FINANCIAL OPTIONS RESEARCH CENTRE

University of Warwick

Implied Volatility Surfaces: Uncovering Regularities for Options On Financial Futures

Robert G Tompkins

September 1998

Financial Options Research Centre
Warwick Business School
University of Warwick
Coventry
CV4 7AL
Phone: (01203) 524118

FORC Preprint: 98/93

IMPLIED VOLATILITY SURFACES: UNCOVERING REGULARITIES FOR OPTIONS ON FINANCIAL FUTURES

Robert G. Tompkins
Visiting Professor
Department of Finance
Vienna University of Technology* and
Permanent Visiting Professor
Department of Finance
Institute for Advanced Studies#

ABSTRACT

While it is now generally accepted that implied volatilities of European options differ across strike prices and time, what has not been examined in the literature is the characteristics of the strike price biases between different assets and asset classes and the variability of these surfaces over time. This research examines twelve options markets on financial futures (comprising three asset classes) and compares the strike price biases both for the same markets and across all markets.

When implied volatility surfaces are standardised, their patterns are smooth, well behaved and display the asymmetrical and convex patterns previously described in the literature. Furthermore, there appears to be consistency in how the implied volatility surfaces evolve over time for individual markets and we uncover consistencies in the behaviours of implied volatility surfaces between markets. To test the significance of these empirical regularities, we develop a model based upon a polynomial expansion across strike price and time. The polynomial expansion approach allows us to examine each effect separately and include time dependent interactions. Coefficients from an OLS regression model allow direct comparisons.

Cross-sectional comparisons suggest that the first order strike price biases (skewness) differs between asset classes. However, the first order strike price biases display similar time dependencies within the same asset class. Regarding the second order strike price effect (kurtosis), consistencies exist between the three financial asset classes examined. These models are stable overtime and lose little explanatory power outside of sample. This suggests the paradigm used by option participants to adjust option prices away from Black-Scholes-Merton prices is consistent over time.

JEL classifications: G13

Keywords: Implied Volatility Surfaces, Volatility Smiles, Shocks, Risk Neutral Processes, Skewness, Kurtosis, Heterokurtosis

---

* Floragasse 7/4, A-1040 Wien, Austria, Phone: +43-1-726-0919, Fax: +43-1-729-6753, Email: rtompkins@ins.at

Stumpergasse 56, A-1060 Wien, Austria, Phone: +43-1-599-91-125,Fax: +43-1-597-0635 Email: rtompkins@ihs.ac.at

1. INTRODUCTION

If the Black-Scholes-Merton model [Black and Scholes (1973) and Merton (1973)] accurately describes conditions in actual options markets, then the volatilities implied from the market price for options on the same underlying asset would be constant regardless of the strike price of the option or its maturity. However, it is now generally accepted that the implied volatilities differ across strike prices for the same maturity and across diverse expiration periods.[1]

A recent trend in the empirical investigation of implied volatilities has been to concentrate on understanding the behaviour of implied volatilities across strike prices and time to expiration [see Jackwerth and Rubinstein (1996) and Dumas, Fleming and Whaley (1996)]. This line of research assumes implicitly that these divergences provide information about the dynamics of the options markets. Another approach [Dupire (1992, 1994), Derman & Kani (1994) and Rubinstein (1994)] suggests that the divergences of implied volatilities across strike prices may be providing information about the expected dispersion process for underlying asset prices.

A separate line of research that has not yet been examined is whether these implied volatility surfaces display consistencies for individual option markets and whether consistencies in the implied volatility surfaces exist between markets. This research will deal with these issues and do so for twelve financial markets drawn from three asset classes: stock indices, fixed income and foreign exchange. Furthermore, this analysis will be done for a longer period of analysis than has previously been examined in the literature.

---

[1] The existence of the strike price effect has been pointed out extensively in the literature. Early examples include: Black (1975), MacBeth and Merville (1979), Galai (1983,1987) and Rubinstein (1985). Recent examples include: Xu and Taylor (1993) and Heynen, Kemna and Vorst (1994).

In this paper, we will demonstrate that if the implied volatilities are standardised, it is possible to compare their behaviours over time. Furthermore, there appears to be consistency in how the implied volatility surfaces evolve over time for individual markets. We find that the implied volatility surfaces for all markets examined are smooth and well behaved. This fact allows us to capture most of the dynamics of the surfaces using a polynomial expansion. Statistical testing of models based upon this polynomial expansion suggests that the dynamics of implied volatility surfaces are stable over time. Furthermore, the coefficients of the models allow comparisons of implied volatility surface dynamics between markets.

The paper is organised as follows. The second section will discuss the data sources and the steps required preparing the data for further analysis. The third section discusses the process of standardisation of implied volatilities and strike prices. In section four, the results are presented for the twelve markets under examination. The final sections of this paper (five and six) will assess the extent to which the implied volatility structures for the twelve markets under investigation can be captured using a simple polynomial functional form. This will be done using an ANCOVA approach with the coefficients of the statistical model quantifying the strike price effects and facilitating comparisons between markets. The final section provides conclusions and suggested areas for further research.

2. DATA SOURCES

The options examined in this research are options on futures[2]. The underlying markets include four fixed income futures contracts: US Treasury Bond Futures, UK Gilt Futures, German Bundesanleihen Futures (Bunds), and Italian Government Bond

[2] For the DAX options and FTSE options, these were actually on the cash index. However, these products were European style options expiring on the same day as the Futures for these markets. Thus, these options can be considered as de facto options on futures.

Futures (BTPs), four equity index futures: S&P 500 Futures, FTSE 100 Futures, DAX Index Futures, and Nikkei Dow Futures, and four currency futures: US Dollar/Deutsche Mark, US Dollar/British Pound, US Dollar/ Japanese Yen and US Dollar/Swiss Franc.

For all of the option markets, the analysis period extends back in time to either the introduction of these contracts or to include all the publicly available data. For the following underlying options, the following time periods of analysis were examined:

Underlying Asset	Time Period of Analysis
Stock Index Options	
S&P 500 Futures	25/03/1986 - 24/12/1996
FTSE Futures	02/01/1985 - 20/12/1996
Nikkei Dow Futures	25/09/1990 - 16/12/1996
DAX Futures	02/01/1992 - 20/12/1996
Fixed Income Options	
Bund Futures	20/04/1989 - 21/11/1996
BTP Futures	11/10/1991 - 21/11/1996
Gilt Futures	13/03/1986 - 22/11/1996
US T-Bond Futures	02/01/1985 - 15/11/1996
Currency Options	
Deutsche Mark /US Dollar	03/01/1985 - 09/12/1996
British Pound / US Dollar	25/02/1985 - 09/12/1996
Japanese Yen / US Dollar	05/03/1986 - 09/12/1996
Swiss Franc / US Dollar	25/02/1985 - 09/12/1996

Since our aim is to examine the characteristics of implied volatility surfaces over time, it is critical that the sample period of analysis extends as far back in time as is possible. In most cases more than 10 years was available, which will allow us enough observations to conduct meaningful analysis.[3]

The data for the options and futures contracts was obtained from the exchange where these trade. The London International Financial Futures Exchange (LIFFE)

[3] In total, the number of option prices examined for all twelve markets was 1,263.317. Given that we also had the underlying futures prices for the same dates (and at the same time) as the options, we were able to assure that both time series were consistent to each other. From this analysis, we were able to clean both series and assure our analysis was minimally impacted by errors in data.

provided information on the BTPs, Bunds, Gilts and the FTSE 100[4]. The Chicago Board of Trade provided data on the US T-Bond Futures and Options. The Deutsche Terminbörse (DTB) provided data on the DAX[5] futures and options and the Chicago Mercantile Exchange provided the future and options data for the S&P 500, Nikkei Dow, Deutsche Mark, British Pound, Swiss Franc and Japanese Yen. For most of these markets, the data obtained included the closing prices of the options, the strike prices, the price of the underlying futures and in some instances, the implied volatility of the options estimated by the exchange.

Eight of the options under examination in this research were American style options on Futures. These included the FTSE 100 (prior to 1992), S&P 500 and Nikkei 225, the US T-Bond, and all the currency options. These options are paid for up-front and are American style. To estimate the implied volatilities correctly, we chose to use the Barone-Adesi and Whaley (1987) model. According to Clewlow and Xu (1994), this model has been shown to be very accurate for option maturities shorter than twelve months. As we will restrict our research to only those options with three months or less to expiration, this model will be suitable for our purposes. For this model, an interest rate parameter must be included in order to estimate the implied volatilities. For all US Dollar based options, we used the US Treasury Bill interest rate for that day whose maturity fell most closely to the actual expiration date of the options. This data was obtained directly from the Federal Reserve Bank in

---

[4] For the FTSE 100, options data was only available from 1992 from the LIFFE and only for the European options. To extend the analysis, we obtained additional data from Professor Gordon Gemmill, of City University Business School, who had compiled data for American options on the FTSE 100 (from the financial press) from 1985 to 1992.

[5] It included all the tick by tick prices of the futures and options contracts during each trading day. Given that the rest of the analysis was done on closing prices, we sorted the data by time of trade and chose the options and futures prices within one hour of the close. For this, we would only select options for analysis if a futures trade occurred within 3 minutes of the option trade, otherwise, we ignored the option. In most instances, the options selected were within the last 30 minutes of the close and the accompanying underlying futures contract traded within 1 minute of the option.

Washington D.C. These contracts that required the US Dollar interest rates included all the currencies, the US T-Bond, the S&P 500 and the Nikkei 225 contracts.

The FTSE 100 option (after 1992) was a European style option with stock type settlement (this means the premium was paid up-front). Since the options expire on the same date as the FTSE 100 futures contract, these options can be considered de facto options on futures and the implied volatilities can be determined using the Black (1976) model. Most of the remaining options under examination in this research were American style options on Futures. All these options are traded at the LIFFE. These include the Bund, BTP and Gilt options. For these options the mechanics of margining of both the underlying futures and options removed the possibility of early exercise. For these options, it is possible to estimate the implied volatilities using the Black (1976) model and interest rates can be ignored.

These approaches addressed all the markets under examination except the DAX index options. For this contract, the option was based upon the cash index. This was selected because the option on the DAX futures is extremely illiquid and the date of settlement is exactly the same as the expiration of the futures. Thus, this contract is also de facto an option on the futures. This option contract is European style and has stock type settlement. This means that the appropriate model is the Black and Scholes (1973) model with an interest rate input. No dividend yield is required because the DAX index is a total return index where the dividends are assumed to be automatically reinvested in the index. The interest rate input was the average weekly 3-month LIBOR for Deutsche Marks obtained from The Bank of England.

Finally, given that this research is empirical in nature it was of utmost importance that the data we examined was carefully screened to remove errors. This was achieved in a number of ways. Firstly, we compared the futures price series with the options price series for the same days to identify obvious errors in recording either price series. This comparison was achieved by comparing the put-call parity values of the options with the underlying futures prices for every single date in our database (and for all twelve markets). A screening procedure was imposed such that if futures or options prices diverged by more than the normal bid/offer spread (of one tick), the observations were flagged. Once this was done, each price was compared with the original daily price sheets to confirm if a 'keypunch' error had occurred. We discovered that only 1-2% of the data had such errors. Nevertheless, these errors were of a sufficient magnitude that they did influence the results and therefore required correction.

To avoid spurious results from option prices that appear to admit arbitrage, we eliminated all options prices (and the accompanying implied volatilities) that were traded at the minimum level at the relevant market or allowed a butterfly arbitrage. Jackwerth and Rubinstein (1996) also used this approach. Furthermore, to address the potential problem of nonsynchronous prices for the options and underlying futures, only those implied volatilities from the available out-of-the-money option contracts (not admitting arbitrage) were examined. Thus, we only examined put options with strike prices below the current futures price and call options above.

Finally, a non-trivial issue is the term to maturity of the options selected for analysis. We observed from the collection of the option and futures prices that the liquidity for both assets is not uniform over time. While options and futures could be offered with up to nine months to expiration, these contracts rarely traded until they became the nearest contracts to expiration. Furthermore, for the options markets, we must have enough strike prices to be able to fit a meaningful polynomial model to ascertain the shapes. Therefore, we selected both futures and options contracts that were the nearest contracts to expiration and traded on the quarterly expiration schedule of March, June, September and December maturities. This assured that we had prices that represented actual trading by market participants. By restricting our analysis solely to the nearby options and only on the same quarterly trading cycle as the underlying futures, the options we examined tended to have a maximum maturity of approximately 90 days to expiration.

3. STANDARDISED IMPLIED VOLATILITY SURFACES: METHODOLOGY

While a number of approaches have been proposed to standardise the implied volatilities, the simplest approach is to create an index where the implied volatilities at each strike price are expressed as the ratio to the implied volatility of the option closest to the at-the-money level. Fung and Hsieh (1991), Tompkins (1994) and Natenberg (1994) have all used this approach. All these approaches take the simple ratio between the implied volatility at each strike price divided by the ATM implied volatility and multiply the result by 100 (or express the result in percentages). This transformation is required because the levels of volatility are not constant over time. The logic behind this approach is that the relative relationships between the volatilities and not the absolute levels are of interest.

