Risk-Neutral Densities Estimation: performance of Non-Structural Methods in a "true" world marked by jumps in asset returns

André Santos
Banco BPI, S.A.,
Largo Jean Monnet, 1 - 5, Lisboa,1269-067 Lisboa, Portugal
andre.duarte.santos@bancobpi.pt

João Guerra
CEMAPRE, ISEG, Universidade de Lisboa,
Rua do Quelhas 6, 1200-781 Lisboa, Portugal
jguerra@iseg.utl.pt

February 27, 2015

Abstract

The option prices can be used to extract the implied Risk-Neutral Density functions (RNDs) of the future underlying asset prices and returns. These market expectations provide valuable information that can be helpful to policy makers and investors.

In this work, we tested the accuracy and stability of four non-structural models in estimating the "true" RNDs. These models are the Density Functional Based on Confluent Hypergeometric function (DFCH), the Mixture of Lognormal distributions (MLN), the Smoothed Implied Volatility Smile (SML) and the Edgeworth expansions (EE). The "true" RND is unknown, so it was generated using a structural model called CGMY Gamma-OU, that is able to produce RNDs that are closer to reality (more leptokurtic and with a higher probability of extreme events). The CGMY Gamma-OU is a Lévy process that models the asset returns and volatility as two stochastic processes with jumps, so it is more appropriate to the "true" RND generation than stochastic processes based on diffusion processes, like the Heston model which was used in previous studies.

In order to build "true" RNDs that incorporate the characteristics of emerging market currency options, we used parameters which were estimated from the USD/BRL (US Dollar to Brazilian Real exchange rate) currency option prices.

We observed that the DFCH and MLN outperformed the SML and the EE models in capturing the "true" RND. The SML had the best performance in terms of stability.

JEL Classification: G13; C13; C15; F31

Keywords: Risk–neutral density, Option pricing, Lévy processes, Natural spline, Hypergeometric function, Brazilian real exchange rate.

1. Introduction

The prices of financial derivatives are mainly influenced by the expectations of the future underlying asset prices and by its uncertainty. Within this framework, the observed option prices play a very important role due to its capacity to provide information about implied Risk-Neutral Densities (RNDs).

The RNDs estimated from the cross-sections of observed option prices provide information that is vital for policy makers, investors and risk managers because they can identify and monitor valuable information, such as market concerns and probabilities about extreme movements in key asset and commodity prices (tail events). Recent steps have been taken with a view to increase the public visibility of RNDs output and its usage by policymakers, as is the case of the Federal Reserve Bank of Minneapolis, which provides commentary on its website regarding the moments and tail probabilities of the RNDs from various markets. In order to improve the RNDs interpretation, it is important to extract them using a model that closely reflects the “true” RNDs.

The standard model in option pricing, the Black and Scholes (B&S) model, has very strong assumptions, such as describing the asset returns as a stochastic process that evolves according to a geometric Brownian motion, with a constant expected return and a constant volatility. It is clear that the second assumption contradicts empirical evidence, as shown in two observed real world phenomena: different implied volatilities across maturities and across strike prices. The first phenomenon indicates that the volatility is stochastic over time. The second phenomenon is called volatility smile pattern and indicates higher volatilities for strike prices out-of-the-money. The use of Gaussian models in option pricing and RNDs estimation can be dangerous; such practice could lead to the underestimation of extreme losses and mispriced derivative products.

In order to tackle these problems, various parametric models have been suggested to extract RNDs from option prices and several studies have been carried out to examine the robustness of these estimates and their information power. Jondeau *et al.* [2007] divide the alternative models into two categories: structural and non-structural. A structural model assumes a specific dynamic for the price or volatility process (examples: B&S model and jump processes). A non-structural model allows the estimation of a RND without describing any stochastic process for the price or volatility of the underlying asset, which can give more flexibility in the RND estimation. The non-structural approaches can be divided into three subcategories: parametric (propose a form for the RND without assuming any price dynamics for the underlying asset), semi-parametric (suggest an approximation of the “true” RND) and non-parametric models (do not propose an explicit form for the RND).

In this work, we compare the performance of four non-structural models in the extraction of the “true” RND: the Density Functional Based on Confluent Hypergeometric function (DFCH), the Mixture of Lognormal distributions (MLN), the Smoothed Implied Volatility Smile (SML), and the Edgeworth expansions (EE).

The “true” RND is unknown, so it was simulated using a structural model called CGMY Gamma-OU, a rich and sophisticated stochastic volatility model that generates jumps for asset returns and volatility and, consequently, is able to generate plausible “true” RNDs, with statistical properties similar to the ones observed in the real world (leptokurtic and with fat tails). In Cooper [1999] and Santos and Guerra [2014], the “true” RND was build using the

Heston stochastic volatility model, a diffusion-based model based on two Brownian motion processes, one for the asset return and the other for the volatility. As is well known, diffusion processes cannot generate jumps, which damage their capacity to generate realistic short-term RNDs. The adoption of a more flexible process that generates RNDs with patterns that deviate from the Gaussian forms, addresses the criticism of Bliss and Panigirtzoglou [2002] regarding the dependence from the accuracy tests of the model used to simulated the "true" RND.

The characteristics of the USD/BRL options market were incorporated in the "true" RNDs through the application of the CGMY Gamma-OU parameters estimated from each month's currency options (end of month prices). At the end of this work, we added an analysis about a methodology to transform the RND into a Subjective Probability Density Function (SPDF) and about the goodness of fit of the SPDF to the observed values of the USD/BRL.

The remainder of this work is organized into six sections. Section 2 describes how option prices can provide information about the implied probabilities given by market participants to future events. Sections 3 and 4 describe the five models used in this work: CGMY Gamma-OU, DFCH, MLN, SML and EE. Section 5 presents the measures that evaluate the performance of these models in terms of accuracy and stability, as well as the results of the Monte Carlo simulation experiments that allow the comparison between the non-structural models. At the end of section 5 we also investigate whether the SPDFs obtained through non-structural models have a good fit to the observed returns of the exchange rate USD/BRL. Finally, section 6 presents the conclusions.

2. Relation between option prices and the extraction of RNDs

The value of a call option is given by the discounted value of its expected payoff on the expiration date *T*,

$$C(X,T) = e^{-rT}\int_{X}^{\infty} f(S_T)(S_T - X)dS_T\,, \qquad (1)$$

where *X* is the exercise price, S_T is the price of the underlying asset at time *T* and *r* is the risk-free interest rate. Under risk neutrality, the expectation is taken with respect to the risk neutral probabilities.

Breeden and Litzenberger [1978] were the first to realize that the RND can be recovered from the option prices, taking the second derivative of equation (1) with respect to the exercise price. Indeed,

$$\frac{d^2C(X,T)}{dX^2} = e^{-rT} f(X). \qquad (2)$$

3. Structural Models

The structural models are very important due to their capacity of modeling the paths of asset returns through time. The tractability of these stochastic processes is very useful for pricing and hedging purposes. For this reason, we generate the "true" RND using a structural model called CGMY Gamma-OU, which belongs to the class of Lévy processes with stochastic volatility. We consider that this process is capable of producing RNDs with statistical properties that are close to reality.

