confluent hypergeometric functions (DFCH). This method is solidly founded in the theory of statistical density functionals and is particularly appealing for its semi-nonparametric nature. It is more efficient than fully nonparametric estimation but less restricted than fully parametric methods. It encompasses a large class of

traditional densities, such as the normal, gamma, inverse gamma, Weibull, Pareto and their mixtures. Therefore, the possibility of misspecification is expected to be small. They showed that this method performed uniformly well in their two applications. Although the DFCH method appears to be an appealing alternative, surprisingly it did not attract any noticeable follow-up, at least to our best knowledge. The main contribution of this paper is the comparison of the SML method and the DFCH method for estimating option implied RNDs.

To compare the two methods, we conduct two Monte Carlo experiments. Both experiments are based on the pseudo-prices methodology. In the first experiment, we follow Cooper (1999) and use the Heston (1993) stochastic volatility model to generate true RNDs. Different sets of parameters are selected for this model so that our true RNDs incorporate various market conditions¹. In the second experiment, the true RND is specified as a mixture of three lognormals. In order to generate the true RND that is representative to the observed world, we calibrate the model using observed prices of a typical cross section of S&P500 Index options traded at Chicago Board Option Exchange (CBOE). It should be stressed that the two experiments can be regarded as complementary to each other in the sense that when combined they represent a broader setting for making comparison between alternative methods². In both experiments, we examine the ability of the two methods for recovering the true RNDs in the presence of small pricing errors. We evaluate the performance of the two RND estimators by focusing on the criterion of Root Mean Integrated Squared Error (RMISE). Results from both experiments show that the DFCH method outperforms the SML method for the overall quality of the estimated RNDs concerning both the accuracy and the stability defined in this paper.

As an illustration of how the two methods perform in real world, we apply both methods to OTC currency option data. For the SML method, three different objective choices of the smoothing parameter are examined and issues of goodness of fit by the two methods are discussed. We show how distribution summary statistics calculated from the estimated RNDs can be used to analyze the effects

¹The setup of the first experiment is similar to Cooper (1999) only in the choice of the true RNDs. They differ in all remaining aspects, including the number of strikes and strike intervals used in each cross section, the error specification, and most importantly the criteria used for assessing the quality of the estimated RNDs. These improved features of our simulation experiment are expected to enhance the validity of the results from the comparison. Details of these features are discussed in Section 3.1.

²We believe that our study is in this respect more general than any of the previous studies of similar nature.

of major events on the market's expectations of future exchange rate movements.

The remainder of the paper is organized as follows. Section 2 sets out the technical details of the two estimation methods and we improve the SML method by providing an analytic expression for the RND estimator. In Section 3, we present the two Monte Carlo simulation experiments and discuss the results. In section 4, we apply the two methods to a set of OTC currency option data. Concluding remarks are given in Section 5.

2. Methods for Estimating Implied RNDs

2.1. Option Prices and Risk-Neutral Densities

Prices of European call options at time zero on the underlying asset S with expiration at T and strike price K are related to the risk-neutral probability density function (RND), $f(\cdot)$, through the following expression:

$$C(K) = e^{-rT} \int_K^{\infty} (S_T - K) f(S_T) dS_T,$$

where r is the continuously compounded risk-free interest rate. Breeden and Litzenberger (1978) noticed that differentiating the integral with respect to strike price K gives

$$\frac{\partial C(K)}{\partial K} = -e^{-rT} \int_K^{\infty} f(S_T) dS_T = -e^{-rT} [1 - F(K)], \quad (1)$$

where $F(\cdot)$ is the cumulative distribution function (CDF) corresponding to the risk-neutral probability density function (PDF), $f(\cdot)$. The second derivative is given by

$$\left. \frac{\partial^2 C(K)}{\partial K^2} \right|_{K=S_T} = e^{-rT} f(S_T), \quad (2)$$

which reveals the required RND, $f(S_T)$. It follows that the implied RND can be recovered by calculating the compounded second partial derivative of the call pricing function with respect to the strike price. In practice, however, some approximating or smoothing method has to be used to construct such a function due to the limited number of observed call prices.

2.2. Smoothed Implied Volatility Smile

2.2.1. General Procedure

The smoothed volatility smile method was originally developed by Shimko (1993). The method is an approximating function method applied to the implied volatility smile. Option prices are first converted to implied volatilities using Black-Scholes option pricing formula³. A continuous smoothing function is then fitted to the implied volatilities against the strike prices. The reason for smoothing the volatility smile instead of interpolating the call pricing function directly is that it is technically difficult to fit accurately the shape of the latter and small fitted price errors tend to have large effects on the resulting RNDs, particularly in the tails. Shimko (1993) used a quadratic functional form to interpolate across the implied volatilities. The continuum of fitted implied volatilities are then converted back to a continuum of fitted option prices. The implied RNDs can be obtained by applying equation (2). Malz (1997a,b) also used a low-order polynomial as the smoothing function, but fitted the implied volatility against the Black-Scholes option delta ($\delta = \partial C / \partial S$). Campa et al. (1998) introduced the use of a smoothing spline for fitting implied volatility curves. They also applied this to smoothing the implied volatility/strike function.

The smoothed implied volatility smile (SML) method considered in this paper was developed by Bliss and Panigirtzoglou (2002). The method follows Malz (1997a,b) in smoothing in implied volatility/delta space and Campa et al. (1998) in using a natural spline to smooth the function. The natural spline minimizes the following objective function:

$$\min_{\Theta} \left\{ (1 - \lambda) \sum_{i=1}^N w_i \left(IV_i - \widehat{IV}_i(\delta_i, \Theta) \right)^2 + \lambda \int_0^{e^{-r^*T}} g''(x; \Theta)^2 dx \right\}, \quad (3)$$

where Θ is the matrix of polynomial parameters of the cubic spline; $g(\Theta)$ is the cubic spline function; and $\widehat{IV}_i(\delta_i, \Theta)$ is the fitted implied volatility at δ_i given the spline parameters Θ . Relative weights to each observation are determined by the values of w_i .⁴ The smoothness of the spline is controlled by the smoothing

³The use of the Black-Scholes formula is to convert data from one space to another, where smoothing can be done more efficiently. It does not assume that the underlying price process is lognormal.

⁴Bliss and Panigirtzoglou (2002) discussed different types of weighting schemes and how the weighting can account for different sources of pricing errors.

parameter, λ , which multiplies a measure of the degree of curvature in the function — the integral of the squared second derivative of the function over its range. It should be recalled that $0 \leq \delta_i \leq e^{-r^*T}$, where r^* is the dividend rate of the underlying asset.

Meanwhile, the objective function in (3) suggests that the degree of freedom for the estimation is also related to the smoothing parameter. In particular, the maximum degree of freedom is achieved when $\lambda = 1$, which amounts to fitting a straight line to the data; whereas, when $\lambda = 0$, the cubic spline provides an exact fit to the data. Fisher et al. (1995) gives a rigorous definition of the effective number of parameters of the regression.

A natural spline is superior to a low-order polynomial because it allows for more flexibility in the shape of the fitted volatility smile and it also allows the user to control the smoothness of the fitted function. Using the option delta rather than the strike price as the function argument has the advantage that away-from-the-money implied volatilities are grouped more closely together than near-the-money implied volatilities. This permits greater flexibility in the shape of the approximating function near the center of the distribution where data is more reliable, without having to use a variable smoothing parameter⁵. In addition, since possible values in the delta space always range from 0 to e^{-r^*T} , the extrapolation area becomes relatively smaller.

Once the natural spline is constructed, the fitted volatility smile is then converted back to the fitted call pricing function. The implied RNDs are obtained by applying equation (2).

2.2.2. Smile Conversion vs Point Conversion

The construction of the SML method proposed by Bliss and Panigirtzoglou (2002) requires that the implied volatility smile be smoothed in delta space. Two different ways of converting a strike into its delta have been suggested in the literature, differing in their choice of the volatility in the delta function. The original one is to use the implied volatility that corresponds to the strike price (Malz 1997a,b). This is achieved by converting strike prices into deltas using the Black-Scholes

⁵See Waggoner (1997) for more discussions on variable smoothness penalties in spline regression.

delta given by the following equation:

$$\delta_K = e^{-r^*T} \Phi \left(\frac{\ln S_0 - \ln K + \left(r - r^* + \frac{\sigma_K^2}{2} \right) T}{\sigma_K \sqrt{T}} \right).$$

The subscript of σ_K emphasizes that a particular strike price K in a given cross section is converted into δ_K through its corresponding implied volatility σ_K on the volatility smile. We call it “smile conversion”. An alternative is to use a single at-the-money implied volatility to convert all strike prices in a given cross section (Bliss and Panigirtzoglou 2004). We define it as “point conversion”. This was in fact accomplished through the following equation:

$$\delta_K = e^{-r^*T} \Phi \left(\frac{\ln S_0 - \ln K + \left(r - r^* + \frac{\sigma_A^2}{2} \right) T}{\sigma_A \sqrt{T}} \right), \quad (4)$$

where σ_A is the at-the-money volatility. Transforming each strike into a delta using the at-the-money implied volatility has the advantage that the ordering of deltas is always the same as that of the strikes. Panigirtzoglou and Skiadopoulos (2004) pointed out that using the implied volatilities that correspond to each strike could change the ordering in the delta space, in cases where steep volatility skews are observed. This would result in generating volatility smiles with artificially created kinks. As a result, they applied equation (4) to convert strikes in their study. In this paper, we do the same.

2.2.3. Analytic CDF and PDF

In previous studies, once the natural spline function is fitted, a large number of δ -spaced points on the function are computed. These are then converted to K -spaced values in price/strike space. These in turn are used to compute the implied CDF or PDF numerically⁶. In this paper we show that another remarkable advantage of “point conversion” is that the implied CDF and PDF can be evaluated analytically. In this section we improve the computational efficiency of the SML method by providing the analytic expression of the estimated option implied CDF and PDF for “point conversion”.

⁶See Bliss and Panigirtzoglou (2004) and Panigirtzoglou and Skiadopoulos (2004).