The strike prices must also be standardised to allow comparisons to be drawn. A simple approach suggested by Tompkins (1994) was to take the ratio of the strike price to the underlying price. Jackwerth and Rubinstein (1996) used the same approach. A similar approach was used by Fung and Hsieh (1991), the difference being that they inverted this ratio. While this has practical advantages for market participants (namely it is simple to reverse the equation to obtain actual strike prices), the approach is somewhat misleading as the ratios are time-independent, while options prices are not. A better approach, which is more consistent with the time-dependency of option prices and is also consistent with the assumptions of the Black-Scholes-Merton option pricing models is to use the approach suggested of Natenberg (1994). In this approach the strike price is expressed as the natural logarithm of the ratio of the strike price X of the option relative to the underlying futures price F, divided by the square root of the time remaining until the expiration of the option. This will express the strike prices as the relative percentage movement in the underlying futures required to reach an exercise price. We have modified this approach by modifying the denominator of the formula by multiplied the square root of time by the level of the at-the-money implied volatility. This will be expressed as:

$$\frac{\ln(X_\tau / F_\tau)}{\sigma\sqrt{\tau / 365}} \qquad \textbf{(1)}$$

where X is the strike price of the option, F is the underlying futures price and the square root of time factor reflects the percentage in a year of the remaining time until the expiration of the option. The sigma (σ) is the at-the-money volatility. For this analysis, we will assume that the relevant time is calendar days and will express time as the percentage of calendar days remaining in the options life to the total trading time in a year (which we assume is 365 days).

This transformation is consistent with the assumptions of the B-S-M model, where the movement in the underlying asset is measured on a logarithmic scale. From the first term in the Black-Scholes model [N(d1)], the relationship between the exercise price of an option and the current underlying price is expressed as the logarithm of the exercise price divided by the underlying price. At the same time, GBM assumes that movement over time is governed by a square root relationship, so that in the Black-Scholes model, the relative amount of movement (for the underlying asset) to reach a strike price is fully expressed by the above formula 1. Finally, the inclusion of the at-the-money volatility will allow us to express the strike prices in standard deviation terms. This will allow us to compare the smile relationships between and within markets more directly and consistently

4. STANDARDISED IMPLIED VOLATILITY SURFACES: RESULTS

To gain an overview of the smile estimation process and the implications for this research, we determined the implied volatilities for the twelve markets under investigation on a single date: May 7, 1996. Inverting the appropriate option-pricing model for each market (including the other known variables) and solving for the free parameter yielded the implied volatility. These implied volatility point estimates were connected by linear interpolation to draw implied volatility patterns. These can be seen in Figures 1, 2 and 3 for the stock index, fixed income and currency options.

For the options on the stock index futures (in Figure 1), the implied volatility patterns generally display a convex shape and are skewed to the left. This implies that the market prices of options with lower strike prices are higher than the theoretical prices obtained from the option-pricing model using the at-the-money volatility. In addition, the market prices of options with higher strike prices are lower than what would be predicted based on the assumptions of the option-pricing model. This result has been identified extensively in the literature[6]. While for the S&P 500, FTSE 100 and DAX options the smile relationships are skewed to the lower strike prices, the Nikkei displays a relatively more symmetrical shape for the September 1996 maturity. Thus, all the stock index options display the skew behaviour previously identified in the literature. Regarding the curvature of the implied volatility patterns, it is now clear why this pattern has been referred to as a 'smile'.

For the options on the fixed income futures (in Figure 2), we also observe both skewness and curvature effects. Unfortunately, we did not have multiple option maturities to examine the effects of time on the implied volatility patterns. The only market that allows us some comparison is the Gilt with two available maturities. This is due to the fact that there is an even greater concentration of trading activity on the options that are closest to maturity. Nevertheless, it appears from the Gilt options that a similar pattern exists for those observed for the stock index options. The curvature of the implied volatility patterns becomes more extreme the closer we are to expiration of the option.

For the options on the foreign exchange futures (in Figure 3) we fortunately have more option maturities to examine. This will allow us to gain some insights in the effects of time on the implied volatility patterns. For all four markets, there does not appear to be a systematic skewed relationship between strike prices and implied volatilities. However, one can clearly see that as the options are closer to maturity, there is much greater curvature. Furthermore, the levels of the implied volatilities differ significantly across maturities.

Apart from the obvious conclusion that the smile patterns are not flat, this method of presentation makes comparison between markets difficult. The scales for both the strike prices and implied volatilities vary widely between the markets. Comparison is easier if both the implied volatilities and strike prices are standardised such that both the strike price and implied volatilities are expressed in the same metric.

To accomplish this, we converted the levels of the implied volatilities into index form. The denominator (in the ratio) is the at-the-money volatility. Given that it

---

[6] See references in Footnote 5.

is unlikely that the underlying futures price will be exactly equal to an available strike price, we were required to estimate the level of the at-the-money volatility. After examining various approaches to estimate this, we found that a simple linear interpolation for the two implied volatilities of the strike prices that bracketed the underlying asset price (one below and one above) had much less errors than other methods suggested in the literature[7]. This transformation means that all the implied volatilities are now expressed as an index with the numeraire equal to the at-the-money implied volatility. The second standardisation was to index the strike prices to the level of the underlying futures price. This standardisation was achieved using formula 1. This allows all strike prices to be transformed into a standard deviation measure allowing consistent comparison both overtime and cross-sectionally.

By standardising these two variables, a metric has now been constructed that allows comparisons not only between markets but also for individual markets over time. It is acknowledged that whenever some method of standardisation is employed, a loss of information (detail) results. However, given our objective is to compare *relative* smile behaviours both cross-sectionally and across time, we believe the loss of information by standardising is more that made up by the ability to compare smile dynamics within and between markets more directly. One potential problem is that the level of the expected variance (as measured by the level of the at-the-money implied volatility) is lost and might be important to the dynamics of the implied volatility surfaces. This will be examined in the final portion of this paper.

Given that we restricted our analysis to the quarterly expiration schedule of March, June, September and December maturities, we have only examined implied

[7] The first approach used was to determine the quadratic functional form that fits the volatility smile. This used a quadratic approach suggested by Shimko [see Shimko (1991,1993)]. We found two major problems with this approach. The first is that for many days, we had barely enough degrees of freedom

volatility surfaces with a maximum term to expiration of approximately 90 days. We further pruned the data required for this analysis by restricting our analysis of the implied volatility surfaces to eighteen time points from (the date nearest to) 90 calendar days to expiration to (the date nearest) 5 calendar days to expiration in 5-day increments.

Finally, given the standardisation of the strike prices implied that they are expressed in standard deviation terms, we had to choose a reasonable range including the available options. From casual observation, we found that the options we examined would not be excluded from the analysis over a range that was 4.5 standard deviations away from the underlying asset price.

Using this approach, all the implied volatility surfaces were examined for the twelve markets and covered the entire period of analysis. While it would be possible to present each of the standardised smile patterns for each expiration cycle for each of the twelve markets (and forty eight expiration periods), this would only allow for a qualitative comparison. Our interest is to assess the general results of the standardised implied volatility surfaces to provide insights into the form of the model that would capture the general mannerisms.

To this end, an index of the implied volatility surfaces was constructed for each market under investigation. This was achieved by pooling all the standardised implied volatilities with the same period to expiration for each market. With these (time to expiration) homogenous samples, the following regression was run:

$$VSI = \alpha + \beta_1 \cdot \frac{\ln(X_\tau / F_\tau)}{\sigma\sqrt{\tau/365}} + \beta_2 \cdot \left[\frac{\ln(X_\tau / F_\tau)}{\sigma\sqrt{\tau/365}}\right]^2 + \varepsilon \qquad (2)$$

(options prices) to determine the quadratic form. Secondly, many of our markets (the US T-Bond market in particular) were not well described by a quadratic function.

The results of this regression were repeated for each of the eighteen expiration periods and examined for all twelve options markets. With the coefficients from this quadratic regression a fitted line was produced for each expiration period and for each market. These results are graphed in Figures 4 for the four stock index options, 5 for the four fixed income options and 6 for the four foreign exchange options.[8]

In Figure 4, one can see that some consistent relationships seem to exist among the four stock index options markets. The S&P 500 and the DAX options display almost identical patterns with similar dynamics over time. Both markets start with 90 days to expiration with a relatively linear left skewed shape. This shape flattens somewhat as expiration is approached and the patterns display progressively greater curvature. One can see that the FTSE and the Nikkei also display similar dynamics. The implied volatility patterns are more convex at 90 days compared to the other stock index options and this pattern becomes more convex as expiration is approached.

In Figure 5, there are also consistencies among the four fixed income options markets. The Bund, BTP and Gilt options display similar patterns over time. In some ways, these shapes are reminiscent of those observed for the FTSE and Nikkei surfaces. All markets start with 90 days to expiration with a relatively linear left skewed shape with a slightly convex shape. As expiration is approached, the skew flattens and the convexity increases. The odd market out is the US T-Bond, where there is little evidence of a skew but the convexity can clearly be seen as increasing as expiration approaches.

---

[8] Later in this paper, we will demonstrate that to correctly understand the characteristics of implied volatility surfaces a simple quadratic model of this form is inadequate. However, this goal here is to generate implied volatility surfaces which will provide qualitative insights into the nature of the complete model that will be developed later.

In Figure 6, there appears to be even greater consistency among the four foreign exchange options markets. All four markets display almost identical patterns with similar dynamics over time. In many ways, there appears to be a constant shape across time for all the four markets. The only time dependent variation that seems to occur is that the convexity becomes more extreme as the time to expiration is reduced.

These graphical presentations suggest that smile patterns both for individual markets, within asset classes and indeed among all markets display some similar characteristics. What remains to be examined is the degree of similarity that exists and whether these regularities are stable over time.

5. STATISTICAL COMPARISONS OF IMPLIED VOLATILITY SURFACES

Our approach is related to the work by Dumas, Fleming and Whaley (1996), who tested a number of arbitrary models based upon a polynomial expansion across strike price and time. They found that the model that best explained the dynamics of the implied volatility surface was an expansion to degree two. However, they rejected the existence of a purely deterministic functional form of the implied volatility surface.

In this research, we will extend the polynomial expansion to degree three and include additional factors, which may influence the behaviours of volatility surfaces. We will also test the finding of Rubinstein (1994) that the overall shapes of volatility surfaces changed after the 1987 stock market crash. Furthermore, we will examine how other individual market specific shocks may affect the shapes of implied volatility surfaces for these markets.

Consider a Taylor's series expansion to degree three.

Expanding the function $\sigma = f(x,t)$ with Taylor's expansion series

$$\sigma = \sum_{i,j=0}^{\infty} \frac{1}{(i+j)!} \frac{\partial^{(i+j)}\sigma}{\partial x^i \partial t^j} x^i t^j \qquad \textbf{(3.1)}$$

and stopping the expansion at the third degree ($i + j \leq 3$), we obtain

$$\begin{aligned} \sigma \approx & \frac{\partial \sigma}{\partial x} x + \frac{\partial \sigma}{\partial t} t + \frac{1}{2!} \frac{\partial^2 \sigma}{\partial x^2} x^2 + \frac{1}{2!} \frac{\partial^2 \sigma}{\partial t^2} t^2 + \frac{1}{2!} \frac{\partial^2 \sigma}{\partial x \partial t} xt + \\ & + \frac{1}{3!} \frac{\partial^3 \sigma}{\partial x^3} x^3 + \frac{1}{3!} \frac{\partial^3 \sigma}{\partial t^3} t^3 + \frac{1}{3!} \frac{\partial^3 \sigma}{\partial x^2 \partial t} x^2 t + \frac{1}{3!} \frac{\partial^3 \sigma}{\partial x \partial t^2} xt^2 \end{aligned} \qquad \textbf{(3.2)}$$

Given that we have nine derivatives in the expansion, we will construct nine variables to capture these effects.

To capture the 1987 crash effect pointed out by Rubinstein (1994), we constructed a dummy variable, which assumes a value of 0 prior to the crash and 1 thereafter. To assess the impacts on the strike price effect, this dummy variable will be multiplied by the first and second order strike price variables from equation 3.2. Similar dummy variables were constructed to capture idiosyncratic shocks for individual markets. To determine when such a shock had occurred, we examined the exponentially weighted unconditional volatility time series for the twelve markets of interest. The two most extreme spikes in the unconditional volatility that occurred over the period of analysis were chosen as shock events.

Finally, there is concern that important information has been removed from the analysis by the process of standardising the strike prices and implied volatilities. Thus, if the strike price bias is a function of the level of the implied volatility, this can be tested by including the level of the at-the-money implied volatility in the model. Once again, since our objective is to understand the strike price effects, combination variables will be estimated which are the products of the first and second order strike price effects in equation 3.2 with the level of the at-the-money implied volatility.

