Stock Index Futures Markets: Stochastic Volatility Models and Smiles

ROBERT G. TOMPKINS

This study examined whether the inclusion of an appropriate stochastic volatility that captures key distributional and volatility facets of stock index futures is sufficient to explain implied volatility smiles for options on these markets. I considered two variants of stochastic volatility models related to Heston (1993). These models are differentiated by alternative normal or nonnormal processes driving log-price increments. For four stock index futures markets examined, models including a negatively correlated stochastic volatility process with nonnormal price innovations performed best within the total sample period and for subperiods. Using these optimal stochastic volatility models, I determined the prices of European options. When comparing simulated and actual options prices for these markets, I found substantial differences. This suggests that the inclusion of a stochastic volatility process consistent with the objective process

This piece of research was partially supported by the Austrian Science Foundation (FWF) under Grant SFB#10 (Adaptive Information Systems and Modelling in Economics and Management Science). This article has benefited from comments by Stewart Hodges, Walter Schachermayer, Friedrich Hubalek, Stefan Pichler, Ole Barndorff-Nielsen, Neil Shephard, and attendees at the University of Aarhus minisymposium on volatility and the Austrian Working Group on Banking and Finance. The author would also like to thank the editor of this journal and the two referees for extremely helpful suggestions. As always, the author is responsible for all remaining errors.
For correspondence, Robert G. Tompkins, Doeltergasse 33, A-1220 Wien, Austria; e-mail: rtompkins@ins.at

Received May, 1999; Accepted March, 2000

■ *Robert Tompkins is a University Dozent in the Department of Finance at the Vienna University of Technology and in the Department of Finance at the Institute for Advanced Studies.*

alone is insufficient to explain the existence of smiles. Jrl Fut Mark 21:43–78, 2001

INTRODUCTION

Soon after the publication of Black and Scholes' (1973) article on option pricing, Black (1975) pointed out that the constant volatility assumption may be incorrect, noting that the volatility may be a function of the underlying price level. Subsequent research has identified the existence of different volatilities implied from option prices for different strike prices and terms to maturity (see Jackwerth & Rubinstein, 1996, for a review of this literature). This effect has been commonly called the *implied volatility smile* (for options with the same term to expiration) and the *term structure of volatility* (for options with different terms to expiration).

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

This article examines the nature of the objective dispersion processes for stock index futures and proposes the inclusion of stochastic volatility to explain these empirical facets of the objective price processes. Once this is done, the implications for option pricing are considered. Previous research has examined the effects of stochastic volatility (in the underlying objective price process) on option pricing (see Heston, 1993; Hull & White, 1987; Johnson & Shanno, 1987; Melino & Turnbull, 1990; Scott, 1987; Stein & Stein, 1991; Wiggins, 1987). This article extends this line of research.

The models considered include nonnormality in the underlying log-price process and correlated price and volatility processes. Because of the rich nature of these models, they do not lend themselves to parameterization using standard methods (e.g., maximum likelihood methods), and I utilized a simulated method of moments approach. I chose target empirical moments to address most stylized results previously pointed out in the literature. Specifically, I considered higher moments of the return series and dynamics of the volatility process and leverage effects. Once I found suitable parameter values for alternative models and derived a feasible risk-neutral drift adjustment, I determined options prices numerically, assuming the drift adjustment was unique.[1]

To allow comparisons between the simulated and actual option prices, all prices were reexpressed as standardized implied volatilities and were plotted as surfaces. I found substantial differences between these surfaces for the Standard & Poor's (S&P) 500 futures market, and this suggests that the inclusion of stochastic volatility for the objective process alone is insufficient to explain the existence of implied volatility smiles. However, I found consistent divergences between simulated and actual implied volatility surfaces across for four markets, and this may provide insights into the nature of market imperfections and risk premia. Scott (1987) completed similar research on stock options and found similar results. He suggested that further research should test the effects of stochastic volatility on the pricing of options for a larger sample. This study achieved this by the examination of four stock index futures and options markets for a period encompassing the decade of the 1990s.

The article is organized as follows. In the first part, I briefly review previously presented empirical evidence that indicates that stock index futures prices do not follow an independent and identically distributed (i.i.d.) geometric Brownian motion (GBM) process. I consider two possible models that have been proposed in the literature to explain these results. Seven empirical attributes were selected to capture key aspects of nonnormality and interdependence. A description of the data sources used for this research follows. After this, I examine both possible models, using a simulated method of moments approach to assess their ability to explain the key attributes. Having determined the best model, I estimate option prices consistent with these processes and then express them as implied volatilities, using the Black (1976) formula. I compare these simulated implied volatility surfaces to the implied volatility surfaces from traded options on the S&P 500 futures markets. Finally, I present conclusions and suggestions for further research.

PRICE PROCESSES FOR STOCK INDEX FUTURES UNDER THE EMPIRICAL MEASURE

It is well established that the unconditional return series for individual stocks, stock indexes, and stock index futures do not conform to the as-

[1] A number of previous researchers have made similar assumptions. Scott (1987), Hull and White (1987), and Johnson and Shanno (1987) assumed either that the market price of volatility risk was zero or that the existence of a marketable asset perfectly correlated to volatility.

sumptions of an i.i.d. log-normal dispersion process (see Stoll & Whaley, 1990). Returns for stock index futures usually display excess kurtosis and (for many periods) significantly negative skewness compared to a normal distribution. Returns for futures also display intertemporal dependence. In the examination of absolute returns for stock index markets, significant positive autocorrelations have been found. This result led Ding, Granger, and Engle (1993) to conclude that "It is clear that the S&P 500 stock market return process is not an i.i.d. process" (p. 87).

Another line of empirical research has examined the unconditional volatility series directly. A typical approach to better understand the volatility process is to examine the statistical moments of the process. Burghardt and Lane (1990) examined the variability of the unconditional volatility process with a volatility cone approach. I 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 nonoverlapping data, I obtained an average estimate of the unconditional volatility and the standard deviation of this average. Under the assumption that an i.i.d. price process is generating these volatilities, the expected coefficient of variation of the 20-day volatility is known. This measures the variability of volatility for a given time horizon.

However, the time-varying dynamics of the variability of volatility as the time horizon of estimation is extended is of additional interest. I chose a simple log-linear form to capture these dynamics. The rate of decay in the standard deviation of volatility implies that the maturity structure of historical volatility experiences long-term memory (interdependence). This can be seen as complimentary to long-term memory effects for absolute returns identified by Ding et al. (1993). An additional and important feature of these price processes is the leverage effect pointed out by Christie (1982) among others.

Although a number of theories have been proposed to explain these results, three alternative hypotheses are examined here. The constant elasticity of variance (CEV) model of Cox and Ross (1976), a non-GBM price process model (e.g., the jump diffusion model proposed by Merton, 1976), and stochastic volatility models are 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 Bakshi, Cao, & Chen, 1997; Bates, 2000; Heston, 1993). Variants of the Heston model proposed by

Barndorff-Nielsen (1997) and Barndorff-Nielsen and Shephard (1999) are also considered.

Model 1

$$dF = \mu F dt + \sigma F dZ(t) \tag{1}$$

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. GBM.

It is clear that Model 1 is a straw man, given the well-known results that stock index futures returns display both nonnormality and interdependence. 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 consider stochastic volatility ($\hat{\sigma}$), which is evaluated in terms of a stochastic variance process ($\sqrt{V}$).

Model 2

$$dF = \mu F dt + \hat{\sigma} F dZ_1 \tag{2.1}$$

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 volatility of the process, respectively. This model is called SV (Stochastic Volatility) in this article.

