FINANCIAL OPTIONS RESEARCH CENTRE

University of Warwick

Stock Index Futures Markets: Stochastic Volatility Models and Smiles

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: 2000/105

STOCK INDEX FUTURES MARKETS: STOCHASTIC VOLATILITY MODELS AND SMILES†

Robert G. Tompkins
Visiting Professor, Department of Finance
Vienna University of Technology* and
Permanent Visiting Professor, 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).

* Böcklinstrasse 7/1/7, A-1020 Wien, Austria, Phone: +43-1-726-0919, Fax: +43-1-729-6753, Email: rtompkins@ins.at

STOCK INDEX FUTURES MARKETS: STOCHASTIC VOLATILITY MODELS AND SMILES

ABSTRACT

This paper asks why implied volatility smiles exist for options on Stock Index Futures. In the literature, two possible reasons have been proposed to explain this phenomenon: the existence of market imperfections and that the underlying price process differs from the lognormal diffusion process assumed by the Black Scholes (1973) model. This paper examines the latter hypothesis.

Four stock index futures markets were considered. From the return and unconditional volatility series for these markets, seven key attributes were chosen capturing the empirical non-normal and non-IID characteristics of these markets. Using these attributes, comparisons were made with attributes consistent with the Heston (1993) stochastic volatility model. Four variants of the stochastic volatility model were examined: two models assumed the underlying price series follows GBM and two models assumed a Normal Inverse Gaussian (NIG) distribution (as a surrogate for a skewed jump process). For both alternative price processes, uncorrelated or negatively correlated subordinated stochastic volatility processes were considered.

The alternative models were evaluated using a simulated method of moments approach. Model comparisons were based upon minimisation of the sum of squared errors for the seven attributes (consistent with the model and observed empirically). For all four stock index futures markets, models including a negatively correlated stochastic volatility process with non-normal price innovations performed best within the total sample period. This result is time invariant, consistently observed when the data set is split into sub-periods.

Price series were simulated from these optimal stochastic volatility models and Monte Carlo values of European options determined. These option prices were then expressed as implied volatility surfaces and compared to the actual implied volatility surfaces for options on these same four Stock Index Futures. Comparisons between simulated and actual implied volatility surfaces for the four established Stock Index options markets suggest that the alternative price process hypothesis is insufficient to explain the existence of smiles. However, consistent divergences between simulated and actual implied volatility surfaces were found. This may provide insights into the nature of market imperfections and risk premia.

JEL classifications: C15, G13

Keywords: Price Dispersion Processes, Stochastic Volatility, Skewed Jump Processes, Normal Inverse Gaussian Distributions, Methods of Moments Estimation, Monte Carlo Simulations, Option Pricing, Implied Volatility Smiles.

1. INTRODUCTION

This paper examines the nature of the objective dispersion processes for stock index futures contracts. This paper considers alternative stochastic volatility models to explain the statistical and time series properties of the unconditional price and volatility processes and explores the implications for option pricing.

The paper is organised as follows: the first part briefly reviews previously presented empirical evidence, which indicates that stock index futures prices do not follow an i.i.d. GBM process. Four possible models, which have been proposed in the literature to explain these results, were considered. 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, each of the possible models were examined using a simulated method of moments approach to assess their ability to explain the key attributes. Having determined the best model, option prices were estimated consistent with these processes and 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 four stock index futures markets. Finally, conclusions and suggestions for further research appear.

2. PRICE PROCESSES FOR STOCK INDEX FUTURES UNDER THE 'EMPIRICAL' MEASURE

It is well established that the unconditional return series for individual stocks, stock indices and Stock Index Futures do not conform to the assumptions of an i.i.d. lognormal dispersion process [see Stoll & Whaley (1990)]. As with many financial series, returns for stock index futures display excess kurtosis and (for many periods) significantly negative skewness compared to a normal distribution. Regarding the serial correlation of returns, it is now well established that for stock markets, returns themselves contain little serial correlation [Fama (1970), Taylor (1986)]. This has also been found for stock index futures returns [Randolph and Najand (1991)]. However, Ding, Granger and Engle (1993) point out that a return series can be serially uncorrelated but can still remain dependent. They extend the findings of Taylor (1986) by examining the serial correlation of absolute returns. This can be interpreted as a measure of volatility clustering. Ding, Granger and Engle (1993) conclude, "It is clear that the S&P 500 stock market return process is not an i.i.d. process" (page 87).

Another approach to better understand the volatility process is to examine the statistical moments of the unconditional volatility series. Burghardt and Lane (1990) 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. If the empirical coefficients of variation of volatility were higher, this could be interpreted as a rejection of an i.i.d. process[1]. 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. A simple log-linear form was chosen to capture these dynamics. The rate of decay in the variability 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 leverage effect pointed out by Christie (1982) among others.

While a number of theories have been proposed to explain these results, three alternative hypotheses will be examined here. The Constant Elasticity of Variance (COV) model of Cox and Ross (1976), a non-GBM price process model [such as the Jump diffusion model proposed by Merton (1976)] and Stochastic Volatility Models were considered. These three models can be nested within a stochastic volatility model. Given the wide range of stochastic volatility models proposed in the literature [see Taylor (1994)], it was not obvious which model to select. As models with correlated processes were considered, 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 (1999)].

Specifically, the following five models were considered:

MODEL 1 $$dF = \mu F dt + \sigma F dZ(t) \quad (1)$$

Where Z(t) is a standard Wiener Process and μ and σ are constants, the return series r_t is normally distributed with $r_t = \mu + \sigma Z_t$ and $Z_t \sim$ N(0,1). This is the assumption of the Black (1976) model of i.i.d. Geometric Brownian Motion and will be referred to as GBM.

It is clear that Model 1 is a straw man, given the well-known results that Stock Index 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 ($\sqrt{V}$).

MODELS 2 & 3 $$dF = \mu F dt + \hat{\sigma} F dZ_1 \tag{2.1}$$

With the variance process defined by:

$$dV = \kappa(\theta - V)dt + \xi\sqrt{V}dZ_2 \tag{2.2}$$

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 $\sqrt{V}$ represent the variance and the volatility of the process, respectively. Model 2 assumes no correlation between the Wiener processes and Model 3 allows for non-zero correlations. Model 2 is referred to as SV, and model 3 as SVρ. These are two variants of the model proposed by Heston (1993).

MODELS 4 & 5 $$dF = \mu F dt + \hat{\sigma} F d\mathrm{N} \tag{3.1}$$

With the variance process defined by:

$$dV = k(\theta - V)dt + \xi\sqrt{V}dZ \tag{3.2}$$

where N represents a non-normal price process for the underlying price series and in this research is the Normal Inverse Gaussian distribution (NIG). This model is related to that of Barndorff-Nielsen (1997) 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, 1999) and Ho, Perraudin and Sørensen (1996), who assumed the volatility process is subordinated in a non-normal price process. In this model, the stochastic variance process is assumed to follow a standard Wiener Processes, Z, with correlation ρ existing between the two processes.[2] The notation remains as in Models 2 and 3. As with the previous models, Models 4 and 5 differ in that the former assumes no correlation between the Wiener processes and the latter allows for non-zero correlations. Model 5 extends and tests a model for incorporating a leverage effect in a stochastic volatility model proposed by Barndorff-Nielsen and Shephard (1999). Model 4 is referred to as NIGSV and model 5 as NIGSVρ. The COV model of Cox and Ross (1976) which allows for the inclusion of negative correlations between the two stochastic processes is nested in Models 3 and 5. This allows the leverage effect to be captured.

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 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 key attributes, simulates price processes consistent with the alternative models and assesses the sum of squared errors between simulated and empirical attributes.[3] 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 Model 1). 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 critical 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 negative skewness is an important attribute for describing the returns of stock indices and stock index futures. Both Theodossiou (1998) and Harvey and Siddique

(1998) chose to examine the skewness in addition to the excess kurtosis. This attribute would provide evidence for the existence of asymmetric jumps and/or leverage effects.

To examine both of these moments, price changes were estimated based upon continuously compounded returns in the standard manner [$r_t = \ln(P_t / P_{t-1})$]. With P_t and P_{t-1} referring to the level of the underling instrument at date t and t-1, respectively. Extreme care was taken to assure that returns were estimated using only futures prices with the same expiration date[4]. With this time series of daily returns, the unconditional skewness and kurtosis statistics were determined.

Clearly, given that this research examined four variants of the Heston (1993) stochastic volatility model, 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. This can be expressed as: $\sigma_N = \sqrt{\frac{\sum_{t=1}^{N}(r_t)^2}{N}}$.[5] This measure was annualised using the number of trading days in a year (√252) and N is the number of observations (horizon period). Even though alternative methods have been suggested to estimate volatility, closing prices were selected as the basis for determination.[6]

However, a critical assumption made when estimating the (sample) volatility is that the ultimate population is stationary. If the return series are heteroscedastic, the volatility will no longer be an unbiased estimate of the true standard deviation. The fact that such heteroscedasticity is commonly observed, has been a key motivation for the development of stochastic volatility models.

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[7].

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.[8] From statistical theory, the expected decay in the standard deviation of a volatility estimate ($\sigma*$) 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:

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

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

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

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.[9]

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 the leverage 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.[10] 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:

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

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.[11] This correlation coefficient serves as an attribute to measure the leverage effect.

With these seven target attributes, the stochastic volatility models were parameterised via simulation. Following the five 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 stock index 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[12]. 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 stock index 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 stock index futures markets for the relevant period of analysis.

