FINANCIAL OPTIONS RESEARCH CENTRE

University of Warwick

The Dynamics of Implied Volatility Surfaces

George Skiadopoulos
Stewart Hodges
and
Les Clewlow

May 1998

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

FORC Preprint:1998/86

The Dynamics of Implied Volatility Surfaces*

George Skiadopoulos† Stewart Hodges‡ Les Clewlow§

May 1998

Abstract

Motivated by the papers of Dupire (1992) and Derman and Kani (1997), we want to investigate the number of shocks that move the whole implied volatility surface, their interpretation and their correlation with percentage changes in the underlying asset. This work differs from Skiadopoulos, Hodges and Clewlow (1998) in which they looked at the dynamics of smiles for a given maturity bucket.

We look at daily changes in implied volatilities under two different metrics: the strike metric and the moneyness metric. Since we are dealing with a three dimensional problem, we fix ranges of days to maturity, we pool them together and we apply the Principal Components Analysis (PCA) to the changes in implied volatilities over time across both the strike (moneyness) metric and the pooled ranges of days to maturity.

We find similar results for both metrics. Two shocks explain the movements of the volatility surface, the first shock being interpreted as a shift, while the second one has a Z-shape. The sign of the correlation of the first shock with percentage changes in the underlying asset depends on the metric that we look at, while the sign is positive under both metrics regarding the second shock.

The results suggest that the number of shocks, their interpretation and the sign of their correlation with changes in the underlying asset is the same for the whole implied volatility surface as it is for the smile corresponding to a fixed maturity bucket.

JEL Classification: G13

Keywords: Implied Volatility Surface, Number of Shocks, Interpretation, Correlation, Metrics, Principal Components Analysis.

---

*We would like to thank Russell Grimwood, Wojtek Krzanowski, Joao Pedro Nunes, Juan Carlos Mehia Perez and Chris Strickland for helpful discussions and comments. Part of this paper was written while the first author benefited from the Human Capital & Mobility Programme of the European Commission. Financial support from the Corporate Members of FORC is also gratefully acknowledged. Any remaining errors are our responsibility alone.

†Research Fellow at the Financial Options Research Centre, Warwick Business School University of Warwick, Coventry CV4 7AL, UK; Tel: 01203-528211, Fax: 01203-524167, e-mail: Georgios.Skiadopoulos@warwick.ac.uk

‡Director of the Financial Options Research Centre, Warwick Business School, University of Warwick, Coventry CV4 7AL, UK; Tel: 01203-523606, e-mail: S.D.Hodges@warwick.ac.uk

§Associate Research Fellow at the Financial Options Research Centre; e-mail: Les_Clewlow@compuserve.com

1 Introduction

Empirical evidence (see for example Rubinstein (1985), Sheikh (1991) and Heynen (1993)) shows that implied Black-Scholes (B-S) volatilities strongly depend on the maturity and strike of the European option under scrutiny and therefore in three dimensions we have a volatility surface which is not flat as the B-S (1973) model assumes. The former type of dependence is called the term structure of volatility and the latter type is called the smile-effect.

One of the ways that researchers tried to explain the non-flat implied volatility surface was by introducing stochastic volatility models (see Hull and White (1987), Johnson and Shanno (1987), Scott (1987), Wiggins (1987)). Effectively, by doing this, they introduced non-traded sources of risk and inevitably they incorporated the market prices of volatility something which does not allow arbitrage, but only equilibrium pricing. Moreover, these models do not fit the observed smiles well.

As a solution to this, Dupire (1992) and Derman and Kani (1997), proposed models similar in spirit, which allow for arbitrage pricing with stochastic volatility, with no need for any volatility risk premium to be specified. In order to do this, they used a methodology which is similar to that of the arbitrage arguments used in the interest rate literature (see Heath, Jarrow, Morton (1992)), i.e they took the current implied volatility surface as given and they assume a process for the forward variance of the form

$$\frac{dV_{K,T}(t,S)}{V_{K,T}(t,S)} = a_{K,T}(t,S)dt + \sum_{i=1}^{n} b^i_{K,T}(t,S)dW^i_t \tag{1}$$

where K is the strike price and T is the time to maturity. Their idea was that we can price and hedge standard and exotic options by Monte Carlo simulation where the process of the forward variance is required either implicitly (Dupire), or explicitly (Derman and Kani).

This "arbitrage pricing" methodology can be very promising if we specify correctly the process for the forward variance and more specifically if we answer the following three questions: (a) what is n, i.e. the number of shocks which affect the forward volatility surface, (b) what do these shocks look like. The interpretation of the shocks is necessary so that to be able to identify the function $b^i_{K,T}(t,S)$ and subsequently to estimate it. In fact, the simpler the interpretation, the easier it will be to identify the functional form of $b^i_{K,T}(t,S)$ and (c) what is the correlation between the shock of the process for the underlying asset and the shocks of the process for the forward variance since knowledge of the correlation is necessary when we simulate jointly the two processes.

In Skiadopoulos, Hodges and Clewlow (1998, S-H-C hereafter) the number of shocks turned out to be two, their interpretation in general was a shift for the first one and a Z-shape for the second, while the correlation of the futures price with the first shock was either positive or negative depending on the metric and the correlation with the second was positive. These results came from applying Principal Components Analysis (PCA) to the first differences of implied volatilities on the strike and moneyness metric for given ranges of days to maturity. In other words, they analysed the dynamics of the smile phenomenon, for individual maturity buckets[1].

[1]Hereafter, we will refer to this paper as the smile analysis.

On the contrary, in this paper we look at the dynamics of implied volatilities across *both* the strike (moneyness) metric and the ranges of days to maturity. In other words by doing a sort of pooling we look at the whole implied volatility surface and we try to answer the three questions[2]. This is the main contribution of the paper since no other researcher so far, has dealt with these questions.

Answering (a), (b) and (c) is necessary in order to cope with the vega risk in portfolios which consist of derivatives with many different expiry dates and strikes. The difference with S-H-C is that there, by construction, different shocks were moving the different slices of the implied volatility surface, while here by pooling the maturity buckets, the same shock moves the whole implied volatility surface.

The paper is organised as follows : in the second section we describe the data that we use and the filtering that we applied to it. In the third section we describe briefly the PCA analysis. In the fourth and fifth sections we apply the technique on the first differences of the whole implied volatility surface under the strike and moneyness metric, respectively. We test for how many "pooled" components we should retain, we try to interpret them and we look at how they affect the volatility surfaces across the different maturity buckets. In the sixth section we calculate the correlations between the changes of the "pooled" principal components and the underlying asset price. We end up with the conclusions.

2 The Data Set and the Filtering Method

The data that we use are daily data on futures options on the Standard and Poor index (S&P 500) from the Chicago Mercantile Exchange (CME) for the years 1992-95.

The primary source database for this study is the transaction report "Stats Database", compiled daily by CME. This database contains the following information for each option traded during a day: the date, the style (call or put), the options and futures expiration month, the exercise price, the number of contracts traded, the opening, closing, low and high future's and option's price, the opening, closing, low and high bid-ask future's and option's price and the settlement price. We are going to use the closing options and futures prices, since we regard those as containing more information than the opening, low, or high ones.

The futures contracts have maturities every March, June, September and December and their last trading day is the business day prior to the third Friday of the contract month. The options have maturities March, June, September, December and serial months. For the serial months expiry options, the underlying instrument is the future with the nearest expiry. The last day of trading for March, June, September, December expiring options, is the same day as the underlying futures contract.

In addition we are using London Euro-Currency interest rates (middle-rates) on the US dollar, collected from Datastream in order to approximate to the riskless rate in our option pricing model. We collected daily interest rates for 7-days, one-month, three months, six months and one year and for the other maturities we interpolated linearly.

[2] By analyzing implied volatilities, we implicitly assume that implied volatilities give us all the information that we need for the specification of the process of the forward volatility. This is a reasonable assumption because forward volatilities can be extracted from implied volatilities (see Derman and Kani (1997)).

Given that futures options on the S&P 500 are American-style options, implied volatilities σ_{imp} were calculated by the Barone-Adesi, Whaley (1987) quadratic approximation which explicitly takes into account the early exercise premium.

However, even if the market uses the Barone-Adesi Whaley model, the calculation of implied volatilities will be subject to other sources of errors and therefore we should screen the data by imposing several "filtering" constraints.

A preliminary screening of the data is done so as to exclude the data which violate the arbitrage boundary conditions. An American-style futures option is like an American-style option on an asset which pays no dividends. So, the arbitrage boundary conditions are:

$$C > F - K \tag{2}$$

$$P \geq K - F \tag{3}$$

where C (P) is the American call (put) option price, F is the futures price and K is the strike price. We also exclude options having a price of less than 0.01 cents, i.e.

$$C > 0.01 \tag{4}$$

$$P > 0.01 \tag{5}$$

We also eliminate short term options with less than 10 days to maturity,

$$\tau > 10 \tag{6}$$

because they are very sensitive to small errors in the option price.

Since it is documented that non-synchroneity and the bid-ask spread induce spurious negative serial correlation in the volatility changes if closing option prices are used (see Harvey and Whaley (1991) and Roll (1984)), we have also to impose the appropriate constraints[3].

In order to deal with the non-synchroneity problem we exclude in-the-money (ITM) calls and puts and we construct our smiles by using out-of-the-money (OTM) options; OTM puts for the left part of the smile and OTM calls and for the right part[4]. This is not going to affect the results of our analysis, because the "smiles" for both calls and puts should look similar if put-call parity is to be respected[5].

Coping with the bid-ask spread problem, we impose the "vega constraint", as we call it, which is interpreted in the following way: if we want the given measurement error in the option price ΔC to produce an error in the calculation of implied volatility σ_{imp} less than $\Delta\sigma_{imp}$, then vega has to be greater than a certain cut-off point. So, we keep in our analysis

[3] The trading hours for the S&P 500 Index are 8.30 a.m. till 3.15 p.m. (Chicago time), while the trading hours for the S&P 500 Index Options on futures are 8.30 a.m. till 3.15 p.m. (Chicago time), but the option can be exercised until 7.00 p.m. (Chicago time) on any business day the option is traded. The closing times for the option's and the underlying asset's market coincide, but deep in-the money and out-of the money options are rather illiquid and they trade less frequently than those nearer to the-money, something which means that for those options the option market has closed some time before the asset's market.

[4] We do so, because since ITM calls and puts have a very high delta, their prices and therefore the calculated implied volatilities are very sensitive to the non-synchroneity problem.

[5] Even though we deal with American-style options, the divergence from put-call parity should be very small.

the implied volatilities which are calculated from options having vega bigger than a certain cut-off point. This cut-off point is set to a value of eight for reasons that are explained in S-H-C (1998)[6].