Model 3

$$dF = \mu F dt + \hat{\sigma} F dN \tag{3.1}$$

With the variance process defined by

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

where N represents a nonnormal price process for the underlying price series and, in this research, is the normal inverse Gaussian (NIG) distribution. 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, 1999b). This approach extends the findings of Bates (1996, 2000) and Ho, Perraudin, and Sørensen (1996), who assumed the volatility process is subordinated in a nonnormal price process. In this model, the stochastic variance process is assumed to follow a standard Wiener process, Z, with correlation ρ existing between the two processes.[2] When correlated processes are considered, this variant of Model 3 is related to the model for incorporating a leverage effect in a stochastic volatility model proposed by Barndorff-Nielsen and Shephard (1999). This model is called *NIGSV* in this article. The CEV model of Cox and Ross (1976), which allows for the inclusion of negative correlations between the two stochastic processes, is nested in both Models 2 and 3. This allows the leverage effect to be captured.

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) provided a good review of the approaches used to parameterize stochastic volatility models. Many of the models considered here do not lend themselves to estimation by traditional maximum likelihood methods. Andersen et al. proposed the use of the efficient method of moments approach suggested by Gallant and Tauchen (1996). Their primary contributions were to understand the sample properties of this estimator and show that the method was robust in larger sample sizes. Recently, Andersson (1999a) 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. Because of the complexities of my models, it is not clear whether estimation of parameter inputs by maximum likelihood techniques is feasible. Given that these articles have demonstrated that parameter estimation via simulation can be equally robust and that this approach may provide a better intuitive understanding of the estimation procedure, a simulated method of moments approach similar in spirit to the general method of moments (GMM) was utilized.

[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 an NIG distribution and that the volatility process follows GBM.

In this research, I used a hybrid between the GMM approach of Melino and Turnbull (1990) and the simulated method of moments approach of Duffie and Singleton (1993). This approach subjectively selects essential attributes, simulates price processes consistent with the alternative models, and assesses the sum of squared errors (SSE) between simulated and empirical attributes. I examined and optimized alternative parameterizations of the models. Similarly to Andersen et al. (1999), I investigated the sample properties of this estimator technique (using Model 1). This allowed me to make comparisons between the attributes of each market and the models and to draw conclusions regarding the overall fit of each model.

At the heart of this estimation technique is the judicious choice of key attributes. It is crucial that the choice of the attributes jointly considers the relevant elements of empirical interdependence and nonnormality and provides a means by which the salient features of the alternative models can be captured. Because 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 indexes and stock index futures. Both Theodossiou (in press) 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, leverage effects, or both.

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

Clearly, because this research examined two variants of a 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 with squared returns, and these results were annualized (with the assumption of 252 trading days in a year). A number of attributes were selected to capture critical dynamics of the volatility process.

The first attribute examines the volatility of volatility. I estimated a time series of unconditional volatilities on a daily basis for a time horizon of 20 days (using nonoverlapping data). From this series, I estimated the average and standard deviation. Because of 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 i.i.d. $[1/\sqrt{(2 \times 19)}]$.

Although 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 instead achieved by the estimation of the coefficient of variation of volatility with time horizons from 20 days to 200 days in 20-day increments.[3] From statistical theory, the expected decay in the standard deviation of a volatility estimate (σ^*) is a square-root function of the number of observations used for estimation ($SE = \sigma^*/\sqrt{2N}$). Given that the average level of volatility remains constant (which was found empirically for these four markets), the decay in the coefficient of variation should follow the same functional form. If the first volatility observation is at the 20-day horizon, the decay in the standard deviation of volatility from that point forward can be expressed as $\sigma_N \cdot \sqrt{N/(N + 1)}$. 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 $SE = \sigma^*/\sqrt{2} \times N^{-0.5}$. A linear regression of (the natural logarithms of) the time horizon of estimation regressed on 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 this 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.

[3]The estimation of the volatilities was done with overlapping data. This introduced a bias in the estimation of the standard deviation of volatility. This bias was corrected with the Hodges and Tompkins (2000) approach.

Another salient feature of empirical return volatility series is that subsequent realizations may not be independent. To capture serial correlation in absolute returns, autocorrelation dynamics were examined directly rather than with alternative methods that rely on maximum likelihood. This was achieved through examination of the autocorrelograms of absolute returns previously employed by Taylor (1986) and Ding et al. (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 were required. Because 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 that allow for a simple parametric confidence interval test. Therefore, to assess the sample characteristics of these composite measures, nonparametric confidence intervals were determined via simulation. These two attributes provide additional information relevant to the calibration of a stochastic volatility model, as they capture both short-term and medium-term evidence of volatility clustering.

The final attribute must measure the leverage effect and provide a means for a correlation between the stochastic processes to be captured. There is, however, one problem with the determination of the leverage

effect: Even if the volatility is a stationary series, the prices are not. To solve this problem, a new variable was constructed that measures recent price movements and is stationary.[4] This variable is an exponentially weighted return series that indicates whether recent price movements are relatively high or low. It can be shown that with 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

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

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

[4]The author would like to thank Stewart Hodges (University of Warwick) for suggesting the use of this variable.

(W) for all markets imposed at 0.03, the correlation between the two variables was estimated.[5] This correlation coefficient serves as an attribute to measure the leverage effect.

With these seven target attributes, the stochastic volatility models were parameterized via simulation. Following the three models proposed previously, price series of 1,500 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 with an SSE statistic. To reduce scaling impacts resulting from different levels of the attributes, the squared errors were divided by the standard deviation of the attributes across the four markets.[6] This test statistic can be written as

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

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 1,500 prices consistent with Model 1 (GBM) were drawn to better understand the sample properties of all the attributes and the test statistic. This provided a nonparametric estimation of the standard errors of the attributes and allowed test statistics for comparisons between markets and models.

DATA SOURCES

The futures markets examined include the S&P 500 and Nikkei 225 traded at the Chicago Mercantile Exchange, the DAX (Deutsche Aktien Index) index futures traded at the Deutsche Terminbörse, and the FTSE (Financial Times Stock Exchange Index) 100 futures traded at London International Financial Futures Exchange. Daily closing futures prices were analyzed and were obtained directly from the relevant exchange, and

[5] To simplify the selection of the attribute, a fixed weight of 0.03 was applied to all markets and for all periods of analysis to allow the leverage correlation factor to not be subject to differing weights. This was close to the optimal weights for each market with a maximum likelihood estimation procedure.

[6] In addition, the natural logarithm of the unconditional kurtosis was examined rather than the absolute levels. This was done because of wide variations in this statistic across the four markets. However, all results are presented as absolute levels.

TABLE I

Markets Included in the Research, Time Periods of the Data, and Number of Observations

Underlying Asset	*Time Period of Analysis*		
Entire period	*First Date*	*Last Date*	*Number of Observations*
S&P 500 futures	23/11/90	17/12/98	2,042
FTSE 100 futures	23/11/90	18/12/98	2,042
DAX futures	23/11/90	18/12/98	2,020
Nikkei 225 futures	23/11/90	11/12/98	2,037
First period	*First Date*	*Last Date*	
S&P 500 futures	23/11/90	15/12/94	1,029
FTSE 100 futures	23/11/90	16/12/94	1,029
DAX futures	23/11/90	16/12/94	1,016
Nikkei 225 futures	23/11/90	08/12/94	1,023
Second period	*First Date*	*Last Date*	
S&P 500 futures	19/12/94	17/12/98	1,013
FTSE 100 futures	19/12/94	18/12/98	1,013
DAX futures	19/12/94	18/12/98	1,004
Nikkei 225 Futures	09/12/94	11/12/98	1,014

Note: Futures prices are the closing daily levels for the nearest-to-expiration contract. S&P = Standard & Poor's; FTSE = Financial Times Stock Exchange Index; DAX = Deutsche Aktien Index.

the period of analysis was restricted to a period in which all four stock index futures were traded. The underlying assets, the time period of analysis, and the number of observations used in the analysis appear in Table I. To assess if results were period specific, the data were also split into roughly equal time periods from 1990 to 1994 and from 1995 to 1999. The time periods and number of observations in these split periods also appear in Table I.

