

Modelling the Dynamics of Implied Volatility Surfaces

Ihsan Ullah Badshah *

May 2008

Abstract

The objective of this study is to model implied volatility surfaces and identify risk factors that account for most of the randomness in the volatility surface. The approach is similar to the Dumas, Fleming and Whaley (DFW) (1998) study; we use moneyness (e.g., in forward price) and out-of-the-money (OTM) put-call options on FTSE100 index. After these adjustments, the nonlinear parametric optimization technique is employed to estimate different models of DFW in order to characterize the implied volatility surfaces and produce smooth implied volatility surfaces. Next, principal component analysis (PCA) is applied to the implied volatility surfaces to extract principal components that account for most of the dynamics in the shape of the surface. Hence, we estimate and obtain smooth implied volatility surfaces with the parametric models that account for both smile and time to maturity, and, therefore, the constant volatility model fails to explain the variations in the surfaces. Finally, we find that the first three principal components can explain about 69-82% of the variances in the implied volatility surfaces. The applications of our study are hedging of derivatives positions, trading, risk management, and policy making.

Keywords: implied volatility surface, smile, implied volatility, principal component analysis, options

JEL Classification: G13

*HANKEN-Swedish School of Economics and Business Administration, Department of Finance and Statistics, P.O. Box 287, FIN-65101 Vaasa, Finland. Phone: +358-6-3533 721, Fax: +358-6-3533 703, Email: ihsanullah.badshah@hanken.fi

1 Introduction

Modelling and exploring the dynamics of financial assets' volatility has been the main issue and central focal point among finance academics and practitioners. Volatility is fundamental for risk management, option pricing, hedging of derivatives positions and policy making. Most previous research has focused on historical volatility, but recently implied volatility has received attention due to some breakthrough studies, such as Christensen and Probhala (1998), Fleming (1998), and Dumas, Fleming and Whaley (1998). They show that implied information is superior to historical when forecasting volatility and further point out that a precise notion of the market's judgment and expectation of volatility is vital. For instance, to develop expectations about volatility based on past behaviour of stock prices and other relevant information is backward-looking and called historical volatility. In contrast, implied volatility is forward looking, that is implied by market prices of options. Option prices are the common consensus of the market participants about the expected future volatility. Therefore, implied volatility is the market expectation about the average future volatility of the underlying asset over the remaining life of an option. The volatility expectation of market participants can be recovered by inverting the option-pricing formula. However, it is well known that after the October 1987 market crash, implied volatility computed from the options' prices on stock indexes appears to be different across strikes and term structure, if examined at the same time, for example, the same day. This implies that implied volatilities present a two-dimensional surface whose dynamics across strikes and over time has to be examined.

However, the Black and Scholes (1973) model assumes that underlying assets follow a geometric Brownian motion and constant volatility. This implies that options with different strikes and maturities on the same underlying asset should have the same implied volatility. In the real world, however, we observe departure from this assumption. Mostly asset prices are

influenced by risk factors such as jumps, stochastic volatility, and costs occurs during transactions (Carr et al 2001, Heston 1993 and Leland 1985). In order to account for these deviations, traders and practitioners use different implied volatilities for different strikes and maturities. Hence, the implied volatility of an option reflects determinants of the option value that are not captured by the BS-model; this volatility structure illustrates discrepancies between theoretical prices and the market. Thus, a volatility surface that includes all these features is essential in practice. A surface helps particularly in pricing illiquid options and in hedging exotic derivatives.

On the other hand, financial markets present a high degree of correlations, which is high dependence among market-risk factors. When few important sources of information are common to market risk factors, we find a high degree of correlation among risk factors. For risk management systems that hedge and price huge portfolios, these portfolios might consist of different securities. Hence, they depend on hundreds of underlying correlated risk factors. Principal component analysis (PCA) is a tool that extracts the most important independent risk factors from these correlated systems that can explain most of the dynamics. Hence, PCA makes it easier for risk managers to reduce the dimensionality of the systems and enhance computational efficiency. Similarly, implied volatility surfaces can be viewed as highly correlated, that could be explained by just few independent risk factors.

There are few studies that estimate and obtain smooth implied volatility surfaces and further study the dynamics of these surfaces by applying principal component analysis. However, we can find abundant studies examining either Smile/Skew (strike) or term structure of volatilities. One of the famous studies on implied volatility surfaces was conducted by Cont and Fonseca (2002), who examined the dynamics of implied volatility surfaces of SP500 and FTSE100 index options. The prices of an index option at a given date are represented by a corresponding implied volatility surface, having smile/skew and term

structure features. They used different methods to obtain smooth implied surfaces. However, the main findings were, first, that an implied volatility surface has a non-flat shape and displays both a strike and term structure. Second, the shape of the implied volatility surface changes with time. Third, implied volatilities have positive autocorrelation and behave mean reversion. Fourth, the variance of the daily log variations in implied volatility can be explained by two or three principal components. Finally, the movements in the underlying assets are not correlated with movements in the implied volatility. Another study on implied volatility surfaces was conducted by Roux (2006). First, he estimated surfaces and then examined the dynamics of these surfaces. For that he developed an econometric model to estimate the implied volatility surface and then explored the dynamics of the implied volatility surface by the principal component analysis (PCA) technique. However, his data consist of the VIX index and options on the SP500 index; he found that 75.2% of the variations of the implied volatility surface can be explained by the first principal component (PC) and another 15.6% by the second PC. Moreover, he found that processes for these factors appear to be independent of the process for the VIX index. Alentorn (2004) adapted a parametric form of the Dumas, Fleming and Whaley (DFW) (1998) approach and estimated implied volatility surfaces for options on FTSE100 futures. He used three of the DFW models. He further suggests how to implement this methodology in real time. This study, however, differs from the above two. This study only estimates implied volatility surfaces, whereas the other two further study the dynamics of implied volatility surfaces, an addition to estimation of implied surfaces.

However, many studies have been conducted that study the dynamics of either smile or surfaces by applying PCA; one famous study by Skiadopoulos, Hodges, and Clewlow(1999) investigated the number and shape of shocks that move implied volatility smiles and surfaces of S&P500 index options by applying principal component analysis. They formed maturity

buckets in the surfaces, average implied volatilities of options whose maturities fall into them and apply PCA to each bucket covariance matrix separately. They identified two factors that explain about 60% of the variance. They suggest that for effective pricing hedging of future options, one needs only three risk factors: one for the underlying asset, and the other two for implied volatility. Another study was conducted by Alexander (2001), who developed a new principal component model of fixed strike volatility deviations from at-the-money volatility. She focuses on its application to the skew and used FTSE100 index options with maximum three-month maturities. She investigated if the index moves, the volatility surface will move continuously and furthermore, if second and higher PCs have non-zero conditional correlation with the index changes, there will be non-parallel movements in the surface as the index moves. She found that most of the variation in the implied volatility of the FTSE100 index option can be explained by just three key risk factors: parallel shifts, tilts and curvature changes. However, the latter study differs from the former in a manner that it studied only the skew, whereas the former studied both skew and surfaces.