3 Principal Components Analysis and the Implied Volatility Surface

The natural technique to identify the number of stochastic shocks that move the volatility surface, is Principal Components Analysis (P.C.A.). We will describe it briefly to explain why it is appropriate and to introduce the notation we will use.

P.C.A. has been constructed to answer the following question: How can we explain the systematic behavior of the observed variables by means of a smaller set of computed but unobserved latent random variables?

From a purely mathematical viewpoint the purpose of a population principal component (PC) model is to transform p correlated random variables to an orthogonal set which reproduces the original variance-covariance structure. This is equivalent to rotating a pth dimensional quadratic form to achieve independence between the variables.

Consider n observations on p random variables represented by the $(n \times 1)$ vectors $x_1, x_2, \ldots, x_p$, where for simplicity we assume that the means of these vectors are zero, i.e. $\overline{x_1} = \overline{x_2} = \ldots = \overline{x_p} = 0$ and an estimator S of the variance covariance matrix is given by $S = (\frac{1}{n-1})X'X$ where X is a $(n \times p)$ matrix[7]. Construct now linear transformations Z_i for $i = 1, 2 \ldots, p$ so that

$$Z = XP \tag{7}$$

where Z is a $n \times p$ matrix and P is a $p \times p$ matrix with the ith column the $\underline{P_i'}$ vector $\underline{P_i'} = (p_{1i}, p_{2i}, \ldots p_{pi})$, for $i = 1, 2, \ldots, p$. The ith column Z_i is the *ith principal component* (PC hereafter) and the elements of the vector $\underline{P_i'}$ are the coefficients of the ith PC and they are called the *loadings*.

The variance of the Z_i PC is given by $Var(Z_i) = \underline{P_i'}X'X\underline{P_i}$ and the PCs are constructed such that $Cov(Z_i, Z_j) = 0$ for $Z_i \neq Z_j$.

The purpose of the P.C.A. is achieved by determining the unknown fixed loadings, so as to maximize sequentially the variance of the PC's, starting from the first one up to the pth, under the constraint that $P'P = I$ where I is the identity matrix.

[6] The advantages of the vega-constraint over the constraint on $\frac{F}{K}$ that other researchers have used so far, are that (a) the cut-off point is not determined exogenously, but it has been delivered endogenously by the data and (b) the ratio $\frac{F}{K}$ remains constant regardless of the days to maturity of the option, while the calculation of the vega-constraint takes explicitly into account the days to maturity. For a more detailed description of the vega-constraint see S-H-C (1998).

[7] Since degrees of freedom represent a scalar constant, they are usually omitted, and the analysis is based on $X'X$ rather than on the sample covariance matrix S.

The first order condition of this maximization problem results in

$$(X^{'}X - lI)P = 0 \tag{8}$$

where l_i are the langrange multipliers.

From equation (8) it is evident that the P.C.A. boils down to the calculation of the eigenvalues l_i and the eigenvectors of the variance-covariance matrix $X^{'}X$. S is now represented as

$$S = P^{'}LP \tag{9}$$

where L=$diag(l_1, l_2, \ldots, l_p)$ and the sum of the variances of the PCs equals the sum of the variances of the X variables. Moreover, it can be proved that if the covariance matrix is diagonal, then there is no gain in performing the PCA.

As Basilevsky explains (Basilevsky (1994), page 143), it is better to work with the standardized variables of Xs and therefore with the correlation matrix, rather than the covariance matrix of Xs. *From now on we will be working with standardized variables and PCs* (unless otherwise stated) for the reason that will become immediately obvious.

From Equation (7) standardizing the PCs to unit length we get a new matrix Z^* which is $Z^* = ZL^{-\frac{1}{2}} = XPL^{-\frac{1}{2}}$. Hence,

$$X = Z^*A^{'} \tag{10}$$

where

$$A^{'} = L^{\frac{1}{2}}P^{'} \tag{11}$$

When both variables and components are standardized to unit length, the elements a_{ij} of A' are correlations between the variables and PCs and they are called *correlation loadings*. If we retain $r < p$ PCs then

$$X = Z^*_{(r)}A^{'}_{(r)} + \varepsilon_{(r)} \tag{12}$$

where $\varepsilon_{(r)}$ is a $(n \times p)$ matrix of residuals and the other matrices are defined as before having r rather than p columns. The percentage of variance of X_i for $i = 1, 2, \ldots, p$ which is explained by the retained PCs (communality of X_i) is calculated from the correlation loadings. Hence, after retaining $r < p$ components which explain a sufficient amount of the original variance-covariance structure we look at equation (12) to see how big the communalities are and what is the meaning of the retained components.

Since we want to determine the number n of shocks that appear in the stochastic differential equation (1) we are going to apply the PCA on the *first differences of implied volatilities*, because this is the natural thing to look at, if we work with a discretized version of equation (1).

Moreover, we are going to apply it on two different measuring units (metrics): (a) we will consider as variables the first differences of implied volatilities classified across strikes (strike metric) which is the natural measuring unit and (b) we will consider as variables the first differences of implied volatilities classified across moneyness i.e. $\frac{K-F}{F} * 100$ (moneyness metric). The reason for choosing this metric is that there are theoretical reasons (see Heynen (1993), Taylor and Xu (1993)) for believing that smiles are a function of moneyness.

In order to look into the dynamics of implied volatilities as a surface for a given year, we will determine first the buckets, then the variables for each bucket and finally we will pool the variables from all the buckets together, so as to apply the PCA on them. Notice that the imposition of the constraints created many missing observations which we do not try to fill in because as Anderson et al. (1983) note "The only real cure for missing data is to not have any". Instead, we are going to apply *listwise deletion* i.e. we will delete the whole day for which at least the observation for one variable (strike or moneyness) is missing. This means that pooling the ranges and applying listwise deletion will decrease considerably the number of observations if the maturity ranges are too fine.

It turns out that the intervals of days to maturity that give us a satisfactory number of observations (not less than 100) and permit us to measure smiles across a wide range (not less than 20 variables) when they are pooled together, are 90-10, 180-90 and 270-90 and consequently these are the ranges that we are going to work with[8].

4 PCA on the Strike Metric

4.1 Number of Retained Principal Components and a First Interpretation

Under the application of listwise deletion, the variables for each year, within a given range, on which we will perform the PCA for the strike metric appear in Table 1. We can see that we have a wide range of strikes and the number of observations does not fall below 100. Moreover, the overall correlation between the variables to which we apply the PCA, as this is measured by the Kaiser-Meyer-Olkin measure (KMO), is between 0.77 and 0.89, something which guarantees that the application of the PCA is feasible.

One of the contributions of the paper is that instead of relying on rules of thumb, we will decide on the number of PCs to be retained, by following three steps[9]. As a first step we will apply Velicer's (1976) non-parametric criterion, and in the second and third step we will look iteratively at the communalities and at the interpretation of the retained, according to Velicer's criterion, PCs plus one more until the additional PC is just noise[10].

[8]In S-H-C (1998) the maturity ranges were six : 30-10, 60-30, 90-60, 150-90, 240-150 and 360-240. However, when we tried to pool these ranges together, the listwise deletion gave us a too small number of observations (around 50) for the application of PCA.

[9]Such rules of thumb are to keep the components which have eigenvalues bigger than the mean of all the eigenvalues, or omit correlation loadings which are smaller than 0.20, or look at the scree plot to decide on how many components we should retain, or keep the components which explain 90% of the total variance (see for example Litterman and Sheinkman [14]). For the drawbacks of all these ad hoc rules see Jackson [12].

[10]The reason that we apply a non-parametric test is that the usual tests are parametric based on the assumption of multivariate normality (For a review of these tests see Basilevsky (1994)). We checked this assumption by applying Bera-Jarque test where the null hypothesis is that of univariate normality. The results showed that for all the maturity buckets the null-hypothesis of univariate and hence of multivariate normality was rejected (this conclusion comes from the well-known theorem that multivariate normality implies univariate normality). This was expected, since the implied volatility can not be negative, and therefore is a random variable truncated at zero. This leaves us with no alternative choice, than the use of a non-parametric test.

Velicer proposed a non-parametric method for selecting nontrivial PCs, i.e. components which have not arisen as a result of random sampling, measurement error, or individual variation, based on the partial correlations of the residuals of the PCs model, after $r < p$ components have been extracted. The idea of the criterion can be described briefly as follows: Let

$$X = Z_{(r)}A_{(r)}^{'} + \varepsilon_{(r)} \tag{13}$$

Basilevsky (1994) shows (theorem 3.13, page 132) that $X^{'}X = AA^{'}$. Hence, the variance-covariance matrix of the residuals $\varepsilon_{(r)}$ is given by the following expression

$$\varepsilon_{(r)}^{'}\varepsilon_{(r)} = X^{'}X - A_{(r)}A_{(r)}^{'} \tag{14}$$

Let $D = diag(\varepsilon_{(r)}^{'}\varepsilon_{(r)})$. Then, $R^{*} = D^{-\frac{1}{2}}\varepsilon_{(r)}^{'}\varepsilon_{(r)}D^{-\frac{1}{2}}$ is the matrix of partial correlations of the residuals[11]. If r_{ij}^{*} represents the ith row, jth column element of R^{*}, then the Velicer statistic is given by

$$f_r = \sum_{i \neq j}\sum \frac{r_{ij}^{2*}}{p(p-1)} = \sum_{i=j}\sum \frac{r_{ij}^{2*} - p}{p(p-1)} \tag{15}$$

and lies in the interval 0 to 1.

One might expect that as we retain more PCs, $f_r \rightarrow 0$, but this is not the case. The behavior of f_r is that it is decreasing until a number r^{*} and then it increases again.

Velicer suggests that $r = r^{*}$ should be the number of components to be retained. The logic behind this is that as long as f_r is declining, there is still space for the additional components to capture part of the covariance of the residuals. Since it is obvious from the formula of the partial correlation that f_r is declining as long as the the partial covariances are declining faster than the residual variances, the Velicer's procedure will terminate when, on the average, additional PCs would explain more of the residual variances than their covariances.

In Table 2 we see the number of PCs that should be retained according to Velicer's criterion[12]. We see that the retained number of components varies across the years, but it is never bigger than two[13]. This is very different to what the number of retained PCs is according to the mean eigenvalue rule of thumb $\bar{l}$, which retains the PCs corresponding to the eigenvalues which are greater than the mean eigenvalue. From Table 2 we can see that the mean eigenvalue rule retains either five or six PCs. We also show a second statistic which is f_0 and it is the usual Velicer's statistic when no PCs have been retained. If $f_1 > f_o$ then no

[11] Velicer is talking about partial correlations in the sense that we have kept out of the analysis $p - r$ components.