The final form of the model can be expressed as:

$$VSI = \alpha + STRIKE \cdot (\beta_1 + \beta_2 \cdot TIME + \beta_3 \cdot TIME^2 + \beta_4 \cdot CRASH + \beta_5 \cdot SHOCK1 + \beta_6 \cdot SHOCK2 + \beta_7 \cdot ATMVOL)$$
$$+ STRIKE^2 \cdot (\beta_8 + \beta_9 \cdot TIME + \beta_{10} \cdot TIME^2 + \beta_{11} \cdot CRASH + \beta_{12} \cdot SHOCK1 + \beta_{13} \cdot SHOCK2 + \beta_{14} \cdot ATMVOL)$$
$$+ \beta_{15} \cdot STRIKE^3 + \beta_{16} \cdot CRASH + \beta_{17} \cdot SHOCK1 + \beta_{18} \cdot SHOCK2 + \beta_{19} \cdot ATMVOL + \beta_{20} \cdot TIME + \beta_{21} \cdot TIME^2$$
$$+ \beta_{22} \cdot TIME^3 + \varepsilon$$

(4)

Equation 4 differs from equation 3.2 in a number of significant ways. The first is that the implied volatility is now the standardised indexed form rather than the absolute level. Secondly, additional variables have been added to test the impact of market shocks on the implied volatility surface and finally, the level of the at-the-money volatility has been included to assess how much information has been lost by the standardisation procedure.

Given the model is a mixture of normal variables and dummy variables, an analysis of covariance (ANCOVA) was utilised with the standardised implied volatilities (VSI) as the independent variable. The computer programme used for the analysis was STATISTICA for Windows (version 5.0). [See STATISTICA (1995)].

One potential criticism of this approach is that this has led a fairly complex model with a large number of variables. Normally when evaluates an equation with so many independent variables, there are too many parameters to identify satisfactorily. We are fortunate here that the number of observations is extraordinarily high. In addition, the results have high degrees of explanatory power (adjusted R squared). We will demonstrate that these models retain this high degree of explanatory power regardless of how the simple OLS approach is corrected for potential biases in the regression and also remains high (and has consistent coefficients for the independent variables of interest) for different periods of analysis. Therefore, for many regression models over-fitting is endemic when the number of independent variables is large relative to the number of observations or the models fail to predict outside of sample. However, in this case, the use of our enormous database and the fact that the models are stable for different periods of analysis suggest that this approach is not over-fitting the data series.

6. RESULTS OF THE STATISTICAL ANALYSIS

Prior to the presentation of the results of the analysis, we will provide an economic interpretation of the variables. The first order strike price variables can be interpreted roughly as measuring the third moment of a statistical distribution (skewness). The Beta coefficients from β_1 to β_7 capture these dynamics. If a skew effect does exist and it is not time varying, only the first Beta would be significant. If it is time varying, then β_2 and β_3 should be significant. If any of the shocks caused a change in the nature of the skew, then β_4, β_5 and β_6 will capture these effects. Finally, the coefficient of β_7 will indicate an interaction between the expected level of future variance (at-the-money implied volatility) and the first order strike price effect.

The second order strike price effect can be seen as related to the fourth moment of the distribution (kurtosis). These dynamics can be understood by examining the Beta coefficients from β_8 to β_{14}. Similarly to the analysis of the first order strike price effects, if excess kurtosis effects exist and are not time varying, only β_8 would be significant. If these effects are time varying, then β_9 and β_{10} would indicate this. If any of the shocks caused a change in the nature of the implied kurtosis, then this will be indicated by significant coefficients for β_{11}, β_{12} and β_{13}. Finally, β_{14} will allow us to examine the interaction between the expected level of future variance (at-the-money implied volatility) and the implied kurtosis.

The remaining variables exist for two primary purposes. The first is that the previous variables will provide insights into the average slope relationships that exist for the strike price effects. These variables will allow us to understand the dynamics of the process. The second reason is that for a number of these variables we have a prior expectation that they should be insignificant in the regression. These would include the dummy variables that identify the occurrence of the 1987 stock market crash and the individual market shocks. We would also expect that the level of the at-the-money implied volatility would not be significant. This expectation is due to the standardisation of the implied volatilities (indexing them to 100). If these variables are significant, this may serve as a diagnostic for potential errors in variables. Furthermore, the degree of significance will provide a guide to the severity of the errors in variables problems.

It may very well be that further moments are required to understand the implied dispersion process beyond skewness and kurtosis. For this reason, β_{15} will indicate if this is the case or not ($STRIKE^3$). We will choose to refer to this effect as hetero-kurtosis.

The first approach was to run an ordinary least squares (OLS) step-wise regression with dummy variables (ANCOVA). [9] This was done separately for all twelve markets and used all the available data. The results of these statistical procedures can be seen in Tables 1, 2 and 3 for the three asset classes, stock index futures, fixed income futures and foreign exchange futures.

[9] This is to be expected as a number of the variables are highly correlated by design. For example, many of the variables are products of other included variables and are by design highly correlated.

In these tables, the coefficients of the regression are presented along with the standard error of the estimates and the t-statistic.[10] For all variables that have a significant t-statistic (at a 95% level), the results are presented in **bold**. For all results that are in normal text, these were not significantly different from zero for the independent variables or 100 for the intercept. In the instance that the variable was not selected in the forward stepwise regression, this is represented by "-.--". We have also included the number of observations included in the analysis, the adjusted R-squared statistic and the Durbin-Watson statistic to measure possible problems with serial correlations in the residuals.

In Table 1, we find that almost all of the included independent variables are statistically significant for the four stock index options. Exceptions include the variables that include SHOCK1 for the S&P and FTSE and variables that include CRASH for the DAX and Nikkei. This is hardly surprising since these variables were not appropriate for these markets. For the S&P and FTSE, the CRASH and SHOCK1 were the same event. For the DAX and Nikkei, the available observations were only available after the CRASH.

The extremely high explanatory power of each of the models is somewhat surprising. The adjusted R-squared statistic is between 0.9084 (for the Nikkei) to 0.9573 (for the S&P). These results suggest that the models are explaining almost all the variance in the implied volatility surfaces. Nevertheless, we must be careful with interpreting these results since the regressions might be biased and therefore the statistics could be misleading. For example, the Durbin-Watson statistics suggest that (for three of the four markets) some of the variance explained by the regressions is

[10] The t-statistics for all the independent variables indicate whether the coefficient is statistically significantly different than zero. For the intercept, the t-statistic indicates whether the coefficient (alpha) is statistically significantly different than 100.

due to serial correlations in the residuals. However, when the regressions were corrected for serial correlations in the residuals using a weighted least squares approach, the differences in the Beta coefficients for the independent variables was slight (and in most cases insignificant).[11] Given that this is the case, we shall interpret the regression results in Tables 1, 2 and 3 knowing that alternative regression approaches correcting for violations in regression assumptions will not substantially alter our conclusions.

At this point, we will examine in detail the results of the regressions for the stock index options markets and provide an economic interpretation. For the first order strike price effects, we find that most of the first order strike price related variables are consistently of the same sign and relative magnitude. However, the pure strike price effect varies between the four markets. For the S&P, the coefficient is *positive*. This suggests that controlling for all the other variables, the skew for this market is positive. However, it is clear that the negative sign and magnitude of the combination of the first order strike price effect with the crash is causing an overall negative skew for the S&P. For the FTSE, DAX and Nikkei, the pure first order strike price effect is negative. For the DAX and Nikkei, this is hardly surprising given that both of these markets only had observations after the 1987 and 1989 stock market crashes. This suggests that the crashes had impacts for all stock index options markets. For the options markets observed both before and after these events the effect is found in the coefficients for the STRIKE*CRASH variable (β_4) and for the STRIKE variable (β_1) for those markets only observed after the crash.

[11] Corrections were also done for Heteroscedasticity, Omission of Key Variables and the issue of Multicollinearity were also examined. Furthermore, the regressions were run for separate samples of the data to assess stability. All these tests demonstrated that the regression results remained similar and that the overall model was robust.

Caution must be exercised in the interpretation of the model; the overall strike price effects are an aggregate of a number of variables. To gauge the overall first order strike price effect, one must compare the STRIKE related variables for all markets and include the STRIKE related variables associated with the dummy variables for each market.

Apart from this difference, the interaction variables that combine STRIKE and TIME have similar coefficients among all stock index options (both first order and second order effect). The interpretation of these variables is that as the expiration of the option is approached the negative skew begins to flatten. In some ways, it is difficult to interpret the coefficients of these two variables between the markets. This is because they are a combined effect. Nevertheless, the signs and magnitudes of the effects are of a similar dimension.

Regarding the second order strike price effect, again many of the coefficients for the stock index options (and the aggregated file) share the same sign and similar magnitudes. For all the markets, the coefficient of the Beta for the pure kurtosis effect ($STRIKE^2$) is positive and of a similar magnitude. The first order impact of $STRIKE^2$ with TIME is negative for all the models. Again, this suggests that the curvature of the surfaces becomes more extreme as expiration is approached. When considering the second order time effect on the curvature, the impacts were not consistent across all the markets. Even so, the effect tends to be a positive one (for those markets where the impact was significant).

An important result is the statistically significant relationship between the expected kurtosis implied in the implied volatility patterns and the expected level of variance. We find that a significantly negative coefficient for this variable

($STRIKE^2*ATMVOL$) exists (β_{14}). This suggests that the higher the level of the expected variance, the flatter the curve of the implied volatility pattern.

Another significant result is that for all four stock index options and for the overall aggregate of all stock index options, there are significant third order strike price effects. The coefficient for $STRIKE^3$ variable (β_{15}) is positive for all four markets and roughly of the same order of magnitude. This variable can be seen as the interaction of the first and second order strike price effects and this suggests that the effect of excess kurtosis is asymmetrical. One interpretation of this result is that the higher the strike price of the option, the greater the level of the excess kurtosis. We will choose to coin this effect hetero-kurtosis.

This result has not previously appeared in the literature and may suggest that additional higher order terms are required to model the dynamics of stock index options implied volatility patterns correctly. Furthermore, even though the magnitude of these coefficients is small, the t-statistics are among the highest in the models (and can thus be interpreted as being among the most significant effects).

The other non-strike price related variables are for the most part insignificant apart from the time-related variables. While the fact that these variables may suggest that an error in variables problem may exist, subsequent regression models eliminated these effects without significantly affecting the sign or relative size of the coefficients for the other independent variables.[12]

For the fixed income options markets, there is a similar degree of consistency in the importance of the independent variables to each model. In Table 2, there are five first order STRIKE-related independent variables and the coefficients of the Betas are fairly consistent for all bond option markets. These include both the two

[12] See footnote 9.

TIME interaction variables. As with the stock index options, these indicate that the skewness has a negative relationship with TIME and a positive relationship with $TIME^2$. Again, the interpretation of this coefficient is that if a negative skew exists, it will flatten as the expiration date of the option is approached. Another first order strike price related variable that is consistent across all fixed income option markets exists for the level of the at-the-money volatility (ATMVOL). The consistent negative coefficient indicates that the higher the level of the expected variance, the more the volatility pattern is skewed. Finally, it appears that market specific shocks have an important impact on the nature of the first order strike price effect. For almost all shocks (apart from the second shock for the US T-Bond), the impact was to increase the negative skew of the implied volatility pattern.

As with the stock index options, there is a consistently positive kurtosis effect measured by the coefficient for $STRIKE^2$. There is also consistency in the first order impact of TIME. This result suggests that the curvature of the implied volatility patterns becomes more extreme as the options expiration date is approached. The only exception exists for the BTP. However, the increased curvature effect comes from the second order time factor, $TIME^2$. Finally, the only other consistent independent variable (for the second order strike price effect) is the relationship between the curvature and the level of the at-the-money implied volatility. As with the stock options markets, the significantly negative coefficient for this variable ($STRIKE^2$*ATMVOL) suggests that the higher the level of the expected variance, the flatter the curve of the implied volatility pattern.

As with the stock index options, we observe that higher strike price effect moments are significant in modelling the implied volatility dynamics. The third order strike price variable ($STRIKE^3$) is positive and of a similar magnitude for all the fixed income options examined. Most of the other independent variables do not display the same degree of consistency either in terms of the sign or the magnitude. This may suggest that if an error in variables problem exists the effect is not systematic across all the four markets.