In the next sections we introduce Lévy processes and three models: the Variance Gamma, the CGMY and the CGMY Gamma-OU (adds stochastic volatility to the CGMY model).

3.1. Lévy processes and the Lévy-Khintchine formula

The most famous structural model used in option pricing is the B&S model. Due to its limitations, other models that allow more flexible distributions should be used. A structural model should be able to generate not only a flexible static distribution, but also a flexible stochastic process. Within this scope, it is important to use processes that capture not just small moves but also large moves (jumps). In fact, the non-gaussian nature and fat-tails phenomenon is mainly influenced by rare large moves in asset returns.

The class of Lévy processes, to which the Brownian motion belongs, encompasses sophisticated distributions and allows small moves and large jumps in asset returns and, at the same time, retains much of the tractability of geometric Brownian motion. The Merton jump-diffusion model, the Variance Gamma (VG), the Normal Inverse Gaussian (NIG), the CGMY, the Generalized Hyperbolic model and the Meixner model are examples of Lévy processes (see Schoutens, 2003).

A Lévy process $\{X_t, t \geq 0\}$ is essentially a stochastic process with independent and stationary increments, which also satisfies the property of stochastic continuity.

The well known Lévy-Khintchine formula provides a complete characterization of the distribution associated to a Lévy process (see Cont and Tankov, 2004). This formula states that the characteristic function of a Lévy process is given by

$$\Phi_X(u) = \exp(t\psi(u)),\ u \in \mathbf{R}, \tag{3}$$

where the characteristic exponent in equation (3) is called Lévy exponent and is given by

$$\psi(u) = i\gamma u - \frac{1}{2}\sigma^2 u^2 + \int_{-\infty}^{+\infty}(e^{iux} - 1 - iux1_{|x|\leq 1})\nu(dx), \tag{4}$$

where $\gamma \in \mathbf{R}$, $\sigma \in \mathbf{R}_+$ and ν is a positive Radon measure on $\mathbf{R}\setminus\{0\}$, called the Lévy measure and satisfying:

$$\int_{\mathbf{R}\setminus\{0\}} \min\{1, x^2\}\nu(dx) < \infty.$$

From the Lévy-Khintchine formula, it is seen that in general, the Lévy process can be decomposed into three independent parts represented by the Lévy triplet (γ, σ^2, ν):

1. The linear deterministic part γt. This component is analogous to a drift, but depends on the chosen truncation function (for more details see Cont and Tankov, 2004, Chapter 3),

2. The Brownian part σW_t, where W_t is a standard Brownian motion.

3. The pure jump part, which is characterized by ν. The Lévy measure $\nu(dx)$ dictates how jumps occur. Jumps of sizes in the set $A \in \mathbf{R}$ occur according to a Poisson process with intensity parameter $\int_A \nu(dx)$.

The Lévy measure also characterizes the activity and path properties of a Lévy process. When $\nu(\mathbf{R}) = \infty$, the Lévy process has infinite activity. This means that almost all the paths of *X* have an infinite number of jumps on $[0,T]$. If $\nu(\mathbf{R}) < \infty$, the Lévy process has finite activity. This means that almost all paths of *X* have only a finite number of jumps on $[0,T]$. If the Lévy process has no Brownian part $(\sigma = 0)$ and $\int_{|x|\leq 1} |x| \nu(dx) < \infty$, then almost all the paths of *X* have finite variation. If the Lévy process has a Brownian part $(\sigma \neq 0)$ or $\int_{|x|\leq 1} |x| \nu(dx) = \infty$, then almost all paths of *X* have infinite variation.

In finance, the price of an underlying asset S_t can be modeled by an exponential Lévy process and the risk-neutral dynamics is given by:

$$S_t = S_0 \exp(X_t), \tag{5}$$

where X_t is a Lévy process under the equivalent martingale measure **Q**. The use of different types of exponential Lévy models in finance depends mainly on the choice for the Lévy measure ν. Lévy processes used in finance can be divided into two categories:

1. Jump diffusion models, where the evolution of prices is given by a diffusion process punctuated by jumps at random intervals (rare events).

2. Infinite activity pure jump models. These models are able to capture both frequent small jumps and rare large jumps. Geman *et al.* [2001] have suggested that price processes of financial assets must have a jump component but there is no need to introduce a diffusion or Brownian component.

3.1.1. Variance Gamma process

The Variance Gamma (VG) process has been extensively used in the modeling of asset returns due to its capacity in capturing a wide range of realistic distributional and stochastic properties, while remaining relatively tractable (see Madan *et al.*, 1998 and Carr *et al.*, 2002). The VG is an infinite activity model with finite variance paths. It has no diffusion component, given the ability of its high activity to capture both frequent small moves and rare large jumps. The process has two possible representations: as a difference of two gamma processes or as a subordinating Brownian motion.

In the subordinating approach, the VG is defined as a Brownian motion with drift θ at a random time given by a gamma process,

$$X_{VG}(t) = \theta G_t + \sigma W(G_t), \tag{6}$$

where G_t is a gamma process, with mean t and variance rate νt, independent of the Brownian motion *W*. The characteristic function of the VG process is given by

$$\Phi_{X_{VG}}(u,t) = \left(\frac{1}{1 - i\theta\nu u + \sigma^2 \nu u^2/2} \right)^{t/\nu}, \tag{7}$$

The definition of the VG process as a difference of two gamma processes is useful because it determines the representation of its Lévy measure (see Madan *et al.*, 1998 and Carr *et al.*, 2002):

$$\nu_{VG}(dx) = \begin{cases} \dfrac{C \exp(G|x|)}{|x|} \text{for} x < 0 \\ \dfrac{C \exp(-Mx)}{x} \text{for} x > 0 \end{cases}, \tag{8}$$

where

$$C = 1/\nu > 0,$$
$$G = \left(\sqrt{\frac{1}{4}\theta^2\nu^2 + \frac{1}{2}\sigma^2\nu} - \frac{1}{2}\theta\nu \right)^{-1} > 0,$$
$$M = \left(\sqrt{\frac{1}{4}\theta^2\nu^2 + \frac{1}{2}\sigma^2\nu} + \frac{1}{2}\theta\nu \right)^{-1} > 0.$$

3.1.2. CGMY process

Carr *et al.* [2002] generalized the VG Lévy density in equation (8) in order to obtain the CGMY process (named after Carr, Geman, Madan and Yor), a pure jump process that generalizes the VG process allowing finite activity and infinite variation. The Lévy density of the CGMY process is given by

$$f_{CGMY}(dx) = \begin{cases} \dfrac{C \exp(-G|x|)}{|x|^{1+Y}} \text{for} x < 0 \\ \dfrac{C \exp(-Mx)}{x^{1+Y}} \text{for} x > 0 \end{cases}, \tag{9}$$

where $Y < 2$ is also a model parameter.