Denote the fitted natural spline function in the implied volatility/delta space as $\widehat{g}(\delta, \Theta)$, which is a piecewise function of the form

$$\widehat{g}(\delta, \Theta) = \begin{cases} \widehat{g}_0(\delta) & \text{if } \delta < \delta_1 \\ \widehat{g}_1(\delta) & \text{if } \delta_1 \leq \delta < \delta_2 \\ \vdots & \\ \widehat{g}_{n-1}(\delta) & \text{if } \delta_{n-1} \leq \delta < \delta_n \\ \widehat{g}_n(\delta) & \text{if } \delta \geq \delta_n \end{cases}, \quad (5)$$

where $\widehat{g}_i(\delta)$ is a third degree polynomial defined by

$$\widehat{g}_i(\delta) = a_i(\delta - \delta_i)^3 + b_i(\delta - \delta_i)^2 + c_i(\delta - \delta_i) + d_i,$$

for $i = 0, 1, 2, \dots, n$ where n is the number of strikes in the cross section. The first and second derivatives of these n equations are

$$\widehat{g}'_i(\delta) = 3a_i(\delta - \delta_i)^2 + 2b_i(\delta - \delta_i) + c_i,$$

and

$$\widehat{g}''_i(\delta) = 6a_i(\delta - \delta_i) + 2b_i.$$

For any terminal asset price $S_T = K$, at which the option implied CDF or PDF is to be evaluated, the corresponding delta point, δ_K , in the delta space can be directly calculated through equation (4). The fitted implied volatility σ_K is then evaluated at δ_K by the spline function $\widehat{g}_i(\delta)$ in (5) for $\delta_i \leq \delta_K < \delta_{i+1}$. Finally, the fitted European call option price, $\widehat{C}(K)$, can be calculated by substituting both K and σ_K into the Black-Scholes call option pricing formula. Under this framework, the fitted European call option pricing function implied by ‘‘point conversion’’ can be written as

$$\widehat{C}(K) = e^{-r^*T} S_0 \Phi(d_1) - e^{-rT} K \Phi(d_2), \quad (6)$$

where

$$d_1 = \frac{\ln S_0 - \ln K + \left(r - r^* + \frac{\sigma_K^2}{2}\right) T}{\sigma_K \sqrt{T}}, \quad (7)$$

$$d_2 = \frac{\ln S_0 - \ln K + \left(r - r^* - \frac{\sigma_K^2}{2}\right) T}{\sigma_K \sqrt{T}}, \quad (8)$$

$$\sigma_K = \widehat{g}_i(\delta_K), \quad (9)$$

$$\delta_K = e^{-r^*T} \Phi(d_A), \quad (10)$$

$$d_A = \frac{\ln S_0 - \ln K + \left(r - r^* + \frac{\sigma_A^2}{2}\right) T}{\sigma_A \sqrt{T}}, \quad (11)$$

where $\Phi(\cdot)$ is the standard normal distribution function.

We noticed from (6) to (11) that the fitted call pricing function, $\widehat{C}(K)$, though seemingly complicated, is in fact explicit as a function of K . Thus, the derivatives of $\widehat{C}(K)$ with respect to K can be obtained analytically by implementing the chain rule of differentiation. It can be easily verified that the implied risk-neutral CDF, $F(S_T)$, and PDF, $f(S_T)$, can be written as

$$F(S_T) = S_0 e^{(r-r^*)T} \sqrt{T} ABC - \Phi(d_2) + 1 \Big|_{K=S_T}, \quad (12)$$

and

$$f(S_T) = S_0 e^{(r-r^*)T} \sqrt{T} (A'BC + AB'C + ABC') - \phi(d_2) \frac{\partial d_2}{\partial K} \Big|_{K=S_T}, \quad (13)$$

respectively, where

$$\begin{aligned} A &= \phi(d_1), & A' &= -d_1 \frac{\partial d_1}{\partial K} A, \\ B &= \widehat{g}'_i(\delta_K), & B' &= \widehat{g}''_i(\delta_K) C, \\ C &= -\frac{e^{-r^*T} \phi(d_A)}{K \sigma_A \sqrt{T}}, & C' &= \frac{d_A - \sigma_A \sqrt{T}}{K \sigma_A \sqrt{T}} C, \\ \frac{\partial d_1}{\partial K} &= \left(\sqrt{T} - \frac{d_1}{\sigma_K} \right) BC - \frac{1}{K \sigma_K \sqrt{T}}, & \frac{\partial d_2}{\partial K} &= \frac{\partial d_1}{\partial K} - \sqrt{T} BC, \end{aligned}$$

where $\phi(\cdot)$ is the standard normal probability density function. More details on the derivation of (12) and (13) can be found in Bu (2006). Both $F(S_T)$ and $f(S_T)$ can be directly evaluated at any terminal asset price S_T . We therefore use the analytic formula in (13) in our study.

2.3. Density Functionals Based on Confluent Hypergeometric Functions

Abadir and Rockinger (2003) proposed a semi-nonparametric approach for estimating density related functionals. This method is essentially built on the basis of the confluent hypergeometric function, also known as the Kummer's function. It is defined as

$${}_1F_1 \equiv \sum_{j=0}^{\infty} \frac{(\alpha)_j}{(\beta)_j} \frac{z^j}{j!} \equiv 1 + \frac{\alpha}{\beta} z + \frac{\alpha(\alpha+1)}{\beta(\beta+1)} \frac{z^2}{2} + \dots,$$

where

$$(a)_j \equiv (a)(a+1)\dots(a+j-1) = \frac{\Gamma(a+j)}{\Gamma(a)},$$

in which $\Gamma(\nu)$, for $\nu \in \mathbb{R}$, is the gamma function and $-\beta \notin \mathbb{N} \cup \{0\}$. The ${}_1F_1$ can be used to represent a variety of density-related functions⁷.

The functional in the context of option implied RNDs is based on a couple of confluent hypergeometric functions. We call it the density functionals based on confluent hypergeometric functions (DFCH). In contrast to the SML technique, the DFCH method is an approximating function method applied to the option prices. It specifies the European call pricing function as a mixture of two confluent hypergeometric functions:

$$\begin{aligned} C(K) \equiv & c_1 + c_2 K + 1_{K > m_1} a_1 (K - m_1)^{b_1} {}_1F_1(a_2; a_3; b_2 (K - m_1)^{b_3}) \\ & + (a_4) {}_1F_1(a_5; a_6; b_2 (K - m_2)^2), \end{aligned} \quad (14)$$

where $-a_3, -a_6 \notin \mathbb{N} \cup \{0\}$ and $b_2, b_4 \in \mathbb{R}_-$. The indicator function is required to represent a component of the density with bounded support. It can be shown that the first ${}_1F_1$ function in $C(\cdot)$ covers the double integrals of the gamma and other asymmetric generalizations and the second covers the double integrals of a family of symmetric quadratic exponential densities such as the normal. Thus, the DFCH approach encompasses many known distributions in statistics and their mixtures. Examples of special cases giving integrals of known density functions include Gamma, Inverse Gamma, Weibull, Normal and Pareto, etc. See Abadir and Rockinger (2003) for more details.

⁷The usefulness of the Kummer's function in econometrics and dynamic economics is elaborated in Abadir (1999).

Differentiating (14) twice⁸ with respect to K yields the implied RND given by

$$\begin{aligned}
f(S_T) &\equiv e^{rT} \left. \frac{\partial^2 C(K)}{\partial K^2} \right|_{K=S_T} \\
&= e^{rT} \left\{ 1_{K>m_1} a_1 (K - m_1)^{b_1-2} \left[b_1 (b_1 - 1) {}_1F_1 \left(a_2; a_3; b_2 (K - m_1)^{b_3} \right) \right. \right. \\
&\quad + \frac{a_2}{a_3} b_2 b_3 (2b_1 + b_3 - 1) (K - m_1)^{b_3} \\
&\quad \times {}_1F_1 \left(a_2 + 1; a_3 + 1; b_2 (K - m_1)^{b_3} \right) \\
&\quad \times \frac{a_2 (a_2 + 1)}{a_3 (a_3 + 1)} b_2^2 b_3^2 (K - m_1)^{2b_3} \\
&\quad \left. \left. \times {}_1F_1 \left(a_2 + 2; a_3 + 2; b_2 (K - m_1)^{b_3} \right) \right] \right. \\
&\quad + 2a_4 \frac{a_5}{a_6} b_4 \left[{}_1F_1 \left(a_5 + 1; a_6 + 1; b_4 (K - m_2)^2 \right) \right. \\
&\quad \quad + 2 \frac{a_5 + 1}{a_6 + 1} b (K - m_2)^2 \\
&\quad \quad \left. \left. \times {}_1F_1 \left(a_5 + 2; a_6 + 2; b_4 (K - m_2)^2 \right) \right] \right\} \Big|_{K=S_T}. \quad (15)
\end{aligned}$$

Given observations of call option prices, the parameters of the implied RND can be estimated using several different methods which include maximum likelihood, generalized least squares, generalized methods of moments⁹, and so on. In this paper, we use non-linear optimization methods¹⁰ to minimize the sum of

⁸An important feature of the ${}_1F_1$ function is that iterated integrals or derivatives of ${}_1F_1$ gives mixtures of ${}_1F_1$, which makes it a natural tool to model option prices and, more generally, functionals of densities.

⁹Abadir and Rockinger (2003) showed that subject to functions in (15) being nondegenerate (i.e., the existence condition for the moments), the moments of the implied RND can be explicitly characterized.

¹⁰Abadir and Rockinger (2003) proved that the nonlinear LS estimators are consistent and asymptotically normal for any of the parameters that have a nonzero impact on the function.

squared fitted pricing errors, i.e.

$$\min_{\Theta} \left\{ \sum_{i=1}^N w_i \left[C(K_i) - \widehat{C}(K_i|\Theta) \right]^2 \right\}, \quad (16)$$

where as before w_i represents the relative weights placed on each observation.