The objective of this study is to estimate and obtain smooth surfaces while using parametric models, and further study the dynamics of these surfaces by applying PCA. However, our study is related to the studies of Dumas, Fleming and Whaley (DFW) (1998) in terms of functional forms of their estimated models for different implied volatility surfaces, while the second part of our study is closer to the study by Skiadopoulos, Hodges, and Clewlow (1999), in which they apply PCA on implied volatility surfaces. Nevertheless, our study differs in some ways from the two cited studies; first, we estimate implied volatility surfaces by using out-of-the-money (OTM) puts for low strikes, and out-of-the-money calls for high strikes, which are not accounted in DFW. Second, we study implied volatility surfaces in term of moneyness and time to maturity, whereas DFW incorporate strike instead. Third, in our next analysis that employs principal component analysis (PCA) to extract those

risk factors that explain most of the dynamics in the smile surfaces and determine what these factors look like and how many factors can explain these surfaces. We study the dynamics of the surfaces of FTSE100 index options (i.e. recent three years), whereas Skiadopoulos, Hodges, and Clewlow (1999) study future options on the S&P index.

Our main findings for this study are that the DFW constant volatility model fails to capture the variations in the implied volatility surfaces and, consequently, generate flat volatility surfaces. Model 2 which captures the smile (skew) effects is good enough; however, there are some variations that can be attributed to the time dimension in the implied volatilities. Therefore, models such as 3 and 4, which account for variations in both dimensions, fit well to the surfaces. They produce smooth and tighter volatility surfaces to the observed data. Finally, we find that the first three principal components can explain about 69-82% of the variances in the implied volatility surfaces. However, when applying PCA to maturity groups, we find that parallel shifts are evident and important for shorter and longer maturities of option implied volatilities.

This paper is organized as follows. Section 2 consists of the methodology; it presents the implied volatility model, optimization of implied volatility surfaces, and principal component analysis of implied volatility surfaces. Section 3 describes our data. The empirical results are presented in Section 4. Section 5 concludes our study.

2 Methodology

We first present in Section 2.1 our model for implied volatility. Section 2.2 describes the parametric form of models and the optimization technique for implied volatility surface. Section 2.3 presents the principal component analysis of the implied volatility surface.

2.1 Model for Implied Volatility

Merton (1973) extended the Black-Scholes (BS) model for European call and put options on dividend paying index S_t with strike price K and time-to-maturity $\tau = T - t$, a call with payoff of $\text{Max}(S_t - K, 0)$ and put $\text{Max}(K - S_t, 0)$. The Black-Scholes formula for the value of these call and put options are

$$C(S_t, K, \tau, \sigma) = S_t e^{-q\tau} N(d_1) - K e^{-r\tau} N(d_2), \quad P(S_t, K, \tau, \sigma) = K e^{-r\tau} N(-d_2) - S_t e^{-q\tau} N(-d_1),$$

$$d_1 = \frac{\ln(S_t/K) + \tau(r - q + \sigma^2/2)}{\sigma\sqrt{\tau}}, \quad d_2 = \frac{\ln(S_t/K) + \tau(r - q - \sigma^2/2)}{\sigma\sqrt{\tau}}$$

where q is dividend yield, r is the risk free interest rate, and $N(d)$ is the cumulative unit normal density function with upper integral d . Now consider that in a market where the Black-Scholes assumptions do not hold as previously empirically found, the quoted market prices of call and put options are $C_t^*(K, T)$, and $P_t^*(K, T)$. Now the implied volatility of the BS-model $\sigma_t^{\text{BS}}(K, T)$ of the option is the value of the volatility parameter which equates the theoretical option price with the market price.

$$C_{\text{BS}}(S_t, K, \tau, \sigma_t^{\text{BS}}(K, T)) = C_t^*(K, T), \quad P_{\text{BS}}(S_t, K, \tau, \sigma_t^{\text{BS}}(K, T)) = P_t^*(K, T).$$

The Black-Scholes implied volatility can be found uniquely because of the monotonicity of the Black-Scholes formula in the volatility parameter,

$$\frac{\partial \text{BS}}{\partial \sigma} > 0.$$

For the fixed (K, T) , $\sigma_t^{\text{BS}}(K, T)$ is a stochastic process and for fixed t , its value depends on the characteristics of the option such as the maturity T and the strike K , the function

$$\sigma_t^{\text{BS}} : (K, T) \rightarrow \sigma_t^{\text{BS}}(K, T).$$

This is the implied volatility surface at date t . However, with moneyness of the option the implied volatility surface is a function of moneyness and time-to-maturity

$$I_t(M, \tau) = \sigma_t^{\text{BS}}(M(S(t), t + \tau)).$$

This depiction is appropriate because there is a range of moneyness nearby ATM position where options are liquid and empirical data are easily available. Cont and de Fonseca (2002) used a similar technique but their emphasis was only on call implied volatility, whereas this study includes both call and put implied volatilities.

To compute implied volatilities we use the Bisection Method. This method is efficient and fast and without the knowledge of Vega we can estimate implied volatilities. The Bisection method requires two initial volatility estimates (seed values) and the algorithm goes as follows:

- A low estimate of the implied volatility, σ_{Low} , corresponding to an option value, C_{Low} .
- A high volatility estimate, σ_{High} , corresponding to an option value, C_{High} .

The market price C_t^* lies between C_{Low} and C_{High} . The bisection estimate is then given as the linear interpolation between the two estimates. We use a linear interpolation and with support of an iterative process, given a desired degree of accuracy. With this iterative search method, we finally obtain the volatility figure that is the most accurate one. We use this approach in order to obtain both out of the money put and call implied volatilities.

2.2 Modelling Implied Volatility Surface

In order to model and obtain a smooth implied volatility surface, we follow the approach proposed by Dumas, Fleming and Whaley (DFW) (1998). According to Christoffersen and Jacobs (2004), this approach is known as the Practitioner Black-Scholes model. Its simplicity and application motivate us to study it thoroughly and extend it further. However, we approach the problem somewhat differently than the original Dumas, Fleming and Whaley (1998) work. First, we use out-of-the-money put and call data. Second, we obtain smooth surfaces with more practitioner-friendly moneyness, which takes into account forward prices as well as time dependence. We need to be able to update the volatility surface in order to take into account changes in the market prices and implied volatilities for the traded contracts. The implied volatility as a function of strike does not adequately capture volatility market movements, whereas the implied volatility as a function of moneyness parameter does. Third, we use data of FTSE100 index options, while they use S&P500 data. Fourth, we employ the parametric optimization technique to estimate surfaces, which is more common in practice.

One simple moneyness approach suggested by Fung and Hsieh (1991) and Jackwerth and Rubenstein (1996) is to take the ratio of the strike price and underlying (with the former inverting this ratio). They can easily reverse the equation in order to obtain the actual strike price. However, rather than absolute moneyness, many traders express moneyness based on the log of the forward price versus the strike. For a strike K and a forward price $Se^{r_t \tau}$, where r_t is interest rate until maturity of the contract, and τ is time to maturity. This ratio has the shortcoming that it is time independent. The time dependence is important when updating the volatility surface. There are many approaches in the market which take into account different information in the market. We follow Gross and Waltner (1995), Cassese and Guildolin (2006) and Alentorn (2004) in using an alternative to the strike over the underlying ratio.