For the foreign exchange options markets, we observe less consistency among the four markets we examined, and there is less consistency compared with the results of the previous two asset classes. In Table 3, we find that that overall first order strike price effect is not consistent (or even significant for the four markets). It would appear that skewness in not normally endemic for these assets. While the pure first order effect (skew) does tend to be positive (in three of four markets), the effect for all foreign exchange options is insignificant. There also does not seem to be any consistency in the relationship between the first order strike price effect with either of the two TIME interaction variables or for the level of the at-the-money volatility. The only first order strike price effect somewhat consistent across the foreign exchange markets is that the CRASH tended to cause a negative skew to appear in the implied volatility patterns. However, this effect is slight. An opposite effect exists for the occurrence of the second shock for each market. This tended to cause the first order effect to be slightly more positive. One could claim that our interpretations of these coefficients are misleading given that we have chosen the foreign currency as our numeraire.[13]

For the second order strike price effects, there is a great deal of consistency both between the four currency options examined and with the two previous asset

[13] All these options (and the underlying futures) are expressed as the number of US Dollars per unit of foreign currency. If these were expressed as the inverse, then the coefficients would reverse sign. However, given that the coefficients are of different signs for the four markets and in many instances insignificant, even with this transformation of the price series, the first order strike price effect is not as important for foreign exchange as it is for the fixed income and stock index markets.

classes. As with both the two previous asset classes, there is a consistently positive kurtosis effect measured by the coefficient for $STRIKE^2$. There is also consistency in the first and second order impacts of TIME. Again, this result suggests that the curvature of the implied volatility patterns becomes more extreme as the options expiration date is approached. Finally, the only other consistent independent variable (for the second order strike price effect) is the relationship between the curvature and the level of the at-the-money implied volatility. As with the other two asset classes, the significantly negative coefficient for this variable ($STRIKE^2$*ATMVOL) suggests that the higher the level of the expected variance, the flatter the curve of the implied volatility pattern.

There also appears to be a higher order strike price effect required for understanding the dynamics of foreign exchange implied volatilities. The third order strike price variable, ($STRIKE^3$), is also significant for all the four markets. However, while for the previous two asset classes this relationship was positive, this relationship is a negative one for all markets apart from the British Pound (which was barely significant). Of the other independent variables, the only variable that is consistent across all the four markets and overall is the ATMVOL variable.

Overall, it would appear that foreign exchange options are not as homogeneous in the nature of the implied volatility process as stock index or fixed income options. Nevertheless, the adjusted R-squared statistic suggests that the majority of the variation in the implied volatility surfaces is being explained. Thus, it could be said that while fewer variables are required to understand the nature of the foreign exchange options implied volatility dynamics, the significant variables provide relatively more explanatory power.

From this analysis, we can compare the factors that describe the implied volatility process across asset classes. It is clear that individual option markets have idiosyncratic features that cause them to differ even within their own asset class. Nevertheless, there does appear to be consistency within an asset class for those independent variables that explain the implied volatility dynamics.

A key concern in any such regression analysis is that the high degree of explanatory power is due to over-fitting within a sample period. To address this issue a number of tests were done. The initial test was to the regressions with all contracts included as dummy variables. This was done to assess if the regression results would hold across time. If the strike price effects were significantly different over time, this effect should be indicated by significant coefficients for the contract dummy variables and important changes in the coefficients of the key variables in our regressions. What we observed was that few of the contract dummy variables were statistically significant and there appeared to be minor impacts on changes in the coefficients of the key variables of interest to this research. Nevertheless, it is important to assess if the overall conclusions we have drawn from the regression results would hold if the regressions were run for sub-periods of the data set.

To test the stability of the regression equations, the observations for each of the twelve markets was split into two sub-periods. These periods were divided roughly into halves. With these divided sets of observations, we re-ran the OLS regression solely for the first period of the data set and then used these results to predict the standardised implied volatilities in the second period. The form of the regression model appears in equation 5.

$$VSI = \alpha + \beta \cdot VSI^{*} + \varepsilon \qquad (5)$$

Where VSI is the standardised implied volatility outside of sample and VSI* is the predicted standardised implied volatility using the results from the regression (equation 4) using the first half of the data sample.

To assess the effectiveness of the predictions, we ran a simple OLS regression and were interested in the R squared, and the coefficients of the regression equation. The results of this test can be seen in Table 4 for all twelve markets.

We observe that in most cases, the predicted standardised implied volatilities are well explained by the models solely determined using the first portion of the options data set. Thus, the results do not appear to be period specific.[14]

7. CONCLUSIONS AND IMPLICATIONS

In this paper, we examined the implied volatility surfaces for twelve option markets. The underlying assets were financial futures representing a cross section of Stock indices, Fixed Income and Currency markets.

A standardisation method was developed that allowed for the comparison of implied volatility patterns over time. These patterns were drawn over time to yield volatility surfaces. From a comparison of all the figures, it appears that consistencies exist for the options in the same asset class. For example, the stock index options all display a similar first order and second order strike price effect. For all the fixed income options, the shapes are almost identical. This also appears to be the case for the currency options. Given that the implied volatility surfaces for all twelve markets are smooth and well behaved, a functional form was tested for these surfaces based upon a polynomial expansion to degree three on the strike price and time.

[14] Regressions were also run for both periods and compared to the regression results for the entire period. It was found that most of the Beta coefficients were of a similar sign and order of magnitude for all three periods of analysis. Results available from the Author upon request.

To examine quantitatively whether the implied volatility surfaces were consistent over time for individual markets and to compare the surfaces between markets, a model was developed explained the majority of the variance in the implied volatility surfaces for the twelve option markets under investigation. Thus, we conclude implied volatility surfaces can be explained well by a relatively simple model and that the effectiveness of this model is fairly stable overtime.

While the first order strike price effect (skewness) is idiosyncratic, it appears that the second order strike price effect for all markets is similar both in absolute effect and time dependency. All markets experience more implied kurtosis as the expiration of the option is approached. In addition, the strike price effects (both first and second order) for all markets are inversely related to the level of the expected variance. An increase in the at-the-money volatility serves to increase a negative skew in the implied volatility pattern and decreases the curvature of the implied kurtosis in the pattern.

These findings suggest that market participants are using a similar algorithm over-time to adjust option prices away from Black-Scholes-Merton values. What remains for further research is to understand what this algorithm is. In previous work, which examined the objective dispersion processes for financial assets, Scott (1994) and Bates (1996) have found that these markets are best described by models which include both stochastic volatility and jump processes. One could interpret the first order and second order strike price effects we have identified in this paper as reflecting both these factors.

The existence of a consistent negative (skewed) first-order strike price effect might suggest that market participants are concerned about negative jumps. There are a number of possible explanations for this. The most obvious is that the 1987 stock market crash led participants to include a risk premium of such a negative jump. This research confirms the finding by Rubinstein (1994). Thus, it would be appropriate to expect jumps in order to understand this effect.

Furthermore, extensive evidence has appeared in the literature that the implied volatilities are stochastic. Thus, it is consistent that an understanding of the dynamics of implied volatilities requires a stochastic volatility model. One possible area of future research is to examine the link between models for the objective dispersion process and the risk neutral process associated with options prices. Given that both stochastic volatility and jump processes may require risk premia to exist, another area of research would be to examine the nature of this risk premia.

References:

Barone-Adesi, G. and R.E. Whaley, 1987, Efficient Analytical Approximation of American Option Values, The Journal of Finance 42, 301-320.

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

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. and M. Scholes, 1973, The Pricing of Options and Corporate Liabilities, Journal of Political Economy 81, 637-654.

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

Derman, E. and I. Kani, 1994, Riding on the Smile, RISK 7, 32-39.

Dumas, B., J. Fleming and R.E.. Whaley, 1996, Implied Volatility Functions: Empirical Tests, Working Paper, Fuqua School of Business, Duke University, February 1, 1996.

Dupire, B., 1992, Arbitrage Pricing with Stochastic Volatility, Working Paper, Société Générale Options Division.

Dupire, B. 1994, "Pricing with a Smile," RISK 7, 18-20.

Fung, W.K.H., and D.A. Hsieh, 1991, Empirical Analysis of Implied Volatility: Stocks, Bonds and Currencies, Paper presented at the Fourth Annual Options Conference of the Financial Options Research Centre, University of Warwick, Coventry, England, 19-20, July, 1991.

Galai, D., 1983, A Survey of Empirical Tests of Option Pricing Models, in: M. Brenner, Option Pricing: Theory and Applications, (Lexington Books, Lexington, Mass) 45-80.

Galai, D., 1987, Empirical Studies of Options Prices: An Updated Review, AMEX/SOFFEX Seminar, Zurich, September.

Heynen, R., A. Kemna and T. Vorst, 1994, Analysis of the Term Structure of Implied Volatilities, Journal of Financial and Quantitative Analysis 29, 31-56.

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

MacBeth, J. and L. Merville, 1979, An Empirical Examination of The Black-Scholes Call Option Pricing Model, The Journal of Finance 34, 1173-1186.

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

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

Rubinstein, M., 1985, Nonparametric Tests of Alternative Options Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 through August 31, 1978, The Journal of Finance 40, 455-480.

Rubinstein, M., 1994, Implied Binomial Trees, The Journal of Finance 49, 771-818.

Scott, L.O., 1994, Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Applications of Fourier Inversion Methods, Working Paper, Department of Finance, University of Georgia.

Shimko, D., 1991, Beyond Implied Volatility: Probability Distributions and Hedge Ratios Implied by Option Prices, Working Paper, University of Southern California.

Shimko, D., 1993, Bounds of Probability, RISK 6, 33-37.

STATISTICA for Windows, 1995, General Conventions and Statistics I (2nd Edition), Volume 1 (StatSoft, Inc., Tulsa, Oklahoma).

Tompkins, R.G., 1994, Options Explained^2 (Macmillan Press, Basingstoke, England).

Xu, X. and S. Taylor, 1993, Conditional Volatility and the Informational Efficiency of the PHLX Currency Options Market, Working Paper, Financial Options Research Centre, University of Warwick.

