The University of Reading

THE BUSINESS SCHOOL
FOR FINANCIAL MARKETS

Principal Component Analysis of Implied Volatility Smiles and Skews

Discussion Papers in Finance 2000-10

Carol Alexander

Chair of Risk Management,
ISMA Centre,
University of Reading, PO Box 242,
Reading RG6 6BA , UK

The University of Reading • ISMA Centre • Whiteknights • PO Box 242 • Reading RG6 6BA • UK
Tel: +44 (0)118 931 8239 • Fax: +44 (0)118 931 4741
Email: research@ismacentre.rdg.ac.uk • Web: www.ismacentre.rdg.ac.uk

Abstract

The construction of appropriate scenarios for movements in the implied volatility smile surface corresponding to movements in the underlying price is facilitated by using only a few key independent risk factors. The empirical model presented here is applied to equity index option markets to identify the current market regime and the price-volatility scenarios that should be applied. The framework is quite general and has applications to implied volatilities of many types of financial assets. It builds on the regime models of volatility introduced by Derman (1999) in two ways. First it provides an empirical investigation into the existence of the regimes that were hypothesized by Derman, and secondly it extends the linear parameterization of the skew that is implied by Derman's models to allow non-linear movements in fixed-strike implied volatilities as the underlying price changes.

JEL Classifications: C13, C22, C51, G19

Introduction

Following the Basle Accord Amendment in 1996 for the calculation of market risk capital using internal models, the Basle Committee on Banking Supervision have recommended two methods for generating a unified set of risk measures on a daily basis. These methods have become industry standards for measuring risk not only for external regulatory purposes, but also for internal risk management.

The first approach is to calculate a Value-at-Risk (VaR) measure, which is a lower percentile of an unrealized profit and loss distribution that is based on movements of the market risk factors over a fixed risk horizon. Central to most VaR models are large covariance matrices that encompass all risk positions - even historical simulation may employ covariance matrices for portfolio stress testing and scenario analysis. The efficient computation of large positive semi-definite covariance matrices is a difficult problem and simplifying assumptions are common.[1] Although large covariance matrices that are based on generalized autoregressive conditional heteroscedasticity (GARCH) models would have clear advantages[2] unfortunately multivariate GARCH models of more than a few dimensions are impossible to apply in practice. However Alexander (2000) uses key risk factor methods, similar to those used in this paper, to generate very large GARCH covariance matrices in an efficient and robust manner.

The focus of the present paper is on the second approach to modelling market risks that was recommended in the 1996 Amendment. That is to quantify the maximum loss of portfolios over a large set of scenarios for movements in the risk factors. The applicability of maximum loss measures depends on portfolio revaluation over all possible scenarios, including movements in both prices and implied volatilities of all risk factors. But given the huge number of market risk factors affecting the positions of a large financial institution, scenario-

[1] For example the RiskMetrics methodologies designed by JP Morgan use either simple equally weighted moving averages, or exponentially weighted moving averages with the same smoothing constant for all volatilities and correlations of returns. There are substantial limitations with both of these methods, described in Alexander (1996).

[2] Unlike moving average models, GARCH volatility term structures will converge to the long-term average volatility level. But the real beauty of GARCH stems from the fact that a stochastic volatility is built into the model, which is closer to the real world, yet it does not introduce an additional source of uncertainty and therefore delta hedging is still sufficient. See Engle (1982) and Bollerslev (1986).

based loss models may become very complex indeed. In fact their implementation becomes extraordinarily cumbersome, if not impossible, without making assumptions that restrict the possibilities for movements in the risk factors. In complex portfolios the computational burden of full revaluation over thousands of scenarios would be absolutely enormous, and certainly not possible to achieve within an acceptable time frame unless analytic price approximations and advanced sampling techniques are employed in conjunction with a restriction of the possibility set for scenarios.

Jamshidian and Zhu (1996) have shown how principal components may be used to improve computational efficiency for scenario based risk measures in large multi-currency portfolios. This paper uses similar ideas but applies the analysis to the construction of scenarios for an implied volatility smile or skew surface. An empirical model of equity index option markets shows that the most likely scenarios for volatility skew surfaces as the underlying price moves will depend on the current market regime.

At the heart of the model is the identification of a few key market risk factors that capture the most important independent sources of information in the data. Such an approach is computationally efficient because it allows an enormous reduction in the dimension of the problem whilst retaining a very high degree of accuracy. For example, in some of the FTSE100 index option data that is analyzed here there are sixty different fixed-strike fixed-maturity volatility series. However these time series are highly co-dependent and in fact they may be modelled, to a very high degree of accuracy, by only three series: the trend, tilt and curvature principal components of fixed-strike deviations form at-the-money volatility. Any movements that are not captured by these factors are deemed to be insignificant 'noise' in the system, and by cutting out this noise the empirical model for smile scenarios becomes more stable.

How should fixed-strike volatilities be changed as the underlying price moves?

Scenario based maximum loss calculations require at least the definition, if not the joint distribution, of scenarios for implied volatilities and underlying asset prices. In the absence of

Figure 1a: At-the-Money Volatility vs FTSE 100
(Daily Changes) May and June 1998

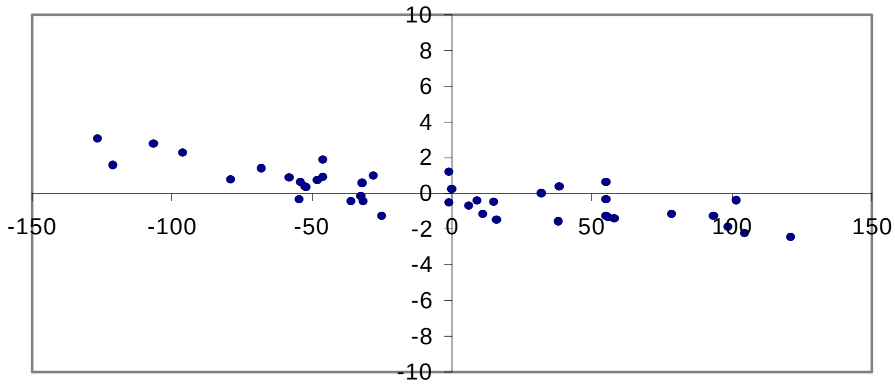

Figure 1b: At-the-Money Volatility vs FTSE 100
(Daily Changes) February and March 1998

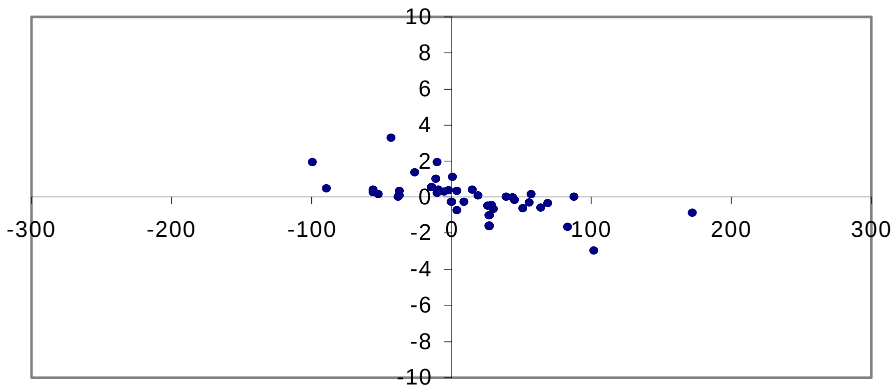

Figure 1c: At-the-Money Volatility vs FTSE 100
(Daily Changes) August and September 1998