Therefore, we use the moneyness that is more common in practice as well as suitable to this kind of study. Hence, we enter this moneyness separately to the models below:

$$m = \frac{\log(S_t e^{r\tau_t}/K)}{\sqrt{\tau_t}}.$$

We estimate and obtain the smooth implied volatility surface while using four different structural forms, each of which differs in the parametric structure that serves to characterize implied volatility:

Model 1: $I(m, \tau) = \beta_0 + \varepsilon$

Model 2: $I(m, \tau) = \beta_0 + \beta_1 m + \beta_2 m^2 + \varepsilon$

Model 3: $I(m, \tau) = \beta_0 + \beta_1 m + \beta_2 m^2 + \beta_3 \tau + \beta_4 \tau m + \varepsilon$

Model 4: $I(m, \tau) = \beta_0 + \beta_1 m + \beta_2 m^2 + \beta_3 \tau + \beta_4 \tau m + \beta_5 \tau^2 + \varepsilon$

Model 1 is the volatility function of the constant volatility model of Black-Scholes that gives constant volatility, without taking into account the smile (skew) and term structure implied volatility effects. Model 2 attempts to capture the quadratic volatility smile (e.g., moneyness slope and curvature) across moneyness. Model 3 captures additional variation attributable to the time (e.g., time to maturity slope), and a combined effect of the moneyness and time dimension. Lastly, Model 4 adds up a parameter that captures the time to maturity curvature effect.

Further details of our models are as follows: $I(m, \tau)$ represents the implied volatility dependence on the moneyness level as well as time to maturity. While β_0 is the level parameter which is constant for all models, β_1 captures the moneyness slope, the smile effect (skew). β_2 captures the moneyness curvature across moneyness, while β_3 is the parameter for the time to maturity slope, β_4 captures the time to maturity and moneyness combined, and finally β_5 represents the time to maturity curvature.

We estimate and obtain smooth surfaces through the above four structural forms by using the optimization technique; first, we choose optimization because this is more common in the market, particularly in the multifactor models, such as the famous LIBOR Market Models. Second, optimization is more flexible. Third, one needs to calibrate parameters to the market information continuously. Fourth, this can capture the nonlinear effects nicely. Nevertheless, we optimize the surface with a nonlinear least square function (lsqnonlin) in Matlab. This is a built in function; it follows an algorithm that minimizes the sum of squared errors between the actual and the prediction for a given vector of parameter. However, to test for the goodness of fit for these models, therefore, we employ two important loss functions proposed by Christoffersen and Jacobs (2004), which are relative root mean squared errors (%RMSE) and implied volatility root mean squared errors (IVRMSE). The relative root mean squared error loss function is the following

$$\%RMSE(\theta) = \sqrt{\frac{1}{N} \sum_{i=1}^N (e_i(\theta) / optionprice_i)^2}$$

The %RMSE loss function minimizes the relative difference between market (i.e. Black-Schole volaitities) and Practitioner Black-Scholes model volatilities. Hence %RMSE assign more weights to deep out-of-the-money options, as these values have small values. As we can notice that in the %RMSE loss function in the denominator includes market price of options, when market price of option is small then the difference between the volatilities become amplified. On the other hand, implied volatility root mean squared errors (IVRMSE) loss function assign equal weights to volatilities as follows

$$IVRMSE(\theta) = \sqrt{\frac{1}{N} \sum_{i=1}^N (\sigma_i - \sigma_i(\theta))^2}.$$

Thus IVRMSE loss function minimizes the error between the given and model obtained implied volatilities.

2.3 Principal Component Analysis of the Implied Volatility Surface

We show how to further simplify the implied volatility surface models to reduce the whole surface dimensions to just two or three factors, while using the technique of principal component analysis (PCA). The rationale is to extract independent factors that explain most of the dynamics of the implied volatility surface. Previously many researchers have attempted to extract those risk factors that explain most of the variations in the implied volatility, skew and surfaces; for FTSE100 index implied volatility, Alexander (2000) and Cont and da Fonseca (2002); for the S&P500 index implied volatility, Cont and da Fonseca (2002), and Skiadopoulos, Hodges, and Clewlow (1999); and for the DAX index implied volatility, Fengler, Härdle and Villa (2004). Their findings confirmed that about 70-90 percent of the total variations in the volatility surface or skew can be attributed to just three risk factors: parallel shift, tilt, and curvature, changes that are captured by the first three principal components. Our approach is, therefore, similar to that of Skiadopoulos, Hodges, and Clewlow (1999) but our data differ, i.e. FTSE100 index implied volatility surfaces are used, whereas they used SP500 index implied volatilities data.

Principal component analysis is a more suitable technique to identify those random shocks that influence the implied volatility surfaces which are independent risk factors. By applying this technique, we can identify important principal components without any assumption. Nonetheless, the data input to PCA must be stationary, since implied volatilities are often nonstationary. We, therefore, compute the first difference implied volatilities across different moneyness and time to maturities. For instance, if we do not make series stationary, we can end up with results that are mostly influenced by input factors with the greatest volatility. Moreover, we assemble data into different groups on the basis of trading days left to expiry, i.e. 8-30 days, 30-60, 60-90, 90-150, 150-220, and 220-above. After grouping, we then pool all groups to analyze the dynamics of the implied volatility surface. However, for a more

detailed investigation we apply PCA on each group, i.e. 30-60, 60-90, 90-150, and so forth.

A detailed illustration of principal component analysis is the following: Principal component analysis is a method of matrix decomposition into eigenvectors and eigenvalues matrices. We apply this method to the implied volatility surface, which in our case is a pool of implied volatilities across moneyness level and term structure. Thus, we decompose $p \times p$ matrix $B = V' * V$ between the variables V , which is a variance-covariance matrix of input data. We have time $\tau = 1, \dots, T$, and variable V represents the number of variables, i.e. volatilities. As V_i is $T \times 1$ vector of V . As we have T volatility observations of p variables. Now V is a $T \times p$ data matrix of volatilities. The objective is to find a linear combination of the first difference volatilities that explains maximum of variation in the implied volatility surface. Now the weights in the linear combination can be chosen from a set of eigenvectors of the correlation matrix, represented by W the $p \times p$ matrix of eigenvectors Z . Therefore, we have

$$ZW = W\Lambda \text{ or } Z = W\Lambda^{-1}$$

where Λ is $p \times p$ diagonal matrix of eigenvalues of Z . The column of W , (w_{ij}) for $i, j = 1, \dots, p$, are then sorted to the size of the corresponding eigenvalue that the r th column of W , denoted $w_r = (w_{1r}, \dots, w_{pr})'$, is the $p \times 1$ eigenvector corresponding to eigenvalue λ_r where the columns of Λ are sorted so that $\lambda_1 > \lambda_2 > \dots > \lambda_p$.

