**Extracting expectations from currency option prices:
a comparison of methods**

Marian Micu*

15 September 2004

Abstract. This paper compares the goodness-of-fit and the stability of six methods used to extract risk-neutral probability density functions from currency option prices. We first compare five existing methods commonly employed to recover risk-neutral density functions from option prices. Specifically, we compare the methods introduced by Shimko (1993), Madan and Milne (1994), Malz (1996), Melick and Thomas (1997) and Bliss and Panigirtzoglou (2002). In addition, we propose a new method, namely the piecewise cubic Hermite interpolation of the implied volatility function. We use data on 12 emerging market currencies against the US dollar and find that the piecewise cubic Hermite interpolation method is by far the method with the best accuracy in fitting observed option prices. We also find that there is a relative tradeoff between the goodness-of-fit and the stability of the methods. Thus, methods which have a better accuracy in fitting observed option prices appear to be more sensitive to option pricing errors, while the most stable methods have a fairly disappointing fitting. However, for the first two PDF moments as well as the quartiles of the risk-neutral distributions we find that the estimates do not differ significantly across methods. This suggests that there is a large scope for selection between these methods without essentially sacrificing the accuracy of the analysis. Nonetheless, depending on the particular use of these PDFs, some methods may be more suitable than others.

JEL Classifications: C52, F31, G13.

1. Introduction

This paper contributes to a new and growing literature on option-based approaches to modelling exchange rate expectations. Interest in this topic has burgeoned in recent years, driven by the growth of derivative markets and a greater appreciation of the information imbedded in the prices of options. The key tool used to extract the information contained in exchange rates is the risk-neutral probability density function (PDF). Risk-neutral PDFs provide the probabilities attached by a risk-neutral agent to particular outcomes for future values of exchange rates.

This paper compares several methods of estimating risk-neutral probability density functions with the aim of determining which method fits more accurately observed market option prices.[1] Specifically, it presents the methods introduced by Shimko (1993), Madan and Milne (1994), Malz (1996), Melick and Thomas (1997) and Bliss and Panigirtzoglou

*Monetary and Economic Department, Bank for International Settlements, Tel: +4161 280 92 79, Fax: +4161 280 91 00, E-mail: Marian.Micu@bis.org. The views expressed in this paper are those of the author and not necessarily the views of the BIS.

[1] Most of the previous studies on the extraction of risk-neutral density functions have concentrated on estimation issues whereas less emphasis has been placed on the formal assessment of the goodness-of-fit of such estimates. The most common instruments used to test the accuracy of fit are the pricing errors, computed as difference between observed and theoretical option prices. These pricing errors are then averaged to compute aggregate indicators, such as the mean squared error (MSE), and the mean squared percentage error (MSPE).

(2001). In addition, it proposes a new technique for recovering risk-neutral density functions based on the piecewise cubic Hermite interpolation of the implied volatility function. The comparison of these techniques is made on the basis of summary statistics derived from the estimated PDFs, using a sample of five years of daily data for 12 emerging market currencies against the US dollar. The summary statistics refer to measures of location, dispersion, asymmetry, fat-tailness and various tail percentiles. The average values of these statistics are then compared across methods. Most of the authors found similar values for the first two moments across various models, and large discrepancies in the higher moments of the distributions, which seems to suggest that the measures of skewness and kurtosis are highly model-dependent. Our empirical tests confirm most of the results from previous studies. However, we find that the quartiles of the distributions were quite similar across methods, except for Malz's (1996) approach. Moreover, the estimates of the PDF quartiles appear to be less sensitive to option pricing errors.

Before giving a detailed description of the methods used to extract risk-neutral probabilities from the prices of foreign exchange options, we present the advantages and limitations of risk-neutral density functions with respect to traditional time series methods.

2. Time series methods versus risk-neutral probability density functions

Exchange rate prices depend on many factors. Empirical models, based on both macroeconomic equilibria and market microstructure, fail to explain the changes in these prices. The failure of these models is due to the difficulty of modelling two basic factors: the changing nature of investors' expectations and the change in their risk attitudes.

The analysis of investors' expectations is usually accomplished with time series methods. However, such models have a couple of drawbacks. One of these is that they are by nature backward-looking. This means that they assume the whole set of information about future exchange rates is fully included in their historical prices. Another inconvenience is that, in most of these models, the parameters do not change with the arrival of new information, ie the parameters are not time-varying. Furthermore, the change in exchange rate regimes is usually not easily captured.[2] A third limitation of time series models is that they only offer some information about a (central) point estimate of future exchange rate prices or volatilities but not about the whole future distribution.

The forecasts based on risk-neutral distributions give the probability value of a set of future possible exchange rate prices. These probabilities are "risk-neutral", which means that investors are assumed not to charge an additional premium for a change in the uncertainty about future exchange rates.[3] They also mean that all investors have similar risk preferences. However, in the real world market participants are not risk-neutral. The effect of risk on investors' decision-making process depends on their wealth, their utility

[2] Some regime-switching models have been proposed to deal with this issue (see Engle and Hamilton (1990)). However, these models are in general over-parameterised, which make them difficult to be empirically implemented. Moreover, they only have a limited number of possible regimes. This may turn out to be problematic if the exchange rate goes though a set of multiple intermediary regimes, which is in fact the case with most crawling peg arrangements.

[3] Technically, the fact that dynamic hedging is possible creates the possibility of the perfect replication of an option's payoff, which determines the cancelling out of its writer's risk exposure. If this is true, the assumption of risk-neutrality is a natural one. However, hedging incurs costs and is usually not feasible at very high frequencies, at least in less liquid foreign exchange markets.

function and the perceived level of risk in the market (see Pratt (1964)). Moreover, investors have heterogeneous risk preferences and therefore risk-neutral distributions might not perfectly match the "real" distributions used by market participants to price foreign exchange options. Some recent papers investigate the differences between the risk-neutral and subjective distributions.[4]

Various techniques have been proposed to extract risk-neutral PDFs from option prices. These techniques can be classified in four broad categories (Bahra (1996)): I) the probability density function is estimated by assuming a particular stochastic process for the underlying asset (eg Bates (1991), Malz (1996)); II) a functional form for the PDF is assumed and its parameters are estimated by minimizing the difference between actual and predicted option prices (eg Rubinstein (1994), Melick and Thomas (1997), Bahra (1996)); III) the probability density function is implied from some parametric specification of the call pricing function or the implied volatility smile (see eg Shimko (1993), Madan and Milne (1994), Malz (1997)); IV) the risk-neutral density function is estimated non-parametrically (see eg Aït-Sahalia and Lo (1998), Bondarenko (2003), etc.). However, the implied PDFs by these methods are not always similar. Moreover, there is no consensus on which of these techniques provides a better fit.

3. The Garman-Kohlhagen model

Before pricing any derivative security, we need to be able to characterise the behaviour of its underlying asset. Usually, the price changes in the underlying asset are described by making some assumptions about their future distribution. One of these assumptions traces its roots to Bachelier (1900). Bachelier assumed that bond prices follow an *arithmetic Brownian motion.* He used this assumption to model options on French government bonds. However, such a stochastic process can also take negative values. This led to anomalous approximations of the prices of bonds with long maturities. An alternative to this hypothesis was proposed by Samuelson (1965). He considered long-term equity options, and used the *geometric Brownian motion* to model the random behaviour of the underlying stock price. Based upon this, he modelled the random value of the option at exercise. The model required two assumed parameters. The first was the expected rate of return α for the stock price. The second was the rate β at which the option's value at exercise should be discounted back to the pricing date. These two factors depended upon the unique risk characteristics of, respectively, the underlying stock and the option. Neither factor was observable in the market. Depending upon their degree of risk aversion, different investors might propose different values for the factors. Accordingly, Samuelson's formula was largely arbitrary. It offered no means for a buyer or a seller with different risk aversions to agree on a price for an option.

Black and Scholes (1973) proposed a completely new approach. They considered an options trader who is about to sell an option. The trader is assumed to *dynamically hedge* the exposure until the option expires.[5] Therefore, the trader needs to evaluate the cost of

[4] See for example Ait-Sahalia and Lo (2000), Bliss and Panigirtzoglou (2004), Jackwerth (2000), Rosenberg and Engle (2002). This area of research looks very promising for the re-evaluation of the role the risk-premium hypothesis plays in explaining the forward premium puzzle.

[5] Dynamic hedging is a procedure for hedging an option position by periodically changing the exposure in the underlying asset, so that the change in the price of the option is fully compensated by the change in the price of the underlying asset.

dynamically hedging the short option. Given certain simplifying assumptions, they found that this cost could be known in advance. One of the basic assumptions in their model is that the asset price is part of a continuous process.[6] More precisely, the price of the underlying asset is assumed to follow a random walk in continuous time, with the variance proportional to the square of the underlying asset price change. Thus, the distribution of the underlying price at the end of any finite interval is lognormal and the variance of the return on the underlying asset is constant. Another hypothesis of their model is that the risk-free interest rate is known and constant through time, which means that if an investor holds a risk-free security for τ years, the value of every unit of this security will grow to $\exp(r\tau)$, where r is the continuously compound risk-free constant interest rate. Moreover, Black and Scholes assumed that there are no arbitrage opportunities, every investor can trade continuously with no transaction costs and can borrow or lend any fraction of a security at the risk-free interest rate. A final assumption was that short-selling was allowed.[7]

Black and Scholes (1973) started with the observation that if the frequency of re-hedging increases, dynamic hedging becomes more predictable. Using stochastic calculus, they took the limiting case as the frequency of re-hedging approaches infinity. In that limiting case, the cost of dynamic hedging is independent of the actual path taken by the price of the underlying asset. It depends only upon the price's volatility. If that volatility is constant and known in advance, the cost of dynamic hedging an option is certain. Being certain, it entails no risk, so it can be discounted at a risk-free rate to obtain the price of the option.

Based upon this approach, Black and Scholes derived a partial differential equation for valuing claims contingent on a traded stock. The equation is, however, general. By applying different boundary conditions, it can be solved to price any such contingent claim.

Garman and Kohlhagen (1983) adapted the Black and Scholes (1973) model for the pricing of European-type currency options. They used similar assumptions as Black and Scholes (1973), and considered that the spot exchange rate S_t follows a geometric Brownian motion, ie the differential representation of spot price movements is

$$dS_t = \mu S_t dt + \sigma S_t dW_t \tag{1}$$

where μ and σ represent the instantaneous mean and volatility and W_t is a standard Wiener process. Additionally, they considered that markets are frictionless and interest rates, both in foreign and domestic markets, are constant.

Under these assumptions they showed that in a risk-neutral world the process S_t can be written as:

$$dS_t = (r - r^*)\, S_t dt + \sigma S_t dW_t \tag{2}$$

where r and r^* are the domestic and foreign continuously compounded risk-free interest rates. Using Itô's lemma, this can be further developed as:

[6] This means that asset prices do not exhibit jumps, which in fact is not the case in many markets, including the foreign exchange market.

[7] Short-selling means selling some securities the seller does not own. Specifically, a seller who does not own a security accepts the price of the security from a buyer and agrees to settle the buyer on some future date by paying him an amount equal to the price of the security on that date.

$$\ln(S_T) = \ln(S_t) + \left(r - r^* - \frac{1}{2}\sigma^2\right)dt + \sigma dW_t \tag{3}$$

where T denotes the expiration date of an option contract and S_T is the exchange rate at the maturity of the contract. The diffusion process in equation (3) can again be represented as:

$$\ln(S_T) = \ln(S_t) + \left(r - r^* - \frac{1}{2}\sigma^2\right)\tau + \sigma(W_T - W_t) \tag{4}$$

where $\tau = T - t$ is the time to maturity of an option contract, expressed in years. Knowing that $W_T - W_t$ is a random variable with mean 0 and variance $\tau = T - t$, we can express the mean of the diffusion process as $\ln(S_t) + \left(r - r^* - \sigma^2/2\right)\tau$ and its variance as $\sigma^2\tau$. Under the hypothesis of log-normality of the underlying asset, the risk-neutral density function is given by:

$$p(S_T) = \frac{1}{\sigma S_T\sqrt{2\pi\tau}}\exp\left\{-\frac{1}{2}\left[\frac{\ln(S_T) - \ln(S_t) - \left(r - r^* - \sigma^2/2\right)\tau}{\sigma\sqrt{\tau}}\right]^2\right\} \tag{5}$$

Garman and Kohlhagen (1983), following the methodology outlined by Black and Scholes (1973) and Merton (1973) derived the following formulae for European-type foreign exchange option prices:

$$C = e^{-r^*\tau}S_tN(d_1) - e^{-r\tau}KN(d_2) \tag{6}$$

$$P = Ke^{-r\tau}N(-d_2) - S_te^{-r^*\tau}N(-d_1) \tag{7}$$

$$d_1 = \frac{\ln(S_t/K) + \left(r - r^* + \sigma^2/2\right)\tau}{\sigma\sqrt{\tau}} \tag{8}$$

$$d_2 = \frac{\ln(S_t/K) + \left(r - r^* - \sigma^2/2\right)\tau}{\sigma\sqrt{\tau}} \tag{9}$$

where C is the price of a call option, P is the price of a put option, K is the exercise price and $N(x)$ is the cumulative normal distribution function for a standardised normal random variable given by:

$$N(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-\frac{1}{2}z^2}dz, \quad z \sim N(0,1) \tag{10}$$

Option prices calculated with the Garman-Kohlhagen formulae depend on six elements: (1) the spot price of the underlying asset S_t, (2) the price at which the currency can be bought K, (3) the time at which the option expires T, (4) the risk-free interest rate in domestic currency r, (5) the risk-free interest rate in foreign currency r^*, and (6) the volatility of the exchange rate σ. The only unknown parameter in the Black and Scholes model is the volatility of the underlying asset. This is usually not a constant parameter, nor is there even general agreement on the best procedure for estimating it.