A process increment over a time interval of length t is distributed according to the distribution $CGMY(tC, G, M, Y)$, with characteristic function

$$\Phi_{X_{CGMY}}(u, tC, G, M, Y) = \exp(Ct\Gamma(-Y)((M-iu)^Y - G^Y)). \qquad (10)$$

The Y parameter controls the integrability of the Lévy density, characterizing the structure of the process by controlling the behavior of small jumps. The case of $Y = 0$ is the special case of the VG process. If $Y < -1$, the Lévy density is not monotonic, which means that the small jumps are less frequent than the large jumps (finite activity and finite variation). If $Y > -1$, the Lévy density becomes monotonic. The density has finite activity if $Y < 0$. When $0 < Y < 1$, the process has infinite activity but is still of finite variation. Finally, if $Y > 1$, the process has infinite activity and infinite total variation. The parameter C can be viewed as a measure of the overall level of activity, scaling the expected number of jumps of all sizes. The parameters G and M control the rate of exponential decay on the right and left tails of the Lévy density, respectively.

3.1.3. CGMY Gamma-OU process

The introduction of jumps proved to be efficient in generating implied volatility patterns only for a single short maturity. Nevertheless, due to the independence of log-returns, the exponential-Lévy models fail to reproduce the implied volatility smiles and skews observed in options market prices over a range of different maturities. This drawback is due to the volatility clustering phenomenon. A solution to this problem is to introduce a stochastic process for volatility.

The introduction of a volatility stochastic process adds another dimension to the analysis. Now, we are modeling log returns using changes in prices and changes in volatility. In order to capture the positiveness and volatility clustering, the volatility evolution should be captured by a process that is positive and mean-reverting. We can build this process making time stochastic: in periods of high volatility time runs faster than in periods of low volatility. Thus, random changes in volatility can be captured by random changes in time. One of the alternatives to model this time change is through the Gamma-OU process. The concept of stochastic time in asset pricing was first used in Clark [1973].

The Ornstein-Uhlenbeck (OU) process was introduced by Barndorff and Shephard [2001] as a model to describe volatility in finance. It is defined as the solution of the SDE

$$dy_t = -\lambda y_t dt + dz_{\lambda t}, y_0 > 0,$$

where $\{z_t, t \geq 0\}$ is a subordinator (strictly positive and increasing Lévy process) and $\lambda > 0$. The process z_t is also called a Background Driving Lévy Process (BDLP).

The D-OU process $\{y_t, t \geq 0\}$ is a nonnegative nondecreasing Lévy process with no Brownian part, nonnegative drift and only positive increments. This process is strictly stationary on the positive half-line, that is, there exists a law D, called the stationary law or the marginal law, such that y_t will follow the law D for every t if y_0 is chosen according to D. In the case of a Gamma-OU process, the BDLP $\{z_t, t \geq 0\}$ is a Gamma process that has a Lévy density of the form

$$w(x) = ab\exp(-bx)\mathbf{1}_{\{x>0\}},$$

where a and b are positive parameters. The Gamma-OU process has the characteristic function

$$\Phi_{Y_{Gamma-OU}}(u,t,\lambda,a,b,y_0) = \exp(iuy_0\lambda^{-1}(1-\exp(-\lambda t)))\times \quad (11)$$

$$\exp\left(\frac{\lambda a}{iu-\lambda b}\left(b\log\left(\frac{b}{b-iu\lambda^{-1}(1-\exp(-\lambda t))}\right)-iut\right)\right). \quad (12)$$

The stochastic process CGMY Gamma-OU transforms the CGMY process X_t into a stochastic volatility process Z_t through a random time change process, modeled by an increasing Lévy process (subordinator) Y_t independent of the original process. In this case, we are subordinating a CGMY process (X_t) to a Gamma-OU process (Y_t). This type of stochastic volatility for Lévy processes was introduced in Carr *et al.* [2003].

The characteristic function of the CGMY Gamma-OU process is given by

$$E(\exp(iuZ_t)) = E(\exp(Y_t\psi_{X_{CGMY}}(u))) \quad (13)$$

$$= \Phi_{Y_{Gamma-OU}}(-i\psi_{X_{CGMY}}(u),t,\lambda,a,b,y_0)\,. \quad (14)$$

For more details about the subordination of Lévy processes, see Cont and Tankov [2004] and Sato [1999].

Given the characteristic function of the Lévy process, the RND can be extracted from option prices with a method based on the Fourier Transform. For this calculation we follow Carr and Madan [1999]. The method consists of two parts. First, we obtain an analytical expression for the Fourier transform of an option price. Then, we recover the option price by inverting the Fourier transform using the Fast Fourier Transform (FFT) algorithm.

4. Non-Structural Models

The non-structural models do not assume any stochastic process for the underlying asset price and focus directly on the RND, which provides more flexibility and makes them more appropriate in recovering the "true" RND.

4.1. Parametric models

4.1.1. Density Functional Based on Confluent Hypergeometric function

This model allows the estimation of a RND without assuming a specific functional form for it. It consists on the use of a formula that encompasses various densities, such as the normal, gamma, inverse gamma, Weibull, Pareto and mixtures of these probability densities.

Abadir and Rockinger [2003] developed a functional based on the confluent hypergeometric function ${}_1F_1$. These authors suggest that the usefulness of ${}_1F_1$ relies on the fact that it includes special cases of the incomplete gamma, normal and mixtures of the two distributions. In fact, this function has the advantage of being more efficient than fully nonparametric estimation methods for small samples and more flexible than other parametric methods, since it does not restrict functional forms.

The confluent hypergeometric function can be defined by

$$
\begin{aligned}
&{}_1F_1(\alpha;\beta;z) \equiv \sum_{j=0}^{\infty} \frac{(\alpha)_j}{\beta_j}\frac{z^j}{j!} \equiv 1 + \frac{\alpha}{\beta}z + \frac{\alpha(\alpha+1)}{\beta(\beta+1)}\frac{z^2}{2} + \ldots, \\
&(a)_j \equiv (a)(a+1)\ldots(a+j-1) \equiv \frac{\Gamma(a+j)}{\Gamma(a)},
\end{aligned}
\tag{15}
$$

with $\Gamma(v)$, for $v \in \mathbf{R} \setminus \{\ldots,-2,-1,0\}$, being the gamma function and $-\beta \notin \mathbf{N} \cup \{0\}$.

The functional considered for option pricing is called density functional based on confluent hypergeometric function (DFCH) and specifies the European call price as a mixture of two confluent hypergeometric functions:

$$
\begin{aligned}
&C(X) = c_1 + c_2 X + \mathbf{1}_{\{X>m_1\}} a_1((X-m_1)^{b_1})\, {}_1F_1(a_2;a_3;b_2(X-m_1)^{b_3}) \\
&+ (a_4)\, {}_1F_1(a_5;a_6;b_4(X-m_2)^2),
\end{aligned}
\tag{16}
$$

where $-a_3,-a_6 \notin \mathbf{N} \cup \{0\}$, $b_2, b_4 \in \mathbf{R}^-$, $a_1, a_2, a_4, a_5, b_1, b_3 \in \mathbf{R}$ and $\mathbf{1}$ represents the indicator function.

