L Bonney, G Shannon and N Uys*

Modelling the Top40 volatility skew : A principal component analysis approach

1. INTRODUCTION*

Hedge Funds, Asset Managers and Traders that participate in option markets are all exposed to changes in implied volatility, as these changes directly affect the value of an option. In general, implied volatility $\sigma_K(t,s)$ is a forecast of an asset's return uncertainty over a specified future time period t, implied from the price of an option (strike level: K, market level: S).

Alexander (2001a&b) considers the risk implications of these 'changes in implied volatility'. She constructs a model that explains volatility change (particularly volatility skew) as a result of three dominant effects: trend, slope and convexity. Her model has some appeal in that these independent factors are readily understandable and their effect on an option's price can be calculated independently.

Decomposing volatility risk in this way benefits both option traders and risk managers. The intention is to explore these benefits in an emerging market setting by applying Alexander's implied volatility model to the less liquid South African Top40 option market.

2. VOLATILITY SKEW

The volatility change process mentioned above refers specifically to changes in implied volatility **relative** to that of an at-the-money option (an at-the-money option is defined to have a strike that equals the underlying futures price at inception of the option transaction concerned). This relative measure of implied volatility is often referred to as **volatility skew**.

Mathematically, volatility skew represents an *implied volatility difference* i.e., for a given vanilla option the associated skew can be defined to represent the difference between its implied volatility $\sigma_K(t,s)$ and that of an at-the-money option $\sigma_{ATM}(t,s)$, for the same expiry. We denote daily volatility skew for a fixed strike option as:

$$\sigma_K(t,s) - \sigma_{ATM}(t,s) \qquad \ldots (1)$$

Alexander's volatility model is unique in that she chooses to model implied volatility using *daily changes in volatility skew*, instead of daily implied volatility changes for fixed strike options (see: Kamal and Derman (1997), Skiadopoulos *et al.* (1999), Ané and Labidi (2001), Fengler *et al.* (2003)).

The daily change in volatility skew is expressed as:

$$\Delta\left(\sigma_K(t,s) - \sigma_{ATM}(t,s)\right) \qquad \ldots (2)$$

Her motivation for using **daily skew** change data is that it is less noisy than **daily implied volatility** change data. Moreover, skew data is highly stationary i.e., it has finite variance and is mean reverting. Stationarity is a necessary requirement for the successful application of many time-series-based models particularly the Principal Component Analysis (PCA) method chosen by Alexander (2001a&b).

3. THE ROLE OF PRINCIPAL COMPONENT ANALYSIS

An observable characteristic of skew change data is its high degree of collinearity. That is, there are common sources of information driving the daily evolution of skew change, across all strikes. In constructing the skew model it is useful to have the most **important** and the most **independent** of these common sources of information exposed. This would then explain the daily changes in skew, for any given strike, using the smallest number of relevant independent variables.

Principal Component Analysis is a mathematical technique that can be used to construct such a framework of independent factors for a given set of sufficiently correlated data. In fact, PCA has been successfully applied to highly correlated fixed-income data, to model the yield curve. The results show that the dynamics of the yield curve can be explained by three dominant independent components i.e., the 'trend', 'slope' and 'convexity' effects. Interestingly enough, Alexander's PCA methodology produces a volatility skew model whose variability is dependent on these same factors.

*Respectively, Programme in Advanced Mathematics of Finance, School of Computational and Applied Mathematics, University of the Witwatersrand, Johannesburg, Republic of South Africa; Derivative Research: GES London; Programme in Advanced Mathematics of Finance, School of Computational and Applied Mathematics, University of the Witwatersrand, Johannesburg, Republic of South Africa. The authors are indebted to Mr Manoshan Pillay of Cadiz FSG for help with the gathering and construction of the data and for his input at the start of the project. Any errors or omissions, however, remain the authors' responsibility.
Email: Lisa.Bonney@students.wits.ac.za; grantp.shannon@gmail.com; Nadia.Uys@wits.ac.za

4. SALIENT FEATURES OF PRINCIPAL COMPONENT ANALYSIS

Principal Component Analysis can be expressed as a maximisation problem that results in the transformation of k correlated variables into a set of p independent components or factors. The variance-covariance matrix (or correlation matrix) of these components closely matches, if not equals, that of the original time series data, where $p \leq k$. If the information content of the p components, as measured by the correlation matrix, say, still **closely** resembles that of the original data structure containing k variables, then the p principal components can be used to replace the initial dataset. In this way the dimensionality of the problem may be reduced without introducing any significant information loss.