Therefore, we define the r th principal component as follows

$$P_r = w_{1r}V_1 + w_{2r}V_2 + w_{3r}V_3 + \dots + w_{pr}V_p$$

or in the matrix form,

$$P_r = VW_r.$$

Thus, each PC is a time series of the transformed v factors. The full $T \times p$ matrix of PCs is

$$P = VW.$$

To know that the PCs are orthogonal, we note that

$$P'P = W'V'XW = W'VW = W'W\Lambda .$$

Since W is uncorrelated, therefore, $W'W = WW' = I$. Thus,

$$PP' = \Lambda$$

where Λ is diagonal matrix, then the columns of P are orthogonal, and the variance of r th PC is λ_r . Nonetheless, the proportion of the total variation in V that is explained by the r th PC is

$\frac{\lambda_r}{p}$ where p is the sum of eigenvalues is the trace of Λ , the diagonal matrix of eigenvalues of

V . The trace of Λ is the number of variables in the system, which is in our case different maturity groups of implied volatilities. Therefore, the proportion of variation explained by the

first m PCs is

$$\sum_{i=1}^m \frac{\lambda_i}{p} .$$

Because of the choice of column labelling in W , the PCs is ordered so that P_1 belongs to the largest eigenvalue λ_1 , P_2 belongs to second largest eigenvalue λ_2 and so forth. In highly correlated implied volatilities, the first eigenvalue would be much larger than the others; as a result, the first PC can alone explain much of the variation in the volatilities. If the first two or three PCs explain most of the variation in the implied surface, then these PCs can replace the original p variables without loss of much information. Since $W' = W^{-1}$, we can write

$$V = PW'$$

that is

$$V_i = w_{i1}P_1 + w_{i2}P_2 + w_{i3}P_3 + \dots + w_{ip}P_p .$$

Thus, each of original volatility input data may be written as a linear combination of the principal components, which reduces the dimensions of the system. As we have noticed that once the data input V has been transformed into the principal component system, then the original system can be transformed back by the above equation.

3 Data

We use daily data on European type options on the FTSE100 index from Euronext for the period 1 January, 2004, to 31 October, 2006. The main data for our study are the end-of-the-day prices updated daily in the Euronext database. This database includes the following daily information for each call and put traded options: the style (call/put), the date, delivery date, strike price, volume, open interest; the opening, high, low and closing option prices; and underlying price. First, we use options data ranges from eight days up to one year maturities, and options across different strike prices. The second type of data is discount curve data with different maturities to proxy for a risk-free rate downloaded from Thomson DataStream.

We merely consider data for out-of-the-money (OTM) put and call options: OTM puts for low strikes and OTM calls for high strikes. OTM options have most of the information and exhibit high volatilities. Since a put option is a hedging instrument against the downward movements of the underlying asset, and investors want to hedge their positions with puts, they are willing to pay more for them. Hence, put options volatilities would be higher than calls. However, the data on in-the-money (ITM) puts and calls are not used because they are sensitive to the non-synchronicity problem. Moreover, options with less than 8 days to expiry are excluded because they are very sensitive to the small errors in the option price.

Figure 1 displays the scatter graph of the FTSE100 Index level and ATM implied volatility. On the x-axis we have ATM implied volatility, whereas on the horizontal axis the FTSE100 index level. We observe noticeable negative correlation between the two historical series. This evidence is consistent with previous research, which found similar evidence that the price is negatively correlated with implied volatility.

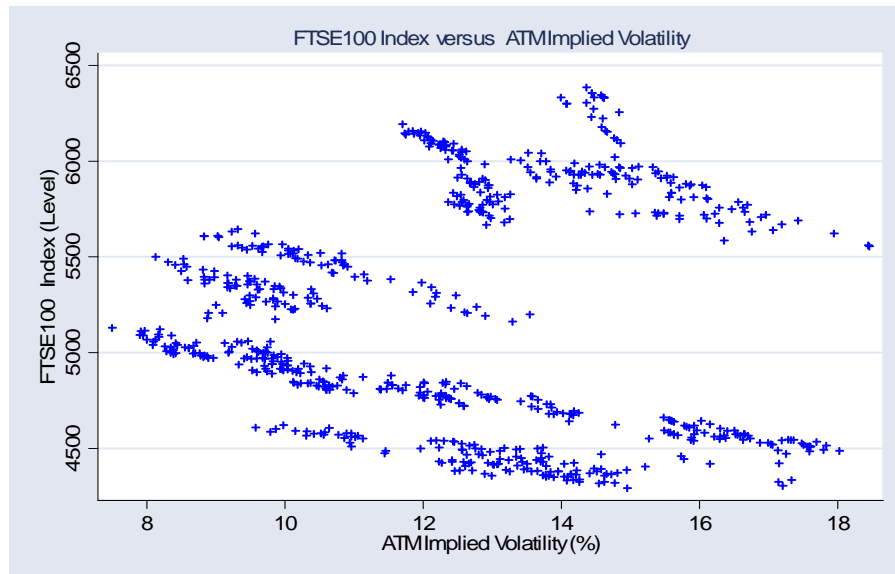


Figure 1: FTSE100 Index (level) versus ATM Implied Volatility from 1 January, 2004, to 31 October, 2006.

Table 1 presents the descriptive statistics for the FTSE100 index and ATM implied volatilities, both for level and log differences 1 January, 2004, to 31 October, 2006. Panel I of Table 1 shows descriptive statistics for the level. The mean ATM volatility is 12.5% annually, maximum 18.4% and minimum 7.5 %, whereas for the FTSE100 index the mean is 5178 points. The maximum level it reaches during the sample period is 6385 points and minimum 4292 points. However, Panel II of Table 1 documents the descriptive statistics on the log differences. We can observe from the skewness and kurtosis values that ATM-implied volatility is non-normally distributed. ATM implied volatility is negatively skewed, in order to be normally distributed; the skewness must be equal to zero and kurtosis 3, which is not the case here; thus, we observe here non-normality. Similarly, we observe non-normality in the underlying series. Previously researchers documented that non-normality in the underlying triggers non-constant implied volatilities. Panel III of Table 1 shows correlation between the FTSE100 index and ATM implied volatility. Thus, correlation between the FTSE100 index and ATM implied volatility is -0.48, which is significantly negative and, therefore, confirms Figure 1 results.

Table 1**Descriptive Statistics I**

At-the-Money Implied volatilities and FTSE100 index both level and log differences

Panel I: Level		
	ATM Implied Volatility	FTSE100 Index
Mean	12.4831	5178.7524
Median	12.5500	5040.5000
Minimum	7.5000	4292.0000
Maximum	18.4500	6385.5000
Panel II: Logarithmic 1st Differences		
	ATM Implied Volatility	FTSE100 Index
Mean	0.0000	-0.0005
Median	0.0032	-0.0005
Minimum	-0.3426	-0.0313
Maximum	0.1557	0.0293
Standard Deviation	0.0415	0.0069
Skewness	-2.3094	0.2559
Kurtosis	15.9252	1.9701
Panel III: Correlation		
	ATM Implied Volatility	FTSE100 Index
ATM Implied Volatility	1.0000	-0.4841
FTSE100 Index	-0.4841	1.0000

However, Table 2 displays descriptive statistics of the implied volatilities for the out of the money put and call options weekly data from 2004-2006 for the months of March and October. The mean level of the implied volatility is about 17% annually, whereas the log difference mean is zero. The Maximum level of the volatility is 41%, with a minimum level of 7.5 %. However, there is noticeable non-normality in the data, which can be noticed from the values of skewness and kurtosis for log differences of the implied volatility.