It is important to note that not all the parameters in (14) are free to vary unrelatdly. For the function to be the integral of a CDF, at least three restrictions in general and a possible seven in the problem at hand can be imposed on the parameter space¹¹. The restrictions imposed by Abadir and Rockinger (2003) are given by (17)-(21) which include the martingale condition in (21). As a result, the actual number of parameters to be estimated is reduced to seven¹².

$$a_5 = -\frac{1}{2}, a_6 = \frac{1}{2}, b_1 = 1 + a_2 b_3, \quad (17)$$

$$c_1 = -c_2 m_2, \quad (18)$$

$$c_2 = -e^{-rT} + a_4 \sqrt{-b_4 \pi}, \quad (19)$$

$$a_4 = \frac{1}{2\sqrt{-b_4 \pi}} \left[e^{-rT} - a_1 (-b_2)^{-a_2} \frac{\Gamma(a_3)}{\Gamma(a_3 - a_2)} \right], \quad (20)$$

$$E(z) = e^{rT} a_1 \frac{\Gamma(a_3)}{\Gamma(a_3 - a_2)} (-b_2)^{-a_2} (m_1 - m_2) + m_2. \quad (21)$$

3. Monte Carlo Experiments

3.1. Monte Carlo Experiment Based on Heston's Stochastic Volatility Model

In this section, we perform a Monte-Carlo experiment to compare the ability of the two techniques to recover option implied RNDs. As pointed out previously, a good RND estimation technique should be able to recover the true RNDs whatever the complexity of their shapes. Therefore, for comparison purposes the choice of the true RNDs in a simulation should itself be able to take on a wide range of

¹¹See Abadir and Rockinger (2003) for derivations.

¹²Further restrictions may be imposed on the parameters to get more restrictive distributions in cases where insufficient number of observations are available.

different shapes reflecting various empirical features of asset distributions¹³: high or low volatility, positive or negative skewness, excess kurtosis, and cater for the full range of maturities that are encountered in practice. To generate risk-neutral densities that incorporate these features, we follow Cooper (1999) and use Heston's (1993) stochastic volatility model to generate true RNDs and fitted option prices. Under Heston's model, the underlying asset price dynamics are described by the following stochastic differential equations:

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{v_t} S_t dz_1, \\ dv_t &= \kappa (\theta - v_t) dt + \sigma_v \sqrt{v_t} dz_2. \end{aligned} \quad (22)$$

Here the volatility of the underlying asset $\sqrt{v_t}$ is also stochastic. The conditional variance v_t follows a mean reverting process such that the volatility mean-reverts to a long run of $\sqrt{\theta}$ at a rate dictated by κ . The term σ_v sets the volatility of the volatility. The two Wiener process dz_1 and dz_2 have a correlation given by ρ . By changing the correlation parameter we can generate skewness in asset returns¹⁴. Heston shows that the European call option price on an asset that behaves according to (22) has a closed form solution, which is given by

$$C(S_0, v_0, K, T) = S_0 e^{-r^* T} P_1 - K e^{-r T} P_2, \quad (23)$$

where for $j = 1, 2$

$$P_j(\ln(S_0), v_0, T; \ln(K)) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{e^{-i\phi \ln(K)} f_j(\ln(S_0), v_0, T; \phi)}{i\phi} \right] d\phi,$$

$$f_i(\ln(S_0), v_0, T; \phi) = e^{C(T; \phi) + D(T; \phi) v_0 + i\phi \ln(S_0)},$$

$$C(T; \phi) = (r - r^*) \phi i T + \frac{a}{\sigma_v^2} \left\{ (b_j - \rho \sigma_v \phi i + d) T - 2 \ln \left[\frac{1 - g e^{dT}}{1 - g} \right] \right\},$$

$$D(T; \phi) = \frac{b_j - \rho \sigma_v \phi i + d}{\sigma_v^2} \left[\frac{1 - e^{dT}}{1 - g e^{dT}} \right],$$

$$g = \frac{b_j - \rho \sigma_v \phi i + d}{b_j - \rho \sigma_v \phi i - d},$$

¹³Note that option implied RNDs are risk-neutral and thus different from the empirical asset distributions. But one can justifiably suppose a rough similarity between the risk-neutral and the objective distributions. See Rubinstein (1994).

¹⁴See Cooper (1999) for an illustration of the effect of ρ on the implied RNDs.

$$d = \sqrt{(\rho\sigma_v\phi i - b_j)^2 - \sigma_v^2 (2u_j\phi i - \phi^2)},$$

$$u_1 = \frac{1}{2}, u_2 = -\frac{1}{2}, a = \kappa\theta, b_1 = \kappa + \lambda - \rho\sigma_v, b_2 = \kappa + \lambda.$$

As in Cooper (1999), we test performance across a range of six scenarios, which correspond to combinations of low and high volatility and three levels of skewness. We then generate European-style call and put option prices with 4 different contract maturities from 2 weeks up to 6 months. Thus, a total of 24 different pairs of scenario and maturity are generated. The Heston's model parameters used for each scenario are set out in Table 1. These are chosen to generate true RNDs that correspond to situations of negative skewness, and weak and strong positive skewness in the terminal asset price and also conditions of low and high volatility. To generate these levels of skewness in the terminal asset price distributions, three different values for the correlation parameter -0.9, 0 and 0.9 are used. The long run volatilities for the high volatility scenarios are chosen on the basis of the levels of implied volatility typically observed within equity markets. The low volatility scenarios are used to mimic data from Stock Index, FX and interest rate markets. It is important to note that the Heston model is used here simply as a convenient tool to generate underlying RNDs that incorporate the empirical features discussed above, namely, different levels of spread, skewness and excess kurtosis of the implied RNDs. Doing so does not presume that equation (22) correctly describes the asset price dynamics in the real world. For this reason, it is innocuous to assume for simplicity that the market price of volatility risk is zero and that the conditional volatility at time zero is equal to the long run volatility, as long as the resulting RNDs serve our purpose.

[Table 1 here]

To obtain the true RND for each scenario and maturity pair, we generate a large number of theoretical prices of Heston call options using the pricing formula in equation (23) and calculate the RND numerically by applying equation (2). Table 2 presents the following descriptive statistics for the true RNDs used in this experiment. They are: (a) the mean, $\hat{\mu}$; (b) the standard deviation, $\hat{\sigma}$; (c) the skewness (the third central moment), *Skew*; as well as (d) the Kurtosis (the fourth central moment), *Kurt*.

For simplicity, the mean of the true RNDs are set equal to 100. It can be seen that our true RNDs take on a wide range of different shapes, with the standard deviations ranging from 2.038 (small spread) to 23.060 (large spread), skewness

from -0.474 (large negative skew) to 1.964 (large positive skew), and kurtosis from 2.770 (thin tails) to 10.847 (fat tails). Specifically, scenario 1, 2, and 3 are low volatility cases and scenario 4, 5, and 6 are high volatility cases; Scenario 1 and 4 represent strong negative skewness, while others are positive cases. As we would expect, all three measures including the kurtosis increase with time-to-expiry.

[Table 2 here]

For each of the 24 cases generated above, we compute theoretical option prices at a number of different strikes. We assume that the strike range is just wide enough to cover the area between the 1st and the 99th percentiles of the true RND, while at the same time the number of and the interval between strike prices are monitored so that they reflect the typical features of the markets. Specifically, we construct strike interval equal to 1 for low volatility scenarios and 5 for high volatility scenarios, respectively. The final number of strikes for each cross section is set out in Table 3. It can be seen that the numbers of strikes used in this study reflect the real world situation in the following two senses. Firstly, they are increasing in time-to-expiry. Secondly, they are close to the actual numbers of strikes one may observe in reality for corresponding maturities¹⁵.

[Table 3 here]

As discussed earlier, accuracy and stability are both desirable properties of a good RND estimator. To test the robustness of alternative methods to small errors embedded in option prices, we add noise ε_i to the theoretical prices computed above. Noise ε_i is introduced to model observational errors that arise from market imperfections such as nonsynchronicity, bid-ask spread, and discreteness, etc. In this paper, we adopt the same error specification as in Bondarenko (2003). Under such specification, ε_i is independently and uniformly distributed on $[-0.5s_i, 0.5s_i]$, where s_i is the spread. The value of the spread s_i depends on the strike K_i . The spread is larger for in-the-money options and smaller for out-of-the-money options. Specifically, the value of spread is defined as

$$s_i = c \min (M (C_i), M(P_i)),$$

¹⁵We believe that this is an improvement over Cooper's experiment in which 71 strikes with interval equal to 1 are used universally across all scenarios and maturities.

where C_i and P_i are corresponding call and put prices; $M(\cdot)$ is the maximum bid-ask spread permitted by the exchange, which is a function of the option prices¹⁶; c is a scale parameter. The advantages of such specification for ε_i are that noise is smaller in the absolute terms but larger in the relative terms for far-from-the-money strikes, which reflects the empirical feature of heteroskedastic pricing errors¹⁷. The presence of the scale constant c allows us to proportionally increase or decrease the level of noise across all strikes. However, such error specification does not guarantee nonnegative option prices after perturbation. Recall that the SML method requires option prices to be first converted to the implied volatilities which are not defined for negative option prices. Failure to monitor the nonnegativity of the option prices will lead to failure of the SML method, although this is not a problem for the DFCH method. An obvious solution to this problem is to choose an arbitrarily small value of c . Nevertheless, too small noises may invalidate the test for the robustness. In order to reconcile these two, we select the maximum possible value of c that still guarantees nonnegativity of the option prices after perturbation.

Comparisons of RNDs based on visual inspections are usually not reliable. Besides, it is not easy to present and compare more than a few RNDs in the same graph. For this reason, we analyze the perturbed-price RNDs by examining the distribution of certain summary statistics. Both Cooper (1999) and Bliss and Panigirtzoglou (2002) studied the mean, standard deviation, skewness and kurtosis of their estimated RNDs, but Bliss and Panigirtzoglou (2002) examined a larger set of summary statistics which also include median-based skewness, mode-based skewness, quartile-based skewness as well as tail percentiles. It is believed that the quality of these summary statistics derived from the estimated RNDs, when taken together, reflects the quality of the entire distribution. However, it is highly debatable to use higher moments and tail quantiles in examining option implied RNDs. The reason is that these quantities are fairly sensitive to the tails of the distribution where option data provide little information outside the range of the available strike prices. As discussed in Melick and Thomas (1998), there is an infinite variety of probability masses outside the strike range that can be consistent

¹⁶The maximum bid-ask spread permitted by the exchange is linked to the option quotes. For instance, the CBOE rules state that the maximum bid-ask spread is $\frac{1}{4}$ for options with bid quote below \$2, $\frac{3}{8}$ for bid quotes between \$2 and \$5, $\frac{1}{2}$ for bid quotes between \$5 and \$10, . . . , and so on. The function $M(\cdot)$ is constructed to represent such rules. See Bondarenko (2003) for details on the construction of $M(\cdot)$.

¹⁷This is in contrast to Cooper (1999) where homoskedastic pricing errors were used.

with the observed option prices. The allocation of tail probability mass is specific to each estimating technique. For this reason, results from previous stability tests based on quantities derived from tail distributions should be interpreted with caution.