Finally, 500 samples of 1500 prices consistent with Model 1 (GBM) were drawn to understand the sample properties of all the attributes and of the test statistic better. 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 S&P 500 and Nikkei 225 traded at the Chicago Mercantile Exchange (CME), the DAX index futures traded at the Deutsche Terminbörse (DTB) and the FTSE 100 futures traded at London International Financial Futures Exchange (LIFFE). Closing futures prices were obtained directly from the relevant exchange and the time period of analysis was restricted to a period in which all four stock index futures were traded.[13] 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 equal time periods from 1990-1994 and from 1995-1999. 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 stock index futures and options markets were examined was to allow a cross-sectional comparison. Apart from the benefit of assessing general tendencies across markets, it is also possible to use anomalous results as an additional check on data validity. This assured that the data series employed in this research was as accurate as is humanly 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.[14] 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.[15] 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. However, the skewness is neither consistently negative for all markets nor always significant. 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. Recently, Theodossiou (1998) examined the return series for a variety of financial assets. He reports significant excess kurtosis for all the assets but neither consistently negative nor significant levels of skewness. Harvey and Siddique (1998) found similar results when examining the return statistics for stock market indices for the US, German and Japanese markets.

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 (based on non-overlapping observations). 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 i.i.d. 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 each market, 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 (determined by an exponential moving average of returns). Consistent with the negative leverage effects Christie (1982) pointed out for individual stocks, stock index futures markets also seem to have a significant negative leverage effect. The exceptions are the FTSE 100 and S&P 500 Futures over the first period (1990-1994). 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 and for all time periods, these averaged autocorrelations are statistically significantly positive. However, the degree of autocorrelations seems to be higher in the second period relative to the first 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 i.i.d. process. With these attributes, 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.

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. Specifically, this entailed the generation of 500 samples of 1500 draws from an independent normal distribution. According to standard procedures, the random number generation process used a standard Box-Muller method and the anti-thetic approach suggested by Boyle (1977). With these 500 samples, price series were constructed that conformed to GBM with constant variance (using equation 1). The distributional and time series attributes of each series were assessed and compared to the theoretical moments of an i.i.d. GBM process. Utilising formula 6, the sum of squared errors between each of the 500 samples and the true attributes of an i.i.d. GBM process were determined. The two distributions with the lowest sums of squared errors relative to the priors were selected as representative normal distributions. For the sake of convenience, the normal disturbances for the underlying price process will be referred to as Z_1 and the normal disturbances for the volatility process will be referred to as Z_2. These two series were uncorrelated. Thereafter, whenever simulation used either of these distributional forms, the same sets of random numbers were used (to reduce errors introduced by the selection of random numbers). 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.

To generate the NIG distributions, random numbers were generated using the method suggested by Rydberg (1997). These simulations required the input of the four moments of the distribution. 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 in Table III. This was done, due to the fact that the stochastic volatility will interact with the NIG distribution and yield simulated moments that are amplified compared to the NIG moments. Given these effects could not be ascertained prior to the simulation, four NIG distributions were simulated and all were examined as potential candidates for the underlying price process.

[Table V appears here]

This approach contains an apparent inconsistency: while the random numbers estimated conform to a NIG distribution, prices were simulated assuming that a risk-neutrality condition exists. When the underlying price process follows an alternative process, such as a NIG distribution, a correction to the drift term is required to allow risk neutral evaluation. An appropriate risk-neutral drift adjustment is:

$$a_t = \frac{\delta}{\Delta t} * \left[\sqrt{\alpha^2 - (\beta^2 + \sigma_{t-1} \cdot \sqrt{\Delta t})^2} - \sqrt{\alpha^2 - \beta^2} \right] - \frac{\mu \cdot \sigma_{t-1}}{\sqrt{\Delta t}} \quad (7)$$

Where a_t is the risk neutral adjusted drift, the Greek letters, α, β, δ, and μ are the parameters of the NIG distribution and σ_{t-1} is the random volatility (from stochastic volatility process in equation 3.2) for the previous discrete observation.[16] This drift is then used in equation 3.1 with the drift term (μ) equal to a_{t-1} from equation 7 and the disturbances are drawn from the NIG distribution.

Finally, to simulate correlated processes, we generated a new set of random numbers (Z') for the volatility process using the usual method for drawing samples from a standardised bivariate distribution:

$$Z' = Z \cdot \rho + \sqrt{(1-\rho^2)} \cdot Z_2 \quad (8)$$

Where Z represents the random disturbances for the underlying price process (in the case of GBM, the Z_1 set of normal draws, in the case of a NIG process, the draws from one of the four samples). The term, Z_2 , refers to the representative normal distribution selected for the volatility process and the term, ρ, refers to the correlation coefficient between the two processes[17].

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 paramater 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). Of the four variables in the Heston (1993) model, it was found that only three of the factors were critical. For a given long-term volatility, θ (taken as the average 20-day volatility in Table III), the crucial factors are the rate of mean reversion, κ, the volatility of the volatility, ξ and the correlation between the price and volatility process, ρ. It was found that small changes in the level of the long-term volatility had almost no effect on the sum of squared errors. 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. When no further improvement in the sum of squared errors is possible, the final combination of parameter values is deemed the optimal estimation of the Heston (1993) model.[18]

Tables VIa, VIb, VIc and VId display the results of the models for each of the stock index futures markets and split into the three periods of analysis. In each of these tables, the three time periods of analysis appear. On the left-hand side, the parameterisation of each of the five 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. Finally, in the furthest right-hand column appears the SSE of the attributes of each model compared to the empirical attributes.[19]

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

As was expected, the GBM model with constant variance is rejected. Recalling that the standard deviation of the SSE for the GBM simulations was 4.5646, it is clear that rejection occurs at well above a 99% confidence interval. All the stochastic volatility models have substantially lower SSEs. If we assume that the model with the lowest SSE is the optimal one, for the twelve data sets (four markets for three periods), model 5 (NIGSVρ) is the best in eight, model 3 (SVρ) is the best for three and model 4 (NIGSV) is the best for the remaining one.

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. In addition, the volatility of volatility (CoV attribute) and the decay of this over time (Line fit) are also captured. This model fails to address skewness in returns or the leverage effect and fails to generate sufficient excess kurtosis. The correlated stochastic volatility model (SVρ), does allow the leverage effect and the negative skewness in the returns to be captured. However, this model still fails to generate sufficient excess kurtosis. The NIGSV model captures the excess kurtosis and the volatility related attributes, but generates too excessive negative skewness and fails to capture the leverage effect. The final model (NIGSVρ) is able to capture all seven attributes.

The SSEs are compared between the alternative models to allow a more meaningful comparison among the models and relying upon the simulated standard deviation of this statistic, T-statistics are estimated to differentiate between models. This appears in Table VII.

[Table VII appears here]

In this table, the starting point for comparison is the number of standard deviations in improvement found relative to the first model (GBM)[20]. When this improvement is significant at a 95% level or above, the T-statistics appear in bolded text. For all these tables, the columns represent the initial model being compared and the rows the subsequent model. The first column shows the number of standard deviations in improvement that has occurred in the SSEs by the inclusion of a stochastic volatility model relative to a GBM assumption. For all time periods and for all markets, all the stochastic volatility models lead to a significant reduction in the sum of squared errors. For all stock index futures markets in the entire period and the second period, Models 3, 4 and 5 are a significant improvement over the simple SV model. In the first period, only for the DAX futures are these models an improvement, the reason being that either a leverage effect and/or extreme excess kurtosis existed requiring the alternative models. In no instance is Model 5 (NIGSVρ) significantly better (in a statistical sense) relative to Model 3 (SVρ). Thus, it would appear that correlated processes are relatively more important than the inclusion of a non-normal price process in explaining the empirical attributes. However, in a number of instances, the NIGSVρ model is a significant improvement over the NIGSV model.

7. IMPLICATIONS FOR OPTION PRICING

From the preceding section, a more realistic price process for the four stock index futures markets has thus been uncovered. This will serve as a prior process for the estimation of option values.[21] The next step is to examine the implications for options based upon these assets. Given that the parameter estimation of the stochastic volatility models relies upon simulation, it is a simple matter to use a similar simulation technique to estimate option prices. This simulation approach determines European call and put options numerically (a Monte Carlo approach) over a variety of strike prices and times to expiration. This is similar in spirit to the approach used by Johnson and Shanno (1987).

In the immediately preceding section, it was demonstrated that in all instances either a stochastic volatility model with correlated processes and/or a NIG model was optimal, therefore, only the NIGSVρ model was considered.[22] The choice of the appropriate NIG distributions and the parameter values for the stochastic volatility models are taken from the previous section. The simulations were based on the results for the entire period of analysis.

An apparent inconsistency for this approach is that the use of Monte Carlo simulations to price options assumes that a risk-neutrality condition exist. This suggests that the state space is continuous and spanned across that space by existing securities. However, all of the models simulated, introduce stochastic volatility, correlated price processes and include (negative) jumps and into the state space. However, no securities exist allowing the state space to be spanned when the volatility displays such dynamics. Thus (in the strictest sense), these models do not permit us to use the risk-neutrality argument to price the options. This is the apparent theoretical inconsistency. However, the determination of the risk-neutral drift adjustment for the NIGSVρ process in equation 7 does allow limited comparisons of the simulated options prices to those estimated by the Black (1976) formula to be made. The equivalent risk-neutral measure in equation 7 is certainly feasible, however, it may not be unique. With this proviso, we will proceed to compare options prices from both models assuming both are determined under an equivalent risk-neutral measure. This will allow us to examine whether an alternative price process (also under an equivalent risk-neutral measure) is sufficient to explain the existence of implied volatility smiles.

In this simulation price series of one year 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. The 252 trading days were examined at 20 points. These points represent weekly (5 business days) expirations out to three month, which correspond to the term to maturity of the most liquid exchange traded options. Thereafter, the time points are chosen to correspond to monthly maturities out to one year. Options at these maturities are more likely to correspond to over-the-counter (OTC) options.