Some restrictions must be set in order to guarantee that f in equation (2) integrates to 1. Through these restrictions, we obtain that

$$
\begin{aligned}
c_2 &= -1 + a_4\sqrt{-b_4\pi}, \\
a_4 &= \frac{1}{2\sqrt{-b_4\pi}}\left[1 - a_1(-b_2)^{-a_2}\frac{\Gamma(a_3)}{\Gamma(a_3-a_2)}\right].
\end{aligned}
\tag{17}
$$

Using assumptions $c_1 = -c_2 m_2$, $b_1 = 1 + a_2 b_3$, $a_5 = -\frac{1}{2}$, $a_6 = \frac{1}{2}$, formula (16) can be further simplified (see Abadir and Rockinger, 2003) and the final number of parameters to estimate in the calculation of the theoretical price in equation (16) is reduced to seven.

4.1.2. Mixture of lognormal distributions

The mixture of lognormal distributions (MLN) was proposed by Bahra [1997] and Melick and Thomas [1997] and assumes a functional form for the RND that accommodates various stochastic processes for the underlying asset price. Instead of specifying the dynamics for the underlying asset price (which leads to a unique terminal value), it is possible to make assumptions about the functional form of the RND function itself and then obtain the parameters of the distribution by minimizing the distance between the observed option prices and those that are generated by the assumed parametric form. According to the authors, this makes this model more flexible than the B&S model and increases its ability to capture the main contributions to the smile curve, namely the skewness and the kurtosis of the underlying distribution.

It is well known that the price of a European call option at time *t* can be expressed as

$$c(X,\tau) = e^{-r\tau}\int_X^{\infty} f(S_T)(S_T - X)dS_T, \tag{18}$$

where $\tau = T - t$. Bahra [1997] assumed the lognormal distribution as the functional form of $f(S_T)$ and considered that it would be plausible to use a weighted sum of two lognormal density functions,

$$f(S_T) = wL(\alpha_1, \beta_1, S_T) + (1-w)L(\alpha_2, \beta_2, S_T), \tag{19}$$

where $L(\alpha_i, \beta_i, S_T)$ is the i-th lognormal density with parameters α_i and β_i:

$$\begin{gathered} L(\alpha_i, \beta_i, S_T) = \frac{1}{S_T \beta_i \sqrt{2\pi}} e^{[-(\ln(S_T)-\alpha_i)^2/2\beta_i^2]}, \\ \alpha_i = \ln(S_t) + (\mu_i - \frac{1}{2}\sigma_i^2)(T-t), \\ \beta_i = \sigma_i \sqrt{(T-t)}. \end{gathered} \tag{20}$$

The price of an European call option is given by (see Jondeau *et al.*, 2007)

$$\begin{gathered} c(X,\tau) = e^{-r\tau}\{w[e^{\alpha_1 + \frac{1}{2}\beta_1^2} N(d_1) - XN(d_2)] \\ + (1-w)[e^{\alpha_2 + \frac{1}{2}\beta_2^2} N(d_3) - XN(d_4)]\}, \end{gathered} \tag{21}$$

where

$$\begin{gathered} d_1 = \frac{-\ln(X) + \alpha_1 + \beta_1^2}{\beta_1}, \\ d_2 = d_1 - \beta_1, \\ d3 = \frac{-\ln(X) + \alpha_2 + \beta_2^2}{\beta_2}, \\ d_4 = d_3 - \beta_2. \end{gathered} \tag{22}$$

4.2. Semi-parametric models

4.2.1. Edgeworth expansions

The B&S assumption of a lognormal distribution for the underlying asset is relaxed in this method by using a more flexible distributional form, based on an Edgeworth series expansion around a lognormal distribution. This technique was developed by Jarrow and Rudd [1982]. Edgeworth expansions are conceptually similar to Taylor expansions, but are applied to densities instead of points.

In order to show how an Edgeworth expansion can be obtained, consider the "true" cumulative distribution function (cdf) F, the approximating cdf L, the probability density function (pdf) *f*, the approximating pdf *l*, the random variable S_T and the characteristic function of *F*, $\phi_F(u) = \int_{-\infty}^{\infty} e^{ius} f(s)ds.$ Given that n moments $u_j(F)$ exist, the first $n-1$ cumulants $k_j(F)$ also exist and are defined by the expansion (the cumulant-generating function)

$$\log(\phi_F(u)) = \sum_{j=1}^{n-1} k_j(F)\frac{(it)^j}{j!} + o(u^{n-1}).$$

Thus, if the characteristic function $\phi_F(u)$ is known, it is possible to obtain the cumulants by taking an expansion of its logarithm around $u = 0$. Jarrow and Rudd [1982] showed that after imposing the equality of the first moment of the approximating density and true density, the implied probability density function can be written as

$$f(S_T) = l(S_T) + \frac{k_2(F)-k_2(L)}{2!}\frac{d^2 l(S_T)}{dS_T^2} - \frac{k_3(F)-k_3(L)}{3!}\frac{d^3 l(S_T)}{dS_T^3} + \frac{(k_4(F)-k_4(L))+3(k_2(F)-k_2(L))^2}{4!}\frac{d^4 l(S_T)}{dS_T^4} + \varepsilon(S_T), \tag{23}$$

where $\varepsilon(S_T)$ captures the neglected terms of the fourth order expansion.

Jarrow and Rudd [1982] suggested the lognormal distribution $l(S_T)$ as the approximating function and imposed the condition $k_2(F) = k_2(L)$. After imposing the equalities of the first and second moments of the approximating and true densities and using the pdf given in equation (23) into equation (2), we have the following option price formula:

$$C(F) = C(L) - e^{-rT}\frac{k_3(F)-k_3(L)}{3!}\frac{dl(X)}{dS_T} \tag{24}$$

$$+e^{-rT}\frac{(k_4(F)-k_4(L))+3(k_2(F)-k_2(L))^2}{4!}\frac{d^2l(K)}{dS_T^2}$$

Corrado and Su [1986] rewrote equation (26) in order to estimate directly the skewness, $\gamma_1(F)$, and the kurtosis, $\gamma_2(F)$, using the relationships

$$\gamma_1(F)=\frac{k_3(F)}{k_2(L)^{3/2}}, \qquad \gamma_2(F)=\frac{k_4(F)}{k_2(L)^2}, \tag{25}$$

The parameters to be estimated are the implied volatility σ^2, the skewness $\gamma_1(F)$ and the kurtosis $\gamma_2(F)$.

4.3. Non-parametric models

4.3.1. Spline methods

These methods are based on the derivation of the RND using the results of Breeden and Litzenberger [1978], but with a preliminary process of smoothing the volatility smile. The first Spline method approach was developed by Shimko [1993]. In this work, we used the method proposed in Bliss and Panigirtzoglou [2002], and applied a natural cubic spline in the volatility/delta space and a smoothness parameter λ, which weights the degree of curvature of the spline function. This method consists in connecting the adjacent points (Δ_i,σ_i), $(\Delta_{i+1},\sigma_{i+1})$, using the cubic functions $\hat{\sigma}_i, i=0,...,n-1$, in order to piece together a curve with continuous first and second order derivatives. Thus, $\hat{\sigma}_i$ is a third order polynomial defined by

$$\hat{\sigma}_i(\Delta)=d_i+c_i(\Delta-\Delta_i)+b_i(\Delta-\Delta_i)^2+a_i(\Delta-\Delta_i)^3, \tag{26}$$

with $\Delta\in[\Delta_i,\Delta_{i+1}]$.

