FINANCIAL OPTIONS RESEARCH CENTRE

University of Warwick

Stochastic Volatility Models With Jumps: Implications for Smiles in Foreign Exchange Markets

Robert G Tompkins

September 2000

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

FORC Preprint: 2000/106

STOCHASTIC VOLATILITY MODELS WITH JUMPS: IMPLICATIONS FOR SMILES IN FOREIGN EXCHANGE MARKETS†

Robert G. Tompkins
University Dozent, Department of Finance
Vienna University of Technology* and
Department of Finance, Institute for Advanced Studies

† This piece of research was partially supported by the Austrian Science Foundation (FWF) under grant SFB#10 (Adaptive Information Systems and Modelling in Economics and Management Science). This paper has benefited from joint work with Friedrich Hubalek, and comments by Walter Schachermayer, Ole Barndorff-Nielsen, Neil Shephard and attendees at the University of Technology, Vienna symposium on Stochastic Volatility and Lévy Processes. As always, I am responsible for all remaining errors.

* Doeltergasse 33, A-1220 Wien, Austria, Phone: +43-1-726-0919, Fax: +43-1-729-6753,
Email: rtompkins@ins.at

STOCHASTIC VOLATILITY MODELS WITH JUMPS:
IMPLICATIONS FOR SMILES IN FOREIGN EXCHANGE MARKETS

ABSTRACT

A convenient starting point for option pricing has been to assume that markets conform to Geometric Brownian Motion (GBM). Extensive empirical evidence has shown that Foreign Exchange markets reject this assumption. In addition, the prices of options on these markets systematically diverge from prices assumed by the Black-Scholes (1973), Garman-Kohlhagen (1983) and Black (1976) models. Patterns of implied volatilities for these markets display the now familiar convex shape as a function of striking prices, which is commonly known as the smile.

This paper examines whether alternative models that include jumps and stochastic volatility can account for the empirical dynamics of Foreign Exchange Futures and explain the implied volatility smiles for options on Foreign Exchange markets.

In this paper, we consider a rich class of stochastic volatility models that include jumps and correlated processes. This work extends Bates (1996b) by considering more markets over a longer period of analysis and alternative jump process models. Two alternative Jump models are considered which include jumps either in the underlying asset price or the stochastic volatility processes.

For the four Foreign Exchange futures markets examined, either class of jump models can explain the empirical features of these markets. For both types of jump models, prices of European options are determined. Comparisons of simulated and actual options prices for these markets find substantial differences. These results concur with previous studies that the inclusion of a model consistent with the objective process alone is insufficient to explain the existence of smiles. This points to the existence of a risk premium when pricing options on Foreign Exchange. This research identifies the dynamics of this risk premium across strike price and time and finds consistencies of this risk premium for all four Foreign Exchange markets. Such consistency suggests market agents are assuming a similar functional form for the risk premium.

JEL classifications: C15, G13

Keywords: Stochastic Volatility, Normal Inverse Gaussian Distributions, Methods of Moments Estimation, Implied Volatility Smiles.

1. INTRODUCTION

Soon after the publication of the Black-Scholes (1973) paper on option pricing, Black (1975) pointed out that the constant volatility assumption for stocks may be incorrect; noting that the volatility may be a function of the underlying price level. Black (1976) and Garman & Kohlhagen (1983) showed that the Black-Scholes (1973) formula could be modified to price options on futures contracts and foreign exchange, respectively. Subsequent research for both futures options and foreign exchange options has identified the existence of different volatilities implied from option prices for different strike prices and terms to maturity [see Ben Khelifa (1991), Cao (1992) and Shastri & Wethyavivorin (1987) for evidence of smile effects for foreign exchange options]. This effect has been commonly referred to as the implied volatility smile (for options with the same term to expiration) and the term structure of volatility (for options with different terms to expiration).

Broadly speaking, two possible reasons have been proposed to explain these effects. The first approach assumes that market imperfections exist that systematically prevent option prices from taking their true Black-Scholes values. Such market imperfections include the introduction of errors associated with discrete hedging, transactions costs, incomplete markets and heterogeneous market agents with diverse expectations. Alternatively, a number of papers have examined the implications for option prices when the underlying asset price process differs from the lognormal diffusion process and/or the volatility is neither constant [as in Black-Scholes (1973)] nor a deterministic function [as in Merton (1973)]. Mayhew (1995) provides a brief review of both approaches. This research examines the latter hypothesis.

This paper examines the nature of the objective dispersion processes for foreign exchange futures and proposes the inclusion of jump processes and stochastic volatility to explain these empirical facets of the objective price processes. Once this is done, the implications for option pricing are considered.

Previous research has included jumps in the underlying price process. Bodurtha and Courtadon (1987), Jorion (1988) and Vlaar and Palm (1993) have suggested that the excess kurtosis they observed in foreign exchange returns could be explained by the presence of jumps. Merton (1976) first examined the inclusion of jumps in option price by using Poisson processes. Subsequent research has examined alternative stationary fat-tailed distributions to Geometric Brownian Motion (GBM) as a proxy for jumps to both explain the time series properties of foreign exchange markets and consider the implications for pricing options. A non-trivial issue is the choice of the fatter tailed distribution, which will mimic jump processes. Rogalski and Vinso (1978), Nelson (1991), Liu & Brorsen (1995), and Ho, Perraudin & Sørensen (1996) have previously examined a symmetric Student-t distribution. However, this is not always flexible enough in either the tails and/or the centre of the distribution. In addition, it cannot account for the skewness often found in the innovation processes [Mittnik and Rachev (1993)]. Within the family of generalised exponential distributions (GED) and t distributions, a wide variety of potential candidates exist. For symmetrical distributions, the GED nests both the normal and Laplace distributions. McDonald and Newey (1988) proposed the generalised Student's t, which was suggested for use in conjunction with GARCH models by Bollerslev, Engle and Nelson (1994). For the estimation of asymmetric fat-tailed distributions, Steel (1998) proposed the skewed exponential power distribution and Theodossiou (1998) used a skewed Generalised Student t distribution.

Outside the GED-t family, a number of alternative distributions have been proposed. McDonald and Xu (1995) examined a Generalised Logistic distribution (they refer to as EGBII). Symmetric and asymmetric Stable Paretian distributions have been proposed since Fama (1963, 1965) [also see Westerfield (1977) for analysis of foreign exchange markets] but these have generally been dismissed, as the variance is not finite. Recently, the family of generalised hyperbolic distributions introduced by Barndorff-Nielsen (1977,1978) has been considered.

Recent studies examining the unconditional return processes for European equity markets [see Eberlein and Keller (1995)] and for US equity markets [see Rydberg (1997)] show that hyperbolic distributions fit the data well. Nested within this class of distributions is the Normal Inverse Gaussian (NIG) distribution. Barndorff-Nielsen (1995) introduced this distribution as a possible model for financial returns data and Blæsild (1995) showed that the NIG distribution fit German stock returns better than the generalised hyperbolic distribution. Rydberg (1997) also examined this distribution and provided a simple approximation technique for simulation purposes. This research will examine this class of distributions for Foreign Exchange markets for the first time.

Another line of research has examined the effects of stochastic volatility (in the underlying objective price process) on option pricing. Many models assume some mean reverting process with the disturbances driven by a Wiener Process. Examples of such models include Hull and White (1987), Bailey and Stulz (1989), Johnson and Shanno (1987), Scott (1987,1991) and Wiggins (1987), Chesney and Scott (1989) and more recently Stein and Stein (1991) and Heston (1993). Taylor (1994) probably provides the best survey of this work. Bates (1996b) chose to model the variance of Foreign Exchange futures using a square root process similar in spirit to Heston (1993). In this research we will also examine the Heston (1993) model as the subordinated volatility process.

Both lines of research have been able to describe key empirical facets of either the objective time series behaviour of Foreign Exchange futures or of options prices. However, neither can address all aspects of these markets. For example, stochastic volatility models can address many of the volatility clustering behaviours but fail to explain the higher moments of the log price increments. Non-normal distributions, which are independent and identically distributed (i.i.d.) fail to capture time series facets. To address these problems, another class of models combines both models. Examples of these include the Student's *t*/GARCH (stochastic volatility) model of Baillie and Bollerslev (1989) and the jump diffusion/ARCH model of Jorion (1988). The work most closely related to this research is that of Bates (1996b) that considered stochastic volatility and jump diffusion processes under systematic jump and volatility risk. This paper extends this line of research.

Recently, a number of researchers have examined stochastic volatility models with the disturbances drawn from non-normal price processes. Essentially, this allowed the jumps to impact the price process via the volatility process. Herzel (1998), Perrod (1999), Duffie, Pan & Singleton (1999) have all considered this. Barndorff-Nielsen & Shephard (2001) have proposed a model class where the volatility is an Ornstein-Uhlenbeck type process driven by a Lévy process and a standard Wiener process drives the underlying asset.

The objective of this paper is to compare and contrast three alternative models for the unconditional price processes of Foreign Exchange Futures. The first and second models assume Heston (1993) describes the volatility process with the first model assuming that a Wiener process drives the underlying asset and the second model assumes a NIG Distribution. Tompkins (2000) previously examined this latter model for Stock Index Futures markets. The third model is that of Barndorff-Nielsen & Shephard (2001). Of additional interest are the implications for option pricing, with the ultimate goal of better understanding implied volatility surfaces for traded options.

The paper is organised as follows: the first part briefly reviews previously presented empirical evidence, which indicates that Foreign Exchange futures prices do not follow an GBM process. The models examined will then be introduced. Seven empirical attributes were selected to capture key aspects of non-normality and inter-dependence. A description of the data sources used for this research follows. After this, the alternative models were parameterised using a simulated method of moments approach. Optimal parameters for each model minimised the sum of squared errors between empirical and simulated attributes. Using this model and the optimal parameters, option prices were estimated consistent with these processes. These were then expressed as implied volatilities using the Black (1976) formula. These simulated implied volatility surfaces were compared to the implied volatility surfaces from traded options on these same Foreign Exchange futures markets (for the same period of analysis). Finally, conclusions and suggestions for further research appear.

2. PRICE PROCESSES FOR FOREIGN EXCHANGE FUTURES UNDER THE 'EMPIRICAL' MEASURE

It is well established that the unconditional return series for Foreign Exchange markets do not conform to the assumptions of an i.i.d. lognormal dispersion process. Returns for foreign exchange futures usually display excess kurtosis and (for many periods) significantly negative skewness compared to a normal distribution [see Hsieh (1988), Meese (1986) and Bates (1996b)]. Returns for futures also display inter-temporal dependence. In the examination of absolute returns for Foreign exchange markets, significant positive autocorrelations have been found. This result has led Ding, Granger and Engle (1993) to conclude, "It is clear that the ... market return process is not an i.i.d. process" (page 87).

Another line of empirical research has examined the unconditional volatility series directly. It is undisputed that the unconditional volatility process is time varying and has been extensively examined in the ARCH and GARCH frameworks [see Bollerslev, Chou and Kroner (1992) for a survey of the ARCH/GARCH literature with specific consideration of foreign exchange rates]. Given the wide varieties of such models proposed in the literature and the lack of agreement on the appropriate ARCH/GARCH model for foreign exchange, we will examine the statistical moments of the volatility process directly. This alternative approach to capturing key aspects of the volatility process was first proposed by Burghardt and Lane (1990) who examined the variability of the unconditional volatility process using a volatility cone approach. We extended this by looking at the sampling properties (restricted to the standard deviation) of the unconditional volatility measured at a 20-day time horizon. Using non-overlapping data, we obtained an average estimate of the unconditional volatility and the standard deviation of this average. Under the assumption that an i.i.d. price process is generating these volatilities, the expected coefficient of variation of the 20-day volatility is known. This measures the variability of volatility for a given time horizon.

However, the time varying dynamics of the variability of volatility as the time horizon of estimation is extended, is of additional interest. To yield sufficient observations to determine statistical moments, overlapping data was used. Biases introduced by overlapping estimation were corrected using Hodges and Tompkins (2000). A simple log-linear form was chosen to capture these dynamics. The rate of decay in the standard deviation of volatility implies that the maturity structure of historical volatility experiences long-term memory (interdependence). This can be seen as complimentary to long-term memory effects for absolute returns identified by Ding, Granger and Engle (1993). An additional and important feature of these price processes is the existence of correlations between levels of volatility and the underlying foreign exchange rate. For stock and stock index markets, this correlation has been coined the leverage effect and has been pointed out by Christie (1982) among others.

While a number of theories have been proposed to explain these results, three alternative Stochastic Volatility Models were considered. Given the wide range of stochastic volatility models proposed in the literature [see Taylor (1994)], it was not obvious which model to select. As we consider correlated processes, the Heston (1993) model was the obvious choice. This model has the additional benefit of having a closed form solution for the pricing of options [see Heston (1993), Bakshi, Cao and Chen (1997) and Bates (2000)]. Two additional classes of stochastic volatility models were also examined which include jumps. The first model examined is a variant of the Heston (1993) model proposed by Tompkins (2000), which includes jumps (captured by a NIG process) in the underlying price process. The final model examined is the OU volatility process with innovations driven by an Inverse Gaussian process. Barndorff-Nielsen and Shephard (2001) have proposed this class of models. For the sake of comparison, the standard Wiener process with constant variance will also be examined.

These four cases can be expressed as:

GBM CASE $$dF = \mu F(t)dt + \sigma F(t)dZ(t) \qquad \textbf{(0)}$$

Where Z(t) is a standard Wiener Process and μ and σ are constants, the return series r_t is normally distributed with $r_t = (\mu - \frac{1}{2}\sigma^2)t + \sigma Z_t(t)$ and $Z_t \sim$ N(0,1). This is the assumption of the Black-Scholes (1973), Garman-Kohlhagen (1983) and Black (1976) models of i.i.d. Geometric Brownian Motion and will be referred to as GBM.

It is clear that the GBM case is a straw man, given the well-known results that foreign exchange futures returns display both non-normality and inter-dependence. However, this can serve as a benchmark for the relative effectiveness of the alternative models and is used to assess the sampling properties of the evaluation approach. Subsequent models will consider stochastic volatility ($\hat{\sigma}$) which will be evaluated in terms of a stochastic variance process ($\sqrt{V}$).

MODEL 1 $$dF(t) = \mu F(t)dt + \hat{\sigma} F(t)dZ_1(t) \quad (1.1)$$

With the variance process defined by:

$$dV(t) = \kappa(\theta - V(t))dt + \xi\sqrt{V(t)}dZ_2(t) \quad (1.2)$$

Where Z_1 and Z_2 are standard Wiener processes with correlation ρ. The term κ indicates the rate of mean reversion of the variance, θ is the long-term variance and ξ indicates the volatility of the variance. The terms V and √V represent the variance and the volatility of the process, respectively. This model will be referred to as SV in this paper.

MODEL 2 $$dF(t) = \mu F(t-)dt + \hat{\sigma}(t)F(t-)dN(t) \quad (2.1)$$

With the variance process defined by:

$$dV(t) = k(\theta - V(t))dt + \xi\sqrt{V(t)}dZ(t) \quad (2.2)$$

where N(t) is a purely discontinuous martingale corresponding to log-returns driven by a NIG Lévy process.[1] Tompkins (2000) considered an extension of this model where Z is a linear combination of a Wiener process and the NIG Lévy process driving the price process, thus allowing dependence between the two processes. As the NIG Lévy process nests the Wiener process as a limiting case, this models nests both Model 1 [Heston (1993)] and by taking $\kappa=\xi=0$, the model proposed by Barndorff-Nielsen (1995). who was the first to propose a stochastic volatility model of this form [subsequently extended by Andersson (1999a)]. This approach extends the findings of Bates (1996, 2000) and Ho, Perraudin and Sørensen (1996), who assumed the volatility process is subordinated in a non-normal price process. When correlated processes are considered, this variant of model 2 is related to the model for incorporating a leverage effect in a stochastic volatility model proposed by Barndorff-Nielsen and Shephard (1998). This model will be referred to a NIGSV in this paper.[2]

Barndorff-Nielsen and Shephard (2001) propose the final model considered here. This can be expressed as:

MODEL 3 $$dF(t) = \mu F(t-)dt + \hat{\sigma}(t-)F(t-)dZ(t) + dM(\lambda t) \quad (3.1)$$

With the variance process defined by:

$$dV(t) = -\lambda V(t)dt + d\grave{Z}(\lambda t) \quad (3.2)$$

The process $\grave{Z}$ called the background driving Lévy process (BDLP) is an *increasing* Lévy process, independent of Z, the process M is a purely discontinuous martingale resulting from the incorporation of a leverage effect in that model [see Tompkins and Hubalek (2000)].

A useful (and somewhat more flexible) extension suggested in Barndorff-Nielsen and Shephard (2001), is to replace the variance process (3.2) by a superposition of k independent components $V_1, \ldots, V_k$ and use:

$$dF(t) = \mu F(t-)dt + \hat{\sigma}(t-)F(t-)dZ(t) + \sum_{i=1}^{k} dM(\lambda_i t) \quad (3.3)$$

$$dV(t) = \sum_{i=1}^{k} w_i dV_i(t) \quad (3.4)$$

$$dV_i(t) = -\lambda_i V_i(t)dt + d\grave{Z}_i(\lambda_i t) \quad (3.5)$$

Barndorff-Nielsen and Shephard (2001) introduce not a particular model, but rather a general framework, which is flexible and analytically very tractable. They also suggest various concrete specifications for the Lévy process Z and the corresponding distributions. Among those is the BDLP that yields a stationary Gamma distribution for the variance. The superposition of two variance processes with stationary Gamma distributions is investigated by Hubalek and Tompkins (2000) yielding a closed form solution for this type of OU variance process. This model will be referred to as Γ2-OU in this paper.