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: $\dfrac{\ln(X_\tau / F_\tau)}{\sigma\sqrt{\tau / 252}}$ **(9)**

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 variance. 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.[23]

For the simulations, the volatility parameter chosen was equal to the level of volatility used in the calibration of the optimal NIGSVρ 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 will allow more direct comparisons to be drawn.

For the Monte Carlo simulation, 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.[24] These were also stored and used for all analysis using that particular NIG distribution. Then, the bivariate distribution was estimated for the volatility series using formula 7 and the optimal parameters of the NIGSVρ models for each stock index futures in Tables VIa, VIb, VIc, and VId (for the entire period of analysis).

With the appropriate NIG distribution and the estimated bivariate distribution for the stochastic volatility, volatilities and prices were estimated using an Euler approach (discrete form of formula 3.1 and 3.2). With the prices of each model estimated, the payoffs of the fifteen options (at each of the 20 points 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.

In parallel, we estimated the prices of all the options 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. 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[25]. Given that the underlying asset was assumed to be futures, the interest rate and dividend yield was set to zero. The same assumptions were made when estimating the NIGSVρ option prices[26].

The standardisation of the strike prices (as a function of both time and the volatility level of each market) make direct comparisons of the biases in the models possible. Unfortunately, the sheer amount of information in the tables and the fact that four markets are compared makes such comparisons cumbersome. To allow more direct comparisons, the simulated option prices were expressed as implied volatilities using the Black (1976) formula, indexed to the level of the long-term volatility [constant in the Black (1976) case] and graphed relative to the standardised strike prices and time to expiration. The indexed implied volatilities are then presented as a continuous surface. These surfaces are presented in Figure 1 for the four stock index futures markets.

In Figure 1, the implied volatility surfaces display many of the features of implied volatility surfaces found for actual stock index options. When comparisons are made between the surface for the S&P 500 simulated options to those reported by Derman & Kani (1994) and Corrado and Su (1996), similar patterns are found. There is the curvature characteristic of volatility smiles, and the higher levels for the implied volatilities of the lower strike price (put) options and relatively lower levels of implied volatilities for the higher strike price options (call) are consistent with a volatility skew. This implies that significant differences exist between the implied volatilities consistent with the NIGSVρ model and the Black (1976) implied volatilities (which should be constant for all strike prices and maturities). 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. The NIGSVρ volatilities for out-of-the-money options tend to have higher implied volatilities relative to the Black (1976) implied volatilities (especially for lower strike prices).

[Figure 1 appears here]

One can observe that there is a significant range of strike prices and maturities where the Black (1976) model over-values the options. It is also the case that the overpricing bias tends to be increasingly as the term to expiration is extended. These results are consistent with the biases of stochastic volatility on option prices found elsewhere [Hull & White (1988)].

For three of the four established markets (all apart from the S&P 500), the implied volatility surfaces are relatively linear for the term nearest to expiration and become more curved as the term to expiration is extended. This contrary to the behaviour of implied volatility surfaces estimated from actual option prices that appear in Figure 2.

[Figure 2 appears here]

In Figure 2, implied volatilities were estimated using the Black (1976) model for all the nearest to expiration traded options for the four established markets from November 1990 to December 1998. The analysis was restricted to those options with the same expiration cycles as the futures examined (a maximum of 3 months to expiration). The strike prices were then converted to the same standardised form as for the simulated implied volatility surfaces (using formula 9) and the levels of implied volatility were also 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 quadratic fitting procedure [see Shimko (1993)] was implemented to yield an implied volatility smile. These were then combined and the volatility surface was drawn.

The comparison between Figures 1 and 2 must take into account that differences exist in the basis of the volatility indexation and the term to maturity of the options. The implied volatilities in Figure 1 are expressed relative to the Black (1976) implied volatility, while the implied volatilities in Figure 2 are expressed relative to the at-the-money implied volatility associated with option prices. Secondly, in Figure 1 the maturities for the simulated options extend to one year and in Figure 2 are restricted to 3 months. However, in Figure 1, the first twelve time points correspond directly to those in Figure 2. In Figure 2, it is clear that the degree of curvature is most extreme the shorter the term of the options life and becomes more linear for longer expirations. Another significant difference is that the implied volatility surfaces for the actual options market displays more curvature that is observed in the simulated surfaces. For all four markets, the relative shapes are similar but the amplitude of the effect is dampened in the simulated smiles.

This divergence can be understood by re-examining the models for the unconditional returns. The NIG distribution was included to capture the effects of daily excess kurtosis and daily skewness. From the Central Limit Theorem, this distribution will rapidly approach a normal stochastic process. This being the case, the impact of the NIG distribution is short-lived. In the absence of other factors, this would lead to a skewed and curved implied volatility smile for the shortest time period to expiration. However, the inclusion of correlated price processes will counteract this effect and increase the degree of negative skewness as a function of the time horizon. The fact that the surfaces become more curved with the time horizon is a directly related to the impacts of stochastic volatility increasing with the term to expiration. This can be seen in all four surface graphs with the curvature tending to increase with the term to expiration of the option.

A puzzle is introduced by the divergence between the patterns of implied volatility surfaces estimated using the simulated option prices (consistent with the unconditional prices process of established Stock Index futures) and actual option prices on these futures. The most obvious conclusion is that the NIGSVρ model is inappropriate as a candidate for describing the unconditional price process for the underlying futures. However, it is unclear what alternative approaches could be tried. The existence of time varying volatility requires some stochastic volatility model and ample evidence exists for the mean reversion of volatility (suggesting some sort of Ornstein-Uhlenbeck process would be appropriate). These models (such as SV) are not able to generate sufficient excess kurtosis consistent with observed unconditional kurtosis, thus some non-normal price process is suggested. Finally, given that ample evidence suggests that a leverage effect exists, it is reasonable to consider the two stochastic processes are correlated.

A second possibility is that the models we have fit to the unconditional futures return series are mis-specified. It may be that relevant attributes have been excluded. It is not clear what these excluded attributes would be. Alternatively, it may be that due to the incompleteness of markets implied by our findings, the prices of options contain risk premia. Therefore, even if the price process of the underlying asset were correctly estimated, 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 therefore be incomplete.

To understand the nature of the divergence between the simulated and actual implied volatility surfaces better, we determined the simple differences between the surfaces in Figure 1 and 2. This appears in Figure 3. In Figure 3, the simulated implied volatility surface (from Figure 1) was re-estimated to conform to the same periods to expiration as appear for the actual option implied volatility surface in Figure 2.

[Figure 3 appears here]

It is apparent that the differences between the simulated and actual implied volatility surfaces are remarkably similar among the four Stock Index futures & options markets. Once we have controlled for the implied volatility surfaces consistent with the price processes of the underlying assets, whatever factors are causing the divergences in the shapes of the implied volatility surfaces are generally similar across these Stock Index option markets.

Given that the implied volatilities have been standardised, the differences in the levels can be interpreted as the percentage differences between the two implied volatility surfaces. One interesting result is that the divergences are relatively symmetrical relative to the ATM level for three of the four markets (exception is the DAX futures). Thus, it would appear that the skew effect pointed out in the literature for options on stock index futures might be due primarily to an alternative price process (asymmetrical jumps and a leverage effect). The fact that the difference in the DAX smiles is skewed towards lower strike prices may suggest another factor may be at work.

Given all implied volatilities have been standardised, comparisons are simplified (independent of scaling). The levels of divergence are everywhere positive and at similar levels. For the OTM options, the difference is in the range of 40-60%. This is most extreme for options with short time periods to expiration. However, for both the FTSE and the Nikkei markets, the divergence seems relatively time invariant. The fact that everywhere the difference is positive could be interpreted as evidence for the existence of a risk premia, be associated with transaction costs or could be caused by other sources of market imperfections or frictions.

8. SUMMARY AND SUGGESTIONS FOR FUTURE RESEARCH

This paper examines the nature of the objective dispersion processes, which can be observed for futures contracts on four stock indices. 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, four alternative stochastic volatility models were examined. It was found that all four markets are best understood with a NIGSVρ model, which assumes the price process follows a Normal Inverse Gaussian process and the stochastic process driving volatilities is negatively correlated to this NIG process.

Based upon an optimal parameterisation of the NIGSVρ model, European options were determined numerically for each market. Significant divergences from the Black (1976) values were observed. For all markets, the NIGSVρ 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 research points out for the first time the relative biases associated with the NIGSVρ option pricing model proposed by Barndorff-Nielsen and Shephard (1999). This extends the research of Hull & White (1988) to consider a richer class of stochastic volatility models.

These NIGSVρ option prices were then expressed as standardised implied volatility surfaces and compared to actual implied volatility surfaces from options on these same four markets. While the NIGSVρ implied volatility surfaces resemble actual surfaces, differences are found that appear to be systematic across the four markets. For the entire surfaces, the actual option prices have higher implied volatilities than for option prices consistent with the model. It was found that these divergences appear to display similar dynamics across the four markets. One interesting result is that for three of the four markets the divergence seems to be symmetric relative to the underlying price. This suggests that the skewness effect associated with implied volatility smiles for these three stock index options is primarily due to the unconditional skewness of the underlying price process and a leverage effect. For most of the markets, the divergence is relatively time invariant. The consequences of this are unknown.

The nature of this difference remains for future research. The research could follow these lines: it could incorporate early exercise features for American options, transactions costs, segmentation of hedging strategies across strike prices, option type and time (by heterogeneous agents) and examination of the risk premia (due to incomplete markets).