The natural spline minimizes the objective function

$$\min_\theta (1-\lambda)\sum_{i=1}^{N}w_j(\sigma_i-\hat{\sigma}_i(\Delta_i,\theta))^2+\lambda\int_{-\infty}^{\infty}(\sigma''(\Delta;\theta))^2d\Delta, \tag{27}$$

where N is the number of quoted deltas, $\hat{\sigma}_i(\Delta_i,\theta)$ is the implied volatility corresponding to the spline parameters represented by vector θ and w_i represents the weight attributed to each observation. The first term measures the goodness of fit and the second term measures the smoothness of the spline.

The variable regarding the weight parameter w in equation (27) is described by Bliss and

Panigirtzoglou [2002] as a source of price error. It is known that in the context of the B&S formula, the only unobservable parameter is volatility (σ), which means that the uncertainty regarding the probability density function lies in σ. The greek vega (v) measures the sensitivity of the option price with respect to σ and reflects the uncertainty concerning the volatility. As in Bliss and Panigirtzoglou [2002], we use this v weighting when fitting the volatility smile because this weighting scheme places more weight on near-the-money options and less weight on away-from-the-money options. We tested this method using the value λ that minimizes the RMISE (root mean integrated squared error).

5. Accuracy and Stability analysis

5.1. Data

The data used in this work was the currency OTC option prices with the underlying USD/BRL (price of US dollars in terms of Brazilian reals). The quotes were taken from the daily settlement bid prices in Bloomberg for Offshore USDBRL FX Options[1]. The data covers the period from June 2006 (half a year before the problems regarding the subprime crisis started to worsen) to February 2010 (seven months after the Brazilian general election) and comprises the monthly quotes (end of month prices).

The calls and puts used are of the European type and are priced in volatility as a function of delta. The grid of quoted deltas is 0.05, 0.1, 0.15, 0.25, 0.35 and 0.5 deltas. This means that we only considered out-of-the money options (calls and puts) and at-the-money options[2], due to the general understanding that out-of-the-money options tend to be more liquid than in-the-money options. We estimate the RNDs using one, three and six months to maturity options.

5.2. Testing RND estimation techniques using a Monte Carlo approach

In order to test the accuracy of the referred non-structural models at capturing the risk-neutral density functions, we have to see how closely they fit the “true” RND. For each month, we build a “true” RND from just that month’s option prices (using the calibrated CGMY Gamma-OU parameters). This way, we generate 45 “true” RNDs for the whole period (because we follow this process for each maturity – 1, 3 and 6 months – we have ended up with 135 “true” RNDs) between June 2006 and February 2010. With this process, we simulated “true” RNDs that incorporate the characteristics of the USD/BRL option market over an extended period, full of market events. For each “true” RND we obtained the corresponding “true” option prices.

In order to test the robustness of the DFCH, MLN, SML and EE models, we add a uniformly distributed random noise in the “true” option prices of size between minus half and half of the tick size (according to BM&FBOVESPA, the minimum tick size is 0.001) as in

[1] Information provided by Bloomberg for the OTC Market. The USDBRL is quoted in volatility in terms of delta according to international conventions (does not use the specific maturity of BM&F calendar and a day count of business days/252 just like other financial instruments traded in BM&F)

[2] The delta value varies from 0 for very out-the-money options to 1 for deeply in-the-money options. At the money options have a delta clclose to 0.5.

Bliss and Panigirtzoglou [2002]. Given these shocked option prices, we use the four non-structural models to estimate the "true" RNDs. This process of first shocking prices and then fitting the RND is repeated 500 times for the 135 combinations of maturities and dates. The performance of these models is then evaluated by the root mean integrated squared error (RMISE) as in Bu and Hadri [2007]. We did not examine the bias of the higher statistical moments, skewness and kurtosis as in Cooper [1999] and Santos and Guerra [2014], due to the poor reliability of these measures, which are extremely sensitive to the tails of the distribution. The adoption of the CGMY Gamma-OU process in the "true" RND generation allows us to avoid the criticism of Bliss and Panigirtzoglou [2002] regarding the dependence from the accuracy tests of the model used to simulated the "true" RND: the CGMY Gamma-OU model is much more flexible than the Heston model (used in Cooper, 1999, Bu and Hadri, 2007, and Santos and Guerra, 2014) and can produce RNDs with more forms and patterns, allowing leptokurtic distributions that deviate from the Gaussian paradigm and are closer to reality.

By considering $\hat{f}(S_t)$ as the estimator of the "true" RND, we define the Root Mean Integrated Squared Error (RMISE) by

$$RMISE(\hat{f}) = \sqrt{E[\int_0^\infty (\hat{f}(S_t) - f(S_t))^2 dS_t]}, \tag{28}$$

representing a measure of the average integral of the squared error over the support of the RND. It is a measure of the quality of the estimator. The square of the RMISE can also be broken down into the sum of the square of the RISB (root integrated squared bias) and the square of the RIV (root integrated variance):

$$RMISE^2(\hat{f}) = RISB^2(\hat{f}) + RIV^2(\hat{f}), \tag{29}$$

$$RISB(\hat{f}) = \sqrt{\int_0^\infty (E[\hat{f}(S_t)] - f(S_t))^2 dS_t}, \tag{30}$$

$$RIV(\hat{f}) = \sqrt{\int_0^\infty E[(\hat{f}(S_t) - E[\hat{f}(S_t)])^2] dS_t}. \tag{31}$$

The model with the best overall performance will have the lower RMISE: a higher accuracy is represented by a lower RISB and a higher stability is represented by a lower RIV.

5.3. Performance of the non-structural models in capturing "true" RNDs

In order to summarize the information contained in the series of the monthly statistical measures between June 2006 and February 2010 (45 data points for each maturity and statistical indicator), we present tables with the quartiles of these distributions (Tables A.1, A.2 and A.3 in the appendix).

Besides summarizing the statistical measures for the whole period, the results were also analyzed looking separately to three subperiods: the first is from June 2006 to July 2007 (after the subprime crisis became apparent in August 2007 - hereafter called "normal period"), the second is between August 2007 and August 2008 (a period characterized by the uncertainty regarding the seriousness of the problem - hereafter called "rumors period") and the third is between September 2008 and February 2010 (the period of higher turbulence in financial markets, where the main events regarding the subprime crisis took place - hereafter called "peak period").

5.3.1. Analysis using RMISE

The results regarding the evolution of the RMISE values for the analyzed period are shown in Figures 1, 2 and 3. These values, as well as the RISB and RIV measures, are summarized in Tables A.1, A.2 and A.3 in the appendix.