3. CHOICE OF ATTRIBUTES AND FITTING PARAMETER VALUES

A key problem in the empirical testing of stochastic volatility models is the estimation of optimal input parameters into the model. Andersen, Chung and Sørensen (1999) provide a good review of the approaches used to parameterise stochastic volatility and jump models. Many of the models considered here do not lend themselves to estimation by traditional maximum likelihood methods. Andersen, Chung and Sørensen (1999) propose the use of the efficient method of moments approach suggested by Gallant and Tauchen (1996). Their primary contributions are to understand the sample properties of this estimator and show that the method is robust in larger sample sizes. Recently, Andersson (1999b) also examined the maximum likelihood estimator of the NIGSV model of Barndorff-Nielsen (1997) and by Monte Carlo simulation demonstrated that a simulated method of moments approach is equally robust. Due to the complexities of our models, it is not clear whether estimation of parameter inputs by maximum likelihood techniques is feasible. Given that these papers have demonstrated that parameter estimation via simulation can be equally robust and this approach may provide a better intuitive understanding of the estimation procedure, a simulated method of moments approach similar in spirit to General Method of Moments (GMM) was utilised.

This research uses a hybrid between the Generalised Method of Moments (GMM) approach of Melino and Turnbull (1990) and the Simulated Method of Moments (SMM) approach of Duffie and Singleton (1993). This approach subjectively selects essential attributes, simulates price processes consistent with the alternative models and assesses the sum of squared errors between simulated and empirical attributes. Alternative parameterisation of the models was examined and optimised. Similar to Andersen, Chung and Sørensen (1999), we investigated the sample properties of this estimator technique (using the GBM case). This allowed comparisons to be made between the attributes of each market and the models and conclusions to be drawn regarding the overall fit of each model.

At the heart of this estimation technique is the judicious choice of key attributes. It is crucial that the choice of the attributes jointly considers the relevant elements of empirical interdependence and non-normality and provides a means by which the salient features of the alternative models can be captured. Given that both alternative price process (to GBM) and stochastic volatility models have been proposed to explain excess kurtosis in returns, the unconditional kurtosis for daily returns is a logical attribute to choose. Furthermore, some research has indicated that skewness has been an important attribute for describing the returns of foreign exchange markets. Both Bodurtha and Courtadon (1987) and Bates (1996a) chose to examine the skewness in addition to the excess kurtosis. However, these effects were relevant only for periods from 1983-1986. As this study will examine foreign exchange futures and options over this period and will extend analysis to the millennium, skewness will be examined. This attribute would provide evidence for the existence of asymmetric jumps.

To examine both of these moments, price changes were estimated based upon continuously compounded returns in the standard manner (using log-price increments). Extreme care was taken to assure that returns were estimated using only futures prices with the same expiration date. With this time series of daily returns, the unconditional skewness and kurtosis statistics were determined.

Clearly, given that this research examines three alternative stochastic volatility models, attributes had to be selected that would allow the salient features of these models to be captured. This required attributes capturing the volatility of volatility parameter, the rate of mean reversion and the correlation between the processes. For this research, the estimation of the standard deviation of the returns was determined using squared returns and these results were annualised (assuming 252 trading days in a year). A number of attributes were selected in order to capture critical dynamics of the volatility process.

The first attribute examines the volatility of volatility. A time series of unconditional volatilities was estimated on a daily basis for a time horizon of 20 days (using non-overlapping data). From this series, the average and standard deviation were estimated. Given widely different levels of these sample statistics, a coefficient of variation statistic was chosen as the key attribute. As a basis for comparison, the coefficient of variation of volatility measured at the 20th lag would be approximately equal to 0.1622 if the underlying return series were independent and identically distributed [1/√(2*19)].

While this measured the variability of volatility at a single time horizon, the time varying dynamics of the standard deviation of volatility could not be captured. This was achieved by the estimation of the coefficient of variation of volatility with time horizons from 20 days to 200 days in 20-day increments.[3] From statistical theory, the expected decay in the standard deviation of a volatility estimate (σ^*) is a square root function of the number of observations used for estimation ($SE = \sigma^* / \sqrt{2N}$). Given that the average level of volatility remains constant (which was found empirically for these four markets), the decay in the coefficient of variation should follow the same functional form. If the first volatility observation is at the 20-day horizon, the decay in the standard deviation of volatility from that point forward can be expressed as $\sigma_N \cdot \sqrt{\frac{N}{N+1}}$. In this formula, σ_N is the standard deviation of the volatility at the Nth observation and N+1 is the number of observations in the next time horizon (initially N=20 and N+1=40). The functional form of this decay can be expressed as a power function of the form $SE = \frac{\sigma^*}{\sqrt{2}} \times N^{-0.5}$. A linear regression of (the natural logarithms of) the time horizon of estimation regressed upon the levels of the coefficient of variation was used to capture the empirical rate in decay. This can be expressed as:

$$\ln(\hat{\sigma}_N) = \alpha + \beta(\ln(N)). \quad \textbf{(4.1)}$$

From this regression equation, the decay attribute, follows the following exponential form:

$$e^{\alpha+\beta(\ln(N))} \quad \textbf{(4.2)}$$

Observed divergences in the Beta coefficient of equation 4.2 from -0.5 give an indication of the degree of the volatility of volatility persistence observed in the unconditional process. This also provides an attribute to capture the interaction between the rate of mean reversion and volatility of volatility in the stochastic volatility process.

Another salient feature of empirical return volatility series is that subsequent realisations may not be independent. To capture serial correlation in absolute returns, autocorrelation dynamics were examined directly rather than using alternative methods relying on maximum likelihood. This was achieved by examining the autocorrelograms of absolute returns previously employed by Taylor (1986) and Ding, Granger and Engle (1993). This approach has the additional benefit of known sampling properties. Thus, a simple confidence interval test can be used to reject the null hypothesis of independence.

For the purposes of this research, composite measures of the autocorrelations are required. Given that the markets differ in the manner that the autocorrelations decay, the averages of the autocorrelations from lag 1 to 20 and from lag 51 to 70 were both selected. The first average represents the short-term autocorrelation. The medium-term average provides an indication of how quickly the autocorrelations die out. Unfortunately, such measures may no longer have known sampling properties allowing for a simple parametric confidence interval test. Therefore, to assess the sample characteristics of these composite measures, nonparametric confidence intervals were determined via simulation. These two attributes provide additional information relevant to the calibration of a stochastic volatility model, as they capture both short-term and medium-term evidence of volatility clustering.

The final attribute must measure the leverage effect and provide a means for a correlation between the stochastic processes to be captured. There is, however, one problem with the determination of this effect: Even if the volatility is a stationary series, the prices are not. To solve this problem, a new variable was constructed which measures recent price movements and is stationary.[4] This variable is an exponentially weighted return series, which indicates whether recent price movements are relatively high or low. It can be shown that given some exponential weighting scheme: where $W = 1 - \theta$ and $\theta \approx e^{-W\Delta t}$, we can define a new series, ω_i, that can be expressed as:

$$\omega_i = \omega_{i-1} + W(r_{i-1} - \omega_{i-1}) \quad (5)$$

where ω_i is the exponentially weighted price movement, r_i is the daily return and the initial ω_i is set to zero. W represents the weight used in the weighting scheme. This new variable was created and the series of 20-day unconditional volatility were compared. With an arbitrary weight (W) for all markets imposed at 0.03, the correlation between the two variables was estimated.[5] This correlation coefficient serves as an attribute to measure the correlation between volatility and price processes.

With these seven target attributes, the stochastic volatility models were parameterised via simulation. Following the three models proposed previously, price series of 1500 observations were generated consistent with these models. The return and volatility characteristics of these simulated price series were estimated in exactly the same manner as was done for the four foreign exchange futures. The resulting simulated attributes were then compared to the empirical attributes using a sum of squared errors statistic. To reduce scaling impacts due to different levels of the attributes, the squared errors were divided by the standard deviation of the attributes across the four markets[6]. This test statistic can be written as:

$$\min \sum \left(\frac{M_i - X_i}{\sigma_i} \right)^2 \quad \textbf{(6)}$$

where M_i is the attribute for the foreign exchange futures market, X_i is the attribute of the price series generated by the model and σ_i is the standard deviation of the attributes across all the foreign exchange futures markets for the relevant period of analysis.

Finally, 500 samples of 1500 prices consistent with the GBM straw man model were drawn to better understand the sample properties of all the attributes and of the test statistic. This provided a non-parametric estimation of the standard errors of the attributes and allowed test statistics for comparisons between markets and models.

4. DATA SOURCES

The futures markets examined include the Deutsche Mark, Japanese Yen, British Pound and Swiss Franc futures (all versus the US Dollar) traded at the Chicago Mercantile Exchange (CME). Daily closing futures prices were analysed and were obtained directly from the relevant exchange and the period of analysis was restricted to a period in which both futures and options for all four markets were traded. The underlying assets, the time period of analysis and the number of observations used in the analysis appear in Table I. To assess if results are period specific, the data was also split into roughly three equal time periods from 1985-1990, 1990-1995 and from 1995-2000. The time periods and number of observations in these split periods also appear in Table I.

[Table I appears here]

Given that this portion of the research is empirical in nature, a major effort was made to assure the validity of the data used in the analysis and to verify that the analytic methods employed were correct. This was achieved in a number of ways. Firstly, the futures price series were compared with the options (on the futures) price series for the same days to identify obvious errors in recording either price series. This comparison was achieved by comparing the put-call parity values of the options with the underlying futures prices for every single date in our database (and for all four markets). A screening procedure was imposed: If futures or options prices diverged by more than the normal bid/offer spread (of one tick), the observations were flagged. Once this was done, each price was compared with the original daily price sheets to confirm if a 'keypunch' error had occurred. We discovered that only 1-2% of the data had such errors. Nevertheless, these errors were of a sufficient magnitude that they did influence the results and therefore required correction.

The most arduous of the data cleaning process was the ongoing examination of the data as results of the analysis were obtained. One reason why four foreign exchange futures and options markets were examined was to allow a cross-sectional comparison. Apart from the benefit of assessing general tendencies across markets, it was also possible to use anomalous results as an additional check on data validity. This assured that the data series employed in this research were as accurate as possible.

5. ATTRIBUTES FOR INDIVIDUAL MARKETS

For each market, the return statistics for daily returns were determined for the entire period of analysis and for each sub-period. The results of these analyses can be seen in Table II.

[Table II appears here]

The first column describes the market under investigation and indicates the time period examined. The second and third columns present the mean and standard deviation of the return distribution. The fourth and fifth columns present the statistics for the unconditional skewness and these provide a significance level relative to a null hypothesis of normality.[7] All skewness significance statistics that are significantly different from the normal assumption at a 95% level (±1.96) appear in bold type. In the sixth and seventh columns, the unconditional kurtosis statistic and the significance level relative to a null hypothesis of normality appear.[8] When this statistic does not reject the null hypothesis of normality at the 95% level or above, the statistic and the significance levels are in bolded type. The statistic in the eighth column is the Bera-Jarque (BJ) statistic for detecting departures of the data from normality. Under the null hypothesis of normality, the BJ statistic is distributed as χ^2 with 2 degrees of freedom. The critical value at the one-percent level is 9.21. When the BJ statistic exceeds this level, this statistic also appears in bolded type.

For all four markets, the dispersion of returns is not well described by a normal distribution. For many of the four markets, the skewness statistic tends to be significantly different than that of a normal distribution. In almost all instances, when the skewness is significantly, it is positive. On the other hand, for all four markets (and for all time periods), the daily returns always display significant excess kurtosis. These two factors lead to all BJ values exceeding their critical values. These results are consistent with previous empirical examination of return series for equity markets by Theodossiou (1998) and Harvey & Siddique (1998) and were pointed out for foreign exchange markets by Bates (1994, 1996).

As the remaining attributes all capture characteristics of the volatility process, these will be summarised in a single table, Table III.

[Table III appears here]

In this table, the first column describes the individual market examined and the time period of the analysis. The next three columns display the average annualised volatility measured at a 20-day time horizon. At the bottom of these columns are the expected attributes from a GBM dispersion process with constant variance. The attribute of interest to this research is the Coefficient of Variation statistic. For all four markets and for all time periods, we can compare this measure of the volatility of volatility to what would be expected under the GBM assumption. By determining a standard error of this attribute by simulation, we can reject the hypothesis that the volatility process conforms to the GBM assumption.

In the fifth column appears the beta of the regression of the relationship between the [natural logarithms of the] time horizon of the estimation period against the coefficient of variation of volatility. If markets conform to a GBM i.i.d. process, a decay coefficient of -.50 (seen at the bottom of the column) would be observed. For almost all markets (the exception is for the Japanese Yen for the 1985-1990 period), the rate of decay is (statistically) significantly less that this decay function (the standard error of this attribute is estimated in a non-parametric manner by simulation). In Column six, the leverage correlation coefficient appears. This measures the relationship between the 20-day unconditional volatility and the recent relative prices. For ten of the sixteen comparisons, this leverage effect is statistically significant. In most instances, this is a positive relationship. One interpretation of this is that the higher the level of recent prices the higher the level of the volatility. A comparison between the unconditional skewness in Table II with the leverage correlation in Table III suggests that some similar mechanism is at work. The nature of this mechanism is unknown and remains for future research. These results should be interpreted with care given that the simulated standard error of this attribute is fairly high (at 0.0999). While these effects remain significant, it is only (barely) at the 95% level.

The final two columns represent the average autocorrelations of absolute returns for lagged periods from 1-20 days and 51-70 days. Underlying the assumption of an i.i.d. price process, these have a prior expectation of zero. For all four markets over the shortest time period (1-20 days), these averaged autocorrelations are statistically significantly positive. For the longer time horizon (51-70 days), the autocorrelations tend to be significantly positive (apart from the D-mark, Japanese Yen and Swiss Franc futures in the 1985-1990 period).

The results from both tables II and III indicate that both the return series and the volatility process are significantly divergent from a prior assumption of a GBM process. With these attributes as target conditions, the proposed alternative models were examined.

6. FITTING ALTERNATIVE MODELS

In previous sections the models to be tested were presented. It only remains to discuss technical issues in the simulation process and the parameterisation of the stochastic volatility models before proceeding directly to the results.

6.1 SMALL SAMPLING PERFORMANCE SIMULATION FOR THE ESTIMATION METHOD

The simulated method of moment's approach was done in a two-stage process. The first stage was to determine representative distributions for the later simulations. To assess the sampling performance of the estimation method 500 samples of 1500 random variates from a standard normal distribution were drawn using the Box-Müller method. With these 500 samples, price series were constructed that conformed to daily observed GBM. The distributional and time series attributes of each series were computed and compared to the theoretical values. Utilising equation (6), the sum of squared errors between each of the 500 samples and the theoretical attributes of a GBM process were determined. Table IV details the results of the simulated GBM price series and provides the sample standard deviation of the 500 simulated series.
[Table IV appears here]

In this table, the theoretical attribute values for a GBM process are listed as are the average attribute values and the standard deviations of the attributes across the 500 simulations. Of crucial interest are the sampling properties of the attributes and especially the sum of squared errors (SSE) statistic. The standard deviation of this statistic was found to be equal to 4.5646 and will be used subsequently as a means to establish confidence intervals for the comparison of the alternative models. Finally, the characteristics of the two representative draws of the GBM process appear in the bottom two rows.

6.2 Parameter Estimation of Heston (1993) and Tompkins (2000)

Let us recall the estimation of Model 1 and 2 from Tompkins (2000) briefly. Two independent sequences (ε_i) and $(\varepsilon_i^{\perp})$ of iid standard normal random numbers were generated. Another iid sequence of normals (ε_i') with $corr(\varepsilon_i, \varepsilon_i') = \rho$ was generated by

$$\varepsilon_i' = \rho\varepsilon_i + \sqrt{1-\rho^2}\,\varepsilon_i^{\perp} \tag{7}$$

To reduce errors introduced by the selection of random numbers the same sequences were used as innovations for all simulations. A discrete time approximation was employed, namely

$$F_i = F_0 \exp(X_i), \qquad \sigma_i = \sqrt{V_i} \tag{8}$$

with an Euler approximation for log-returns as:

$$\Delta X_i = (\mu - \tfrac{1}{2}V_{i-1}) + \sqrt{V_{i-1}}\sqrt{\Delta t}\,\varepsilon_i \tag{9}$$

and variance [9]

$$\Delta V_i = \kappa(\theta - V_{i-1})\Delta t + \xi\sqrt{V_{i-1}}\sqrt{\Delta t}\,\varepsilon_i' \tag{10}$$

To simulate the NIG random numbers required for Model 2 Tompkins (2000) followed Rydberg (1997, Section 4), which uses the mean-variance mixture representation of the NIG distribution and the multiple root transformation method of Michael et al. (1976), see also Atkinson (1982) and Shuster (1968). He decided not to optimise his model over all dynamic and distributional parameters, but rather to choose from four representative NIG distributions as potential candidates, which appear in Table V as NIG#1 to NIG#4. The four parameters of the $NIG(u, \delta, \alpha, \beta)$ distribution have been chosen to get zero mean and unit variance, as location and scale are governed by the other ('dynamic') model parameters, and, to match empirical skewness and kurtosis[10]. In this way four iid sequences $(\varsigma_i^{(1)}),\ldots,(\varsigma_i^{(4)})$ from the distributions NIG#1 to NIG#4, independent of the normal draws, were generated.

[Table V appears here]

To address correlation, Tompkins (2000) mimicked the standard construction for Gaussian variables (7), but taking NIG and Wiener process increments to generate the sequence (ς_i') by: [11]

$$\varsigma_i' = \rho\varsigma_i + \sqrt{1-\rho^2}\,\varepsilon_i \tag{11}$$

with ς_i equal to one of $\varsigma_i^{(1)}),\ldots,\varsigma_i^{(4)}$. This gives again $corr(\varepsilon_i, \varsigma_i') = \rho$. Tompkins (2000) used the Euler approximation:

$$\begin{aligned} \Delta X_i &= (\mu - \psi(\sqrt{\sigma_{i-1}}\sqrt{\Delta t})) + \sqrt{V_{i-1}}\sqrt{\Delta t}\,\varsigma_i, \\ \Delta V_i &= \kappa(\theta - V_{i-1})\Delta t + \xi\sqrt{V_{i-1}}\sqrt{\Delta t}\,\varsigma_i' \end{aligned} \tag{12}$$

for his stochastic volatility model. Here ψ is the cumulant function for the NIG distribution given in Footnote 2. For approximate simulation of SDEs driven by Lévy processes see Protter & Talay (1997).