REFERENCES:

Andersen, T.G., Chung, H-J & Sørensen, B.E.(1999). Efficient method of moments 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. 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. 1-11.

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

Barndorff-Nielsen, O.E. & N. Shephard (1999). Incorporation of a Leverage Effect in a Stochastic Volatility Model. Working Paper, The Centre for Mathematical Physics and Stochastics, University of Aarhus.

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

Bates, D.S. (1999). Post-'87 Crash Fears in the S&P 500 Futures Options Markets. Forthcoming in Journal of Econometrics.

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

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

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

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

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

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

Clewlow, L & Xu, X. (1994). The Dynamics of Stochastic Volatility. Financial Options Research Centre, Warwick Business School, FORC Preprint 94/53.

Corrado, C.J. & Su, T. (1996). Skewness and kurtosis in S&P 500 Index returns implied by option prices. Journal of Financial Research, 19, 175-192.

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

Derman, E. & Kani, I. (1994). Riding on the Smile. RISK, 7 (2). 32-39.

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.

Fama, E. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work", The Journal of Finance, 25, 383-417.

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

Garman, M. & Klass, M. (1980). On the Estimation of Security Price Volatilities from Historical Data. Journal of Business, 53, 67-78.

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.

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.

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.

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

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

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

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

Neuberger, A. (1994). The Log Contract. Journal of Portfolio Management, 20, 74-80.

Parkinson, M. (1980). The Extreme Value Method for Estimating the Variance of the Rate of Return. Journal of Business, 53, 61-66.

Randolph, W.L. & Najand, M. (1991). A Test of Two Models in Forecasting Stock Index Futures Price Volatility, Journal of Futures Markets, 11, 179-190.

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

Shimko, D. (1993). Bounds of Probability, RISK, 6, Number 4, 33-37.

Stoll, H.R. & Whaley, R.E. (1990). The Dynamics of Stock Index and Stock Index Futures Returns. Journal of Financial and Quantitative Analysis, 25, 441-468.

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, London, England, RISK Publications.

FOOTNOTES:

[1] It may appear at first glance that this measure is related to a Variance Ratio test. However, the nature of this link is not obvious and is left for future research to consider.
[2] Barndorff-Nielsen (1997) actually assumed that the underlying price process follows GBM and the volatility process follows an Inverse Gaussian distribution. However, it can be shown that this is equivalent to the underlying price process following a NIG distribution and the volatility process follows GBM.
[3] As is pointed out in Melino and Turnbull (1990), the choice of moments is critical to the estimation procedure. Given that the choice of moments may be subjective, it is critical that these be chosen judiciously. Assuming relevant attributes have been chosen, the previous references suggest that a GMM approach will be as robust as a maximum likelihood technique.
[4] One problem when estimating the return series of futures is that each contract has a fixed expiration date and one must exercise care that returns are determined correctly. Essentially, the problem is to decide when to rollover the contracts in the analysis and assure at rollover that returns are not estimated from contracts with different expiration dates. It was decided to restrict the analysis to the nearest to expiration contracts, as these tend to have the most liquidity and the most representative prices. The returns of the nearest to expiration futures price series were determined up to the day prior to expiration and on that date determined the next day's return by using the prices of the futures contract for the next expiration period. This assured that each return series was calculated correctly.

[5] Neuberger (1994) points out that if the drift, $\bar{r}$, is close to zero (as it should be for a futures price), the sample second moment provides a better estimate of the population standard deviation than the sample standard deviation
[6] These include estimators by Parkinson (1980) and Garman and Klass (1980). Both estimators incorporate high and low prices, and have been shown to be statistically more efficient that the commonly used close-to-close return standard deviation. However, recent evidence has suggested that using close to close data for the estimation of volatility is as good a measure as using the extreme methods [see Clewlow and Xu (1994)].
[7] This is equal to $1/\sqrt{(2*19)}$.
[8] Given the limited number of observations in this study, it was difficult to have enough data points on an overlapping basis to determine meaningful statistical moments. Therefore, overlapping observations were used and the bias introduced by overlapping data was corrected using the procedure pointed by Tompkins (1997). In some ways, this is not necessary given that the bias of overlapping data will impact the empirical time decay and the simulated time decay, as both series have roughly the same number of observations and the bias due to overlapping is directly related to the number of observations.
[9] One potential problem with inclusion of both the measures of the average autocorrelations of absolute returns and of the time decay of the coefficient of variation of the volatility for different time horizons is that both sets of attributes are redundant. To assess this possibility, two tests were completed. The first was to estimate the correlation coefficients between these attributes for the stock index futures observed over various time periods. The second test was to estimate the correlation coefficients for a series of 500 random price series of 1500 observations that conformed to a GBM world with constant variance. In both cases, the levels of correlation were significantly below 1.0 and suggest that these attributes are measuring different dynamics of the volatility process.
[10] Stewart Hodges (University of Warwick) suggested this and the full proof of this result is available upon request.
[11] Using a maximum likelihood method, the optimal weight was selected for each market than minimised the correlation between the two series. For most of the markets, the optimal weighting factor was found to be between 0.03 and 0.035. To simplify the selection of the attribute, a fixed weight of 0.03 was applied to all markets and for all time periods of analysis, to allow the leverage correlation factor to be not be subject to differing weights.
[12] 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.
[13] Given the DAX futures were the last contract introduced on November 23rd 1990, this is the starting date for the analysis.
[14] 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.
[15] 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.
[16] A complete proof of the derivation of this risk neutral drift adjustment appears in Appendix 1 at the end of this paper. It is noted there that this is only one manner to adjust this drift to achieve risk-neutrality. While it is

acknowledged that given markets are incomplete, there may not exist a unique martingale measure, this is certainly a feasible one and will allow for relative comparisons to be drawn.

[17] Given the manner the bivariate distribution is determined, the risk-neutral drift adjustment to the NIG distribution is unchanged.

[18] For Models 2 and 4, the correlation between the processes was assumed to be zero. This meant only four alternative parameter sets needed to be considered and the grid search was in two dimensions.

[19] Given that the results have been presented by individual markets rather than by time periods, the standard deviation of attributes across the stock index futures (which is required for formula 6) does not appear here. However, to save space, these are not presented. The reader can determine these simply by aggregating the attributes for each time period and estimating the standard deviation.

[20] The test is the difference in the SSEs of the alternative models divided by the standard deviation of the SSE statistic found in Table IV.

[21] This assumes that the price processes for the stock index futures will be similar under the risk-neutral measure for option prices. Clearly this assumes complete markets exist. Given that the optimal models for stock index futures prices imply this is not the case, an apparent contradiction exists. The aim of this research is to consider whether smiles can be explained solely by an alternative price process hypothesis. Under the assumption that the market price of risk is zero, we will test this hypothesis. If the price process under P is able to adequately explain the existence of implied volatility smiles, the null hypothesis of the existence of a risk premium is accepted. Alternatively, rejection of the null hypothesis may provide insights into the nature of the market price of risk.

[22] While for three of the twelve data sets, the SVρ model was marginally superior to the NIGSVρ model, this difference was not found to be statistically significant.

[23] 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.

[24] As a technical point, it is not possible to use the anti-thetic method when determined asymmetrical distributions as this eliminates the asymmetry.

[25] These average prices were not significantly different than 100.

[26] Due to the sheer volume of results, these tables do not appear in this paper. However, they are available upon request. In addition, for the purposes of independent replication of our results, NIG distributions must be estimated. While referring to the work of Rydberg (1997) can do this, the author can also provide a Visual Basic © programme running under EXCEL© 97, which implements this procedure.

APPENDIX 1

Derivation of the Risk Neutral Drift Adjustment for a Normal Inverse Gaussian Process..

The process S, generated by the NIG distribution should be a martingale. In the discrete time setting this means:

$$E\left[S_{t_j} \mid \Im_{t_{j-1}}\right] = S_{t_{j-1}}, j = 1,2,3,\ldots,n \qquad 1.1$$

with $\Im_{t_{j-1}}$ denoting the information at time t_{j-1}. Since both the drift term, a, and the volatility, σ, are known at time t_{j-1}, and the variance of the series V_j is independent from the history leading up until t_{j-1}, this is equivalent to:

$$e^{a_{t_{j-1}}\Delta t} m\left(\sigma_{t_{j-1}}\sqrt{\Delta t}\right) = 1, \qquad 1.2$$

where $m(x) = E[e^{xV_j}]$ is the moment generating function of the $NIG(\mu, \delta, \alpha, \beta)$ distribution.

This function is:

$$m(x) = \exp(\delta(\sqrt{\alpha^2 - \beta^2} - \sqrt{\alpha^2 - (\beta + x)^2}) + \mu x) \qquad 1.3$$

So the appropriate choice of drift is:

$$a_{t_{j-1}} = \frac{1}{\Delta t} \ln\left(\frac{1}{m\left(\sigma_{t_{j-1}}\sqrt{\Delta t}\right)}\right) \qquad 1.4$$

Combining equations 1.3 and 1.4 leads to

$$a_t = \frac{\delta}{\Delta t} * \left[\sqrt{\alpha^2 - (\beta^2 + \sigma_{t-1} \cdot \sqrt{\Delta t})^2} - \sqrt{\alpha^2 - \beta^2}\right] - \frac{\mu \cdot \sigma_{t-1}}{\sqrt{\Delta t}} \qquad 1.5$$

As a final note, adjusting the drift is only one way to obtain a martingale (for risk-neutral valuation). Given the models examining suggest an incomplete market, there are alternative approaches. For example, for the constant volatility NIG process in continuous time, the natural approach would be to estimate the parameters, μ, δ, α $and \beta$ from the observations of S and adjust β to obtain a risk neutral valuation.