Because 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. First, 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 the comparison of 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 key punch error had occurred. We discovered that only 1 to 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 part 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 was humanly possible.

ATTRIBUTES FOR INDIVIDUAL MARKETS

For each market, the return statistics for daily returns were determined for the entire period of analysis and for each subperiod. The results of these analyses can be seen in Table II. The first column describes the market under investigation and indicates the time period examined. The second and third columns present the mean and standard deviation of the return distribution. The fourth and fifth columns present the statistics for the unconditional skewness, and these provide a significance level relative to a null hypothesis of normality.[7] All skewness significance statistics that are significantly different from the normal assumption at a 95% level (± 1.96) appear in bold type. In the sixth and seventh columns, the unconditional kurtosis statistic and the significance level relative to a null hypothesis of normality appear.[8] When this statistic does not reject the null hypothesis of normality at the 95% level or above, the statistic and the significance levels are in bold 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 chi-square with two degrees of freedom. The critical value at the 1% level is 9.21. When the BJ statistic exceeds this level, this statistic also appears in bold type.

For all four markets, the dispersion of returns was not well described by a normal distribution. For many of the four markets, the skewness

[7]Under the null hypothesis of normality, the skewness statistic is asymptotically normally distributed with standard errors: $SE = \sqrt{(6/T)}$, where T is the sample size.

[8]Under the null hypothesis of normality, the excess kurtosis statistic is asymptotically normally distributed with standard errors: $SE = \sqrt{(24/T)}$, where T is the sample size. This statistic is equal to the kurtosis statistic appearing in Table II minus 3.0.

TABLE II

Statistics of the Daily Returns for Four Stock Index Futures Markets (1990–1998)

Underlying Asset	*M*	*SD*	*Unconditional Skewness*	*Skewness Significance Level*	*Unconditional Kurtosis*	*Kurtosis Significance Level*	*BJ Statistic*	*Number of Observations*
Standard & Poor's 500 futures								
Overall period	0.00055	0.00907	**−0.4397**	**−8.11**	**11.401**	**77.47**	**6067.18**	2,041
First period	0.00030	0.00717	**0.2074**	**2.72**	**6.078**	**20.14**	**413.14**	1,028
Second period	0.00079	0.01066	**−0.6477**	**−8.42**	**10.884**	**51.22**	**2694.11**	1,013
Financial Times Stock Exchange Index 100 futures								
Overall period	0.00038	0.01000	−0.0944	−1.74	**5.979**	**27.47**	**757.87**	2,041
First period	0.00014	0.00950	0.1322	1.73	**4.457**	**9.54**	**93.94**	1,028
Second period	0.00063	0.01048	**−0.2779**	**−3.61**	**6.983**	**25.88**	**682.76**	1,013
Deutsche Aktien Index futures								
Overall period	0.00043	0.01272	**−0.6466**	**−11.86**	**11.744**	**80.20**	**6572.19**	2,019
First period	0.00011	0.01200	**−0.8826**	**−11.48**	**19.675**	**108.44**	**11891.68**	1,015
Second period	0.00075	0.01340	**−0.4865**	**−6.29**	**6.408**	**22.04**	**525.49**	1,004
Nikkei 225 futures								
Overall period	−0.00037	0.01538	**0.1161**	**2.14**	**4.919**	**17.68**	**317.04**	2,036
First period	−0.00039	0.01516	**0.1943**	**2.54**	**4.447**	**9.44**	**95.58**	1,022
Second period	−0.00035	0.01560	0.0437	0.57	**5.352**	**15.29**	**234.01**	1,014

Note: These returns were based on closing futures prices for the nearest-to-expiration futures contracts. The first column describes the market under investigation and the period of analysis (entire period, 1990–1998; first period, 1990–1994; second period, 1994–1998). The second and third columns present the first (mean) and second (standard deviation) moments of the return distribution. The fourth and sixth columns present the statistics for the unconditional skewness and kurtosis. Under the null hypothesis of normality, the skewness statistic is normally distributed with standard errors, *SE* = *sqrt*(6/*T*), where *T* is the sample size. This appears in the last column. The fifth column indicates a significance statistic testing a normality hypothesis. If the skewness statistic rejects this at a 95% level or greater, this is indicated in bold text. Under the null hypothesis of normality, the excess kurtosis statistic is normally distributed with standard errors, *SE* = *sart*(24/*T*). This statistic is equal to the kurtosis statistic appearing in the table minus 3.0. Column 7 indicates a significance statistic testing the null hypothesis of normality. If the kurtosis statistic rejects this at a 95% level or above, this is indicated in bold text. The statistic in the sixth column is the Bera-Jarque (BJ) statistic for detecting departures of the data from normality. Under the null hypothesis of normality, the BJ statistic is distributed as chi-square with two degrees of freedom. The critical value at the 1% level is 9.21. When the BJ statistic is above this level, this statistic appears in bold text.

statistic tended to be significantly different from that of a normal distribution. However, the skewness was 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 displayed significant excess kurtosis. These two factors led to all BJ values exceeding their critical values. These results are consistent with previous empirical examination of return series for stock markets by Theodossiou (in press) and Harvey and Siddique (1998).

As the remaining attributes all capture characteristics of the volatility process, these are summarized in a single table, Table III. In this table,

TABLE III

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

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)*
Standard & Poor's 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**
Financial Times Stock Exchange Index 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**
Deutsche Aktien Index 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
SE of attributes	3.244%	0.032%	0.010	0.0778	0.0999	0.0051	0.0061

Note: The first column, the stock index futures markets and the time periods of the analysis appear. In the bottom two rows appear the attribute values expected from an independent and identically distributed (i.i.d.) price process associated with geometric Brownian motion (GBM). The standard error of the attributes was determined for a series of 500 draws from an i.i.d. GBM process of 1,500 observations. These nonparametrically estimated standard errors allowed significance testing between the empirical attributes and the assumption of an i.i.d. GBM process. When a *T* statistic rejects this assumption at a 95% level or above, the attributes are bold. In Columns 2, 3, and 4, an analysis of the 20-day volatility estimated on a nonoverlapping basis appears. With all available observations, daily returns (differences in the logarithm of daily closing stock index futures prices) were estimated. With these returns, the standard deviation was estimated for a fixed time horizon of 20 days and then annualized with the *sqrt*(252). The estimation of the 20-day volatility was done on a nonoverlapping basis. In the fifth column appears the relationship between the standard deviation of volatility and the time horizon of estimation. The time-factor term is determined by a regression of the logarithm of the time horizon (*T*) on the logarithm of the unbiased standard deviation of volatility. Only the beta coefficients of the regression appear. Observed divergences in the beta coefficient from −0.5 give an indication of the degree of the volatility of volatility persistence observed in the unconditional process. The sixth column provides information about the leverage relationship between the levels of stock index futures prices and volatility. This is the correlation between the unconditional volatility estimated with daily log price increments and estimated on a 20-day rolling time horizon and an exponentially weighted return series, which indicates whether recent price movements are relatively high or low. For the FTSE (Financial Times Stock Exchange Index) 100 futures in the first period and the Nikkei 225 futures in the first period, we cannot reject the assumption of a leverage effect. In the last two columns, information appears about the autocorrelations in absolute returns (and volatility clustering): composite measures of autocorrelograms were required. Given the relatively rapid decay of the autocorrelations, we averaged the autocorrelations for each lag from lag 1 to 20 and separately averaged the autocorrelations for each lag from lag 51 to 70.