In this study, however, we focus our attention on the criterion of root mean integrated squared error (RMISE). If $\hat{f}(S_T)$ is the RND estimator of the true RND $f(S_T)$, then the RMISE is defined as

$$\mu(\hat{f}) := \text{RMISE}(\hat{f}) = \sqrt{E \left[\int_{-\infty}^{\infty} (\hat{f}(S_T) - f(S_T))^2 dS_T \right]}. \quad (24)$$

It can be seen that since the RMISE is a measure of the average of the integral of the squared deviation over the whole support of the distribution, it is not particularly sensitive to the tails of the distribution. Another advantage of RMISE is that it can be conveniently represented as

$$\mu^2(\hat{f}) = \mu_1^2(\hat{f}) + \mu_2^2(\hat{f}),$$

$$\mu_1(\hat{f}) := \text{RISB}(\hat{f}) = \sqrt{\int_{-\infty}^{\infty} (E[\hat{f}(S_T)] - f(S_T))^2 dS_T}, \quad (25)$$

$$\mu_2(\hat{f}) := \text{RIV}(\hat{f}) = \sqrt{\int_{-\infty}^{\infty} E \left[(\hat{f}(S_T) - E[\hat{f}(S_T)])^2 \right] dS_T}, \quad (26)$$

where RISB is the root integrated squared bias and RIV is the root integrated variance. We define the RMISE as our measure of the overall quality of the estimator, RISB as our measure of the accuracy, and RIV as our measure of the stability. Since both accuracy and stability are important properties of a good estimator, the decomposed representation allows us to easily study the relative contributions of the bias μ_1 and the variability μ_2 to the RMISE of different methods. For each cross section, we repeat the procedure of shocking the prices and then fitting the RND for 500 times¹⁸. The RMISE, RISB and RIV are then obtained by applying equation (24) to (26).

In the estimation of the RNDs, each observation is weighted by the inverse of the variance of ε_i to account for the heteroskedasticity introduced by the error

¹⁸We found that by using more than 500 replications the changes in RMISE, RISB and RIV become negligible.

specification. Meanwhile, in order to examine the impact of the choice of weighting schemes on the relative performance of the two methods, we also redo the experiment by using option vega weighting and equal weighting in the estimation.

The results of the simulation for the two methods based on the inverse variance weighting are presented in Table 4¹⁹. The results for the SML method are displayed on the left panel. Recall that for the SML method the smoothing parameter λ is a free parameter which allows the user to control the trade-off between the smoothness and the goodness of fit. In this study, we search for the optimal parameter λ that minimizes the RMISE. It is important to note that this is only possible in simulation studies where the true RND is known. In real world where the true RND is unknown the smoothing parameter has to be selected by the user. Different values of λ will result in different RNDs. In the following comparison with the DFCH, we use the minimum RMISE corresponding to such optimal λ for the SML²⁰. For the DFCH method, however, the RMISE result is unique for each cross section.

[Table 4 here]

Examining the RMISE values from the two methods, we find that in more than two thirds of the 24 cases the DFCH provides lower RMISE than the SML does, indicating better overall quality of the DFCH as an RND estimator. Specifically, for scenario 1 and 4, which represent negative skewness cases, the DFCH dominates the SML method across all maturities by a substantial margin.

For those cases where the DFCH underperforms the SML, the differences in RMISE are relatively small, with the largest being the scenario 6 - 6 month maturity case. Note that as far as the shape of the true RND is concerned, scenario 2 and 5 represent very weak positive skewness, especially for short maturities. The two long maturity cases in scenario 6 represent very strong positive skewness. The DFCH method appears to fair better in all remaining cases.

Investigation of the RISB reveals that the DFCH is often less biased than the SML, suggesting the flexibility of the confluent hypergeometric functions. The RIV show that the SML is relatively stable for cases where the true RND is not strongly skewed (close to being symmetric).

¹⁹As expected, because the fitted price errors are generally small, the weights used to multiply them have little impact on the estimation. To save time and space, the results from the other two weighting schemes are not reported.

²⁰We believe this is biased in favour of the SML method, because the RMISE would have been larger if we chose the λ as if we did not know the true RND.

Figure 1 gives an example of the differences in the estimated RNDs from the two methods. For each method, the 500 RND estimates are plotted against the true RND. It can be seen that the SML is significantly biased, particularly on the left tail of the distribution. It also shows relatively larger variations in the center of the distribution. In contrast, the DFCH fits the true RND fairly well and exhibits relatively smaller variations. It should be noted that because of the highly nonlinear functions and the larger number of parameters, the DFCH method need higher computational time than the SML method.

[Figure 1 here]

3.2. Monte Carlo Experiment Based on Mixture of Lognormals

A criticism over such simulation approaches as above is that the performance of a particular RND estimating method may be related to the choice of the true RNDs. Results obtained from estimating RNDs of some particular functional form may not be generalized to RNDs outside the set examined. Such concern would be eased to a large extent if one method could perform consistently well in different settings. As a cautionary step, we conduct a further Monte-Carlo experiment in which the true RND is specified as a mixture of three lognormals. Compared to the RNDs implied by the Heston's model, the three-lognormal specification has the advantage that the RND has an explicit analytic form and it also allows for multiple modes in the density should the user or the option data used for estimating the RND demands it.

To reflect real world conditions as closely as possible, the parameters of this specification are chosen to describe a typical cross section of the S&P 500 Index options traded at the Chicago Board Options Exchange (CBOE). They are calibrated from the closing prices on March 21, 1995 of the S&P 500 options with the maturity date on April 21, 1995²¹. The resulting RND as modeled by the three-lognormals are depicted by the dashed lines in Figure 2, which exhibits negative skewness and bimodality. This RND is then used to generate pseudo option prices which are to be estimated by the two methods for comparisons.

In this experiment, we apply the same error specification as in the first experiment. The RMISE, RISB, and RIV values (obtained based on 500 replications) for the two methods under three different weighting schemes are presented in Table

²¹The same dataset and density specification were used by Bondarenko (2003).

5. As before, for the SML method the RMISE reported here are the minimum obtained by searching for the optimal parameter λ^{22} . Our results once again suggest remarkable dominance of the DFCH over the SML. As shown in this table, the RMISE values provided by the SML are substantially larger than those given by the DFCH across all three weighting schemes. Examining the RISB and the RIV reveals that the large bias from the true RND is the main cause of the relatively poor performance of the SML method, not the variance as the RIV values are fairly close. This is evidence that results from pure stability test without concern for accuracy could be very misleading. These findings are invariant across all weighing schemes. To visualize such differences we plot the estimated RNDs from both methods against the true RND in Figure 2. Both methods are able to recover negatively skewed RNDs and produce the second mode implied by the data, and as suggested by the RIV values both methods show similar level of variability. But whereas the SML is apparently biased in many parts of the density, the DFCH recovers the true RND with superior precision, further evidence of the flexibility of the confluent hypergeometric functions.

[Table 5]
[Figure 2 here]

4. Application

4.1. The data

As an illustration of the two RND estimating methods in real settings, we apply the two methods to OTC data²³ of European French franc/Deutsche mark (FF/DM) rate options of the two dates: 17 May 1996, a day when the exchange rate markets are considered to be calm, and on 25 April 1997, a few days after the French President Chirac announced dissolution of the National Assembly, which implied nation-wide elections. This type of options are quoted in terms of delta. For the first date, we have 13 strikes corresponding to delta values 5, 10, 15, 20, 30, 40, 50 (at the money), 60, 70, 80, 85, 90, 95. For the second date, we have strikes corresponding to deltas from 10 to 90 (11 strikes). All possible deltas are used in this study. For each option of a given maturity, the actual strike prices

²²These optimal values are 3.3727×10^{-2} , 2.3791×10^{-2} , and 1.3816×10^{-2} for inverse variance weighting, equal weighting and vega weighting, respectively.

²³We thank Professor Karim Abadir and Professor Michael Rockinger for this data.

are extracted by using a numerical procedure. Following the literature, the average between the bid and ask prices are used. Although we have options with six different time-to-maturities (1, 2, 3, 6, 9, and 12 month), we select to report the results for 1, 3 and 12 month maturity, which are typical, representing short, medium and long horizons, respectively. The spot exchange rate as well as the domestic (French) and foreign (German) Eurocurrency interest rates are all readily available.

4.2. Summary of the Results

As discussed earlier, for the SML method the presence of the smoothing parameter λ allows the user to control the trade-off between the goodness of fit and smoothness of the estimated RNDs. In the two Monte Carlo experiments conducted in the previous section, because the true RND is known we were able to search for the optimal λ that minimizes the RMISE in each case. In the face of observed option data, the choice of λ has to be decided in some ad-hoc ways. A suitable smoothing parameter can be obtained by simply plotting the distribution for different smoothing parameters and choose the one which yields the “best” result. The main disadvantage of this method is that the shape of the estimated RND relies on subjective judgement. Therefore, two researchers may come up with different RNDs for the same data. Besides, this is a cumbersome method for studies where a large number of distributions need to be estimated.

Several procedures for automatically choosing an optimal smoothing parameter have been proposed in the spline regression literature. The most popular class of these methods is based on cross validation (CV) proposed by Craven and Wahba (1979). The basic principle of cross validation is to leave out the data points one at a time and to choose the smoothing parameter for which the missing data points are best predicted by the remainder of the data. One of the criteria used in this study is the generalized cross validation (GCV), which reweights the contribution of deletion residuals to the total score. The smoothing parameters chosen by CV and GCV are in general similar in cases where the size of the pricing errors is small.

Another criterion of interest was suggested by Bliss and Panigirtzoglou (2002). It is to select the λ such that the maximum fitted price error is approximately equal to one half of the tick size by which the options are quoted. They believe that by doing so one can effectively fit the data within the precision of option price measurement.

In addition, since the SML method is always capable of providing an exact fit to the data, for comparison purposes, it is useful to see whether or not the SML method can lead to reasonable RNDs while at the same time gives as good a fit as the DFCH method does. Because the sum of the squared residuals (SSR) is a monotonically decreasing function of the smoothing parameter λ , we are able to find the value of λ such that the SSR from the two methods are exactly equal. We apply both RND estimating methods to the OTC currency data discussed above. The performance of the SML method is examined under all three smoothing parameter selection criteria.