On the contrary, actual volatility, and also the market's volatility estimate, appear to vary randomly over a wide range of methods.[8]

4. The volatility smile: straddles, strangles and risk reversals

The volatility derived from quoted option prices is called the *implied volatility*.[9] The Garman-Kohlhagen model implies that all options on the same underlying currency have identical implied volatilities, regardless of time to maturity and moneyness. However, there are well known stylised facts from the practice of options trading showing that the Garman-Kohlhagen formula does not hold. One of these is that options with the same exercise price but different tenures often have different implied volatilities, giving rise to a term structure of implied volatility. This indicates that market participants expect the implied volatility of short-term options to change over time. A rising term structure suggests that market participants expect short-term implied volatility to rise or a willingness to pay more to protect against longer-term exchange rate volatility (see Campa and Chang (1995)). Another well-known stylised fact is that out-of-the money options on currencies with flexible exchange rates often have higher implied volatilities than at-the-money options, indicating that the market perceives the distribution of exchange rate changes to be leptokurtic, that is, the likelihood of large exchange rate moves is greater than its theoretical value obtained from the lognormal distribution.[10] This is often referred to as the "volatility smile".

An interesting feature of the implied volatility is related to the asymmetry between the volatilities of out-of-the-money put and call options with the same deltas. In general, out-of-the-money call options have implied volatilities which differ from those of equally out-of-the money put options, suggesting that the market perceives the distribution of the future exchange rate to be skewed.[11] This is very common in the stock markets and it is called the "volatility smirk". In foreign exchange markets, a persistent volatility smirk might suggest a "peso problem".

In the foreign exchange over-the-counter market options are frequently sold in combinations. The most common in the interbank currency option markets is the *straddle*. This is a combination of an at-the-money forward call and an at-the-money forward put with the same time to maturity. This combination has a V-shaped payoff function: the further the spot price is away from the strike price at the time of expiry, the higher is the payoff. The price of a straddle conveys information about the expected variance of the exchange rate: the higher the variance is expected to be, the higher is the profit expected from holding a straddle, and as a result, the higher its price. The standard quoted over-the-counter straddle contract is the *at-the-money straddle*, where the strike price for both options is equal to the current forward price. This also means that the delta for

[8] See Poon and Granger (2003) for an extensive survey on the differences in the forecasting performance of different methods.

[9] In organised exchange markets where currency option prices are known, the implied volatility can be easily deduced from observed option prices by solving for the parameter σ in the Garman-Kohlhagen model. In the foreign exchange over-the-counter market, the volatility is directly quoted by traders.

[10] Such patterns may be generated if implied volatility is stochastic or if the exchange rate follows a jump-diffusion process (see Heynen (1994) or Taylor and Xu (1994, 1995)).

[11] Such patterns are generated if the exchange rate follows an asymmetric jump-diffusion (eg Bates (1991) or Bates (1996a, 1996b)).

both options that make a standard straddle is approximately equal to 0.5. Straddles are quoted in terms of volatility, therefore, the quoted price of a straddle gives us the implied volatility of an option with an approximate delta of 0.5. The straddle price therefore is often referred to as the at-the-money implied volatility.

There is also active trading in option combinations with which participants can take directional ("informed") positions in the foreign exchange options market. The most representative are the *strangles* and the *risk reversals*, both consisting of out-of-the money put and call options. Thus, a strangle consists of a purchase or sale of an out-of-the-money put and call option on the same underlying asset, with the same expiration date. To buy a strangle is similar to "buying the volatility" because this strategy leads to profits only if there is a drastic move in the price of the underlying asset, outside the two strike prices for call and put options.

Dealers usually quote strangles by stating the implied volatility at which they want to buy or sell the options and record strangle prices as the spread of the strangle volatility over the at-the-money forward volatility. If market participants were convinced that exchange rates move lognormally, the out-of-the-money options would have the same implied volatility as at-the-money options and strangle prices would be zero. Strangles, then, indicate the degree of curvature of the volatility smile. The most popular strangle contracts traded on the foreign exchange over-the-counter market correspond to levels of the delta of 25% and 10%. The midpoint of the 25% strangle price can be expressed as:

$$str_t^{25\Delta} = 0.5\left[\sigma_t^{(\Delta_{0.25}^c)} + \sigma_t^{(\Delta_{0.25}^p)}\right] - \sigma_t^{ATM} \tag{11}$$

where $\sigma_t^{(\Delta_{0.25}^c)}$ and $\sigma_t^{(\Delta_{0.25}^p)}$ are the implied volatilities of 25% delta call and 25% delta put options, and σ_t^{ATM} is the at-the-money forward implied volatility. The higher the dispersion of implied volatilities for out-of-the-money call and put options, the higher the price of the strangle.[12]

A *risk reversal* consists of a simultaneous purchase of an out-of-the-money call option and sale of an equally out-of-the-money put option on the same underlying asset. Since prices of over-the-counter options are quoted in terms of implied volatility, and these options are equally out-of-the money, the price of a risk reversal is by convention quoted as the volatility of the long option minus the volatility of the short option. Risk reversals are a measure of the relative value of options with strikes above or below the current at-the-money forward rate. They consequently express the skew that may exist in the volatility smile. The price of a risk reversal reflects market participants' view of the balance of risk towards an appreciation (if the difference is positive) or a depreciation (if the difference is negative) of the domestic currency. If the directional view is correct, the strategy makes a profit. However, if the directional view is incorrect, there is unlimited risk in the short open position. The degree of the "out-of-the-moneyness" of risk reversals is expressed in terms of delta. Thus, a purchase of a 25% delta call accompanied by a sale of a 25% delta put would be referred to as a 25% delta risk reversal. Usually, in the FX over-the-counter market, risk reversals are traded for levels of delta equal to 25% and 10%. In practice, the market is more liquid for 25% deltas, and to a lesser degree for 10% deltas. As a rough guideline, 25% delta call and 25% delta put options have approximately a 25% probability of finishing in-the-money. In comparison, 10% delta

[12] For a more analytical description of this instrument see Malz (1996, 1997).

options have approximately a 10% probability of finishing in-the-money. The 1-month 25% delta risk reversal is the reference risk reversal for most interbank market makers.

The price of a risk reversal is given in terms of the difference between call and put volatilities. Thus, the formula for a 25% delta risk reversal is

$$rr_t^{25\Delta} = \sigma_t^{(\Delta^c_{0.25})} - \sigma_t^{(\Delta^p_{0.25})} \tag{12}$$

A plausible model for explaining the skewness in the distribution of exchange rates, and therefore the pricing of risk reversal spreads, is to assume that changes in volatility are correlated with changes in exchange rates. For example, suppose that the Mexican peso against the US dollar exchange rate is correlated with its volatility. Then, if the peso depreciates, volatility is more likely to go up, and if the peso appreciates, volatility is more likely to go down. Thus, if market participants expect a peso depreciation, the out-of-the-money put options are sold for a higher implied volatility than out-of-the-money call options. This would lead to a skewed distribution. Thus, if the market is postured defensively against significant exchange rate moves in one direction, the risk reversal prices are non-zero.[13]

In the foreign exchange over-the-counter market, the implied volatility for out-of-the-money options can be easily inferred from the prices of straddles, strangles and risk reversals. Thus, the volatility of the 50% delta call and put options is equal to the price of the straddle. The volatility of the 25% delta call option is

$$\sigma_t^{\Delta^c_{0.25}} = \sigma_t^{ATM} + str_t^{25\Delta} + 0.5 rr_t^{25\Delta} \tag{13}$$

and the corresponding volatility of the 25% delta put option is

$$\sigma_t^{\Delta^p_{0.25}} = \sigma_t^{ATM} + str_t^{25\Delta} - 0.5 rr_t^{25\Delta} \tag{14}$$

To close a straddle deal, the counterparties use the Garman-Kohlhagen formulae to translate the straddle price expressed in volatilities into currency units. In the same way, to close a strangle or a risk reversal deal, the exercise prices of the individual components must first be set, which in part requires the counterparties to agree on $\sigma_t^{\Delta^c_{0.25}}$ and $\sigma_t^{\Delta^p_{0.25}}$, the implied volatilities of the 25% delta call and the 25% delta put. The counterparties then translate the price expressed in volatilities into currency units by using equations (6) and (7).

5. Comparison of methods used to infer risk-neutral density functions from option prices

The existence of a multitude of methods to extract risk-neutral distributions from option prices raises the question of which method is better. Two criteria are usually used to compare these methods: the goodness-of-fit and the stability.

The goodness-of-fit means how well theoretical option prices calculated with each of these methods fit observed market prices. Practically, synthetic measures for the pricing

[13] Risk reversal prices of zero do not indicate that no skew is perceived in exchange rates, since a modest skew is implied by lognormality. For a more technical analysis of this instrument, see McCauley and Melick (1996a, 1996b).

errors are used: mean squared error, mean absolute error, mean squared percentage error, mean absolute percentage error, etc. The best method is considered to be the one which has the smallest pricing errors.

The stability of implied PDFs refers to their robustness to small changes in option prices. In general, the stability is tested by using bootstrapping or Monte-Carlo methods to perturb the option prices or the parameters of estimated distributions.

One of the earliest papers which compared the goodness-of-fit of risk-neutral PDFs extracted from option prices was the paper by Sherrick, Garcia and Tirupattur (1996). They compared two PDF approximating function approaches, using double lognormal and Burr III specifications. They performed a series of tests to compare the two functional forms including: evaluation of average pricing errors, comparisons of implied PDF mean prices with future prices, and assessment of the predictability of resulting futures price variability. In the first two cases, they found evidence of marginal improvement in the Burr III method. Moreover, their estimates showed that the Burr III method substantially outperformed the mixture of lognormals in depicting future price variability.

Campa, Chang and Reider (1998) compared the performance of three methods used to extract exchange rate expectations from option prices: cubic splines, an implied binomial tree and a mixture of lognormals, and showed that, despite their methodological differences, the three approaches led to similar probability density functions, clearly distinct from the lognormal benchmark, and typically characterised by skewness and leptokurtosis. They also documented a striking positive correlation between the skewness and the spot rate, which they interpreted as a rejection of the hypothesis that innovations in exchange rates are independent of the position of spot exchange rates within the ERM band.

In a similar paper, Jondeau and Rockinger (2000) used over-the-counter options data for the FRF/DEM exchange rate for various dates between May 1996 and June 1997, and considered a mixture of lognormal distributions, a Hermite polynomial approximation, Malz's (1996) jump-diffusion model, Heston's (1993) stochastic volatility model and a maximum entropy model based on the Edgeworth expansion of Jarrow and Rudd (1982). First, they empirically tested the lognormal assumption. The hypothesis of normality of exchange rate returns was clearly rejected by the presence of concave implied volatility functions with respect to the strike price and of a non-linear term structure of implied volatilities. These features confirmed that more complicated models than Garman-Kohlhagen should be considered. Their estimates from the jump diffusion model showed an upward sloped term structure of implied volatilities, which indicated that investors were more uncertain about price movements in the long run. Their estimated probability of a jump before maturity varied from 0.0399 to 0.0699, suggesting that, for that period, the likelihood of an exchange rate jump was rather small.[14]

Another interesting finding of their study was that the expected jump size decreases with the option's tenure. Recent empirical studies by Andersen and Andreasen (1999) and Carr and Wu (2003) confirmed this pattern and suggested that pure jump models seem to be more appropriate to explain the behaviour of short term option prices, whereas a diffusion price coupled with a diffusion in volatility are needed to describe the behaviour of long term options.

Jondeau and Rockinger's estimates of Heston's (1993) stochastic volatility model confirmed their results obtained with the geometric jump-diffusion model. In this case, their

[14] This might also be a feature of the jump-diffusion model, given that it has been confirmed by other similar studies using other currency pairs and time periods (see, for instance, Bates (1996a) and Malz (1996)).

results showed unambiguously that on a normal day the term structure of implied volatilities was upward sloped, while a flat or slightly decreasing pattern was observed on a day with agitated markets. This was consistent with the mean reversion of volatility.[15]

They also compared the estimated moments of risk-neutral distributions and found that the first two moments, the mean and the variance, were quite homogeneous across methods. However, the volatility implied by the lognormal model appeared to be systematically smaller than the one obtained with the other approaches. Their estimates of skewness and kurtosis were much more contrasted. The skewness obtained from Hermite polynomial and Edgeworth expansion models was systematically lower than the one obtained with other models, although the difference appeared to be small. Nonetheless, concerning kurtosis, they noticed pronounced differences between models, with the mixture of lognormals and stochastic volatility approaches giving generally very large estimates of excess kurtosis.

One of the papers which focused explicitly on the goodness-of-fit of implied distributions was the paper by McManus (1999). He used options on Eurodollar futures traded at the Chicago Mercantile Exchange (CME) and applied the Black (1976), the mixture of lognormals, the jump-diffusion, the Hermite polynomials and the maximum entropy methods to estimate risk-neutral probability density functions. He found that the double lognormal method ranked first in terms of pricing errors. However, the mixture of lognormals method appeared to be slow to converge, especially if the true risk-neutral PDF was close to being lognormal. The Hermite polynomial method ranked second yielding similar results to the mixture of lognormals, and was quicker to converge. The only drawback with the Hermite polynomials method was that the estimation of the risk-neutral PDF occasionally yielded negative probability values.[16] He also found that the higher moments, skewness and kurtosis, vary widely across methods.

Most of the empirical papers tested the properties of the methods used to extract risk-neutral distributions by using observed options data. However, Cooper (1999) compared two approaches to estimating implied risk-neutral probability density functions from the prices of European-style options by using simulated artificial options price data. He examined the mixture of lognormals and a method based on the cubic spline interpolation of implied volatility and tested the ability of these techniques to recover the risk-neutral distributions. Cooper tested both the accuracy and stability of the estimated PDFs. He also performed a Monte Carlo test to evaluate the ability of each of these techniques to match simulated distributions based on Heston's (1993) stochastic volatility model. He found that the second method performed better in fitting the first two moments of the risk-neutral PDFs. Moreover, the cubic spline method was systematically more stable than the mixture of lognormals. His empirical results showed that the higher moments of the distribution appeared to be much more difficult to estimate accurately with both techniques, often resulting in estimates that are quite far from the simulated ones.