[12] The minimum point in Velicer's criterion is not obvious in some cases because of the truncation of the displayed results up to 4 decimal places. The same applies for Velicer's results in the moneyness metric.

[13] In terms of the communalities, there are some cases that the two retained PCs seem to perform poorly (for instance they explain 19.35% for the strike price of 450 for 180-90 days to maturity for the year 93), but it is not clear whether the inclusion of the third PC would improve the communalities significantly (for the case that we mentioned the inclusion of the third PC increases the communalities to 25.63%). Therefore, before reaching any conclusions about the number of PCs to be retained, we should look at the interpretation of the shape of the first three PCs, as well.

components would be extracted. We can see that $f_1 < f_o$ in all the cases and consequently the PCA can reduce the dimensionality of the original variables[14].

Table 2 also shows the amount of variance that it is explained by the two retained components and the additional amount of variance that it would have been explained if we include and the third PC[15]. An interesting point is that the explained by the first PC variance declines over the years, while for year 95 the cumulative variance has been distributed almost equally between the two PCs[16].

The interpretation of the retained PCs can be revealed by looking at the correlation loadings matrix A' as equation (10) reveals. So, the first three columns of A' give the way that the first three PCs, respectively, affect the implied volatilities.

It is clear from Figures 1 and 2 that the first PC has a Z-shape, even though in the year 95 we do not have many variables for the range 90-10 (only three), so that to be able to see the shape of the first PC clearly for this particular range[17]. This interpretation is quite similar to what we discovered from the smile analysis for the first PC in the strike metric for the longer maturities.

Regarding the interpretation of the second PC, we can see that for the years 92 and 93 it looks the same in a given range, but then in the remaining two years the interpretation changes. To make things more concrete, for all the years in the range 90-10, it has a shift interpretation which nevertheless deviates a lot from a parallel shift.

In the other two ranges and for the years 92 and 93 it has a triangular shape where the correlation loadings for the very low strikes become negative. However, in the year 94 the 2nd PC has been flipped over for the longer ranges (preserving still the triangular shape), while in the year 95 the shape is no longer triangular, but it is a Z-shape instead. The strong triangular shape that we get now from the pooled buckets should not surprise us, since from the smile analysis we saw that for the maturities with more than 90 days, the 2nd PC has a triangular shape as well[18].

Hence, it seems that the appropriate number of PCs to be retained in the strike metric is two, something which is consistent with Velicer's criterion and this is similar to what S-H-C found.

However, the shape of the two retained PCs contrasts to the intuition coming from the Taylor series expansion, while especially in the longer maturities it does not have a simple interpretation[19]. Since a simple interpretation of the retained PCs helps in implementing

[14]If the null hypothesis of Velicer's test is the reduction of dimensionality, then as Reddon (1985) showed, the type-I error rates of Velicer's test go to zero, as long as the number of observations is far more than twice the number of variables. Since this is the case here, our results from the comparison f_0 and f_1 are credible.

[15]This is done so that we can compare it to the results found by Litterman and Scheinkman (1988), even though the retainment of the third PC there had been justified only on the grounds of explained variance and of interpretation. Besides following our selection of number of PCs procedure, we still have to look at the interpretation of the third PC, before deciding whether we should keep it, or not.

[16]The average, over the years, cumulative variance explained by the two first PCs is 52.5%, in contrast to the 98% that Litterman and Sheinkman found.

[17]To make the interpretation easier, we have interpolated for the missing variables correlation loadings (and later loadings) across the variables for a given maturity bucket.

[18]Coming to the third PC, the graphs revealed that this is just noise since its shape is different across years for a given maturity bucket and it is also different accross the maturity buckets for a given year.

[19]Intuitively thinking, any well-behaved function can be approximated by a Taylor series expansion of first

equation (1) we should seek for it through a rotation of the retained PCs[20].

4.2 Interpretation of the Rotated PCs

Given that intuitively from the Taylor series expansion we would expect the first PC to have a shift interpretation and the second to have a slope interpretation, we use a "Procrustes" type of rotation[21 22].

The idea behind "Procrustes" type rotation is that given two $p \times r$ matrices, A and B, what $r \times r$ transformation matrix T will best transform A into B? In other words, what matrix T will make AT most like B? B is called the target matrix, and in our case it will be the $p \times 2$ matrix of rotated loadings which will give the parallel and the slope character for the first and second PC respectively.

The way that we construct our rotation method is by using the general properties of orthogonal rotations as outlined in Basilevsky (1994) and then we regress a vector of constants, say of ones, on the vector of loadings of the first two PCs. The latter determines the elements of the first row of the matrix T and combining this with the former we get the elements of the second row[23]. Doing so, we force the first rotated PC to look like a shift but we do not know a priori what the second rotated PC is going to look like.

In Figures 3 and 4 we show what the first two rotated PCs look like for the years 92 up to 95. We can see that in the range 90-10 the rotated first PC has a shift interpretation, even though this is not entirely true for some years if we look at the extreme strikes. In the other two ranges and for all the years the first PC has a triangular shape which is flipped over for the year 94. In general, the shape of the first rotated PC is very similar to what the second unrotated PC looked like.

The second PC has clearly a Z-shape for the years 92 and 93 and for all the ranges. For the year 94 and for the shortest maturity it has a rather noisy shape, while for the other two ranges it has a Z-shape. Finally, for the year 95 in the range 90-10 the shape is not clear, since we do not have many variables, while for the other two ranges there is a Z-shape.

In general, the interpretation of the rotated PCs is the same as the one that we found in S-H-C for the separate maturity buckets. This implies that the implied volatility surface, in

order, where the zero order expansion is the level, and the first order expansion is the slope.

[20] In our case of course, in order to implement equation (1) it does not matter whether or not we have a simple (or even any) interpretation of the retained PCs, since the correlation loadings are the estimates of the b_i's coefficients. However, if we want to estimate the b_i's by another econometric technique, we should know their functional form. A simple interpretation of the PCs will help us in achieving this. Intuitively thinking, it is possible to interpret the first PC as a level and the second as a slope because any well-behaved function can be approximated by a Taylor series expansion of first order, where the zero order expansion is the level, and the first order expansion is the slope.

[21] We use the terminology "shift" and "slope" for the interpretation of the PCs, as Litterman and Scheinkman have already established.

[22] Before trying to apply a "Procrustes" type of rotation, we applied the most popular rotation methods, i.e. the varimax, the quartimax and the oblique method. However, none of these methods produced the desired interpetation because of the way that they are constructed (for more details about these rotation methods see Basilevsky (1994)).

[23] Notice that we apply the rotation on P and not on A', because we use in the construction of our method the fact that the rotated eigenvectors remain orthogonal, while this is not true for rotated correlation loadings (see Basilevsky (1994)).

the strike metric, is driven not only by the same number, but also by the same shocks, as the smiles.

Looking at the shape of the PCs, before and after the rotation, it seems as if the rotation has changed the names between these two, i.e there is a swap effect, as we call it; before the rotation the first PC had a Z-shape and the second had positive loadings, but after the rotation the first PC was the shift and the second the Z-shape. In order to be sure about this effect we should look at the percentage of variance that the two rotated PCs explain. Table 3 shows the amount of the variance that the two rotated PCs explain. Comparing Tables 2 and 3 it is clear that the variance that the first PC explained before the rotation, is explained now by the second rotated PC, and similarly for the second PC. Therefore, the swap effect is confirmed.

As a final part of the analysis on the strike metric we should look at whether the effect of the retained components is bigger on the shorter, or the longer maturity options' implied volatilities. Figures 3 and 4 show that the size of the effect from the first PC varies over the years. For instance in year 93 the effect is bigger for the range 270-180 than for the range 180-90, while this is reversed in 94. On the other hand the second PC moves the volatilities of the longer ranges more than the volatilities of the shorter ranges and this is true for all the years.

5 PCA on the Moneyness Metric

5.1 Preliminary Testing

Coming to the application of PCA under the moneyness metric, the problem is that the moneyness variables $\frac{K_i - F_t}{F_t} * 100$ (we will call this the "natural" moneyness metric), for $i = 1, 2, \ldots, s$ where s is the number of strikes traded for a given day t , are different from day t to day $t+1$, as the futures price changes from F_t to F_{t+1}. This prevents the application of the technique. The only solution to this is to determine (or better to fix) the moneyness variables before starting the analysis (we will call these fixed moneyness variables the "artificial" moneyness metric) and then, for each day, to interpolate across the implied volatilities for these "fixed" variables.

The spacing between the variables of the "artificial" moneyness metric (step-size) was set so that between any two consecutive variables of the "natural" moneyness metric, there will be only one variable of the "artificial" moneyness metric. Otherwise, the implied volatilities created by the interpolation of both of the "fixed" variables, would depend on the *same* two consecutive values of the "natural" moneyness metric and consequently they would exhibit spurious dependence, something which of course would distort the results of our subsequent analysis[24].

Table 4 shows the number of variables, the number of observations and the KMO that we

[24]The step-size was set to 1.2% for the 90-10 range and it varied from 2.2% to 2.4% for the other two ranges across years. We had also to think of the range of the moneyness variables, so as to have as many observations as possible after the listwise deletion. The range of moneyness that we look at, varies by the years, but the extreme values are -18% for the lower and 6% for the higher. For more details see S-H-C.

have for the moneyness metric, when we pool the maturity buckets together. The demonstrated results lead us to expect reliable results from the PCA.

5.2 Number of Retained Principal Components and a First Interpretation

In order to see how many PCs we should retain we are going to apply again the three step procedure as we did in the strike metric.

The results from applying Velicer's criterion on the moneyness metric are shown in Table 5 [25]. We can see that for all the years the test tells us that we should keep two PCs, while the reduction in the dimensionality of the variables is legitimate, since $f_0 > f_1$. Again, we can see the discrepancy between the results from Velicer's procedure and the mean eigenvalue rule. The latter retains between four and six PCs.

Compared to the corresponding results that we had when we worked in the strike metric (see Table 2) the cumulative variance is bigger in the moneyness metric with the difference ranging from 4% up to 10%. This difference stems from the higher amount of variance explained by the second PC in the moneyness metric. The fact that the explained by the first two PCs variance is bigger in the moneyness metric, than it is in the strike metric, is similar to what S-H-C found. Another point is that we do not observe any longer the declining behavior of the explained variance by the first PC across the years, as it was the case in the strike metric.

The interpretation of the first two PCs in the moneyness metric is shown in Figures 5 and 6 [26]. The first PC for the years 92 and 93 and for all the ranges, has a Z-shape. In years 94 and 95 and for the range 90-10 it has a shift interpretation, while for the other two ranges it preserves the Z-shape. This is very close to what S-H-C found in the analysis of the smiles in the moneyness metric, where we got a shift shape for the ranges 60-30 and 30-10 and a Z-shape for the remaining ranges over the years.