Our results show that all three criteria provide too loose smoothing parameters and the resulting RNDs exhibit unreasonably large fluctuations across all maturities for both dates. The inconsistency of the GCV criterion indicates that even if the spline function is optimal according to the GCV procedure in the implied volatility/delta space, it is not necessarily "optimal" after the transformation required to obtain the RND. The fact that the second criterion failed to provide reasonable result suggests that the real size of the pricing errors, at least in this market, is far greater than that of those imposed by the discreteness of option quotes. In fact, implementing such criterion relies on the assumption on the size of the pricing errors. As the real size of pricing errors is unknown, this criterion is arbitrary. Moreover, the smoothing parameter that provides the same SSR is also too loose to generate reasonable RNDs. A much tighter λ is required to generate a plausible RND. This indicates that at least in the set examined the SML method could not provide as good fit as the DFCH while at the same time maintaining proper shape of the RNDs. In contrast, the nonparametric nature of the DFCH method enables it to give a high goodness of fit and at the same time the parametric property ensures that the estimated RNDs are proper density functions.

As an example, Figure 3 shows the RNDs estimated from the DFCH method and the SML method with the three different choices of λ for the 1 month maturity options on date 17 May 1996. The values of the λ selected by the three criteria are 1.0696×10^{-4} , 2.2451×10^{-5} , and 3.3606×10^{-4} , respectively.

[Figure 3 here]

The above results show that at least within the samples examined, we have found no objective ways for the choice of the smoothing parameter. In the end, by relying on visual inspection we are able to select a value of λ so that the RNDs

estimated by the SML method are smooth enough²⁴ but with the best possible fit (smallest possible SSR). For the DFCH method, however, the estimation is unique. We report the parameter estimates as well as the bootstrapping standard errors²⁵ in Table 6. The estimated RNDs for the two dates are plotted in Figure 4 and Figure 5, respectively.

[Table 6 here]
[Figure 4 here]
[Figure 5 here]

Distribution summary statistics of the implied RNDs such as mean, standard deviation, skewness and kurtosis are useful (to a certain degree) when analyzing changes in the shape of the implied RNDs. To study the implied RNDs estimated by the two methods, we calculate these statistics and report them, together with the (degree of freedom adjusted) goodness of fit measure, the adjusted R^2 (or \bar{R}^2), and the values of λ chosen in the SML estimation, in Table 7. In this study, the \bar{R}^2 is defined as

$$\bar{R}^2 = 1 - \frac{(N - k)^{-1} \sum_{i=1}^N (C_i - \hat{C}_i)^2}{(N - 1)^{-1} \sum_{i=1}^N (C_i - \bar{C})^2},$$

where $\bar{C} = 1/N \sum_{i=1}^N C_i$. For the DFCH method, k is the number of actual parameters, whereas for the SML method k is replaced by the equivalent number of parameters of the spline regression defined by Fisher et al. (1995). It can be seen that the DFCH method provides very high values of \bar{R}^2 across all estimations, indicating great flexibility of the confluent hypergeometric functions. As shown earlier, a λ that provides the same goodness of fit is too loose to generate non-oscillating RNDs. Thus, the value of λ is selected by visual inspection to provide the best possible fit while ensuring minimum acceptable smoothness in the resulting RND. It turned out that under such required λ the \bar{R}^2 for the SML method are lower by a sizable magnitude than that of the DFCH. However, it should be noted that the \bar{R}^2 reported here corresponds to a non-linear relation and therefore gives only a broad indication of goodness of fit. Figure 6 displays the plot of the original quotes and the fitted volatility function for the two methods,

²⁴In the sense of being nonnegative everywhere and free of peculiar zigzags.

²⁵The asymptotic distribution of the parameter estimates provided by Abadir and Rockinger (2003) has limited relevance here due to the generally small number of observations in the data. These standard errors were thus obtained by bootstrapping the fitted pricing errors.

which is representative of the fit across all estimations we attempted within the dataset. We notice an excellent fit from the DFCH method as compared to the SML method, particularly for away-from-the-money options²⁶.

[Table 7 here]
[Figure 6 here]

By construction, the SML method will always fit the mean of the implied RND to the forward price. Thus, calculating the mean of the implied RND does not provide additional information. Since we also imposed the mean-forward equality as a constraint in the DFCH procedure, for any particular date and maturity we obtain the same values of the mean from both methods. But as expected all remaining statistics are different. We notice that the differences in skewness and kurtosis are much more remarkable than the difference between the two standard deviations. The reason is, as discussed earlier, that higher moments are sensitive to the tails of the distribution where observed option prices do not provide sufficient information. Nevertheless, we may consider giving more confidence to the DFCH RNDs by recalling that the DFCH apparently provides better fit to away-from-the-money options than the SML does, as Figure 6 vindicates.

Interesting conclusion can be drawn from these estimated summary statistics. We find these statistics reflect the influence of major events upon this market. Comparing the mean for the two dates reveals an overall shift to the left for all maturities. This has come from the fact that the FF had appreciated against the DM. When we compare the standard deviations for the two dates, there is a larger spread for the implied RNDs at all maturities for the second date. It is an indication that for the second date there is a greater uncertainty among the markets participants about how the exchange rate will evolve towards maturity, following Present Chirac's announcement of the dissolution of the National Assembly. Specifically, across all maturities the right tail of the RND decays more slowly for the second date, suggesting that the market is contemplating a non-negligible probability of subsequent depreciation of the FF. Market participants' uncertainty about the exchange rate movement and fear of large price changes are also reflected by the large values of Kurtosis from both methods on both dates, as they are willing to pay a high premium for protection against such large price changes.

²⁶The SML cannot provide better fit than the one depicted in the figure without incurring excessive curvature in the middle of the fitted volatility smile, which leads to oscillating RNDs.

5. Conclusion

In this paper, we compared the widely known SML method with a new semi-nonparametric DFCH method for estimating option implied RNDs. We conducted two Monte Carlo experiments based on the pseudo-prices methodology. This methodology consists of re-estimating implied RNDs from randomly perturbed cross sections of fitted theoretical option data based on presumed true RNDs. In the first experiment, the true RNDs were generated by Heston's stochastic volatility model. These RNDs were selected to represent various empirical features of asset distributions. In the second experiment, an alternative specification of the true RND was considered. It was based on a mixture of three lognormals, and the parameters of this specification were calibrated from a typical cross section of S&P500 Index option data.

To compare the two RND estimating methods, we focused on the RMISE criterion, which is a measure of average distance between the true RND and the estimated ones. Results from both experiments have provided strong evidence of the superiority of the DFCH method over the SML method under both accuracy and stability considerations. In particular, for a large majority of the cases the DFCH has more closely recovered the implied RNDs than the SML. We also found that our results are insensitive to the weighting schemes applied in the estimation.

We also applied the two methods to OTC currency option data. The statistical analysis conducted in our Monte Carlo experiments ignored the problem related to the choice of the smoothing parameter in the SML method and assumed that the theoretically optimal λ is applied. In this empirical study, however, as the true RND is unknown we attempted three different objective choices of λ . We found that all three criteria failed to generate RNDs with reasonable shapes. In particular, we found that in these examples the SML could not provide as good fit to the data as the DFCH while still generating non-oscillating RNDs. In the end, arbitrary choices of λ were used. The main disadvantage of this is that two researchers may come up with different RNDs for the same data. A set of distribution summary statistics of the estimated RNDs were then examined. We showed how these RNDs summary statistics can be used to analyze the effects of major events on the market's expectations of future exchange rate movements.

In summary, we found that the SML method is attractive as it is a practically very efficient way of extracting implied RNDs from option data, particularly with its analytic form of the implied risk-neutral CDF and PDF provided in this paper. Previous studies also suggested that it is both absolutely and relatively more ro-

bust to errors embedded in daily option prices than the most widely used mixture of lognormals technique. The DFCH is a theoretically well-founded statistical density functional model. Because of its semi-nonparametric nature, the DFCH technique is less data-intensive than those fully nonparametric methods, and more flexible than purely parametric methods in generating abundant potential probability density shapes. In this paper, we have shown that the DFCH method outperforms the widely used SML method for estimating option implied RNDs, which makes it an attractive alternative to many of the existing RND estimating methods, despite its requirement for comparatively higher computational time.

References

- Abadir, K.M., 1999. An introduction to hypergeometric functions for economists. *Econometric Reviews* 18 (3), 287-330.
- Abadir, K.M., Rockinger M., 2003. Density functionals, with an option-pricing application. *Econometric Theory* 19, 778-811.
- Ait-Sahalia, Y., Lo, A.W., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53, 499-547.
- Ait-Sahalia, Y., Lo, A.W., 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9-51.
- Ait-Sahalia, Y., Wang, Y., Yared, F., 2001. Do option markets correctly price the probabilities of movement of the underlying asset? *Journal of Econometrics* 102, 67-110.
- Aparicio, S., Hodges, S., 1998. Implied risk-neutral distribution: A comparison of estimation methods. Working Paper, Warwick University.
- Bahra, B., 1997. Implied risk-neutral probability density functions from option prices: Theory and application. Working Paper, Bank of England.
- Bates, D., 1996. Jumps and stochastic volatility: Exchange rate process implicit in Deutsche mark options. *Review of Financial Studies* 9, 69-107.
- Bliss, R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. *Journal of Banking and Finance* 26, 381-422.
- Bliss, R., Panigirtzoglou, N., 2004. Option implied risk aversion estimates. *Journal of Finance* 59 (1), 407-446.