Weinberg (2001) examined how useful is the information contained in option prices for predicting future price movements of the underlying assets. He developed a semi-parametric cubic spline methodology for estimating risk-neutral PDFs and applied this technique to options on the dollar-yen, dollar-mark, and S&P 500 index futures. He compared the second and third moments of the risk-neutral distribution with the corre-

[15] However, their empirical estimates of the parameter which describes the mean-reversion showed high variability, which illustrates the difficulties their model has in pinning down the mean-reversion pattern.

[16] These negative probability values can occur because the fourth and sixth order Hermite polynomial approxinations involve truncating an infinite polynomial series.

sponding moments of daily returns and found that implied volatility predicts reasonably well future realised volatility. Interestingly, he found that at-the-money volatility outperformed the volatility implied by the risk-neutral PDF in predicting the realised volatility. However, the PDF-implied skewness measure failed to predict the realised skewness.[17]

More recently, Bliss and Panigirtzoglou (2002) used option contracts written on sterling futures and the FTSE 100 index to test the stability of mixture of two lognormals and natural spline methods. They perturbed randomly option prices by no more than plus or minus one half of the quotation tick size and found that the double lognormal method was systematically less stable than the smoothed implied volatility smile method, even when the latter was calibrated to have the same goodness-of-fit. They motivated these results by the sensitivity of the mixture of lognormals method to computational problems, such as local optima, non-convergence, spikes in the estimated PDFs, etc. They also found that the confidence intervals for the higher moments of the distribution were very wide for both methods, which may suggest a high sensitivity of these statistics to option pricing errors.

Anagnou et al. (2002) also tested the ability of risk-neutral PDFs extracted from the S&P 500 index options and from options on the sterling-dollar exchange rate to predict future asset price distributions. They proposed two new approaches to estimate the risk-neutral PDFs: the generalised beta and the normal inverse Gaussian, and compared the estimates with those obtained from mixture of two lognormals and B-spline approaches. They found that the normal inverse Gaussian approach provided slightly better estimates than the generalised beta or the mixture of lognormals. Despite its high flexibility, the B-spline approach does not recover the tails of the risk-neutral distribution outside the range of available strike prices, which requires either an extrapolation or a truncation of the risk-neutral PDF. Anagnou et al. (2002) opted for the second solution and truncated and then rescaled the estimated distributions for the range of available strike prices. Their empirical tests showed that the implied risk-neutral distributions extracted from option prices do not represent an appropriate forecast of the true underlying asset's distribution at expiry, regardless of the functional form (parametric or non-parametric) chosen to model them. Not surprisingly, they found that the main reason for rejecting the predictive ability of risk-neutral distributions is given by misspecification of the tails outside the range of available strike prices. The second reason of the rejection appeared to be the bias in the variance, ie an underestimation of the underlying asset's realised variance.

6. Data

While options on currencies and currency futures of developed countries are traded both on exchanges and over-the-counter, the options written on currencies of emerging countries are almost exclusively traded over-the-counter. Over-the-counter options are European-type. They are usually traded for standardised maturities. Price quotes are expressed as implied volatilities corresponding to defined levels of options' delta, which traders by agreement substitute into the Garman-Kohlhagen formula to determine the option premium. Since the volatility is the only unobservable parameter in the Garman-Kohlhagen formula, these volatilities uniquely determine the options' prices. However,

[17] The inability of implied skewness to translate into realised price movements indicates that there were profit opportunities available from betting against the market.

this does not necessarily mean traders believe that the Garman-Kohlhagen formula gives a fair evaluation of option prices. This market convention simply allows a direct mapping from implied volatility quotes into option prices. One reason for this market convention is that if option prices were expressed as a premium for a fixed strike, most intra-day changes in the option premium would result from innovations in the spot rate, requiring much greater coordination between spot and option markets. Another possible reason is that quoting volatility, ie the perceived risk in the foreign exchange market, facilitates the task of foreign exchange risk managers in hedging their exposures.

Our data consists of market quotes of over-the-counter options on 12 emerging market currencies against the US dollar. These currencies are: the Brazilian real, Chilean peso, Czech koruna, Indonesian rupiah, Malaysian ringgit, Mexican peso, Polish zloty, South African rand and South Korean won. In particular, we use the information from straddles, strangles and risk-reversals corresponding to 50%, 25% and 10% levels of delta, as well as Eurocurrency interest rates recorded by currency option traders at a major global foreign exchange dealer bank. Observations are for options with maturities of one, three and six months, with a daily frequency from 10 November 1997 to 10 November 2002.

Although the Garman-Kohlhagen model assumes constant volatilities across exercise prices, the implied volatility quoted by option traders for our selected currencies typically varies as a function of options' strike prices. This reflects a departure from the Garman-Kohlhagen formula's assumptions, implying that traders assume a non-lognormal distribution for the future exchange rates when they price these options. In general, our data shows that the implied volatility is lowest for at-the-money options, increasing for both in- and out-of-the money options. This pattern is referred to as the "volatility smile" and is consistent with the leptokurtosis of the distribution of future exchange rate returns. In general, the probability of future exchange rate realisations is not symmetrically distributed around the at-the-money strike price. For most of the selected currencies in our sample, we noticed a greater probability of a large depreciation. This is translated into positive risk-reversal quotes and positive skewness of the risk-neutral distributions.

The prices used as inputs for estimating risk-neutral probability density functions are subject to various errors that may cause the observed prices to deviate from those we would expect in a frictionless world.[18] One of these possible errors may be due to mistakes in the recording and reporting of these prices. We visually inspect all data series and eliminate from our sample obvious typos and other data records which appear implausible.[19] Another potential data problem is the non-synchronicity of these prices, which arises from the need to use multiple simultaneous variables as inputs to the model. However, this was not a concern for our study as all prices used are collected at the same time by our data provider. Another potential concern about the data is the liquidity of the options contracts. If the daily quotes on these options are illiquid, the information content of the implied PDFs may become noisy. A graphical inspection of our data series suggests that for some short periods stale prices are present, especially for risk-reversals and strangles. This may reduce the information content of estimated PDFs.

[18] See Bliss and Panigirtzoglou (2002) for an in-depth analysis of possible data errors in option prices.
[19] This kind of data errors accounted for less than 0.05% from the total number of observations.

7. METHODS USED TO EXTRACT RISK-NEUTRAL PROBABILITY DENSITY FUNCTIONS FROM OPTION PRICES

In this study we estimate the risk-neutral distributions using six different methods: the quadratic interpolation of the implied volatility (Shimko (1993)), Hermite polynomials (Madan and Milne (1994)), geometric jump-diffusion (Malz (1996)), mixture of lognormals (Melick and Thomas (1997)), cubic spline interpolation (Bliss and Panigirtzoglou (2002)), and the piecewise cubic Hermite interpolation method. The estimation procedure follows two steps. First, we calculate the option prices from at-the-money volatilities, risk reversals and strangles using the market convention described in Section 4. Second, we recover the risk-neutral density functions using our selected estimation methods.

7.1. Shimko's method. Shimko's (1993) method consists of fitting the volatility smile with a parabola to obtain a continuous function of call prices within the range covered by the data. Shimko estimated by ordinary least squares the following regression:

$$\sigma = \beta_0 + \beta_1 K + \beta_2 K^2 + \varepsilon \tag{15}$$

where σ is the vector of implied volatilities, K is the vector of strike prices corresponding to these volatilities and ε is the residual of this OLS regression. For strike prices outside the observed range, one can extrapolate by extending the quadratic function outwards. The coefficients of this parabola are used to obtain a continuous call function with respect to the strike price. Breeden and Litzenberger (1978) showed that the second derivative of the option price function with respect to the strike price yields the implied risk-neutral probability density function $p^{SM}(S_T)$:

$$\frac{\partial^2 C}{\partial K^2} = e^{-r\tau} p^{SM}(S_T) \tag{16}$$

The implied probability density function can be used to calculate numerically the implied moments of the distribution. The first moment, the mean of the distribution is calculated as:

$$\mu_{p^{SM}(S_T)} = \int_{K=0}^{\infty} K p^{SM}(S_T)\, dK \tag{17}$$

In order to obtain the higher moments of the distribution, the first four central moments are used. The n-th central moment of the implied empirical distribution is defined as:

$$\mu^{(n)} = \int_{K=0}^{\infty} \left[K - S_t e^{(r-r^*)\tau}\right]^n p^{SM}(S_T)\, dK \tag{18}$$

Thus, the implied standard deviation of the fitted risk-neutral density function is given by:

$$\sigma_{p^{SM}(S_T)} = \sqrt{\frac{\ln\left[\frac{\mu^{(2)}}{\left[\mu^{(1)}\right]^2} + 1\right]}{\tau}} \tag{19}$$

and the implied skewness and kurtosis are defined as follows:

$$Skew_{p^{SM}(S_T)} = \frac{\mu^{(3)}}{\left[\mu^{(2)}\right]^{3/2}} \tag{20}$$

$$Kurt_{p^{SM}(S_T)} = \frac{\mu^{(4)}}{\left[\mu^{(2)}\right]^{2}} \tag{21}$$

Shimko used this method to recover risk-neutral density functions from the S&P 500 index options. He showed that the results of the analysis are very sensitive to the choice of the interpolation method. Call option prices can be interpolated directly. However, Shimko suggested that, in general, the option price function can be made smoother by interpolating instead implied volatility. The results of his analysis showed that for the S&P 500 index options the implied volatility profile was negatively sloped, implying a negative skew to the distribution of future index values. This suggests that the market was placing a relatively greater weight on a fall in the index than on a rise. He also documented a slight convexity of the implied volatility function, which indicates that the market was placing a greater weight on larger movements (positive or negative) than those implied by the Black and Scholes model.

Shimko calculated the distribution only between the endpoints from the smoothed volatility values. However, this did not cover the entire risk-neutral distribution. Outside the available strike prices there is a mass of probability which has to be approximated. To solve this, Shimko assumed that the tail distributions were lognormal.

The main advantage of this method is given by the stability of the optimisation algorithm, which has an analytic solution. The major drawback is that irregular volatility smile functions are (sometimes) poorly approximated with a parabola.[20]

7.2. Madan and Milne's method. Madan and Milne (1994) modelled the prices of contingent claims as elements of a separable Hilbert space that has a countable orthogonal basis. They noticed that one may think of the basis elements as analogous to factors in asset pricing. Thus, pricing in terms of a Hilbert space is analogous to the use of discount bonds as a basis for pricing fixed income securities or the construction of branches of a binomial tree in pricing options. However, a Hilbert space basis is in general difficult to construct because it requires a knowledge of the stochastic process of the underlying asset prices. Madan and Milne showed that, under fairly general conditions, one can specialise the Hilbert space basis to the family of Hermite polynomials. Using this assumption, one can infer the underlying risk-neutral density from traded security prices. This model has been applied to extract risk-neutral probability distributions from options written on stock index futures (Madan and Milne (1994), Coutant (1999)), interest rate futures (Abken, Madan and Ramamurtie (1996), McManus (1999), Coutant et al. (2001)) and exchange rates (Jondeau and Rockinger (2000)). In the case of exchange rates, the Garman-Kohlhagen model is a parametric special case of the Madan and Milne (1994) model. Thus, the Hermite polynomials approximation is equivalent to performing a Fourier expansion to the baseline lognormal solution obtained

[20] A direct implication of this drawback is on the value of skewness, which is sometimes close to zero, ie its theoretical value from a lognormal distribution.

from the Garman-Kohlhagen model. More precisely, the risk-neutral distribution is obtained through successive orthogonal perturbations to a normalised density function.[21] Thus, in the case of European currency options, the Hermite polynomial adjustments are constructed with respect to the normalised stochastic variable:

$$z = \frac{\ln\left(\frac{S_T}{F_t}\right) - \left(\mu - \frac{1}{2}\sigma^2\right)\tau}{\sigma\sqrt{\tau}}, \qquad z \sim N(0,1) \tag{22}$$

where $F_t = S_t e^{(r-r^*)\tau}$ is the forward price and $\mu = \ln(S_t) + \left(r - r^* - \sigma^2/2\right)\tau$ is the mean of the diffusion process from the Garman-Kohlhagen model. Thus, the risk-neutral normal distribution used as reference for the Hermite polynomials approximation is:

$$n(z) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right) \tag{23}$$

The Hermite polynomials approximation of the risk-neutral density function, $p^{HP}(z)$, can be written as:

$$p^{HP}(z) = \lambda(z)\, n(z) \tag{24}$$

where $\lambda(z)$ denotes the departures from the reference distribution $n(z)$, which are captured by an infinite summation of Hermite polynomials, that is:

$$\lambda(z) = \sum_{k=0}^{\infty} b_k \phi_k(z) \tag{25}$$

where b_k are constants which have to be estimated and

$$\phi_k(z) = \frac{(-1)^k}{k!} \frac{1}{n(z)} \frac{\partial^k n(z)}{\partial z^k} = -\frac{1}{\sqrt{k}} \frac{\partial \phi_{k-1}(z)}{\partial z} + \frac{1}{\sqrt{k}} z \phi_{k-1}(z) \tag{26}$$

is an orthogonal system of standardised Hermite polynomials. The first four standardised Hermite polynomials are:[22]

$$\begin{aligned}
\phi_0(z) &= 1 \\
\phi_1(z) &= z \\
\phi_2(z) &= \frac{1}{\sqrt{2}}\left(z^2 - 1\right) \\
\phi_3(z) &= \frac{1}{\sqrt{6}}\left(z^3 - 3z\right) \\
\phi_4(z) &= \frac{1}{\sqrt{24}}\left(z^4 - 6z^2 + 3\right)
\end{aligned} \tag{27}$$

[21] In other words, rather than assuming specific expressions for *the change* in the risk-neutral probabilities, as one does under the martingale approach for option valuation, Madan and Milne (1994) assume a parametric structure for the risk-neutral density function itself.