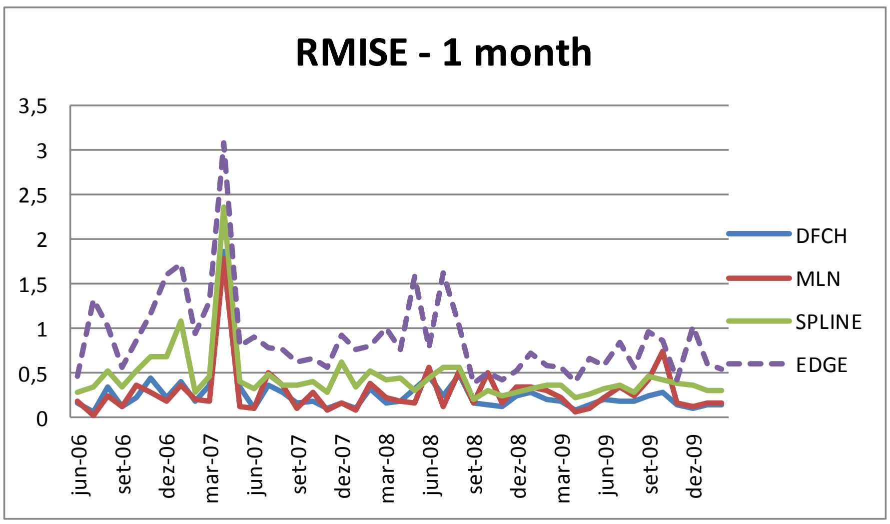

Best method	nr. times [a]
DFCH	26
MLN	18
SPLINE	1
EDGE	0

Figure 1: One-month RMISE for the period between June 2006 and February 2010.
a) number of times of each model as the best model.

In general terms, the DFCH and MLN models outperformed the other models as overall estimators of the "true" RND. The DFCH had the lowest RMISE in 78 pairs of dates and maturities and the MLN model had inferior RMISE values 56 times. The worst RND estimator was the EE model, having the higher RMISE values in the majority of dates.

In Table A.2 it is noticed a slightly better performance of the DFCH over the MLN model in terms of accuracy, having lower quartiles RISB values most of times. The DFCH model had the smallest RISB 71 times (29 dates in one month term, 18 dates in three months term and 24 dates in six months term) and the MLN model had an inferior RISB 64 times (16 dates in one month term, 27 dates in three months term and 21 dates in six months term).

Analyzing the results in more detail, we observe that the MLN model had a slightly better accuracy in the "normal period" for one month and three months terms. Nevertheless the DFCH model had a better fit for the six months term.

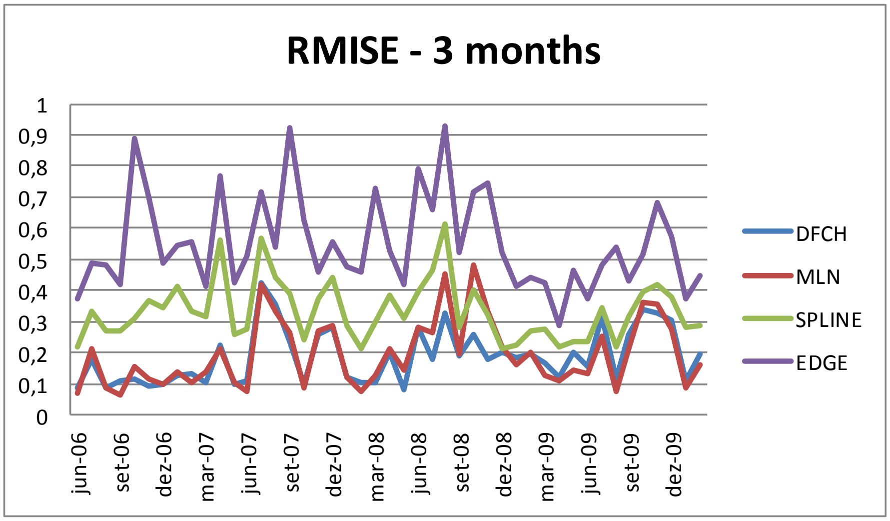

	Best method nr. times [a]
DFCH	25
MLN	20
SPLINE	0
EDGE	0

Figure 2: Three-month RMISE for the period between June 2006 and February 2010.
a) number of times of each model as the best model.

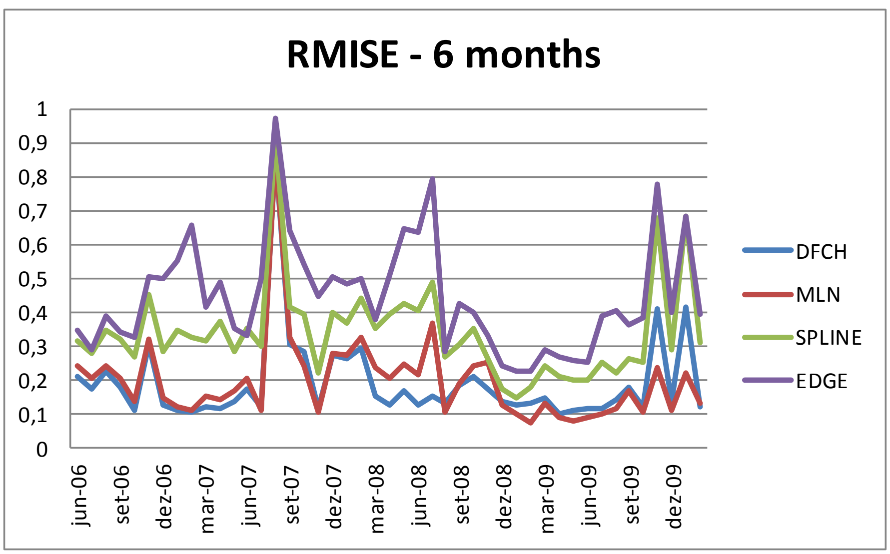

	Best method nr. times [a]
DFCH	27
MLN	18
SPLINE	0
EDGE	0

Figure 3: Six-month RMISE for the period between June 2006 and February 2010.
a) number of times of each model as the best model.

For the “rumors period”, the DFCH outperforms the MLN as the “true” RND estimator for all the maturities, as demonstrated in the lower RMISE values. It is interesting to note that in the first months of the “peak period”, the DFCH model was capable of producing the RNDs with the lower bias in comparison to the true ones for all the maturities. Let us recall that the four months between September 2008 and December 2008 were marked by a sequence of negative events and defined the peak of the subprime crisis: in September 2008, the Government-sponsored enterprises Fannie Mae and Freddie Mac, which owned or guaranteed about half of the U.S mortgage, were taken over by the US government, Lehman

Brothers filed for bankruptcy and the Bank of America purchased Merril Lynch; in October 2008 the US government bailed out Goldman Sachs and Morgan Stanley. This indicates that during this stressing period, characterized by the maximum uncertainty levels and highest probability of extreme events, the DFCH showed a higher capacity and flexibility to capture these abnormal "true" RNDs. This can be observed in Figure 4, where the RND produced through the DFCH model is closer to the "true" RND.

For the remainder "peak period" the DFCH model showed the highest accuracy for the one month term and the MLN model performed better in the longer terms.