- Breeden, D.T., Litzenberger, R.H., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621-651.
- Bondarenko, O., 1997. Testing rationality of financial markets: An application to S&P500 index options. Working Paper, Caltech.
- Bondarenko, O., 2003. Estimation of risk-neutral densities using positive convolution approximation. *Journal of Econometrics* 116, 85-112.
- Bu, R., 2006. Essays in financial econometrics and time series analysis. Ph.D Thesis, University of Liverpool.
- Buchen, P.W., Kelly, M., 1996. The maximum entropy distribution of an asset inferred from option prices. *Journal of Finance and Quantitative Analysis* 31, 143-159.
- Campa, J.M., Chang, P.H.K., Reider, R.L., 1998. Implied exchange rate distributions: Evidence from OTC option markets. *Journal of International Money and Finance* 17, 117-160.
- Cooper, N., 1999. Testing techniques for estimating implied RNDs from the prices of European and American options. Working Paper, Bank of England.
- Craven, P., Wahba, G., 1979. Smoothing noisy data with spline functions. *Numerische Mathematik* 31, 377-403.
- Fisher, M., Nychka, D., Zervos, D., 1995. Fitting the term structure of interest rates with smoothing splines. Working Paper, Finance and Economics Discussion Series, Federal Reserve Board.
- Garcia, R., Gencay, R., 2000. Pricing and hedging derivative securities with neural networks and a homogeneity hint. *Journal of Econometrics* 94, 93-115.
- Gottschling, A., Haefke, C., White, H., 2000. Closed form integration of artificial neural networks with some applications to finance. Working Paper, UC San Diego.
- Heston, S., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options, *The Review of Financial Studies* 6, 327-343.
- Jackwerth, J.C., 1999. Implied binomial trees: A literature review. *Journal of Derivatives* 7, 66-82.
- Jackwerth, J.C., 2000. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13, 433-451.
- Jarrow, R., Rudd, A., 1982. Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics* 10, 347-369.

- Jondeau, E., Rockinger, M., 2000. Reading the smile: The message conveyed by methods which infer risk neutral densities. *Journal of International Money and Finance* 19, 885-915.
- Lim, G.C., Martin, G.M., Martin, V.L., 2005. Parametric pricing of higher order moments in S&P500 options. *Journal of Applied Econometrics* 20, 377-404.
- Malz, A.M., 1997a. Option-implied probability distributions and currency excess returns. Staff Reports, Federal Reserve Bank of New York.
- Malz, A.M., 1997b. Estimating the probability distribution of the future exchange rate from options prices. *Journal of Derivatives* 5 (2) 18-36.
- McManus, D.J., 1999. The information content of interest rate futures options. Working Paper, Bank of Canada.
- Melick, W.R., Thomas, C.P., 1997. Recovering an asset's implied PDF from option prices: An application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91-115.
- Melick, W.R., Thomas, C.P., 1998. Confidence intervals and constant maturity series for probability measures extracted from option prices. Paper Presented at the Bank of Canada Conference on Information Contained in Prices of Financial Assets.
- Panigirtzoglou, N., Skiadopoulos, G., 2004. A new approach to modeling the dynamics of implied distribution: Theory and evidence from the S&P500 options. *Journal of Banking and Finance* 28, 1499-1520.
- Rosenberg, J., 1998. Pricing multivariate contingent claims using estimated risk-neutral density functions. *Journal of International Money and Finance* 17, 229-247.
- Rosenberg, J., Engle, R., 2002. Empirical pricing kernels. *Journal of Financial Economics* 64, 341-372.
- Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49, 771-818.
- Rubinstein, M., 1998. Edgeworth binomial trees. *Journal of Derivatives* 5 (3), 20-27.
- Shimko, D.C., 1993. Bounds of probability. *Risk* 6, 33-37.
- Söderlind, P., 2000. Market expectations in the UK before and after the ERM crisis. *Economica* 67,1-18.
- Söderlind, P., Svensson, L.E.O., 1997. New techniques to extract market expectations from financial instruments. *Journal of Monetary Economics* 40, 383-429.
- Stutzer, M., 1996. A simple nonparametric approach to derivative security valuation. *Journal of Finance* 51, 1633-1652.

- Waggoner, D.F., 1997. Spline methods for extracting interest rate curves from coupon bond prices. Working Paper, Federal Reserve Bank of Atlanta.
- Wu, L., Huang, J., 2004. Specification analysis of option pricing models based on time-changed Levy processes. *Journal of Finance* 59, 1405-1439.

TABLE 1

Model Parameters Used under Each Scenario

	Strong Negative Skew	Weak Positive Skew	Strong Positive Skew
	Scenario 1	Scenario 2	Scenario 3
Low Volatility	$\kappa = 2, \sqrt{\theta} = 0.1$	$\kappa = 2, \sqrt{\theta} = 0.1$	$\kappa = 2, \sqrt{\theta} = 0.1$
	$\sigma_v = 0.1, \rho = -0.9$	$\sigma_v = 0.1, \rho = 0$	$\sigma_v = 0.1, \rho = 0.9$
	Scenario 4	Scenario 5	Scenario 6
High Volatility	$\kappa = 2, \sqrt{\theta} = 0.3$	$\kappa = 2, \sqrt{\theta} = 0.3$	$\kappa = 2, \sqrt{\theta} = 0.3$
	$\sigma_v = 0.4, \rho = -0.9$	$\sigma_v = 0.4, \rho = 0$	$\sigma_v = 0.4, \rho = 0.9$

TABLE 2
Descriptive Statistics of the True RNDs

	Scenario	2 weeks	1 month	3 months	6 months
$\hat{\mu}$	1	100.000	100.000	100.000	100.000
	2	100.000	100.000	100.000	100.000
	3	100.000	100.000	100.000	100.000
	4	100.000	100.000	100.000	100.000
	5	100.000	100.000	100.000	100.000
	6	100.000	100.000	100.000	100.000
	$\hat{\sigma}$	1	2.038	2.877	4.956
2		2.041	2.887	5.003	7.081
3		2.045	2.898	5.052	7.200
4		6.085	8.555	14.529	20.127
5		6.130	8.677	15.094	21.491
6		6.175	8.802	15.702	23.060
<i>Skew</i>		1	-0.206	-0.281	-0.418
	2	0.062	0.089	0.159	0.231
	3	0.331	0.459	0.743	0.956
	4	-0.172	-0.229	-0.304	-0.275
	5	0.188	0.273	0.505	0.762
	6	0.551	0.781	1.362	1.964
	<i>Kurt</i>	1	3.045	3.082	3.180
2		3.046	3.088	3.223	3.356
3		3.178	3.346	3.931	4.602
4		2.983	2.966	2.888	2.770
5		3.135	3.270	3.821	4.678
6		3.532	4.081	6.487	10.847

TABLE 3
Number of Strikes Used in the First Experiment

Scenario	2 Weeks	1 Month	3 Months	6 Months
1	12	16	24	34
2	11	15	26	36
3	12	15	26	36
4	7	10	15	20
5	8	10	17	23
6	8	11	17	25

TABLE 4
RMISE, RISB and RIV Results from the First Experiment (Inverse Variance Weighting)

		SML						DFCH					
Scenario		2 weeks	1 month	3 months	6 months	2 weeks	1 month	3 months	6 months	2 weeks	1 month	3 months	6 months
RMISE	1	0.0063	0.0067	0.0087	0.0085	0.0016	0.0012	0.0033	0.0026	0.0029	0.0025	0.0018	0.0022
	2	0.0021	0.0021	0.0025	0.0029	0.0035	0.0053	0.0055	0.0079	0.0013	0.0012	0.0022	0.0012
	3	0.0060	0.0076	0.0084	0.0088	0.0011	0.0033	0.0023	0.0029	0.0046	0.0057	0.0073	0.0075
	4	0.0046	0.0057	0.0073	0.0075	0.0024	0.0028	0.0023	0.0029	0.0010	0.0020	0.0024	0.0028
	5	0.0010	0.0020	0.0024	0.0028	0.0024	0.0040	0.0122	0.0213	0.0044	0.0057	0.0069	0.0037
	6	0.0044	0.0057	0.0069	0.0037	0.0024	0.0040	0.0122	0.0213	0.0044	0.0057	0.0069	0.0037
RISB	1	0.0057	0.0065	0.0082	0.0083	0.0009	0.0010	0.0024	0.0023	0.0012	0.0008	0.0006	0.0005
	2	0.0012	0.0014	0.0022	0.0025	0.0013	0.0008	0.0006	0.0005	0.0058	0.0069	0.0083	0.0087
	3	0.0058	0.0069	0.0083	0.0087	0.0033	0.0043	0.0054	0.0078	0.0041	0.0054	0.0068	0.0074
	4	0.0041	0.0054	0.0068	0.0074	0.0005	0.0004	0.0008	0.0009	0.0009	0.0015	0.0023	0.0028
	5	0.0009	0.0015	0.0023	0.0028	0.0009	0.0010	0.0020	0.0026	0.0043	0.0057	0.0061	0.0032
	6	0.0043	0.0057	0.0061	0.0032	0.0024	0.0040	0.0121	0.0212	0.0043	0.0057	0.0061	0.0032
RIV	1	0.0027	0.0016	0.0029	0.0019	0.0013	0.0007	0.0022	0.0012	0.0017	0.0015	0.0017	0.0021
	2	0.0017	0.0015	0.0013	0.0014	0.0026	0.0024	0.0017	0.0021	0.0015	0.0032	0.0010	0.0011
	3	0.0015	0.0032	0.0010	0.0010	0.0011	0.0030	0.0012	0.0011	0.0021	0.0019	0.0026	0.0011
	4	0.0021	0.0019	0.0026	0.0011	0.0012	0.0011	0.0020	0.0008	0.0006	0.0014	0.0006	0.0006
	5	0.0006	0.0014	0.0006	0.0006	0.0007	0.0031	0.0012	0.0013	0.0003	0.0001	0.0031	0.0018
	6	0.0003	0.0001	0.0031	0.0018	0.0004	0.0005	0.0020	0.0019	0.0003	0.0001	0.0031	0.0018

TABLE 5

RMISE, RISB and RIV Results from the Second Experiment

	Inverse Variance		Equal		Vega	
	Weighting		Weighting		Weighting	
	SML	DFCH	SML	DFCH	SML	DFCH
RMISE	0.0222	0.0048	0.0200	0.0043	0.0181	0.0038
RISB	0.0222	0.0048	0.0200	0.0042	0.0180	0.0036
RIV	0.0008	0.0006	0.0011	0.0009	0.0015	0.0013