[22] Higher-order Hermite polynomials can be easily calculated using the recurrence relationship: $\phi_k(z) = \frac{z}{\sqrt{k}}\phi_{k-1}(z) - \sqrt{\frac{k-1}{k}}\phi_{k-2}(z)$. The polynomials are orthogonal because $\int_{-\infty}^{\infty} \phi_k(z)\,\phi_j(z)\,n(z)\,dz$ equals one if $k = j$ and zero otherwise.

With the above notations and $d_2 = \frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S_t}{K}\right) + \left(r - r^* - \frac{\sigma^2}{2}\right)\tau\right]$, one can write the Garman-Kohlhagen formula for the European-style call currency option as:[23]

$$\begin{aligned} C_{GK}(z) &= e^{-r\tau}\int_{-d_2}^{\infty}\left[F_t e^{\left(\mu-\frac{\sigma^2}{2}\right)\tau+z\sigma\sqrt{\tau}} - K\right] n(z)\,dz \\ &= e^{-r^*\tau}S_t\int_{-d_2}^{\infty} e^{\left(\mu-\frac{\sigma^2}{2}\right)\tau+z\sigma\sqrt{\tau}} n(z)\,dz - e^{-r\tau}K\int_{-d_2}^{\infty} n(z)\,dz \end{aligned} \tag{28}$$

By substituting the risk-neutral density function $p^{HP}(z)$ from equations (24) and (25), we obtain the Hermite polynomial approximation of the call option price:

$$\begin{aligned} C_{HP}(z) &= e^{-r^*\tau}S_t\int_{-d_2}^{\infty} e^{\left(\mu-\frac{\sigma^2}{2}\right)\tau+z\sigma\sqrt{\tau}}\sum_{k=0}^{\infty} b_k\phi_k(z)\,n(z)\,dz \\ &\quad -e^{-r\tau}K\int_{-d_2}^{\infty}\sum_{k=0}^{\infty} b_k\phi_k(z)\,n(z)\,dz \end{aligned} \tag{29}$$

To evaluate the price of this option, one needs to replicate its payoff and estimate the coefficients b_k. The expected payoff of the option can be expressed as:

$$g_{C_{HP}}(z) = \int_{-d_2}^{\infty}\left[F_t e^{\left(\mu-\frac{\sigma^2}{2}\right)\tau+z\sigma\sqrt{\tau}} - K\right]\sum_{k=0}^{\infty} b_k\phi_k(z)\,n(z)\,dz \tag{30}$$

and the call option price from equation (29) can be represented as follows:

$$C_{HP}(z) = e^{-r\tau}\sum_{k=0}^{\infty}\alpha_k b_k \tag{31}$$

where

$$\alpha_k = \int_{-\infty}^{\infty} g_{C_{HP}}(z)\,\phi_k(z)\,n(z)\,dz \tag{32}$$

Madan and Milne (1994) showed that, given the assumed probability model, the Hermite polynomial coefficients α_k are well defined and hence the b_k can be inferred from the observed option prices. The α_k coefficients are defined as:

$$\alpha_k = \frac{1}{\sqrt{k!}}\frac{\partial^k\Phi(u)}{\partial u^k}\,|_{u=0} \tag{33}$$

where the generating function $\Phi(u)$ is given by:

$$\begin{aligned} \Phi(u) &= S_t e^{(r-r^*)\tau} e^{\mu\tau+\sigma\sqrt{\tau}u} N[d_1(u)] - KN[d_2(u)] \\ d_1(u) &= \frac{\ln(S_t/X) + (r-r^*)\tau}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau} + u \\ d_2(u) &= d_1(u) - \sigma\sqrt{\tau} \end{aligned} \tag{34}$$

[23] To simplify the presentation we only give the derivation for the call option. The analysis for the put option valuation with a Hermite polynomial approximation is straightforward and follows the same reasoning as for the call option.

For practical estimation purposes, the infinite sum of Hermite polynomials must be truncated at a finite order in z. In our study we truncate it at the fourth order.[24] In order to ensure that the risk-neutral PDF for the Hermite polynomial approximation behaves as a density function, the following restrictions are usually imposed (see Abken, Madan and Ramamurtie (1996)): $\beta_0 = 1$, $\beta_1 = 0$ and $\beta_2 = 0$. Under these restrictions, the risk-neutral probability density function for the fourth-order Hermite polynomials approximation is given by:

$$p^{HP}(z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} \left(1 + \frac{3b_4}{\sqrt{24}} - \frac{3b_3}{\sqrt{6}} z - \frac{6b_4}{\sqrt{24}} z^2 + \frac{b_3}{\sqrt{6}} z^3 + \frac{b_4}{\sqrt{24}} z^4 \right) \tag{35}$$

and the risk-neutral PDF for the exchange rate at the maturity of the call option is:

$$p^{HP}(S_T) = \frac{1}{S_T \sigma \sqrt{\tau}} p^{HP} \left[\frac{\ln(S_T/F_t) - \left(\mu - \frac{1}{2}\sigma^2\right)\tau}{\sigma\sqrt{\tau}} \right] \tag{36}$$

The mean μ, the standard deviation σ and the parameters b_3 and b_4 are estimated by minimising the sum of squared differences between theoretical and observed option and forward prices. The mean of the risk-neutral PDF is equal to the theoretical forward price and is given by:

$$\mu_{p^{HP}(S_T)} = F_t e^{\mu t} \left[\sum_{k=0}^{4} \frac{b_k}{\sqrt{k!}} \left(\sigma\sqrt{\tau}\right)^k \right] \tag{37}$$

The implied standard deviation of the PDF is $\sigma_{p^{HP}(S_T)} = \sigma\sqrt{\tau}$ and the higher order moments, skewness and kurtosis, are calculated directly from the estimated parameters b_3 and b_4, namely:

$$\begin{aligned} Skew_{p^{HP}(S_T)} &= b_3\sqrt{6} \\ Kurt_{p^{HP}(S_T)} &= 3 + b_4\sqrt{24} \end{aligned} \tag{38}$$

A drawback of this method is that it occasionally provides unreliable estimates for the tails of the risk-neutral PDFs. This is because of the truncation of the infinite sum of Hermite polynomials, which may lead to anomalous negative values for the estimated tails of the risk-neutral distributions. Jondeau and Rockinger (2001) proposed a way of applying positivity constraints on the estimated risk-neutral density functions. They used a constrained expansion referred to as a Gram-Charlier density. Gram-Charlier expansions allow for additional flexibility over the normal density because they introduce the skewness and kurtosis of the distribution as parameters. However, being polynomial approximations, they have the drawback of yielding negative values for certain parameters. Jondeau and Rockinger (2001) showed how it is numerically possible to restrict the parameters of such densities. Their essential contribution was to delimit the domain in the skewness-kurtosis space over which the Gram-Charlier expansion is positive. Based on these boundaries, they presented a mapping which transforms the constrained estimation problem into an unconstrained one. For the empirical part of their paper they used foreign exchange over-the-counter FRF/DEM options and carried out two simulation exercises.

[24] Given the data restrictions, higher-order approximations are practically infeasable for foreign exchange over-the-counter options.

In the first simulation experiment they considered as the true data generating process random variables distributed according to the Gram-Charlier density. They found that the estimates for the first and second moments were very close to the theoretical ones. They also found that, on average, the estimates for skewness and kurtosis matched fairly well the simulated higher-order moments. However, they noticed that the estimates for kurtosis differ by a larger percentage from the theoretical values than the other moments. In particular, they found that for a given level of theoretical kurtosis, the smaller the skewness, the worse the average of the estimated kurtosis. Moreover, they noticed that as kurtosis increases, the dispersion of the parameters decreases. These results indicate that the Gram-Charlier density provides better results the more the distribution differs from the normal one. In their second simulation exercise, they assessed the ability of the Gram-Charlier density to capture correctly the moments of the data simulated from a mixture of two lognormal distributions. They showed that the first two moments of the simulated data were on average, up to the third decimal, identical with the theoretical ones. Moreover, the higher moments, skewness and kurtosis, were also close to the theoretical moments.

7.3. Malz's method. One of the basic assumptions used in the Garman-Kohlhagen model is the continuity and the normal distribution of the stochastic process that characterises the exchange rate changes. However, several empirical studies showed that nominal exchange rate distributions at high frequencies are significantly leptokurtic, which means that the probability of larger jumps in exchange rates is higher than that predicted by the lognormal assumption.[25] Empirical studies also suggested that there is an inverse relationship between excess kurtosis and the length of the holding period.[26] If jumps in either direction are equally likely, then there will be an excess kurtosis without having an impact on the skewness of the distribution. However, if jumps in one direction are larger and more frequent, the distribution will also be skewed. The skewness and kurtosis of such a stochastic process appear to be successfully captured with a class of jump-diffusion models proposed by Merton (1976).[27] Bates (1996a, 1996b) also found strong evidence that flexible exchange rate returns follow jump-diffusions, that is, a sum of independent identically distributed normal and Poisson distributed jump components. In his model the evolution of exchange rate returns is represented by two components, a diffusion process and a Poisson jump process:

$$S_T = S_t + \int_0^T \left(r - r^* - \lambda E\left[\kappa\right]\right) S_t dt + \int_0^T \sigma_w S_t dW_t + \int_0^T S_t \kappa dq_{t,T} \tag{39}$$

where σ_w denotes the diffusion volatility of the exchange rate, $q_{t,T}$ is a Poisson counter over the interval (t, T) with intensity λ, denoting the average rate of occurrence of jumps, and κ is the average random jump size.[28]

[25] See, inter alia, Boothe and Glassman (1987), Hsieh (1988), and de Vries (1994).

[26] For instance, Hsieh (1988) estimated unconditional kurtosis of 12.8 for daily changes in the USD/DEM exchange rate, while Meese (1986) estimated kurtosis of 4.2 for monthly returns.

[27] See Akgiray and Booth (1988), Tucker and Pond (1988), Jorion (1988) and Bates (1996a).

[28] The option valuation formuale for this model are derived by Merton (1976) and Bates (1991). An important issue in deriving the option value is that the risk generated by an increase in the option price following a jump in the underlying asset price cannot be managed by a continuous-adjusted hedging strategy. Thus, the option might jump further in-the-money, in which case the option writer will be

Malz (1996) proposed a simplified version of the jump-diffusion model to estimate the realignment probabilities for the sterling-mark exchange rate, for the period 31 March - 16 September 1992. During this period sterling came several times under speculative attacks. This seems to have been successfully reflected in option prices. Thus, the densities based on the jump-diffusion model had fatter tails than the lognormal, and the left tail was fatter than the right, indicating higher probabilities of exchange rate realignments, ie of sterling pound depreciation beyond the 4.5% band level established by the ERM. Following Ball and Torous (1983, 1985), Malz (1996) considered κ to be non-stochastic, ie it was specified as either zero or one over the life of the option.[29] In this model, option pricing formulae consist of a weighted sum of the lognormal solutions where the weights are given by the probability of no jumps occurring and one jump occurring over the lifetime of the option. Thus, the formulae for call and put option prices are:

$$\begin{aligned} C_M &= (1-\lambda\tau)\left[\frac{S_t e^{-r^*\tau}}{1+\lambda\kappa\tau} N(d_1) - K e^{-r\tau} N(d_2)\right] \\ &\quad + \lambda\tau \left[\frac{S_t e^{-r^*\tau}}{1+\lambda\kappa\tau}(1+\kappa) N(d_3) - K e^{-r\tau} N(d_4)\right] \end{aligned} \tag{40}$$

$$\begin{aligned} P_M &= (1-\lambda\tau)\left[K e^{-r\tau} N(-d_2) - \frac{S_t e^{-r^*\tau}}{1+\lambda\kappa\tau} N(-d_1)\right] \\ &\quad + \lambda\tau \left[K e^{-r\tau}(1+\kappa) N(-d_4) - \frac{S_t e^{-r^*\tau}}{1+\lambda\kappa\tau} N(-d_3)\right] \end{aligned} \tag{41}$$

where

$$d_1 = \frac{\ln(S_t/X) - \ln(1+\lambda\kappa\tau) + \left(r - r^* + \sigma_w^2/2\right)\tau}{\sigma_w\sqrt{\tau}} \tag{42}$$

$$d_2 = d_1 - \sigma_w\sqrt{\tau} \tag{43}$$

$$d_3 = \frac{\ln(S_t/X) - \ln(1+\lambda\kappa\tau) + \ln(1+\kappa) + \left(r - r^* + \sigma_w^2/2\right)\tau}{\sigma_w\sqrt{\tau}} \tag{44}$$

$$d_4 = d_3 - \sigma_w\sqrt{\tau} \tag{45}$$

The jump-diffusion model postulates that the parameters σ_w, λ and κ are constants. These parameters are estimated by minimising the sum of squared deviations of predicted from actual option prices.

The cumulative distribution function is given by the following expression:

underhedged. If an option seller attemps to hedge in advance of jumps, he will be overhedged unless a jump occurs. Therefore, in contrast to the Garman-Kohlhagen model, the jump-diffusion model does not permit risk-neutral pricing techniques without additional assumptions (see Ball and Torous (1983, 1985) and Jarrow and Rudd (1982)).

[29] This is often referred to as the Bernoulli distribution version of the jump-diffusion model. In the Bernoulli distribution model, the Poisson counter $q_{t,T}$ is zero with probability $1-\lambda$ and one with probability λ.

$$\Pr(S_T \leq K) = (1-\lambda) N \left[\frac{\ln\left(\frac{K}{F_t}\right) + \ln(1+\lambda\kappa) + \frac{\sigma_w^2}{2}}{\sigma_w} \right] + \lambda N \left[\frac{\ln\left(\frac{K}{F_t}\right) + \ln\left(\frac{1+\lambda\kappa}{1+\kappa}\right) + \frac{\sigma_w^2}{2}}{\sigma_w} \right] \quad (46)$$

The mean of the PDF is equal to the forward price and its variance is equal to $\sigma_w^2 \tau$. The higher moments of the distribution, skewness and kurtosis, are calculated using equations (18), (20) and (21), where $p^{SM}(S_T)$ is replaced by the probability density function corresponding to the distribution in equation (46).