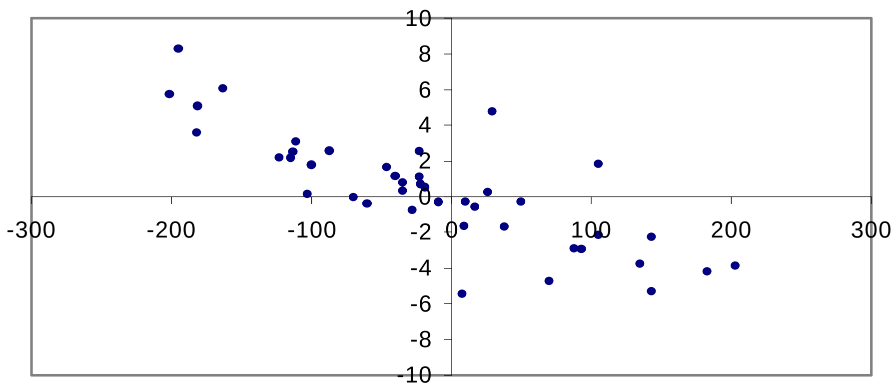

an effective model of how implied volatilities change with market price, these scenarios may be rather simplistic. The base scenario that the smile surface remains unchanged over all risk horizons is often augmented by a only a few simple scenarios. For example the 1996 Basle Amendment recommends parallel shifts in all volatilities that are assumed to be independent of movements in underlying prices.

But for equity options there is often a negative correlation between at-the-money volatility and the underlying price. This is clear from figure 1 which shows, for three different two month periods during 1998, a scatter plot of the daily changes in 1mth at-the-money volatility vs daily changes in index price for the FTSE100 European option. The periods chosen were (a) May and June 1998; (b) February and March 1998; and (c) August and September 1998.[3]

Casual observation of these scatter plots indicates a significant negative correlation between the 1mth implied volatility and the index price, but the strength of this correlation depends on the data period. Period (b) when the UK equity market was very stable and trending, shows less correlation than period (a), when daily movements in the FTSE100 index were limited to a 'normal' range; but the negative correlation is most obvious during the mini-crash period (c) that followed the LTCM crisis in July 1998. These observations are not peculiar to the 1mth at-the-money FTSE100 volatilities, and not just during the periods shown: negative correlations, of more or less strength depending on the data period, are also evident in other fixed term at-the-money volatilities and in other equity markets.

So realistic scenarios for at-the money volatility and index prices would be for movements in at-the-money volatility to occur in the opposite direction to the index price movements. But how large should these movements be in relation to each other? Does the answer depend on current market conditions? If so, how can we model the current market conditions to quantify the correlation effect? And what about the fixed-strike volatilities? Since positions are likely

[3] The fixed maturity implied volatility data used in this section have been obtained by linear interpolation between the two adjacent maturity option implied volatilities. However this presents a problem for the 1mth volatility series because often during the last few working days before expiry data on the near maturity option volatilities are totally unreliable. So the 1mth series rolls over to the next maturity, until the expiry date of the near-term option, and thereafter continues to be interpolated linearly between the two option volatilities of less than and greater than 1 month.

to move in- or out-of-the-money during the risk horizon, we need to know what scenarios are most probable for the whole volatility skew.

Derman's volatility regimes

Figure 2a shows the 1mth implied volatilities for European options of all strikes on the FTSE100 index for the period 4th January 1998 to 31st March 1999. The bold red line indicates the at-the-money volatility and the bold black line the FTSE100 index price (on the right-hand scale). Look at the movements in the index and the way that at-the-money volatility is behaving in relation to the index during the three different periods chosen in figure 1.

Figure 2a: Fixed-Strike Volatilities, At-the-Money Volatility and the Index Level

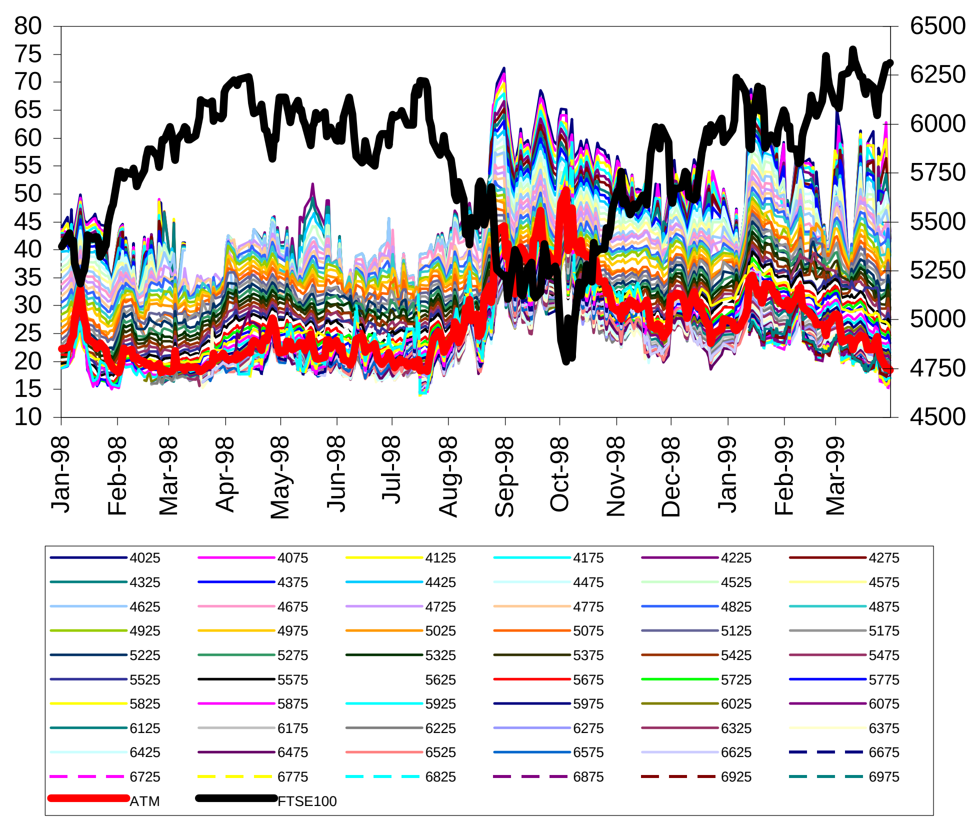

Observation of data similar to these, but on the S&P500 index option 3mth volatilities, has motivated Derman (1999) to formulate three different market regimes:

(a) *Range-bounded*, where future price moves are likely to be constrained within a certain range and there no significant change in realized volatility;
(b) *Trending*, where the level of the market is changing but in a stable manner so there is again little change in realized volatility in the long run; and
(c) *Jumpy*, where the probability of jumps in the price level is particularly high so realized volatility increases.