A PCA model produces output in the form of eigenvalues and the eigenvectors (see Sections 4.1 and 4.2). The required PCA model input in this study is a correlation matrix generated from the 'daily skew change' time series data.

4.1 Eigenvalues

An *eigenvalue* is an estimate of the variance explained by a particular component. If most of the total population variance can be attributed to the first two or three ordered principal components, then these components can replace the original k variables without much error or loss of information.

4.2 Eigenvectors

An *eigenvector* is typically a column vector comprising the weights for each strike, associated with a given principal component. These weights are conceptually similar to the coefficients in a multivariate regression. Each eigenvector has a corresponding eigenvalue.

Inspecting the eigenvectors is useful when attempting to interpret the principal components.

4.3 The principal components

The principal components are column vectors that resemble time series data in a regression model. These components are, by construction, not correlated with each other and are typically ordered by decreasing explanatory power. As such, they each represent a column of 'synthetic' time series data. Consequently, the components may not always be interpretable into known market variables. Their interpretation in the context of the Top40 skew model however, is discussed in more detail in Section 6.

Each column of daily skew change in the original input dataset can be expressed as a linear combination of the independent principal components, weighted by the coefficients in the corresponding eigenvector.

5. DATA DESCRIPTION

A Cadiz FSG volatility database is used to model the Top40 skew surface. This dataset comprises cleaned daily implied volatility data $\sigma_K(t,s)$ on near-dated Top40 options covering seven fixed strikes: K $\in$ {16000, 16500, 17000, 17500, 18000, 18500, 19000}, over the six month period 2nd February 2006 to 24th July 2006.

Applying Equation (2) to each strike generates daily skew changes as required by Alexander's model. The initial implied volatility database includes the near-dated futures mark-to-market level i.e., a proxy for the at-the-money strike as well as the corresponding at-the-money implied volatility level. As a noise reduction measure, consideration was given to building the Top40 volatility surface using weekly skew change data. There was, however, not enough weekly data to satisfy the initial model requirements.

6. METHOD

The construction of the Top40 skew model follows Alexander's methodology in a step-by-step fashion, starting with the calculation of the correlation matrix:

6.1 The correlation matrix

Define $Y_{T\times k}$ to be the matrix of daily skew changes for all seven strikes (i.e., the number of strike levels under consideration is given by: k = 7 and T represents the number of days for the time period under investigation) and calculate its symmetric correlation matrix $V_{k\times k}$. This is presented in Table 1.

Recall that the objective of PCA is to find the eigenvectors and eigenvalues of the input correlation matrix V.

6.2 The eigenvalues and eigenvectors

This eigenvector-eigenvalue decomposition of V is represented mathematically as:

$$V = W \Lambda W^T \qquad \text{... (3)}$$

where

$$\Lambda_{k\times k} = \text{diag}(\lambda_1, \lambda_2, \ldots, \lambda_k) \qquad \text{... (4)}$$

Λ is a diagonal matrix of eigenvalues relating to the eigenvector matrix W, with:

$$W_{k\times k} = (w_1, w_2, \ldots, w_k) \qquad \text{... (5)}$$

The m^{th} column of W, is the $k \times 1$ eigenvector corresponding to the m^{th} eigenvalue λ_m, denoted:

$$w_m = (w_{1m}, w_{2m}, \ldots, w_{km})^T \quad \ldots (6)$$

Note that the column labelling has been chosen such that $\lambda_1 > \lambda_2 > \cdots > \lambda_m > \cdots > \lambda_k$.