The advantage of this method resides in its high capacity for characterising the evolution of exchange rate prices. Its main disadvantage is that the results may (sometimes) depend on the initial values used in the optimisation algorithm. This is often referred to as the *local optima problem.*

7.4. Melick and Thomas' method. This method is based on a mixture of lognormal distributions. Its main advantage is that it imposes little structure on the process by which exchange rates evolve and permits the estimation of relatively flexible functional forms for their distributions. A reasonably flexible functional form, such as the mixture of a finite number of lognormal distributions, can easily accommodate a wide variety of shapes for the terminal distribution. As Melick and Thomas (1997) pointed out, starting with an assumption about the terminal risk-neutral distribution function, rather than the stochastic process by which the underlying price evolves, has the advantage of being a more general approach. This is because a given stochastic process implies a unique terminal distribution, but the converse is not true, that is, any given risk-neutral distribution function may be consistent with many different stochastic processes. However, placing structure on the terminal distribution rather than on the stochastic process is not without costs. Thus, the recovered distribution does not give any guidance about the evolution of the asset price prior to expiration. This means that the resultant distribution cannot be directly used for constructing dynamic hedges or replication strategies for the option.

The specification of the distribution function for a mixture of n lognormals is as follows:

$$p^{MLN}(S_T) = \sum_{i=1}^{n} \lambda_i p_i^{MLN}(S_T) \quad (47)$$

where

$$p_i^{MLN}(S_T) = \frac{1}{\sigma_i S_T \sqrt{2\pi}} \exp\left[-\frac{1}{2} \left(\frac{\ln(S_T) - \mu_i}{\sigma_i} \right)^2 \right] \quad (48)$$

and

$$\sum_{i=1}^{n} \lambda_i = 1, \quad \text{and } 0 < \lambda_i \leqslant 1,\ i = 1, 2, \ldots, n \quad (49)$$

The mixture of lognormals method also has the advantage of retaining the Garman-Kohlhagen model as a special subcase.[30]

The number of lognormals is usually dictated by the data constraints. On the foreign exchange over-the-counter market only five unique strike prices can be identified. This makes very difficult the estimation of the mixture of lognormals. We choose a mixture of two lognormals which has five parameters to be estimated. As the number of parameters is equal to the number of unique strike prices, it is practically impossible to carry out the estimation without imposing additional restrictions. To solve this, we adopt a two-stage estimation strategy. First, we define a grid of probabilities λ between 0 and 0.5 with a step size of 0.01 and estimate the parameters μ_1, μ_2, σ_1 and σ_2.[31] Second, we calculate the difference between theoretical and observed prices and choose the pair of parameters $\theta = (\mu_1, \mu_2, \sigma_1, \sigma_2, \lambda)$ which minimises the loss function:

$$\min_{\theta} \sum_{j=1}^{m} \left[C - C(\theta)\right]^2 + \sum_{j=1}^{m} \left[P - P(\theta)\right]^2 + \sum_{j=1}^{m} \left[S_t e^{(r-r^*)\tau} - F(\theta)\right]^2 \tag{50}$$

where $m = 5$ is the number of unique strike prices used to estimate each individual probability density function, C and P are observed call and put option prices and $C(\theta)$, $P(\theta)$ and $F(\theta)$ are theoretical call, put and forward prices, expressed as follows:

$$\begin{aligned} C(\theta) &= \lambda \left[\exp\left(\mu_1 + \frac{1}{2}\sigma_1^2\right) N(d_1) - KN(d_2)\right] + \\ &\quad + (1-\lambda)\left[\exp\left(\mu_2 + \frac{1}{2}\sigma_2^2\right) N(d_3) - KN(d_4)\right] \end{aligned} \tag{51}$$

$$\begin{aligned} P(\theta) &= \lambda \left[-\exp\left(\mu_1 + \frac{1}{2}\sigma_1^2\right) N(-d_1) + KN(-d_2)\right] \\ &\quad + (1-\lambda)\left[-\exp\left(\mu_2 + \frac{1}{2}\sigma_2^2\right) N(-d_3) + KN(-d_4)\right] \end{aligned} \tag{52}$$

$$F(\theta) = \lambda \exp\left(\mu_1 + \frac{1}{2}\sigma_1^2\right) + (1-\lambda)\exp\left(\mu_2 + \frac{1}{2}\sigma_2^2\right) \tag{53}$$

with

$$d_1 = \frac{1}{\sigma_1}\left[\mu_1 + \sigma_1^2 - \ln(K)\right] \tag{54}$$

$$d_2 = d_1 - \sigma_1 \tag{55}$$

$$d_3 = \frac{1}{\sigma_2}\left[\mu_2 + \sigma_2^2 - \ln(K)\right] \tag{56}$$

$$d_4 = d_3 - \sigma_2 \tag{57}$$

[30] The Garman-Kohlhagen model is given by: $i = 1$, $\lambda = 1$, $\mu_1 = \ln(S_t) + \left(r - r^* - \sigma^2/2\right)\tau$ and $\sigma_1 = \sigma\sqrt{\tau}$.

[31] It is not necessary to define a grid of probabilities from 0.5 to 1 because $\pi_2 = 1 - \pi_1$, so that after 0.5 the lognormal distributions are reversed.

The moments of the distribution are calculated following the same approach as in the case of Shimko's (1993) method. The cumulative distribution function for the mixture of two lognormal distributions is expressed as:

$$\Pr\left(S_T \leq K\right) = \lambda N\left[\frac{\ln\left(K\right) - \mu_1}{\sigma_1}\right] + \left(1 - \lambda\right) N\left[\frac{\ln\left(K\right) - \mu_2}{\sigma_2}\right] \tag{58}$$

Melick and Thomas (1997) applied this methodology to extract implied risk-neutral density functions from the prices of American-style options on crude oil futures during the Persian Gulf war. They assumed that the terminal price distribution is a mixture of three independent lognormal distributions. They found that the option markets were consistent with the market commentary at the time, in that they reflected a significant probability of a major disruption in oil prices. They also showed that the standard lognormal assumption performed poorly compared with the mixture of lognormals in characterising the observed option prices.

7.5. Bliss and Panigirtzoglou's method. Following Shimko's (1993) approach, Bliss and Panigirtzoglou (2002) fitted the implied volatility function. However, they used a spline instead of the parabolic function proposed by Shimko (1993). The risk-neutral PDF is then easily obtained by taking the second derivative of fitted option prices with respect to the strike price. The essence of the method is to smooth implied volatilities rather than option prices and then convert the smoothed implied volatility function into a smoothed price function, which can be numerically differentiated to produce the probability density function. In their study, Bliss and Panigirtzoglou (2002) used the Newton-Raphson algorithm to recover the implied volatilities from observed option prices. The implied volatility function is then smoothed with a natural spline. The spline function can also be weighted. A natural weight would be the trading volume, as the information content of option prices is directly linked to the liquidity of these instruments. However, detailed data on options trading volumes is quite scarce and unreliable for index options, and nonexistent for currency options. As an alternative the *vega* of the option could be used. The vega is the first derivative of option's price with respect to its implied volatility. Vega weighting places less weight on deep out-of-the money options, which is consistent with the observed lower liquidity of such options. The natural spline minimises the following function:

$$\min_{\theta} \sum_{i=1}^{m} w_i \left[\sigma_i - \sigma\left(\Delta_i, \theta\right)\right] + \lambda \int_{-\infty}^{\infty} g^{''}\left(x; \theta\right)^2 dx \tag{59}$$

where σ_i is the implied volatility of the i-th option in the cross-section, $\sigma\left(\Delta_i, \theta\right)$ is the fitted implied volatility which is a function of the i-th option delta Δ_i, and the parameters, θ that define the smoothing spline $g\left(x; \theta\right)$, and w_i is the weight applied to the i-th option's squared implied volatility error.[32] The parameter λ is a smoothing parameter that controls the tradeoff between goodness-of-fit of the fitted spline and its smoothness measured by the integrated squared second derivative of the implied volatility

[32] Following Bliss and Panigirtzoglou (2002), we initially use the vega parameter to weight the implied volatility function. However, the impact on option pricing appears to be very small. Therefore, in our study, we decide not to weight the spline function. Thus, w_i is equal to one for all i.

function. Bliss and Panigirtzoglou (2002) tested various values of λ, between 0.99 and 0.9999 and found that the results were insensitive to the choice of λ. They chose 0.99 for their final estimations. In our study we also choose $\lambda = 0.99$.

When fitting a PDF it is necessary to extrapolate the spline beyond the range of available strike prices. Since one can rarely observe extreme realisations of the underlying asset, there is little information as to the appropriate shape to impose on the tails of the density function. Bliss and Panigirtzoglou (2002) used a horizontal extrapolation of the spline function outside the range of available data points. Another alternative was proposed by Anagnou et al. (2002) who truncated the implied density function to the range of available strikes and then rescaled. This unusual procedure avoids extrapolating the tails of the PDF, but cannot handle realisations falling outside the range of strikes available when the PDF was constructed.

7.6. Piecewise cubic Hermite interpolation method. Fitting a single polynomial to a large number of datapoints is likely to yield unsatisfactory behaviour of the PDFs, especially in the tails of the distribution.[33] Piecewise polynomial interpolation provides an alternative to the practical and theoretical difficulties associated with high-degree polynomial interpolation. Since the early 1960s, the subject of piecewise polynomial functions has become increasingly popular. These functions have been used in a large variety of ways in approximation theory, data fitting, numerical integration and differentiation. The basic idea is to decompose a function in a number of subintervals and interpolate on each subinterval. The points where the function is delimited are called knots, breakpoints or nodes. The simplest example of such an interpolation is piecewise linear interpolation, in which successive data points are connected by straight lines.

In Hermite interpolation the derivatives as well as the values of the interpolating function are specified at the data points. Specifying derivative values simply adds more equations to the linear system that determines the parameters of the interpolating function. In order to have a well-defined solution, the number of equations and the number of parameters to be determined must be equal.

To provide adequate flexibility while maintaining simplicity and computational efficiency, piecewise cubic polynomials are the most common choice for Hermite interpolation. A Hermite cubic interpolant is a piecewise cubic polynomial interpolant with a continuous first derivative.[34]

Although piecewise cubic Hermite interpolation eliminates the problems of excessive instability of the tails of implied risk-neutral density functions, it appears to sacrifice smoothness of the interpolating function. The moments of the distribution are obtained through numerical integration of implied density functions, as in the case of Shimko's (1993) method.

[33] See for example McManus (1999) for a discussion on the instability of the tails of implied risk-neutral density functions.

[34] See Appendix A for a mathematical description of the piecewise cubic Hermite interpolation method.

8. Statistical tests for comparison of methods

Our empirical comparison of the methods used to extract risk-neutral PDFs from option prices focuses on three main issues. First, we test whether there are significant differences between the PDF measures (moments and quartiles) across methods. Second, we check the accuracy of our selected methods in replicating observed option prices. Third, we rank the methods based on their robustness to option pricing errors.

In order to verify whether our selected methods give significantly different estimates for the implied PDF measures, we apply Student's t tests. The null hypothesis of these tests is that the unconditional mean of the differences between the PDF moments or quartiles estimated with two different methods is not significantly different from zero. If we reject the null hypothesis, we conclude that the methods provide significantly different estimates for the tested PDF moment or quartile.

One approach to test this is to compute the confidence interval for the sample mean differences. If we let X_t denote the difference between the estimates of an implied PDF measure with two different methods, then the t statistics is calculated as follows:[35]

$$t = \frac{\overline{X} - \mu}{s/\sqrt{n}} \tag{60}$$

where $\overline{X}$ is the unconditional mean of the estimated differences, μ is the theoretical mean assumed by the null hypothesis (in our case zero), s is the sample standard deviation, and n is the number of observations. In essence, the approach assumes that if we subtract the population mean from the sample mean, and then divide by the estimated standard deviation of the sample, we get a standard Gaussian curve (with mean zero and variance one). This assumption is incorrect when sample sizes are small.[36] However, if the sample size is sufficiently large, the Student's t test provides accurate results, even when the distributions are not normal.[37] In general, we test the validity of the null hypothesis at certain confidence levels. The confidence level represents the probability of rejecting the null hypothesis when this hypothesis is true. The usual confidence levels are 0.1, 0.05 and 0.01. If we denote with α the confidence level, the null hypothesis that the mean differences between an estimated PDF measure with two different methods is zero is rejected when

[35] For the sake of making the samples comparable across methods, we only selected those observations for which both methods provided meaningful estimates of the risk-neutral density function.

[36] A sample is considered small if it has less than 30 observations, when the empirical distribution is close to being normal. However, if the distribution displays leptokurtosis and non-zero skewness, the sample size should be of at least 100 observations to ensure proper asymptotic properties of the t test (Wilcox (2001)).

[37] We also use a bootstrap technique described by Efron and Tibshirani (1993) to check whether the results differ significantly from those obtained by using the Student's t distribution. We suppose that the observations sampled for differences are $X_1, X_2, \ldots, X_n$ and their unconditional mean and standard deviation are $\overline{X}$ and $\widehat{\sigma}$ respectively. The bootstrap test is based on the t-statistic: $t = \sqrt{n}\left(\overline{X}/\widehat{\sigma}\right)$. We define $\widetilde{X}_t = X_t - \overline{X}$, for $t = 1, 2, \ldots, n$. The null hypothesis of the test is that the distribution of these differences corresponds to the distribution where $\widetilde{X}_1, \widetilde{X}_2, \ldots, \widetilde{X}_n$ are equally likely. We refer to this distribution, which has the mean of zero, as the null distribution. We sample 1000 times with replacement from the null distribution and calculate $t^B = \sqrt{n}\left(\overline{X}^B/\widehat{\sigma}^B\right)$, where $\overline{X}^{\ B}$ and $\widehat{\sigma}^B$ are the sample bootstrapped mean and standard deviation. By comparing t with the appropriate percentile of this distribution, we are able to test whether the null hypothesis can be rejected at a particular confidence level. We find that the results do not differ significantly from those obtained by using the Student's t distribution.

$$\left|\overline{X}\right| > \left[t_{\nu}^{1-\alpha/2}\sqrt{n}\right] s \tag{61}$$

where $\nu = n - 1$ denotes the number of degrees of freedom, equal to the number of observations minus the number of tested parameters. If we cannot reject the null hypothesis, we conclude that there is no systematic difference between estimated PDF measures across methods.