Different linear parameterizations of the volatility skew for pricing and hedging options apply in each regime. These are known as Derman's 'sticky' models, because each parameterization implies a different type of 'stickiness' for the local volatility in a binomial tree.[4] Denote by $\sigma_K(t)$ the implied volatility of an option with maturity t and strike K, $\sigma_{ATM}(t)$ the volatility of the t-maturity at-the-money option, S the current value of the index and σ_0 and S_0 the initial implied volatility and price used to calibrate the tree:

(a) In a range bounded market Derman proposes that skews are parameterized by the 'sticky strike' model:

$$\sigma_K(t) = \sigma_0 - b(t)\,(K-S_0) \tag{1a}$$

So fixed strike volatility $\sigma_K(t)$ is independent of the index level S.
Since $\sigma_{ATM}(t) = \sigma_0 - b(t)\,(S-S_0)$ this model implies that σ_{ATM} decreases as index increases.

(b) For a stable trending market skews are parameterized by the 'sticky delta' model:

$$\sigma_K(t) = \sigma_0 - b(t)\,(K-S) \tag{1b}$$

[4] The 'sticky strike' is so called because local volatilities are constant with respect to strike, changing only with moneyness; the 'sticky delta' model has local volatilities that are not constant with strike, but are constant with respect to moneyness or delta; and only in the 'sticky tree' model is there one, unique tree for all strikes and moneyness.

So fixed strike volatility $\sigma_K(t)$ increases with the index level S.

Since $\sigma_{ATM}(t) = \sigma_0$ this model implies that $\sigma_{ATM}(t)$ is independent of the index.

(c) In jumpy markets skews are parameterized by the 'sticky tree' model:

$$\sigma_K(t) = \sigma_0 - b(t)\,(K+S) + 2b(t)S_0 \tag{1c}$$

So fixed strike volatility $\sigma_K(t)$ decreases as the index increases.

Since $\sigma_{ATM}(t) = \sigma_0 - 2b(t)\,(S-S_0)$, the at-the-money volatility $\sigma_{ATM}(t)$ also decreases as index increases, and twice as fast as the fixed strike volatilities.

Fixed-strike volatility deviations from at-the-money volatility

Time series data such as that shown in figure 2a should contain all the information necessary to estimate the skew parameterization that is appropriate for the current market regime. But there are around 60 different strikes represented there, and their volatilities form a correlated, ordered system that is similar to a term structure. It is therefore natural to consider using principal component analysis to identify the main independent sources of information. Both analytic simplicity and computational efficiency would result from a model that is based only on these key risk factors.

Principal component analysis of the volatility skew has been used before, by Derman and Kamal (1997). However their work is based on quite different data to that shown in figure 2a.[5] Time series data on fixed strike or fixed delta volatilities often display very much negative autocorrelation, possibly because markets over-react, so the ‘noise’ in daily changes

[5] Dermand and Kamal use weekly mid-market volatility of S&P500 index options from May 1994 to September 1997 where the surface is specified by 12 numbers corresponding to three different deltas for 1mth, 3mth, 6mth and 12mth maturities; and daily Nikkei 225 index volatility from September 1994 to May 1997 for 9 deltas and 5 different maturities. For each of these markets they analyze the principal components of the changes in the whole implied volatility surface.

of fixed strike volatilities is a problem. Therefore a principal components analysis of daily changes in fixed-strike volatilities may not give very good results.

Figure 2b: Deviations of Fixed-Strike Volatility from At-the-Money Volatility (1mth)

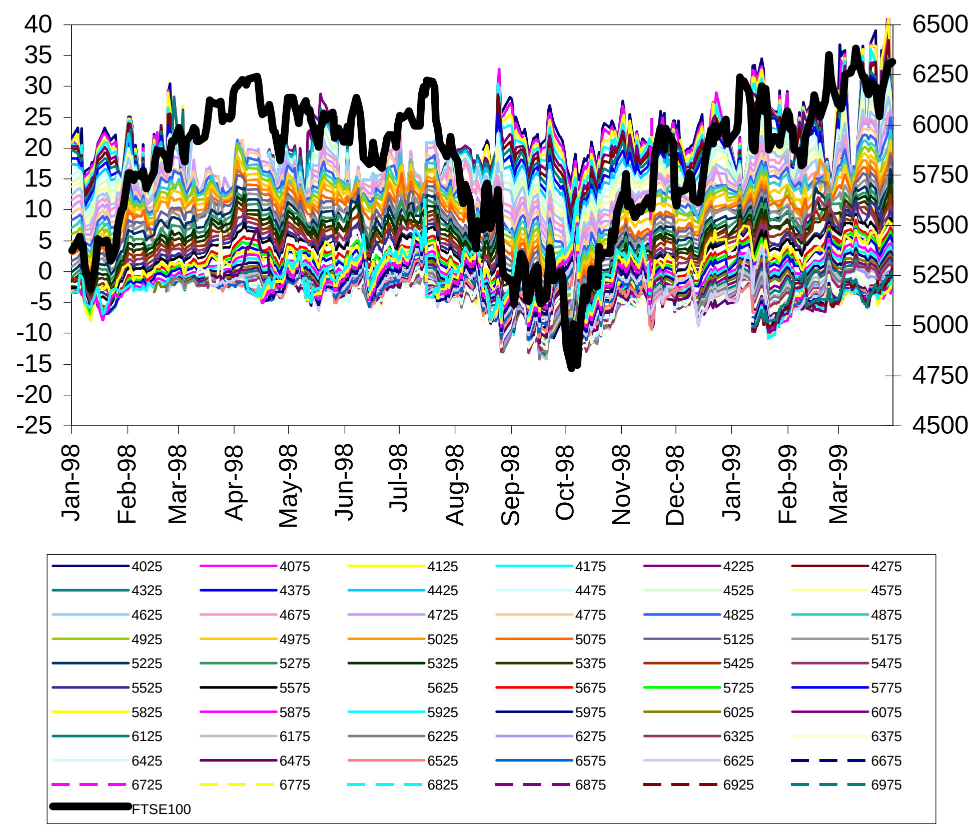

But look at the deviations of fixed strike volatilities from at-the-money volatility, shown in figure 2b. These display less negative autocorrelation, they are even more highly correlated and ordered than the fixed strike volatilities themselves, and their positive correlation with the index is very evident indeed during the whole period.

The reason for this becomes evident when (1a) – (1c) are rewritten in terms of fixed-strike volatility deviations from at-the-money volatility $\sigma_K(t) - \sigma_{ATM}(t)$. Each of Derman's models yields the same relationship between fixed-strike volatility deviations from at-the-money volatility and the current index price, viz.:

$$\sigma_K(t) - \sigma_{ATM}(t) = -b(t)\,(K-S) \tag{2}$$

So all three models imply the same, positive correlation between the index and the skew deviations $\sigma_K(t) - \sigma_{ATM}(t)$. In fact an alternative formulation of Derman's sticky models is (2) with a different specification for the behaviour of at-the-money volatility in relation to the index in each regime, viz.