Regarding the shape of the second PC the general impression is that it has a shift interpretation for all the years and all the ranges, while the triangular shape that we observed in the strike metric is not present any longer. This is quite different to what we found in the smile analysis, where we had a Z-shape in the ranges 60-30 and 30-10, a triangular shape for the ranges 360-240 and 240-150 and a shift interpretation for 150-90[27].

Hence, given the interpretation of the PCs and the previous two steps of our selection procedure, *it makes sense to retain two PCs for the moneyness metric.* This is exactly the same result that we found from the smile analysis.

[25] We checked again the null hypothesis of multivariate normality by using Bera-Jarque's test and we found that this was rejected. Therefore, we had to use the non-parametric Velicer's criterion.

[26] We looked at the communalities of the retained, according to Velicer's procedure, PCs in the moneyness metric and we found that they are satisfactory. The only cases that they need the assistance of a third PC are in year 94 for the at-the-money variable in the ranges 180-90 and 270-180 (the communalities there were 4.71% and 1.23% respectively). If we add and the third PC, then the explained communalities would have been again only 17.82% and 4.02% respectively.

[27] The graphs for the third PC showed that even though there is some evidence of consistency in its behavior across the years for a given range, the way that it moves the implied volatilities of the different ranges, for a given year, is different and therefore it can be regarded as noise.

Given that the shape for the first PC is either a Z-shape or a shift and for the second is a shift, we have to perform again our rotation method. In this way we hope that we will get a shape for the first PC which will not alter and which will have the shift (ideally parallel) interpretation and that we will get the Z-shape (ideally slope) for the second PC. This will be consistent with the Taylor series intuition as well.

5.3 Interpretation of the Rotated PCs

In Figures 7 and 8 we show the shape of the first and second rotated PCs in the moneyness metric. We can say that the first PC moves the implied volatility surface consistently across ranges and across years and its shape "approaches" the shift interpretation. Therefore, the rotation is very successful in this case. This is what we found from the smile analysis.

The second PC in general has a Z-shape across years, even though for the years 94 and 95 it has a shift interpretation in the range 90-10. Hence, the evidence for a Z-shape interpretation of the second rotated PC in the moneyness metric, is not as unanimous, as it was in the smile analysis.

Inspecting the shape of the two retained PCs before and after the rotation, it seems as if the swap effect is present once again. In order to be sure about this, we depict in Table 6 the amount of the variance that the two rotated PCs explain on the moneyness metric. Comparing Tables 5 and 6 we see that the first rotated PC explains the variance that was explained by the second PC and similarly for the second rotated PC. Hence, the rotation has produced a swap effect in the moneyness metric, as well. We also calculated the angle with which we rotated the original axes. We found that the angle was close to 90 degrees for all the years, which is another way of stating the swap effect. In fact, the only cases where the angle was different than 90 degrees were for the years 93 and 95 in the strike metric (64.71 and 29.29 degrees, respectively). Moreover, Table 6 shows that after the rotation it is not the shift which has the dominant effect on the implied volatility surface, but the second PC. This is a very interesting result and it is similar to what we found in the strike metric.

Finally, from Figures 7 and 8 we can see that the first rotated PC affects the shorter implied volatilities more than the longer ones, in contrast to what we found in the strike metric. It is worth mentioning that the effect of the first rotated PC is the same as the one from the smile analysis. The impact of the second rotated PC on the implied volatilities is bigger for the longer maturity options, in contrast to the effect from the first PC.

6 Correlations between the Futures Price and the Principal Components

Apart from knowing the number and shape of the shocks which appear in equation (1), it is necessary to know the sign and the size of the correlations, between the Brownian motion of the process for the underlying asset and the Brownian motions of the process for the forward volatility[28]. One way of solving this problem is by looking at the correlations between the

[28] For example in Derman and Kani (1997) there are two processes : one for the underlying asset and one for the forward volatility. Then, a joint simulation of the two processes (which requires the knowledge of

percentage changes in the futures price and the changes of each one of the two principal components.

This can be done by using the percentage changes of the futures price from each bucket and the differenced-rotated PC since it is the rotated PCs that deliver to us the desired interpretation[29]. We measure correlations by using the Pearson correlation coefficient, where two asterisks show significance at 10% and one asterisk shows significance at 5% level.

In Table 7 we show the correlations between the changes in the futures price and the changes in the rotated first and second PCs for the strike and moneyness metric, respectively. We can see that in the strike metric, the correlation is positive for both the PCs, while in the moneyness metric the correlation for the first PC has reversed its sign[30], i.e. it is negative now something wich can be thought of as a leverage effect (see Christie (1982)) and the correlation for the second PC remains positive, apart from year 93, where it is negative[31].

Regarding the size of the correlation, this changes over the years, something which supports the idea that the size of the correlation changes in a stochastic way because the variance varies stochastically.

These results are almost identical to what we found from the smile analysis in terms of the sign of the correlation, something which demonstrates that the sign of the correlation is not affected by whether we examine the whole implied volatility surface, or just the smiles[32].

7 Conclusions

In this paper we looked at the dynamics of implied volatilities surfaces under two different metrics, so that to enable the implementation of Dupire's and Derman and Kani's models for the pricing and hedging of standard and exotic options. In particular, we wanted to answer three questions, namely the number of shocks that appear in the forward volatility process,

---

the correlation of the shock in the asset process with the shocks in the forward volatility process) can reveal the forward (local as they call it) volatility surface which makes the valuation and hedging of standard and exotic options, possible.

[29]Inspection of the correlations between the changes of the unrotated third PC and the changes of the futures price revealed that these were insignificant in the strike metric and significant for the years 92 and 94 in the moneyness metric. However, its shape for these years does not make any sense and therefore our choice of treating the third PC as noise, is fully justified.

[30]Since the correlation has the same sign as the covariance, it is easy to show that the correlation depends on the metric that we work on, by looking at the covariance between $\Delta\sigma_t = \sigma_{t+1} - \sigma_t$ and ΔF under both metrics. Say that the covariance in the strike metric is $Cov_{strike} = Cov(\Delta\sigma_t(K), \Delta F)$. Then, in the moneyness metric, for a given moneyness we have $Cov_{mon} = Cov[\sigma_{t+1}(K + \Delta F) - \sigma_t(K), \Delta F]$. Expanding $\sigma_{t+1}(K + \Delta F)$ as a Taylor series of order one around a point K we get : $Cov_{mon} = Cov[\sigma_{t+1}(K) + \Delta F \sigma'_{t+1} - \sigma_t(K), \Delta F]$ $= Cov_{strike} + \sigma'_{t+1} Var(\Delta F)$. Therefore, whether or not the correlation sign is going to alter as we change metrics, depends on the slope of the skew, the point around which we do the expansion and the variance of ΔF.

[31]Inspecting scatterplots of $\frac{\Delta F}{F}$ with changes in the second PC for the year 93, revealed that there were some outliers in $\frac{\Delta F}{F}$. When we removed the outliers, the correlation for the range 90-10 became insignificant, but for the other ranges the negative correlation remained significant (even though it was reduced).

[32]We also calculated the correlations by using the non-parametric Spearman coefficient in order to capture any non-linear association, but this did not produce different results compared to Pearson.

their interpretation and their correlation with percentage changes in the underlying asset.

Application of Velicer's procedure retained two PCs in both metrics and for all the years. Looking also at the communalities and at the interpretation of the first three PCs, we decided that we should retain two PCs, in contrast to what researchers did in the interest rate literature, since the third PC did not increase significantly the explained communalities and it was just noise interpretation-wise. The shape of the retained PCs and the fact that they explain more of the variance in the moneyness metric than in the strike metric, is similar to what we found in the smile analysis. However, their interpretation was not a simple one, something which would make the implementation of the discussed models difficult. Therefore, we tried to simplify the interpretation of the retained PCs by applying a "Procrustes" type of rotation.

The desired interpretation was a shift (ideally parallel) for the first PC and a Z-shape for the second PC. Such an interpretation would be consistent with the intuition coming from a Taylor series expansion.

Looking at the shape of the rotated PCs we could say that the rotation met the Taylor series expansion intuition, but it simplified their interpretation only in a few cases (in contrast to the smile analysis). This is because it achieved a swap-effect between the two unrotated PCs which is confirmed by the variance that the rotated PCs explain (compared to the variance that the unrotated PCs explained) and the angle of the rotation.

The interesting point is that after the rotation, it is the Z-shape which dominates the shift in terms of explaining the movement of the implied volatility surface. Considering the size of the effect of the rotated PCs on the implied volatilities across the maturity ranges, in the strike metric the first PC has an effect which varies over years, but in the moneyness metric it affects more the shorter range implied volatilities than the longer range. The second PC moves more the implied volatilities of the longer ranges than the volatilities of the shorter ranges in both metrics.

The correlations between the percentage changes in the futures price with the changes in the rotated PCs were positive for both the PCs in the strike metric, while in the moneyness metric they were negative for the first PC and either positive or negative for the second PC. The fact that the sign of the correlation for the first PC alters as we change metric, should not surprise us, but on the contrary it is something that we should expect.

Comparing to Skiadopoulos, Hodges and Clewlow, the implication of our results is that the number of shocks that move the whole volatility surface, their interpretation and their correlation with changes in the underlying asset, is the same as those which move the smiles. This means that the specification of the implied volatility process when we look the whole implied volatility surface, is the same as with the specification for the implied volatility process when we examine only the smiles for one maturity bucket at a time. In a way, intuitively this was expected, since the latter specification turned out in S-H-C to be the same regardless of the maturity bucket that we were looking at. In this paper what we have done is to pool together (which can be thought of as a way of averaging) maturity buckets and therefore we pooled together similar processes.

To summarise, the current paper has investigated and characterised the nature of empirical shocks to the implied volatility surface of futures options. This is the first key step to the implementation of Dupire's and Derman and Kani's models for the pricing and hedging of standard and exotic options.

References

[1] Basilevsky, A., 1994, Statistical Factor Analysis and Related Methods Theory and Applications (Wiley Series in Probability and Mathematical Statistics).

[2] Barone-Adesi, G. and R. Whaley, 1987, Efficient Analytic Approximation of American Option Values, Journal of Finance 42, 301-320.

[3] Black, F. and M. Scholes, 1973, The Pricing of Options and Corporate Liabilities, Journal of Political Economy 81, 637-654.

[4] Anderson, A., Basilevsky, A. and A. Hum, 1983, Missing Data : A Review of the Literature, in Rossi, P., Wright, J., and Anderson, A., Handbook of Survey Research, (Academic Press).

[5] Christie, A., 1982, The Stochastic Behavior of Common Stock Variances, Value, Leverage and Interest Rate Effects, Journal of Financial Economics 10, 407-432.