In order to test the goodness-of-fit of our selected methods, we construct some synthetic measures of option pricing errors. These measures are based on cross-sectional differences between theoretical and observed option prices. In particular, we use the mean squared error, the mean squared percentage error, the mean absolute error and the mean absolute percentage error. The pricing errors are calculated for each day across strike prices, and then averaged across time. In order to rank the methods we apply again Student's t tests. In particular, if we denote with X_{1t} and X_{2t} the daily pricing errors of two PDF estimation methods and with $X_t = X_{1t} - X_{2t}$ the difference between these errors, the null hypothesis that the pricing errors of the first method are not higher than those of the second method is rejected if

$$\overline{X} > \left[t_{\nu}^{1-\alpha/2}\sqrt{n}\right] s \tag{62}$$

The tests are carried out for all currencies and maturities. The tables present the average p-values of these tests. If we reject the null hypothesis, we conclude that the second method has smaller pricing errors than the first method. A method is considered better if it has smaller pricing errors.

In order to test the robustness of the PDF estimation methods to option pricing errors, we perturb implied volatilities with random uniformly distributed amounts of maximum $\pm 10\%$ of their actual level. This approach is similar to that followed by Bliss and Panigirtzoglou (2002). The difference between our method and theirs is that we perturb implied volatilities rather than actual option prices. Our approach has the advantage of eliminating by construction possible violations of the arbitrage conditions.[38] We repeat our simulation 100 times and recalculate the PDF moments, quartiles and pricing errors. A method is considered stable if the estimated PDF moments or quartiles do not change dramatically when implied volatilities are randomly perturbed. In order to test the stability of the PDF measures, we calculate for each day the variance of the estimates and apply one-sided Student's t tests. The null hypothesis of the tests is that the cross-sectional variances of the first method are not higher than those of the second method. If we cannot reject the null hypothesis, we conclude that the first method is more stable to option pricing errors than the second method.[39]

[38] This might have a substantial impact on the results of the Monte-Carlo simulation. Indeed, much of the difference between the cubic spline and the mixture of lognormals approaches found by Bliss and Panigirtzoglou (2002) might be due to the fact they did not control for the arbitrage conditions violation of the simulated option prices.

[39] In the literature on implied risk-neutral density functions, the analysis of the stability of a PDF estimation technique to errors in option prices is much more complex. The stability of an estimated PDF is considered to have two components: 1) the theoretical stability at the solution; and 2) the stability of the convergence to a solution (see Bliss and Panigirtzoglou (2002)). Söderlind and Svensson (1997) and Melick and Thomas (1998) examined the stability at the solution. Their method assumed a normal distribution of the estimated parameters. However, actual parameter distributions may not be normal. The stability of the convergence to the "true" distribution has been analysed by Cooper (1999). He created European-type artificial option prices based on Heston's (1993) stochastic volatility model and tested how accurately a mixture of lognormals and a cubic spline method can fit the simulated distributions. Bliss

9. Empirical results

In order to compare the implied PDFs across methods, we look at conventional measures of location, dispersion, asymmetry, fat-tailness and various tail percentiles of the estimated risk-neutral distributions. On average, we find that the distributions exhibit slightly positive skewness and excess kurtosis. At times, the parametric methods (mixture of lognormals, Hermite polynomials and the geometric jump-diffusion) fail to converge. In order to eliminate possible distortions in the estimated means induced by the occasional non-convergence of the optimisation algorithm, we exclude from our sample those observations for which we could not obtain a PDF for the methods we compare. We then apply the three statistical tests described above. The results of the tests are presented in Appendix B.

9.1. Tests of the equality of the estimated distributions across methods. Table B.1 shows the tests of the mean equality of PDF means. The tests are carried out for all currencies and maturities. The table shows the average p-values of the tests. We are unable to reject the null hypothesis that the means of the estimated PDFs are equal across methods. This result is not surprising, since it was already confirmed by other studies that used currencies of developed countries or other underlying assets. It suggests that, regardless of the method we choose, we obtain similar estimates for the mean of the implied risk-neutral density functions. However, in the case of the variance, the results are slightly different. Table B.2 shows that variances estimated with Hermite polynomials and the geometric jump-diffusion methods are significantly different from those estimated with other methods. One possible reason of this is that the Hermite polynomials method tends to underestimate the variance of the PDFs compared with other methods, whereas the geometric jump-diffusion method tends to overestimate it.

Tables B.3 and B.4 present the tests of the mean equality of skewness and kurtosis across methods. We find that the higher moments of implied risk-neutral density functions (skewness and kurtosis) vary widely across methods. The parametric methods (mixture of lognormals, Hermite polynomials and the geometric jump-diffusion) exhibit higher excess kurtosis than the semi-parametric ones (quadratic, cubic spline or piecewise cubic Hermite interpolation). The sign of the skewness estimate is rather consistent across methods, albeit its magnitude varies substantially, being more or less model-dependent.

Tables B.5 - B.9 show the tests of the mean equality of the estimated PDF quartiles. Interestingly, we find that the quartiles are very similar across methods, except for the geometric-jump diffusion approach. The reason might be that the estimated jump-diffusion volatilities, which are used for the computation of cumulative distribution functions, are significantly different from those obtained with other methods.

9.2. Tests of the goodness-of-fit. We use four synthetic pricing error measures to test for the goodness-of-fit of the estimated risk-neutral PDFs: the mean squared error, the mean squared percentage error, the mean absolute error, and the mean absolute percentage error. The results are presented in Tables B.10-B.13.

We find that the piecewise cubic Hermite interpolation method ranked first, with the smallest pricing errors. This may be explained by the fact being a local interpolation it

and Panigirtzoglou (2002) critisised Cooper's (1999) approach, arguing that perturbing fitted rather than actual prices may influence significantly the convergence behaviour of the optimisation algorithm.

is very flexible and can easily capture irregular changes in the implied volatility function. However, the piecewise cubic Hermite interpolation sacrifices smoothness for its goodness-of-fit performance.

The second best method in terms of pricing errors appears to be the Hermite polynomials. The Hermite polynomials method is quite flexible to changes in the shape of implied risk-neutral PDFs. However, one drawback of this method is that, at times, the tails of the implied distributions might take negative values. Jondeau and Rockinger (2001) proposed an algorithm to constrain the higher order parameters of the implied PDFs and avoid negative probability estimates. However, their algorithm requires constrained optimisations, which have a much higher likelihood of non-convergence.

The next ranked method is the cubic spline. This method is theoretically more robust to estimation errors, as it does not require a numerical optimisation to obtain the implied density functions. Moreover, it has the capacity to control the smoothness of the distribution, by changing the parameter λ. Despite these advantages, for a small number of cross-sectional observations the fitting might not be very good. This is because the cubic spline polynomial requires at least 5 observations to get a solution and imposes constant derivatives at the endpoints of the distribution, which may hamper the goodness-of-fit.

The geometric jump-diffusion method ranked fourth in terms of pricing errors. Given that the exchange rate distributions exhibit stylised facts consistent with a jump-diffusion process (ie leptokurtosis, non-zero skewness and discrete jumps) the jump-diffusion method would seem a preferred candidate for the estimation of risk-neutral density functions. However, this method has a couple of disadvantages. First, the expected jump-size is relatively dependent on the initial values provided to the optimisation algorithm. We use as initial values a 15% jump size and a 10% probability of such a jump over the lifetime of an option contract. The estimated jump-size ranges between 5% and 32% whereas the probability of a jump varies more widely, between 3% and 44%, depending on the currency and the time period. Usually, for fixed exchange rates, the expected jump size is higher and the jump probability is lower than for flexible exchange rates. Second, the quartiles of the distribution appear to be highly dependent on the estimated jump-diffusion volatility, which suggests that the goodness-of-fit of the distribution is crucially determined by the estimated volatility parameter.

The mixture of lognormals performed rather poorly in fitting the observed option prices. The advantage of the mixture of lognormals is that it can accommodate more flexible functional forms for the implied risk-neutral distributions (eg. bi-modal distributions). However, given that the number of parameters is equal to the number of cross-sectional observations, the estimation of this method is time-consuming. Moreover, this method is quite sensitive to the optimisation algorithm and initial values of parameters. We use a combination of the Nelder-Mead simplex and BFGS algorithms for the optimisation. As initial values we use the estimated parameters from the Garman-Kohlhagen model.

The quadratic interpolation method performs equally bad in fitting observed option prices. However, this is not surprising given that this method is less flexible to changes in the shape of the implied volatility function. Despite its relative inflexibility, the quadratic interpolation method is very stable to errors in option prices and gives meaningful estimates for the implied PDFs even when the majority of the previous methods fail to converge to a solution.

9.3. Tests of the stability. The results of the tests are presented in Tables B.14-B.22. A method is considered stable if the standard deviation of a PDF measure obtained from randomly perturbed implied volatilities is small. The null hypothesis of the tests is that the mean value of the daily standard deviations of a PDF measure estimated with the row method is not higher than the mean value of the daily standard deviations estimated with the column method. The tables present the average probability of the t tests. The average is across currencies for PDFs with 1-month time to expiry.

We find that quadratic interpolation and cubic spline methods are the most robust to errors in option prices for the estimation of the mean and quartiles of implied PDFs, whereas Hermite polynomials, mixture of lognormals and geometric jump-diffusion methods are the most stable for the estimation of the variance, skewness and kurtosis. There seems to be a tradeoff between the goodness-of-fit and the stability of a PDF estimation method. Thus, methods which have a better accuracy in fitting observed option prices appear to be more sensitive to option pricing errors, while the most stable methods have a fairly disappointing fitting.

Given the tradeoff between goodness-of-fit and stability, the choice of a PDF estimation method is not clear cut. To make a selection, we need to know how big are the pricing errors and how sensitive are the PDF estimation methods to errors in option prices.

Given that we need comparable measures of option pricing errors across methods, we only examine the mean absolute and mean squared percentage errors. We find that the average errors vary substantially across methods and tend to increase with the maturity of option contracts. Thus, the average mean absolute percentage errors for the piecewise cubic Hermite interpolation method is 0.24% for options with 1-month time to maturity, and increases to 0.26% and 0.28% for 3- and 6-month option contracts respectively. The mean squared percentage errors are much smaller, namely 0.002%, 0.002%, and 0.003% for 1-, 3- and 6-month time to maturity options. The spectrum of errors for the other methods is relatively wide compared with the piecewise cubic Hermite interpolation method, ranging from 1.86% to 20.98%, in the case of the mean absolute percentage errors, and from 0.41% to 9.88%, in the case of mean squared percentage errors. The quadratic interpolation method exhibits the largest pricing errors.

In order to check how sensitive are our selected PDF estimation methods to possible errors in the implied volatility quotes, we perturb the implied volatilities with random uniformly distributed numbers of maximum ±10% of their level. We then test for each day whether the average of the simulated PDF measures is equal to the original estimate for that day.[40] We find that, for all methods, we cannot reject the null hypothesis that the average of the simulated PDFs is equal to the unperturbed PDF estimates, which means that errors of up to 10% in the quoted implied volatility do not significantly affect the PDF estimates, regardless of the method used. This result suggests that our selected methods are relatively robust to possible errors in option prices.

10. Conclusions

In this paper we compare the goodness-of-fit and the stability of six methods used to extract risk-neutral probability density functions from currency option prices. We use quadratic interpolation, Hermite polynomials, geometric jump-diffusion, mixture of

[40] To get synthetic measures of these tests, we average the p-values of the t tests across time and currencies, for option contracts with 1-month time to expiry.

lognormals, cubic spline and piecewise cubic Hermite interpolation methods. We find that the piecewise cubic Hermite interpolation method is by far the method with the best accuracy in fitting observed option prices. We also find that there is a relative tradeoff between the goodness-of-fit and the stability of the methods. Thus, methods which have a better accuracy in fitting observed option prices appear to be more sensitive to option pricing errors, while the most stable methods have a fairly poor fitting. However, for the first two PDF moments as well as the quartiles of the risk-neutral distributions we find that the estimates do not differ significantly across methods. This suggests that there is a large scope for selection between these methods without essentially sacrificing the accuracy of the analysis. Nonetheless, depending on the particular use of these PDFs, some methods may be more suitable than others.

References

[1] Abken, Peter A., Madan, Dilip B., and Ramamurtie, Sailesh, (1996), "Estimation of Risk-Neutral and Statistical Densities By Hermite Polynomial Approximation: With an Application to Eurodollar Futures Options", *Federal Reserve Bank of Atlanta Working Paper 96-5.*

[2] Akgiray, Vedat, Booth, Geoffrey G., (1988), "Mixed diffusion-jump process modeling of exchange rate movements", *Review of Economics and Statistics*, Vol. 70, pp. 631-637.

[3] Aït-Sahalia, Yacine, and Lo, Andrew W., (1998), "Nonparametric estimation of state-price densities implied in financial asset prices", *Journal of Finance*, Vol. 53, pp. 499-547, April.

[4] Ait-Sahalia, Yacine and Lo, Andrew W., (2000), "Nonparametric Risk Management and Implied Risk Aversion", *Journal of Econometrics*, Vol. 94, Issue 1-2, pp. 9–51.

[5] Anagnou, Iliana, Bedendo, Mascia, Hodges, Stewart, and Tompkins, Robert, (2002), "The relation between implied and realised probability density functions", *Working paper*, University of Technology, Vienna.

[6] Andersen, Leif, and Andreasen, Jesper, (1999), "Jumping Smiles", *Risk*, Vol. 12, No. 11, November, pp. 65-68.

[7] Atkinson, Kendall E., (1989), "*An introduction to numerical analysis*", John Wiley & Sons, USA.