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

In the fifth column, 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 is shown. 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 was (statistically) significantly less than this decay function (the standard error of this attribute was estimated in a nonparametric manner by simulation). In column 6, the leverage correlation coefficient appears. This measures the relationship between the 20-day unconditional volatility and the recent relative prices. Consistent with the negative leverage effects Christie (1982) pointed out for individual stocks, stock index futures markets also seem to have had a significant negative leverage effect. The exceptions were the FTSE 100 and S&P 500 futures over the first period (1990–1994). These results should be interpreted with care because the simulated standard error of this attribute was fairly high (0.0999). Although these effects remained significant, it was only (barely) at the 95% level.

The final two columns represent the average autocorrelations of absolute returns for lagged periods from 1 to 20 days and from 51 to 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 were statistically significant and positive. However, the degree of autocorrelations seemed to be higher in the second period relative to the first period.

The results listed in 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 as target conditions, the proposed alternative models were examined.

FITTING ALTERNATIVE MODELS

In previous sections, the models to be tested were presented. Only technical issues in the simulation process and the parameterization of the stochastic volatility models remain to be discussed 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 1,500 draws from an independent normal distribution. According to standard procedures, the random number-generation process used a standard Box–Muller method and the antithetic approach suggested by Boyle (1977). With these 500 samples, price series were constructed that conformed to GBM with constant variance (with 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. Utilizing Equation 6, I determined the SSE between each of the 500 samples and the true attributes of an i.i.d. GBM process. The two distributions with the lowest SSEs relative to the priors were selected as representative normal distributions. For convenience, the normal disturbances for the underlying price process are referred to as Z_1 and the normal disturbances for the volatility process are 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. 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 SSE statistic. The standard deviation of this statistic was found to be equal to 4.5646 and was 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, I generated random numbers with 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) was set 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 because the stochastic volatility interacts with the NIG distribution and yields simulated moments that are amplified compared with the NIG moments. Because these effects could not be ascertained before the simulation, four NIG distributions were simulated, and all were examined as

TABLE IV

Attribute Values for Simulated Geometric Brownian Motion (GBM) Processes, Standard Deviations of Attributes, and Two Representative Processes

GBM (i.i.d.) Process	*CoV 20-Day Volume*	*Unconditional Skewness*	*Unconditional Kurtosis*	*Leverage Correlation*	*Autocorrelation (1–20 Lags)*	*Autocorrelation (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
SD 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

Note: The first column indicates various simulations of an independent and identically distributed (i.i.d.) Wiener process. The row titled "Expected Results" is not based on simulation but is what is expected from statistical theory. The next two rows, "Average Results of 500 Simulations" and "Standard Deviation of 500 Simulations," indicate the sampling properties of 500 series of simulated prices of 1,500 observations based on random draws from an i.i.d. GBM process. The next two rows represent the two draws from the i.i.d. GBM process, which had the lowest sum of squared errors relative to the expected theoretical results. These were used in the second stage of the analysis as representative normal distributions and are called Z_1 and Z_2. Columns 2 to 8 indicate the actual attribute values. The last column indicates the sum of squared errors (SSE) statistic for that row. Of importance to a later stage of analysis is the average standard deviation of the SSE. This result allowed comparisons to be made between alternative models.

TABLE V

Sample Statistical Moments of Simulations of Four Normal Inverse Gaussian (NIG) Distributions

NIG Distribution	*M*	*SD*	*Skewness*	*Kurtosis*
NIG1	0.01821	1.00088	−0.0049	4.7807
NIG2	0.00078	0.93955	−0.9289	8.7576
NIG3	−0.00665	1.01818	−0.1936	3.4129
NIG4	0.0035	0.99176	−0.0794	7.627

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

potential candidates for the underlying price process. These four possible NIG distributions appear in Table V as NIG1 to NIG4.

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

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

where a_t is the risk-neutral adjusted drift; α, β, δ, and μ are the parameters of the NIG distribution; and σ_{t-1} is the random volatility (from the stochastic volatility process in Equation 3.2) for the previous discrete observation. A complete proof of the derivation of this drift adjustment appears in the Appendix.[9] This drift was then used in Equation 3.1 with the drift term (μ) equal to a_{t-1} from Equation 7, and the disturbances were drawn from the NIG distribution.

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

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

[9]In the Appendix, it is noted that this is only one way to adjust this drift to achieve risk neutrality. Because this model implies markets are incomplete, there may not exist a unique martingale measure. However, this is certainly a feasible adjustment and will allow relative comparisons to be drawn.

where Z represents the random disturbances for the underlying price process (for GBM, the Z_1 set of normal draws; for an 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.

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 the second was to determine the attribute values for this simulated series. The third step entailed varying the parameter inputs into the models to minimize the SSE relative to the observed empirical attributes. To efficiently perform the third step, I started by examining the sensitivities of the overall SSE to a small incremental change in each of the parameter inputs (holding the other parameters constant). Of the four variables in both of the stochastic volatility models, I 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 were the rate of mean reversion, κ; the volatility of the volatility, ξ; and the correlation between the price and volatility process, ρ. I found that small changes in the level of the long-term volatility had almost no effect on the SSE. Because only three parameters needed to be varied, the parameterization simply compared the SSEs with the initial seeded parameter values to the SSE for the same model (and random numbers) but with variation of the three critical parameters (κ, ξ, and ρ). By varying only three parameters both up and down, I had to compare only eight alternatives to the original results. If one of the new combinations of parameters achieved a lower SSE, 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 (e.g., 1.0 for κ and 0.1 for ξ and ρ). When optimal parameters were found, the increments for the search were progressively reduced (e.g., as low as 0.01 for κ and 0.0001 for ξ and ρ) until no further reduction in the SSE was achieved. When no further improvement in the SSE was possible, the final combination of parameter values was deemed the optimal estimation of the alternative stochastic volatility models.

Tables VI through IX display the empirical attributes (taken from Tables II and III) and the results of the models for each of the stock index futures markets. These results are 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 parameterization of each of the three models appears.