[6] Derman, E. and I. Kani, 1997, Stochastic Implied Trees : Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility Goldman Sachs, Quantitative Strategies Research Notes.

[7] Dupire, B., 1992, Arbitrage Pricing with Stochastic Volatility, Working Paper, Societe Generale Division Options, Paris.

[8] Harvey, C. and R. Whaley, 1991, S&P 100 Index Option Volatility, Journal of Finance 46, 1551-1561.

[9] Heath, D., Jarrow, R., and A. Morton, 1992, Bond Pricing and the Term Structure of Interest Rates : A New Methodology For Contingent Claims Valuation, Econometrica 60, 77-105.

[10] Heynen, R., 1993, An Empirical Investigation of Observed Smile Patterns, Working paper Tinberger Institute, Erasmus University Rotterdam.

[11] Hull, J. and A. White, 1987, The Pricing of Options on Assets with Stochastic Volatilities, Journal of Finance 3 281-300.

[12] Jackson, E., 1991, A User's Guide to Principal Components, (Wiley Series in Probability and Mathematical Statistics).

[13] Johnson, H., and D. Shanno, 1987, Option Pricing when the Variance is Changing, Journal of Financial and Quantitative Analysis 22, 143-151.

[14] Litterman, R. and J. Scheinkman, 1988, Common Factors Affecting Bond Returns, Goldman Sachs, Financial Strategies Group.

[15] Reddon, J., 1985, Monte Carlo type I Error Rates for Velicer's Partial Correlation test for the Number of Principal Components, Criminometrica 1, 13-23.

[16] Roll, R., 1984, A Simple Implicit Measure of the Effective Bid-Ask Spread, Journal of Finance 39, 1127-1139.

[17] Rubinstein, M., 1985, Non-Parametric Tests of Alternative Option Pricing Models, Journal of Finance 40, 455-480.

[18] Scott, L., 1987, Option Pricing when the Variance changes Randomly : Theory, Estimation and an Application, Journal of Financial and Quantitative Analysis 22, 419-438.

[19] Sheikh, A., 1991, Transaction Data Tests of S&P 100 Call Option Pricing, Journal of Financial and Quantitative Analysis 26, 459-475.

[20] Skiadopoulos, G., Hodges, S., and L. Clewlow, 1998, The Dynamics of Smiles, Working Paper, Financial Options Research Centre, University of Warwick.

[21] Taylor, S., and X. Xu, 1993, The Magnitude of Implied Volatility Smiles : Theory and Empirical Evidence for Exchange Rates, Working Paper, Financial Options Research Centre.

[22] Velicer, W., 1976, Determining the Number of Components from the Matrix of Partial Correlations, Psychometrica 41, 321-327.

[23] Wiggins, J., 1987, Option Values under Stochastic Volatility, Journal of Financial Economics 19, 351-372.

Year	Number of Variables	Number of Observations	KMO
92	26	182	0.89402
93	26	198	0.87585
94	32	187	0.89416
95	20	103	0.77463

Table 1: Number of Variables, Number of Observations and their Overall Correlation for the Strike Metric

Year	f_0	f_1	f_2	f_3	r^*	$\bar{l}$	1st PC	2nd PC	3rd PC
92	0.1814	0.1807	0.1807	0.1808	2	6	42.5	12.9	6.8
93	0.1415	0.1410	0.1409	0.1410	2	5	37.5	12.7	7.2
94	0.0513	0.0511	0.0512	0.0512	1	6	40.8	9.8	7.2
95	0.1235	0.1229	0.1226	0.1227	2	5	29.9	23.7	9.3

Table 2: Principal Components in the Strike Metric : r* = the number of components retained under Velicer's criterion (minimum of f0,...f3), l = number of components retained under rule of thumb, with percentage of variance explained by components 1-3

Year	Unrot. 1st PC	1st PC	2nd PC	Cumulative
92	42.5%	13.4%	42.0%	55.4%
93	37.5%	17.3%	33.0%	50.3%
94	40.8%	11.5%	39.0%	50.5%
95	29.9%	28.4%	25.2%	53.6%
Average	37.7%	17.6%	34.8%	52.5%

Table 3: Percentage of Variance Explained by the Unrotated First PC and by the Rotated PCs on the Strike Metric

Year	Number of Variables	Number of Observations	KMO
92	24	125	0.85493
93	23	108	0.85693
94	25	129	0.86127
95	25	105	0.83056

Table 4: Number of Variables, Number of Observations and their Overall Correlation for the Moneyness Metric

Year	f_0	f_1	f_2	f_3	r^*	$\bar{l}$	1st PC	2nd PC	3rd PC
92	0.1847	0.1838	0.1836	0.1838	2	4	38.6	22.7	10.7
93	0.1722	0.1714	0.1711	0.1713	2	4	34.5	26.7	10.2
94	0.1894	0.1884	0.1884	0.1885	2	6	40.6	19.2	10.2
95	0.1726	0.1713	0.1712	0.1714	2	4	39.2	18.3	9.4

Table 5: Principal Components in the Moneyness Metric : r* = the number of components retained under Velicer's criterion (minimum of f0,...f3), l = number of components retained under rule of thumb, with percentage of variance explained by components 1-3

Year	Unrot. 1st PC	1st PC	2nd PC	Cumulative
92	38.6%	22.7%	38.6%	61.4%
93	34.5%	26.9%	34.3%	61.2%
94	40.6%	19.5%	40.4%	59.5%
95	39.2%	18.4%	39.2%	57.5%
Average	38.2%	21.8%	38.1%	60.0%

Table 6: Percentage of Variance Explained by the Unrotated First PC and by the Rotated PCs on the Moneyness Metric

Range	Metric		92	93	94	95
90-10	strike	$\Delta PC1$	0.11	0.09	-0.11	0.25*
		$\Delta PC2$	0.40**	0.33**	0.06	0.25*
	moneyness	$\Delta PC1$	-0.29**	-0.36**	-0.55**	-0.29**
		$\Delta PC2$	0.34**	-0.33**	0.16	0.26**
180-90	strike	$\Delta PC1$	0.11	0.14	0.26**	0.24*
		$\Delta PC2$	0.41**	0.37**	0.38**	0.27*
	moneyness	$\Delta PC1$	-0.29**	-0.28**	0.04	-0.29**
		$\Delta PC2$	0.35**	-0.43**	0.27**	0.26*
270-180	strike	$\Delta PC1$	0.11	0.14	0.27**	0.22*
		$\Delta PC2$	0.41**	0.38**	0.36**	0.27**
	moneyness	$\Delta PC1$	-0.28**	-0.28**	0.06	-0.31**
		$\Delta PC2$	0.35**	-0.44**	0.27**	0.26**

Table 7: Correlations between Percentage Changes of the Futures Price with Changes of the Rotated PCs on the Strike and Moneyness Metrics