[8] Bachelier, Louis, (1900), "Théorie de la spéculation", *Thèse,* Annales scientifiques de l'Ecole Normale Supérieure, 3e série, tome 17, Paris, Gauthier-Villars, pp. 21-86.

[9] Bahra, Bhupinder, (1996), "Probability Distributions of Future Asset Prices Implied by Option Prices," *Bank of England Quarterly Bulletin*, August 1996, pp. 299-311.

[10] Ball, Clifford A. and Torous, Walter N., (1983), "A Simplified Jump Process for Commom Stock Returns", *Journal of Financial and Quantitative Analysis*, Vol. 18, pp. 53-65.

[11] Ball, Clifford A. and Torous, Walter N., (1985), "On jumps in common stock prices and their impact on call option pricing", *Journal of Finance*, Vol. 42, pp. 301-320.

[12] Bates, David S., (1991), "The Crash of '87: Was It Expected? The Evidence from Options Markets", *Journal of Finance*, Vol. 46, pp. 1009-1044, July.

[13] Bates, David S., (1996a), "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options", *Review of Financial Studies*, Vol. 9, No. 1, pp. 69-107.

[14] Bates, David D., (1996b), "Dollar jump fears, 1984-1992: distributional abnormalities implicit in currency futures options", *Journal of International Money and Finance*, Vol. 15, No. 1, pp. 65-93.

[15] Black, Fischer (1976), "The Pricing of Commodity Contracts," *Journal of Financial Economics*, Vol. 3, pp. 167-179.

[16] Black, Fischer, and Scholes, Myron, (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, Vol. 81, pp. 637-59, May-June.

[17] Bliss, Robert R., and Panigirtzoglou, Nikolaos, (2002), "Testing the Stability of Implied Probability Density Functions", *Journal of Banking and Finance*, 26, pp. 381-422, March.

[18] Bliss, Robert R., and Panigirtzoglou, Nikolaos, (2004), "Option-Implied Risk Aversion Estimates", *Journal of Fianance*, Vol. 59, No. 1, pp. 407-446.

[19] Bondarenko, Oleg, (2003), "Estimation of Risk-Neutral Densities Using Positive Convolution Approximation", *Journal of Econometrics*, Vol. 116, pp. 85-112.

[20] Boothe, Paul, and Glassman, Debra, (1987), "The statistical distribution of exchange rates", *Journal of International Economics*, Vol. 22, pp. 297-319.

[21] Boyle, Phelim P., (1976), "Options: A Monte Carlo approach", *Journal of Financial Economics*, Vol 4., pp. 323-338.

[22] Boyle, Phelim P. and Emanuel, David, (1980), "Discretely Adjusted Option Hedges", *Journal of Financial Economics*, Vol. 8, pp. 259-282.

[23] Breeden, Douglas T., and Litzenberger, Robert H., (1978), "Prices of state-contingent claims implicit in option prices", *Journal of Business*, Vol. 51, pp. 621-651, October.

[24] Brennan, Michael J., (1979), "The pricing of contingent claims in discrete time models", *Journal of Finance*, Vol. 34, pp. 53-68.

[25] Campa, José M., and Chang, Kevin P.H., (1995), "Testing the expectations hypothesis of the term structure of volatilities in foreign exchange options", *Journal of Finance*, Vol. 50, pp. 529-547.

[26] Campa, José Manuel, Chang, Kevin P.H., and Reider, Robert, (1998), "Implied Exchange Rate Distributions: Evidence from OTC Option Markets", *Journal of International Money and Finance*, Vol. 17, No. 1, pp. 117-160.

[27] Carr, Peter, and Wu, Liuren, (2003), "What Type of Process Underlies Options? A Simple Robust Test", *Journal of Finance*, Vol. 58, No. 6, pp. 2581-2610.

[28] Coutant, Sophie, (1999), "Implied risk aversion in option prices using Hermite polynomials", in "*Estimating and interpreting probability density functions*", Bank for International Settlements.

[29] Coutant, Sophie, Rockinger, Michael, and Jondeau, Eric, (2001), "Reading PIBOR Futures Options' Smile: The 1997 Snap Election", *Journal of Banking and Finance*, Vol 25, No. 11, pp 1957-1987.

[30] De Vries, Casper G., (1994), "Stylized Facts of Nominal Exchange Rate Returns", in Van der Ploeg, Frederick, "*The Handbook of International Macroeconomics*", Blackwell, Oxford.

[31] Efron, Bradley and Tibshirani, Robert J., (1993), "*An Introduction to the Bootstrap*", Chapman & Hall, New York.

[32] Engel, Charles M., and Hamilton, James D., (1990), "Long swings in the dollar: are they in the data and do markets know it?", *American Economic Review*, Vol. 80, pp. 689-713.

[33] Galai, Dan, (1983), "The Components of the Return from Hedging Options Against Stocks", *Journal of Business*, Vol. 56, pp. 45-54.

[34] Garman, Mark B., and Kohlhagen, Steven W., (1983), "Foreign currency option values", *Journal of International Money and Finance*, Vol. 2, pp. 231-237, May.

[35] Heston, Steven L., (1993), "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", *Review of Financial Studies*, Vol. 6, pp. 327-343.

[36] Heynen, Ronald C., (1994), "An Empirical Investigation of Observed Smile Patterns", *Review of Futures Markets*, Vol. 13, No. 2, pp. 317-353.

[37] Hsieh, David A., (1988), "Statistical properties of daily exchange rates: 1974-1983", *Journal of International Economics*, Vol. 24, pp. 124-145.

[38] Jackwerth, Jens Carsten, (2000), "Recovering Risk Aversion from Option Prices and Realized Returns", *Review of Financial Studies*, Vol. 13, pp. 433-451.

[39] Jarrow, Robert, and Rudd, Andrew, (1982), "Approximate valuation for arbitrary stochastic processes", *Journal of Financial Economics*, Vol. 10, pp. 347-369.

[40] Jondeau, Eric, and Rockinger, Michael, (2000), "Reading the smile: the message conveyed by methods which infer risk neutral densities", *Journal of International Money and Finance*, Vol. 19, pp. 885-915.

[41] Jondeau, Eric, and Rockinger, Michael, (2001), "Gram-Charlier densities", *Journal of Economic Dynamics and Control*, Vol 25, pp. 1457-1483.

[42] Jorion, Philippe, (1988), "On jump processes in the foreign exchange and stock markets", *Review of Financial Studies*, Vol. 1, pp. 427-445.

[43] Madan, Dilip B., and Milne, Frank, (1994), "Contingent claims valued and hedged by pricing and investing in a basis", *Mathematical Finance*, Vol. 4, pp. 223-245.

[44] Malz, Allan M., "Using option prices to estimate realignment probabilities in the European Monetary System: the case of sterling-mark", *Journal of International Money and Finance*, Vol. 15, pp. 717-748, October.

[45] Malz, Allan M., "Estimating the probability distribution of the future exchange rate from options prices", (1997), *Journal of Derivatives*, pp. 18-36, Winter.

[46] McCauley, Robert, and Melick, William R., (1996a), "Propensity and density", *Risk*, Vol. 9, pp. 52-54, May.

[47] McCauley, Robert, and Melick, William R., (1996b), "Risk reversal risk", *Risk*, Vol. 9, pp. 54-57, November.

[48] McManus, Des J., (1999), "The information content of interest rate futures options", *Working Paper, No. 99-15*, Bank of Canada.

[49] Meese, Richard A., (1986), "Testing for Bubbles in Exchange Markets: A Case of Sparkling Rates", *Journal of Political Economy*, Vol. 94, pp. 345-373.

[50] Melick, William R., and Thomas, Charles 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*, Vol. 32, pp. 91-115, March.

[51] Melick, William R., and Thomas, Charles P., (1998), "Confidence Intervals and Constant-Maturity Series for Probability Measures Extracted from Options Prices", Bank of Canada, mimeo.

[52] Merton, Robert C., (1973), "The Theory of Rational Option Pricing", *Bell Journal of Economics and Management Science*, Vol. 4, pp. 141-183.

[53] Merton, Robert C., (1976), "Option Pricing When Underlying Stock Returns are Discontinuous", *Journal of Financial Economics*, Vol. 3, pp. 125-144.

[54] Poon, Ser-Huang, and Granger, Clive W., (2003), "Forecasting volatility in financial markets", *Journal of Economic Literature*, Vol. 41, pp. 478-539.

[55] Pratt, John W., (1964), "Risk Aversion in the Small and in the Large", *Econometrica*, Vol. 32, No. 1/2, pp. 122-136.

[56] Rosenberg, Joshua V., and Engle, Robert F., (2002), "Empirical Pricing Kernels", *Journal of Financial Economics*, Vol. 64, Issue 3, pp. 341–372.

[57] Rubinstein, Mark, (1994), "Implied binomial trees", *Journal of Finance*, Vol. 49, pp. 771-818, July.

[58] Rubinstein, Mark, (2002), "The Valuation of Uncertain Income Streams and the Pricing of Options", in *Asset Pricing Theory and Tests*, edited by Robert Grauer, Elsevier Science.

[59] Samuelson, Paul A., (1965), "Rational Theory of Warrant Pricing", *Industrial Management Review*, Vol. 6, Issue 2, pp. 13-31.

[60] Sherrick, Bruce J., Garcia, Philip, and Tirupattur, Viswanath, (1996), "Recovering Probabilitic Information from Option Markets: Tests of Distributional Assumptions", *Journal of Futures Markets*, Vol. 16, No. 5, pp. 545-560.

[61] Shimko, David C., (1993), "Bounds of probability", *Risk*, Vol. 6, pp. 33-37, April.

[62] Söderlind, Paul, and Svensson, Lars E. O., (1997), "New Techniques to Extract Market Expectations from Financial Instruments", *Journal of Monetary Economics*, Vol. 40, Issue 2, pp. 383-429.

[63] Taylor, Stephen J., and Xu, Xinzhong, (1994), "The term structure of volatility implied by foreign exchange options", *Journal of Financial and Quantitative Analysis*, Vol. 29, pp. 57-74.

[64] Taylor, Stephen J., and Xu, Xinzhong, (1995), "Conditional volatility and the informational efficiency of the PHLX currency options market", *Journal of Banking and Finance*, Vol. 19, pp. 803-821.

[65] Tucker, Alan L, and Pond, Lallon, (1988), "The Probability Distribution of Foreign Exchange Price Changes: Tests of Candidate Processes", *Review of Economics and Statistics*, Vol. 70, No. 4, pp. 638-647.

[66] Weinberg, Steven A., (2001), "Interpreting the Volatility Smile: An Examination of the Information Content of Options Prices", Board of Governors of the Federal Reserve System, *International Finance Discussion Paper no. 706*.

[67] Wilcox, Rand R., (2001), "*Fundamentals of Modern Statistical Methods*", Springer-Verlag, New York.

11. Appendices

A. Piecewise cubic Hermite interpolation method

Hermite interpolation finds a polynomial $p(x)$ that interpolates a function $f(x)$ and a polynomial $p^{'}(x)$ that interpolates $f^{'}(x)$, when $f(x)$ and $f^{'}(x)$ are given at each data point x_i, for $i = 1, \ldots, n$. The basic interpolation problem is given as follows:

$$p(x_i) = y_i, p^{'}(x_i) = y_i^{'}, i = 1, \ldots, n \tag{63}$$

The interpolating polynomials can be stated in an analogous form to Lagrange's formula. The Hermite interpolation polynomial is given by:

$$H_n(x) = \sum_{i=1}^{n} y_i h_i(x) + \sum_{i=1}^{n} y_i^{'} \bar{h}_i(x) \tag{64}$$

where

$$h_i(x) = \left[1 - 2l_i^{'}(x)(x - x_i)\right]\left[l_i^{'}(x)\right]^2 \tag{65}$$

$$\bar{h}_i(x) = (x - x_i)\left[l_i(x)\right]^2 \tag{66}$$

and

$$l_i(x) = \frac{(x - x_1)\ldots(x - x_{i-1})(x - x_{i+1})\ldots(x - x_n)}{(x_i - x_1)\ldots(x_i - x_{i-1})(x_i - x_{i+1})\ldots(x_i - x_n)} \tag{67}$$

To determine the Hermite polynomials, we do not need to determine and evaluate the Lagrange polynomials and their derivatives. An alternative method, based on Newton's divided difference formula, is the following:

$$\begin{aligned} p_{2n-1}(x) &= f(z_1) + (x - z_1) f[z_1, z_2] + \ldots \\ &\quad + (x - z_1)\ldots(x - z_{2n-1}) f[z_1, \ldots, z_{2n}] \end{aligned} \tag{68}$$

Defining a sequence of knots $z_1, z_2, \ldots, z_{2n}$ by

$$z_{2i-1} = z_{2i} = x_i \tag{69}$$

for all $i = 1, 2, \ldots, n$ we obtain

$$\begin{aligned} p_{2n-1}(x) &= f(x_1) + (x - x_1) f[x_1, x_2] + \ldots \\ &\quad + (x - x_1)^2 f[x_1, x_1, x_2] + \ldots \\ &\quad + (x - x_1)^2 (x - xn - 1)^2 (x - x_n) f[x_1, x_1, \ldots, x_n, x_n]. \end{aligned} \tag{70}$$

From equations [69] and [70] we have that

$$f[z_{2i-1}, z_{2i}] = f^{'}(z_{2i-1}) = f^{'}(x_i) \tag{71}$$

Thus we can use the derivatives

$$f^{'}(x_1), f^{'}(x_2), \ldots, f^{'}(x_n) \tag{72}$$

in place of the divided differences

$$f\left[z_1,z_2\right],f\left[z_3,z_4\right],\ldots,f\left[z_{2i-1},z_{2i}\right]. \tag{73}$$