TABLE VI

Empirical Attributes of Standard & Poor's (S&P) 500 Futures and Attributes of Three Alternative Models

Models (Parameters)						Empirical and Simulated Attributes								
						Empirical Attributes								
						CoV	Skewness	Kurtosis	Leverage	Corr(1–20)	Corr(51–70)	Line Fit		
S&P 500		Whole Period				0.5033	−0.4397	11.4006	−0.2174	0.1465	0.0903	−0.1359		
futures	Distribution	LTV	K	VoV	Correlation	Simulated Attributes							SSE	T
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	3.10	0.890	−0.25	0.4471	−0.4685	5.9502	−0.2742	0.1763	0.0765	−0.1086	6.571	54.651
Model 3	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	55.349
						Empirical Attributes								
						CoV	Skewness	Kurtosis	Leverage	Corr(1–20)	Corr(51–70)	Line Fit		
S&P 500		First Period				0.3180	0.2074	6.0778	0.2049	0.0563	0.0391	−0.1943		
futures	Distribution	LTV	K	VoV	Correlation	Simulated Attributes							SSE	T
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	2.39	0.172	0.035	0.2928	0.0161	3.5894	0.1824	0.0852	0.0371	−0.1456	2.686	26.503
Model 3	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	26.868
						Empirical Attributes								
						CoV	Skewness	Kurtosis	Leverage	Corr(1–20)	Corr(51–70)	Line Fit		
S&P 500		Second Period				0.5350	−0.6206	11.2081	−0.3999	0.1541	0.0834	−0.1191		
futures	Distribution	LTV	K	VoV	Correlation	Simulated Attributes							SSE	T
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	0.700	−0.4	0.5072	−0.6220	6.5818	−0.3921	0.2139	0.0796	−0.1171	4.899	85.247
Model 3	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	85.689

Note: This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the S&P 500 futures. These results were split into the three periods of analysis. On the left-hand side, the parameterization of each of the three models appears. Only the parameter values and the normal inverse Gaussian (NIG) distribution that has the lowest sum of squared errors (SSE) are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison. In the column immediately to the right of this, SSE for each model appears. The final column indicates a *T*-test statistic comparing Models 2 and 3 to the geometric Brownian motion (GBM) assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV. LTV represents the long term volatility, θ in equations 2.2 and 2.3. K represents the rate of mean reversion, **K** in equations 2.2 and 2.3. VOV represents the volatility, ξ in equations 2.2 and 2.3. Corr (1–20) represents the average auto correlation of absolute returns for time lags 1 to 20 days. Corr (51–70) represents the average autocorrelation of absolute returns for time lags 51 to 70 days.

TABLE VII

Empirical Attributes of *Financial Times* Stock Exchange Index 100 Futures and Attributes of Three Alternative Models

		Models (Parameters)				*Empirical and Simulated Attributes*								
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
FTSE 100		*Whole Period*				0.3917	−0.0944	5.9793	−0.301	0.1330	0.0679	−0.1600		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	*Simulated Attributes*							*SSE*	*T*
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	3.01	1.001	−0.29	0.4858	−0.4720	5.7005	−0.2829	0.1965	0.0122	−0.1969	3.100	45.670
Model 3	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	46.108
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
FTSE 100		*First Period*				0.2760	0.1322	4.4571	−0.1227	0.0615	0.0274	−0.1584		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	*Simulated Attributes*							*SSE*	*T*
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	2.43	0.251	−0.2	0.2723	−0.0962	3.5016	−0.1034	0.0684	0.0278	−0.1795	0.570	22.445
Model 3	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	21.976
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
FTSE 100		*Second Period*				0.4555	−0.2613	6.5835	−0.4536	0.1725	0.0954	−0.1001		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	*Simulated Attributes*							*SSE*	*T*
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	4.10	0.750	−0.5	0.4706	−0.6577	6.5588	−0.4462	0.1847	0.0832	−0.1088	2.468	86.026
Model 3	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	86.134

Note: This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the FTSE 100 futures. These results are split into the three periods of analysis. On the left-hand side, the parameterization of each of the three models appears. Only the parameter values and the normal inverse Gaussian (NIG) distribution that has the lowest sum of squared errors (SSE) are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison. In the column immediately to the right of this, the SSE for each model appears. The final column indicates a *T*-test statistic comparing Models 2 and 3 to the geometric Brownian motion (GBM) assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV. LTV represents the long term volatility, θ in equations 2.2 and 2.3. K represents the rate of mean reversion, **K** in equations 2.2 and 2.3. VOV represents the volatility, ξ in equations 2.2 and 2.3. Corr (1–20) represents the average auto correlation of absolute returns for time lags 1 to 20 days. Corr (51–70) represents the average autocorrelation of absolute returns for time lags 51 to 70 days.

TABLE VIII

Empirical Attributes of Deutsche Aktien Index Futures and Attributes of Three Alternative Models

Models (Parameters)						Empirical and Simulated Attributes								
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
DAX		*Whole Period*				0.5186	−0.6466	11.7437	−0.35339	0.1806	0.0817	−0.2000		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	Simulated Attributes							*SSE*	*T*
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	3.60	1.090	−0.35	0.4714	−0.5806	6.3360	−0.3619	0.1880	0.0721	−0.1337	7.652	53.348
Model 3	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	54.538
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
DAX		*First Period*				0.4916	−0.8826	19.6753	−0.2395	0.0714	0.0282	−0.2650		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	Simulated Attributes							*SSE*	*T*
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	6.00	1.400	−0.7	0.4105	−0.6305	5.7878	−0.5376	0.1212	0.0217	−0.2302	9.856	18.861
Model 3	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	20.363
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
DAX		*Second Period*				0.5172	−0.4372	6.2881	−0.4614	0.2512	0.1158	−0.1001		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	Simulated Attributes							*SSE*	*T*
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	3.99	0.795	−0.51	0.4657	−0.6172	6.0690	−0.4912	0.1876	0.0733	−0.1242	2.937	94.739
Model 3	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	94.728

Note: This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the DAX futures. These results are split into the three periods of analysis. On the left-hand side, the parameterization of each of the three models appears. Only the parameter values and the normal inverse Gaussian (NIG) distribution that has the lowest sum of squared errors sum of squared errors (SSE) are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison. In the column immediately to the right of this, the SSE for each model appears. The final column indicates a *T*-test statistic comparing Models 2 and 3 to the geometric Brownian motion (GBM) assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV. LTV represents the long term volatility, θ in equations 2.2 and 2.3. K represents the rate of mean reversion, $\mathbf{K}$ in equations 2.2 and 2.3. VOV represents the volatility, ξ in equations 2.2 and 2.3. Corr (1–20) represents the average auto correlation of absolute returns for time lags 1 to 20 days. Corr (51–70) represents the average autocorrelation of absolute returns for time lags 51 to 70 days.

TABLE IX

Empirical Attributes of Nikkei 225 Futures and Attributes of Three Alternative Models

Models (Parameters)						Empirical and Simulated Attributes								
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
Nikkei		*Whole Period*				0.3940	0.1161	4.9192	−0.239	0.1308	0.0397	−0.1934		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	Simulated Attributes							*SSE*	*T*
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	5.35	1.016	−0.315	0.3946	−0.3448	4.9679	−0.2255	0.1409	0.0199	−0.1815	3.021	36.507
Model 3	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	36.862
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
Nikkei		*First Period*				0.3761	0.1943	4.4470	−0.1946	0.1366	0.0355	−0.1958		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	Simulated Attributes							*SSE*	*T*
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	4.48	0.762	−0.21	0.3698	−0.1919	4.2910	−0.0708	0.1353	0.0358	−0.1588	1.604	27.226
Model 3	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	27.137
						Empirical Attributes								
						CoV	*Skewness*	*Kurtosis*	*Leverage*	*Corr(1–20)*	*Corr(51–70)*	*Line Fit*		
Nikkei		*Second Period*				0.4338	0.0528	5.6695	−0.2561	0.1421	0.0574	−0.1491		
futures	*Distribution*	*LTV*	*K*	*VoV*	*Correlation*	Simulated Attributes							*SSE*	*T*
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	4.28	0.750	−0.31	0.4139	−0.3291	4.8356	−0.2273	0.1643	0.0567	−0.1418	2.612	62.524
Model 3	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	62.660

Note: This table displays the empirical attributes (taken from Tables II and III) and the results of the models for the Nikkei 225 futures. These results are split into the three periods of analysis. On the left-hand side, the parameterization of each of the three models appears. Only the parameter values and the normal inverse Gaussian (NIG) distribution that has the lowest sum of squared errors (SSE) are presented. On the right-hand side, the simulated attributes of each model appear with the empirical attributes directly above for comparison. In the column immediately to the right of this, the SSE for each model appears. The final column indicates a *T*-test statistic comparing Models 2 and 3 to the geometric Brownian motion (GBM) assumption. This is computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV. LTV represents the long term volatility, θ in equations 2.2 and 2.3. K represents the rate of mean reversion, **K** in equations 2.2 and 2.3. VOV represents the volatility, ξ in equations 2.2 and 2.3. Corr (1–20) represents the average auto correlation of absolute returns for time lags 1 to 20 days. Corr (51–70) represents the average autocorrelation of absolute returns for time lags 51 to 70 days.