(a) Range-bounded: $\sigma_{ATM}(t) = \sigma_0 - b(t)\,(S - S_0)$

(b) Stable trending: $\sigma_{ATM}(t) = \sigma_0$

(c) Jumpy: $\sigma_{ATM}(t) = \sigma_0 - 2b(t)(S - S_0)$.

Effective methods for identification of the current market regime

The above formulation of Derman's regime models suggests that one might perform an empirical investigation into which regime currently prevails by estimating linear regressions of the form:

$$\Delta\sigma_{ATM}(t) = \alpha(t) + \beta(t)\Delta S + \varepsilon(t) \tag{3}$$

where $\Delta\sigma_{ATM}(t)$ denotes the daily change in at-the-money volatility of maturity t and ΔS is the daily change in the index. In general, due to the negative correlation, each $\beta(t)$ will be negative. But if all the coefficients $\beta(t)$ are insignificantly different from zero the market is stable and trending, so the sticky delta model should be used. A signal that the market has entered a different regime occurs when $\beta(t)$ undergoes a significant change in value. In a jumpy market that is characterized by the sticky tree model, the value of $\beta(t)$ will be approximately twice the value that it takes in a range-bounded market where the sticky strike model is valid.

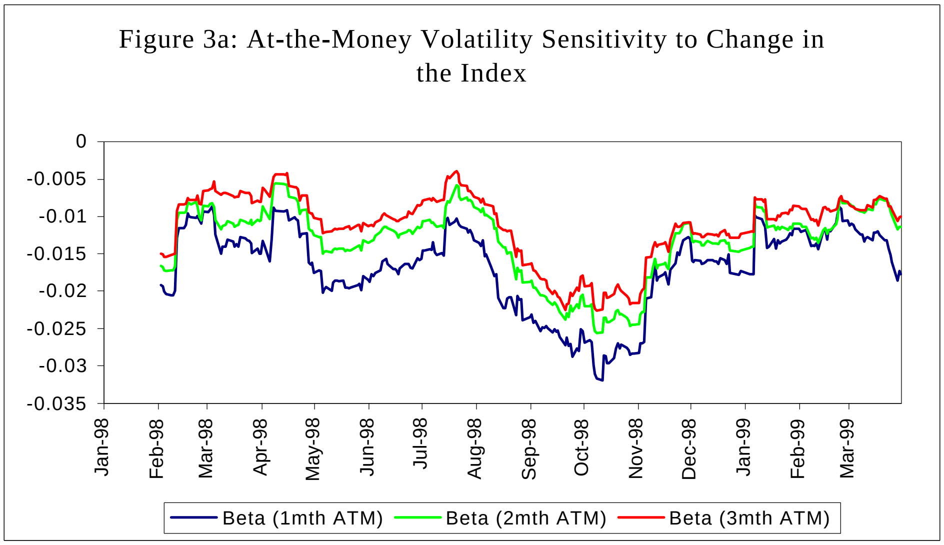

Figure 3a shows the values obtained for β(t) for t = 1mth, 2mths and 3mths. In order to capture the current market conditions one month of daily data is used in each regression. These regressions were rolled over the whole period from 4th January 1998 to 31st March 1999 and each time the coefficient and its t-statistic are recorded.

The response of at-the-money volatility to changes in the underlying index level increases as options approach expiry, and this fact is reflected in figure 3a since at all times

$$|\beta(1mth)| > |\beta(2mth)| > |\beta(3mth)|.$$

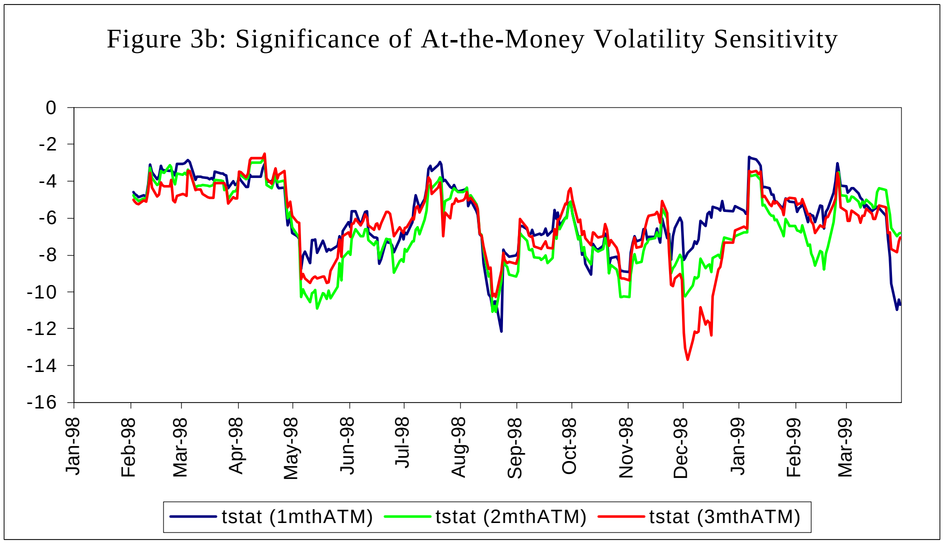

However no such order is apparent in the accompanying t-statistics, shown in figure 3b, so the negative correlation between at-the-money volatility and index price is not a simple function of the maturity of volatility.

Casual observation of figure 2a has indicated that February and March 1998 might be characterized as a stable and trending market. Derman's sticky model for this regime has no correlation between at-the-money volatility changes and underlying price changes, and now figure 3b provides quantifiable evidence of this: During February and March 1998 the t-statistics on the β coefficients are less significant than at other times.[6] Two other periods were picked out in the earlier discussion: May and June 1998, when the market seemed to be operating in a range-bounded regime, and the mini-crash period that began after the LTCM crisis in July 1998 and initiated a very jumpy market until the November of that year. Recall that Derman's sticky models have at-the-money volatility responding to price moves twice as much in the jumpy regime as in the range bounded regime. And again, from figure 3a, it is apparent that the values of the β coefficients during the mini-crash period, although not exactly double their values during May and June 1998, were far greater than at any other time.

A rapid decline in β, for all maturity volatilities, occurred at the end of July 1998, at the time of the equity market mini-crash that was precipitated by the LTCM crisis. Thus the model is providing a leading indicator of a change in market regime. It was not until November 1998 that the level of β returned to more normal levels, when the market appears to pass back into a range-bounded regime.

Using key risk factors to formulate appropriate skew scenarios

The simple regressions just described may be used to identify the current volatility regime, and to forewarn risk managers of any change in market conditions. It is now shown how such information may be put to practical advantage in the construction of the 'most likely' movements in the implied volatility skew surface. We now ask, which type of skew scenarios should accompany the scenarios on movements in the underlying? Are simple static or parallel shift scenarios for the volatility skew appropriate at the moment? If so, is it the volatility by strike that should remain static, so the volatility by moneyness or delta has a parallel shift? Or is it volatility by delta that is static, which is equivalent to a parallel shift in volatility by strike? But perhaps one should be placing more importance on scenarios that encompass changes in the tilt or curvature of the volatility skew? If so, at which end: should in-the-money volatilities be changed as much as out-of-the-money volatilities?