Underlying Asset	Time Period of Analysis		Number of Observations
(Entire Period)	First Date	Last Date	
S&P 500 Futures	23/11/90	17/12/98	2042
FTSE 100 Futures	23/11/90	18/12/98	2042
DAX Futures	23/11/90	18/12/98	2020
NIKKEI-225 Futures	23/11/90	11/12/98	2037
(First Period)	First Date	Last Date	
S&P 500 Futures	23/11/90	15/12/94	1029
FTSE 100 Futures	23/11/90	16/12/94	1029
DAX Futures	23/11/90	16/12/94	1016
NIKKEI-225 Futures	23/11/90	08/12/94	1023
(Second Period)	First Date	Last Date	
S&P 500 Futures	19/12/94	17/12/98	1013
FTSE 100 Futures	19/12/94	18/12/98	1013
DAX Futures	19/12/94	18/12/98	1004
NIKKEI-225 Futures	09/12/94	11/12/98	1014

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
S&P 500 Futures								
Overall Period	0.00055	0.00907	**-0.4397**	**-8.11**	**11.401**	**77.47**	**6067.18**	2041
First Period	0.00030	0.00717	**0.2074**	**2.72**	**6.078**	**20.14**	**413.14**	1028
Second Period	0.00079	0.01066	**-0.6477**	**-8.42**	**10.884**	**51.22**	**2694.11**	1013
FTSE 100 Futures								
Overall Period	0.00038	0.01000	-0.0944	-1.74	**5.979**	**27.47**	**757.87**	2041
First Period	0.00014	0.00950	0.1322	1.73	**4.457**	**9.54**	**93.94**	1028
Second Period	0.00063	0.01048	**-0.2779**	**-3.61**	**6.983**	**25.88**	**682.76**	1013
DAX Futures								
Overall Period	0.00043	0.01272	**-0.6466**	**-11.86**	**11.744**	**80.20**	**6572.19**	2019
First Period	0.00011	0.01200	**-0.8826**	**-11.48**	**19.675**	**108.44**	**11891.68**	1015
Second Period	0.00075	0.01340	**-0.4865**	**-6.29**	**6.408**	**22.04**	**525.49**	1004
Nikkei 225 Futures								
Overall Period	-0.00037	0.01538	**0.1161**	**2.14**	**4.919**	**17.68**	**317.04**	2036
First Period	-0.00039	0.01516	**0.1943**	**2.54**	**4.447**	**9.44**	**95.58**	1022
Second Period	-0.00035	0.01560	0.0437	0.57	**5.352**	**15.29**	**234.01**	1014

Table II, Statistics of the Daily Returns for Four Stock Index Futures Markets (1990-1998)

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)
S&P 500 Futures							
Entire Period	12.50%	6.29%	**0.503**	**-0.1359**	**-0.2174**	**0.1465**	**0.0903**
First Period	10.54%	3.35%	**0.318**	**-0.1943**	**0.2049**	**0.0563**	**0.0391**
Second Period	14.46%	7.74%	**0.535**	**-0.1191**	**-0.3999**	**0.1541**	**0.0834**
FTSE 100 Futures							
Entire Period	14.60%	5.72%	**0.392**	**-0.1600**	**-0.3010**	**0.1330**	**0.0679**
First Period	14.93%	4.12%	**0.276**	**-0.1584**	**-0.1227**	**0.0615**	**0.0274**
Second Period	14.35%	6.54%	**0.455**	**-0.1001**	**-0.4536**	**0.1725**	**0.0954**
DAX Futures							
Entire Period	17.50%	9.08%	**0.519**	**-0.2000**	**-0.3534**	**0.1806**	**0.0817**
First Period	17.30%	8.50%	**0.492**	**-0.2650**	**-0.2395**	**0.0714**	**0.0282**
Second Period	17.77%	9.19%	**0.517**	**-0.1001**	**-0.4614**	**0.2512**	**0.1158**
Nikkei 225 Futures							
Entire Period	23.40%	9.22%	**0.394**	**-0.1934**	**-0.2390**	**0.1308**	**0.0397**
First Period	24.63%	9.26%	**0.376**	**-0.1958**	**-0.1946**	**0.1366**	**0.0355**
Second Period	22.42%	9.73%	**0.434**	**-0.1491**	**-0.2561**	**0.1421**	**0.0574**
Expected GBM Attributes	20.000%	3.244%	**0.162**	**-0.5000**	**0.0000**	**0.0000**	**0.0000**
Standard Error of Attributes	3.244%	0.032%	**0.010**	**0.0778**	**0.0999**	**0.0051**	**0.0061**

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

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

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 (1996b). 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 nine Bond Futures/Forward 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

S&P 500 Futures (Whole Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.5033	-0.4397	11.4006	-0.2174	0.1465	0.0903	-0.1359	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.125	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	255.781
Model 2	GBM	0.125	1.90	0.910	0	0.4834	-0.2676	5.6661	-0.0427	0.2020	0.0344	-0.1417	23.101
Model 3	GBM	0.125	3.10	0.890	-0.25	0.4471	-0.4685	5.9502	-0.2742	0.1763	0.0765	-0.1086	6.571
Model 4	NIG-T2	0.125	0.79	0.405	0	0.5436	-0.9821	10.3968	-0.1931	0.1540	0.1015	-0.1502	3.654
Model 5	NIG-T1	0.125	1.80	0.380	-0.35	0.5239	-0.3702	9.6509	-0.2165	0.1597	0.0759	-0.1819	3.387

S&P 500 Futures (First Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.3180	0.2074	6.0778	0.2049	0.0563	0.0391	-0.1943	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.105	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	123.538
Model 2	GBM	0.105	2.50	0.180	0	0.3250	-0.0047	3.7701	0.1447	0.1090	0.0456	-0.1406	2.785
Model 3	GBM	0.105	2.39	0.172	0.035	0.2928	0.0161	3.5894	0.1824	0.0852	0.0371	-0.1456	2.686
Model 4	NIG-T1	0.105	2.56	0.132	0	0.2970	-0.0797	5.6082	0.0042	0.0413	0.0447	-0.2549	4.311
Model 5	NIG-T3	0.105	2.05	0.170	-0.09	0.2961	-0.1488	4.2690	0.1607	0.0701	0.0377	-0.1875	1.020

S&P 500 Futures (Second Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.5350	-0.6206	11.2081	-0.3999	0.1541	0.0834	-0.1191	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.145	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	393.627
Model 2	GBM	0.145	4.00	0.600	0	0.4804	-0.1251	5.3140	0.0768	0.2092	0.0538	-0.1352	38.586
Model 3	GBM	0.145	4.00	0.700	-0.4	0.5072	-0.6220	6.5818	-0.3921	0.2139	0.0796	-0.1171	4.899
Model 4	NIG-T2	0.145	1.10	0.300	0	0.5016	-1.1189	10.6596	-0.2087	0.1217	0.1042	-0.1309	8.984
Model 5	NIG-T1	0.145	1.99	0.495	-0.61	0.5769	-0.5503	10.6853	-0.4214	0.1833	0.0872	-0.1568	2.885

Table VI a, Empirical Attributes of S&P 500 Futures and Attributes of Five Alternative Models

Models (Parameters) | Empirical and Simulated Attributes

FTSE 100 Futures (Whole Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.3917	-0.0944	5.9793	-0.301	0.1330	0.0679	-0.1600	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.146	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	211.355
Model 2	GBM	0.146	1.11	0.489	0	0.5271	-0.1269	5.0988	-0.0095	0.2458	0.1149	-0.1063	26.843
Model 3	GBM	0.146	3.01	1.001	-0.29	0.4858	-0.4720	5.7005	-0.2829	0.1965	0.0122	-0.1969	3.100
Model 4	NIG-T3	0.146	6.99	0.899	0	0.4444	-0.2207	6.2114	-0.1521	0.1495	0.0182	-0.2280	13.986
Model 5	NIG-T3	0.148	3.45	0.549	-0.64	0.3869	-0.2845	5.1239	-0.3226	0.1439	0.0553	-0.1599	1.101

FTSE 100 Futures (First Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.2760	0.1322	4.4571	-0.1227	0.0615	0.0274	-0.1584	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.149	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	102.921
Model 2	GBM	0.149	2.80	0.250	0	0.2803	-0.0048	3.5250	0.1503	0.0758	0.0278	-0.1662	2.202
Model 3	GBM	0.149	2.43	0.251	-0.2	0.2723	-0.0962	3.5016	-0.1034	0.0684	0.0278	-0.1795	0.570
Model 4	NIG-T3	0.149	2.43	0.225	0	0.2795	-0.1361	4.1981	0.2092	0.0575	0.0300	-0.2049	4.291
Model 5	NIG-T3	0.149	1.76	0.223	-0.266	0.2646	-0.1864	3.9877	0.0522	0.0526	0.0308	-0.2062	2.710

FTSE 100 Futures (Second Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.4555	-0.2613	6.5835	-0.4536	0.1725	0.0954	-0.1001	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.144	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	394.747
Model 2	GBM	0.144	3.79	0.605	0	0.4725	-0.1308	5.3247	0.0753	0.2021	0.0460	-0.1391	39.032
Model 3	GBM	0.144	4.10	0.750	-0.5	0.4706	-0.6577	6.5588	-0.4462	0.1847	0.0832	-0.1088	2.468
Model 4	NIG-T1	0.144	2.98	5.980	0	0.5433	-0.2187	10.6754	-0.4586	0.1691	0.0899	-0.1762	16.704
Model 5	NIG-T3	0.144	2.59	0.549	-0.86	0.4253	-0.3877	4.8504	-0.4268	0.1821	0.1072	-0.1047	1.977

Table VI b, Empirical Attributes of FTSE 100 Futures and Attributes of Five Alternative Models