Table 2
Descriptive Statistics II

	OTM-PutCall Volatility(Level)	OTM-PutCal Implied Volatility (LN)
Mean	16.8708	0.0002
Median	15.5610	0.0289
Minimum	7.5621	-1.4594
Maximum	41.3337	0.1976
Standard Deviation	6.4497	0.1632
Skewness	0.8703	-5.7609
Kurtosis	3.2597	36.1365

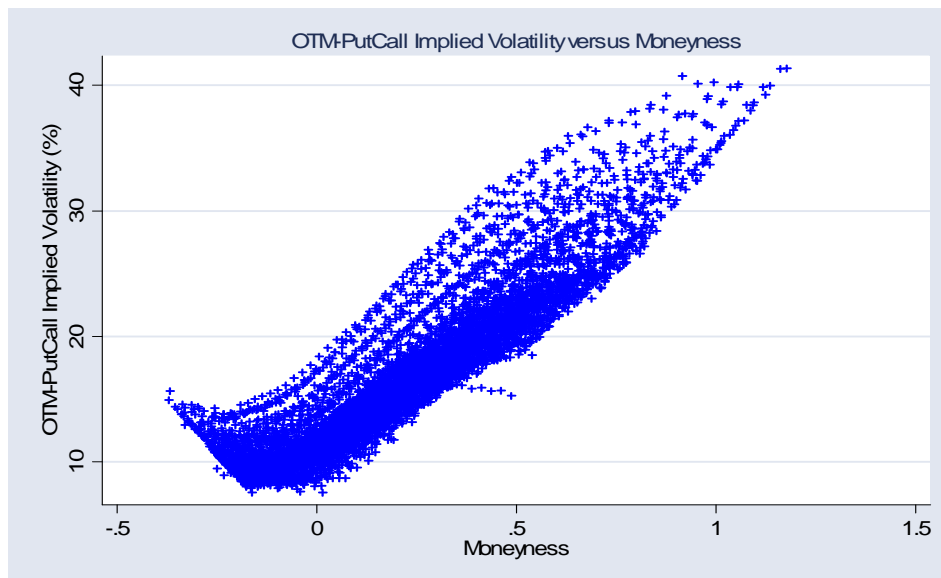


Figure 2 OTM-Put Call implied volatility versus moneyness

Figure 2 shows a plot of OTM put and call implied volatilities against moneyness level from 2004 to 2006 for the months of March and October. There is an almost one to one relation between the implied volatilities and moneyness level, i.e. most of the data lie on a forty-five degree angle. This evidence supports the monotonic relation between implied volatility and the price of the option. As the volatility increases, the price of the option increases.

We obtained two important results, consistent with previous research. First, the underlying

in non-normally distributed which is stock index; as a result, it does not follow lognormal distribution. This non-normality, therefore, causes non-constant volatility; thus, the assumption of the Black-Scholes model of constant volatility fails. Second, there is a one-to-one relationship between volatility and moneyness, which means that the price of the option has monotonic relationship with volatility. This relation is important for investors who hold options who will just observe volatility; if volatility increases, there is high probability that out of the money options will end up in the money and vice versa.

4 Empirical Results

4.1 Results of the Four Estimated Models for Volatility Surfaces

Table 3 reports results for estimated implied volatility surfaces in the months March and October for years 2004, 2005, and 2006. These surfaces are estimated with four parametric models by using nonlinear optimization technique. Panel I of Table 3 shows results for Model 1, the constant volatility model. The goodness of fit criterions IVRMSE and %RMSE ranges from minimum 0.046 (0.046) to maximum 0.078 (0.067) and, therefore, presents very high values. Figure 3 shows the visual display of average implied volatility surface (March, 2004), which corresponds to model 1. As can be seen in Figure 3, the implied volatility surface is flat; therefore, implied volatility is constant at every point of the surface. Nonetheless, later on we use this constant volatility model as a reference model, and merely for comparison purposes.

Table 3

Results for estimated Implied Volatility Surfaces (2004-2006)

Panel-I:		Model 1			$I(m, \tau) = \beta_0 + \varepsilon$						
Date	β_0				<i>IVRMSE</i>	<i>%RMSE</i>					
Mar-2004	0.2082				0.0775	0.0640					
Mar-2005	0.1468				0.0562	0.0639					
Mar-2006	0.1529				0.0458	0.0501					
Oct-2004	0.1752				0.0698	0.0666					
Oct-2005	0.1633				0.0627	0.0689					
Oct-2006	0.1631				0.0493	0.0461					
Panel-II:		Model 2			$I(m, \tau) = \beta_0 + \beta_1 m + \beta_2 m^2 + \varepsilon$						
Date	β_0	β_1	β_2		<i>IVRMSE</i>	<i>%RMSE</i>					
Mar-2004	0.1674	0.2156	0.0077		0.0250	0.0174					
Mar-2005	0.1098	0.1671	0.0794		0.0123	0.0110					
Mar-2006	0.1221	0.1466	0.0669		0.0110	0.0080					
Oct-2004	0.1309	0.1850	0.0679		0.0164	0.0159					
Oct-2005	0.1256	0.1810	0.0614		0.0141	0.0119					
Oct-2006	0.1285	0.1888	0.0006		0.0105	0.0073					
Panel-III:		Model 3					$I(m, \tau) = \beta_0 + \beta_1 m + \beta_2 m^2 + \beta_3 \tau + \beta_4 \tau m + \varepsilon$				
Date	β_0	β_1	β_2	β_3	β_4		<i>IVRMSE</i>	<i>%RMSE</i>			
Mar-2004	0.1570	0.1756	0.0362	0.0186	0.0827		0.0221	0.0152			
Mar-2005	0.1024	0.1535	0.0787	0.0163	0.0384		0.0100	0.0112			
Mar-2006	0.1127	0.1187	0.0937	0.0185	0.0441		0.0068	0.0063			
Oct-2004	0.1223	0.1600	0.0837	0.0182	0.0501		0.0139	0.0148			
Oct-2005	0.1218	0.1491	0.0854	0.0073	0.0858		0.0126	0.0106			
Oct-2006	0.1200	0.1716	0.0200	0.0198	0.0245		0.0076	0.0067			
Panel-IV:		Model 4					$I(m, \tau) = \beta_0 + \beta_1 m + \beta_2 m^2 + \beta_3 \tau + \beta_4 \tau m + \beta_5 \tau^2 + \varepsilon$				
Date	β_0	β_1	β_2	β_3	β_4	β_5	<i>IVRMSE</i>	<i>%RMSE</i>			
Mar-2004	0.1438	0.1818	0.0319	0.0938	0.0739	-0.0666	0.0214	0.0155			
Mar-2005	0.0997	0.1552	0.0772	0.0313	0.0345	-0.0130	0.0100	0.0115			
Mar-2006	0.1137	0.1185	0.0941	0.0133	0.0442	0.0046	0.0068	0.0063			
Oct-2004	0.1271	0.1588	0.0838	-0.0092	0.0525	0.0258	0.0138	0.0144			
Oct-2005	0.1192	0.1503	0.0853	0.0230	0.0823	-0.0152	0.0126	0.0106			
Oct-2006	0.1194	0.1717	0.0198	0.0230	0.0243	-0.0031	0.0076	0.0068			