Table 2 illustrates the output from the PCA model which applies the decomposition in Equation (3) to correlation matrix V.

The decomposition shows the eigenvalues and eigenvectors in decreasing order of explanatory power. These eigenvectors are represented as columns labelled $w_1, w_2, \ldots, w_7$ in Table 2. The first eigenvalue corresponds to eigenvector w_1, the second eigenvalue corresponds to eigenvector w_2, etc.

6.2.1 Interpreting the eigenvectors

The weights as represented by the elements of each eigenvector explain the effect that the relevant component has on driving skew change, for the given strikes.

Analysing these coefficients can help label the principal components. Alexander explains that the elements of the first three eigenvectors are of particular interest in this labelling process, as their corresponding components explain the bulk of the variability in skew change data (see Section 6.4).

Using an S&P implied volatility dataset as input to the skew model, Alexander defines the first three components as the trend, slope and convexity effects. The subsections that follow discuss these effects within the context of the Top40 skew model.

At this point it is worth making a brief mention of the occasional lack of consistency when comparing the behaviour of the eigenvector weights of the first three components in Alexander's S&P-based model to those of the Top40 model: this is largely due to liquidity issues in the Top40 option market data. Moreover, the S&P model incorporates many more liquid strikes, creating depth in Alexander's results.

6.2.1.1 Trend component

In Table 2 most of the weights in the first eigenvector w_1 are positive. The lowest weights are concentrated around the low strikes and they tend to increase in value as the strikes increase. This behaviour is confirmed in Alexander's S&P study.

Table 1: Correlation matrix

Strikes	**16000**	**16500**	**17000**	**17500**	**18000**	**18500**	**19000**
16000	1,00	0,52	0,02	0,34	0,11	0,09	0,01
16500	0,52	1,00	0,28	0,39	0,22	0,28	-0,02
17000	0,02	0,28	1,00	0,55	0,59	0,63	-0,05
17500	0,34	0,39	0,55	1,00	0,46	0,65	0,01
18000	0,11	0,22	0,59	0,46	1,00	0,54	-0,04
18500	0,09	0,28	0,63	0,65	0,54	1,00	0,03
19000	0,01	-0,02	-0,05	0,01	-0,04	0,03	1,00

Table 2: Eigenvalue – Eigenvector decomposition

Contribution	**Eigenvalues**	**Strikes**	**Matrix of Eigenvectors**						
			w_1	w_2	w_3	w_4	w_5	w_6	w_7
42,66%	2,99	16000	0,22	-0,71	-0,01	0,05	0,48	-0,23	-0,4
19,18%	1,34	16500	0,33	-0,53	-0,06	0,21	-0,67	0,26	0,19
14,40%	1,01	17000	0,46	0,3	-0,04	0,13	-0,32	-0,73	-0,23
7,67%	0,54	17500	0,48	-0,05	0,06	-0,54	0,24	-0,12	0,63
7,24%	0,51	18000	0,42	0,25	-0,04	0,7	0,39	0,26	0,21
4,78%	0,33	18500	0,47	0,24	0,09	-0,37	-0,03	0,52	-0,55
4,05%	0,28	19000	-0,01	-0,03	0,99	0,11	-0,04	-0,05	0,02

Positive (negative) movements in the first component typically result in positive (negative) changes in the volatility skew - for each strike. In this way, component one influences the general level of the volatility skew across all strikes. It causes the skew change for each strike to trend up or down relative to the weights in w_1, as it moves up or down. The first factor is thus labelled the **trend** component of the volatility skew.

6.2.1.2 Slope component

Inspecting the coefficients of the second eigenvector w_2 shows that the signs of low strike options are negative, whilst they are predominantly positive or close to zero for the higher strike options. A given change in component two has the opposite effect on high and low strike skew changes, resulting in a **slope** adjustment in the skew profile.