The following discussion illustrates how all these questions can be answered by an empirical model of the relationship between the equity price and the key risk factors of the skew. Derman's models are based a linear parameterization of the skew given by (2). For any given maturity, the deviations of all fixed strike volatilities from at-the-money volatility will change by the same amount b(t) as the index level changes, as shown in figure 4a. Four strikes are marked on this figure: a low strike K_L, the initial at-the-money strike K_1, the new at-the-money strike after the index level moves up K_2, and a high strike K_H. The volatilities at each of these strikes are shown in figure 4b, before and after an assumed unit rise in index level ($\Delta S = 1$). In each of the three market regimes the range of the skew between K_L and K_H, that is $\sigma_L - \sigma_H$, will be the same after the rise in index level. Thus all of Derman's model's imply a parallel shift scenario for the skew by strike.

Figure 4a: Parallel Shift in Skew Deviations as Price Moves Up

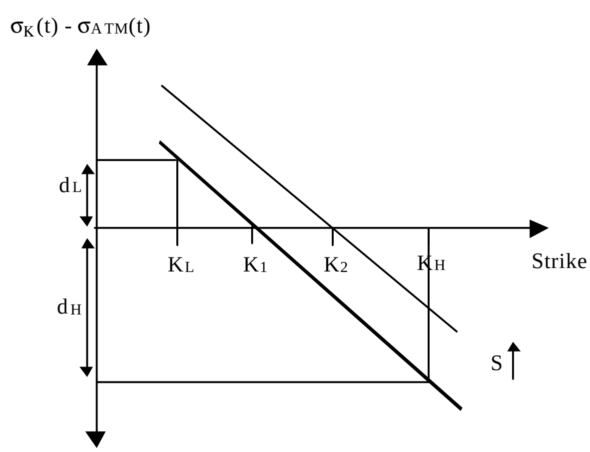

[6] The 99% significance level is approximately 2.5.

The extent of the parallel shift depends on the relationship between the original at-the-money volatility σ_1 and the new at-the-money volatility σ_2, and this will be defined by the current market regime. In a range bounded market $\sigma_2 = \sigma_1 - b(t)$, but fixed-strike volatilities have all increased by the same amount b(t), so a static scenario for the skew by strike should be applied, as depicted in figure 4b. When the market is stable and trending, $\sigma_2 = \sigma_1$ and there is an *upwards* shift of b(t) in all fixed-strike volatilities. Finally, in a jumpy market $\sigma_2 = \sigma_1 - 2b(t)$, so a parallel shift *downwards* of b(t) in the skew by strike should be applied.

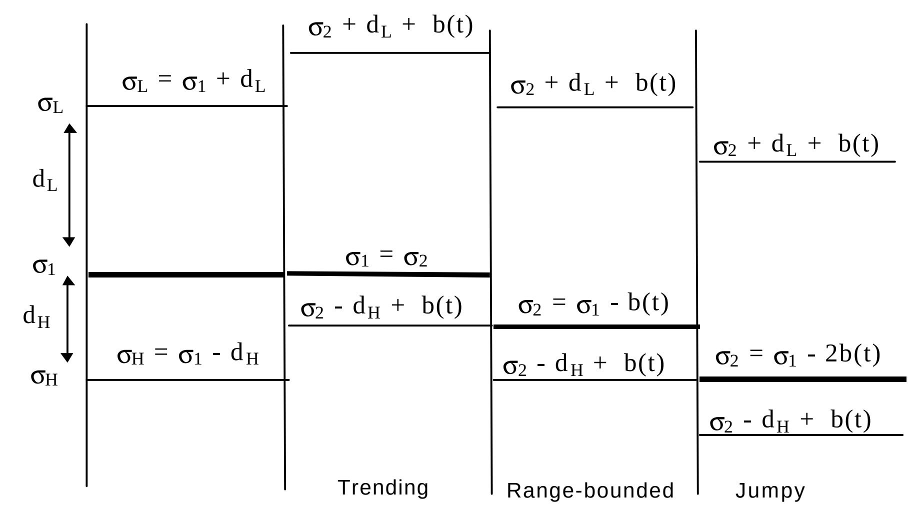

Figure 4b: Parallel Shifts in Fixed-Strike Volatilities as Price Moves Up

Whilst a linear parameterization of the skew may be good approximation for the 3mth or longer maturities, empirical observations show that it may not be very realistic at the shorter end. Figure 5 shows the correlations from simple cross-section regressions based on (2). It is clear that whilst the skew is fairly linear at the 3mth maturity, it becomes quite non-linear at the 1mth maturity, particularly during the summer of 1998. So the parallel shift scenarios for volatility skews that are a consequence of Derman's models may be reasonable for 3mth volatilities, but for shorter-term volatilities a simple, effective non-linear model of the skew would be advantageous.

Such a model can be based on a principal component analysis of $\Delta(\sigma_K(t) - \sigma_{ATM}(t))$, the daily changes in t-maturity fixed-strike volatility deviations from t-maturity at-the-money volatility. In this way the key risk factors for the volatility skew will be identified and consequently used in an empirical justification for skew scenarios that encompass more change at either or both of the wings. Whether one should change volatilities at the out-of-the-money wing or at the in-the-money wing of the skew, or both, will be shown to depend on the current market conditions.

Principal component analysis of $\Delta(\sigma_K(t) - \sigma_{ATM}(t))$ has given some excellent results. For fixed maturity volatility skews in the FTSE100 index option market during most of 1998, the parallel shift component accounted for around 65-80% of the variation, the tilt component explained a further 5 to 15% of the variation, and the curvature component another 5% or so of the variation. The precise figures depend on the maturity of the volatility (1mth, 2mth or 3mth) and the exact period in time that the principal components were measured. But generally speaking 80-90% of the total variation in skew deviations can be explained by just three key risk factors: parallel shifts, tilts and curvature changes.[7]

[7] For example, the principal component analysis for 3mth implied volatility skew deviations over the whole data period gives the following output. Note that sparse trading in very out-of-the money options implies that the extreme low strike volatilities show less correlation with the rest of the system, and this is reflected by their lower factor weights on the first component.