Only the parameter values and the NIG distribution that had 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. In the column immediately to the right of this, the SSE for each model appears. The final column indicates a T-test statistic comparing Models 2 and 3 to the GBM assumption. This was computed by taking the difference in the SSE for the models compared to the GBM case and dividing this by the simulated standard error of the SSE statistic found in Table IV.

As was expected, the GBM model with constant variance was rejected in favor of Models 2 and 3. The T statistics indicate a rejection at well above a 99% confidence interval. If the model with the lowest SSE is assumed optimal, of 12 data sets Model 3 (NIGSV) is the best in 9 and Model 2 (SV) is the best for the remaining 3. For all 12 data sets, correlated processes were indicated.

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

Comparisons between Models 2 and 3 can be made by examination of the T-test statistics in Tables VI through IX. As the T-test statistics provide an indication of the improvement in these models relative to the GBM case, taking the differences between these statistics provides insights into marginal improvement of Model 3 to Model 2. In 9 of the 12 cases, Model 3 had a higher T statistic. However, the differences were small. For all attributes except the excess kurtosis, both models performed equally well. What appears to be most critical is the inclusion of correlations between the underlying price and volatility processes. These capture the leverage and a portion of the skewness effect.

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. The next step is to examine the implications for options based on these assets. Because the parameter estimation of the stochastic volatility models relies on 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 preceding section, I demonstrated that in almost all instances, Model 3 (NIGSV) was optimal; therefore, only this was considered for the estimation of option values.[10] The choice of the appropriate input parameters for this model was taken from the previous section and was 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, for the model examined, stochastic volatility, correlated processes, and (negative) jumps have been introduced into the state space. Because no securities exist that allow the state space to be spanned when the volatility displays such dynamics, these models do not permit us (in the strictest sense) to use the risk-neutrality argument to price the options. This is an apparent theoretical inconsistency. However, the determination of the risk-neutral drift adjustment for the NIGSV process in Equation 7 does allow limited comparisons of the simulated options prices to be made actual option prices, which are evaluated under the risk-neutral measure. Although the equivalent risk-neutral measure in Equation 7 is certainly feasible, it may not be unique. In this research, I examined whether this measure is unique by first assuming that both simulated and actual option prices are evaluated under an equivalent risk-neutral measure. If this is the case, we can assess if this alternative price process alone is sufficient to explain the existence of implied volatility smiles. If this is not the case, this may suggest that the equivalent risk-neutral measure is not unique and markets are incomplete.

In this simulation, price series of three months in length were determined. Because the estimation of the unconditional (historical) dispersion processes was completed for trading days, options were also priced with trading time instead of calendar time. The assumed number of trading days in a year is 252. Option prices were estimated at time horizons from 1 week (5 trading days) to 3 months (in 5-day increments).

[10]Although for 3 of the 12 data sets the SV model was marginally superior to the NIGSV model, this difference was statistically significant.

Such options would correspond to typical terms to maturity of actively traded options on stock index futures.

To gain a better understanding of the impacts of the alternative models across strike prices, I examined 15 strike prices. The median strike price was centered 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 a European put, and when the strike price was above 100, the option evaluated was a European call. A nontrivial problem was the choice of strike prices so that as maturities of options varied, meaningful comparisons could be drawn. In previous articles 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 & Kazemi, 1998).

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

$$\frac{\ln(X_\tau/F_\tau)}{\sigma\sqrt{\tau/252}} \tag{9}$$

where X is the strike price of the option, F is the underlying futures price, and the square root of the time factor reflects the percentage in a trading year of the remaining time until the expiration of the option. 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 in respect to the current price of the underlying, the time to expiration, and the level of expected volatility. This adjustment converts all strike prices into a metric that can be interpreted as a standard deviation. Thus, in this analysis the strike price ranges ± 3.5 standard deviations away from the at-the-money level were examined in 0.5 standard deviation increments. This change in measure allowed a more direct comparison of model impacts on option prices where the time to maturity varied but the relative strike prices remained the same.[11]

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

For the Monte Carlo simulation, random numbers consistent with a GBM process were determined with a Box–Muller technique, and the antithetic approach suggested by Boyle (1977) was employed for both. This series was later used to determine the bivariate distribution used to estimate the stochastic volatilities. This series of random numbers was stored and used for all subsequent estimations of stochastic volatility. For each of the three NIG distributions, 10,000 were drawn with the method suggested by Rydberg (1997). The same approach was used for the estimation of the optimal stochastic volatility parameters in the previous section. These were also stored and used for all analysis that used that particular NIG distribution. Then, the bivariate distribution was estimated for the volatility series with Equation 7 and the optimal parameters of Model 3 for the stock index futures in Tables VI through IX (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 with an Euler approach (discrete form of Equations 3.1 and 3.2). With the prices of each model estimated, the payoffs of the 15 options (at each point in time) were determined, and the result was averaged. As interest rates were assumed to be zero, there was no need to discount the result to present value.

In parallel, I estimated the prices of all the options for each market with the Black (1976) model with the same strike prices as the simulation, the same term to expiration, and 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. Because the

[11]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.

underlying asset was a futures contract, the interest rate and dividend yield was set to zero (the same assumptions were made for the estimation of the Model 3 option prices).

Although standardization of the strike prices simplified comparisons between the four markets, the sheer amount of information made such comparisons cumbersome. Thus, only the results for a single market, the S&P 500, are presented. Furthermore, to simplify comparisons between simulated and actual implied volatility surfaces, the simulated option prices were expressed as implied volatilities. I further standardized these implied volatilities by indexing them to the constant volatility assumed in the Black (1976) model (dividing each of the implied volatilities by the assumed constant volatility and multiplying by 100). The indexed implied volatilities were then presented as a continuous surface relative to the standardized strike prices and time to expiration. This surface for the S&P 500 futures appears in the middle panel of Figure 1 and is titled *Simulated Implied Volatility Surface*. This figure has been scaled to allow direct comparison to actual implied volatility surfaces of options on the S&P 500 futures (which appears in the upper panel).