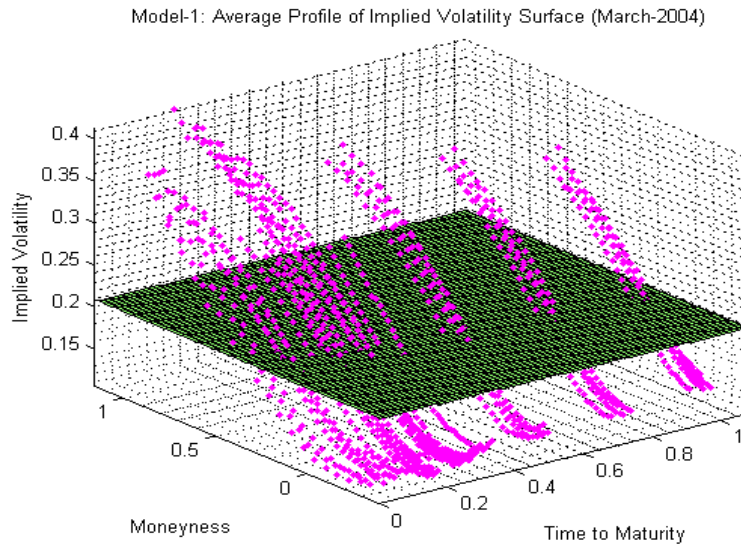


Figure-3 Estimated implied volatility surface with Model 1 for March 2004

Panel II of Table 3 displays results for model 2, the quadratic smile (skew) model, which includes additional parameters for moneyness slope and moneyness curvature. As can be noticed from goodness of fit criteria IVRMSE and %RMSE decreases for all six months ranges from minimum 0.010 (0.007) to 0.025 (0.017), improved by 0.036(0.039) to 0.053 (0.050) points in comparison to model 1. Figure 4 shows the corresponding 3D graph for the March 2004 implied volatility surface. As can be seen, the pink dots present the observed term structure of the smile (also called strings), therefore, skewed in moneyness level, and declining in term structure. Figure 4 presents two important implications. First, as moneyness increases, volatilities also increases, i.e. a monotonic relation. This finding is consistent with previous research. Secondly, we observe high volatilities for options with shorter time to maturities than options with longer time to maturities. Nevertheless, the skewed implied surface of model 2 is contrary to the Black-Sholes (1973) constant volatility assumption.

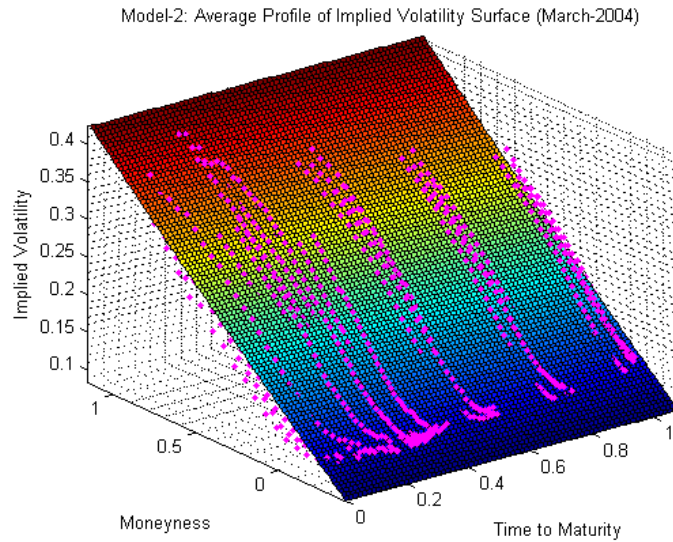


Figure-4 Estimated implied volatility surface with Model 2 for March 2004

Panel III of Table 3 reports results for implied volatility surfaces estimated with model 3, which add on model 2 one additional parameter for combined effect of moneyness and time dimension. As can be observed from the goodness of fit criteria IVRMSE and %RMSE decreases for all six months with ranges from minimum 0.007 (0.006) to 0.022 (0.015), improved by 0.003 (0.001) to 0.0032 (0.002) points in comparison to model 2 (smile model). Figure 5 displays a 3D graph of implied volatilities for March 2004. As can be observed, the implied surface is now tighter to the observed data, with the last corner of the surface stretched upward in comparison to the Figure 4 surface. Figure 5, therefore, suggests that there is a term structure effect in the implied volatilities, as we also noticed in the results of panel III of Table 3.




Figure-5 Estimated implied volatility surface with Model 3 for March 2004

Panel IV of Table 3 represents the results for model 4 which adds one additional parameter to model 3 to capture the curvature effect of the term structure in volatilities. As can be noticed from goodness of fit criteria IVRMSE and %RMSE decreases for all six months with values show minimum 0.007(0.006) to 0.021 (0.015), improved by 0.00(0.00) to 0.001(0.001) points in comparison to model 3. If we compare the goodness of fit of model 4 with model 3, very small improvements can be noticed. Figure 6 presents the 3 D surface; we can notice a slight change in shape, and the last corner of the surface is now stretched somewhat further.

We conclude for Table 3 that model 1 cannot explain variations in the implied volatility surfaces, therefore, gives constant volatilities across both dimensions. However, model 4 explains most of the variations in the implied volatility surface. Thus, it fits to the surface very well. However, the difference with model 3 in explaining variations is very small. The criteria for goodness of fit show significant values, and, therefore, model 3 and model 4 generate smooth implied volatility surfaces. In case if we compare the IVRMSE and %RMSE values, as can be seen that relative root mean squared errors (%RMSE) throughout

presents minimum values in contrast to implied volatility root mean squared error (IVRMSE). Nevertheless, in our case %RMSE is the most suitable one, as we it assign greater weights to out-of-the-money options.




Figure-6 Estimated implied volatility surface with Model 4 for March 2004

4.2 Results for Principal Component Analysis of Implied Volatility Surfaces

Table 4 presents results for PCA on the average implied volatility surfaces for the March and October months for years 2004, 2005 and 2006; from column two to column three, we have results for the three most important principal components (PCs). The first PC is common to the whole surface; we call this parallel shift, the second PC is tilt, and the third PC corresponds to curvature. However, column four represents total variance of the surface explained by three PCs. Figure 7 demonstrates PCA for the March 2004 surface, and Figure 8 shows graphs for the three factor loadings¹. Hence, on average 77% variance of the implied volatility surfaces can be explained by just three principal components for all six surfaces.

¹ Appendix B include figures for PCA of the implied volatility surfaces

When comparing the first three PCs for each of these surfaces, we note that the first three PCs can explain surface variations of the March 2004 surface to the maximum and the October 2006 surface to the minimum. To summarize Table 4, the first three PCs can explain 69-82% of the variance in all six volatility surfaces. These results are consistent with the previous studies on the dynamics of smile surfaces, such as Cont and da Fonseca (2002), and Skiadopoulos, Hodges, and Clewlow (1999) and Roux(2006).