Low strike options are impacted the most by changes in component two, given that their weights are larger in absolute terms than those of the higher strikes. This is reinforced empirically as lower strike options have the highest risk in the equity option markets.

Alexander's S&P results clearly show the same relationship, with the **absolute values** of the weights in the second component going from high (for the low strikes) to low (for the high strikes). The signs of the weights are dependent on the way in which the orthogonal basis vectors are orientated in the PCA model – this may vary depending on the PCA software used.[2]

6.2.1.3 Convexity component

In Alexander's liquid S&P dataset, the coefficients of the third eigenvector, w_3, have the same signs at both ends of the strike range and opposite signs for options in the middle of the strike range, resulting in changes in **convexity** of the skew surface. The Top40 weights, partly match this profile. Whatever the lack of consistency when comparing these Top40 and S&P skew model weights, the overall behaviour of component three for the Top40 model is found to be explained by the convexity effect.

6.2.2 Interpreting the eigenvalues

The eigenvalues in Table 2 indicate that component one, i.e., the trend effect explains 42% of the variability in skew changes, the slope 19% and convexity an additional 14%.

In total, the first three components explain 76,24% of the variability in skew changes. The biggest contributor to skew change is the trend effect. Comparatively, Alexander's S&P model explains approximately 90% of the total variability in the S&P volatility skew changes.

The difference between the 76% and the 90% reflects in part the information loss that can occur between a liquid market skew dataset and a less liquid emerging market equivalent.

6.3 The principal components

The next part of the model building process is to generate the principal component matrix $P_{T\times k}$ governed by the following relationship:

$$P = YW \quad \text{... (7)}$$

where

$$P = (P_1, P_2, \ldots, P_k) \quad \text{... (8)}$$

and

P_i is a $T \times 1$ column vector which represents the i^{th} principal component, for $i = 1, 2, \ldots, k$.

Recall:

$$Y = (Y_1, Y_2, \ldots, Y_k) \quad \text{... (9)}$$

where, $Y_1, Y_2, \ldots, Y_k$ are the $T \times 1$ column vectors of Y, then the m^{th} principal component of the system is given by:

$$P_m = Yw_m = w_{1m}Y_1 + w_{2m}Y_2 + \cdots + w_{km}Y_k \quad \text{... (10)}$$

Due to the column labelling in W, the principal components have been ordered so that P_1 belongs to the first and largest eigenvalue λ_1, P_2 belongs to the second largest eigenvalue λ_2 and so on.

Since:

$$W^T = W^{-1} \quad \text{... (11)}$$

i.e., the transpose of the eigenvector matrix equals its inverse, we can thus write:

$$Y = PW^T \quad \text{... (12)}$$

Hence, for $i = 1, 2, \ldots, k$:

$$Y_i = w_{i1}P_1 + w_{i2}P_2 + \cdots + w_{ik}P_k \quad \text{... (13)}$$

[2] The PCA software used in this study was a dynamic link library (dll) that interfaces with Excel: built by Grant Shannon.

As mentioned in Section 4.3, each column vector of the original skew change dataset may be written as a linear combination of the orthogonal principal components $P_1, P_2, \ldots, P_k$, with coefficients given by the appropriate row in the $k \times k$ eigenvector matrix W. Stated slightly differently, the original implied volatility skew change dataset can now be reconstructed exactly using all seven principal components.

6.4 Dimensionality reduction

Recall that Principal Component Analysis attempts to reduce the dimensionality of the full skew model (seven components) without impacting too much on the accuracy of the resultant implied volatility estimates. Using the eigenvalues as our guide, only those variables that explain a large enough percentage of the variability in the daily skew change dataset are chosen.