Component	Eigenvalue	Cumulative R^2
P1	13.3574	0.742078
P2	2.257596	0.8675
P3	0.691317	0.905906

Factor Weights

	P1	P2	P3
4225	0.53906	0.74624	0.26712
4325	0.6436	0.7037	0.1862
4425	0.67858	0.58105	0.035155
4525	0.8194	0.48822	-0.03331
4625	0.84751	0.34675	-0.19671

Figure 5: R-Squared from Linear Skew Parameterization

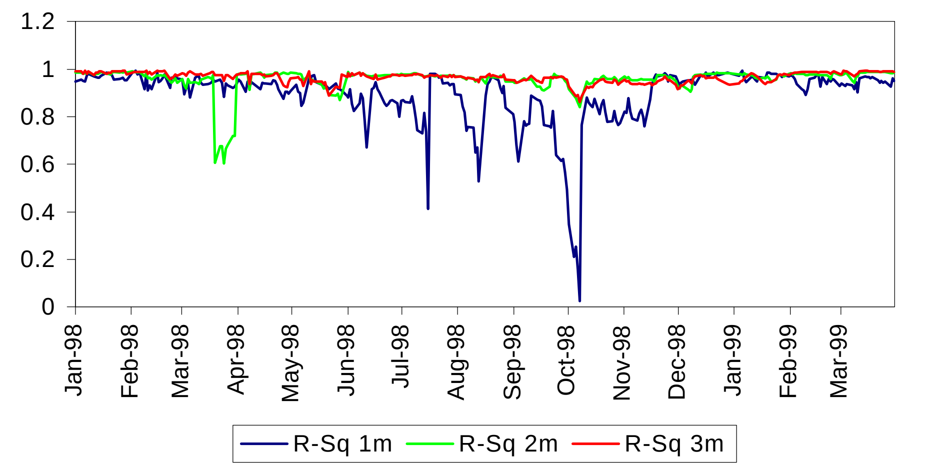

This identification of the important risk factors allows one to quantify the expected movements in the volatility skew as the index moves under different market circumstances. The first stage is to represent fixed-strike skew deviations by three principal components:

$$\Delta(\sigma_K(t) - \sigma_{ATM}(t)) = \omega_{K,1}(t)\ P_1(t)\ + \omega_{K,2}(t)\ P_2(t)\ + \omega_{K,3}(t)\ P_3(t) \qquad (4a)$$

4725	0.86724	0.1287	-0.41161
4825	0.86634	0.017412	-0.43254
4925	0.80957	-0.01649	-0.28777
5025	0.9408	-0.18548	0.068028
5125	0.92639	-0.22766	0.13049
5225	0.92764	-0.21065	0.12154
5325	0.93927	-0.22396	0.14343
5425	0.93046	-0.25167	0.16246
5525	0.90232	-0.20613	0.017523
5625	0.94478	-0.2214	0.073863
5725	0.94202	-0.22928	0.073997
5825	0.93583	-0.22818	0.074602
5925	0.90699	-0.22788	0.068758

The second part of the model employs simple linear regressions of each component P_i (i = 1, 2, or 3) on the daily changes ΔS in the index, viz.:

$$P_i\,(t) = \gamma_{0,i}\,(t) + \gamma_i\,(t)\,\Delta S + \eta_i\,(t) \tag{4b}$$

where t is the volatility maturity (1mth, 2mth or 3mth). Thus the movements a t-maturity volatility at strike K consequent to a change in index level will be determined by the factor weights $\omega_{K,i}$ and the sensitivities of the key risk factors to index movements, $\gamma_i\,(t)$ for i = 1, 2, 3. Note that Derman's models are a special case of this model, where there is just one principal component in the representation (4a) and so in Derman's models a perfect correlation is assumed between all fixed-strike volatility deviations from at-the-money volatility.

In order to capture the current market conditions, the regressions (4b) have been performed using just one month of the FTSE 100 index data. These regressions were rolled over the whole period from 4th January 1998 to 31st March 1999, and each time the coefficients $\gamma_i(t)$ are recorded, for i = 1, 2, and 3 and t = 1mth, 2mths and 3mths. The statistical significance of these coefficients is as interesting as their actual value. In fact it is the significance levels that provide the important information for risk managers when coming to a decision about which types of risk should be the current focus.

Figure 6: Significance of the Key Risk Factors

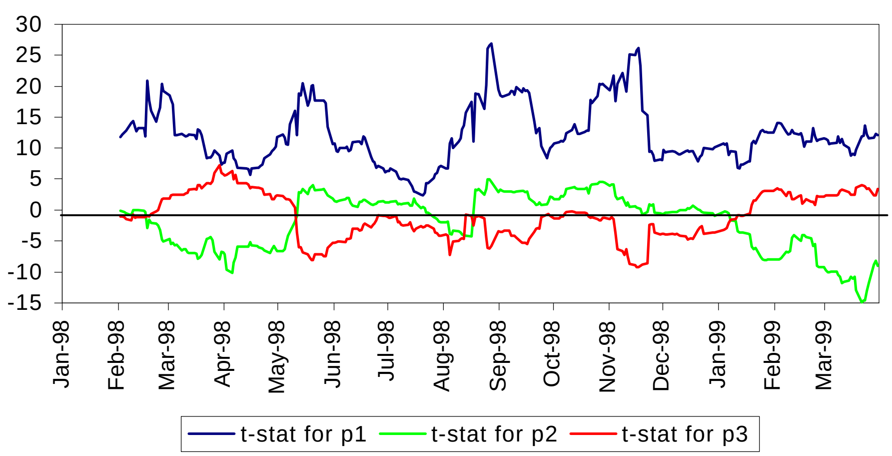

Figure 6 shows the t-statistics on $\gamma_i(t)$ for i = 1, 2 and 3 and t = 1mth from one month rolling regressions (4b). Clearly γ_1, which captures a parallel shift in all fixed-strike volatility deviations, is significant throughout the period, always positive and particularly important during the mini-crash period and the consequent market recovery. But the tilt component γ_2 is much less significant. It is only playing a really important role during the spring of 1998 and again in the spring of 1999. At both these times the tilt has a negative relationship with index moves, indicating that as the index moves up the low strike deviations will decrease and the high strike deviations will increase. It is interesting to see that γ_3, which captures the curvature component of the skew deviations, almost always has the opposite sign to the tilt coefficient.

The implication of these observations, for constructing scenarios to model the likely behaviour of the volatility skew as the index moves will now be explained in the two cases that arise empirically. The first case is when $\gamma_1 > 0$, $\gamma_2 < 0$ and $\gamma_3 > 0$ and the second case is when $\gamma_1 > 0$, $\gamma_2 > 0$ and $\gamma_3 < 0$.

Figure 7a illustrates how the skew deviations move in response to an *upward* movement in the index when $\gamma_1 > 0$, $\gamma_2 < 0$ and $\gamma_3 > 0$. In this case the upward movements in volatility deviations from at-the-money volatility are far greater at high strikes than at low strikes. In fact a result of the upward movement in the index is that one of the high strike deviations, at strike K_2 say, will change from a negative value to a value of zero because the at-the-money strike has moved from K_1 to K_2. Strikes above K_2 will still have volatilities that are lower than the at-the-money volatility, strikes between K_1 and K_2 now have volatilities that are above at-the-money volatility, and strikes below K_1 always have and remain to have volatilities above the at-the-money volatility. For the lowest strikes there will be little change: their volatility deviation from the new at-the-money volatility is about the same as it was before the index move.