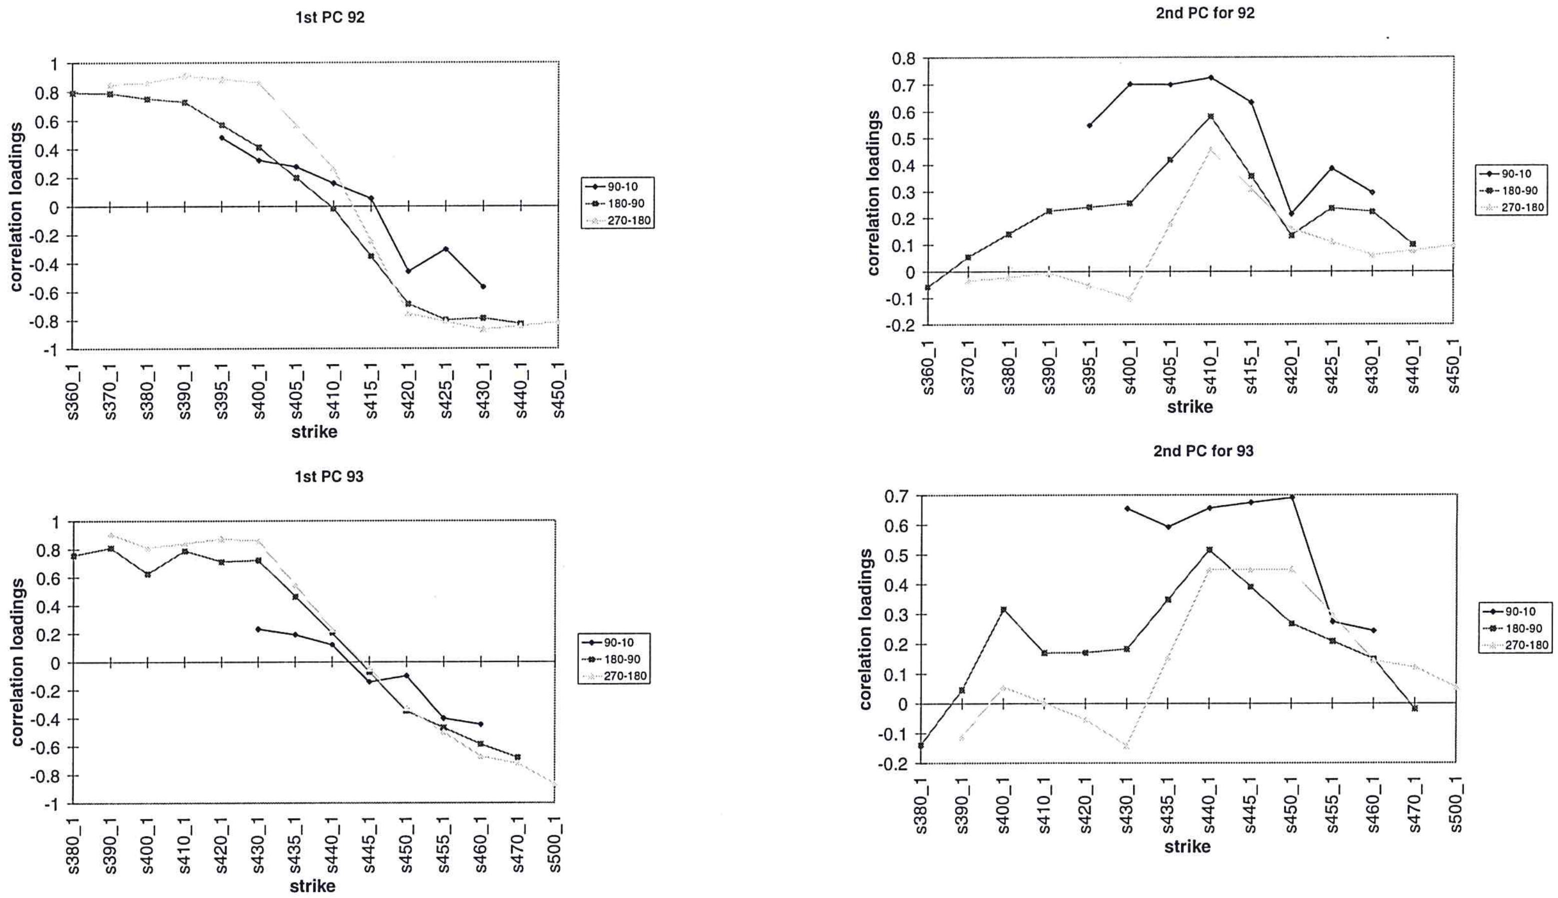

Figure 1 : First and Second PCs for 92 and 93 in the Strike Metric

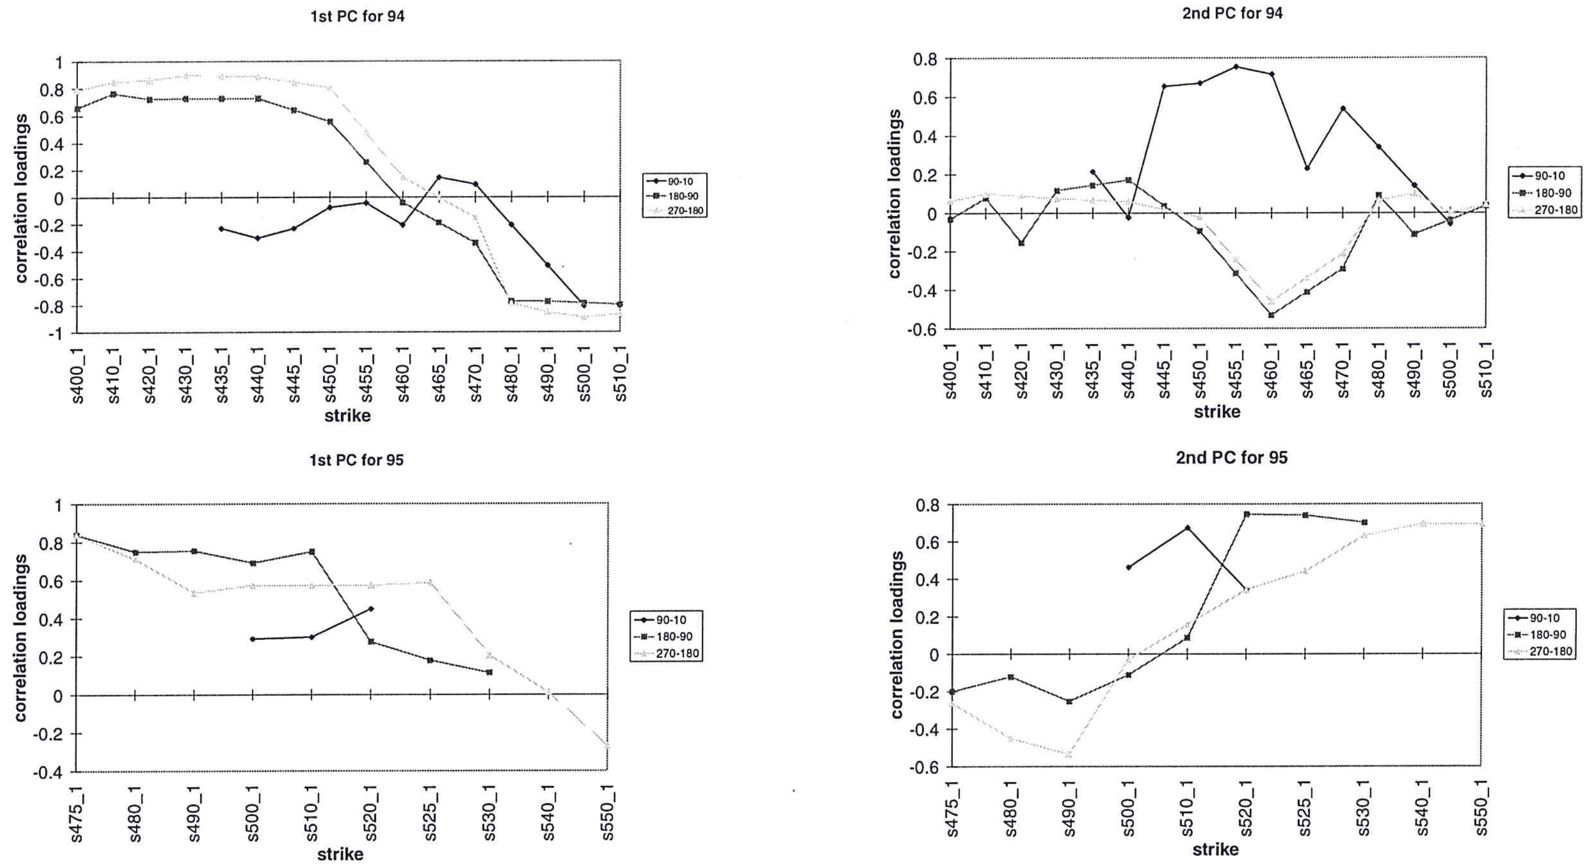

Figure 2 : First and Second PCs for 94 and 95 in the Strike Metric

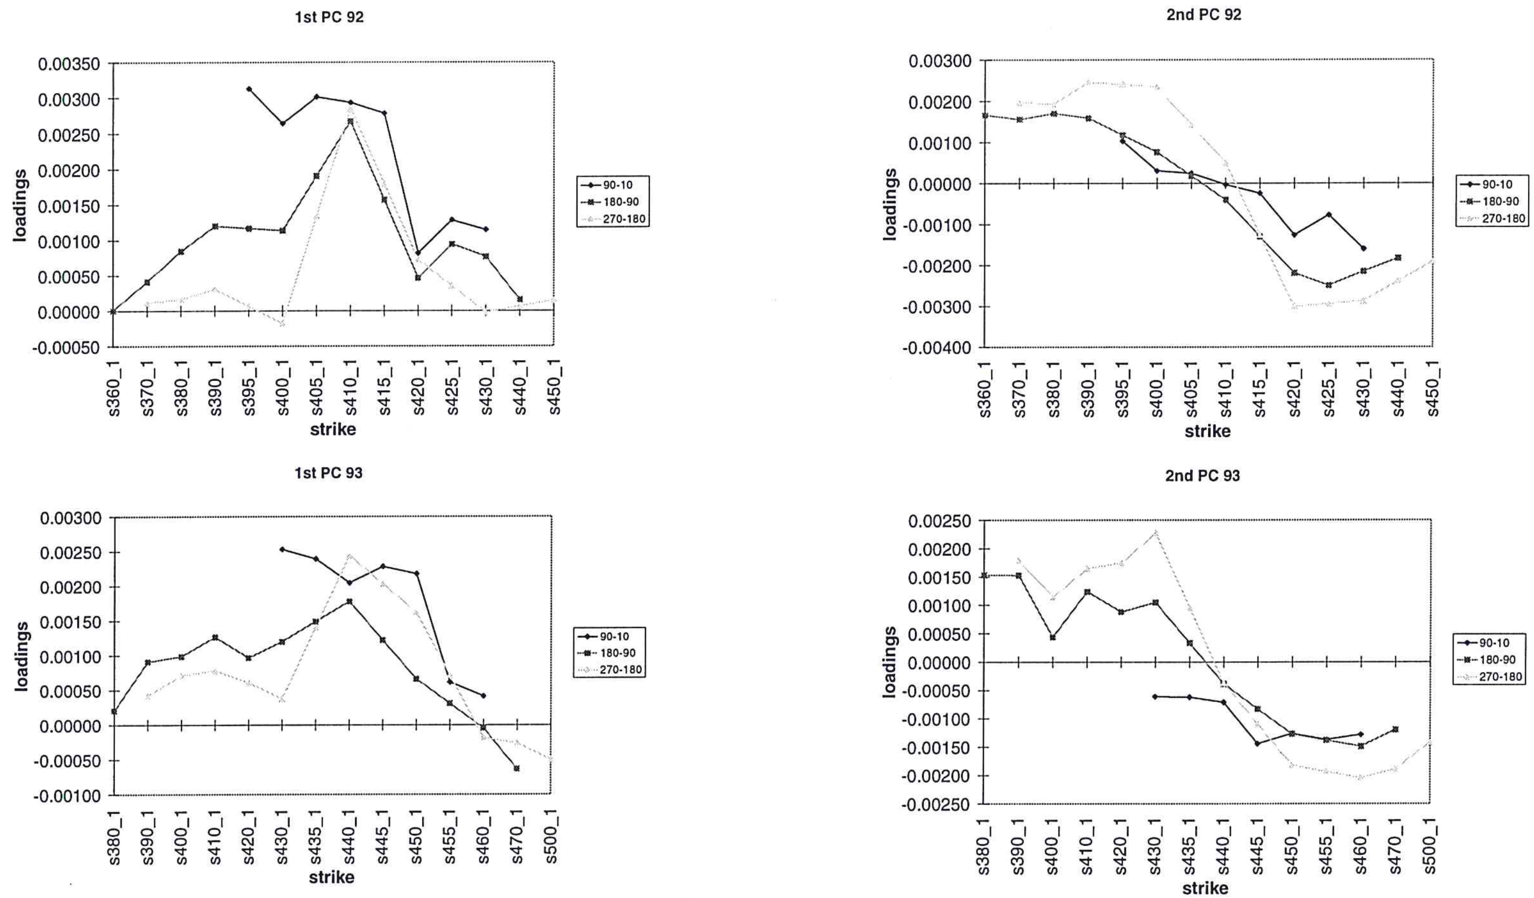

Figure 3 : First and Second Rotated PCs for 92 and 93 in the Strike Metric