The results in Section 6.2.2 show that overall, 76% of the variability in the skew change data, is explained by the first three principal components. To be consistent with Alexander's S&P model, these first three components are used to model Top40 skew change $Y^*_{T \times k}$, i.e., the number of selected principal components is given by: p = 3. Mathematically, this is represented as:

$$Y^* \equiv \Delta(\sigma_K - \sigma_{ATM}) = P_1 w_1^T + P_2 w_2^T + P_3 w_3^T \quad \text{... (14)}$$

The estimated daily skew deviations generated from Equation (14) are translated into fixed strike implied volatilities by calculating the skew for each strike on a given day and then adding this to the corresponding ATM volatility level. This generates daily implied volatility approximations for each of the seven strikes for the full six month period under investigation.

The Top40 implied volatility surface estimated in this way can be compared to the original Top40 implied volatility dataset.

Tables 3 and 4 illustrate data **subsets** from the 'estimated' and 'original' implied volatility surfaces.

Note that the addition of more components to improve the model 'fit' will not impact on the effect that trend, slope and convexity have on skew change.

Moreover, the reduction in the dimensionality of the model, down to three orthogonal components, plays an increasingly significant role as more strikes are included in the description of the original volatility surface.

Alexander's next step is to introduce a functional form for each of the dominant principal components. In essence this adjustment creates a link between the 'market regime', the behaviour of the components and ultimately the behaviour of the volatility skew.

Table 3: Implied volatility data *SUBSET* – *ORIGINAL* volatility surface

Date	16000	16500	17000	17500	18000	18500	19000
2006/05/17	27,52	22,52	21,93	23,18	23,14	20,22	19,40
2006/05/18	25,27	22,21	21,66	22,42	21,29	20,00	19,81
2006/05/19	26,49	26,15	25,75	24,80	23,80	23,49	22,48

Table 4: Implied volatility data *SUBSET* – PCA model *ESTIMATE*

Date	16000	16500	17000	17500	18000	18500	19000
2006/05/17	25,61	25,47	24,81	21,86	20,77	19,74	19,71
2006/05/18	24,31	24,15	24,21	20,93	20,16	19,15	20,15
2006/05/19	26,68	26,61	27,78	23,66	23,72	22,37	22,77

7. MODEL ADJUSTMENT

Alexander (2001a&b) observes that changes in fixed strike volatilities are affected by movements in the principal components, as the underlying asset price changes.

She chooses to model each principal component P_i as a function of the underlying price change ΔS, over time, and includes a time-varying parameter $\gamma_i(t)$ (the gamma coefficient).

The resultant linear equation is:

$$P_i(t) = \gamma_i(t)\Delta S(t) + \varepsilon_i(t) \quad \text{... (15)}$$

for $i = 1, 2, 3$, where ε_i are i.i.d $N(0,1)$.

The gamma coefficients thus represent the sensitivities of the principal components to the underlying asset price movements. In particular γ_1 affects the trend or parallel shifts in the skew, γ_2 the slope, i.e., the range of the skew whether it widens or narrows as the underlying moves, and γ_3 the convexity.

Once these coefficients have been estimated on a daily basis, the Top40 skew surface can be regenerated by replacing the old components in Equation (14) with the new ones in Equation (15).

7.1 Estimating the gamma coefficients

The gamma coefficients are estimated by considering the covariance over time between the principal components and the corresponding daily change in the underlying. Recall that the principal components are by construction, uncorrelated with each other. The covariance is expressed mathematically as:

$$\begin{aligned} cov_t(P_{i,t}, \Delta S_t) &= cov_t(\gamma_{i,t}\Delta S_t + \varepsilon_{i,t}, \Delta S_t) \\ &= \gamma_{i,t} cov_t(\Delta S_t, \Delta S_t) \\ &= \gamma_{i,t} var_t(\Delta S_t) \end{aligned} \quad \text{... (16)}$$

Thus for $i = 1, 2, 3$,

$$\gamma_{i,t} = \frac{cov_t(P_{i,t}, \Delta S_t)}{var_t(\Delta S_t)} \quad \text{... (17)}$$