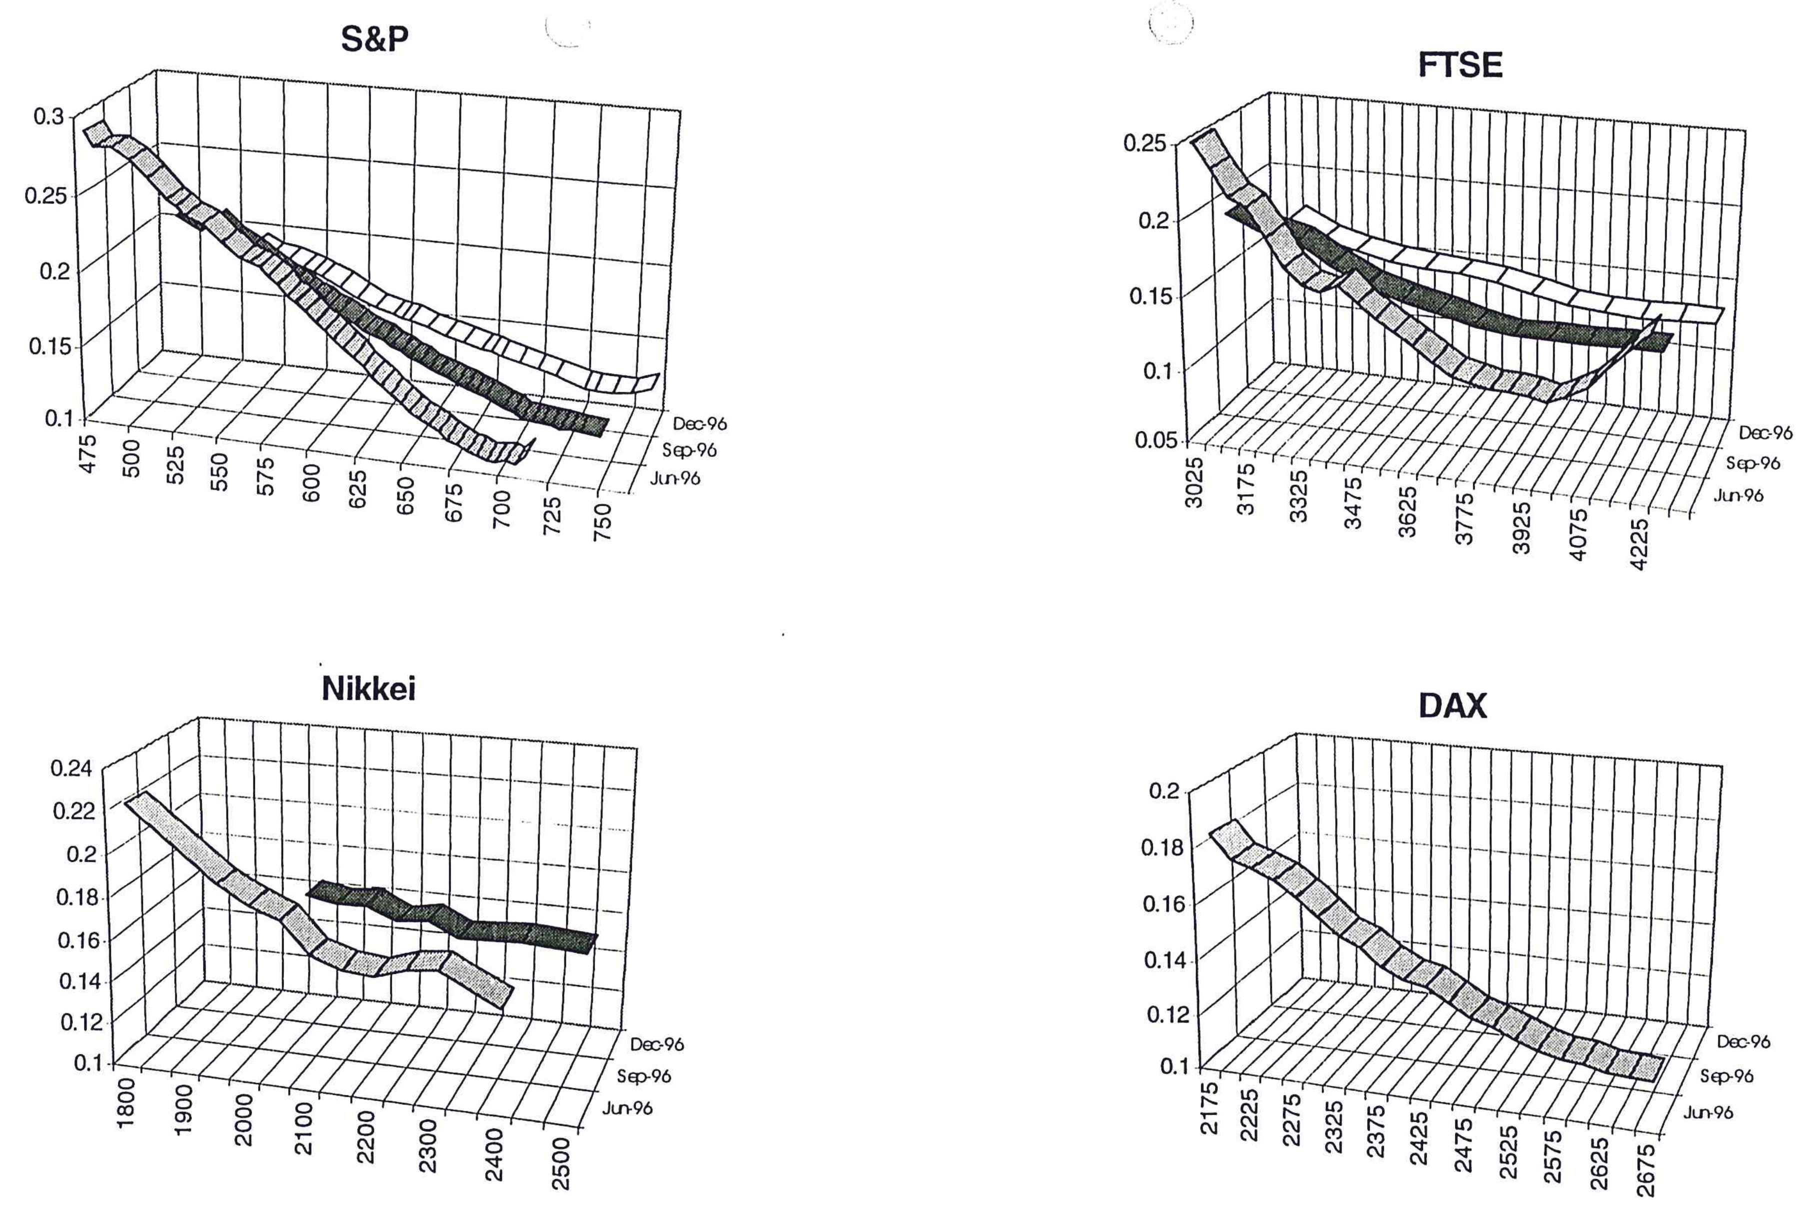

Figure 1 Implied Volatility Smiles for Four Stock Index Options as of May 7^{th}, 1996

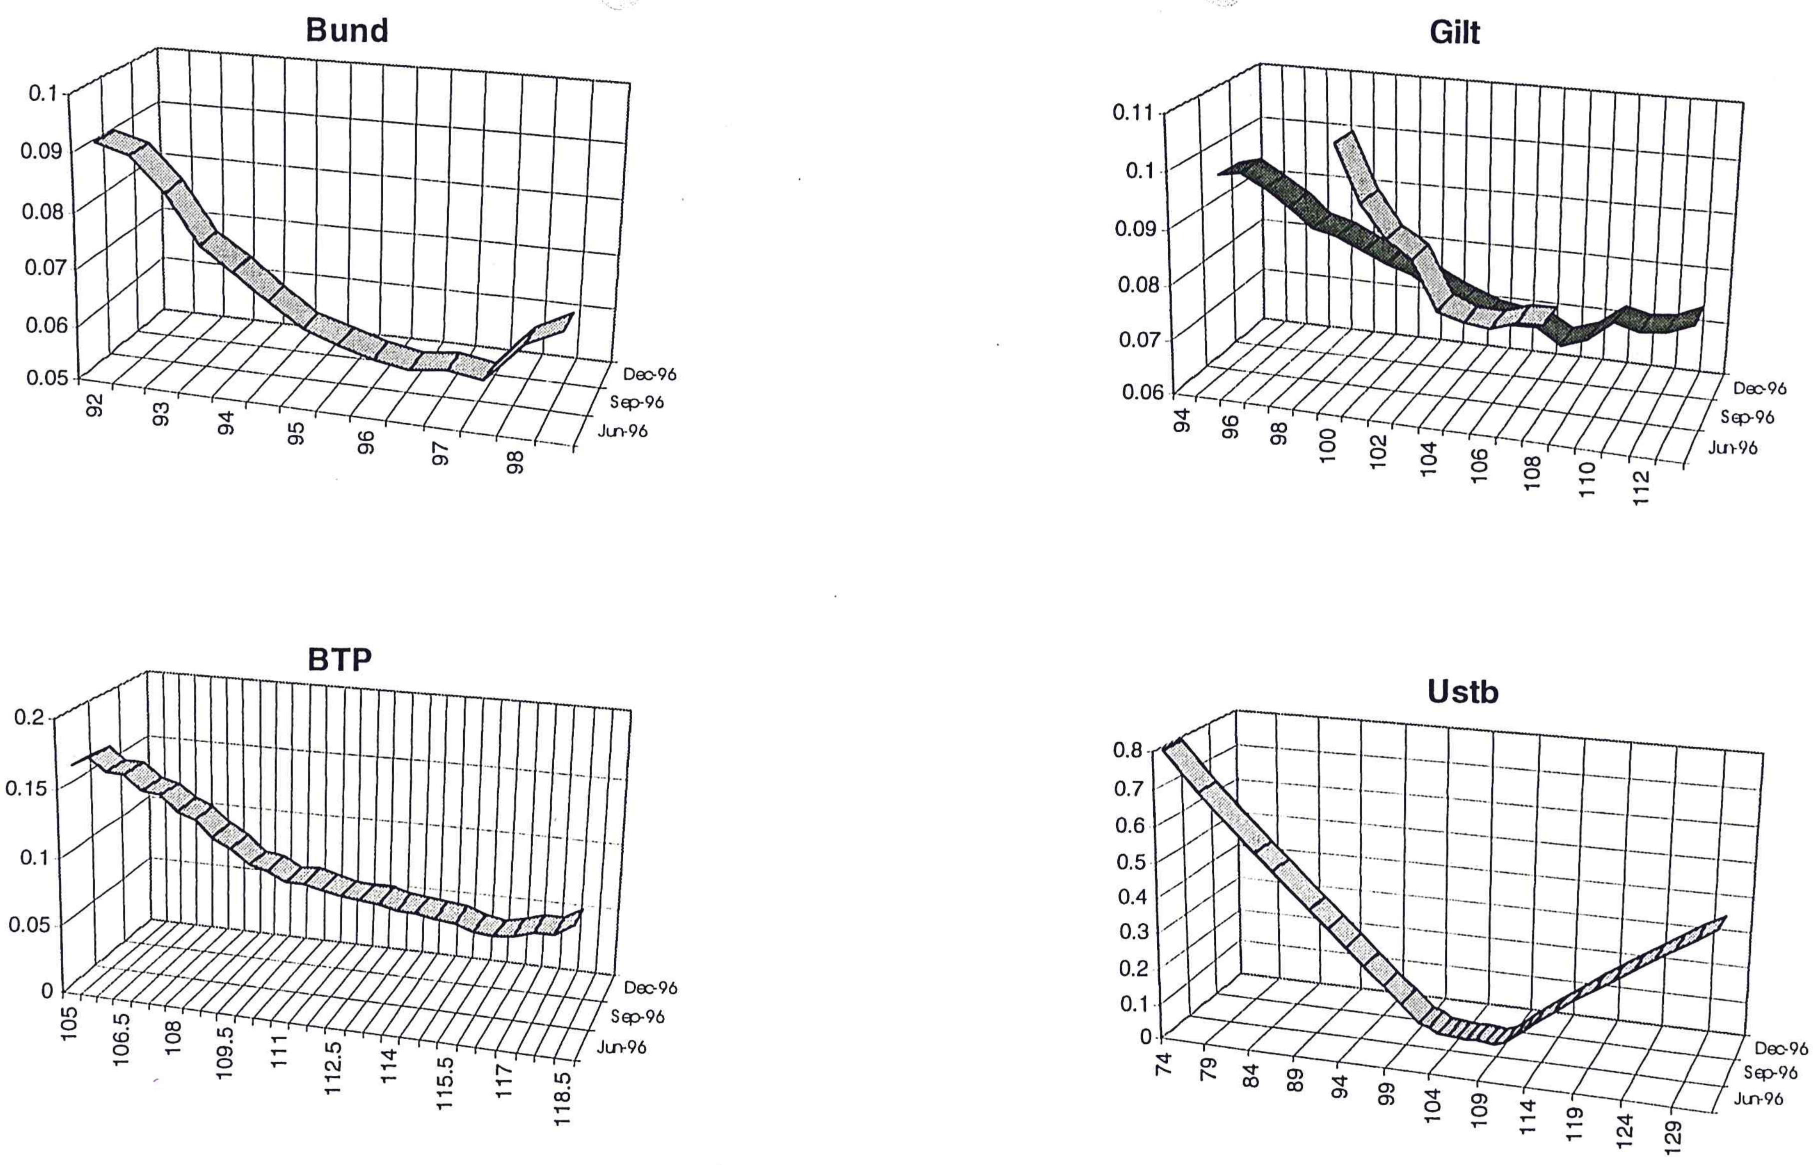

Figure 2 Implied Volatility Smiles for Four Fixed Income Options as of May 7th, 1996