6.3 Parameter estimation of Barndorff-Nielsen & Shephard (1999)

Barndorff-Nielsen & Shephard (2001, Section 2.5) describe a general method to simulate *exact* continuous time increments for their model, based on the methodology for simulating Lévy processes by a random series involving the tail of the Lévy measure, based on Marcus (1987), Rosinski (1991), see also Barndorff-Nielsen & Shephard (2000). This method covers Lévy processes and their stochastic integrals with infinitely many jumps or even infinite variation Lévy processes. In the simple case of a compound Poisson process the random series terminates after a random, but finite number of terms and the method coincides with the intuitive method of simulating the individual jump times and jump sizes. For our application we have to simulate the arrival times of a standard Poisson process, that is, the sum of independent exponential random variables, and the jumps sizes, which are again, exponentially distributed. The exponential random numbers are computed from uniform random numbers by the inverse distribution function, which is simply the logarithm. In more detail: Two fixed sequences of standard exponential random variables (γ_i) and (η_i) are generated that are used for interarrival jump times and jump heights[12]. From the standard exponentials we generate the jump times (τ_i) and jumps for the processes $(z(\lambda t))$ and $(V(t))$. We have $\tau_i = \tau_{i-1} + \gamma_i / \lambda\alpha$ starting with $\tau_0 = 0$. Using the explicit solution of the OU type process for the variance:

$$V(t) = e^{-\lambda t}\left(V(0) + \int_0^t e^{\lambda s} dZ'(\lambda s)\right) \tag{13}$$

we have $$Z'(\lambda\tau_i) = Z'(\lambda\tau_{i-1}) + \frac{\eta_i}{\delta},\quad V(\tau_i) = e^{-\lambda(\tau_i - \tau_{i-1})}V(\tau_{i-1}) + \frac{\eta_i}{\delta} \tag{14}$$

and $$Z'(\lambda\tau_i) = Z'(\lambda\tau_{i-1}),\quad V(\tau_i) = e^{-\lambda(\tau_i - \tau_{i-1})}V(\tau_{i-1}) \text{ for } \tau_{i-1} \le t \le \tau \tag{15}$$

Thus a path from the exact joint distribution of the BDLP and the variance process is simulated and then observed on the discrete equidistant grid corresponding to daily closing prices. For the superposition model two independent copies of simulated BDLP and variance component process increments have been generated by this method.

Conditionally on knowing the process V, the integral $\int_{t_{i-1}}^{t_i} \sqrt{V(s)}ds$ is a centred Gaussian random variable with variance $V*(t) = \int_{t_{i-1}}^{t_i} V(s)ds$. One advantage of the Barndorff-Nielsen & Shephard (2001) OU formulation is that integrated variance can be expressed exactly as

$$V*(t) = \tfrac{1}{\lambda}\left(Z'(\lambda t_i) - V(t) + V(0)\right) \tag{16}$$

Thus we have the exact increment of the log-return simulated by

$$\mu\Delta t - \tfrac{1}{2}\left(V*(t_i) - V*(t_{i-1}) + \sqrt{V*(t_i) - V*(t_{i-1})}\,\varepsilon_i\right) \tag{17}$$

with (ε_i) iid standard normals, as above.

6.4 PARAMETER OPTIMISATION

Once the distributions were drawn, the second stage of the simulated method of moment's approach was done in three steps. The first step was to simulate price and volatility series that were consistent with the proposed models and secondly, determine the attribute values for this simulated series. The third step entailed varying the parameter inputs into the models to minimise the sum of squared errors relative to the observed empirical attributes. To efficiently perform the third step, a starting point was to examine the sensitivities of the overall sum of squared errors to a small incremental change in each of the parameter inputs (holding the other parameters constant). In Model 1 four parameters can vary: the long-term volatility θ, the rate of mean-reversion κ, the volatility of volatility ξ, and the correlation ρ. In a sensitivity analysis of Tompkins (2000), c.f. previous work of Tompkins (1998) he found that small changes in the level of the long-term volatility had almost no effect on the sum of squared errors, and that one can take safely the empirical average 20-day volatility as the estimate (see Table III) for θ. Thus, optimisation was only required over three parameters.

Given that only three parameters needed to be varied, the parameterisation simply compared the sum of the squared errors using the initial seeded parameter values to the sum of squared errors for the same model (and random numbers) but varying the three critical parameters (κ, ξ and ρ). By varying only three parameters both up and down, thus, only eight alternatives to the original results had to be compared. If one of the new combinations of parameters achieved a lower sum of squared errors, this model would replace the previous model. The search routine continued to search the "cube" of eight adjacent alternative parameter combinations until no new combination yielded better results. The initial search procedure first used fairly high increments in the adjacent corner search (for example 1.0 for κ and 0.1 for ξ and ρ). When optimal parameters were found, the increments for the search was progressively reduced (for example, as low as 0.01 for κ and 0.0001 for ξ and ρ) until no further reduction in the sum of squared errors was achieved. This parameterisation estimation method is more commonly known as the direct search polytope type algorithm developed in Tompkins (1998) and successfully applied in Tompkins (2000), [c.f. algorithm UMPOL from the IMSL math library]. For the Tompkins (2000) NIG stochastic volatility model, this procedure was carried out for each of the four representative NIG distributions, and in a final step the best fitting distribution was chosen.

In the superposition version of Model 3 we have nine parameters: The linear drift μ was set to 0 and β= -1/2, so there is $\lambda_1, \alpha_1, \delta_1, \rho_1$ for the first variance 'component', similarly $\lambda_2, \alpha_2, \delta_2, \rho_2$ for the second, and, as $w_2 = 1 - w_1$, one weight parameter describing the mixture. In Tompkins and Hubalek (2000), these nine parameters were shown to collapse to five. This was achieved by mapping the parameters of Model 3 to Model 1. A similar direct search polytope type algorithm was employed across the five parameters.

6.5 PARAMETER ESTIMATES FOR FOREIGN EXCHANGE FUTURES

Tables VIa, VIb, VIc and VId display the empirical attributes (taken from Tables II and III) and the results of the models for each of the foreign exchange futures markets. These results are split into the four periods of analysis. On the left-hand side, the parameterisation of each of the four models appears (the three models examined here plus the GBM case as a benchmark). For model 2, only the parameter values and the NIG distribution that has the lowest SSE are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison's sake. In the column immediately to the right of this the sum of squared errors (SSE) for each model appears. The final column indicates a T-test statistic comparing models 1, 2 and 3 to the GBM assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV.

[Tables VIa, VIb, VIc and VId appear here]

As was expected, the GBM model with constant variance is rejected in favour of models 1, 2 and 3. The T-statistics indicate a rejection at well above a 99% confidence interval. If we assume that the model with the lowest SSE is optimal, of sixteen data sets, Model 2 (NIGSV) is the best in ten and Model 3 (Γ2-OU) is the best for five. In only one instance (Swiss Franc in the period of 1985-1990) was the Heston (1993) model (Model 1) the best. For fifteen of the sixteen data sets, the model with the lowest sum of squared errors required correlated processes (the sole exception was for the Γ2-OU model for Deutsche Marks from 1995-2000).

These comparisons provide insights into which attributes are addressed by the facets of the proposed models. The pure stochastic volatility model (SV) addresses the volatility clustering effects measured by the two autocorrelation attributes as well as other volatility dynamics [the volatility of volatility (CoV attribute) and the decay of this over time (Line fit)]. The inclusion of correlated processes allows for a correlated volatility process and the non-negative skewness in the returns to be captured. However, this model still fails to generate sufficient excess kurtosis. Models 2 and 3 (NIGSV and Γ2-OU) appear equally able to capture all seven attributes.

Comparisons between Models 2 and 3 to Model 1 can be made by examination of the T-test statistics in Tables VIa to VId. As the T-test statistics provide an indication of the improvement in these models relative to the GBM case, taking the differences between these statistics provides insights into marginal improvement of Models 2 and 3 to Model 1. In fifteen of the sixteen cases, Models 2 or 3 has a higher T-statistic than Model 1. This is most notable for Japanese Yen futures, where the difference is significant at above a 95% level (for the overall period and the 1995-2000 period). For the other markets, it is not possible to distinguish in a statistical sense whether Models 2 and 3 are better than Model 1[13]. For all attributes except the excess kurtosis, both models perform equally well. What appears to be most critical is the inclusion of correlations between the underlying price and volatility processes. These capture the leverage and a portion of the skewness effect. Nevertheless, as both Models 2 and 3 still provide better fits than Model 1 for all the empirical attributes, we will qualitatively proclaim these the best models and will examine the implications for option pricing of these two models.

7. IMPLICATIONS FOR OPTION PRICING

From the preceding section, more realistic price processes for the four foreign exchange futures markets have thus been uncovered. These will serve as prior processes for the estimation of option values. The next step is to examine the implications for options based upon these assets. Given that the parameter estimation of the stochastic volatility/ jump models relies upon simulation, it is a simple matter to use a similar simulation technique to estimate option prices. For Model 2, this simulation approach determines European call and put options numerically (a Monte Carlo approach) and for Model 3, a closed form solution proposed by Hubalek and Tompkins (2000) was used. Option prices were determined over a variety of strike prices and times to expiration.

For both models examined, stochastic volatility, correlated processes and jumps have been introduced into the state space. Given that no securities exist allowing the state space to be spanned when the volatility displays such dynamics, these models generate incomplete markets and do not permit us to rely solely on risk-neutrality (no arbitrage) arguments to price the options[14]. Barndorff-Nielsen & Shephard (2001, Section 6.2) have shown that their model does not allow arbitrage by presenting explicitly an equivalent martingale measure. Other risk-neutral measure changes are discussed in Hubalek & Tompkins (2000).

To allow a comparison we have chosen the following approach. We select a measure change, which does not alter the overall structure of the model, but only modifies some parameters. This is done in analogy to several other similar situations: For example, the Hull & White (1987) and the Heston (1993) stochastic volatility models retain their structure after changing the measure. For Lévy processes the widely used Esscher change of measure, advocated for option pricing in Gerber & Shiu (1994), see also Delbaen & Haezendonck (1989) for compound Poisson processes, preservers virtually for all distributions used in finance the parametric class.

As the equivalent risk-neutral measure is not unique, it is certainly feasible, see Kreps (1981) and Harrison & Pliska (1981). This research examines whether this measure yields reasonable option prices. If this is case, we can assess if this alternative price process alone is sufficient to explain the existence of implied volatility smiles. If this is not the case, several reasons have to be investigated further:

1.market frictions not addressed here,

2.the parameter reduction from nine to five parameters,

3.the particular choice of the pricing measure (market price of risk),

4.the choice of the concrete distribution (BDLP),

5.the general structure of the model class.

For both models the starting point was the value of the optimal parameters estimated under the objective probability and reported in Tables VIa, VIb. VIc, and VId (for the entire period of analysis). Option prices will be determined over a variety of strike prices and times to expiration with the ultimate objective to understand the biases (compared to Black-Scholes values) that would result.

7.1 OPTION PRICES IN TOMPKINS (2000)

To determine option prices Tompkins (2000) used the discrete time approximation for Model 2 from Section 6.2. A very simple and natural choice of an equivalent martingale measure in discrete time [15]is to change the drift only. This is achieved by the density process:

$$G_n = \prod_{k=1}^{n} \frac{f(\varsigma_k; u_k; \delta, \alpha, \beta)}{f(\varsigma_k; u_k; \delta, \alpha, \beta)}, \qquad \mu_k = \frac{r-\mu}{\sigma_{k-1}} \tag{18}$$

with

$$f(x;u,\delta,\alpha,\beta)=\frac{\alpha\delta}{\pi}e^{\delta\sqrt{\alpha^2-\beta^2}-\beta(x-u)}\frac{K_1\left(\alpha\sqrt{\delta^2+(x-u)^2}\right)}{\sqrt{\delta^2+(x-u)^2}} \tag{19}$$

the density function of the $NIG(u,\delta,\alpha,\beta)$ distribution. Here K_1 is the modified Bessel function of third kind with index one, see Abramowitz & Stegun (1965). With the pricing measure now chosen, the actual computation of the discounted expectations was done by Monte Carlo simulation, similar in spirit to Johnson & Shanno (1987).

7.2 OPTION PRICES FOR BARNDORFF-NIELSEN & SHEPHARD (2001)

Our approach to option pricing, namely to preserve the model under risk-neutrality, is especially fruitful for Model 3. The Barndorff-Nielsen & Shephard (2001) model allows explicit calculation of the moment generating function of returns, thus an explicit (in the sense of Heston (1993) or Stein & Stein (1991)) option pricing formula [see Barndorff-Nielsen & Shephard (2001, Section 6.2)]. The actual computation of option prices was done by the concrete implementation of this analytic solution is discussed in Hubalek & Tompkins (2000).

We have under the objective probability measure *P:*

$$dx(t)=(\mu+\beta\sigma^2(t-))dt+\sigma(t-)dw(t)+\sum_{j=1}^{m}p_j d\bar{z}_j(\lambda_j t), \tag{20.1}$$

$$d\sigma^2(t)=\sum_{j=1}^{m}w_j d\sigma^2(t), \tag{20.2}$$

$$d\sigma_j^{\,2}(t)=-\lambda_j\sigma_j^{\,2}(t-)dt+d\bar{z}_j(\lambda_j t), \tag{20.3}$$

and by Ito's formula

$$dS(t)=S(t-)\left(\left(\mu+(\beta+\tfrac{1}{2})\sigma^2(t-)+\sum_{j=1}^{m}\lambda_j\bar{\kappa}_j(\rho_j)\right)dt+\sigma(t-)dw(t)+\sum_{j=1}^{m}dm_j(\lambda_j t)\right) \tag{21}$$

where $\bar{\kappa}_j$ is the cumulant function of $\bar{z}_j$ and m_j is the purely discontinuous martingale with jumps $\Delta m_j(t)=e^{\rho_j\Delta\bar{z}_j(t)}-1$. Using the jump measures $\bar{\mu}_j$ and compensated jump measures $\bar{\nu}_j$ associated with $\bar{z}_j$, we can write

$$M_j(t)=\int_0^t\int(e^{\rho_j z}-1)(\bar{\mu}_j-\bar{\nu}_j)(dt,dz) \tag{22}$$

We choose a risk-neutral probability measure *Q*, using a minor generalisation of Nicolato & Prause (1999) and Nicolato & Venardos (2000) with density process

$$G(t)=\exp\left(-\int_0^t\gamma(s)dw(s)-\frac{1}{2}\int_0^t\gamma(s)^2 ds\right), \tag{23.1}$$

where

$$\gamma(t)=\left(\mu-r+\sum_{j=1}^{m}\lambda_j\bar{\kappa}_j(\rho_j)\right)\sigma^{-1}(t-)+(\beta+\tfrac{1}{2})\sigma(t-). \tag{23.2}$$

Then Girsanov's theorem tells us, that

$$w'(t) = w(t) + \int_0^t \gamma(s)ds \tag{24}$$

defines a process w' that is under Q a Brownian motion independent of $\grave{z}$. The processes $\grave{z}_j$ and $\grave{m}_j$ are not affected by the measure change. We get, as desired,

$$dS(t) = S(t-)\left(rdt + \sigma(t-)dw'(t) + \sum_{j=1}^{m} dm_j(\lambda_j t) \right) \tag{25}$$

This means for the returns process x, that under Q:

$$dx(t) = (\mu' + \beta'\sigma^2(t-))dt + \sigma(t-)dw'(t) + \sum_{j=1}^{m} \rho_j d\bar{z}_j(\lambda_j t), \tag{26.1}$$

with

$$\mu' = r - \sum_{j=1}^{m} \lambda_j \bar{\kappa}_j(\rho_j), \quad \text{and} \quad \beta' = -\tfrac{1}{2} \tag{26.2}$$

Concrete for a mixture of BDLPs corresponding to stationary gamma distributions we have

$$\bar{\kappa}_j(\rho_j) = \alpha_j \left(\frac{\delta_j}{\delta_j - \rho_j} - 1 - \frac{\rho_j}{\delta_j} \right) \tag{27}$$

7.3 IMPLIED VOLATILITY SURFACES

In this simulation, price series of three months in length were determined. Given that the estimation of the unconditional (historical) dispersion processes was completed for trading days, options were also priced using trading time instead of calendar time. The assumed number of trading days in a year is 252. Option prices were estimated at time horizons from one week (five trading days) to three months (in 5-day increments). Such options would correspond to typical terms to maturity of actively traded options on Foreign exchange Futures.

To gain a better understanding of the impacts of the alternative models across strike prices, fifteen strike prices were examined. The median strike price was centred at the starting value of the simulation and was equal to 100. As the assumed underlying assets were futures or forward contracts, the interest rate was assumed to zero. This corresponds to an at-the-money option relative to the forward price.

The impacts of the model on options with different strike prices were of additional interest. The analysis was restricted solely to out-of-the-money strike prices. Thus, when the strike price was equal to or below the starting value of 100, the option evaluated was an European put and when the strike price was above 100, the option evaluated was an European call. A non-trivial problem is the choice of strike prices so that as maturities of options vary, meaningful comparisons can be drawn. In previous papers on the impacts of stochastic volatility on option prices, strike price determination has taken one of two forms. Authors have either chosen to fix a single maturity and vary the strike prices in terms of "moneyness" [see Hull & White (1988)] or fixed the degree of moneyness (or strike prices) and examined the impacts across different maturities [see Henker and Kazemi (1998)].

Unfortunately, both methods do not allow meaningful conclusions to be drawn regarding the impacts of the models on option prices across time and a consistent measure of moneyness. Natenberg (1994) and Tompkins (1997) have proposed a more consistent measure of strike price. This was slightly modified to:

$$\frac{\ln(X_\tau / F_\tau)}{\sigma\sqrt{\tau / 252}} \quad \textbf{(28)}$$

where X is the strike price of the option, F is the underlying futures price and the square root of time factor reflects the percentage in a trading year of the remaining time until the expiration of the option. The sigma (σ) is the at-the-money volatility.