Figure 7a: Non-Parallel Shift in Skew Deviations as Price Moves Up

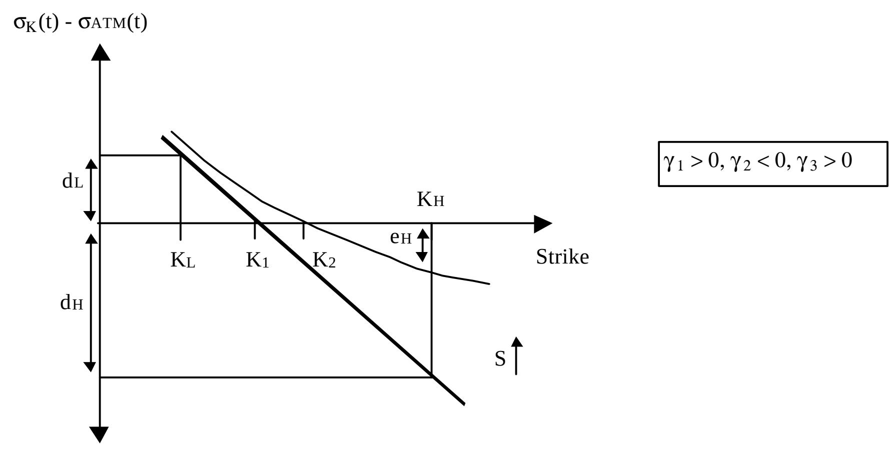

Figure 7b translates the effect of index moves on fixed-strike volatility deviations from at-the-money volatility, into movements in the actual fixed strike volatilities. It is a generalization of figure 6b, using the non-linear model of the skew (4a) and (4b), to accommodate scenarios that are more general than simple parallel shifts.

Figure 7b: Range Narrowing of Fixed-Strike Volatilities as Price Moves Up

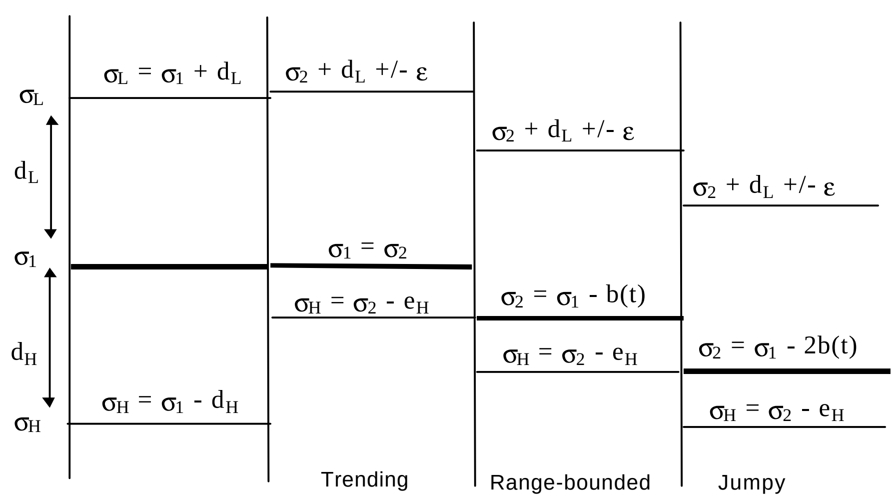

As before the three volatility regimes are shown according as, after a unit rise in the index level, the new at-the-money volatility σ_2 equals the original at-the-money volatility σ_1 (in a stable trending market), or $\sigma_2 = \sigma_1 - b(t)$ (for a range-bounded market), or $\sigma_2 = \sigma_1 - 2b(t)$ (a jumpy market). The difference between this figure and figure 6b is that there is no longer a uniform response b(t) for all fixed-strike volatility deviations when the index level changes. In fact figure 7a shows that there will in fact be little change in low-strike volatility deviations from at-the-money volatility, whereas high strike volatility deviations from at-the-money volatility will change considerably. Therefore the range of the skew between K_L and K_H, that is $\sigma_L - \sigma_H$, will become *narrower* after the rise in index level. Figure 7b also shows that it is the current volatility regime that determines whether the movement should occur at the high in-the-money strikes, the low out-of-the-money strikes, or both.

Similar remarks apply to the effect of a *downward* move in the index. It is left to the reader to depict the effect of a unit decrease in the index level on (a) fixed-strike deviations from at-the-money volatility, and (b) fixed-strike volatilities themselves, again when $\gamma_1 > 0$, $\gamma_2 < 0$ and $\gamma_3 > 0$. The net effect is that the range of the skew will widen as the index moves down with most of the movement in fixed-strike volatilities coming from the low strikes whereas the high strike volatilities move very little.

The model has shown that when $\gamma_1 > 0$, $\gamma_2 < 0$ and $\gamma_3 > 0$ there will be less movement in high strike volatilities and more in the low strikes. If one refers back to figure 2a it is clear that much of the time the low strike volatilities are indeed moving down and up considerably as the index moves up and down. There is much less movement in high strike volatilities, except possibly during the mini-crash period in the late summer of 1998. The widening and narrowing effects in the skew are also quite obvious in figure 2a, particularly during the last few months of the data period. At times like this the simple parallel shift scenarios for the skew, as implied by Derman's model, would not be sufficiently general. Instead, the non-linear model (4a) and (4b) can be used to build non-parallel shift skew scenarios as described above, that are more appropriate for these market conditions.

Now consider what happens when the major risk factor is still the trend component, but when the tilt and curvature components of the skew deviations have the opposite influence to that just discussed. That is when $\gamma_1 > 0$, $\gamma_2 > 0$ and $\gamma_3 < 0$. We have already seen from figure 6 that the parameter values $\gamma_1 > 0$, $\gamma_2 > 0$ and $\gamma_3 < 0$ only occurred during the mini-crash and recovery period, so it may be assumed that the market will be in a jumpy regime.

Figure 8 shows the effect of a unit increase in the index level on (a) fixed-strike volatility deviations from at-the-money volatility, and (b) fixed-strike volatilities themselves, when $\gamma_1 > 0$, $\gamma_2 > 0$ and $\gamma_3 < 0$. The net effect from all three principal components is for high strike volatility deviations from at-the-money volatility to change very little, whereas the low strike deviations will increase further. Thus the range of the skew will *widen* as the index moves up and narrow as the index moves down. The at-the-money volatility response, which depends on the current market regime, will determine whether the movement occurs at low strikes, high strikes or both. Although for completeness all three regimes are shown in figure 8b, it has been observed that the case $\gamma_1 > 0$, $\gamma_2 > 0$ and $\gamma_3 < 0$ only arises during a jumpy market regime, so the fixed-strike volatility movements in the last column of figure 8b are appropriate. When the index level increases high strike volatilities should be adjusted down, about the same amount as the at-the-money volatility. But low strike volatilities should be adjusted less far down, or even upwards.[8] Similarly, if the index level falls, high strike volatilities should be adjusted up, about the same amount as the at-the-money volatility, but low strike volatilities should be adjusted less, and they may even move downwards.

[8] From figure 8b, it is clear that low strike volatilities will move down as the index increases (and up as the index decreases) if and only if $d_L > e_L - 2b(t)$.