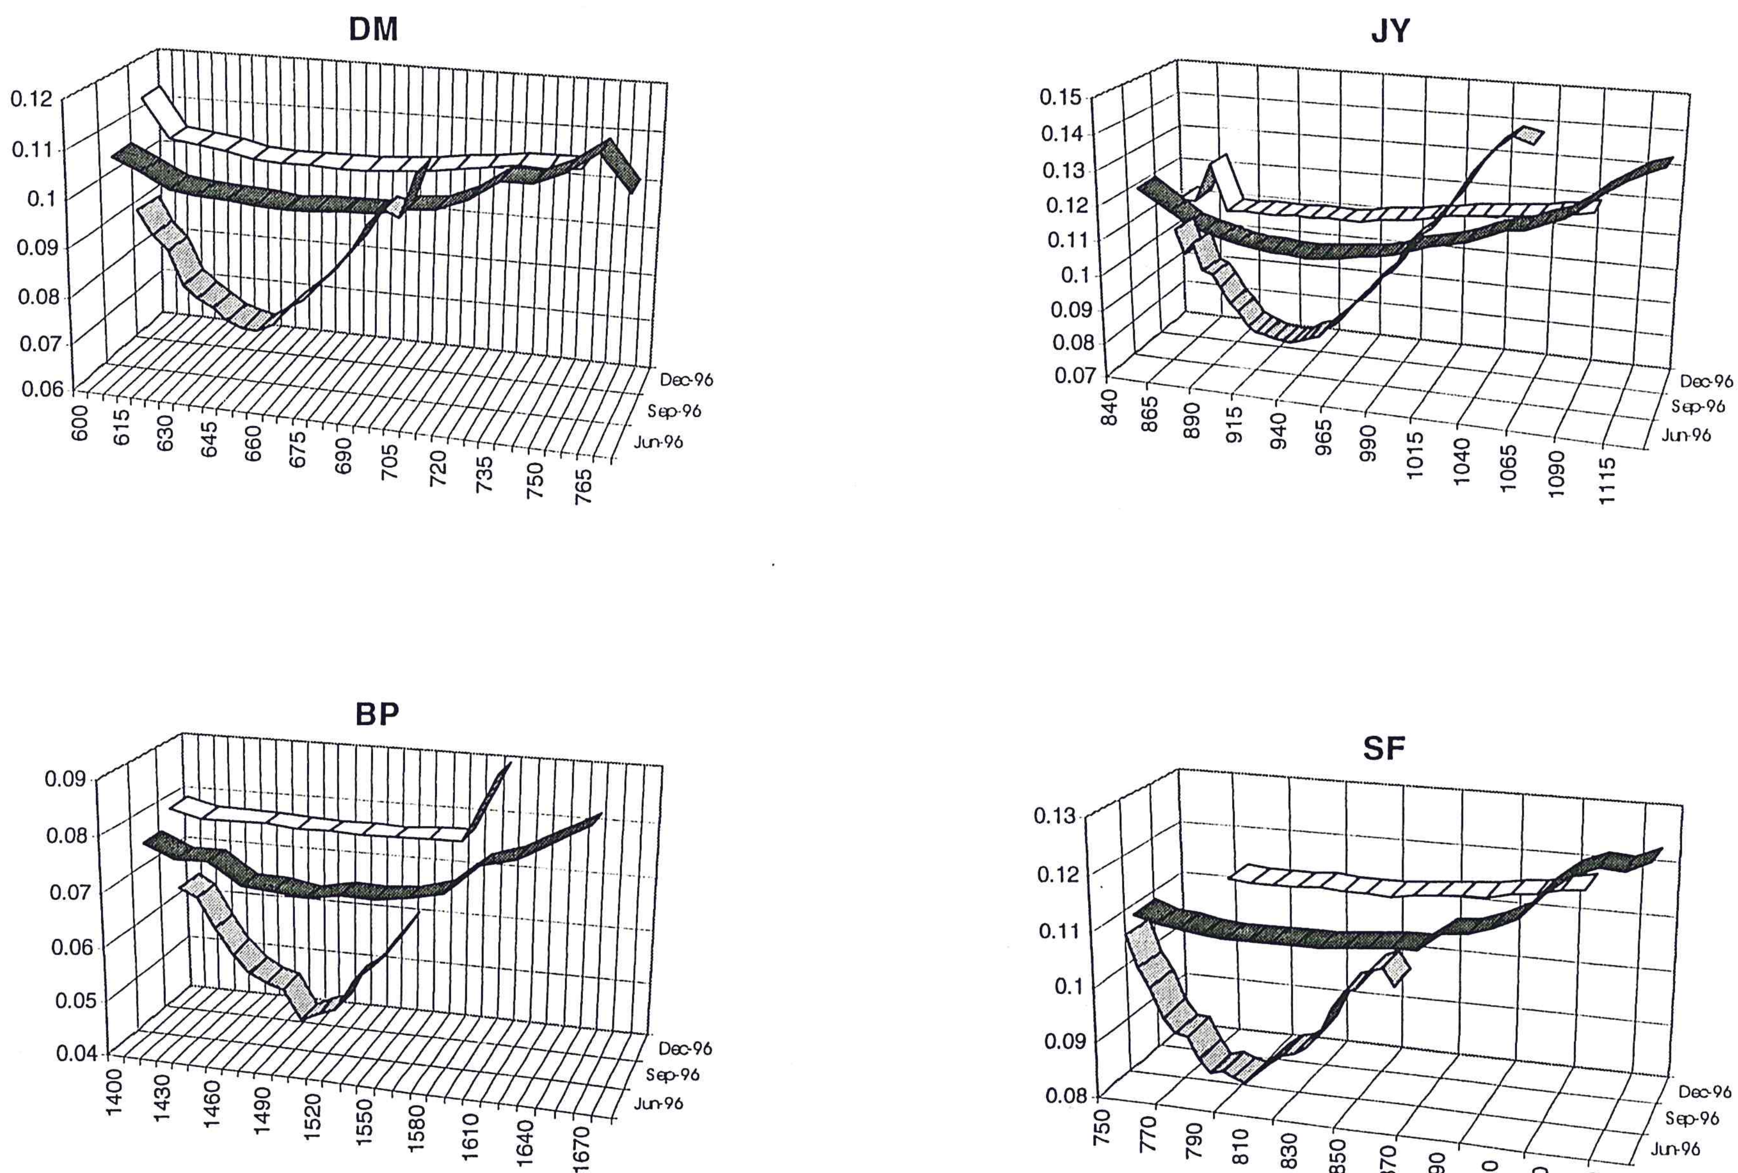

Figure 3 Implied Volatility Smiles for Four Foreign Exchange Options as of May 7th, 1996