Taking the RIV measure as the stability indicator, we conclude that the SML model was the most stable and the EE model was the most unstable. It is also important to compare the stability of the two models that had the best performance in terms of accuracy. The MLN was the most stable model for one month to maturity RNDs and the DFCH had the higher stability in the longer terms.

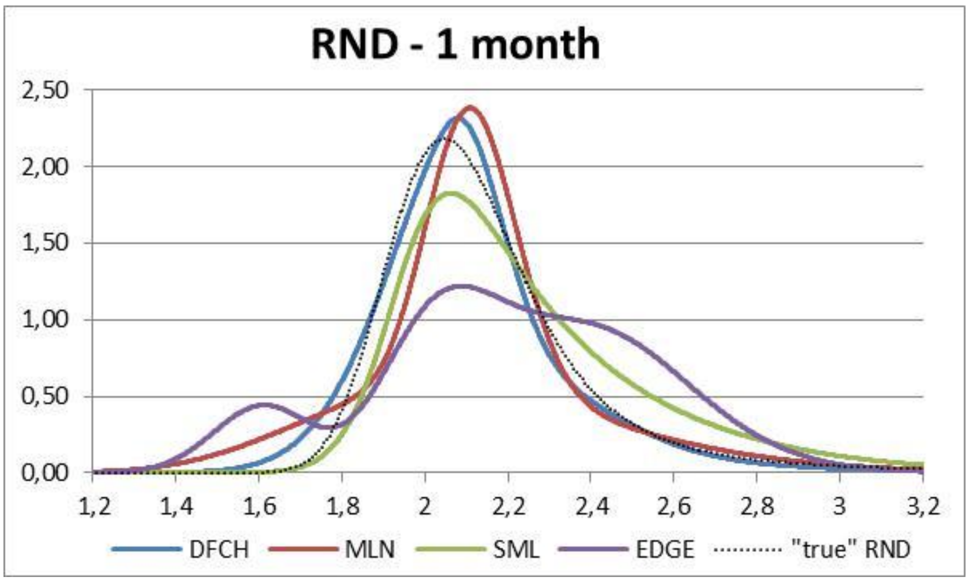

Figure 4: One month "true" RND vs RND estimations - 29th October 2010.

5.4. The Subjective Probability Density Function

An additional analysis about the usefulness of the RNDs as predictive tools was given. For this purpose, the non-structural methods were also used to extract the RNDs directly from observed option prices[3]. The goodness of fit of these RNDs to the observed returns of the exchange rate USD/BRL was then assessed. It should be noted that these measures are risk neutral, so there is a natural inclination to consider its usage inappropriate. However, option prices are also influenced by investor's expectations on the future asset prices, so they also incorporate information about the subjective probabilities that investors attribute to the future realizations. This way, under certain assumptions, we can transform the RND into a Subjective Probability Density Function (SPDF) and compare this estimation with the

[3] These RNDs are different from the RNDs in the previous sections, which were obtained from the "true" option prices (generated by the CGMY Gamma-OU model).

observed values of the exchange rate USD/BRL. The conversion of the RND into the SPDF was made following the approach of Bliss and Panigirtzoglou [2004]: a rational investor can predict, on average, the future realized returns of the underlying asset, meaning that the difference between the RND and the future outcomes is given by the degree of the investor risk aversion. Therefore, as long as we know investors' risk preferences, the SPDF can be estimated through the relation provided by Ait-Sahalia [2000]:

$$\frac{SPDF}{RND} = \lambda \frac{U'(S_T)}{U'(S_t)} = \zeta(S_T, S_t), \tag{32}$$

$$SPDF = \frac{\dfrac{RND}{\zeta(S_T, S_t)}}{\int \dfrac{RND}{\zeta(x, S_t)} dx} = \frac{\dfrac{U'(S_t)}{\lambda U'(S_T)} RND}{\int \dfrac{U'(S_t)}{\lambda U'(x)} RND dx} = \frac{\dfrac{RND}{U'(S_T)}}{\int \dfrac{RND}{U'(x)} dx}, \tag{33}$$

where λ is constant and $\zeta(S_T, S_t)$ is the pricing kernel. The SPDF must be normalized to integrate to 1. For the utility function $U(S_t)$, we assumed a parametric form based on the exponential utility function:

$$U(S_t) = -\frac{e^{-\gamma S_t}}{\gamma}. \tag{34}$$

The SPDF was obtained by searching the parameter γ that minimizes the difference between the SPDF and the realized outcomes. The statistical distance was measured by the statistical test proposed by Berkowitz [2001], an appropriate method that is used to compare density estimations with just one realization in each point of time.

We executed the optimization procedure for the whole period and, additionally, for the subsamples shown in Tables 1 and 2 (these results were obtained for the one month term). The results in Table 2 show the best fit of the SPDF obtained through the DFCH method. In fact, in contrast to the other methods, the SPDFs estimated from the DFCH have p-values that are significantly above 5% for the considered sub-periods and also for the whole sample. It should be noted that the MLN, SML and EE methods perform poorly in the period immediately preceding the subprime crisis (rumors period). This analysis was also made for the longer terms (fit of the 3/6 months SPDF to the observed USD/BRL after 3/6 months term). The p-values for the longer terms were statistically insignificant (very close to zero), which reveals that the SPDF is a biased estimator for the longer terms.

Risk Neutral Density Functions FIT					
Periods	**Período**	**DFCH**	**MLN**	**Spline**	**Edge**
normal period:	Jun2006•Jul2007	43,42%	42,94%	6,61%	3,04%
rumours period:	Aug2007•Aug2008	70,77%	1,51%	5,71%	3,87%
peak period:	Sep2008•Feb2010	49,86%	73,55%	12,27%	1,43%
All sample	Jun2006•Aug2013	26,40%	5,19%	0,71%	3,73%

Table 1: p-values obtained by the RND for the considered calibration samples.

Subjective Probability density Functions FIT					
Periods	**Período**	**DFCH**	**MLN**	**Spline**	**Edge**
normal period	Jun2006•Jul2007	89,70%	79,39%	65,13%	83,94%
rumours period	Aug2007•Aug2008	70,85%	1,58%	7,02%	5,98%
peak period	Sep2008•Feb2010	52,81%	78,12%	44,21%	6,12%
All sample	Jun2006•Aug2013	31,85%	5,94%	8,09%	18,40%

Table 2: p-values obtained by the SPDF for the considered calibration samples.

The results suggest that the DFCH method, besides being able to capture the “true” RND, is also flexible enough to adjust the future observed prices of the currency USD/BRL for the one month term. This capacity was not revealed by the MLN, a method which together with the DFCH produced the unbiased estimators of the “true” RND.

6. Conclusion

In this work, the analysis in Cooper [1999] and Santos and Guerra [2014] was deepened and the performance of the non-structural models in estimating the “true” RNDs was measured through a process that generates “true” RNDs that are closer to reality, due to the use of the CGMY Gamma-OU model, which is a stochastic model with jumps in asset returns and stochastic volatility. We also considered more scenarios by generating a “true” RND that incorporates the characteristics of the USD/BRL currency options for each month between October 2006 and February 2010, a period characterized by important market events.