The simulated implied volatility surface for the S&P 500 appeared to display some of the features reported by Derman and Kani (1994) and Corrado and Su (1996). The characteristic curvature of volatility smiles was found, and there was some degree of negative skewness (especially for longer maturity options). Although I have not explicitly discussed the impacts of this model across a cross section of option prices, this figure implicitly displays these impacts. For all the markets, the Black (1976) pricing model overvalued options that were at-the-money and within a significant range around (and above) the at-the-money level. The shaded areas below 100 represent this in the graphs. For out-of-the-money options, the Black (1976) model tended to undervalue option prices (especially for options with lower strike prices), and the overpricing bias tended to increase the longer the term to expiration was. These results are consistent with the biases of stochastic volatility on option prices found elsewhere (Hull & White, 1988). Similar results were found for the other three stock index futures markets

In the top panel of Figure 1, the actual implied volatility surface associated with options on S&P 500 futures appears. To construct this surface, implied volatilities were estimated with the Black (1976) model for all (out-of-the-money) options on the S&P 500 futures. The period of analysis was from November 1990 to December 1998 (contemporaneous with the period of analysis of the S&P 500 futures). Analysis was restricted solely to options with the same terms to expiration used for the

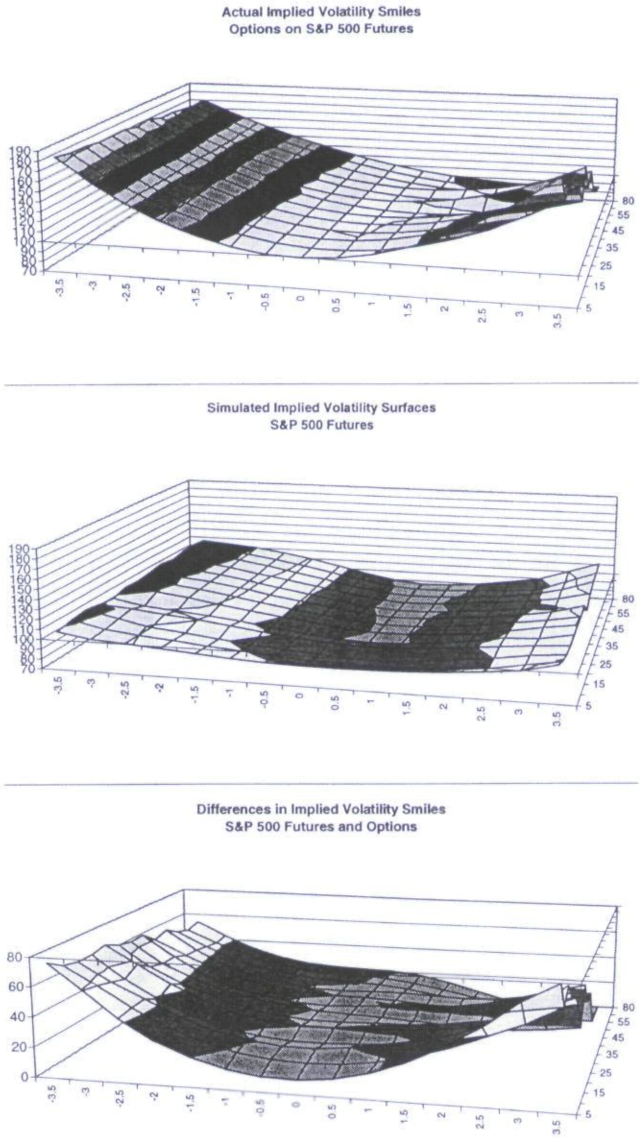

FIGURE 1

Implied volatility surfaces for Standard & Poor's (S&P) 500 futures and options: Actual, simulated, and differences. The top panel presents the actual implied volatility surface associated with options on S&P 500 futures for the period from November 1990 to December 1998. Implied volatilities and strike prices were standardized. The middle panel represents a simulated implied volatility surface consistent with the NIGSV model (see the text), with parameter values determined with the entire period of analysis (1990–1998). Implied volatilities were estimated with the Black (1976) model. The strike prices were expressed in standard deviation terms (away from the average simulated futures price), and the volatilities were indexed to the constant level of 20%. The bottom panel displays the differences between the actual and simulated implied volatility surfaces for the same relative strike prices and terms to expiration.

simulated option prices (5 days–3 months in 5-day increments) and excluded all options prices that allowed arbitrage (see Jackwerth & Rubinstein, 1996). The strike prices were then converted to the same standardized form as was done for the simulated implied volatility surfaces (with Equation 9), and the levels of implied volatility were indexed to the level of the at-the-money implied volatility. Then, the indexed implied volatilities were grouped by the same term to expiration. Finally, for each term to expiration, a polynomial fitting procedure (see Shimko, 1993; Tompkins, 1999) was implemented to yield an implied volatility smile. These were then combined, and the volatility surface was drawn.[12]

For the actual implied volatility surface associated with options on the S&P 500, it is clear that the degree of curvature was most extreme when the options were closest to expiration and became more linear with longer expirations. For the simulated surfaces, the effect was the opposite. Furthermore, the implied volatility surfaces for the actual options market were much more curved that for the simulated surfaces.

Another important difference is that the actual implied volatility surface displayed much greater asymmetry than the simulated surface. Options with strike prices below the current futures price had higher implied volatilities than options with higher strike prices. Similar results were found for the other three markets.

These differences are a consequence of the model chosen under the objective measure. Let us consider each facet of the NIGSV model separately. The NIG process for returns captures the effects of nonnormality in daily returns. Therefore, if subsequent disturbances are independent, this process will rapidly approach that of a normal distribution from the central limit theorem. If the higher moments of the chosen NIG process were sufficiently extreme, a more curved (and skewed) shape would be observed for options with short terms to expiration and would become more linear as the time horizon of the option were extended. For S&P 500 futures, the optimal NIG price process for the S&P 500 (NIG 1) was insufficiently nonnormal to produce the biases found with actual smiles. One interpretation of this result is that the higher moments from the risk-neutral process are different (more extreme) from those for the (risk-neutral adjusted) objective process.

The positive relationship between the degree of curvature in the simulated implied volatility surfaces and the time to expiration was most

[12]Tompkins (1999) used a model expanding Shimko (1993) to include higher moments and time/strike-price interactions. When implied volatility surfaces were standardized (relative to the ATM volatility), this model was able to explain 95% of the variance in the implied volatility surface of the S&P 500, and Tompkins (1999) demonstrated that this degree of explanatory power is retained outside of sample.

likely due to the inclusion of the stochastic volatility process. As pointed out by Hull and White (1988), the degree of bias in option pricing increases with the time to expiration of the option. This would produce an implied volatility surface that would display more curvature with a longer term to expiration. Because the volatility process in Equation 3.2 assumes a Gauss–Wiener process, this effect should be symmetrical. However, the simulated surface displayed an asymmetry relationship with the time to expiration. This asymmetry was due to effects of the negative correlation between the processes. Clearly, the longer the period was, the greater the impact of the negative correlation was. Similar results were found for the other three stock index options markets.

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

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

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

The simple difference between these two surfaces appears in the bottom panel of Figure 1. Standardization of both surfaces simplifies interpretation as the percentage difference in the (standardized) level of implied volatility. The actual implied volatilities were between 10 and 70% higher than those associated with the implied volatilities consistent with the objective price process. For the shortest time-to-expiration options, almost none of the curvature in the actual implied volatility surface was explained by the model of the objective process. Even so, the divergences between the surfaces were relatively symmetrical relative to the ATM (at the money) level for options with a short time period to expiration (out to 15 days to expiration). One possible interpretation of this is that the skew effect for options on stock index futures with short time periods to expiration might have been due primarily to an alternative price process (an NIG process with negative skewness). For longer terms to expiration, approximately one half of the curvature in the actual implied volatility surface was addressed by the objective price process model. However, for longer dated options, the divergence in the surfaces was asymmetric. The inclusion of correlated processes should address this. Even so (and with the use of an appropriate measure of the objective leverage effect), this appears insufficient to completely explain this effect. For options with higher strike prices and a longer term to expiration (35–90 days), the divergence between the two surfaces was small. However, the fact that everywhere the difference between the surfaces was positive could be interpreted as evidence for the existence of a risk premia. Alternatively, this difference could be associated with transaction costs or could be caused by other sources of market imperfections or frictions. For the other three stock index futures markets, the differences in the simulated and actual implied volatility surfaces displayed similar patterns, with the exception that for two of the other markets, the divergence seemed relatively time invariant (Nikkei and FTSE 100 futures).