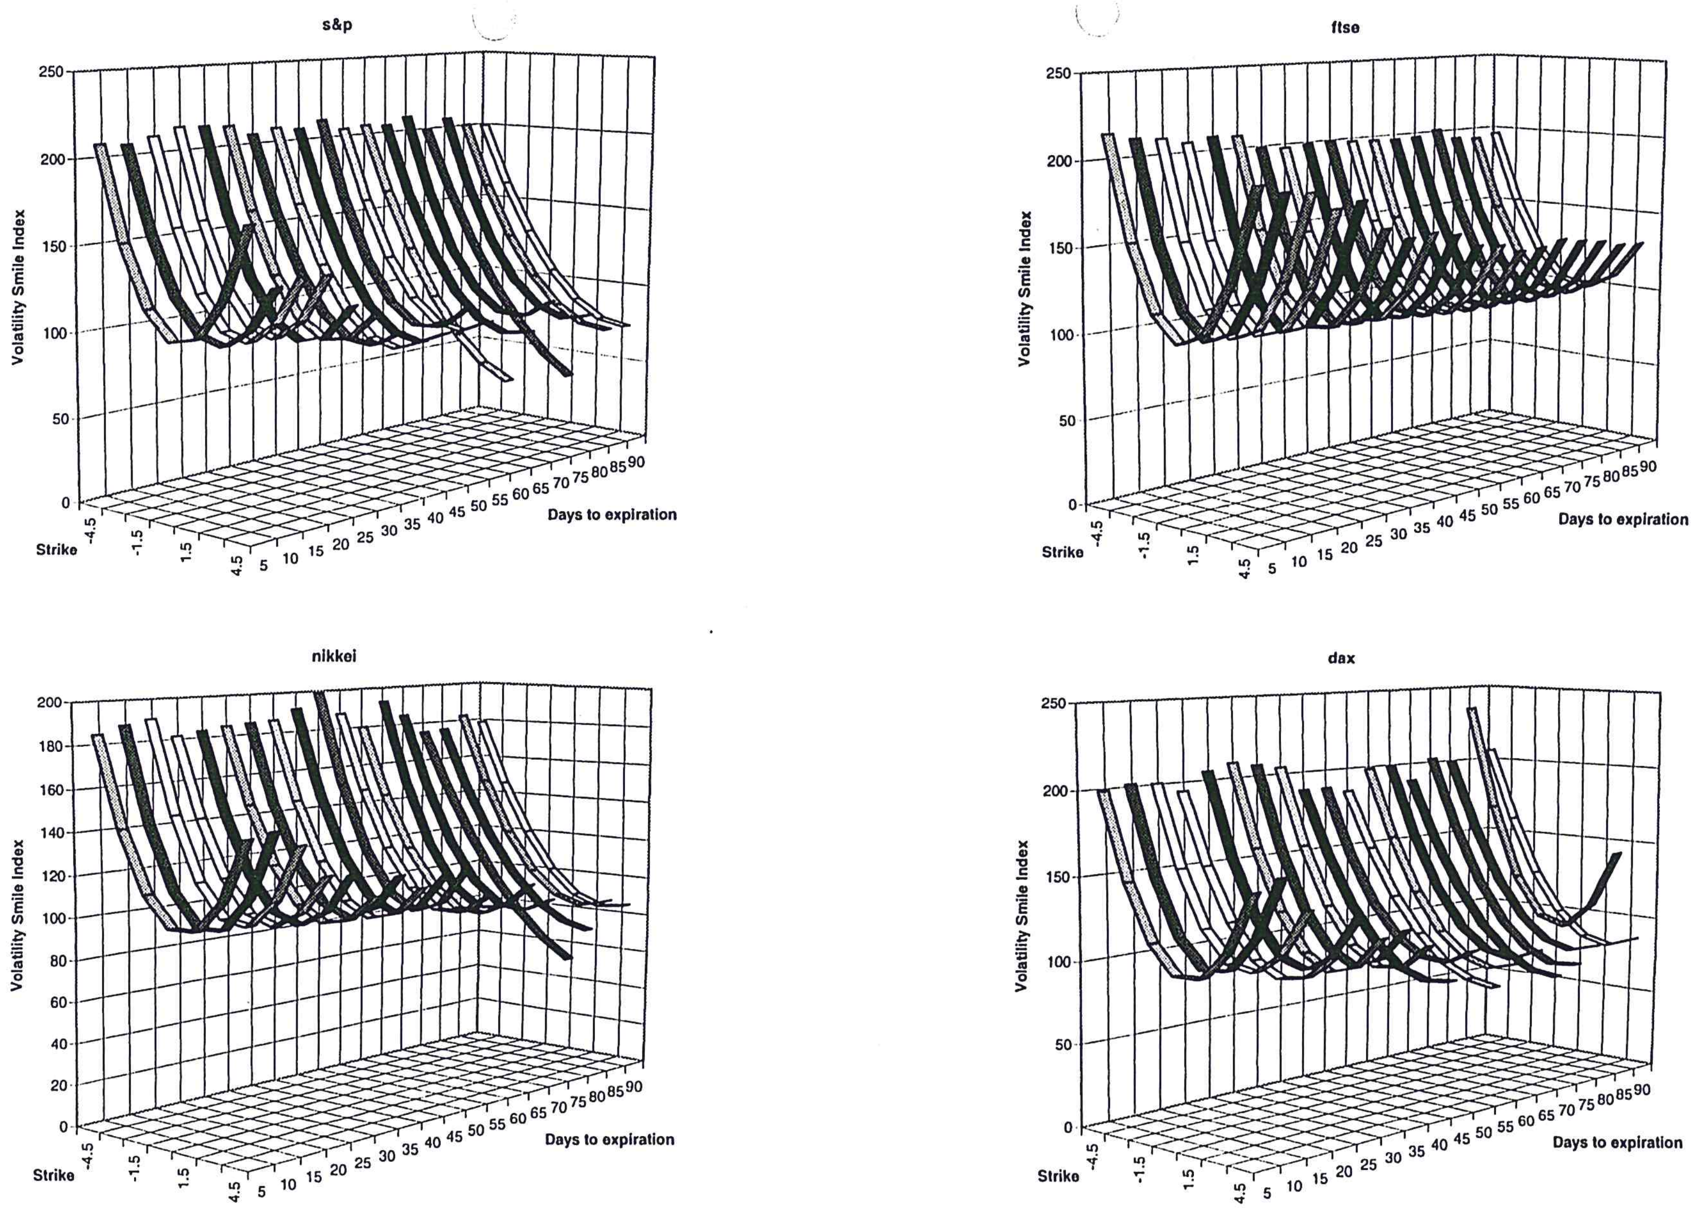

Figure 4 Standardized Volatility Smiles for Four Stock Index Options

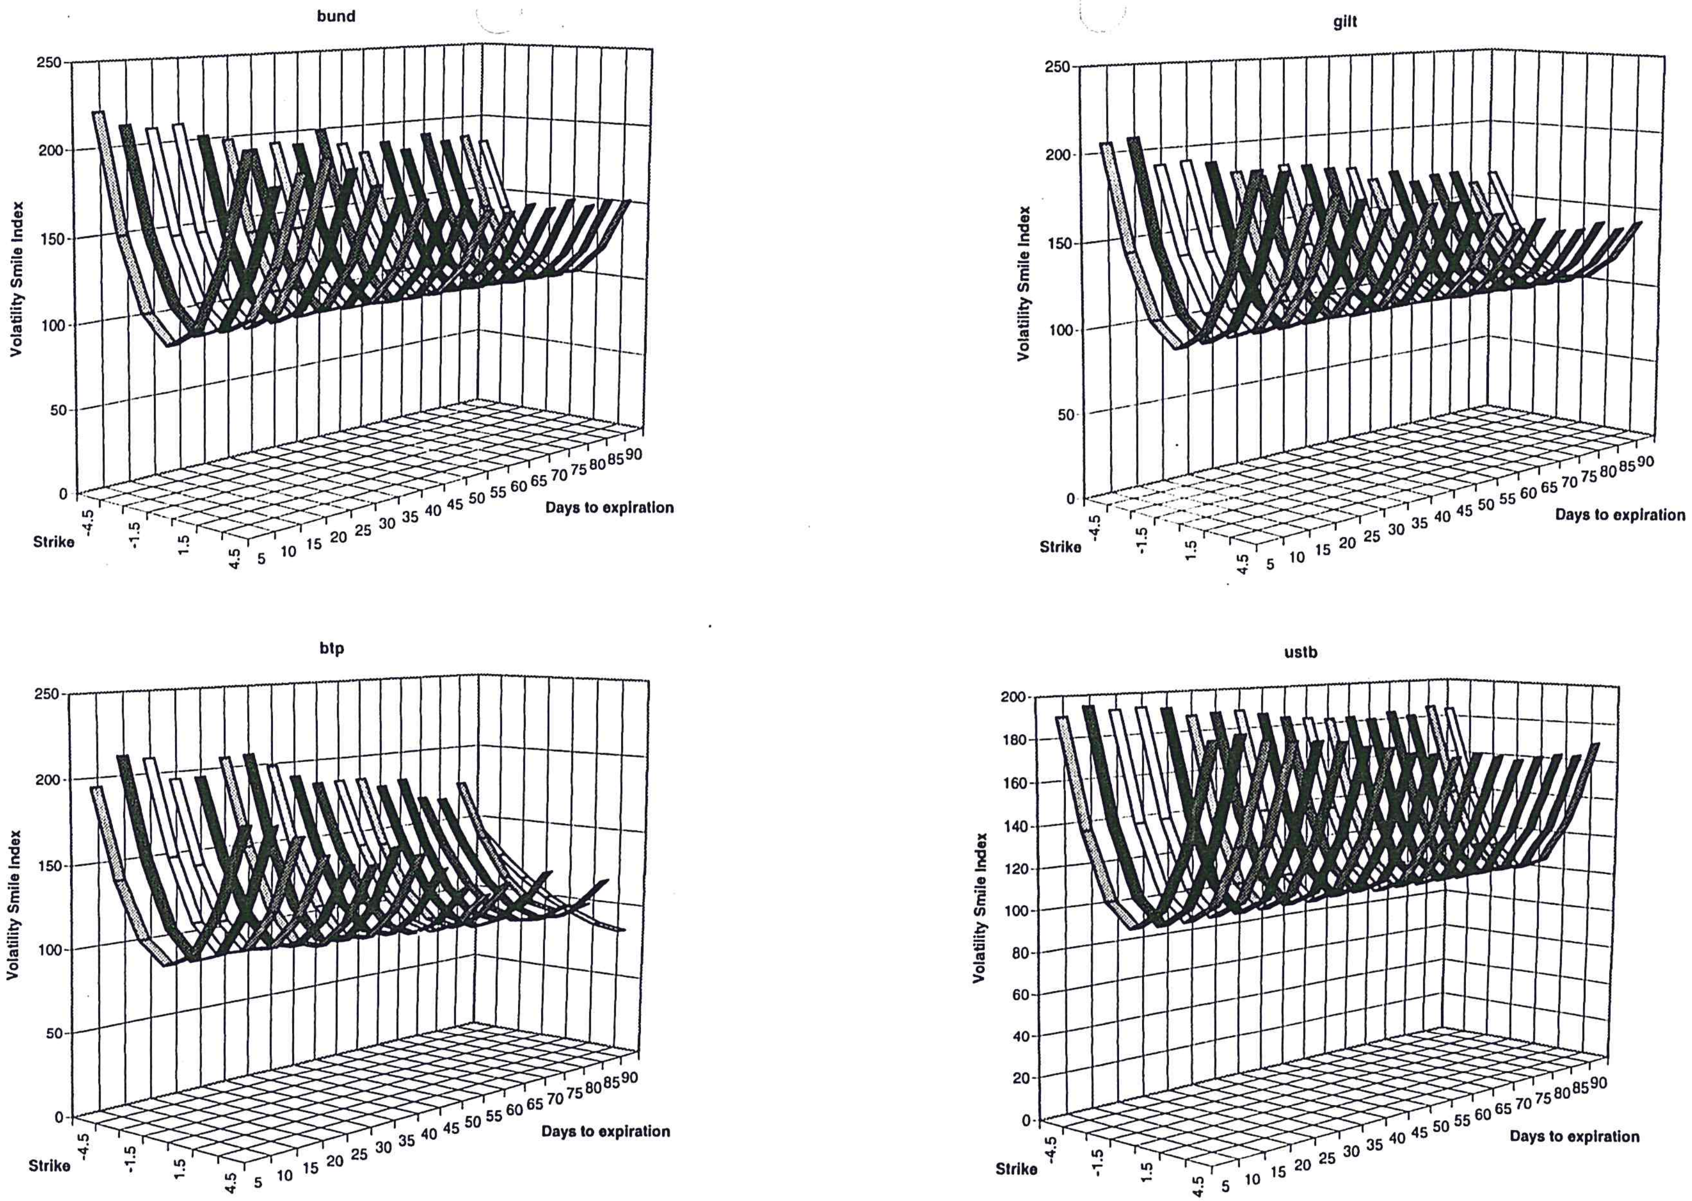

Figure 5 Standardized Volatility Smiles for Four Fixed Income Options

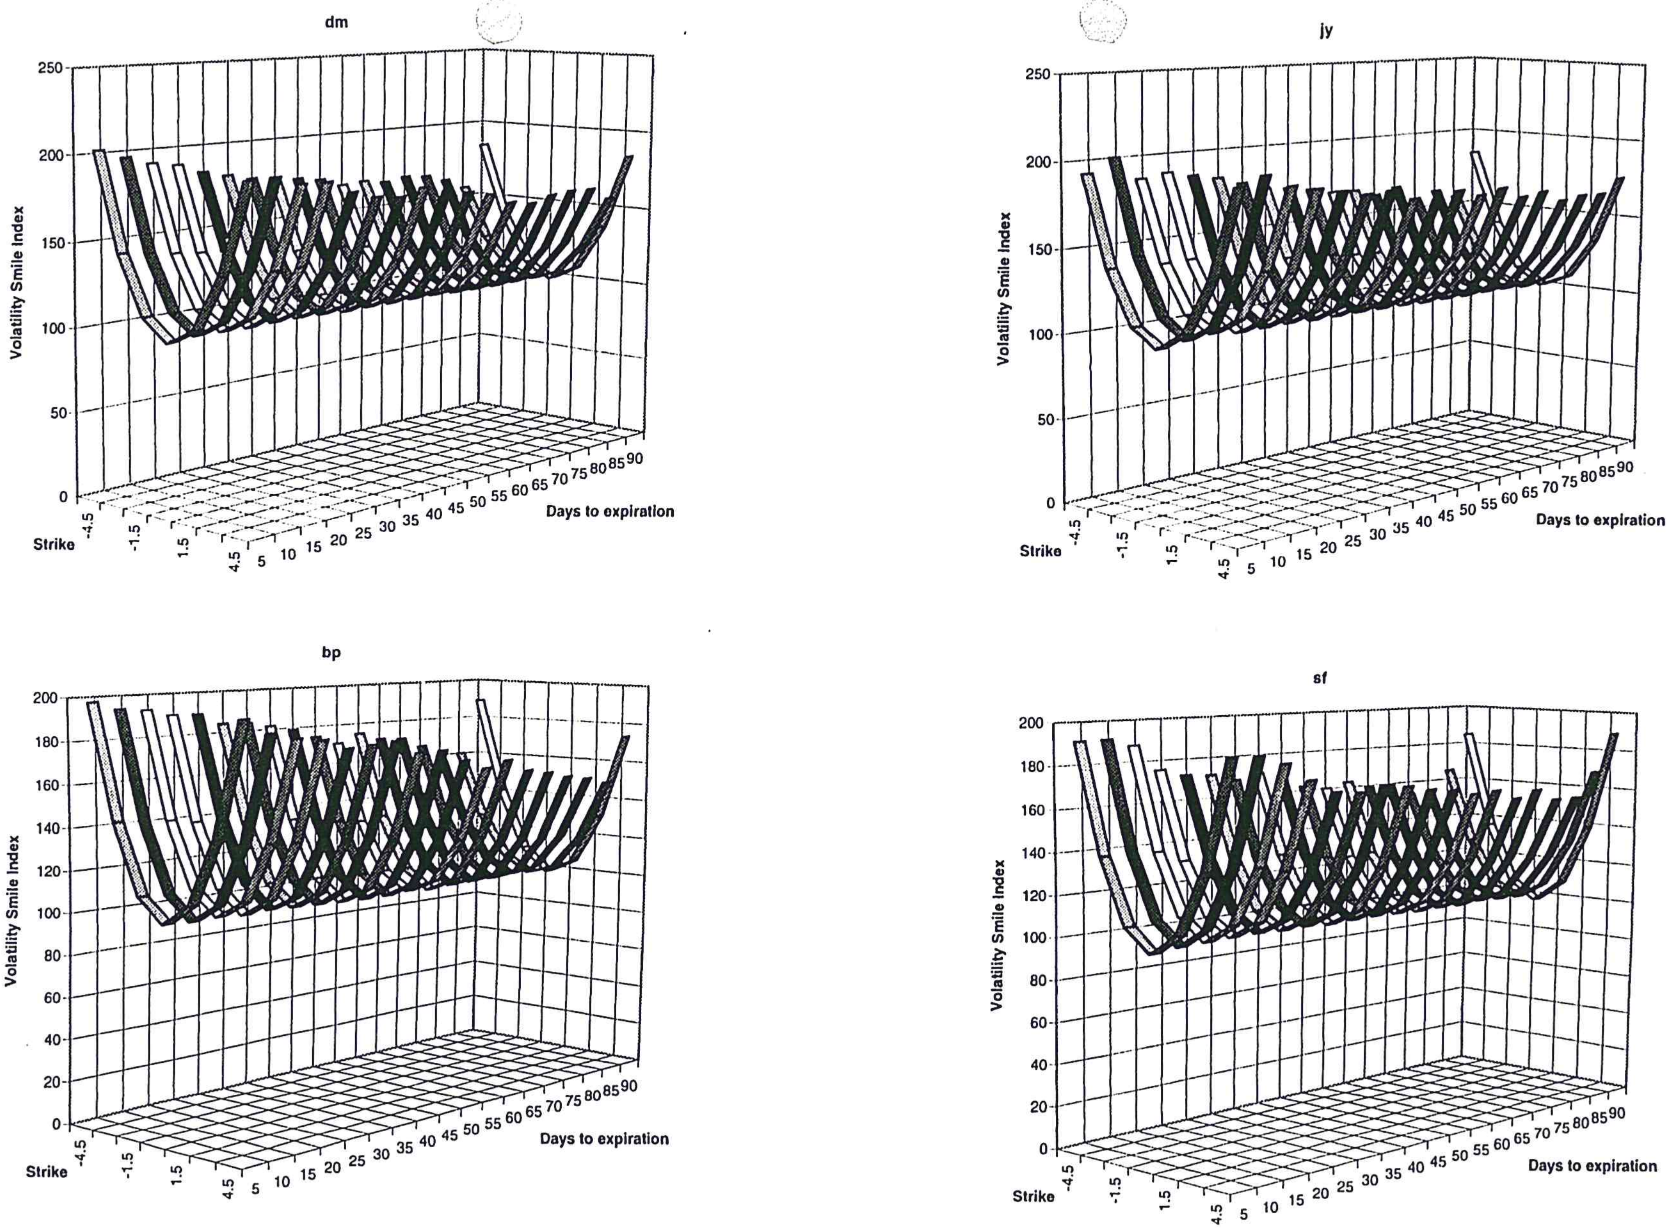

Figure 6 Standardized Volatility Smiles for Four Foreign Exchange Options

FACTOR		S&P			FTSE			DAX			NIKKEI		
		COEFFICIENT	Standard Error	T-Statistic	COEFFICIENT	Standard Error	T-Statistic	COEFFICIENT	Standard Error	T-Statistic	COEFFICIENT	Standard Error	T-Statistic
INTERCEPT	α	97.163	0.4823	-5.8822	98.027	0.3849	-5.1260	98.2797	0.4978	-3.4562	98.115	0.6739	-2.7972
Strike	$\beta 1$	4.725	0.2804	16.8509	-2.062	0.4670	-4.4154	-8.1672	0.6239	-13.0912	-6.731	0.4883	-13.7846
Strike*Time	$\beta 2$	-66.526	2.9015	-22.9281	-22.797	3.4891	-6.5338	-50.35	4.9292	-10.2146	-47.974	4.8015	-9.9915
Strike*$Time^2$	$\beta 3$	159.673	11.2096	14.2443	56.272	13.4858	4.1727	89.7201	19.7946	4.5325	125.628	18.9691	6.6228
Strike*Crash	$\beta 4$	-9.925	0.2373	-41.8247	-5.547	0.4840	-11.4607	-,--	-,--	-,--	-,--	-,--	-,--
Strike*Shock1	$\beta 5$	-,--	-,--	-,--	-,--	-,--	-,--	0.8529	0.4175	2.0431	3.252	0.1776	18.3108
Strike*Shock2	$\beta 6$	-6.083	0.1553	-39.1693	-0.518	0.3890	-1.3316	-2.0974	0.2883	-7.2761	1.213	0.1929	6.2882
Strike*ATMVol	$\beta 7$	-17.17	0.8358	-20.5432	9.58	1.9842	4.8281	14.406	2.9011	4.9657	-9.103	1.2607	-7.2206
$Strike^2$	$\beta 8$	5.899	0.1296	45.5170	8.963	0.1926	46.5369	7.2216	0.2958	24.4146	4.031	0.2067	19.5017
$Strike^2$*Time	$\beta 9$	-12.295	1.2359	-9.9482	-6.491	0.4485	-14.4727	-13.8675	1.9932	-6.9576	-12.205	2.3352	-5.2265
$Strike^2$*$Time^2$	$\beta 10$	15.557	4.8132	3.2322	-,--	-,--	-,--	17.1435	9.1866	1.8661	29.648	9.5672	3.0989
$Strike^2$*Crash	$\beta 11$	0.681	0.1064	6.4004	-1.359	0.2802	-4.8501	-,--	-,--	-,--	-,--	-,--	-,--
$Strike^2$*Shock1	$\beta 12$	-,--	-,--	-,--	-,--	-,--	-,--	-,--	-,--	-,--	1.17	0.0775	15.0968
$Strike^2$*Shock2	$\beta 13$	-0.717	0.0624	-11.4904	-1.317	0.2396	-5.4967	-1.3096	0.1474	-8.8847	-,--	-,--	-,--
$Strike^2$*ATMVol	$\beta 14$	-4.200	0.3970	-10.5793	-9.147	0.7184	-12.7325	-11.3916	1.4880	-7.6556	-4.172	0.5318	-7.8451
$Strike^3$	$\beta 15$	0.801	0.0111	72.1622	0.388	0.0125	31.0400	0.3286	0.0211	15.5882	0.316	0.0207	15.2657
Crash	$\beta 16$	1.465	0.2700	5.4259	2.921	0.4515	6.4695	-,--	-,--	-,--	-,--	-,--	-,--
Shock1	$\beta 17$	-,--	-,--	-,--	-,--	-,--	-,--	-,--	-,--	-,--	-,--	-,--	-,--
Shock2	$\beta 18$	0.441	0.1984	2.2228	-0.886	0.3528	-2.5113	-,--	-,--	-,--	-,--	-,--	-,--
ATMVol	$\beta 19$	-1.835	1.0970	-1.6727	-,--	-,--	-,--	4.0932	3.0059	1.3617	-,--	-,--	-,--
Time	$\beta 20$	54.822	10.7763	5.0873	-,--	-,--	-,--	1.7553	1.4661	1.1973	48.428	19.0513	2.5420
$Time^2$	$\beta 21$	-554.37	86.3968	-6.4166	-291.477	33.2399	-8.7689	-,--	-,--	-,--	-275.02	156.3578	-1.7589
$Time^3$	$\beta 22$	1538.426	207.5037	7.4140	1306.872	132.2727	9.8801	-,--	-,--	-,--	436.426	381.8629	1.1429
			(Observations) (12387)			(Observations) (6980)			(Observations) (2768)			(Observations) (3525)	
		R-Squared	0.9573		R-Squared	0.8875		R-Squared	0.9472		R-Squared	0.9084	
		Durbin-Watson	0.8765		Durbin-Watson	1.119042		Durbin-Watson	1.6069		Durbin-Watson	0.765779	