Ideally, the conditional covariance $cov_t(P_{i,t}, \Delta S_t)$ and the conditional variance $var_t(\Delta S_t)$ would be estimated using a bivariate generalised autoregressive conditional heteroscedasticity model, i.e., GARCH(1,1). For simplicity however, Alexander approximates this GARCH process with an exponentially weighted moving average (EWMA) formulation, using a smoothing constant $\lambda = 0{,}94$. The resulting EWMA conditional variance and covariance representations are:

$$\hat{\sigma}_t^2 = var_t(\Delta S_t) = (1-\lambda)\Delta S_{t-1} + \lambda\hat{\sigma}_{t-1}^2 \quad \text{... (18)}$$

With the first point $\hat{\sigma}_0^2$ estimated

$$\begin{aligned} &cov_t(P_{i,t}, \Delta S_t) = \hat{\sigma}_{P_i,\Delta S;t} \\ &= (1-\lambda)P_{i,t-1}\Delta S_{t-1} + \lambda\hat{\sigma}_{P_i,\Delta S;t-1} \end{aligned} \quad \text{… (19)}$$

with

$$cov_0(P_{i,t}, \Delta S_t) = \hat{\sigma}_{P_i,\Delta S;0} \quad \text{... (20)}$$

for each principal component P_i, $i = 1, 2, 3$.

These variance and covariance estimates are then used to generate daily $\gamma_{i,t}$ estimates: $i = 1, 2, 3$. A sample of the daily evolution of these time-varying parameters is provided in Table 5.

Applying the coefficients in Table 5 to a corresponding sample of Top40 futures price change data generates daily estimates of the principal components. Substituting these into Equation (14) produces daily skew change approximations for the **gamma-adjusted model.**

From an implied volatility surface 'fitting' perspective the **gamma-adjusted model** is not necessarily as robust as the **unadjusted model** as it includes principal component estimates instead of the actual values.

Despite this, the strength of the gamma-adjusted model remains in its ability to decompose implied volatility shocks into trend, slope and curvature effects, for given futures price change scenarios. Moreover, monitoring the daily evolution of the gamma coefficients provides insights into the prevailing market sentiment or market regime.

7.2 Analysing the gamma coefficients

In fact, the results show that with shifts in market regime the gamma estimates become increasingly susceptible to sign changes, particularly the second and third coefficients.

Consider for example the period 2nd Feb 2006 - 30th Mar 2006: both γ_2 and γ_3 are negative suggesting that the range of the skew narrows as the underlying increases with most of the movement in the skew being generated by the low strike volatilities. From 30th Mar 2006 - 11th May 2006, however, γ_2 is positive, widening the range of the skew as the index moves up - most of the skew change is attributable to the high strike volatilities.

In general, the combination of the second component and second gamma coefficient drive the expansion and contraction of the skew as the underlying changes, i.e., when the index moves are negatively correlated with the second principal component γ_2 will be negative and the range of the skew will narrow as the index moves up and widen as the index moves down. Most of the skew movements under this scenario are created by the lower strikes.

When the index moves have zero or positive correlation with the second principal component, the range of the skew narrows as the index moves down and widens as the index moves up, where high strike volatilities move the most. If $\gamma_2 = 0$, the skew will shift parallel as the index moves, i.e., all fixed strike implied volatilities will move by the same amount as the index moves.

Table 5: Daily gamma estimates – Data subset

Date	Gamma 1	Gamma 2	Gamma 3
2006/02/23	0,0045685	-0,0021160	-0,0010195
2006/02/24	0,0044650	-0,0019866	-0,0009995
2006/02/27	0,0041484	-0,0016871	-0,0008789
2006/02/28	0,0041256	-0,0015961	-0,0008659

Overall, for the six month period under consideration, the first coefficient typically maintains a positive value between 0 and 0,005. The second gamma coefficient ranges between -0,003 and +0,003 and the third occupies an even narrower range between -0,0015 and +0,0015.