Table 4

Principal components in the implied volatility surfaces				
Period	1st PC	2nd PC	3rd PC	Total explained variance(%) by 1-3
Mar-2004	64.0%	10.4%	7.7%	82.1%
Mar-2005	53.6%	12.0%	10.0%	75.6%
Mar-2006	43.6%	31.3%	4.8%	79.7%
Oct-2004	55.5%	13.0%	8.5%	77.0%
Oct-2005	59.3%	10.9%	7.0%	77.2%
Oct-2006	46.6%	12.5%	10.6%	69.7%
Mean	53.8%	15.0%	8.1%	77.0%




Figure 7 Principal components of the implied volatility surface for March 2004

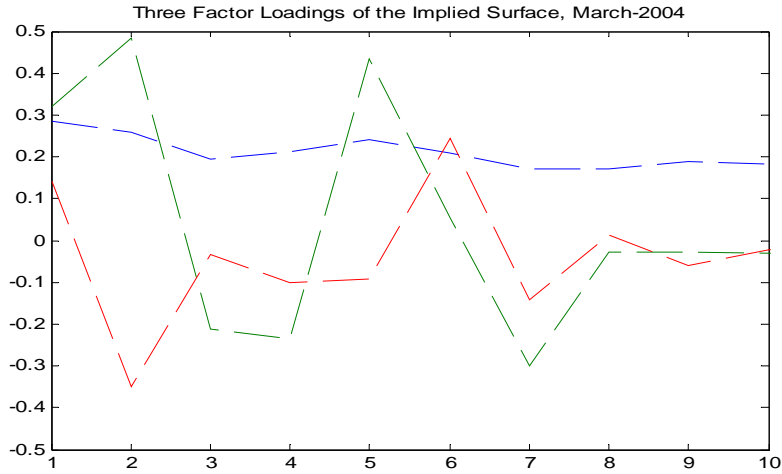


Figure 8 Three factor loadings of the Implied Volatility Surface for March 2004

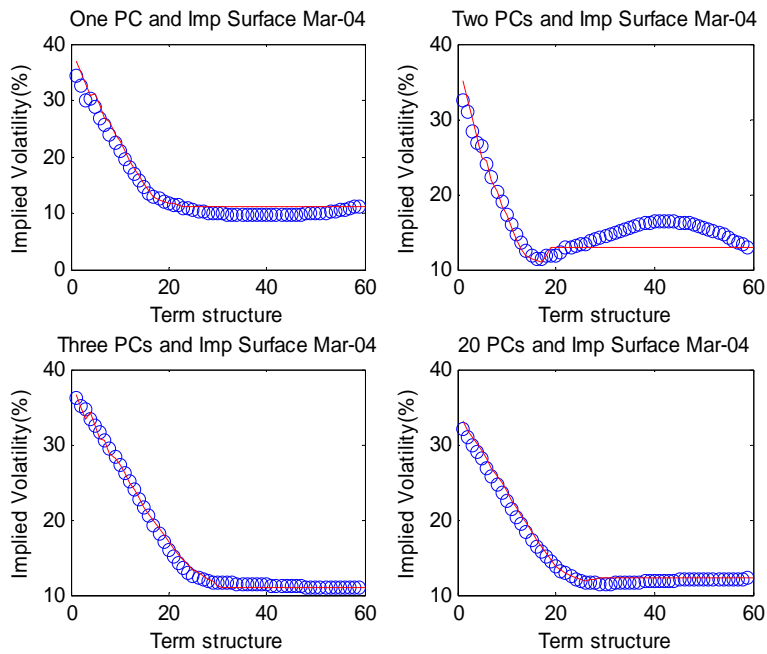


Figure 9 implied Surface versus PCs predicted Surface

To further investigate implied volatilities, we analyze each group of the maturities, i.e. 8-30 days, 30-60, 60-90, 90-150, 150-220 and 220 above. Table 5 documents results for different maturity groups. As can be seen, the first three principal components can capture variations in each group. Thus, we find two distinct patterns in all six groups. The first principal

component which corresponds to parallel shifts, can explain, on average, 60% variance in the first three groups (from 8-90 days). However, for the remaining three groups, results are completely in contrast; here most of the dynamics can be attributed to parallel shifts in the implied surfaces; therefore, about 90 percents of the variations can be explained by parallel shifts. However, the second PC which corresponds to tilt can explain about 18 percent of the variations in the implied volatility for the first three groups, while, for last three groups, this PC, on average, explains 5 percent variance. The third PC is important for the first three groups only that capture the curvature, which explains 9 percent of the variations in the implied volatility, and only 2 percent of the variance it captures in the last three groups.

Upon summarizing Table 5, we can conclude that shocks that are common to the entire surfaces that can be captured by the first PC are very important for the longer (90-above) and for the shortest (8-30) maturities; therefore, this factor is of sole importance for these surfaces. However, the other two factors are essential for (8-90 days) maturities.

Figure 9 presents a graph for actual values (Red) and fitted values (Blue) for the March 2004 implied volatility surface. As can be noticed, there are four subplots in Figure 9; the first subplot shows predicted implied volatility surface values merely from the first principal component, the second subplot in Figure 9 presents predicted values with the first two principal components, similarly third subplot with three PCs, however, the fourth subplot predicted implied surface with twenty principal components. It further confirms our results in Table 4 and Table 5 that the first three PCs can explain most of the dynamics in the implied volatility surface. Thus, the whole correlated system (implied volatility surface) can be reduced to just three risk factors; therefore, we can project implied surface by three principal components.

Table 5
Principal Components in the different maturity groups

Range	Month	1st PC	2nd PC	3rd PC	Total explained variance(%) by 1-3 PCs
8-30	Mar-2004	68.03%	14.35%	7.84%	90.23%
	Mar-2005	68.87%	14.97%	8.53%	92.37%
	Mar-2006	72.79%	14.62%	6.00%	93.42%
	Oct-2004	60.02%	25.98%	8.79%	94.79%
	Oct-2005	66.63%	15.33%	9.97%	91.94%
	Oct-2006	64.71%	19.01%	8.13%	91.85%
	Mean	66.84%	17.38%	8.21%	92.43%
30-60	Mar-2004	69.72%	16.86%	8.31%	94.89%
	Mar-2005	63.28%	14.16%	10.99%	88.43%
	Mar-2006	54.26%	15.23%	10.72%	80.22%
	Oct-2004	55.28%	22.17%	15.24%	92.68%
	Oct-2005	52.15%	22.91%	11.00%	86.07%
	Oct-2006	50.16%	20.39%	15.28%	85.83%
	Mean	57.48%	18.62%	11.92%	88.02%
60-90	Mar-2004	48.65%	21.24%	17.95%	87.84%
	Mar-2005	57.23%	20.09%	14.01%	91.33%
	Mar-2006	55.10%	22.22%	9.76%	87.09%
	Oct-2004	69.06%	11.99%	7.50%	88.55%
	Oct-2005	69.79%	14.25%	7.02%	91.07%
	Oct-2006	52.63%	22.47%	14.33%	89.43%
	Mean	58.75%	18.71%	11.76%	89.28%
90-150	Mar-2004	94.74%	2.72%	1.44%	98.90%
	Mar-2005	92.55%	5.62%	1.13%	99.30%
	Mar-2006	90.98%	7.40%	1.03%	98.81%
	Oct-2004	77.65%	15.71%	5.44%	98.81%
	Oct-2005	94.87%	3.44%	1.10%	99.41%
	Oct-2006	63.30%	18.64%	14.16%	96.10%
	Mean	85.68%	8.92%	4.05%	98.65%
150-220	Mar-2004	95.90%	1.43%	1.27%	98.60%
	Mar-2005	67.29%	29.78%	1.58%	98.67%
	Mar-2006	90.58%	6.01%	1.93%	98.52%
	Oct-2004	94.13%	2.73%	2.41%	99.28%
	Oct-2005	97.91%	0.96%	0.75%	99.63%
	Oct-2006	92.04%	4.08%	2.90%	98.99%
	Mean	89.65%	7.50%	1.80%	98.948%
220-Above	Mar-2004	93.17%	4.230%	1.65%	99.12%
	Mar-2005	97.82%	1.32%	0.60%	99.72%
	Mar-2006	92.52%	6.69%	0.63%	99.84%
	Oct-2004	95.73%	3.15%	0.90%	99.78%
	Oct-2005	96.58%	1.73%	1.10%	99.39%
	Oct-2006	93.10%	3.40%	1.91%	98.91%
	Mean	94.91%	3.43%	1.12%	99.46%