TABLE 6

DFCH Parameter Estimates with Bootstrapping Standard Errors

	17.05.96			28.04.97		
	1 M	3 M	12 M	1 M	3 M	12 M
a_2	4.0049 (0.1342)	2.0124 (0.1981)	4.9108 (0.1332)	0.9342 (0.4816)	3.3819 (1.1657)	2.1954 (1.2051)
a_3	4.2413 (0.1379)	2.1164 (0.2013)	4.9742 (0.1432)	1.0162 (0.4891)	9.2834 (0.7490)	9.3515 (0.5178)
b_2	-0.1245 (0.3343)	-0.6864 (0.3834)	-2.0400 (0.5193)	-0.1868 (0.1174)	-3.4646 (0.6269)	-3.8471 (0.7332)
b_3	2.3999 (0.2753)	1.4529 (0.2767)	1.0404 (0.1505)	1.7578 (0.1866)	2.0873 (0.1868)	1.9945 (0.1349)
b_4	-0.3592 (0.0117)	-0.4355 (0.1204)	-0.4424 (0.0286)	-0.4427 (0.1496)	-0.2564 (0.0082)	-0.3133 (0.0125)
m_1	-3.6789 (0.4141)	-2.3846 (0.4755)	-2.1942 (0.3858)	-3.4576 (0.3142)	-1.2262 (0.0543)	-1.2286 (0.0477)
m_2	-0.4062 (0.0041)	-0.6565 (0.0062)	-0.3708 (0.0033)	-0.8943 (0.0168)	-1.1466 (0.0079)	-1.1615 (0.0078)

TABLE 7
Summary Statistics for OTC Currency Option Data

		17.05.96			28.04.97		
		1 M	3 M	12 M	1 M	3 M	12 M
$\hat{\mu}$	DFCH	3.3896	3.3929	3.4121	3.3743	3.3761	3.3852
	SML	3.3896	3.3929	3.4121	3.3743	3.3761	3.3852
$\hat{\sigma}$	DFCH	0.0193	0.0393	0.0951	0.0250	0.0444	0.0843
	SML	0.0222	0.0437	0.1012	0.0291	0.0507	0.1012
<i>Skew</i>	DFCH	0.6502	1.1511	1.1468	1.6009	1.4277	1.4264
	SML	0.8335	1.2072	1.1710	1.5190	1.7126	1.8315
<i>Kurt</i>	DFCH	4.2705	5.6707	5.5896	7.0460	8.1539	8.2997
	SML	5.2820	5.6405	5.3788	5.9971	6.1997	6.3860
\overline{R}^2	DFCH	0.999910	0.999882	0.999952	0.999897	0.999754	0.999748
	SML	0.995770	0.979888	0.919298	0.996327	0.985903	0.938501
λ	SML	0.0035	0.0045	0.0045	0.0035	0.0150	0.0250

FIGURE 1
True RND and Estimated RNDs from the First Experiment

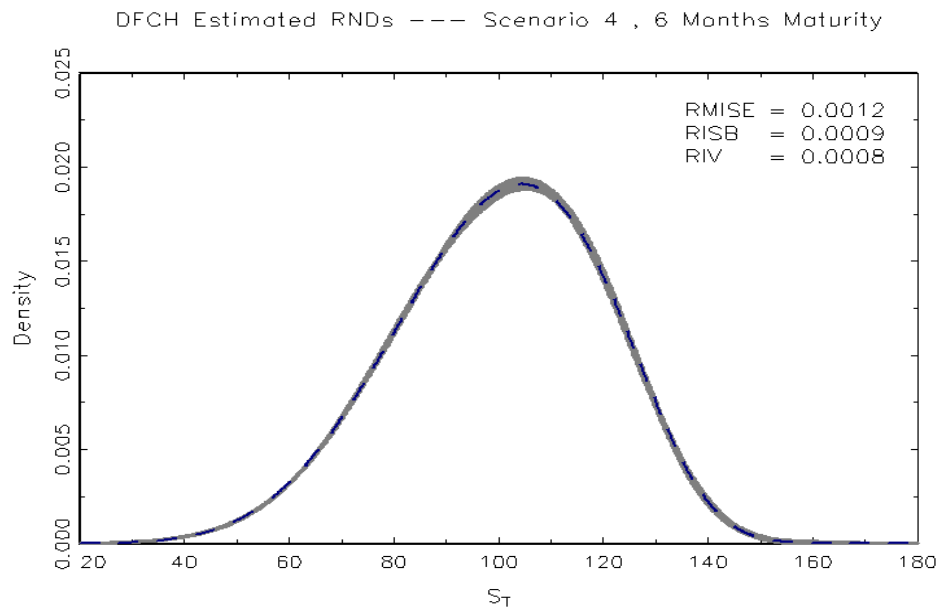
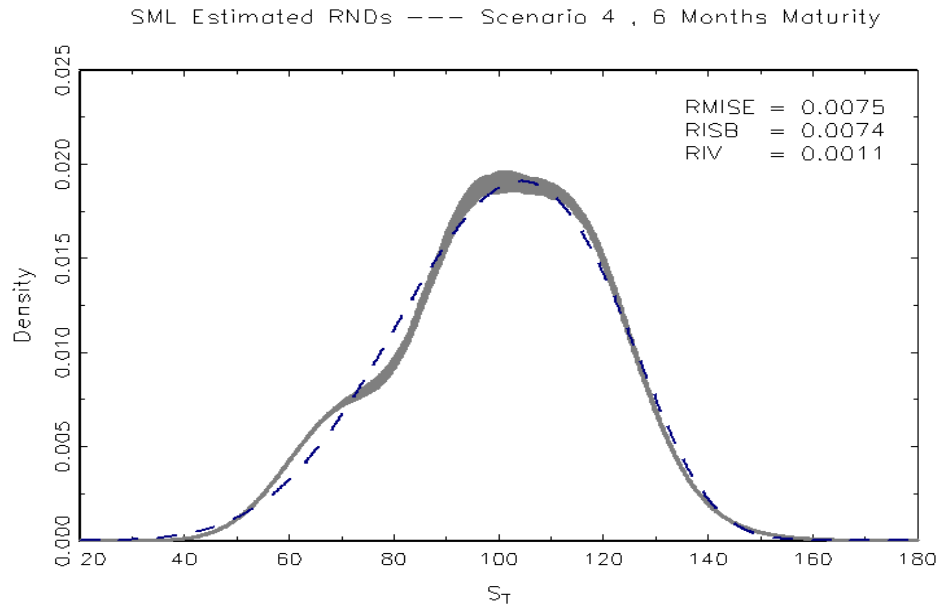