8. CONCLUSION

This study implements Alexander's S&P volatility skew model in the less liquid Top40 options market using Principal Component Analysis. The results show that emerging market volatility can be expressed as a function of three independent factors: trend, slope and convexity. Collectively they explain close to 80% of the variability in Top40 skew data.

The decomposition of volatility risk into these three independent factors enables, for example, the hedging of implied volatility risk to be considered at a component level and not necessarily at an individual strike level. Such a component-based hedge is effective in that it takes into account the correlation structure between implied volatilities of options spanning a range of strikes.

Moreover, the parametric adjustments made to the model allow for a regime-based approach to modelling implied volatility risk. That is, for given changes in the underlying price (reflecting a particular scenario or market regime) the model provides estimates of changes in the implied volatility skew in terms of trend, slope and convexity effects relative to each of these scenarios - this is useful in both risk management and trading.

REFERENCES

Alexander C. 2000(a). Principal component analysis of implied volatility smiles and skews. ISMA Centre Discussion Papers in Finance, University of Reading, UK.

Alexander C. 2000(b). Key market risk factors: Identification and applications. *The Q-Group Seminar on Risk*.

Alexander C. 2001(a). Principles of the skew. *Risk*, 14(1):29-32.

Alexander C. 2001(b). *Market models: A guide to financial data analysis.* John Wiley and Sons, New York.

Ané T and Labidi C. 2001. Implied volatility surfaces and market activity over time. *Journal of Economics and Finance*, 25(3):259-275.

Avellaneda M and Zhu Y. 1997. An E-ARCH model for the term structure of implied volatility of FX options. *Applied Mathematical Finance*, 4:81-100.

Cont R and da Fonseca J. 2001. Deformation of implied volatility surfaces: An empirical analysis. *Empirical Approaches to Financial Fluctuations*, 230-239.

Cont R and da Fonseca J. 2002. The dynamics of implied volatility surfaces. *Quantitative Finance*, 2(1):45-60.

Derman E and Kamal M. 1997. The patterns of change in implied index volatilities. Quantitative Strategies Research Notes, Goldman Sachs.

Derman E. 1999. Regimes of volatility: Some observations on the variation of S&P 500 implied volatilities. Quantitative Strategies Research Notes, Goldman Sachs.

Fengler MR. 2002. The phenomenology of implied volatility surfaces. Masters thesis. Department of Business and Economics. Humboldt University, Berlin.

Fengler MR, Härdle WK and Villa C. 2003. The dynamics of implied volatilities: A common principal components approach. *Review of Derivatives Research*, 6:179-202.

Fengler MR, Härdle WK and Schmidt P. 2002(a). The analysis of implied volatilities. In W. Härdle, T. Kleinow and G. Stahl eds, *Applied Quantitative Finance*, Springer, Chapter 6.

Fengler MR, Härdle WK and Schmidt P. 2002(b). Common factors governing VDAX movements and the maximum loss. *Journal of Financial Markets and Portfolio Management*, 16(1):16-29.

Golub G. 1989. *Matrix computations*. John Hopkins University Press.

Guidolin M and Goncalves S. 2006. Predictable dynamics in the S&P 500 index options implied volatility surface. *Journal of Business*, 79:1591-1635.

Hafner R and Wallmeier M. 2001. The dynamics of DAX implied volatilities. *International Quarterly Journal of Finance*, 1(1):1-27.

Johnson R and Wichern D. 1998. *Applied multivariate statistical analysis*. Prentice Hall.

Le Roux M. 2006. A long term model of the dynamics of the S&P 500 implied volatility surface. *North American Actuarial Journal*, 11(4):61-75.

Skiadopoulos G, Hodges S and Clewlow L. 1999. The dynamics of the S&P 500 implied volatility surface. *Review of Derivatives Research*, 3(3):263-282.

Sylla A and Villa C. 2000. Measuring implied volatility surface risk using principal components analysis. In J. Franke, W. Härdle and G. Stahl ed., *Measuring Risk in Complex Stochastic System.* Springer-Verlag, New York. 131-147.