This adjustment notes that the distance of an option strike price to the level of the underlying asset is relative, both in respect to the current price of the underlying, the time to expiration and the level of expected volatility. This adjustment converts all strike prices into a metric that can be interpreted as a standard deviation. Thus, in this analysis, strike price ranges ± 3.5 standard deviations away from the at-the-money level in 0.5 standard deviation increments were examined. This change in measure will allow more direct comparison of model impacts on option prices where the time to maturity varies but the relative strike prices remain the same. [16]

For the simulations, the volatility parameter chosen was equal to the level of volatility used in the parameter determination of all models for each of the four markets for the entire period of analysis. Given the extremely wide range of volatilities across the four markets, the standardisation of the strike prices allows direct comparisons to be made and allows subsequent comparisons to be made with actual implied volatility surfaces.

To determine the prices from the Tompkins (2000) model, a Monte Carlo method was used. For this random numbers consistent with a GBM process were determined using a Box-Muller technique and employed the anti-thetic approach suggested by Boyle (1977) for both. This series was later used to determine the bivariate distribution used to estimate the stochastic volatilities. This series of random numbers were stored and used for all subsequent estimations of stochastic volatility. For each of the three NIG distributions, 10000 were drawn using the method suggested by Rydberg (1997). The same approach was used for the estimation of the optimal stochastic volatility parameters in the previous section. These were also stored and used for all analysis using that particular NIG distribution. Then, the dependent random numbers were constructed for the volatility series using Equation (11) and the optimal parameters of Model 2 for each of the Foreign Exchange futures markets for the entire period of analysis (in Tables VIa, VIb, VIc and VId).

With the appropriate NIG distribution and the estimated bivariate distribution for the stochastic volatility, volatilities and prices were computed using an Euler approach. With the prices of each model estimated, the payoffs of the fifteen options (at each point in time) were determined and the result averaged. As interest rates were assumed to be zero, there is no need to discount the result to present value. Option prices consistent with the Barndorff-Nielsen & Shephard (2001) model were estimated using the implementation method suggested by Hubalek & Tompkins (2000). The parameter values for the model are drawn from Tables VIa, VIb, VIc and VId for the four markets.

In parallel, we estimated the prices of all the options for each market using the Black (1976) model with the same strike prices as the simulation, the same term to expiration and the volatility equal to the same long-term volatility used in the simulations[17]. The underlying futures prices used in the determination of the

Black (1976) price were equal to the average futures price in the simulation at the same point in time. Given that the underlying asset is a futures contract, the interest rate and dividend yield was set to zero (the same assumptions were made when estimating the option prices for the two alternative models).

Although, standardisation of the strike prices simplifies comparisons between the seven markets, the sheer amount of information makes such comparisons cumbersome. Thus, to simplify comparisons between simulated and actual implied volatility surfaces; the simulated option prices were expressed as implied volatilities. Thus, the simulated option prices were input into the Black (1976) model and the implied volatility was determined in the usual numerical manner. These implied volatilities were further standardised by indexing them to the constant volatility assumed in the Black (1976) model (dividing each of the implied volatilities by the assumed constant volatility and multiplying by 100). The indexed implied volatilities are then presented as a continuous surface relative to the standardised strike prices and time to expiration.

7.3 IMPLIED VOLATILITY SURFACES FOR FOREIGN EXCHANGE OPTIONS

Implied volatility surfaces (for the four Foreign Exchange futures markets) consistent with Tompkins (2000) appears in Figure 1a and for Barndorff-Nielsen & Shephard (2001) in Figure 1b. The scaling of these figures is somewhat extreme to allow subsequent comparisons with actual implied volatility surfaces for options on these four markets (in Figure 2).

[Figures 1a, and 1b appear here]

The simulated implied volatility surface for the four foreign exchange futures markets appears to display some of the characteristic features of implied volatility surfaces associated with options on these markets [see Bates (1996b)]. The characteristic curvature of volatility smiles is found and for the Japanese Yen and Swiss Franc markets, the surfaces are positively skewed. This is most apparent in Figure 1b for the BNS model. The smiles for the Tompkins model (Figure 1a) display more curvature, which tends to be more extreme for longer maturity options, while for the BNS model, there is much less curvature although the degree of skewness appears more pronounced, especially for shorter-term options. The Tompkins smiles are somewhat more jagged compared to the BNS smiles and this is due to errors introduced by numerical estimation, while the BNS model is estimated analytically.

While we have not explicitly discussed the impacts of these models across a cross-section of option prices, this figure implicitly displays these impacts. For all the markets, the Black (1976) pricing model overvalues options that are at-the-money and within a significant range around (and above) the at-the-money level. The shaded areas below 100 represent this in the graphs. For out-of-the-money options, the Black (1976) model tends to undervalue option prices and the overpricing bias tends to increase the longer the term to expiration and depends on whether the models produce skewed patterns. These results are consistent with the biases of stochastic volatility on option prices found elsewhere [Hull & White (1988)]. A significant difference between the two models is that for the Tompkins smiles, there are much more extreme differences compared to the Black (1976) model. For the BNS model, the degree of Black (1976) overvaluation of ATM options is slight. Divergences tend to occur mostly for out-of-the-money options.

Of particular interest in this research is the comparison of these simulated implied volatility surfaces to actual implied volatility surfaces associated with options on these four foreign exchange futures markets. In Figure 2, the actual implied volatility surface associated with options on these four foreign exchange futures appears.

[Figure 2 appears here]

To construct this surface, implied volatilities were estimated using the Black (1976) model for all (out of the money) options on the four foreign exchange futures markets. The period of analysis was contemporaneous with the period of analysis of the four foreign exchange futures markets. Analysis was restricted solely to options with the same terms to expiration as were used for the simulated option prices (5 days to 3 months in 5-day increments) and excluded all options prices allowing arbitrage [see Jackwerth & Rubinstein (1996)]. The strike prices were then converted to the same standardised form as was done for the simulated implied volatility surfaces, using Formula (28), and the levels of implied volatility were indexed to the level of the at-the-money implied volatility. Then, the indexed implied volatilities were grouped by the same term to expiration. Finally, for each term to expiration, a polynomial fitting procedure [see Shimko (1993) and Tompkins (1998b)] was implemented to yield an implied volatility smile. These were then combined and the volatility surface was drawn. The rationale for standardisation of both strike prices and (indexing) the implied volatilities is now clear: direct comparisons of the relative (strike price and time varying) biases associated with the models can be compared to actual deviations from the Black (1976) model.

For the actual implied volatility surface associated with options on these four markets, it is clear that the degree of curvature is most extreme when the options are closest to expiration and becomes slightly less curved with longer expirations. Of the two models compared here, both display curvatures. However, the degree of curvature is much less than is observed for actual option prices, and on average, the actual smile patterns are symmetric, while the simulated patterns often are asymmetric. The BNS model appears to conform better to the time dynamics of actual smile patterns, as the greatest degree of curvature is also for the shortest term to expiration options (5 trading days) and tends to flatten with longer maturities. However, the degree of curvature is simply too low. For the Tompkins model, the curvature for the simulated surfaces tends to be more stable over time (apart from the Japanese Yen, which becomes more extreme for longer-term options). While this model introduces more curvature than the BNS model, the amplitude is also too low. Given the scaling allows direct comparisons, for the actual options markets, the deepest out-of-the money options have standardised implied volatilities of 130% to 160% of the level of volatility of the at-the-money option. In contrast, the simulated implied volatility surfaces have standardised implied volatilities for these same options of only 110%-130% (relative to the constant volatility assumption of the Black (1976) model). In addition, these models introduce a skewness effect that is not found in the (average) actual implied volatility surfaces.

To better understand the divergences in the simulated and actual implied volatility surfaces, simple differences were taken between the standardised volatilities in Figure 2 and Figures 1a and 1b. These appear in Figures 3a and 3b for the Tompkins and BNS models, respectively.

[Figures 3a, and 3b appear here]

Standardisation of both simulated and actual implied volatility surfaces simplifies interpretation: as the percentage difference in the (standardised) level of implied volatility. For the Tompkins (2000) model, the actual implied volatilities are between 10% and 50% higher than those associated with the implied volatilities consistent with this model. The divergences tend to be fairly stable over time. For those markets where skewness and/or correlated processes were important in the underlying Futures, this leads to asymmetric divergences from actual option prices. For all maturities, the Tompkins model of the objective process explains almost none of the curvature in the actual implied volatility surface. This is especially relevant for options with higher strike prices. For higher strike price options, the difference remains at fairly similar levels (approximately 30% - 50%) and appears to be time invariant. Of particular interest is that the differences in these two surfaces seems to be similar across the four markets.

The Barndorff-Nielsen & Shephard (2001) model also fails to explain the actual implied volatility surfaces. In all instances, there is much more curvature in the actual surfaces than those consistent with the BNS model. This is especially the case for options with the shortest term to expiration. As with the Tompkins model, the skewness and/or correlation effects associated with the objective price process are inconsistent with actual implied volatility surfaces. This leads to asymmetric divergences. The major difference between the Tompkins and BNS models is that in many instances the BNS model produces option prices, which are above those observed for actual option markets. This can be seen by divergences below zero.

7.4 Implications of Divergences in Implied Volatility Surfaces

Let us assume that either of the two models examined here for the objective price process of financial futures is appropriate. Why do the option prices consistent with this process diverge in a systematic manner from actual options on these markets?

The divergence between the patterns of implied volatility surfaces suggests that the inclusion of our proposed stochastic volatility models in option pricing is insufficient to explain the existence of implied volatility smiles. The most obvious conclusion is that we have examined the wrong models and that some alternative approach may better describe the unconditional price process for the underlying futures. However, it is unclear what the alternative approaches would be. The models tested here included non-normality in the return process, stochastic volatility and correlated processes. Most models proposed in the literature to explain the objective price and volatility processes are nested in the model examined here.

A second possibility is that models we have examined are correct, but that our method of parameter estimation has yielded incorrect model inputs. The potential problems with parameter estimation could be due either to inadequacies of the simulated method of moments approach or the selection of inappropriate (or incomplete) target moments. As the former has been extensively examined in the literature and dismissed, this is unlikely. The choice of target moments was determined including all relevant facets of non-normality and volatility that have been pointed out in the empirical literature; it is unclear what additional attributes would be included.

The third possibility is even if the price process of the underlying asset was correctly modelled, it is not obvious that option prices would conform to this process. The Cox & Ross (1976) link between the objective and risk-neutral processes would be incomplete in the presence of non-traded sources of risk. For the models examined here, the introduction into the state space of non-constant volatility, jump-processes and correlations between the two processes would introduce just such non-traded sources of risk. Thus, it is clear that our assumption of a unique martingale measure is incorrect. Given this is the case, a logical next step would be to better understand the nature of this risk. Scott (1987) explicitly examined the dynamics of the risk premium when volatility is stochastic and is not a traded security. We contend that the differences in the simulated and actual implied volatility surfaces for both models provide this information.

The fact that such systematic differences are observed may point to alternative explanations for the existence of smiles: market frictions and risk premia and that these dynamics are general to all four options markets. The fact that for the Tompkins model the difference between the surfaces is always positive could be interpreted as evidence for the existence of a risk premium. Recently, Guo (1998) examined the risk premium of implied volatilities in the currency markets. He also examined the Heston (1993) model (Model 1) and compared parameters implied by option prices to parameters estimated by the unconditional futures price process. He estimates a significant risk premium exists which is non-zero, time varying and a substantial component of implied volatilities. Figure 3a also displays similar results and provides evidence of what could be interpreted as a risk premium. For the BNS model, the fact that the simulated option prices exceed actual option prices is divergent with the existence of a positive risk premium. Given this, it would appear that the Tompkins model with the inclusion of risk premium might be a better candidate for understanding actual implied volatility surfaces.

Alternatively differences in these surfaces could be associated with transaction costs or could be caused by other sources of market imperfections or frictions. The nature of these effects warrants further research. The primary contribution of this research is to allow these effects to be stripped out and examined once a feasible alternative price (and volatility) process for the objective price process has been found.

8. SUMMARY AND SUGGESTIONS FOR FUTURE RESEARCH

This paper examines the nature of the objective dispersion processes, which can be observed for four foreign exchange futures contracts. To capture the multi-faceted non-normality and inter-dependence of the empirical dispersion processes, seven attributes were identified. With these attributes for each of the four markets, three alternative stochastic volatility models were examined. It was found that for all four markets two alternative models with stochastic volatility and jumps were able to explain the objective price processes.

Based upon an optimal parameterisation of these models, European options were determined for each market. Significant divergences from the Black (1976) values were observed for both models. For all markets, this option pricing model tended to increase the prices of options with lower strike prices relative to Black (1976) prices and decrease the value of ATM and higher strike price options. This effect is most pronounced for the Tompkins model. When differences in simulated and actual implied volatility surfaces are examined, the Tompkins model indicates the existence of a positive risk premium, which has been identified elsewhere in the literature. The Barndorff-Nielsen & Shephard (2001) model seems to suggest a negative risk premium for a wide range of option prices, which is divergent to this same literature. Thus, we conclude that of these two approaches, it does matter how jumps are incorporated in the unconditional price process. For Foreign Exchange markets, it appears that jumps in the underlying price process are more consistent than jumps in the volatility process.

Nevertheless, we reject the hypothesis that the inclusion of jumps in combination with stochastic volatility models is sufficient to explain the existence of implied volatility smiles. Of interest is that the differences between the implied volatility surfaces consistent with the objective and risk-neutral processes appear to display similar dynamics for all four markets (for both models). The nature of this difference adds more evidence for the existence of a risk premium in options prices and provides insights into the relationship between the objective and risk-neutral processes.

REFERENCES:

Abramowitz, M. & Stegun, I. A. (1965), Handbook of Mathematical Functions, Dover Publications Inc.

Andersen, T.G., Chung, H-J & Sørensen, B.E.(1999). Efficient method of moment's estimation of a stochastic volatility model: A Monte Carlo study. Journal of Econometrics, 91, 61-87.

Andersson, J. (1999a). On the normal inverse Gaussian stochastic volatility model. Essays on Financial Time Series Models, *Stochastic Volatility and Long Memory,* Acta Universitatis Upsaliensis. *Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences,* 82, 1-10.

Andersson, J. (1999b). Maximum Likelihood estimation of the normal inverse Gaussian stochastic volatility model. *Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences,* 82. 1-11.

Atkinson, A. C. (1982), The simulation of generalised inverse Gaussian and hyperbolic random variables, SIAM Journal of Scientific Statistical Computing 3(4), 502-515.

Baillie, R.T. & Bollerslev, T. (1989). The Message in Daily Exchange Rates: A Conditional Variance Tale. Journal of Business and Economic Statistics, 7, 297-305.

Bakshi, G., Cao, C. & Chen, Z. (1997). Empirical Performance of Alternative Option Pricing Models. Journal of Finance, 52, 2003-2049.

Barndorff-Nielsen, O. E. (1977), `Exponentially decreasing distributions for the logarithm of particle size', Proceedings of the Royal Society, London. 353, 401-419.

Barndorff-Nielsen, O. E. (1978), `Hyperbolic distributions and distributions on hyperbolae', Scandinavian Journal of Statistics 5, 151-157.

Barndorff-Nielsen, O. E. (1995), Normal inverse Gaussian processes and the modelling of stock returns, Research Report 300, Aarhus University, Dept. Theoretical Statistics.

Barndorff-Nielsen, O.E. (1997). Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling. Scandinavian Journal of Statistics, 24, 1-13.

Barndorff-Nielsen, O. E. & Shephard, N. (1998), Incorporation of a leverage effect in a stochastic volatility model, MaPhySto Research Report 18, University of Aarhus.

Barndorff-Nielsen, O. E. & Shephard, N. (2000), Modelling by Lévy processess for financial econometrics, in O. E. Barndorff-Nielsen, T. Mikosch & S. Resnick, eds, `Lévy Processes -- Theory and Applications', Birkhauser, Boston.

Barndorff-Nielsen, O. E. & Shephard, N. (2001), `Non-Gaussian OU based models and some of their uses in financial economics', The Journal of The Royal Statistical Society. forthcoming.

Barone-Adesi, G. & Whaley, R. (1987), "Efficient analytic approximation of American option values", Journal of Finance 42, 301-320.

Bates, D.S. (1996a). Dollar Jump Fears, 1984-1992: Distributional Abnormalities Implicit in Currency Futures Options. Journal of International Money and Finance, 15, 65-93.

Bates, D.S. (1996b). Jumps and Stochastic Volatility: Exchange Rate Process Implicit in PHLX Foreign Currency Options. Review of Financial Studies, 9, 69-107.

Bates, D.S. (2000). Post-'87 Crash Fears in the S&P 500 Futures Options Markets. Journal of Econometrics, 94, 181-238.

Ben Khelifa, Z. (1991). Parametric and Nonparametric Tests of the Pure Diffusion Model Adjusted for the Early Exercise Premium Applied to Foreign Currency Options. in Essays in International Finance, Wharton School dissertation, University of Pennsylvania, Chapter 1.

Black, F. & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. Journal of Political Economy, 81, 637-654.

Black, F. (1975). Fact and Fantasy in the Use of Options. Financial Analysts Journal, 31, 36-41,61-72.

Black, F. (1976). The Pricing of Commodity Contracts. Journal of Financial Economics, 3, 167-179.

Blæsild, P. (1995), `Lecture given at a workshop on stochastic processes and financial markets', quoted in Rydberg (1997).

Bodurtha, J. N. & Courtadon, G. R. (1987), `Tests of an American option pricing model on the foreign currency options market', Journal of Financial and Quantitative Analysis, 153-167.

Bollerslev, Engle and Nelson (1994)

Boyle, P.P.. (1977). Options: A Monte Carlo Approach. Journal of Financial Economics, 4, 323-338.

Burghardt, G. & Lane, M. (1990). How to Tell if Options are Cheap. Journal of Portfolio Management, 16, 72-78.

Cao, C. (1992). Pricing Foreign Currency Options with Stochastic Volatility. Working Paper, University of Chicago.