Knowing that the Hermite form of the cubic polynomial which solves

$$\begin{aligned} p\left(x_{i-1}\right) &= f\left(x_{i-1}\right),p^{'}\left(x_{i-1}\right)=f^{'}\left(x_{i-1}\right) \text{ and} \\ p\left(x_{i}\right) &= f\left(x_{i}\right),p^{'}\left(x_{i}\right)=f^{'}\left(x_{i}\right) \end{aligned} \tag{74}$$

is the following (see Atkinson (1989)):

$$\begin{aligned} H_2\left(x\right) &= \left[1+2\frac{x-x_{i-1}}{x_i-x_{i-1}}\right]\left[\frac{x_i-x}{x_i-x_{i-1}}\right]^2 f\left(x_{i-1}\right) \\ &+\left[1+2\frac{x_i-x}{x-x_{i-1}}\right]\left[\frac{x-x_{i-1}}{x_i-x_{i-1}}\right]^2 f\left(x_i\right) \\ &+\frac{\left(x-x_{i-1}\right)\left(x_i-x\right)^2}{\left(x_i-x_{i-1}\right)^2}f^{'}\left(x_{i-1}\right) \\ &-\frac{\left(x-x_{i-1}\right)^2\left(x_i-x\right)}{\left(x_i-x_{i-1}\right)^2}f^{'}\left(x_i\right) \end{aligned} \tag{75}$$

The divided difference formula becomes

$$\begin{aligned} H_2\left(x\right) &= f\left(x_{i-1}\right)+\left(x-x_{i-1}\right)f^{'}\left(x_{i-1}\right) \\ &+\left(x-x_{i-1}\right)^2 f\left[x_{i-1},x_{i-1},x_i\right] \\ &+\left(x-x_{i-1}\right)^2\left(x-x_i\right)f\left[x_{i-1},x_{i-1},x_i,x_i\right] \end{aligned} \tag{76}$$

where

$$f\left[x_{i-1},x_{i-1},x_i\right]=\frac{f\left[x_{i-1},x_i\right]-f^{'}\left(x_{i-1}\right)}{x_i-x_{i-1}} \tag{77}$$

and

$$f\left[x_{i-1},x_{i-1},x_i,x_i\right]=\frac{f^{'}\left(x_i\right)-2f\left[x_{i-1},x_i\right]+f^{'}\left(x_{i-1}\right)}{\left(x_i-x_{i-1}\right)^2} \tag{78}$$

Hermite piecewise polynomial interpolation is a local interpolation, where the polynomial $p\left(x\right)$ on each subinterval $\left[x_{i-1},x_i\right]$ is determined by its interpolating data points. When we have a cubic Hermite interpolation polynomial the piecewise polynomial has four degrees.

The algorithm to find the piecewise polynomial interpolation function based on cubic Hermite interpolation polynomial using divided differences is given as follows: $\forall x\in\left[x_{i-1},x_i\right]$ where $i=1,2,\ldots,n$, if we note with

$$
\begin{aligned}
xx &= \left(x - x_{i-1}\right)^2, F = \frac{f\left(x_i\right) - f\left(x_{i-1}\right)}{x_i - x_{i-1}}, F_{i-1} = \frac{F - f^{'}\left(x_{i-1}\right)}{x_i - x_{i-1}} \text{ and} \\
F_i &= \frac{f^{'}\left(x_i\right) - 2F + f^{'}\left(x_{i-1}\right)}{\left(x_i - x_{i-1}\right)^2}
\end{aligned} \tag{79}
$$

then

$$
H = f\left(x_{i-1}\right) + \left(x - x_{i-1}\right) f^{'}\left(x_{i-1}\right) + xxF_{i-1} + xx\left(x - x_i\right) F_i. \tag{80}
$$

B. Comparison of methods

Table B.1. Test of the mean equality of PDF means across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.41	1.00				
JDM	1.00	0.41	1.00			
MLM	0.46	0.44	0.46	1.00		
CSM	0.48	0.63	0.48	0.48	1.00	
PIM	0.46	0.63	0.46	0.47	0.79	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of the PDF mean estimated with the row method is equal to the mean value of the PDF mean estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.2. Test of the mean equality of PDF variances across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.04	1.00				
JDM	0.03	0.01	1.00			
MLM	0.15	0.08	0.03	1.00		
CSM	0.14	0.08	0.00	0.27	1.00	
PIM	0.34	0.04	0.01	0.22	0.31	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of the PDF variance estimated with the row method is equal to the mean value of the PDF variance estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.3. Test of the mean equality of PDF skewness across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.02	1.00				
JDM	0.01	0.01	1.00			
MLM	0.03	0.10	0.01	1.00		
CSM	0.04	0.05	0.01	0.01	1.00	
PIM	0.03	0.07	0.00	0.01	0.05	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that that the mean value of the PDF skewness estimated with the row method is equal to the mean value of the PDF skewness estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.4. Test of the mean equality of PDF kurtosis across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.02	1.00				
JDM	0.02	0.01	1.00			
MLM	0.02	0.04	0.00	1.00		
CSM	0.04	0.01	0.00	0.04	1.00	
PIM	0.02	0.01	0.00	0.02	0.05	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that that the mean value of the PDF kurtosis estimated with the row method is equal to the mean value of the PDF kurtosis estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.5. Test of the mean equality of PDF 10% quartiles across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.33	1.00				
JDM	0.00	0.11	1.00			
MLM	0.87	0.37	0.00	1.00		
CSM	0.93	0.34	0.00	0.85	1.00	
PIM	1.00	0.32	0.00	0.84	0.98	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of PDF 10% quartiles estimated with the row method is equal to the mean value of PDF 10% quartiles estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.6. Test of the mean equality of PDF 25% quartiles across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.42	1.00				
JDM	0.00	0.12	1.00			
MLM	0.55	0.30	0.00	1.00		
CSM	0.93	0.46	0.00	0.52	1.00	
PIM	0.94	0.43	0.00	0.74	0.88	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of PDF 25% quartiles estimated with the row method is equal to the mean value of PDF 25% quartiles estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.7. Test of the mean equality of PDF 50% quartiles across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.48	1.00				
JDM	0.00	0.07	1.00			
MLM	0.77	0.26	0.03	1.00		
CSM	0.96	0.50	0.00	0.79	1.00	
PIM	1.00	0.48	0.01	0.83	1.00	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of PDF 50% quartiles estimated with the row method is equal to the mean value of PDF 50% quartiles estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.8. Test of the mean equality of PDF 75% quartiles across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.40	1.00				
JDM	0.00	0.09	1.00			
MLM	0.92	0.33	0.00	1.00		
CSM	0.88	0.51	0.00	0.75	1.00	
PIM	0.88	0.49	0.00	0.74	0.86	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of PDF 75% quartiles estimated with the row method is equal to the mean value of PDF 75% quartiles estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.9. Test of the mean equality of PDF 90% quartiles across methods[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	1.00					
HPM	0.38	1.00				
JDM	0.00	0.18	1.00			
MLM	0.77	0.32	0.00	1.00		
CSM	0.67	0.36	0.00	0.83	1.00	
PIM	0.93	0.36	0.00	0.79	0.88	1.00

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of PDF 90% quartiles estimated with the row method is equal to the mean value of PDF 90% quartiles estimated with the column method. The average is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.10. Test of the goodness of fit – mean squared error (MSE)[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.00	0.00	0.06	0.00	0.00
HPM	1.00	–	0.80	1.00	0.66	0.16
JDM	1.00	0.20	–	1.00	0.52	0.16
MLM	0.94	0.00	0.00	–	0.08	0.07
CSM	1.00	0.34	0.48	0.92	–	0.06
PIM	1.00	0.84	0.84	0.93	0.94	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of mean squared errors of the row method is not higher than the mean value of mean squared errors of the column method. The average of the test is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.11. Test of the goodness of fit – mean squared percentage error (MSPE)[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.00	0.00	0.14	0.06	0.00
HPM	1.00	–	0.51	0.94	0.56	0.17
JDM	1.00	0.49	–	1.00	0.48	0.17
MLM	0.86	0.06	0.00	–	0.14	0.10
CSM	0.94	0.44	0.52	0.86	–	0.07
PIM	1.00	0.83	0.83	0.90	0.93	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of mean squared percentage errors of the row method is not higher than the mean value of mean squared percentage errors of the column method. The average of the test is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.12. Test of the goodness of fit – mean absolute error (MAE)[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.00	0.00	0.06	0.00	0.00
HPM	1.00	–	0.78	1.00	0.65	0.10
JDM	1.00	0.22	–	1.00	0.41	0.06
MLM	0.94	0.00	0.00	–	0.08	0.00
CSM	1.00	0.35	0.59	0.92	–	0.47
PIM	1.00	0.90	0.94	1.00	1.00	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of mean absolute errors of the row method is not higher than the mean value of mean absolute errors of the column method. The average of the test is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.13. Test of the goodness of fit – mean absolute percentage error (MAPE)[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.00	0.00	0.16	0.00	0.00
HPM	1.00	–	0.51	0.99	0.53	0.06
JDM	1.00	0.49	–	1.00	0.39	0.09
MLM	0.84	0.01	0.00	–	0.14	0.00
CSM	1.00	0.47	0.61	0.86	–	0.04
PIM	1.00	0.94	0.91	1.00	0.96	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of mean absolute percentage errors of the row method is not higher than the mean value of mean absolute percentage errors of the column method. The average of the test is across currencies and maturities. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.14. Test of the stability of the PDF mean[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.99	1.00	1.00	1.00	1.00
HPM	0.01	–	0.01	0.35	0.87	0.90
JDM	0.00	0.99	–	1.00	1.00	1.00
MLM	0.00	0.65	0.00	–	1.00	1.00
CSM	0.00	0.13	0.00	0.00	–	0.83
PIM	0.00	0.10	0.00	0.00	0.17	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of the PDF mean resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of the PDF mean obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.15. Test of the stability of the PDF variance[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.15	0.75	0.17	0.55	1.00
HPM	0.85	–	0.85	0.85	0.57	0.85
JDM	0.25	0.15	–	0.17	0.08	0.69
MLM	0.83	0.15	0.83	–	0.58	1.00
CSM	0.45	0.43	0.92	0.42	–	1.00
PIM	0.00	0.15	0.31	0.00	0.00	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of the PDF variance resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of the PDF variance obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.16. Test of the stability of the PDF skewness[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.72	0.42	0.46	0.50	1.00
HPM	0.28	–	0.31	0.11	0.36	0.86
JDM	0.58	0.69	–	0.41	0.42	1.00
MLM	0.54	0.89	0.59	–	0.50	1.00
CSM	0.50	0.64	0.58	0.50	–	0.95
PIM	0.00	0.14	0.00	0.00	0.05	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of the PDF skewness resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of the PDF skewness obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.17. Test of the stability of the PDF kurtosis[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.89	0.00	0.89	0.61	1.00
HPM	0.11	–	0.02	0.02	0.11	0.42
JDM	1.00	0.98	–	1.00	1.00	1.00
MLM	0.11	0.98	0.00	–	0.42	0.92
CSM	0.39	0.89	0.00	0.58	–	0.99
PIM	0.00	0.58	0.00	0.08	0.01	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of the PDF kurtosis resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of the PDF kurtosis obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.18. Test of the stability of the PDF 10% quartile[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.99	0.92	1.00	0.75	1.00
HPM	0.01	–	0.18	0.01	0.01	0.01
JDM	0.08	0.82	–	0.08	0.08	1.00
MLM	0.00	0.99	0.92	–	0.41	1.00
CSM	0.25	0.99	0.92	0.59	–	0.92
PIM	0.00	0.99	0.92	0.00	0.08	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of PDF 10% quartiles resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of PDF 10% quartiles obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.19. Test of the stability of the PDF 25% quartile[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.98	0.92	0.83	0.50	1.00
HPM	0.02	–	0.26	0.02	0.02	0.02
JDM	0.08	0.74	–	0.08	0.08	0.08
MLM	0.17	0.98	0.92	–	0.42	1.00
CSM	0.50	0.98	0.92	0.58	–	0.75
PIM	0.00	0.98	0.92	0.00	0.25	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of PDF 25% quartiles resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of PDF 25% quartiles obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.20. Test of the stability of the PDF 50% quartile[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.96	0.92	1.00	0.67	1.00
HPM	0.04	–	0.28	0.04	0.04	0.04
JDM	0.08	0.72	–	0.08	0.08	0.08
MLM	0.00	0.96	0.92	–	0.49	1.00
CSM	0.33	0.96	0.92	0.51	–	0.99
PIM	0.00	0.96	0.92	0.00	0.01	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of PDF 50% quartiles resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of PDF 50% quartiles obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.21. Test of the stability of the PDF 75% quartile[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.98	0.92	1.00	0.58	1.00
HPM	0.02	–	0.25	0.02	0.03	0.03
JDM	0.08	0.75	–	0.08	0.08	0.08
MLM	0.00	0.98	0.92	–	0.42	1.00
CSM	0.42	0.97	0.92	0.58	–	1.00
PIM	0.00	0.97	0.92	0.00	0.00	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of PDF 75% quartiles resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of PDF 75% quartiles obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.

Table B.22. Test of the stability of the PDF 90% quartile[1]

	QIM	HPM	JDM	MLM	CSM	PIM
QIM	–	0.96	0.92	0.91	0.47	1.00
HPM	0.04	–	0.19	0.04	0.04	0.04
JDM	0.08	0.81	–	0.08	0.08	0.08
MLM	0.09	0.96	0.92	–	0.38	1.00
CSM	0.53	0.96	0.92	0.62	–	0.67
PIM	0.00	0.96	0.92	0.00	0.33	–

[1] The table presents average probability values of the *t*-test. The null hypothesis of the test is that the mean value of cross-sectional standard deviations of PDF 90% quartiles resulted from randomly perturbed implied volatilities and estimated with the row method is not higher than the mean value of cross-sectional standard deviations of PDF 90% quartiles obtained with the column method. The average of the test is across currencies for PDFs with 1-month time to expiry. The acronym names of the methods are as follows: QIM – Quadratic interpolation method (Shimko (1993)); HPM – Hermite polynomials method (Madan and Milne (1994)); JDM – Jump-diffusion method (Malz (1996)); MLM – Mixture of lognormals method (Melick and Thomas (1997)); CSM – Cubic spline method (Bliss and Panigirtzoglou (2002)) ; PIM – Piecewise cubic Hermite interpolation method.