Models (Parameters) — Empirical and Simulated Attributes

DAX Futures (Whole Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.5186	-0.6466	11.7437	-0.35339	0.1806	0.0817	-0.2000	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.175	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	250.919
Model 2	GBM	0.175	1.90	1.100	0	0.4928	-0.3228	5.4496	-0.0378	0.2028	0.0106	-0.1738	42.258
Model 3	GBM	0.175	3.60	1.090	-0.35	0.4714	-0.5806	6.3360	-0.3619	0.1880	0.0721	-0.1337	7.652
Model 4	NIG-T2	0.175	1.91	0.809	0	0.5439	-1.0927	10.3649	-0.3439	0.1515	0.0643	-0.2126	4.338
Model 5	NIG-T1	0.175	1.93	0.431	-0.51	0.5036	-0.4030	9.3509	-0.3560	0.1478	0.0722	-0.1772	2.226

DAX Futures (First Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.4916	-0.8826	19.6753	-0.2395	0.0714	0.0282	-0.2650	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.173	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	95.863
Model 2	GBM	0.173	5.95	1.715	0	0.4208	-0.0359	4.9805	0.1233	0.1453	0.0235	-0.1965	17.129
Model 3	GBM	0.173	6.00	1.400	-0.7	0.4105	-0.6305	5.7878	-0.5376	0.1212	0.0217	-0.2302	9.856
Model 4	NIG-T2	0.173	5.80	0.622	0	0.4338	-1.2181	10.8713	-0.3891	0.0752	0.0297	-0.3066	2.990
Model 5	NIG-T1	0.173	6.33	0.708	-0.4884	0.4569	-0.6443	9.3737	-0.4505	0.0959	0.0274	-0.2703	3.006

DAX Futures (Second Period)

						Empirical Attributes							
						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
						0.5172	-0.4372	6.2881	-0.4614	0.2512	0.1158	-0.1001	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.178	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	434.946
Model 2	GBM	0.178	3.70	1.450	0	0.5067	-0.0314	6.1057	0.1901	0.2187	0.0276	-0.1327	64.465
Model 3	GBM	0.178	3.99	0.795	-0.51	0.4657	-0.6172	6.0690	-0.4912	0.1876	0.0733	-0.1242	2.937
Model 4	NIG-T1	0.178	2.05	0.850	0	0.5856	-0.2410	10.9630	-0.1599	0.2036	0.1157	-0.1148	17.186
Model 5	NIG-T3	0.178	2.30	0.700	-0.85	0.4669	-0.3921	5.1255	-0.4106	0.2117	0.1314	-0.0933	2.984

Table VI c, Empirical Attributes of DAX Futures and Attributes of Five Alternative Models

Models (Parameters) — Empirical and Simulated Attributes

Nikkei Futures (Whole Period)

						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
Empirical Attributes						0.3940	0.1161	4.9192	-0.239	0.1308	0.0397	-0.1934	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.234	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	169.494
Model 2	GBM	0.234	1.89	0.890	0	0.4738	-0.2193	5.0073	-0.0293	0.1988	0.0247	-0.1768	23.433
Model 3	GBM	0.234	5.35	1.016	-0.315	0.3946	-0.3448	4.9679	-0.2255	0.1409	0.0199	-0.1815	3.021
Model 4	NIG-T3	0.234	8.34	1.286	0	0.4024	-0.1840	5.8519	-0.0970	0.1192	0.0204	-0.2360	9.285
Model 5	NIG-T3	0.234	3.63	0.813	-0.54	0.3666	-0.2570	5.0156	-0.2674	0.1292	0.0394	-0.1835	1.405

Nikkei Futures (First Period)

						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
Empirical Attributes						0.3761	0.1943	4.4470	-0.1946	0.1366	0.0355	-0.1958	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.246	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	125.754
Model 2	GBM	0.246	3.75	0.652	0	0.3711	-0.0433	4.1324	0.1397	0.1412	0.0357	-0.1617	3.567
Model 3	GBM	0.246	4.48	0.762	-0.21	0.3698	-0.1919	4.2910	-0.0708	0.1353	0.0358	-0.1588	1.604
Model 4	NIG-T3	0.246	3.85	0.671	0	0.3850	-0.1102	5.2489	0.1502	0.1258	0.0352	-0.1981	3.419
Model 5	NIG-T1	0.246	4.10	0.732	-0.121	0.3773	-0.1568	5.1243	0.0423	0.1227	0.0333	-0.2005	2.008

Nikkei Futures (Second Period)

						CoV	Skewness	Kurtosis	Leverage	Corr(1-20)	Corr(51-70)	Line Fit	
Empirical Attributes						0.4338	0.0528	5.6695	-0.2561	0.1421	0.0574	-0.1491	
	Distribution	LTV	K	VoV	Correl	Simulated Attributes							SSE
Model 1	GBM	0.224	0.00	0.000	0	0.1648	0.0004	3.0705	-0.0042	-0.0011	-0.0082	-0.5000	287.723
Model 2	GBM	0.224	3.99	1.495	0	0.4969	-0.1041	5.7954	0.1986	0.2072	0.0245	-0.1396	23.014
Model 3	GBM	0.224	4.28	0.750	-0.31	0.4139	-0.3291	4.8356	-0.2273	0.1643	0.0567	-0.1418	2.612
Model 4	NIG-T1	0.224	2.55	0.750	0	0.5271	-0.1554	9.2605	-0.0860	0.1712	0.0949	-0.1619	11.991
Model 5	NIG-T3	0.224	2.31	0.851	-0.56	0.4243	-0.2431	5.4862	-0.2007	0.1722	0.0683	-0.1560	1.992

Table VI d, Empirical Attributes of Nikkei 225 Futures and Attributes of Five Alternative Models

Comparisons of Models

S&P 500 Futures

(Whole Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**51.026**			
Model 3	**54.651**	**3.625**		
Model 4	**55.291**	**4.265**	0.640	
Model 5	**55.349**	**4.323**	0.698	0.059

(First Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**26.481**			
Model 3	**26.503**	0.022		
Model 4	**26.146**	-0.335	-0.356	
Model 5	**26.868**	0.387	0.365	0.722

(Second Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**77.860**			
Model 3	**85.247**	**7.387**		
Model 4	**84.352**	**6.492**	-0.896	
Model 5	**85.689**	**7.829**	0.442	1.337

FTSE 100 Futures

(Whole Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**40.463**			
Model 3	**45.670**	**5.207**		
Model 4	**43.283**	**2.819**	**-2.387**	
Model 5	**46.108**	**5.645**	0.438	**2.826**

(First Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**22.088**			
Model 3	**22.445**	0.358		
Model 4	**21.629**	-0.458	-0.816	
Model 5	**21.976**	-0.112	-0.469	0.347

(Second Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**78.008**			
Model 3	**86.026**	**8.018**		
Model 4	**82.904**	**4.896**	**-3.122**	
Model 5	**86.134**	**8.126**	0.108	**3.230**

DAX Futures

(Whole Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**45.759**			
Model 3	**53.348**	**7.589**		
Model 4	**54.075**	**8.316**	0.727	
Model 5	**54.538**	**8.779**	1.190	0.463

(First Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**17.266**			
Model 3	**18.861**	1.595		
Model 4	**20.367**	**3.101**	1.506	
Model 5	**20.363**	**3.097**	1.502	-0.004

(Second Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**81.246**			
Model 3	**94.739**	**13.493**		
Model 4	**91.614**	**10.368**	**-3.125**	
Model 5	**94.728**	**13.483**	-0.010	**3.114**

Nikkei Futures

(Whole Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**32.031**			
Model 3	**36.507**	**4.476**		
Model 4	**35.134**	**3.103**	-1.374	
Model 5	**36.862**	**4.831**	0.354	1.728

(First Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**26.795**			
Model 3	**27.226**	0.430		
Model 4	**26.828**	0.033	-0.398	
Model 5	**27.137**	0.342	-0.089	0.309

(Second Period)

	Model 1	Model 2	Model 3	Model 4
Model 2	**58.050**			
Model 3	**62.524**	**4.474**		
Model 4	**60.468**	**2.417**	**-2.057**	
Model 5	**62.660**	**4.610**	0.136	**2.193**

Table VII, Statistical Significance Testing of Sum of Squared Errors of Alternative Stochastic Volatility Models