Chesney, M. & Scott, L.O. (1989). Pricing European Currency Options: A Comparison of the Modified Black-Scholes Model and a Random Variance Model. Journal of Financial and Quantitative Analysis, 24, 267-287.

Christie, A.A., (1982). The Stochastic Behaviour of Common Stock Variances. Journal of Financial Economics, 10, 407-432.

Cox, J.C. & Ross, C.A. (1976). The Valuation of Options for Alternative Stochastic Processes. Journal of Financial Economics, 3, 145-166.

Delbaen, F. & Haezendonck, J. (1989), `A martingale approach to premium calculation principles in an arbitrage free market', Insurance and Mathematical Economics. 8(4), 269-277.

Ding, Z., Granger, C.W.J. & Engle, R.F. (1993). A long memory property of stock returns and a new model. Journal of Empirical Finance, 1, 83-106.

Duffie, D., & Singleton, K.J. (1993). Simulated moments estimation of Markov models of asset prices. Econometrica, 50, 929-952.

Duffie, D., Pan, J. & Singleton, K. (1999), Transform analysis and asset pricing for affine jump-diffusions, NBER Working Paper No. W7105.

Eberlein, E. & Keller, U. (1995), Hyperbolic distributions in Finance, Bernoulli 1(3), 281-299.

Eberlein, E. & Raible, S. (1999), `Term structure models driven by general Lévy processes', Mathematical Finance 9(1), 31-53.

Fama, E. (1963), `Mandelbort and the stable Paretian hypothesis', Journal of Business pp. 420-429.

Fama, E. (1965), `Portfolio analysis in a stable Paretian market', Management Science pp. 404-419.

Gallant, A.R. & Tauchen, G. (1996). Which Moments to Match?. Econometric Theory, 12, 657-681.

Garman, M. & Kohlhagen, S. (1983). Foreign Currency Option Values. Journal of International Money and Finance, 2, 231-237.

Gerber, H. U. & Shiu, E. S. W. (1994). Option pricing by Esscher transforms. Trans. Soc. Actuar. XLVI, 98-140.

Guo, D. (1998). The Risk Premium of Volatility Implicit in Currency Options. Journal of Business and Economic Statistics, 16 (4), 498-507.

Harrison, J. M. & Pliska, S. R. (1981), `Martingales and stochastic integrals in the theory of continuous trading', Stochastic Process. Appl. 11(3), 215-260.

Harvey, C.R. & Siddique, A. (1998). Conditional Skewness in Asset Pricing Tests. Working Paper (June 4, 1998 version). Duke University, Durham North Carolina.

Henker, T. & Kazemi, H.B. (1998). The Impact of Deviations from Random Walk in Security Prices on Option Prices. Working Paper University of Massachusetts, Amherst.

Herzel, S. (1998), `A simple model for option pricing with jumping stochastic volatility', International Journal of Theoretical and Applied Finance 1(4), 487-505.

Heston, S.L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. Review of Financial Studies, 6, 327-343.

Ho, M.S., Perraudin, W.R.M. & Sørensen, B.E. (1996). A Continuous-Time Arbitrage-Pricing Model with Stochastic Volatility and Jumps. Journal of Business & Economic Statistics, 14, 31-43.

Hobson, D. G. & Rogers, L. C. G. (1998), `Complete models with stochastic volatility', Mathematical Finance 8(1), 27-48.

Hodges, S. & Tompkins, R. (2000), The sampling properties of volatility cones, FORC Preprint 00/103, Financial Options Research Centre, University of Warwick, Coventry, England.

Hsieh. D.A. (1988). The Statistical Properties of Daily Foreign Exchange Rates: 1974-1983. Journal of International Economics, 24, 129-145.

Hubalek, F. (1999), On equivalent martingale measures and classical asymptotic statistics for generalised hyperbolic distributions, Talk at Mathematische Stochastik (Finance and Statistics), Mathematisches Forschungsinstitut Oberwolfach, Germany.

Hubalek, F. & Tompkins, R. G. (2000), On closed form option pricing for OU type models, unpublished.

Hull, J. & White, A. (1987), `The pricing of options on assets with stochastic volatilities', Journal of Finance 42, 281-300.

Hull, J. & White, A. (1988). An Analysis of the Bias in Option Pricing Caused by a Stochastic Volatility. Advances in Futures and Options Research, 3, 29-61.

Jackwerth, J.C. & Rubinstein, M. (1996). Recovering Probability Distributions from Option Prices. The Journal of Finance, 51, 1611-1631.

Jacod, J. & Shiryaev, A. N. (1987), Limit theorems for stochastic processes, Springer-Verlag, Berlin.

Johnson, H. & Shanno, D. (1987). Option Pricing When the Variance is Changing. Journal of Financial and Quantitative Analysis 22, 143-152.

Jorion, P. (1988). On Jump Processes in the Foreign Exchange and Stock Markets. Review of Financial Studies 1, 427-445.

Kreps, D. M. (1981), Arbitrage and equilibrium in economies with infinitely many commodities. Journal of Mathematical Economics 8(1), 15-35.

Liu, S.M. & Brorsen, B. W. (1995), `Maximum likelihood estimation of a GARCH-stable model', Journal of Applied Econometrics 10, 273-285.

Marcus, M. B. (1987), "χ-radial processes and random Fourier series', Mem. Amer. Mathematical Soc. 68(368), viii+181.

Mayhew, S. (1995). Implied Volatility. Financial Analysts Journal, 51 (4), 8-20.

McDonald, J. B. & Newey, W.K. (1988). Partially Adaptive Estimation of Regression Models via the Generalized t Distribution. Econometric Theory 4, 428-457.

Meese, R.A. (1986). Testing for Bubbles in Exchange Markets: A Case of Sparkling Rates?. Journal of Political Economy 94, 345-373.

Melino, A., & Turnbull, S.M. (1990). Pricing foreign currency options with stochastic volatility. Journal of Econometrics, 45, 239-265.

Merton, R. (1973). The Theory of Rational Option Pricing. Bell Journal of Economics, 4, 141-183.

Merton, R. (1976). Options pricing when Underlying Stock Returns are Discontinuous. Journal of Financial Economics, 3, 125-144.

Michael, J. R., Schucany, W. R. & Haas, R. W. (1976), `Generating random variates using transformations with multiple roots', Am. Stat. 30, 88-90.

Mittnik, S. & Rachev, S. (1993), `Modelling asset returns with alternative stable models', Econometric Reviews 12, 261-330.

Natenberg, S. (1994). Option Volatility and Pricing: Advanced Trading Strategies and Techniques, Chicago, Illinois: Probus Publishing Company.

Nelson, D. B. (1991), `Conditional heteroskedasticity in asset returns: A new approach', Econometrica 59(2), 347-370.

Nicolato, E. & Prause, K. (1999), Derivative pricing in stochastic volatility models of the Ornstein-Uhlenbeck type, Working paper, University of Aarhus.

Nicolato, E. & Venardos, E. (2000), Derivative pricing in Barndorff-Nielsen and Shephard's OU type stochastic volatility models, Working paper, University of Aarhus.

Perrod, B. (1999), `Jumps in volatility of financial markets', (C.R.D.E.) working paper (Cahier 09999).

Protter, P. & Talay, D. (1997), `The Euler scheme for Lévy driven stochastic differential equations', Ann. Probab. 25(1), 393-423.

Rogalski, R.J. and Vinso, J.D. (1978). Empirical Properties of Foreign Exchange Rates. Journal of International Business Studies 9, 69-79.

Rosinski, J. (1991), On a class of infinitely divisible processes represented as mixtures of Gaussian processes, in `Stable processes and related topics (Ithaca, NY, 1990)', Birkhäuser Boston, Boston, MA, pp. 27-41.

Rydberg, T.H. (1997). The Normal Inverse Gaussian Lévy Process: Simulation and Approximation. Commun. Stat., Stochastic Models, 13, 887-910.

Scott, L.O. (1987). Option Pricing when the Variance Changes Randomly: Theory, Estimation and an Application. Journal of Financial and Quantitative Analysis, 22, 419-438.

Scott, L.O. (1991). Random Variance Option Pricing: Empirical Tests of the Model and Delta-Sigma Hedging. Advances in Futures and Options Research 5, 113-135.

Shastri, K. & Wethyavivorin, K. (1987). The Valuation of Currency Options for Alternative Stochastic Processes. Journal of Financial Research, 10 (4), 283-294

Shimko, D. (1993), `Bounds of probability', RISK 6(4), 33-37.

Shuster, J. (1968), `On the inverse Gaussian distribution function', Journal of the American Statistical Society. 63, 1514-1516.

Steel, M.F. (1998). Bayesian Analysis of Stochastic Volatility Models with Flexible Tails. Economic Reviews 17, 109-143.

Stein, E.M. & Stein, J.C. (1991) Stock Price Distributions with Stochastic Volatility: An Analytic Approach," Review of Financial Studies, 4, 727-752.

Taylor, S.J. (1986). Modelling Financial Time Series, New York: John Wiley & Sons.

Taylor, S.J. (1994). Modelling Stochastic Volatility: A Review and Comparative Study. Mathematical Finance, 4, 183-204.

Theodossiou, P. (1998). Financial Data and The Skewed Generalized T Distribution. forthcoming in Management Science.

Tompkins, R. G. (1997), Measuring equity volatilities, in `Equity Derivatives', RISK Publications, RISK, London, England, chapter 12.

Tompkins, R. G. (1998), Security Price Process Models: Do These Have the Correct Properties for Understanding Option Values?, PhD dissertation, University of Warwick.

Tompkins, R. G. (1998b), Implied volatility surfaces: Uncovering regularities for options of financial futures, FORC Preprint 98/93, Financial Options Research Centre, University of Warwick, Coventry, England.

Tompkins, R. G. (2000), `Stock index futures markets: Stochastic volatility models and smiles', Journal of Futures Markets. forthcoming.

Tompkins, R. G. (2000b), `The Austrian Stock Market: Objective price processes and smiles', Finanzmarkt und Portfolio Management. forthcoming.

Tompkins, R.G. & F. Hubalek (2000). Does it matter where jumps come from? Underlying versus volatility processes, working paper, University of Technology, Vienna.

Vlaar, P. J. G. & Palm, F. C. (1993), The message in weekly exchange rates in the European monetary system: Mean reversion, Conditional Heteroscedasticity, and Jumps, Journal of Business and Economic Statistics. 11, 351-360.

Westerfield, J.M. (1977). Empirical Properties of Foreign Exchange Rates Under Fixed and Floating Regimes. Journal of International Economics 7, 181-200.

Wiggins, J.B. (1987). Option Values Under Stochastic Volatility: Theory and Empirical Estimates. Journal of Financial Economics, 19, 351-372.

FOOTNOTES:

---

[1] As the process F has jumps, left sided limits F(t-) are required for a proper definition of stochastic integrals with respect to infinite variation jump processes.
[2] The price process F in (2.1) satisfies $F(t) = F(0) \exp(X(t))$, where the log-return process, X, satisfies $dX(t) = (\mu - \psi(\sigma(t)))dt + \sigma(t)dZ(t)$ with Z, a standard NIG Lévy process. The cumulant function $\psi(\sigma) = \ln E[e^{\sigma Z(1)}]$ is the pure jump Lévy process analogue of the Itô drift correction $-\sigma^2/2$ for diffusion processes. For the $NIG(\mu,\delta,\alpha,\beta)$ distribution, we have $\psi(\sigma) = \delta(\sqrt{\alpha^2 - \beta^2} - \sqrt{\alpha^2 - (\beta + \sigma)^2} + \mu\sigma$. We note that Z is a purely discontinuous martingale with jumps denoted by ΔZ. By Itô's formula for semi-martingales, we obtain the SDE (2.1), where the N is a purely discontinuous martingale with jumps: $\Delta N(t) = e^{\sigma(t)\Delta Z(t)} - 1$
[3] The estimation of the volatilities was done using overlapping data. This introduces a bias in the estimation of the standard deviation of volatility. This bias was corrected using the Hodges and Tompkins (2000) approach.
[4] The author would like to thank Stewart Hodges (University of Warwick) for suggested the use of this variable.
[5] To simplify the selection of the attribute, a fixed weight of 0.03 was applied to all markets and for all periods of analysis, to allow the leverage correlation factor to not be subject to differing weights. This was found to be close to the optimal weights for each market using a maximum likelihood estimation procedure.
[6] In addition, the natural logarithm of the unconditional kurtosis was examined rather than the absolute levels. This was done due to wide variations in this statistic across the four markets. However, all results are presented as absolute levels.
[7] Under the null hypothesis of normality, the skewness statistic is asymptotically normally distributed with standard errors: se $=\sqrt{(6/T)}$, where T is the sample size.
[8] Under the null hypothesis of normality, the excess kurtosis statistic is asymptotically normally distributed with standard errors: se $=\sqrt{(24/T)}$, where T is the sample size. This statistic is equal to the kurtosis statistic appearing in the table minus 3.0.

[9] The origin is inaccessible for the square-root process with $2\kappa\theta / \xi^2 \geq 1$, which is the case for all our estimates. Yet the discrete time approximation can take negative values with (usually very small) positive probability, which were handled by reflection.

[10] Actually parameters that produce less skewness and kurtosis than observed for unconditional daily returns (see Table III) have been selected, due to the fact that the stochastic volatility will interact with the NIG distribution and yield moments that are amplified compared to moments from a NIG Lévy process with constant variance.

[11] In the Gaussian case there is a natural notion of correlated Wiener processes. This is a bivariate Gaussian process, such that each coordinate is a standard Wiener process and their quadratic covariation is ρ . There is no such natural notion of correlated Wiener and NIG processes. Due to the unbounded jumps of the NIG process now even the continuous time 'variance' process can jump to negative values, and one need to resort to reflection, for example. Another possibility, that, in contrast to the Gaussian case, does not produce the same result, would be Gaussian variance innovation that mix with NIG innovations for the underlying price process.

[12] As the number of jumps in the simulation period now depends on the parameters we do not know a priory how many random innovations will be required. Thus we do not store the random draws, but rather fix them by resetting the random seed in the iterative random generator for every simulation. Varying the parameters in this model amounts now not only to a location-scale change in the state variables as in previous models, but in addition to a (deterministic) time-change.

[13] In Tompkins (2000c), the standard error for comparisons of models was determined for each model class. Thus, instead of using the GBM estimated standard error of the SSE, standard errors were estimated for each SSE for each stochastic volatility model. These were much lower and thus allowed Model 2 to be significantly better than Model 1.
[14] See, However, Hobson & Rogers (1998).

[15] Interestingly this idea does not provide a valid change of measure in continuous time. The reason is, that one can identify the parameters μ and δ with probability 1 from continuous-time observations of jumps with size going to zero, see the Girsanov Theorem for general semi-martingales, eg Jacod & Shiryaev (1987, &)167;3d and Theorem IV.4.32#, cf also Hubalek (1999) and Eberlein & Raible (2000). This is not a deficit of the model, other martingale measures are possible. It is just one aspect, where the analogy with continuous diffusion does not hold. For continuous time change of measure in the Tompkins model see Hubalek & Tompkins (2000).

[16] This manner of expressing the strike price is similar to the d_2 term that appears in the Black Scholes formula. It is common market practice in the currency options market to express strike prices in terms of the delta [$N(d_2)$] and quote implied volatilities relative to this. This approximately expresses equation 1 as a probability.

[17] The reason the Black (1976) model was chosen over the Garman-Kohlhagen (1983) model was that the underlying objective price process was modelled using Futures prices rather than Spot Foreign Exchange prices. Of potentially more concern was that the simulated implied volatility surfaces will be compared to actual implied volatility surfaces based upon American style Options. To examine the significance of this bias, the Barone-Adesi & Whaley (1987) model was used for options on all four markets. When implied volatilities were estimated for out-of-the-money options and then standardised, the degree of bias compared to the use of the Black (1976) model was negligible. Given that most common market practice is to express Foreign Exchange implied volatilities as the solution to the Black (1976) model and for the sake of simplicity, the Black (1976) model was used for all four markets.