Figure 8a: Non-Parallel Shift in Skew Deviations as Price Moves Up

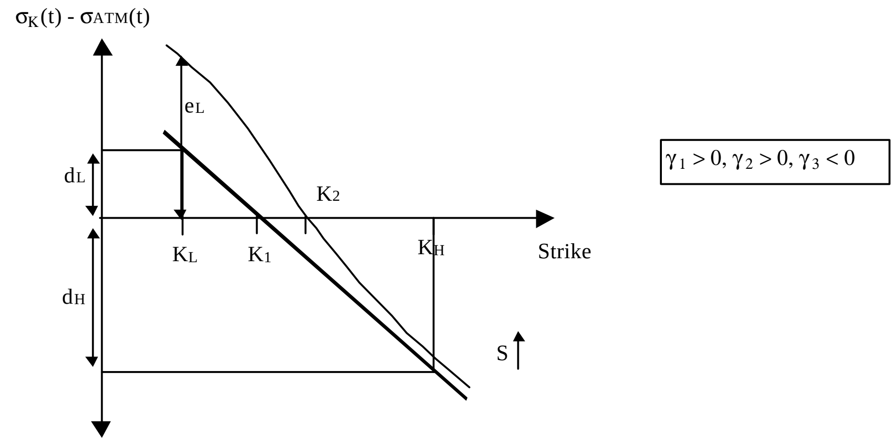

Figure 8b: Range Widening of Fixed-Strike Volatilities as Price Moves Up

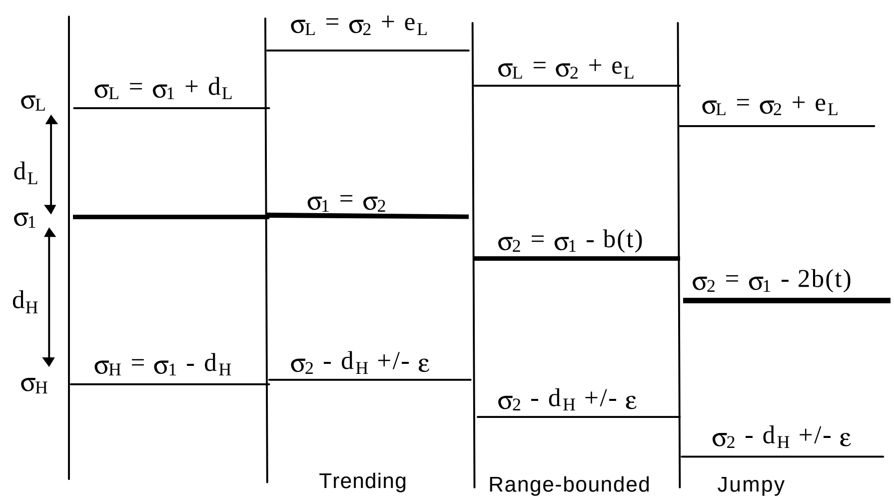

Summary and Conclusions

To summarize the modelling procedure, first regression models of the form (3) that are based only on recent market data are used to indicate which volatility regime is likely to prevail in the near future, and the relevant sensitivity of at-the-money volatility to changes in the index level. Then the key risk factors of a volatility skew are quantified by the trend, tilt and curvature components of the deviations of fixed-strike volatilities from at-the-money volatility, as in the model (4a). The response of fixed-strike volatilities to changes in the index level depends on which of these key risk factors are important in the current market conditions and this information is obtained from regression models of the form (4b). Typically the trend component will always be the most significant risk factor. If it is the only significant risk factor then the parallel shift scenarios that are implied by Derman's models will apply. But when the tilt or curvature are also significant risk factors, adjustments should be made for greater changes at out-of-the-money volatilities and perhaps also at in-the-money volatilities in the skew. The magnitude and direction of such changes are determined by the sensitivities of the three key risk factors to changes in the index level. When measured by models of the form (4b) these sensitivities are found to depend very much on the current market regime.

Application of this model to daily data on the FTSE 100 European index option has produced some likely scenarios for FTSE 100 volatility skews and indicated the circumstances in which they should be applied. Typically the range of volatility in the skew with respect to strike will widen as the index level decreases and narrow as the index level increases. When the market is in a stable trending or range-bounded regime most of the change should be coming from the low strike out-of-the money volatilities. But in a market crash the high strike in-the-money volatilities will also move in the opposite direction to the index.

It is a common problem in risk management today that risk measures and pricing models are being applied to a very large set of scenarios based on movements in all possible risk factors. The dimensions are so large that the computations become extremely slow and cumbersome, so it is quite common that over-simplistic assumptions will be made. For example simple parallel shifts in smile surfaces are often the only scenarios considered, and these may be assumed to be independent of the underlying price movements.

The approach taken in this paper is to simplify the analysis by using only a few key, independent risk factors of the volatility smile surface. Thus dimensions are considerably reduced, from up to sixty fixed-strike volatilities to only three principal components, for each maturity in the smile surface. An empirical model of price-volatility scenarios has been described and implemented using daily data on the FTSE100 index option from 4th January 1998 to 31st March 1999. The analysis is greatly simplified by the fact that it is based on only a few orthogonal risk factors, but these risk factors are still capturing most of the risk, so there is little loss of accuracy.

Non-linear skew parameterizations and non-parallel shift scenarios for the volatility skew are accommodated very easily in this framework. And the empirical nature of the model allows the actual quantification of appropriate moves in the volatility skew as the underlying price changes. The model first provides a leading indicator of the expected market regime. Then, given the expected regime, and for a given change in underlying price, the regression models may be used to provide a numerical forecast of the most likely change in the at-the-money volatility, and in all fixed-strike volatilities, of any maturity.

The general method used here may be applied to other equity index markets and to other types of options, and this is the subject of ongoing research. It is possible that these methods could be used to determine the swaption volatility skew as a function of the key risk factors of cap volatility skews. And the use of three key risk factors in a non-linear model of price-volatility scenarios should be particularly useful in currency option markets, where smile models will be better modelled by a non-linear parameterization.

Acknowledgements

I am greatly indebted to Emmanuel Derman, Goldman Sachs, Head of Risk Management in New York for many very useful discussions; and to Jacques Pezier, General Manager of CAL-FP Bank, London for his insightful comments.

References

Alexander C.O. (1996) "Evaluating the use of RiskMetrics™ as a risk measurement tool for your operation : What are its advantages and limitations" *Derivatives Use, Trading and Regulation* 2:3 pp277-285

Alexander, C.O. (2000) "A primer on the orthogonal GARCH model" Draft in pdf on www.ismacentre.reading.ac.uk

Bollerslev, T. (1986) "Generalized autoregressive conditional heteroskedasticity". *Journal of Econometrics* 31 pp 307-327.

Derman, E. (1999) "Volatility regimes" *RISK Magazine* 12:4 pp55-59

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

Engle, R.F. (1982) "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation". *Econometrica* 50:4 pp987-1007.

Jamshidian, F. and Y. Zhu (1996) "Scenario simulation: Theory and methodology" *Finance and Stochastics* 1:1 pp43-67