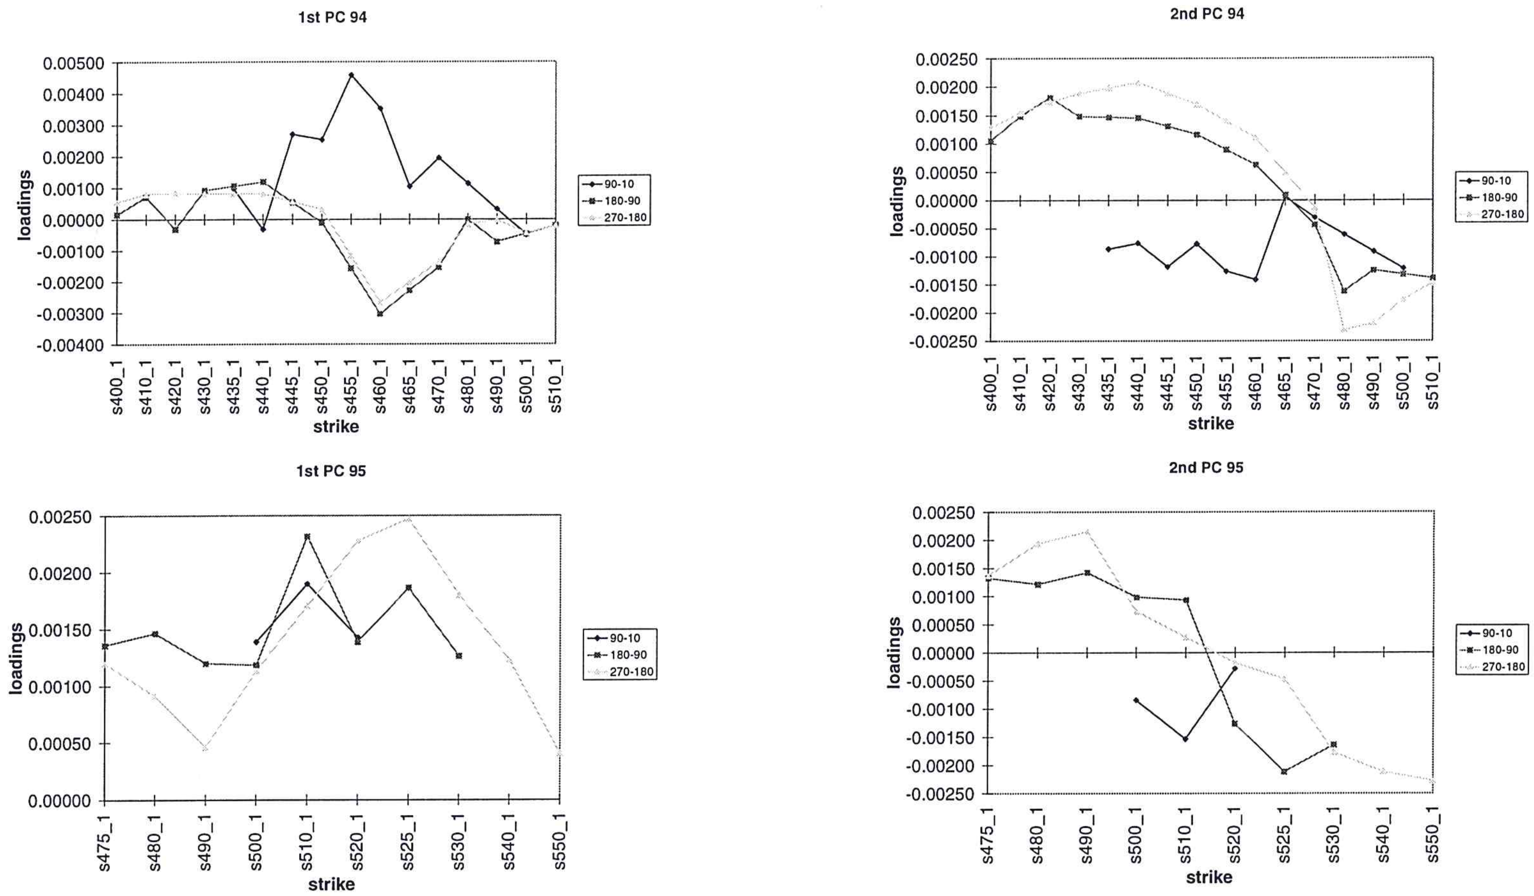

Figure 4 : First and Second Rotated PCs for 94 and 95 in the Strike Metric

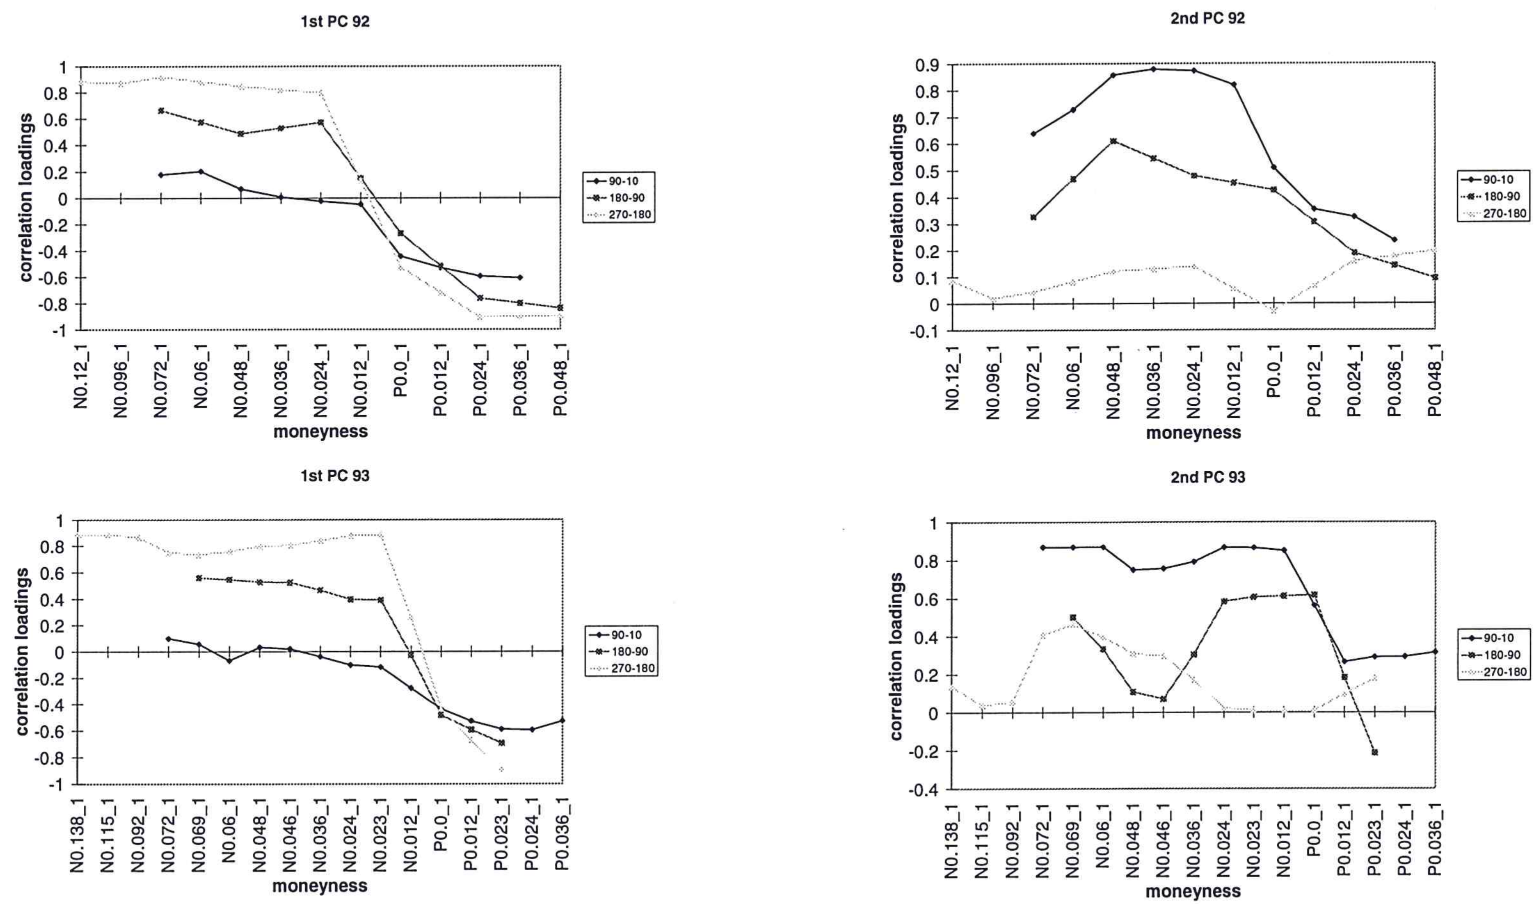

Figure 5 : First and Second PCs for 92 and 93 in the Moneyness Metric

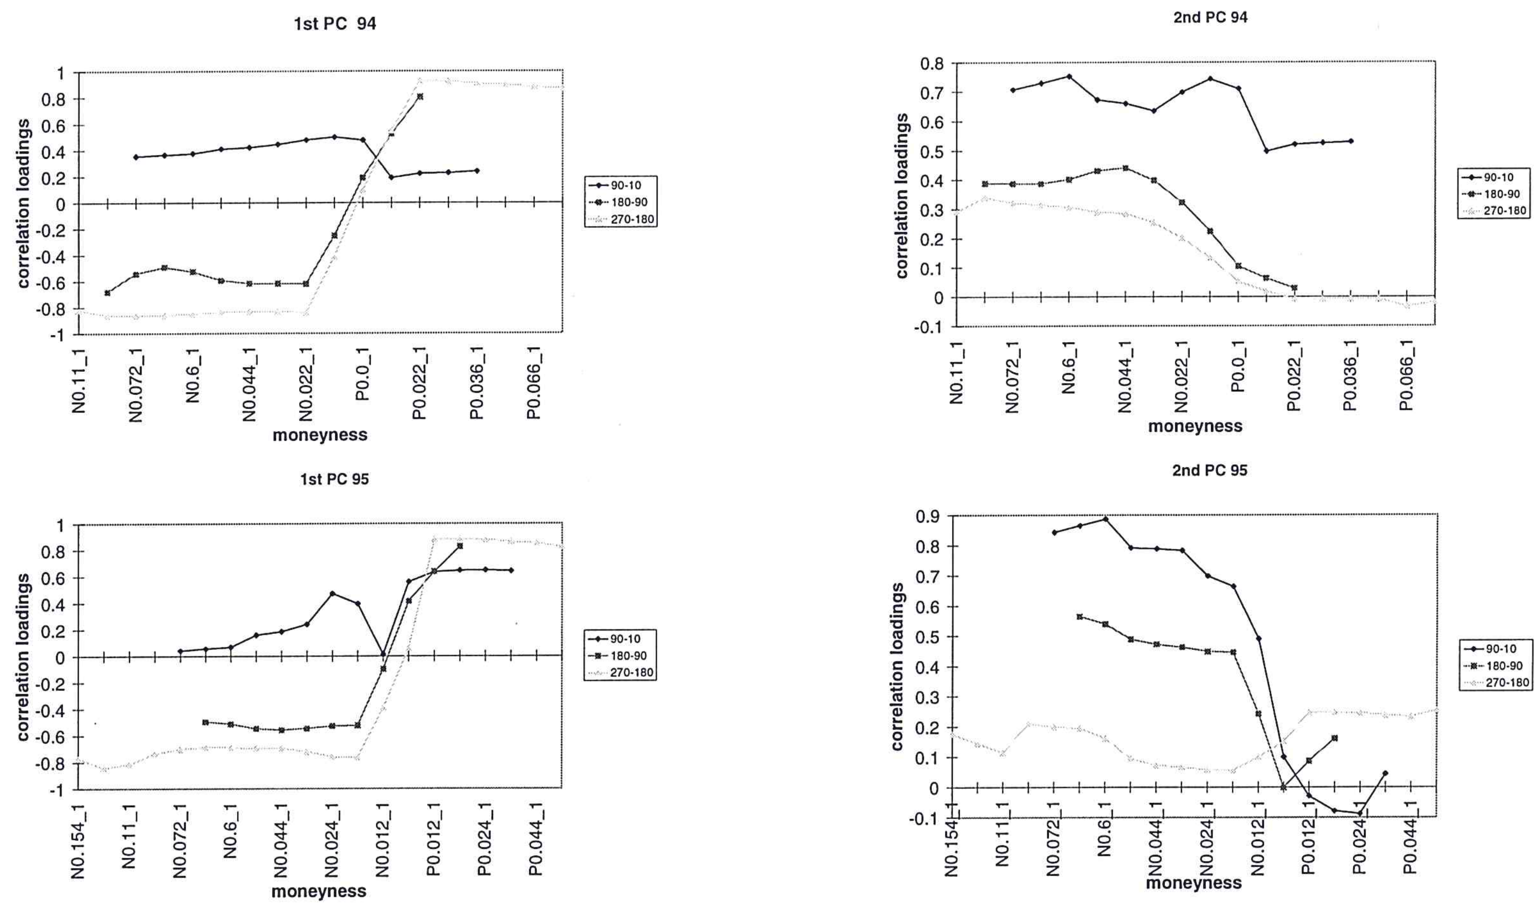