Underlying Asset	Time Period of Analysis		Number of Observations
(Entire Period)	First Date	Last Date	
D-Mark Futures	02/01/85	13/12/99	3778
Japanese Yen Futures	02/01/85	13/12/99	3775
British Pound Futures	02/01/85	13/12/99	3781
Swiss Franc Futures	02/01/85	13/12/99	3779
(First Period)	First Date	Last Date	
D-Mark Futures	02/01/85	18/12/89	1256
Japanese Yen Futures	02/01/85	18/12/89	1251
British Pound Futures	02/01/85	18/12/89	1257
Swiss Franc Futures	02/01/85	18/12/89	1256
(Second Period)	First Date	Last Date	
D-Mark Futures	18/12/89	19/12/94	1267
Japanese Yen Futures	18/12/89	19/12/94	1267
British Pound Futures	18/12/89	19/12/94	1267
Swiss Franc Futures	18/12/89	19/12/94	1267
(Third Period)	First Date	Last Date	
D-Mark Futures	19/12/94	13/12/99	1257
Japanese Yen Futures	19/12/94	13/12/99	1259
British Pound Futures	19/12/94	13/12/99	1259
Swiss Franc Futures	19/12/94	13/12/99	1258

Table I, Markets Included in Research, Time Period of Data, Number of Observations

Futures Prices are the closing daily levels for the Nearest to Expiration Contract

Underlying Asset	Mean	Standard Deviation	Unconditional Skewness	Skewness Significance Level	Unconditional Kurtosis	Kurtosis Significance Level	Bera-Jarque Statistic	Observations
D-Mark Futures								
Overall Period	0.00009	0.00717	**0.1430**	**3.59**	**5.399**	**30.10**	**918.65**	3777
First Period	0.00034	0.00779	**0.2723**	**3.94**	**5.413**	**17.45**	**320.03**	1255
Second Period	0.00018	0.00744	-0.0774	-1.13	**4.724**	**12.52**	**158.10**	1266
Third Period	-0.00025	0.00618	**0.1726**	**2.50**	**5.813**	**20.35**	**420.38**	1256
Japanese Yen Futures								
Overall Period	0.00012	0.00761	**0.7023**	**17.61**	**10.522**	**94.32**	**9207.13**	3774
First Period	0.00033	0.00712	**0.3730**	**5.38**	**7.937**	**35.63**	**1298.35**	1250
Second Period	0.00027	0.00686	**0.3339**	**4.85**	**7.135**	**30.04**	**925.65**	1266
Third Period	-0.00022	0.00870	**1.0846**	**15.71**	**12.319**	**67.47**	**4798.54**	1258
British Pound Futures								
Overall Period	0.00022	0.00706	**0.1020**	**2.56**	**6.831**	**48.08**	**2318.24**	3780
First Period	0.00047	0.00815	**0.3548**	**5.13**	**6.198**	**23.13**	**561.42**	1256
Second Period	0.00015	0.00760	**-0.2347**	**-3.41**	**5.809**	**20.40**	**427.74**	1266
Third Period	0.00006	0.00502	-0.0203	-0.29	**6.326**	**24.08**	**580.08**	1258
Swiss Franc Futures								
Overall Period	0.00005	0.00800	**0.1769**	**4.44**	**5.098**	**26.32**	**712.42**	3778
First Period	0.00025	0.00861	**0.1998**	**2.89**	**5.001**	**14.47**	**217.79**	1255
Second Period	0.00018	0.00817	-0.0064	-0.09	**4.032**	**7.49**	**56.14**	1266
Third Period	-0.00029	0.00713	**0.3523**	**5.10**	**6.682**	**26.65**	**736.20**	1257

Table II, Statistics of the Daily Returns for Four Foreign Exchange Futures Markets (1985-1999)

These returns are based upon closing futures prices for the nearest to expiration futures contracts. The first column describes the market under investigation and the period of analysis (entire period 1985-1999) (1st period 1985-1989) (2nd period 1989-1994) (3rd period 1994-1999)The second and third columns present the first (mean) and second (standard deviation) moments of the return distribution. The fourth and sixth columns present the statistics for the unconditional skewness and kurtosis. Under the null hypothesis of normality, the skewness statistic is normally distributed with standard errors: se =sqrt(6/T), where T is the sample size. This appears in the furthest right-hand column of the table.The fifth column indicates a significance statistic testing a normality hypothesis. If the skewness statistic rejects this at a 95% level or above, this is indicated in ***BOLDED*** *text.Under the null hypothesis of normality, the excess kurtosis statistic is normally distributed with standard errors: se =sqrt(24/T). This statistic is equal to the kurtosis statistic appearing in the table minus 3.0. Column 7 indicates a significance statistic testing the null hypothesis of normality. If the kurtosis statistic rejects this at a 95% level or above, this is indicated in BOLDED text. The statistic in the sixth column is the Bera-Jarque (BJ) statistic for detecting departures of the data from normality. Under the null hypothesis of normality, the BJ statistic is distributed as Chi squared with 2 degrees of freedom. The critical value at the one-percent level is 9.21. When the BJ statistic is above this level, this statistic appears in* ***BOLDED*** *text.*

Markets		20 Day Average Volatility	20 Day SD Volatility	Coefficient of Variation	Line Fit of SD of Volatility vs. Time Horizon	Leverage Correlation (20 Day Volatility vs. Recent Relative Prices)	Average Autocorrelation of Absolute Returns (Lags 1-20)	Average Autocorrelation of Absolute Returns (Lags 51-70)
D-Mark Futures								
	Entire Period	10.81%	3.59%	**0.332**	**-0.2189**	0.1986	**0.0751**	**0.0198**
	First Period	11.73%	3.96%	**0.338**	**-0.3334**	**0.2318**	**0.0726**	**-0.0144**
	Second Period	11.37%	3.30%	**0.290**	**-0.1799**	-0.1633	**0.0576**	**0.0178**
	Third Period	9.32%	2.96%	**0.318**	**-0.1778**	**0.3405**	**0.0563**	**0.0322**
Japanese Yen Futures								
	Entire Period	11.22%	4.41%	**0.393**	**-0.2006**	**0.2429**	**0.0746**	**0.0351**
	First Period	10.66%	3.76%	**0.353**	-0.4453	0.1993	**0.0487**	-0.0021
	Second Period	10.42%	3.33%	**0.320**	**-0.2456**	**0.2141**	**0.0505**	**0.0133**
	Third Period	12.67%	5.49%	**0.434**	**-0.0612**	**0.4292**	**0.0888**	**0.0554**
British Pound Futures								
	Entire Period	10.37%	4.27%	**0.411**	**-0.1496**	-0.0309	**0.1168**	**0.0734**
	First Period	12.21%	4.47%	**0.366**	**-0.1984**	**0.2098**	**0.0882**	**0.0350**
	Second Period	11.31%	4.18%	**0.370**	**-0.1421**	**-0.4743**	**0.0976**	**0.0489**
	Third Period	7.63%	2.39%	**0.314**	**-0.2955**	0.0243	**0.0345**	**0.0255**
Swiss Franc Futures								
	Entire Period	12.14%	3.76%	**0.309**	**-0.2368**	**0.2922**	**0.0605**	**0.0226**
	First Period	13.10%	4.05%	**0.309**	**-0.3013**	**0.3145**	**0.0626**	-0.0025
	Second Period	12.62%	3.05%	**0.242**	**-0.2138**	-0.1096	**0.0338**	**0.0198**
	Third Period	10.72%	3.67%	**0.342**	**-0.2250**	**0.4978**	**0.0509**	**0.0256**
Expected GBM Attributes		10.000%	1.622%	0.162	-0.5000	0.0000	0.0000	0.0000
Standard Error of Attributes		1.622%	0.016%	0.010	0.0778	0.0999	0.0051	0.0061

Table III, Characteristics of the Unconditional Volatility (Standard Deviation) of Returns for Four Foreign Exchange Futures Markets.

The furthest left column, the Foreign Futures market and the time period of analysis appears. In the bottom two rows appears the attribute values expected from an Independent and Identically Distributed (i.i.d.) Price process associated with Geometric Brownian Motion (GBM). The standard error of the attributes is determined for a series of 500 draws from a GBM process of 1500 observations. These non-parametrically estimated standard errors allow significance testing between the empirical attributes and the assumption of a GBM process. When a T-statistic rejects this assumption at a 95% level or above, the attributes are BOLDED. In Columns Two, Three and Four, analysis of the 20 day volatility estimated on an non-overlapping basis appears. Using all available observations, daily returns (differences in the logarithm of daily closing Foreign Exchage futures prices) were estimated. With these returns, the standard deviation was estimated for a fixed time horizon of 20 days and then annualised using the sqrt(252).The estimation of the 20 day volatility was done on a non-overlapping basis. In the fifth column appears the relationship between the standard deviation of volatility and the time horizon of estimation. The Time Factor term is determined by a regression of the logarithm of the time horizon (T) on the logarithm of the unbiased standard deviation of volatility. Only the beta coefficients of the regression appears. Observed divergences in the Beta coefficient from -0.5 give an indication of the degree of the volatility of volatility persistence observed in the unconditional process. The sixth column provides information about the leverage relationship between the levels of Foreign Exchange futures prices and volatility. This is the correlation between the unconditional volatility estimated with daily log price increments and estimated on a 20 day rolling time horizon and an exponentially weighted return series, which indicates whether recent price movements are relatively high or low. In the two furthest right columns appear information about the autocorrelations in absolute returns (and volatility clustering) composite measures of autocorrelograms were required. Given the relatively rapid decay of the autocorrelations, we averaged the autocorrelations for each lag from lag 1 to 20 and separately averaged the autocorrelations for each lag from lag 51 to 70.

GBM (I.I.D) Process	CoV 20-day Vol.	Unconditional Skewness	Unconditional Kurtosis	Leverage Correlation	Auto-Corr. (1-20 lags)	Auto-Corr. (51-70 lags)	Time Decay Line Fit	SSE
Expected Results	0.1622	0.0000	3.0000	0.0000	0.0000	0.0000	-0.5000	0.0000
Average Result of 500 Simulations	0.1617	0.0000	3.0020	0.0000	-0.0014	0.0001	-0.5085	6.9806
Standard Deviation of 500 Simulations	0.0010	0.0607	0.1358	0.0999	0.0051	0.0061	0.0778	4.5646
Representative GBM Price Process (Z_1)	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	0.3636
Representative GBM Volatility Process (Z_2)	0.1586	-0.0160	3.0529	-0.0182	-0.0020	-0.0002	-0.5002	0.3647

Table IV, Attribute Values for Simulated GBM Processes, Standard Deviation of Attributes, and Two Representative Processes

The furthest left column indicates various simulations of an Independent and Identically Distributed (I.I.D.) Geometric Brownian Motion (GBM) price process. The row titled "Expected Results" is not based upon simulation but is what is expected from statistical theory. The next two rows titled "Average Results of 500 Simulations" and "Standard Deviation of 500 Simulations" indicates the sampling properties of 500 series of simulated prices of 1500 observations based upon random draws from an I.I.D. GBM process. The next two rows represent the two draws from the I.I.D. GBM process which has the lowest sum of squared errors relative to the expected theoretical results. These will be used in the second stage of the analysis as representative normal distributions and will be referred to as Z1 and Z2. Columns 2 to 8 indicate the actual attibute values. The furthest right column (the ninth column) indicates the sum of squared errors (SSE) statistic for that row. Of importance to a later stage of analysis is the average standard deviation of the SSE. This result will allow comparisons to be made between alternative models.

NIG Distribution	Mean	Standard Deviation	Skewness	Kurtosis
NIG #1	-0.01821	1.00088	0.0049	4.7807
NIG #2	-0.00078	0.93955	0.9289	8.7576
NIG #3	0.00665	1.01818	0.1936	3.4129
NIG #4	-0.0035	0.99176	0.0794	7.627

Table V, Sample Statistical Moments of Simulations of Four Normal Inverse Gaussian (NIG) Distributions.

We simulated four NIG distributions using the method suggested by Rydberg (1997). This method requires the four moments of the distribution to be input. The first moment (mean) was set to 0.0 and the second moment (variance) to 1.0. The third (skew) and fourth (excess kurtosis) moments were chosen to be less than the observed moments for daily returns of the four Foreign Exchange Futures markets. This was done, due to the fact that the stochastic volatility will interact with the NIG distribution and amplify the resultant simulated moments

Models (Parameters) — Empirical and Simulated Attributes

(Whole Period :1985-1999)

D-Mark Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3321	0.1430	5.3991	0.1986	0.0751	0.0198	-0.2189		
Simulated Attributes															
GBM	GBM	0.111	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	42.984	
Model 1	GBM	0.111	10.10	0.910	0.01	NA	0.3682	0.0329	4.8857	0.3098	0.0901	0.0201	-0.2464	1.390	**9.121**
Model 2	NIG-T3	0.111	13.10	0.690	0.14	NA	0.3602	0.2931	5.4848	0.0752	0.0808	0.0205	-0.2646	1.163	**9.171**
Model 3	Γ2-OU	0.111	12.10	1.110	0.01	0.31	0.3228	-0.2052	4.8290	0.1710	0.0971	0.0134	-0.2583	0.527	**9.311**

(First Period :1985-1989)

D-Mark Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3378	0.2723	5.4132	0.2318	0.0726	-0.0144	-0.3334		
Simulated Attributes															
GBM	GBM	0.119	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	34.719	
Model 1	GBM	0.119	13.90	0.710	-0.09	NA	0.3344	0.0425	4.2028	0.1790	0.0857	-0.0087	-0.3419	1.557	**7.272**
Model 2	NIG-T3	0.119	11.98	0.582	0.24	NA	0.3387	0.3552	5.2047	0.1611	0.0802	-0.0109	-0.3674	0.455	**7.514**
Model 3	Γ2-OU	0.119	13.00	1.000	0.20	0.40	0.2681	0.1307	5.2221	0.1285	0.0302	0.0038	-0.2828	0.672	**7.466**

(Second Period :1989-1994)

D-Mark Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.2901	-0.0774	4.7243	-0.1633	0.0576	0.0178	-0.1799		
Simulated Attributes															
GBM	GBM	0.121	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	34.197	
Model 1	GBM	0.121	10.00	1.000	-0.2	NA	0.3469	-0.2511	4.6752	-0.0252	0.0871	0.0086	-0.2595	4.708	**6.467**
Model 2	NIG-T3	0.121	3.11	0.199	-0.08	NA	0.2856	0.1094	4.2901	-0.2678	0.0604	0.0272	-0.2249	1.110	**7.256**
Model 3	Γ2-OU	0.121	8.00	0.900	0.00	0.30	0.2902	-0.4390	4.8617	-0.3767	0.0587	0.0316	-0.1768	0.687	**7.349**

(Third Period :1994-1999)

D-Mark Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3180	0.1726	5.8131	0.3405	0.0563	0.0322	-0.1778		
Simulated Attributes															
GBM	GBM	0.091	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	43.948	
Model 1	GBM	0.091	12.50	1.250	0	NA	0.3477	0.0609	5.4150	0.3284	0.0560	0.0005	-0.3466	6.283	**8.260**
Model 2	NIG-T1	0.091	0.72	0.100	0.84	NA	0.3275	0.3234	6.4270	0.3628	0.0507	0.0397	-0.2183	0.789	**9.465**
Model 3	Γ2-OU	0.091	12.11	1.111	0.20	0.29	0.3169	0.6039	6.5835	0.0931	0.0586	-0.0023	-0.2347	1.917	**9.217**

Table VI a, Empirical Attributes of D-Mark Futures and Attributes of Four Alternative Models

This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the D-Mark Futures. These results are split into the four periods of analysis. In each of these tables, the four time periods of analysis appear. On the left-hand side, the parameterisation of each of the four models appears. Only the parameter values and the NIG distribution that has the lowest SSE are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison's sake. In the column immediately to the right of this the sum of squared errors (SSE) for each model appears. The final column indicates a T-test statistic comparing models 2, 3 and 4 to the GBM assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV.

Models (Parameters) — Empirical and Simulated Attributes

(Whole Period :1985-1999)

Yen Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3930	0.7023	10.5218	0.2429	0.0746	0.0351	-0.2006		
Simulated Attributes															
GBM	GBM	0.117	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	76.956	
Model 1	GBM	0.117	15.00	1.000	0.11	NA	0.3722	0.2155	5.2408	0.4269	0.0801	0.0096	-0.3048	10.754	**14.518**
Model 2	NIG-T2	0.117	2.80	0.377	-0.28	NA	0.3965	0.7244	8.4363	0.1284	0.0794	0.0404	-0.2501	1.237	**16.605**
Model 3	Γ2-OU	0.117	11.01	1.201	0.30	0.30	0.3663	0.9564	8.3991	0.2493	0.0830	0.0235	-0.1664	2.147	**16.406**

(First Period :1985-1989)

Yen Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3525	0.3730	7.9368	0.1993	0.0487	-0.0021	-0.4453		
Simulated Attributes															
GBM	GBM	0.114	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	34.574	
Model 1	GBM	0.114	14.95	1.250	0.09	NA	0.3486	0.2470	5.1273	0.3170	0.0457	0.0092	-0.4072	2.536	**7.026**
Model 2	NIG-T1	0.114	10.00	1.000	0.10	NA	0.3926	0.1804	7.7378	0.2225	0.0557	0.0049	-0.4383	1.304	**7.296**
Model 3	Γ2-OU	0.114	9.00	1.400	0.10	0.20	0.3410	0.5852	7.7079	0.0317	0.0360	0.0058	-0.4655	1.482	**7.257**

(Second Period :1989-1994)

Yen Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3197	0.3339	7.1354	0.2141	0.0505	0.0133	-0.2456		
Simulated Attributes															
GBM	GBM	0.108	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	38.246	
Model 1	GBM	0.108	17.60	1.240	-0.01	NA	0.3399	0.1952	5.1930	0.2568	0.0441	0.0036	-0.3769	3.150	**7.696**
Model 2	NIG-T1	0.108	1.28	0.137	0.73	NA	0.3198	0.3088	6.3781	0.3183	0.0457	0.0348	-0.2541	1.327	**8.096**
Model 3	Γ2-OU	0.108	9.00	1.200	0.10	0.20	0.3461	0.1669	7.5752	-0.0662	0.0481	-0.0046	-0.3613	0.969	**8.175**

(Third Period :1994-1999)

Yen Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.4336	1.0846	12.3188	0.4292	0.0888	0.0554	-0.0612		
Simulated Attributes															
GBM	GBM	0.131	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	123.442	
Model 1	GBM	0.131	10.00	0.980	0.09	NA	0.4255	0.1908	5.5015	0.4423	0.1287	-0.0136	-0.2609	25.738	**21.426**
Model 2	NIG-T2	0.131	0.98	0.179	0.60	NA	0.4365	1.2909	11.9921	0.3923	0.0773	0.0720	-0.1196	1.718	**26.694**
Model 3	Γ2-OU	0.131	10.96	1.296	0.30	0.39	0.4334	1.2695	9.9088	0.5109	0.1143	0.0069	-0.1588	3.804	**26.236**

Table VI b, Empirical Attributes of Japanese Yen Futures and Attributes of Four Alternative Models

This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the Yen Futures. These results are split into the four periods of analysis. In each of these tables, the four time periods of analysis appear. On the left-hand side, the parameterisation of each of the four models appears. Only the parameter values and the NIG distribution that has the lowest SSE are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison's sake. In the column immediately to the right of this the sum of squared errors (SSE) for each model appears. The final column indicates a T-test statistic comparing models 2, 3 and 4 to the GBM assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV.