Conclusion

We estimate and obtain smooth implied volatility surfaces of FTSE100 index options with parametric structural models proposed by Dumas, Fleming and Whaley (1998). However, we make adjustments to these models by considering moneyness in forward price and out-of-the-money put and call options. Further, we study the dynamics of these implied volatility surfaces by using principal component analysis.

We estimate four parametric models; the constant volatility model 1 fails to capture variations in the implied volatility surfaces; it gives merely flat surfaces. However, model 2 which captures the smile (skew) effect does a good job, and captures most of the smile (skew) effects. When we estimate parametric models 3 and 4, which account for both term structure and smile (skew) effects, they perform better than the one dimensional smile model. Nonetheless, both models fit the data best and, thus, generate smooth and tighter volatility surfaces to the observed data. In the second part of the study, which includes principal component analysis, we, therefore, find that the first three principal components can explain about 69-82% of the variances in the implied volatility surfaces. When applying PCA to the maturity groups, we find that parallel shifts are evident and, hence, important for the shorter and longer maturities of option volatilities.

Implied volatility surface is very important in practice, because in option markets, we usually find every day a limited number of traded options, i.e. about 100 options with different strikes and maturities on the FTSE100 index. When a practitioner needs to quote options with different strikes or maturities, he is generally too restricted, which could be solved by implied volatility surface. With interpolation and extrapolation of observed implied volatilities and generates a whole surface of implied volatilities. Where, he can easily quote new options by taking estimates from implied volatility surfaces. Second, the implied


volatility surface can be used to manage exotic derivatives positions, which can be hedged with plain vanilla options such as European options. Third, the volatility index can easily be constructed from implied volatility surfaces, and, consequently, volatility derivatives can be priced. Fourth, risk can be managed for those securities which underly the FTSE100 index by using independent risk factors from the implied volatility surfaces.

References

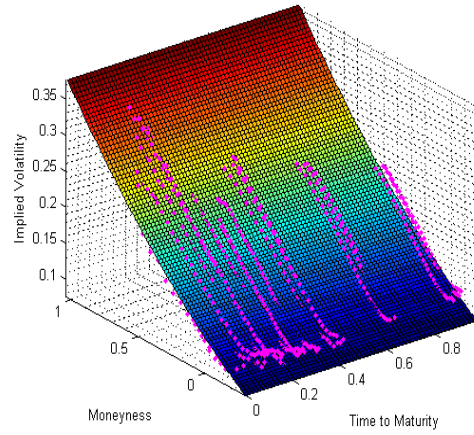
- Alentorn, Amadeo, 2004, Modelling the implied volatility surface: an empirical study for FTSE options, *Working paper, University of Essex, UK*.
- Alexander, Carol, 2001, Principal component analysis of implied volatility smiles and skews, *ISMA discussion paper, University of Reading, UK*.
- Black, F., and Scholes, M, 1973, The pricing of options and corporate liability, *Journal of political economy* 81, 636-654.
- Carr, P., Helyette, G., Delip, M., and Marc Y, 2001, The Fine structure of asset returns, an empirical investigation, *Journal of Business forthcoming*.
- Cassese, G., and Guidolin, M, 2006, Modelling the implied surface: Does market efficiency matter? An application to MIB30 index options” *International review of financial analysis* 15, 145-178.
- Christensen, B. J., and Prabhala, N. R, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50, 125-150.
- Christoffersen, P., and Jacobs, K., 2004, The importance of the loss function in option Valuation, *Journal of Financial Economics* 72, 291-318.
- Cont, R., and Fnseca, J. D, 2002, Dynamics of implied volatility surfaces, *Quantitative Finance* 2, 45-60
- Dumas, B., Fleming, J., and Whaley, R.E, 1998, Implied volatility functions: Empirical tests, *Journal of Finance* 53, 2059-2106.
- Fengler, M., Härdle, W., and Villa, C., 2003, The dynamics of implied volatility: A common principal components analysis approach, *Review of derivative research* 6, 179-202.
- Fleming, J, 1998, The quality of Market volatility forecasts implied by S&P100 index option prices, *Journal of Empirical Finance* 5, 317-345.
- Fung, W., and Hasieh, W, 1991, Empirical Analysis of implied volatility: stocks, bonds and currencies, *paper presented at 4th annual conference of financial options research center*, University of Warwick, 19-20, july, 1991, UK.
- Goncalves, S., and Guidolin, M, 2006, Predictable dynamics in the S&P 500 Index options implied volatility surfaces, *The Journal of Business* 79, 1591-1635.
- Gross, L., and N. Waltner, 1995, S&P 500 options: Put volatility smile and risk aversion, *Salomon Brothers, mimeo*.
- Heston, Steven, 1993, A closed form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327-343.
- Jackwerth, J., and Rubnstein, M, 1996, Recovering probability distributions from options price, *The Journal of Finance* 51, 1611-1631.
- Leland, Hayne, 1985, Option pricing and replications with transaction costs, *Journal of Finance* 40, 1283-1301.
- Merton, R.C, 1973, Theory of rational option pricing, *Bell Journal of Economics and Management sciences* 4, 141-183.
- Skiadopoulos, G., Hodges, S., and Clewlow, L, 1999, The Dynamics of the S&P 500 Implied Volatility Surface, *Review of Derivative Research* 3, 263-282.
- Roux, le. Martin, 2006, A long-term Model of the dynamics of the S&P500 implied volatility surfaces, *working paper ING institutional markets*.

Appendix A Figures of Implied Volatility Surfaces estimated with four Models for Oct-2005


Model-1: Average Profile of Implied Volatility Surface (October-2005)




Model-2: Average Profile of Implied Volatility Surface (October-2005)



Model-3: Average Profile of Implied Volatility Surface (October-2005)



Model-3: Average Profile of Implied Volatility Surface (October-2005)



Appendix B Figures for PCA of the Implied Volatility Surface for October 2005