According to the RMISE criterion, the DFCH model has a higher overall performance as the “true” RND estimator, especially during the first four months of the “peak period”, between September 2008 and December 2008. However, we noticed that the MLN model was slightly better than the DFCH model in capturing the shorter terms “true” RNDs during the “normal period”, before the subprime crisis became apparent in August 2007. In the stability analysis, the SML model showed the best results, having the lowest RIV for almost all pairs of dates and maturities. The EE model estimated the most biased RNDs.

Additionally, we assessed the goodness of fit of the SPDFs (extracted from observed option prices) to the observed returns of the USD/BRL. In this regard, the DFCH showed the best fit to the realized returns of the USD/BRL for the one month term. None of the non-structural methods was able to adjust the observed returns for the three and six months terms. This shows that the SPDF may be a useful risk management tool, but further research is needed to demonstrate the predictive power of these densities in out-of-sample data and across different types of assets.

Acknowledgements

João Guerra was financially supported by FCT – Fundação para a Ciência e Tecnologia through the strategic project PEst-OE/EGE/UI0491/2013.

Appendix A

	RMISE								
	1 month to Maturity			3 months to Maturity			6 months to Maturity		
	Mean	**1st quartile**	**3rd quartile**	**Mean**	**1st quartile**	**3rd quartile**	**Mean**	**1st quartile**	**3rd quartile**
DFCH	0,25683	0,14029	0,28655	0,18710	0,10746	0,25851	0,18996	0,12264	0,21165
MLN	0,27854	0,15134	0,34650	0,19589	0,10875	0,26544	0,19685	0,11157	0,24024
SPLINE	0,44684	0,29727	0,46346	0,33388	0,26983	0,39032	0,33900	0,26183	0,39284
EDGE	0,87609	0,56928	0,99845	0,55168	0,44105	0,65977	0,44517	0,33245	0,50581

Table A.1: Quartiles of the observed RMISE values between June 2006 and February 2010.

	RISB								
	1 month to Maturity			3 months to Maturity			6 months to Maturity		
	Mean	**1st quartile**	**3rd quartile**	**Mean**	**1st quartile**	**3rd quartile**	**Mean**	**1st quartile**	**3rd quartile**
DFCH	0,21687	0,12198	0,22542	0,18364	0,10542	0,23403	0,18862	0,12233	0,21158
MLN	0,24555	0,13926	0,29740	0,18417	0,08982	0,25018	0,19369	0,11040	0,23577
SPLINE	0,44079	0,29408	0,45523	0,33156	0,26809	0,38650	0,33808	0,26106	0,39100
EDGE	0,79547	0,50882	0,89182	0,51498	0,40505	0,61890	0,42482	0,32231	0,47687

Table A.2: Quartiles of the observed RISB values between June 2006 and February 2010.

	RIV								
	1 month to Maturity			3 months to Maturity			6 months to Maturity		
	Mean	**1st quartile**	**3rd quartile**	**Mean**	**1st quartile**	**3rd quartile**	**Mean**	**1st quartile**	**3rd quartile**
DFCH	0,11593	0,06291	0,14496	0,02634	0,01097	0,02910	0,01500	0,00712	0,01723
MLN	0,11573	0,04140	0,12414	0,04963	0,02126	0,05669	0,02755	0,00873	0,03482
SPLINE	0,02488	0,01423	0,03140	0,01043	0,00669	0,01395	0,00612	0,00421	0,00833
EDGE	0,34714	0,22601	0,42875	0,18678	0,14645	0,22832	0,11984	0,07930	0,16116

Table A.3: Quartiles of the observed RIV values between June 2006 and February 2010.

References

Abadir, K. M., and Rockinger, M. (2003). Density functionals, with an option-pricing application. Econometric Theory, 19, 778–811.

Ait-Sahalia, Y., and Lo, A. W. (2000). Nonparametric risk management and implied risk aversion, Journal of Econometrics, 94, 9–51.

Bahra, B. (1997). Implied risk-neutral probability density functions from option prices: Theory and application. Working Paper, Bank of England.

Barndorff, O. E., and Shephard, N. (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. Journal of the Royal Statistical Society B, 63, 167–241.

Berkowitz, J. (2001). Testing Density Forecasts, With Applications to Risk Management. Journal of Business & Economic Statistics, 19, 465–474.

Bliss, R., and Panigirtzoglou, N. (2002). Testing the stability of implied probability density functions. Journal of Banking and Finance, 26, 381–422.

Bliss, R., and Panigirtzoglou, N. (2004). Option-Implied Risk Aversion Estimates. The Journal of Finance, 59, 407–446.

Breeden, D. T., and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. Journal of Business, 51, 621–51.

Bu, R., and Hadri, K. (2007). Estimating option implied risk-neutral densities using spline and hypergeometric functions. Journal of Econometrics, 10, 216–244.

Carr, P., and Madan, D. (1999). Option valuation using the fast Fourier transform. Journal of Computational Finance, 2, 61–73.

Carr, P., Geman, H., Madan, D. H., and Yor, M. (2002). The fine structure of asset returns: an empirical investigation. Journal of Business, 75, 305–332.

Carr, P., Geman, H., Madan, D. H., and Yor, M. (2003). Stochastic volatility for Lévy processes. Mathematical Finance, 13, 345–382.

Madan, D. H., Carr, P., and Chang, E. C. (1998). The variance gamma process and option pricing. European Finance Review, 2, 79–105.

Clark, P. (1973). A subordinated stochastic process model with finite variance for speculative prices. Econometrica, 41, 135-156.

Cont, R., and Tankov, P. (2004). Financial Modelling with Jump Processes. Chapman & Hall/CRC, Boca Raton.

Cooper, N. (1999). Testing technics for estimating implied RNDs from the prices of European and American optins. Working Paper, Bank of England.

Corrado, C., and Su, T. (1996). S&P500 index option tests of Jarrow and Rudd's approximate option valuation formula. Journal of Futures Markets, 6, 611–629.

Geman, H., Madan, D. H., and Yor, M. (2001). Time Changes for Lévy Processes. Mathematical Finance, 11, 79–96.

Jarrow, R., and Rudd, A. (1982). Approximate option valuation for arbitrary stochastic processes. Journal of Financial Economics, 10, 347–369.

Jondeau, E., Poon, S. H., and Rockinger, M. (2007). Financial Modeling Under Non-Gaussian Distributions. Springer-Verlag, London.

Melick, W. R., and 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.

Santos, A., and Guerra, J. (2014). Implied risk neutral probability density functions from option prices: hypergeometric, spline, lognormal and edgeworth functions. Journal of Future Markets (forthcoming 2014). doi: 10.1002/fut.21668.

Sato, K. (1999). Lévy Processes and Infinitely Divisible Distributions. Cambridge Studies in Advanced Mathematics 68, Cambridge University Press, Cambridge.

Schoutens, W. (2003). Lévy Processes in Finance: Pricing Financial Derivatives. John Wiley & Sons Ltd.

Shimko, D. C. (1993). Bounds of probability, Risk, 6, 33–37.