Models (Parameters) — Empirical and Simulated Attributes

(Whole Period :1985-1999)

Pound Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.4114	0.1020	6.8311	-0.0309	0.1168	0.0734	-0.1496		
Simulated Attributes															
GBM	GBM	0.106	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	94.463	
Model 1	GBM	0.106	8.10	0.740	-0.08	NA	0.4080	-0.1738	5.0169	0.1148	0.1312	-0.0070	-0.2424	3.947	**19.850**
Model 2	NIG-T1	0.106	2.20	0.226	0.04	NA	0.4123	0.0895	6.9877	-0.0541	0.1127	0.0813	-0.1873	0.362	**20.636**
Model 3	Γ2-OU	0.106	8.81	0.841	-0.20	0.32	0.3606	-0.0961	5.2228	0.0201	0.1077	0.0424	-0.1389	4.584	**19.710**

(First Period :1985-1989)

Pound Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3660	0.3548	6.1976	0.2098	0.0882	0.0350	-0.1984		
Simulated Attributes															
GBM	GBM	0.119	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	59.901	
Model 1	GBM	0.119	15.90	0.910	0.14	NA	0.3739	0.3199	5.1857	0.4790	0.0813	0.0169	-0.2764	3.153	**12.445**
Model 2	NIG-T3	0.119	10.90	0.651	0.06	NA	0.3623	0.2079	5.3089	0.0178	0.0910	0.0226	-0.2551	1.864	**12.727**
Model 3	Γ2-OU	0.119	11.00	1.200	0.10	0.40	0.3864	0.4755	6.1840	0.1750	0.1156	-0.0024	-0.3541	0.644	**12.995**

(Second Period :1989-1994)

Pound Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3697	-0.2347	5.8086	-0.4743	0.0976	0.0489	-0.1421		
Simulated Attributes															
GBM	GBM	0.125	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	73.220	
Model 1	GBM	0.125	4.90	1.510	-0.19	NA	0.3542	-0.2431	4.7286	-0.0143	0.0865	-0.0316	-0.3483	5.913	**14.760**
Model 2	NIG-T3	0.125	4.70	0.285	-0.70	NA	0.3672	-0.1476	5.4318	-0.5817	0.1008	0.0493	-0.1668	0.423	**15.964**
Model 3	Γ2-OU	0.125	9.95	1.205	-0.11	0.41	0.3388	-0.1653	5.6058	-0.3260	0.0782	0.0273	-0.1301	2.027	**15.613**

(Third Period :1994-1999)

Pound Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
Empirical Attributes							0.3135	-0.0203	6.3264	0.0243	0.0345	0.0255	-0.2955		
Simulated Attributes															
GBM	GBM	0.077	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	28.486	
Model 1	GBM	0.077	12.60	0.760	-0.11	NA	0.3428	-0.0533	4.8869	0.2020	0.0596	-0.0140	-0.3446	2.340	**5.734**
Model 2	NIG-T1	0.077	4.95	0.141	-0.05	NA	0.3022	0.0667	5.8321	0.0263	0.0391	0.0365	-0.3218	0.586	**6.118**
Model 3	Γ2-OU	0.077	9.00	0.900	0.10	0.20	0.3099	-0.3762	6.6930	-0.2065	0.0478	0.0114	-0.2999	0.934	**6.042**

Table VI c, Empirical Attributes of British Pound Futures and Attributes of Four Alternative Models

This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the Pound Futures. These results are split into the four periods of analysis. In each of these tables, the four time periods of analysis appear. On the left-hand side, the parameterisation of each of the four models appears. Only the parameter values and the NIG distribution that has the lowest SSE are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison's sake. In the column immediately to the right of this the sum of squared errors (SSE) for each model appears. The final column indicates a T-test statistic comparing models 2, 3 and 4 to the GBM assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV.

Models (Parameters) — Empirical and Simulated Attributes

Swiss Franc Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
(Whole Period :1985-1999)							Empirical Attributes								
							0.3095	0.1769	5.0977	0.2922	0.0605	0.0226	-0.2368		
							Simulated Attributes								
GBM	GBM	0.124	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	35.057	
Model 1	GBM	0.124	9.90	1.100	-0.13	NA	0.3208	-0.1175	4.6165	0.1823	0.0672	0.0102	-0.2710	1.801	**7.293**
Model 2	NIG-T1	0.124	0.50	0.132	0.71	NA	0.3154	0.2527	6.1567	0.2429	0.0469	0.0375	-0.2540	1.394	**7.382**
Model 3	Γ2-OU	0.124	11.05	0.985	0.22	0.34	0.2805	0.4051	4.6788	0.1620	0.0443	0.0038	-0.3732	1.364	**7.389**

Swiss Franc Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
(First Period :1985-1989)							Empirical Attributes								
							0.3094	0.1998	5.0013	0.3145	0.0626	-0.0025	-0.3013		
							Simulated Attributes								
GBM	GBM	0.132	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	29.298	
Model 1	GBM	0.132	16.70	0.830	-0.04	NA	0.3097	0.0987	4.1226	0.2841	0.0588	-0.0035	-0.3332	0.518	**6.311**
Model 2	NIG-T3	0.132	13.90	0.631	0.21	NA	0.3184	0.2859	4.8025	0.1650	0.0723	-0.0069	-0.3672	1.274	**6.146**
Model 3	Γ2-OU	0.132	10.00	1.100	0.00	0.30	0.3094	0.2470	5.1937	0.1864	0.0682	0.0092	-0.2816	0.698	**6.272**

Swiss Franc Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
(Second Period :1989-1994)							Empirical Attributes								
							0.2417	-0.0064	4.0316	-0.1096	0.0338	0.0198	-0.2138		
							Simulated Attributes								
GBM	GBM	0.132	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	21.198	
Model 1	GBM	0.132	15.00	0.520	-0.21	NA	0.2607	-0.0359	3.5441	0.0949	0.0471	-0.0045	-0.3358	4.505	**3.661**
Model 2	NIG-T3	0.132	1.67	0.153	0.15	NA	0.2404	0.1689	3.8496	-0.1515	0.0348	0.0241	-0.2261	0.427	**4.555**
Model 3	Γ2-OU	0.132	11.00	1.200	-0.10	0.20	0.2417	-0.1375	4.2806	0.0266	0.0360	0.0086	-0.2760	1.285	**4.367**

Swiss Franc Futures	Distribution	LTV	K	VoV	Correl	Weight	CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	SSE	T-Statistic
(Third Period :1994-1999)							Empirical Attributes								
							0.3424	0.3523	6.6824	0.4978	0.0509	0.0256	-0.2250		
							Simulated Attributes								
GBM	GBM	0.106	0.00	0.000	0	NA	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	46.834	
Model 1	GBM	0.106	17.50	1.250	0.1	NA	0.3490	0.1518	5.0570	0.3811	0.0575	0.0175	-0.2547	1.928	**9.848**
Model 2	NIG-T1	0.106	1.73	0.146	0.93	NA	0.3354	0.4093	6.5875	0.5107	0.0509	0.0334	-0.2381	0.205	**10.226**
Model 3	Γ2-OU	0.106	9.59	0.939	0.24	0.26	0.3069	0.5865	5.7056	0.2077	0.0566	0.0223	-0.1520	2.151	**9.799**

Table VI d, Empirical Attributes of Swiss Franc Futures and Attributes of Four Alternative Models

This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the Swiss Franc Futures. These results are split into the four periods of analysis. In each of these tables, the four time periods of analysis appear. On the left-hand side, the parameterisation of each of the four models appears. Only the parameter values and the NIG distribution that has the lowest SSE are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison's sake. In the column immediately to the right of this the sum of squared errors (SSE) for each model appears. The final column indicates a T-test statistic comparing models 2, 3 and 4 to the GBM assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV.

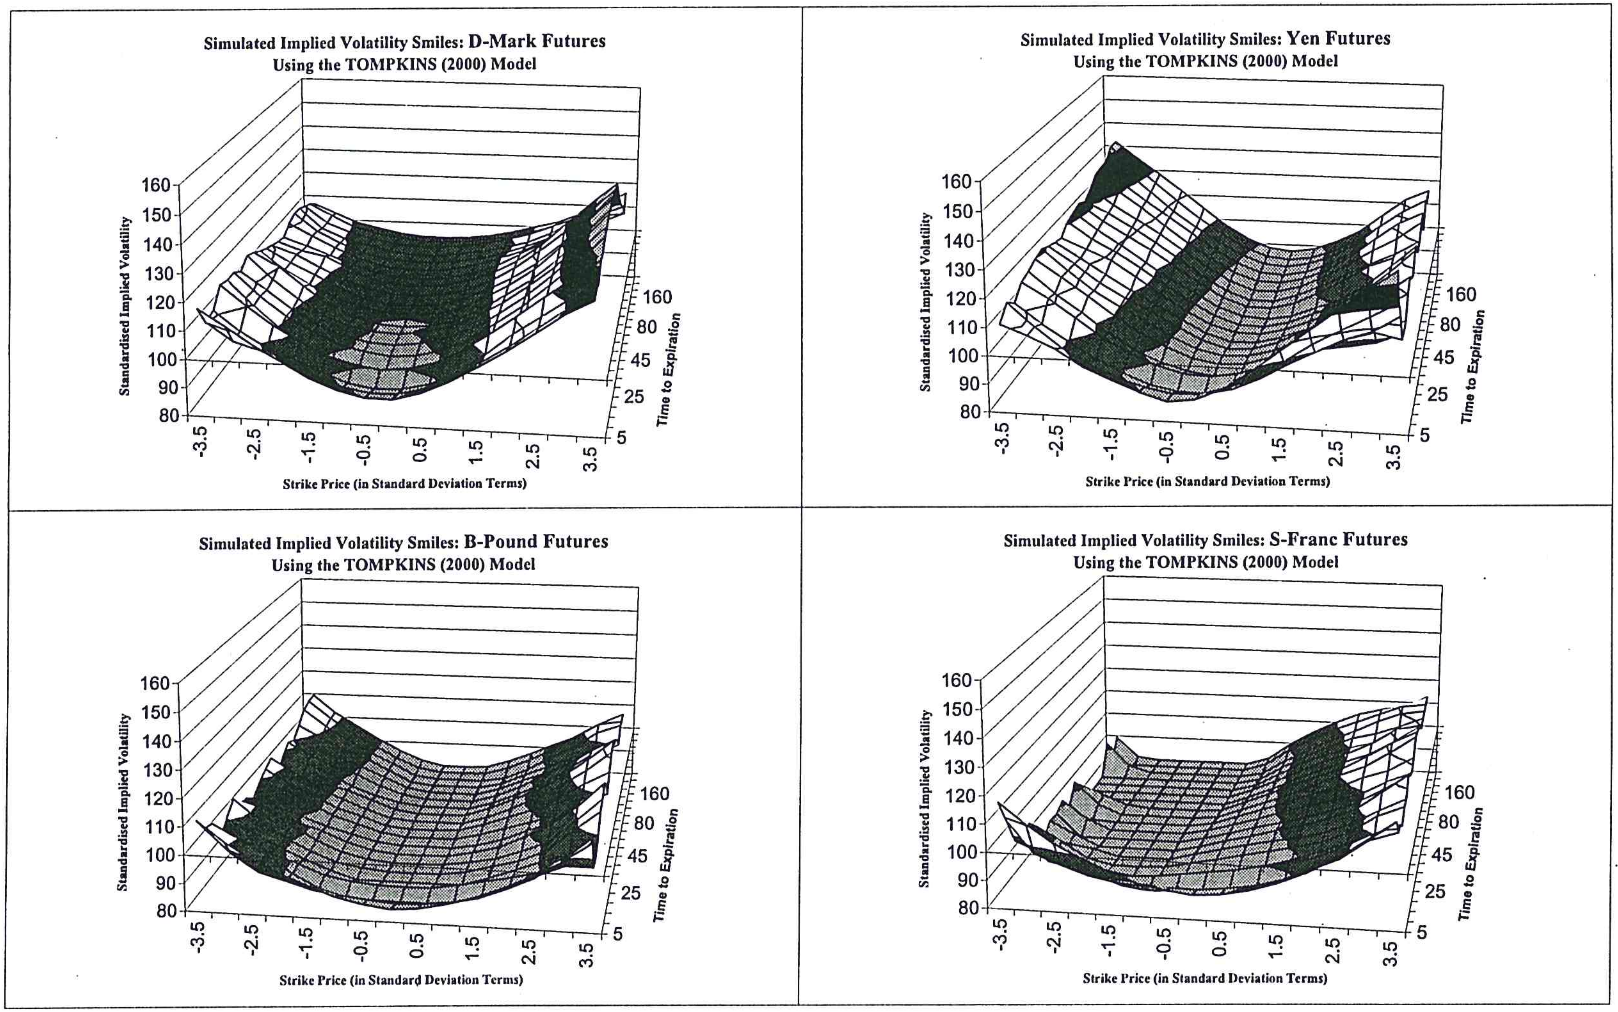

Figure 1a, Implied Volatility Surfaces of Option Prices Following a NIGSVρ Model Compared to the Constant Volatility of the Black (1976) Model for Four Foreign Exchange Futures Markets [100 = Black (1976) volatility].

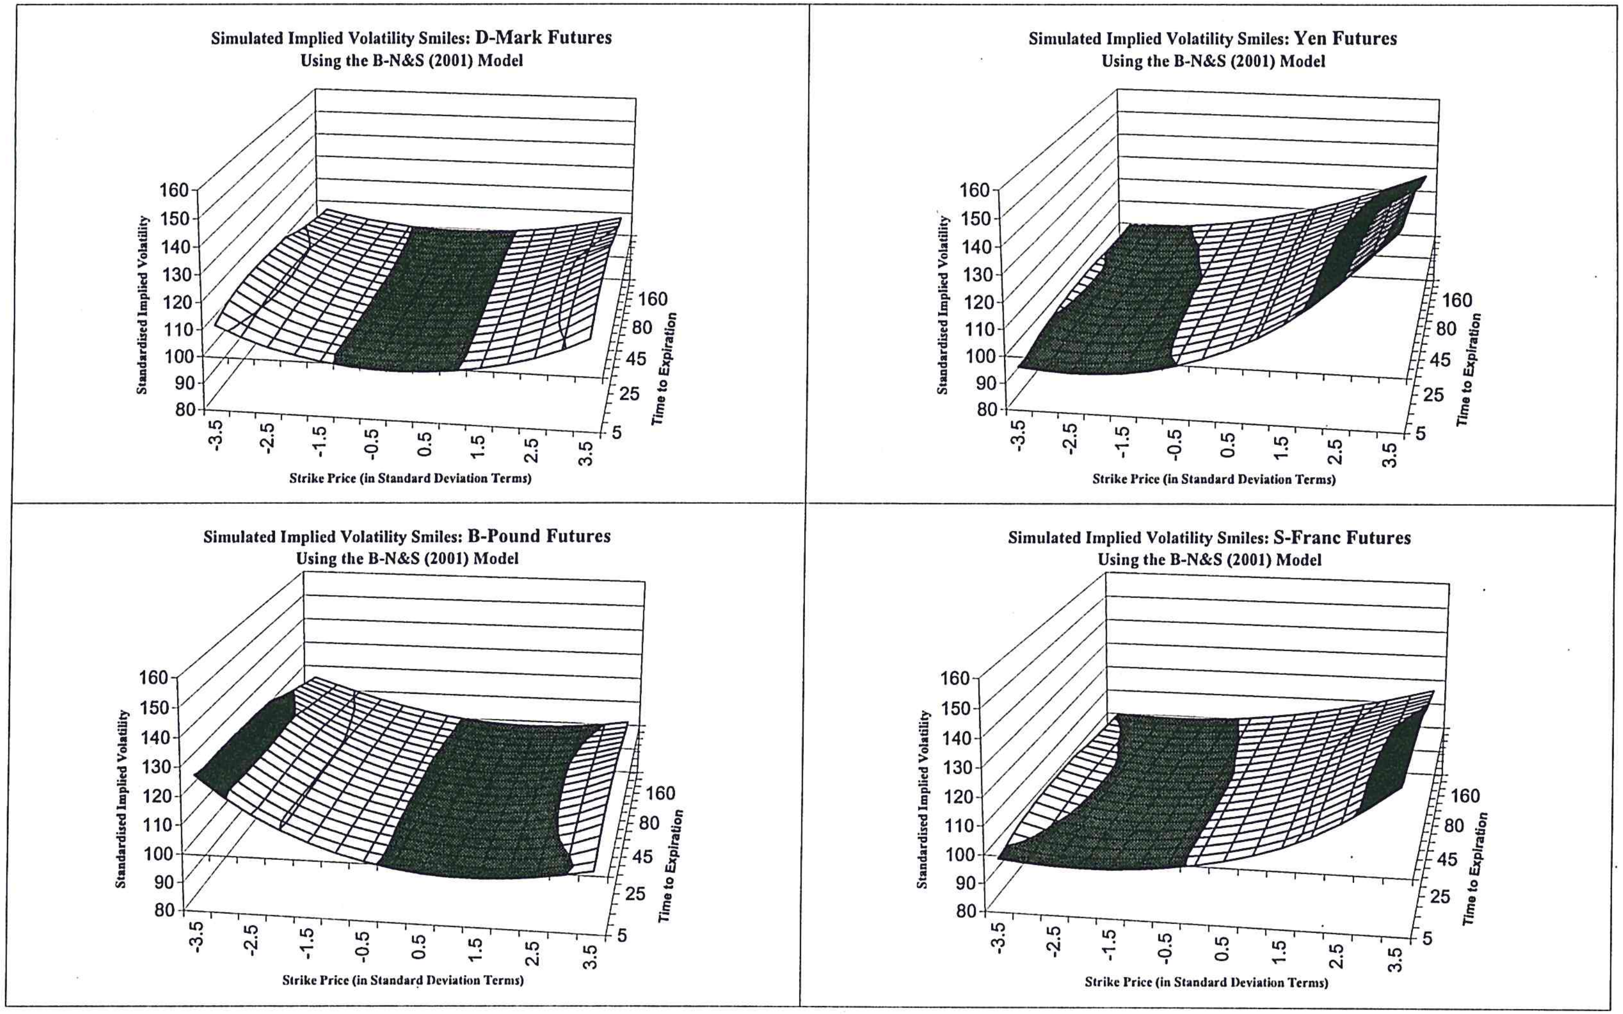

Figure 1b, Implied Volatility Surfaces of Option Prices Following a OU-2Γ Model Compared to the Constant Volatility of the Black (1976) Model for Four Foreign Exchange Futures Markets [100 = Black (1976) volatility].