FIGURE 2

True RND and Estimated RNDs from the Second Experiment

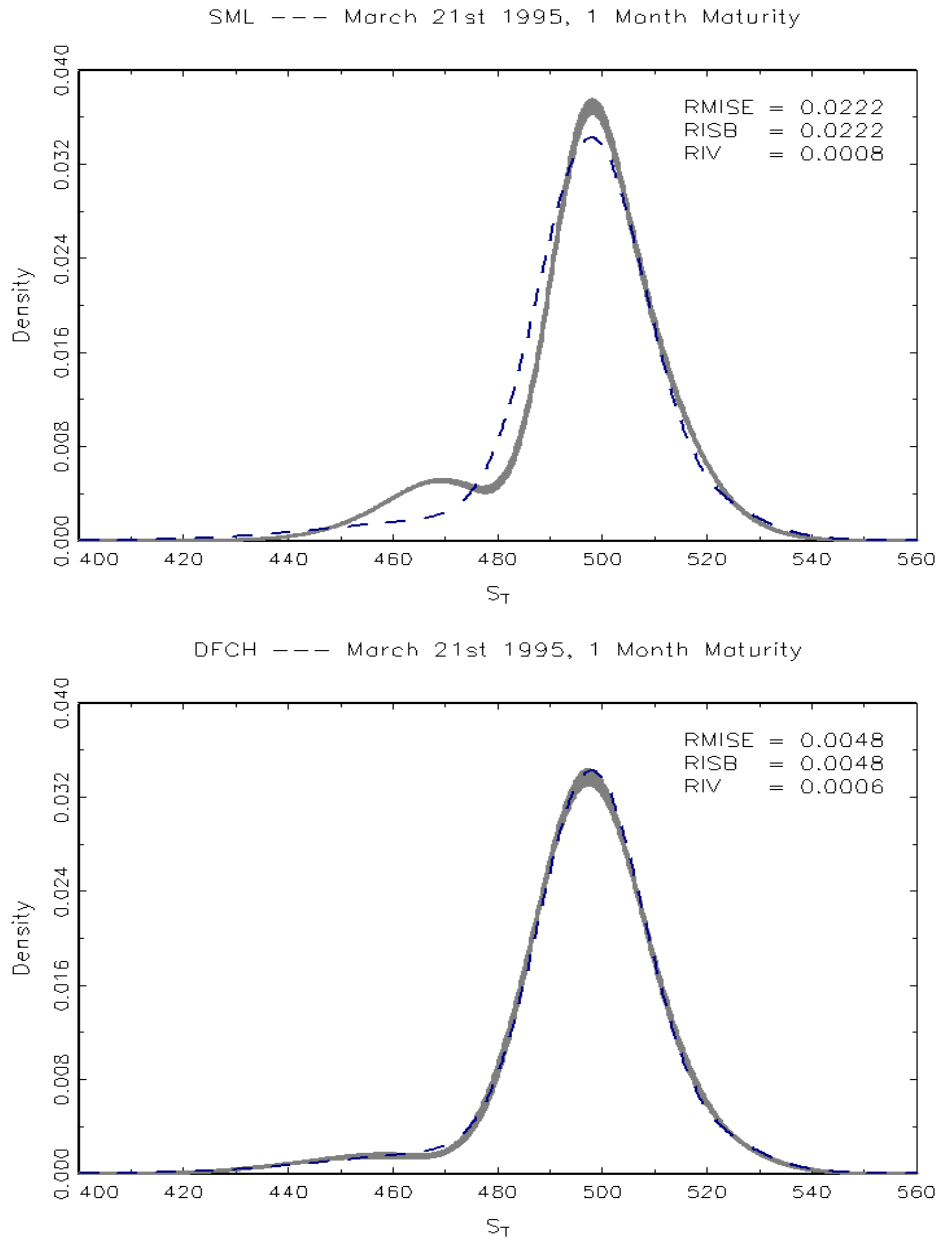


FIGURE 3
DFCH RND and SML RNDs under Different Smoothing Parameters

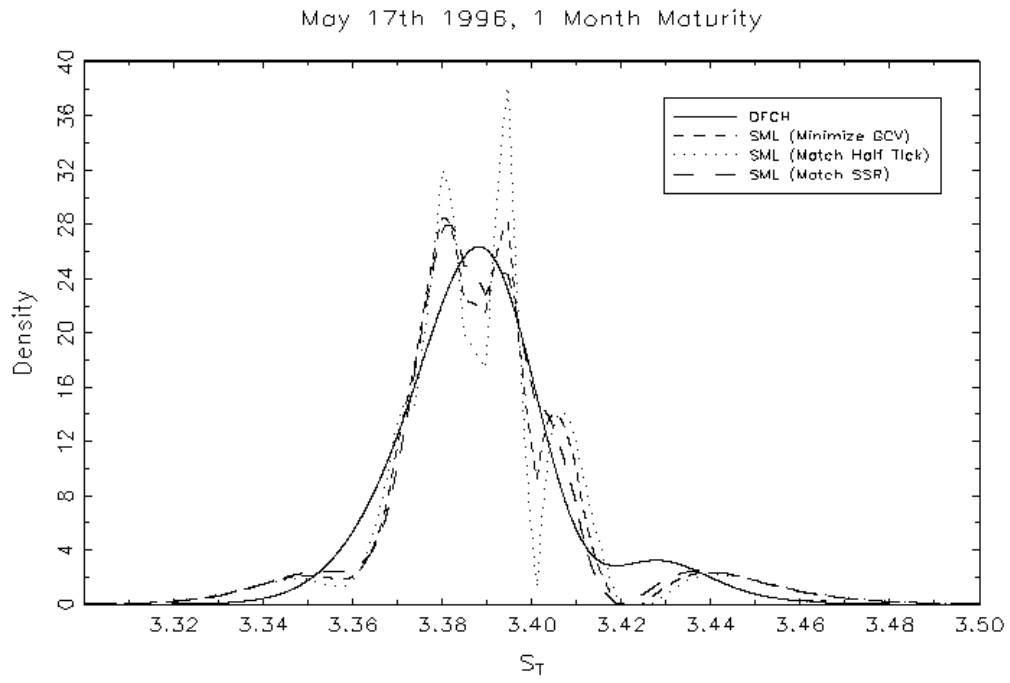


FIGURE 4
Estimated RNDs for the First Date

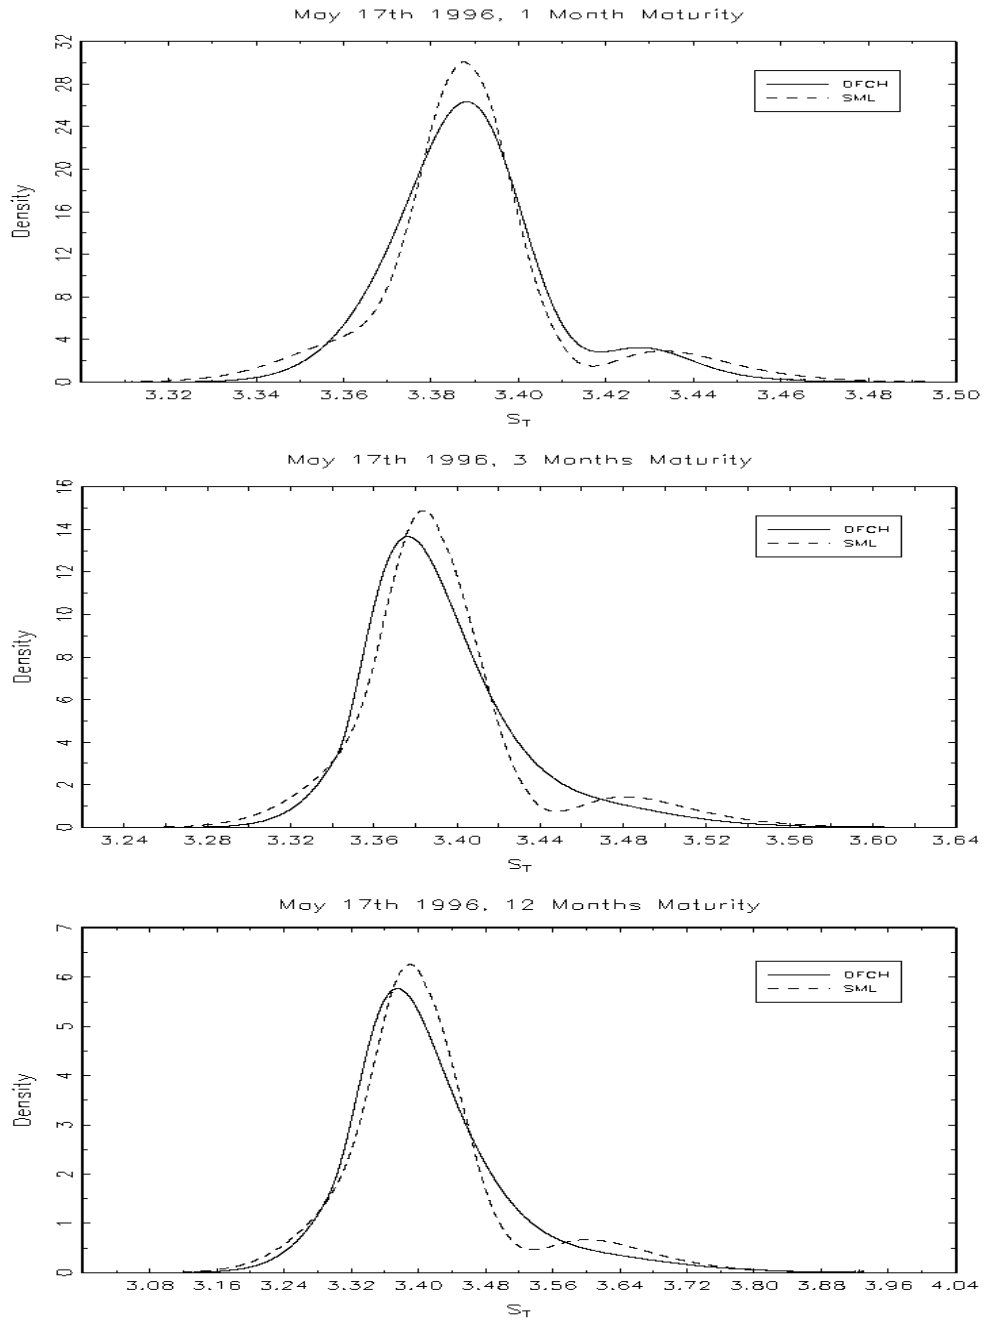


FIGURE 5:
Estimated RNDs for the Second Date

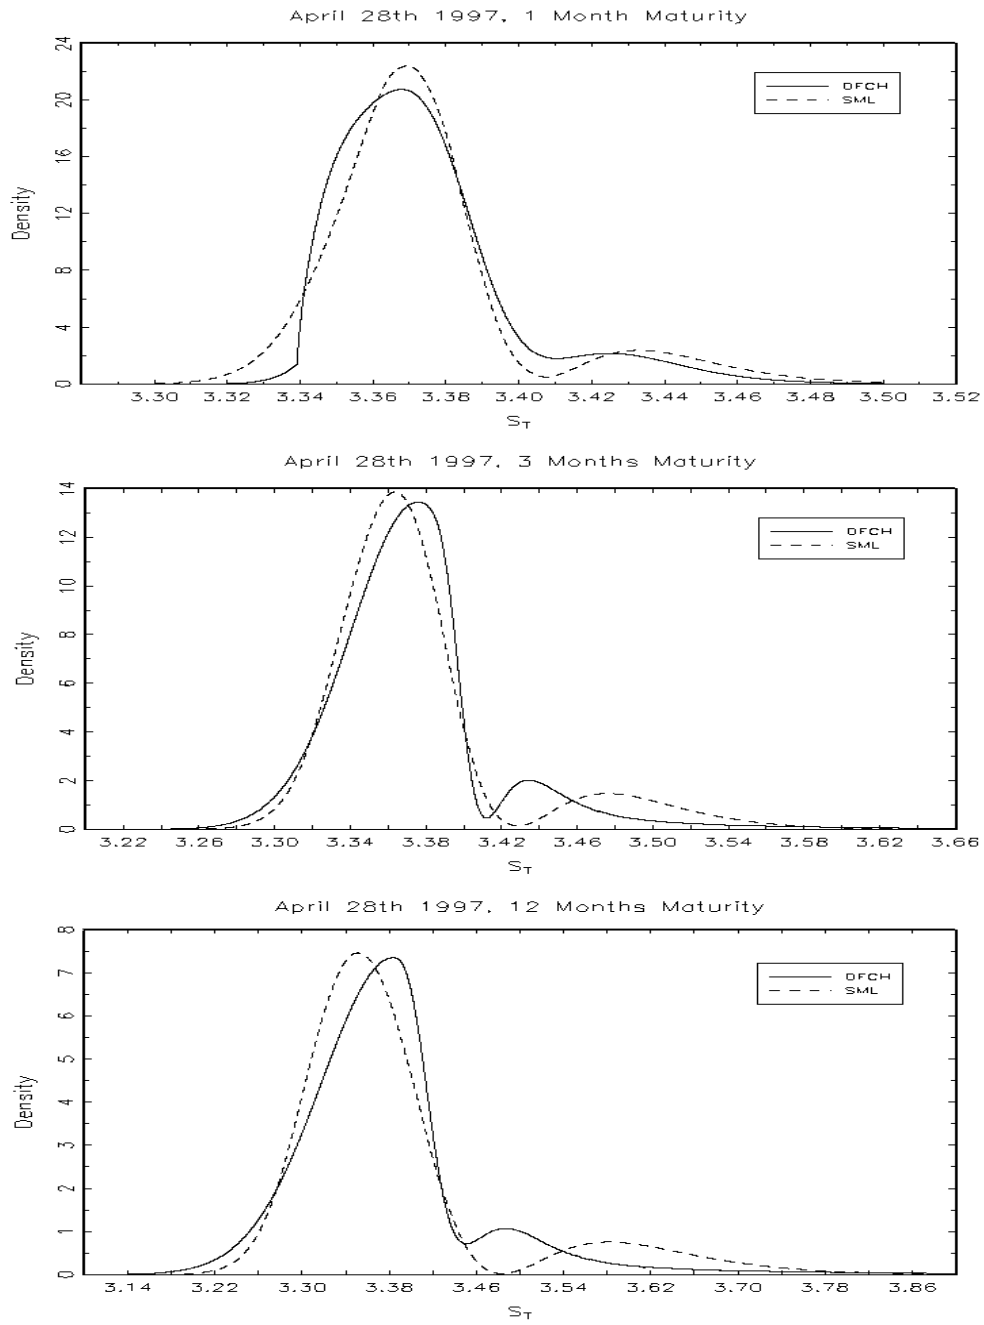


FIGURE 6
Original Volatilities and Fitted Volatility Function