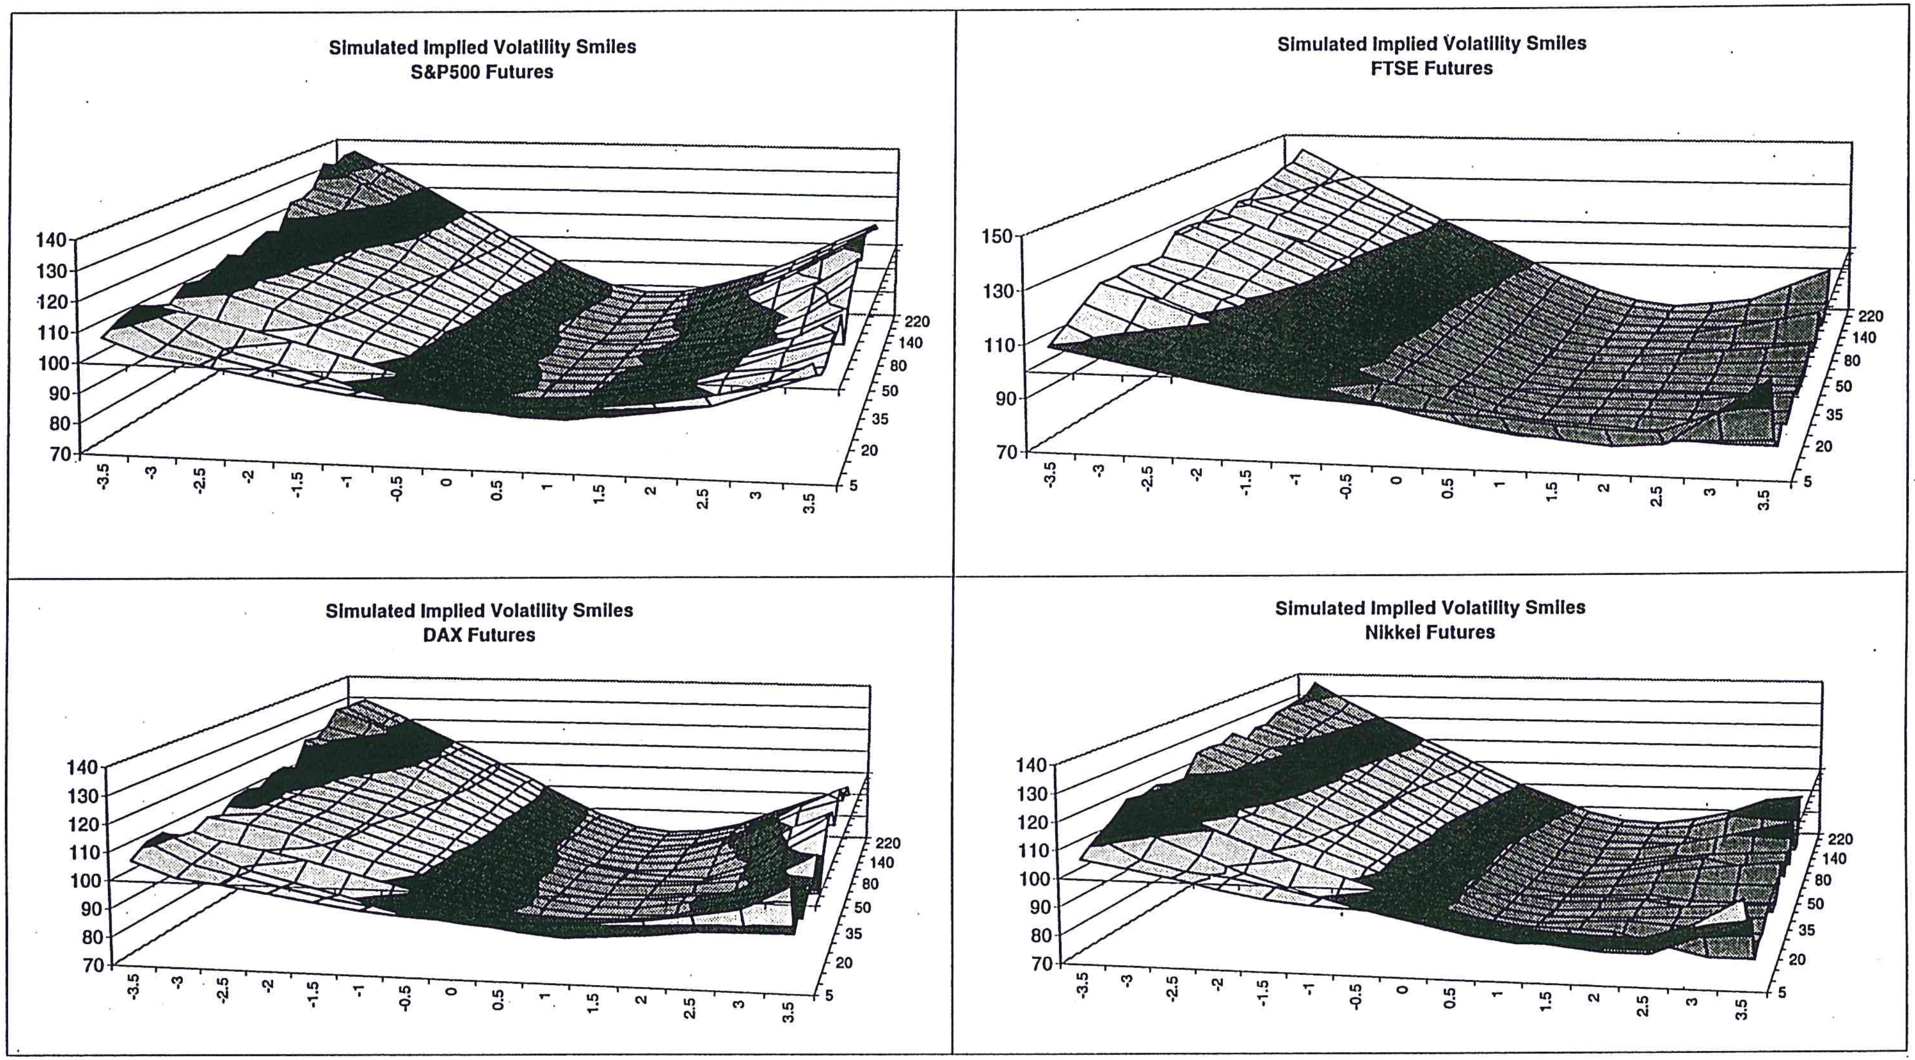

Figure 1, Implied Volatility Surfaces of Option Prices Following a NIGSVρ Model Compared to the Constant Volatility of the Black (1976) Model for Four Stock Index Futures Markets [100 = Black (1976) volatility].