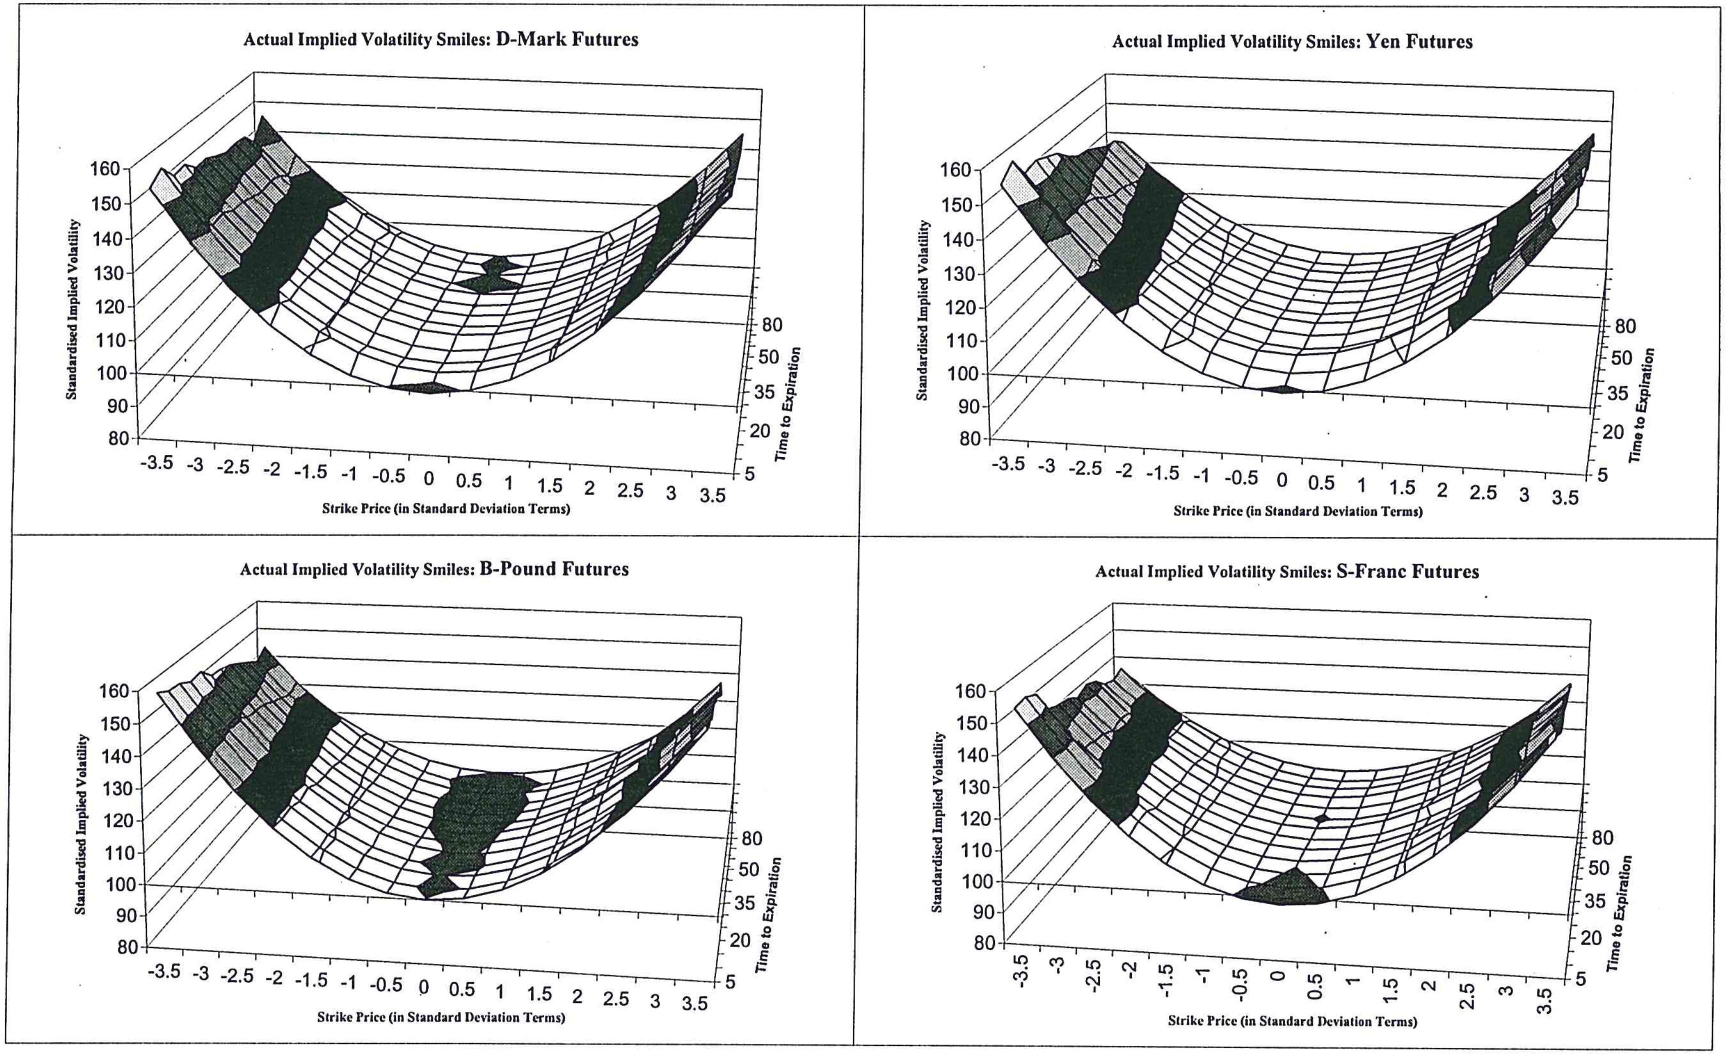

Figure 2, Actual Implied Volatility Surfaces of Option Prices for Four Foreign Exchange Futures Standardised to the Level of the ATM Volatility (1985-2000)

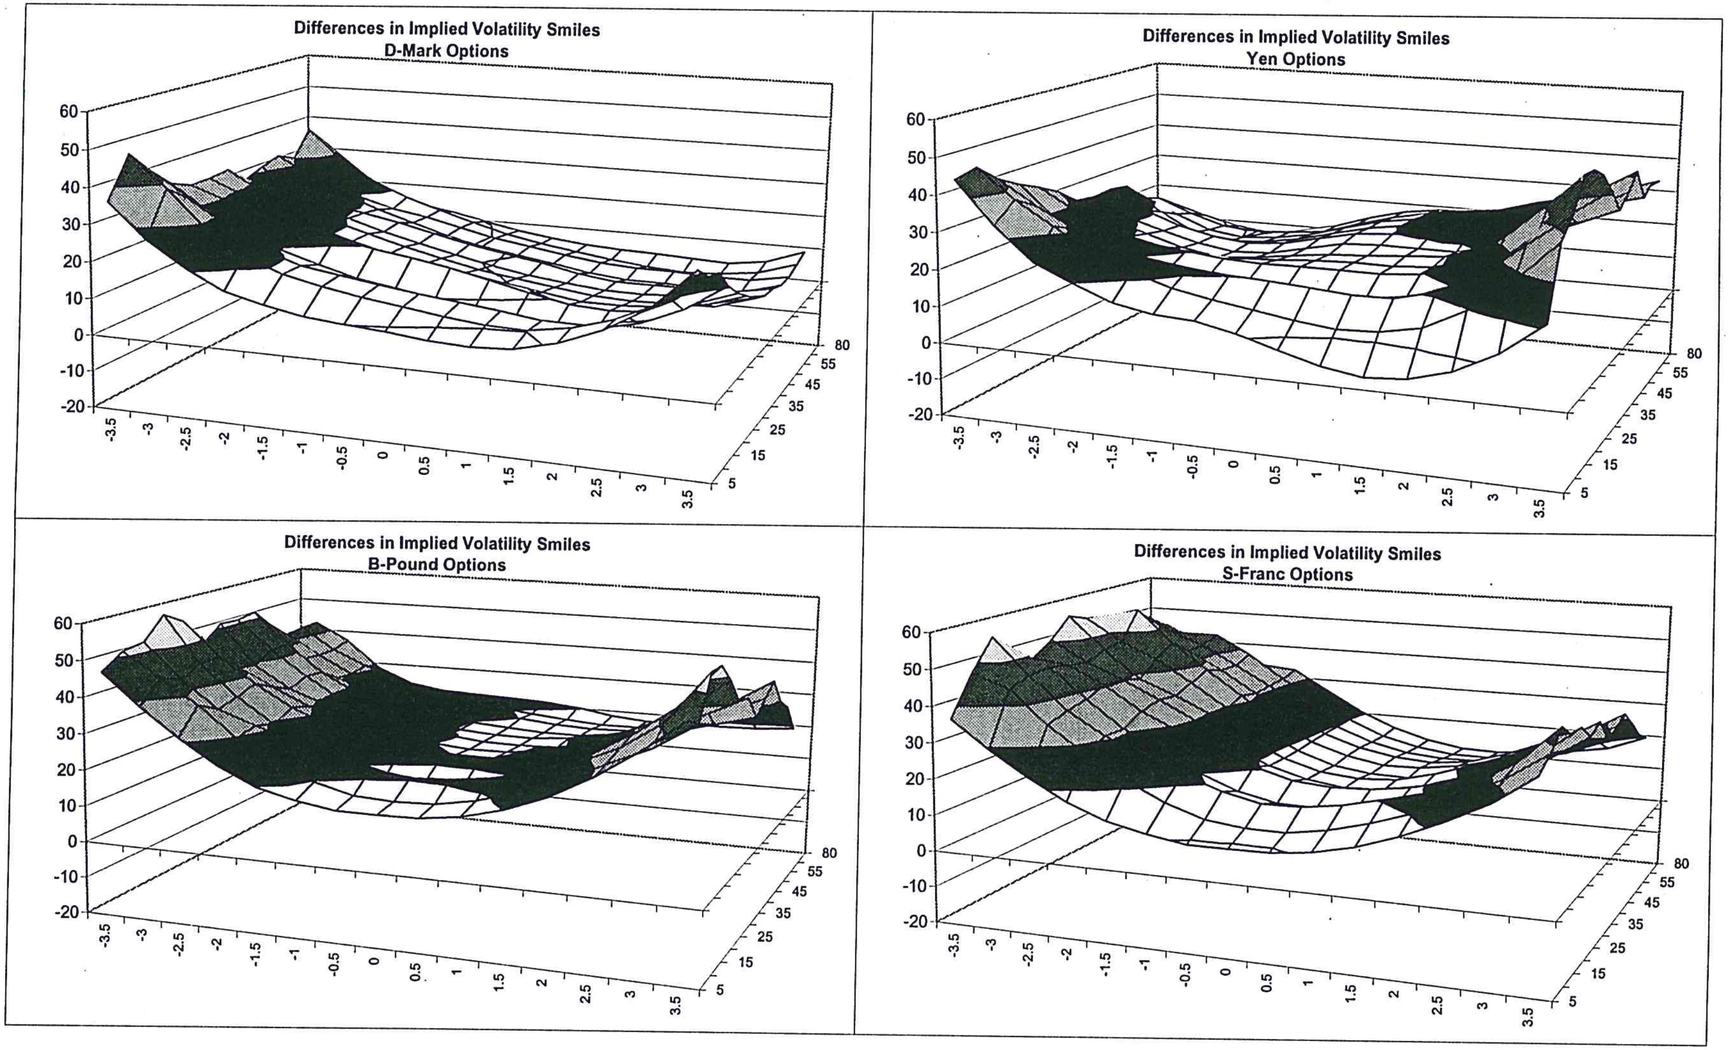

Figure 3a, Differences in Simulated & Actual Implied Volatility Surfaces for Four Foreign Exchange Futures/Options Markets (1985-2000) USING Tompkins (2000) Model

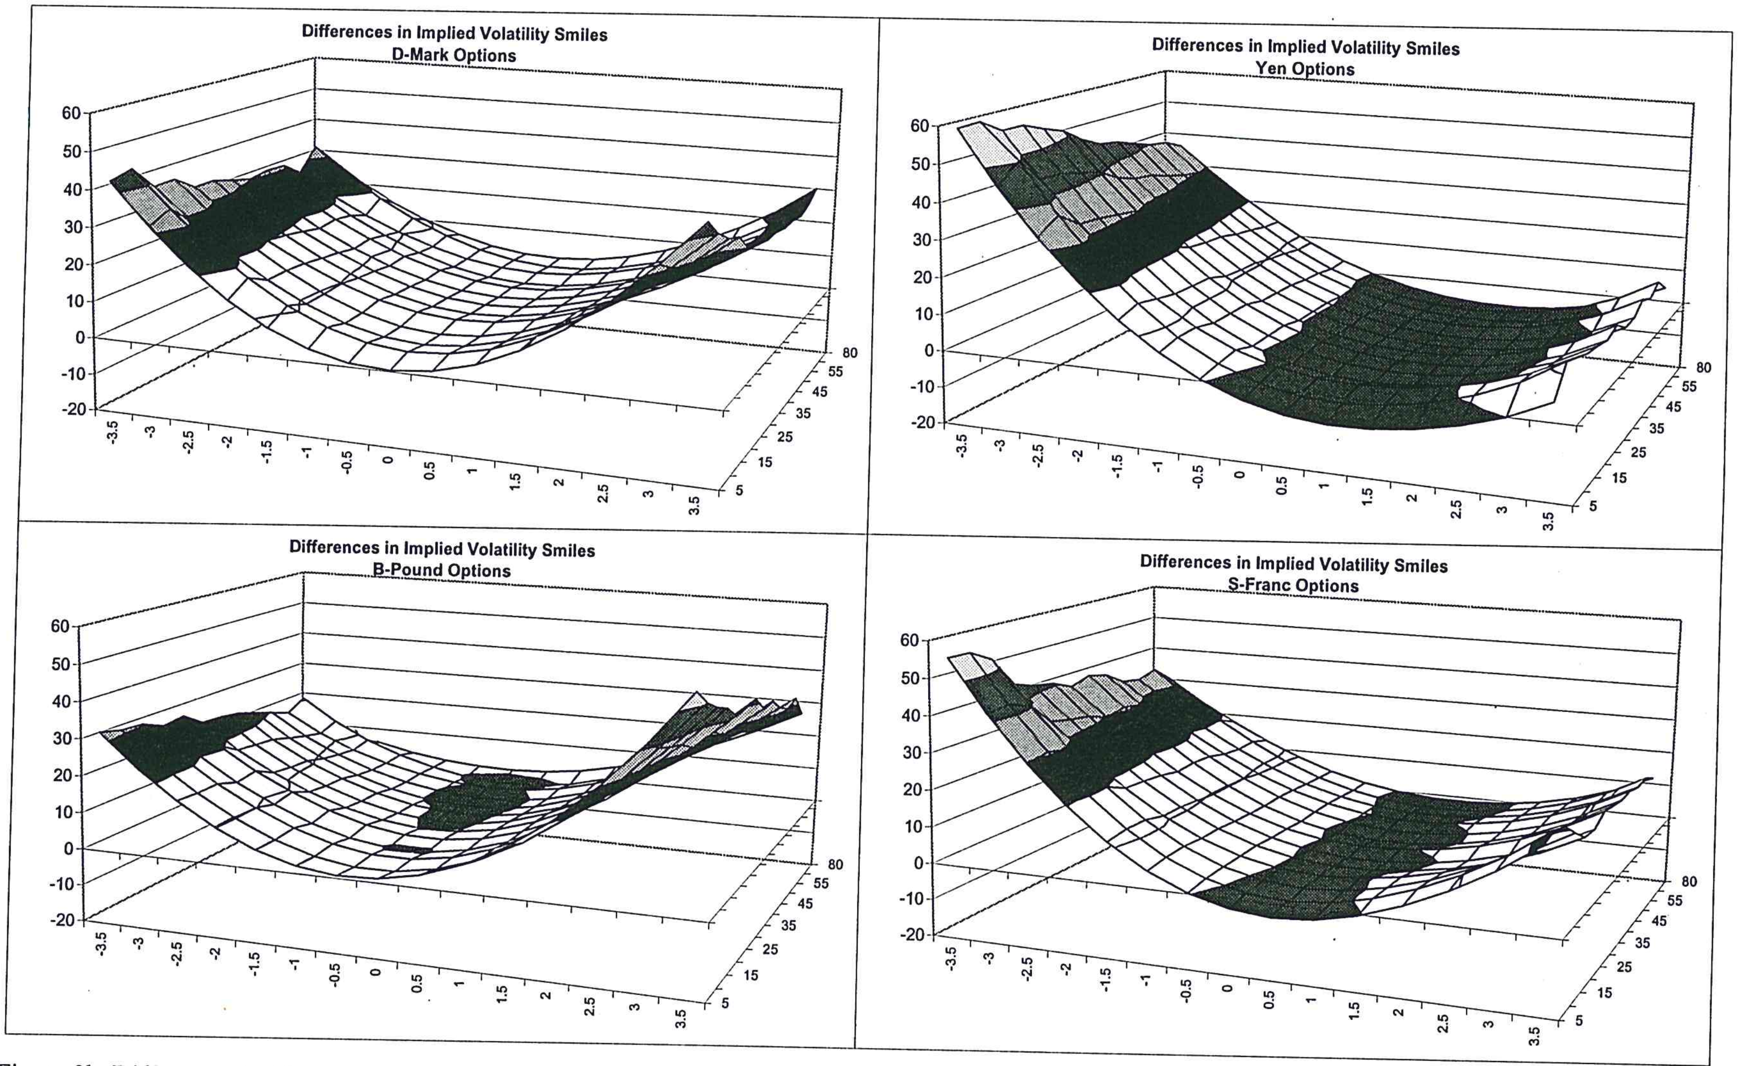

Figure 3b, Differences in Simulated & Actual Implied Volatility Surfaces for Four Foreign Exchange Futures/Options Markets (1985-2000) USING Barndorff-Nielsen & Shephard (2001) Model