SUMMARY AND SUGGESTIONS FOR FUTURE RESEARCH

This article examined the nature of the objective dispersion processes, which can be observed for futures contracts on four stock indexes. To capture the multifaceted nonnormality and interdependence of the empirical dispersion processes, seven attributes were identified. With these attributes for each of the four markets, two alternative stochastic volatility models were examined. I found that all four markets were best understood with a model that assumed the price process follows an NIG process and that the stochastic process driving volatilities was negatively correlated to this NIG process.

On the basis of an optimal parameterization of this model, European options were determined numerically for each market. Significant divergences from the Black (1976) values were observed. For all markets, this option pricing model tended to increase the prices of options with lower strike prices relative to Black (1976) prices and decrease the value of ATM and higher strike price options. The results of this research points out for the first time the relative biases associated with the option pricing model proposed by Barndorff-Nielsen and Shephard (1999). This extends the research of Hull and White (1988) to consider a richer class of stochastic volatility models.

Option prices consistent with the model were then expressed as standardized implied volatility surfaces and compared to actual implied volatility surfaces from options on the S&P 500. Although the simulated implied volatility surfaces display some stylized facts similar to actual surfaces, differences were observed that were similar across all four markets. For the entire surfaces, the actual option prices had higher implied volatilities than for option prices consistent with the model. Therefore, I rejected the hypothesis that the stochastic volatility models proposed were sufficient to explain the existence of implied volatility smiles. Of interest is that the differences between the implied volatility surfaces consistent with the objective and risk-neutral processes appeared to display similar dynamics for all four markets. The nature of this difference should provide a useful starting point for future research into the nature of the risk premium in options prices and should provide insights into the relationship between the objective and risk-neutral processes.

APPENDIX: DERIVATION OF THE RISK-NEUTRAL DRIFT ADJUSTMENT FOR AN NIG PROCESS

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

$$E\lfloor S_{t_j}|\Im_{t_{j-1}}\rfloor = S_{t_{j-1}},\ j = 1, 2, 3 \ldots n \qquad \text{(A1)}$$

with $\Im_{t_{j-1}}$ denoting the information at time t_{j-1}. Because 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

$$e^{a_{t_{j-1}}\Delta t} m(\sigma_{t_{j-1}} \sqrt{\Delta t}) = 1 \tag{A2}$$

where $m(x) = E[e^{xV_j}]$ is the moment generating function of the NIG (μ, δ, α, β) distribution. This function is

$$m(x) = \exp(\delta(\sqrt{\alpha^2 - \beta^2} - \sqrt{\alpha^2 - (\beta + x)^2}) + \mu x) \tag{A3}$$

So the appropriate choice of drift is

$$a_{t_{j-1}} = \frac{1}{\Delta t} \ln \left(\frac{1}{m(\sigma_{t_{j-1}} \sqrt{\Delta t})} \right) \tag{A4}$$

Combining equations A3 and A4 leads to

$$a_t = \frac{\delta}{\Delta t} \times [\sqrt{\alpha^2 - (\beta^2 + \sigma_{t-1} \cdot \sqrt{\Delta t})^2} - \sqrt{\alpha^2 - \beta^2}] - \frac{\mu \cdot \sigma_{t-1}}{\sqrt{\Delta t}} \tag{A5}$$

As a final note, adjusting the drift is only one way to obtain a martingale (for risk-neutral valuation). Given the models examining suggest markets are incomplete, 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 β from the observations of S and adjust β to obtain a risk-neutral valuation.

BIBLIOGRAPHY

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). Maximum likelihood estimation of the normal inverse Gaussian stochastic volatility model. Comprehensive Summaries of Uppsala Dissertations From the Faculty of Social Sciences, 82, 1–11.

Andersson, J. (1999b). 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.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. Journal of Finance, 52, 2003–2049.

Barndorff-Nielsen, O. E. (1997). Normal inverse Gaussian distributions and stochastic volatility modelling. Scandinavian Journal of Statistics, 24, 1–13.

Barndorff-Nielsen, O. E., & Shephard, N. (1999). Incorporation of a leverage effect in a stochastic volatility model. Manuscript in preparation, 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. (2000). Post-'87 crash fears in the S&P 500 futures options markets. Journal of Econometrics, 94, 181–238.

Black, F. (1975). Fact and fantasy in the use of options. Financial Analysts Journal, 31, 36–41, 61–72.

Black, F. (1976). The pricing of commodity contracts. Journal of Financial Economics, 3, 167–179.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81, 637–654.

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

Burghardt, G., & Lane, M. (1990). How to tell if options are cheap. Journal of Portfolio Management, 16, 72–78.

Christie, A. A. (1982). The stochastic behaviour of common stock variances. Journal of Financial Economics, 10, 407–432.

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.

Gallant, A. R., & Tauchen, G. (1996). Which moments to match? Econometric Theory, 12, 657–681.

Harvey, C. R., & Siddique, A. (1998). Conditional skewness in asset pricing tests. Manuscript in preparation, Duke University, Durham, NC.

Henker, T., & Kazemi, H. B. (1998). The impact of deviations from random walk in security prices on option prices. Manuscript in preparation, 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.

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

Hull, J., & White, A. (1987). Hedging the risks from writing foreign currency options. Journal of International Money and Finance, 6, 131–152.

Hull, J., & White, A. (1988). An analysis of the bias in option pricing caused by a stochastic volatility. Advances in Futures and Options Research, 3, 29–61.

Jackwerth, J. C., & Rubinstein, M. (1996). Recovering probability distributions from option prices. Journal of Finance, 51, 1611–1631.

Johnson, H., & Shanno, D. (1987). Option pricing when the variance is changing. Journal of Financial and Quantitative Analysis 22, 143–152.

Mayhew, S. (1995). Implied volatility. Financial Analysts Journal, 51(4), 8–20.

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

Merton, R. (1973). The theory of rational option pricing. Bell Journal of Economics, 4, 141–183.

Merton, R. (1976). Options pricing when underlying stock returns are discontinuous. Journal of Financial Economics, 3, 125–144.

Natenberg, S. (1994). Option volatility and pricing: Advanced trading strategies and techniques. Chicago: Probus.

Rydberg, T. H. (1997). The normal inverse Gaussian Lévy process: Simulation and approximation. Communications of Statistical and Stochastic Model, Stochastic Models, 13, 887–910.

Scott, L. O. (1987). Option pricing when the variance changes randomly: Theory, estimation and an application. Journal of Financial and Quantitative Analysis, 22, 419–438.

Shimko, D. (1993). Bounds of probability. RISK, 6(4), 33–37.

Stein, E. M., & Stein, J. C. (1991) Stock price distributions with stochastic volatility: An analytic approach. Review of Financial Studies, 4, 727–752.

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: Wiley.

Taylor, S. J. (1994). Modelling stochastic volatility: A review and comparative study. Mathematical Finance, 4, 183–204.

Theodossiou, P. (in press). Financial data and the skewed generalized T distribution. Management Science, 44 (12), 1650–1661.

Tompkins, R. G. (1997). Measuring equity volatilities in equity derivatives. London: RISK.

Tompkins, R. G. (1999). Implied volatility surfaces: Uncovering regularities for options of financial futures (FORC Pre-print 98/93). Coventry, England: University of Warwick, Financial Options Research Centre.

Wiggins, J. B. (1987). Option values under stochastic volatility: Theory and empirical estimates. Journal of Financial Economics, 19, 351–372.