Figure 6 : First and Second PCs for 94 and 95 in the Moneyness Metric

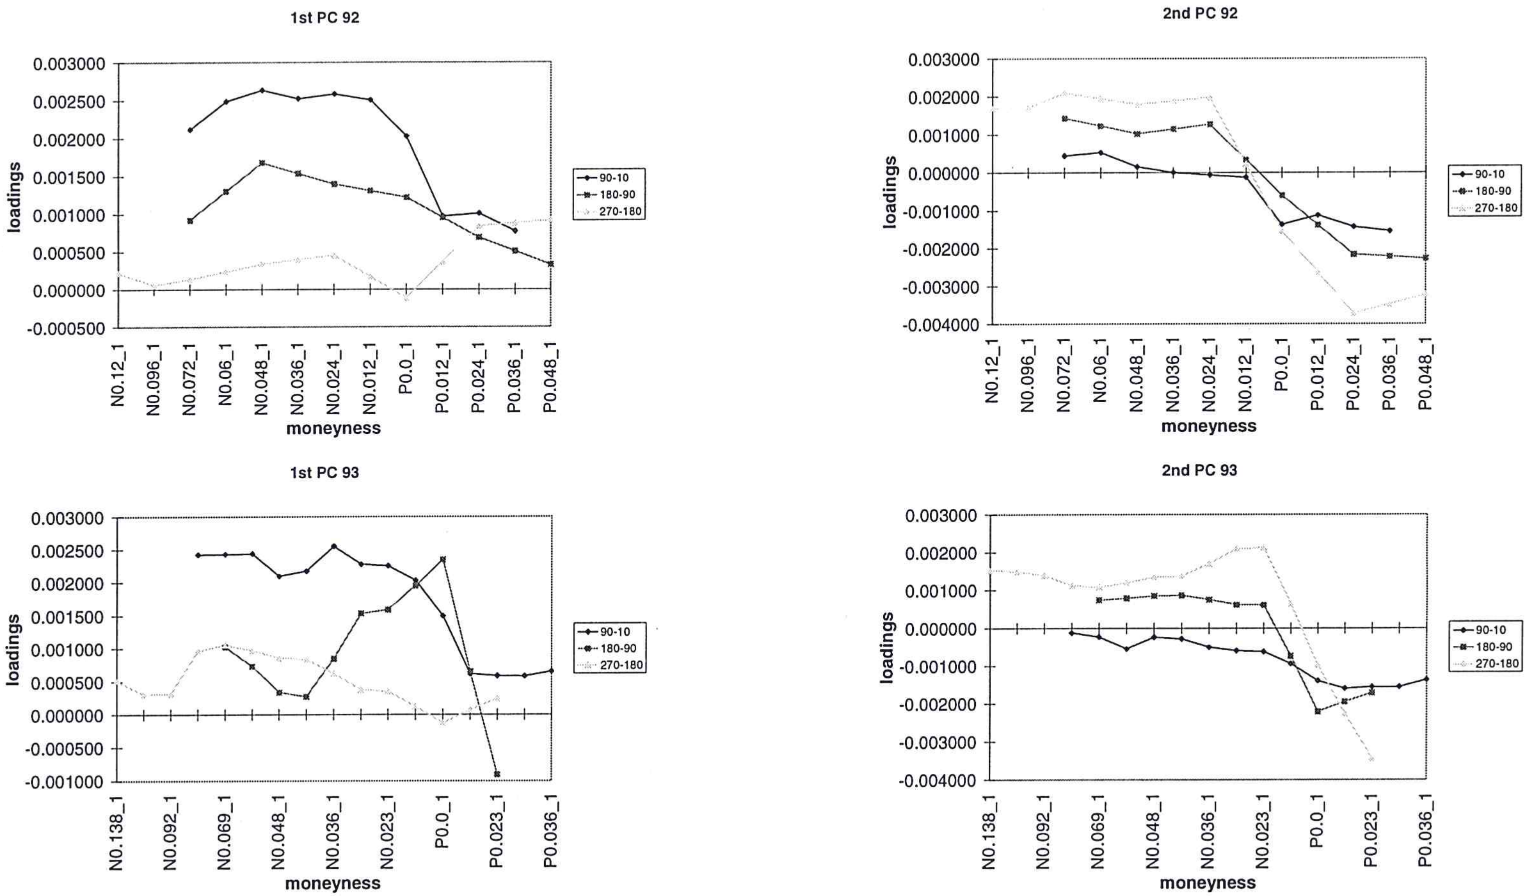

Figure 7 : First and Second Rotated PCs for 92 and 93 in the Moneyness Metric

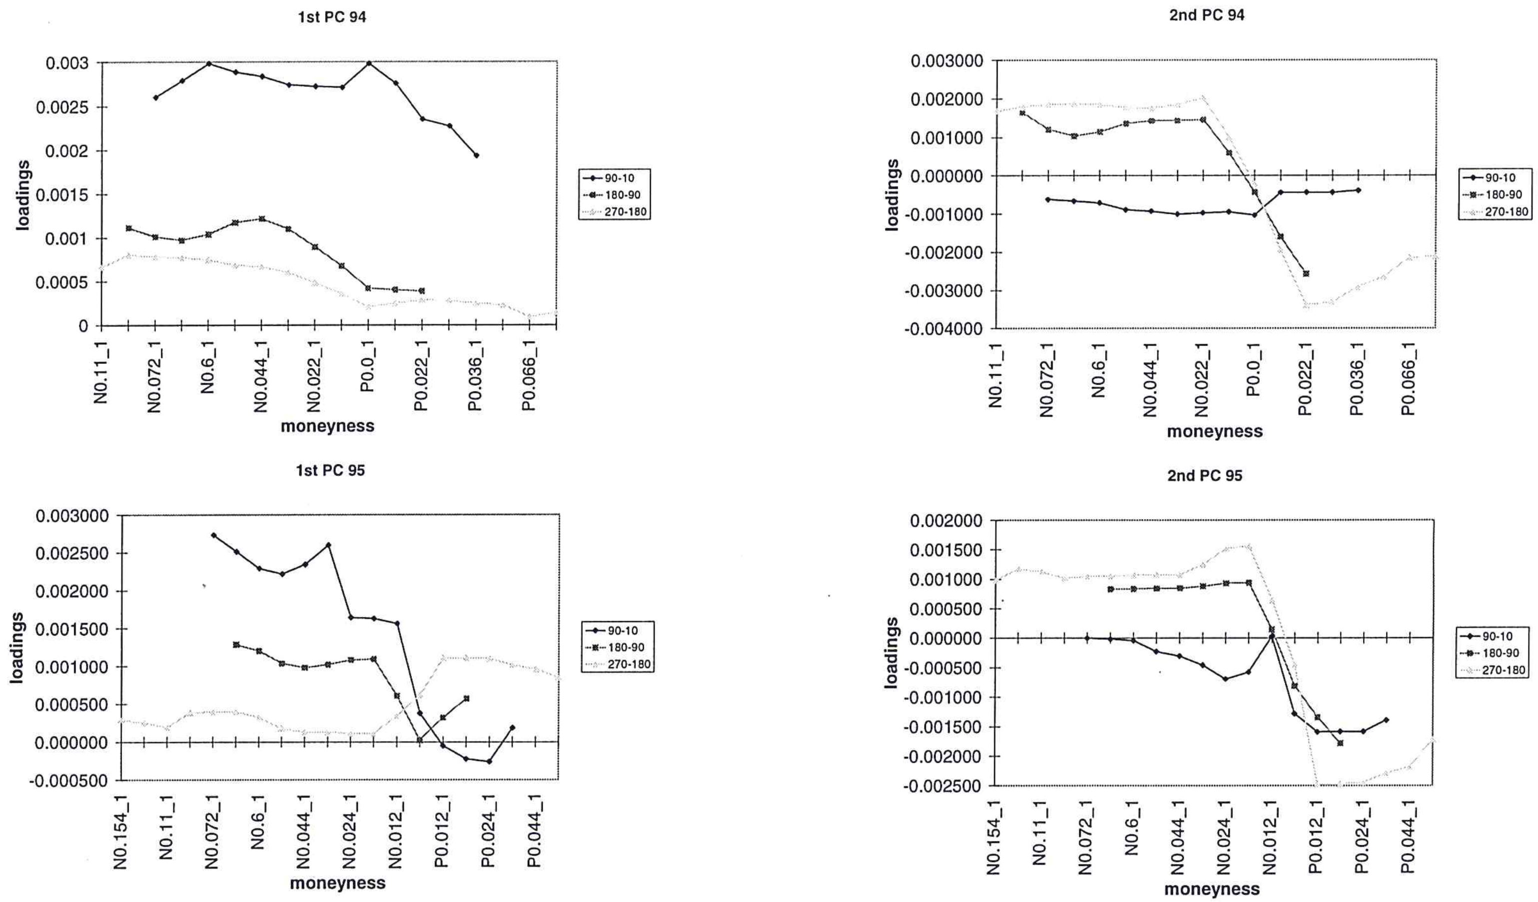

Figure 8 : First and Second Rotated PCs for 94 and 95 in the Moneyness Metric