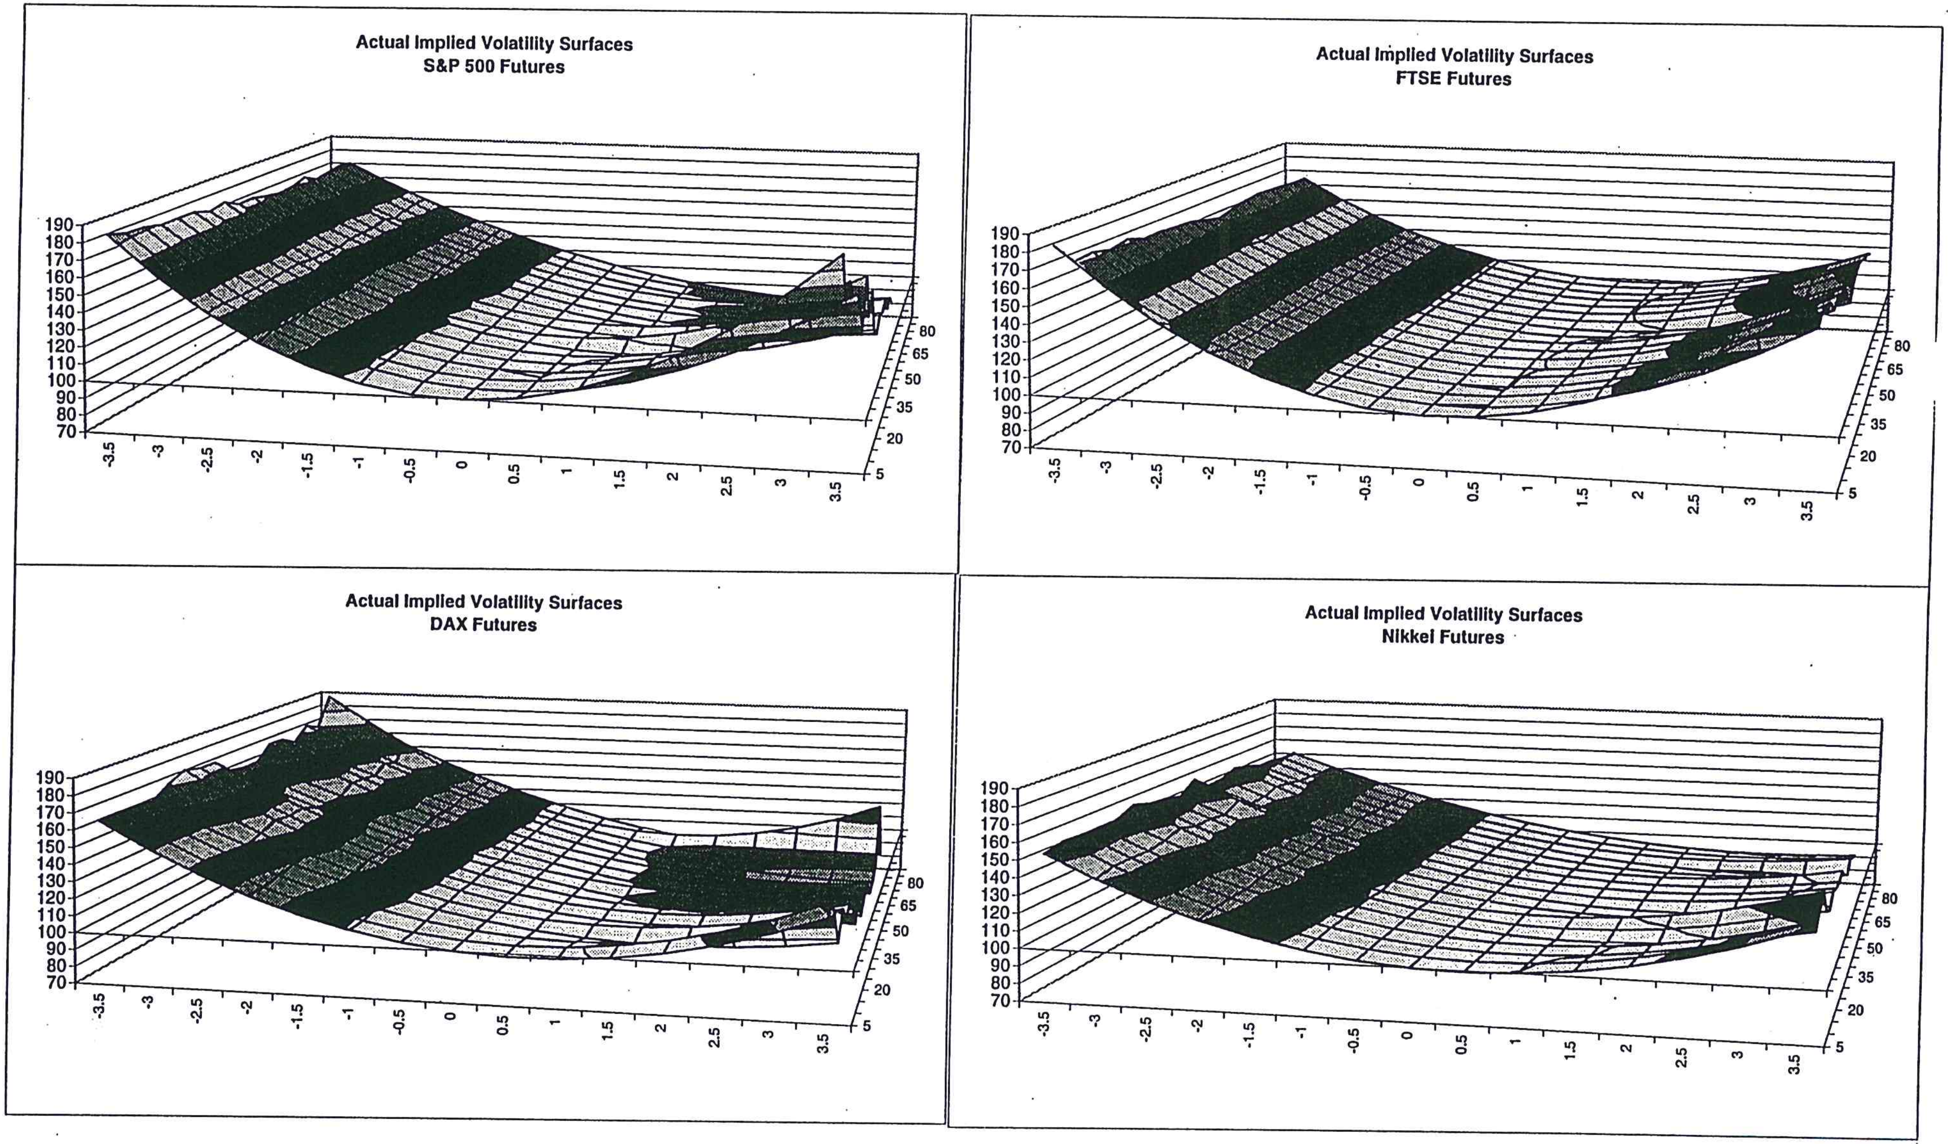

Figure 2, Implied Volatility Surfaces From Options on Four Stock Index Futures Markets: 1990-1998.

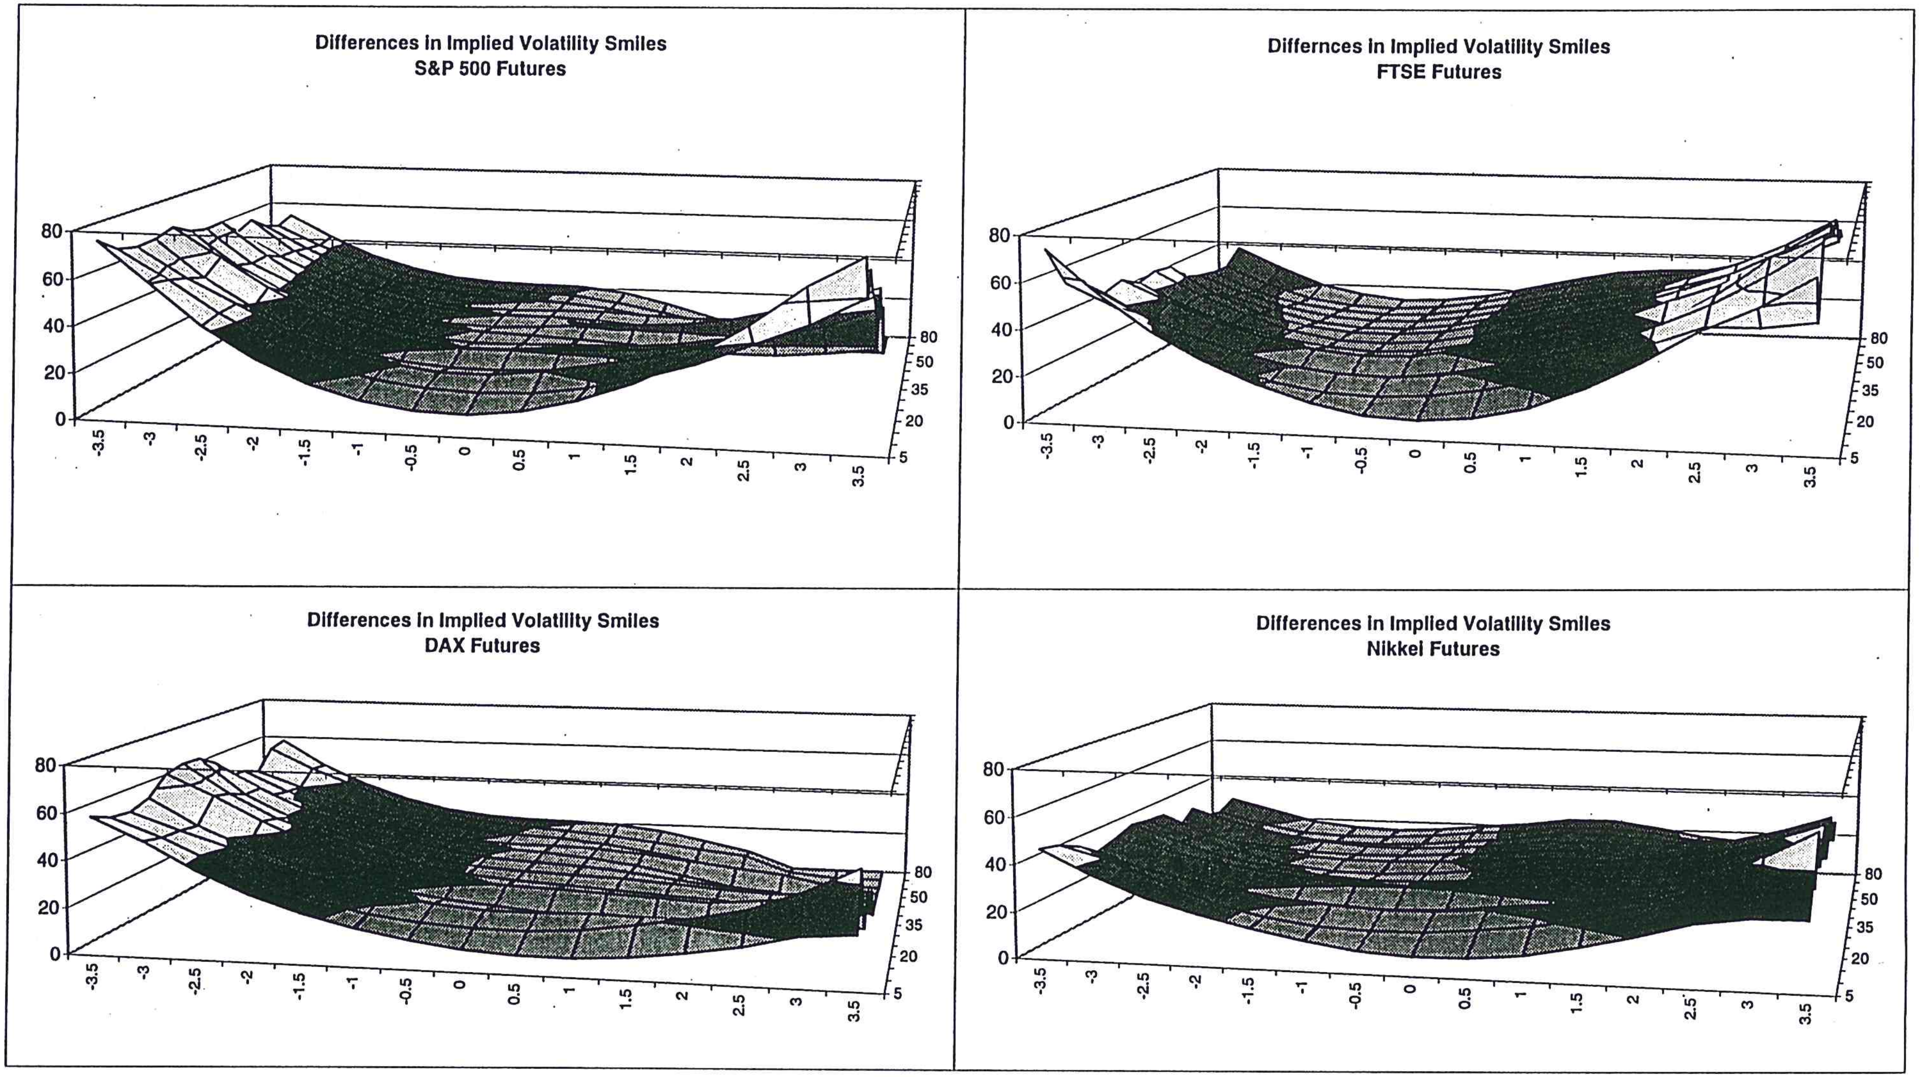

Figure 3, Differences in Simulated and Actual Implied Volatility Surfaces for Four Stock Index Futures/Options Markets: 1990-1998.