Table 1 Ordinary Least Squares Results for Four Stock Index Options

FACTOR		BUND			BTP			GILT			USTB		
		COEFFICIENT	Standard Error	T-Statistic	COEFFICIENT	Standard Error	T-Statistic	COEFFICIENT	Standard Error	T-Statistic	COEFFICIENT	Standard Error	T-Statistic
INTERCEPT	α	98.144	0.4259	-4.3580	98.471	0.3713	-4.1179	97.224	0.5885	-4.7171	93.853	0.8332	-7.3776
Strike	$\beta 1$	5.534	0.2575	21.4938	-0.044	0.2465	-0.1793	2.895	0.3527	8.2081	-0.969	0.3532	-2.7435
Strike*Time	$\beta 2$	-26.088	2.6754	-9.7511	-32.213	3.0246	-10.6502	-9.662	2.2113	-4.3694	-4.085	2.1397	-1.9091
Strike*Time^2	$\beta 3$	85.086	9.6950	8.7763	113.758	11.3120	10.0563	25.352	8.1773	3.1003	0.947	8.0677	0.1174
Strike*Crash	$\beta 4$	-.--	-.--	-.--	-.--	-.--	-.--	-1.585	0.1442	-10.9917	0.781	0.1262	6.1886
Strike*Shock1	$\beta 5$	-1.679	0.1600	-10.4918	-.--	-.--	-.--	-0.657	0.2867	-2.2916	-2.490	0.2658	-9.3679
Strike*Shock2	$\beta 6$	-2.496	0.0897	-27.8292	-2.994	0.1026	-29.1813	-4.156	0.0823	-50.4982	0.452	0.0836	5.4067
Strike*ATMVol	$\beta 7$	-96.526	2.3326	-41.3818	-44.109	2.0069	-21.9784	-22.411	1.7137	-13.0776	-7.141	1.6182	-4.4129
Strike^2	$\beta 8$	4.547	0.1473	30.8669	5.082	0.1358	37.4218	4.411	0.2187	20.1692	5.652	0.1317	42.9157
Strike^2*Time	$\beta 9$	-20.470	1.6384	-12.4937	1.847	1.2639	1.4614	-17.584	1.3549	-12.9781	-8.763	0.8932	-9.8108
Strike^2*Time^2	$\beta 10$	46.894	6.0827	7.7094	-52.302	5.2860	-9.8945	32.372	5.1336	6.3059	15.837	3.3913	4.6699
Strike^2*Crash	$\beta 11$	-.--	-.--	-.--	-.--	-.--	-.--	1.499	0.0901	16.6371	0.350	0.0384	9.1146
Strike^2*Shock1	$\beta 12$	1.538	0.1209	12.7192	0.870	0.1129	7.7100	1.226	0.1651	7.4258	-0.505	0.0928	-5.4418
Strike^2*Shock2	$\beta 13$	-1.325	0.0427	-31.0595	0.167	0.0532	3.1427	-0.757	0.0535	-14.1495	-0.318	0.0355	-8.9577
Strike^2*ATMVol	$\beta 14$	-.--	-.--	-.--	-17.698	1.2902	-13.7169	-22.729	1.0935	-20.7856	-5.578	0.6412	-8.6987
Strike^3	$\beta 15$	0.134	0.0156	8.6174	0.342	0.0122	27.9705	0.151	0.0141	10.7092	0.214	0.0064	33.4375
Crash	$\beta 16$	-.--	-.--	-.--	-.--	-.--	-.--	0.819	0.2431	3.3690	-.--	-.--	-.--
Shock1	$\beta 17$	-0.401	0.3070	-1.3064	-0.772	0.3235	-2.3874	-.--	-.--	-.--	2.134	0.3722	5.7335
Shock2	$\beta 18$	-.--	-.--	-.--	-0.703	0.1801	-3.9036	-0.278	0.1599	-1.7386	-2.552	0.2496	-10.2244
ATMVol	$\beta 19$	9.895	3.2354	3.0583	16.471	3.7509	4.3911	12.233	3.4196	3.5773	10.087	3.8118	2.6463
Time	$\beta 20$	16.737	3.3321	5.0229	-.--	-.--	-.--	-11.485	10.6026	-1.0832	46.468	14.4177	3.2230
Time^2	$\beta 21$	-.--	-.--	-.--	-.--	-.--	-.--	240.888	81.9159	2.9407	-321.256	113.0303	-2.8422
Time^3	$\beta 22$	-185.076	46.1445	-4.0108	70.187	18.1272	3.8719	-700.835	192.7889	-3.6352	789.732	268.8714	2.9372
			(Observations) (8248)			(Observations) (8588)			(Observations) (9058)			(Observations) (9528)	
		R-Squared	0.8098		R-Squared	0.8631		R-Squared	0.8151		R-Squared	0.9016	
		Durbin-Watson	0.8882		Durbin-Watson	0.9224		Durbin-Watson	1.1456		Durbin-Watson	1.0426	

Table 2 Ordinary Least Squares Results for Four Fixed Income Options

FACTOR		D-MARK COEFFICIENT	D-MARK Standard Error	D-MARK T-Statistic	POUND COEFFICIENT	POUND Standard Error	POUND T-Statistic	YEN COEFFICIENT	YEN Standard Error	YEN T-Statistic	S-FRANC COEFFICIENT	S-FRANC Standard Error	S-FRANC T-Statistic
INTERCEPT	α	101.667	0.4569	3.6484	100.214	0.4792	0.4466	103.295	0.3773	8.7327	102.402	0.6149	3.9063
Strike	$\beta 1$	3.933	0.2422	16.2363	-1.002	0.1947	-5.1464	1.247	0.2256	5.5271	2.85	0.3645	7.8189
Strike*Time	$\beta 2$	1.045	0.4019	2.6000	1.929	0.4824	3.9988	-3.428	1.8318	-1.8715	2.177	1.7708	1.2294
Strike*Time^2	$\beta 3$	-.--	-.--	-.--	-.--	-.--	-.--	10.701	6.8146	1.5702	-4.231	6.5244	-0.6485
Strike*Crash	$\beta 4$	-0.595	0.1283	-4.6420	-0.741	0.1401	-5.2891	-1.052	0.5256	-2.0009	-.--	-.--	-.--
Strike*Shock1	$\beta 5$	-1.929	0.1986	-9.7125	-.--	-.--	-.--	1.809	0.5099	3.5470	-2.126	0.2859	-7.4362
Strike*Shock2	$\beta 6$	0.552	0.0598	9.2217	0.333	0.0670	4.9701	1.830	0.0870	21.0383	-.--	-.--	-.--
Strike*ATMVol	$\beta 7$	-15.234	1.1390	-13.3751	4.471	1.0472	4.2695	-6.647	1.3472	-4.9339	-7.234	1.3241	-5.4633
Strike^2	$\beta 8$	7.838	0.1227	63.8766	9.154	0.5791	15.8073	6.340	0.1033	61.3452	8.367	0.1655	50.5559
Strike^2*Time	$\beta 9$	-16.273	0.6849	-23.7600	-14.431	0.9325	-15.4756	-12.943	0.6957	-18.6034	-19.269	0.9969	-19.3289
Strike^2*Time^2	$\beta 10$	36.311	2.6070	13.9284	31.282	3.6792	8.5024	24.408	2.7085	9.0114	53.228	3.7284	14.2764
Strike^2*Crash	$\beta 11$	0.156	0.0641	2.4366	-1.254	0.0782	-16.0358	-0.750	0.1926	-3.8948	0.159	0.1138	1.3972
Strike^2*Shock1	$\beta 12$	-1.087	0.0968	-11.2276	-1.176	0.5676	-2.0719	0.576	0.1838	3.1336	-1.876	0.1237	-15.1657
Strike^2*Shock2	$\beta 13$	-0.127	0.0343	-3.6884	0.391	0.0365	10.7123	0.266	0.0413	6.4547	0.342	0.1084	3.1550
Strike^2*ATMVol	$\beta 14$	-17.166	0.6302	-27.2368	-20.62	0.5570	-37.0197	-12.820	0.6716	-19.0906	-19.014	0.7583	-25.0745
Strike^3	$\beta 15$	-0.089	0.0091	-9.7300	0.032	0.0086	3.7209	-0.169	0.0081	-21.0004	-0.068	0.0100	-6.8000
Crash	$\beta 16$	0.522	0.1912	2.7314	0.517	0.1815	2.8485	3.956	0.6039	6.5510	1.878	0.4488	4.1845
Shock1	$\beta 17$	0.532	0.3495	1.5215	-.--	-.--	-.--	-4.215	0.5737	-7.3476	-0.599	0.3974	-1.5073
Shock2	$\beta 18$	-0.714	0.1180	-6.0456	-0.283	0.1408	-2.0099	1.664	0.1620	10.2706	-2.566	0.4262	-6.0206
ATMVol	$\beta 19$	-13.729	2.2351	-6.1427	-10.127	2.1022	-4.8173	-19.256	2.6602	-7.2384	-25.721	2.6020	-9.8851
Time	$\beta 20$	-.--	-.--	-.--	19.445	9.6800	2.0088	-3.350	0.9182	-3.6487	48.242	8.2293	5.8622
Time^2	$\beta 21$	-.--	-.--	-.--	-223.217	74.9091	-2.9798	-.--	-.--	-.--	-285.325	64.6848	-4.4110
Time^3	$\beta 22$	-.--	-.--	-.--	651.221	175.7895	3.7046	-.--	-.--	-.--	-.--	-.--	-.--
			(Observations) (11079)			(Observations) (9190)			(Observations) (12998)			(Observations) (11834)	
		R-Squared	0.8565		R-Squared	0.8832		R-Squared	0.8634		R-Squared	0.8006	
		Durbin-Watson	1.403844		Durbin-Watson	1.53857		Durbin-Watson	1.33727		Durbin-Watson	1.4904	

Table 3 Ordinary Least Squares Results for Four Foreign Exchange Options

Underlying Asset	Alpha	Beta	R Squared	# Observations
S&P 500 Futures	17.8859	0.8304	0.9655	7622
FTSE Futures	-29.9282	1.3397	0.8870	5180
Nikkei 225 Futures	-0.4213	0.9884	0.9199	1534
DAX Futures	5.2918	0.9631	0.9369	1840
Bund Futures	18.0858	0.8556	0.7563	4671
BTP Futures	3.7422	0.9727	0.8930	5031
Gilt Futures	6.6633	0.9735	0.7384	6739
US T-Bond Futures	33.6890	0.7131	0.6851	7918
Deutsche Mark /US $	14.1239	0.8615	0.8331	7183
British Pound / US $	16.1462	0.8525	0.8905	5353
Japanese Yen / US $	26.6671	0.7812	0.8182	8941
Swiss Franc / US $	24.1125	0.7826	0.7906	7744

Table 4 Regression Results for the Predicted Standardised Implied Volatilities in the Second Period of the Options Data Set using the First Period Model