



Munich Personal RePEc Archive

Higher-order volatility: time series

Carey, Alexander

12 January 2010

Online at <https://mpra.ub.uni-muenchen.de/21087/>

MPRA Paper No. 21087, posted 04 Mar 2010 14:20 UTC

HIGHER-ORDER VOLATILITY: TIME SERIES

Alexander Carey

6 rue de l'Abbé Groult
75015 Paris, France
alexpcarey@gmail.com

January 12, 2010

This paper presents time-series of higher-order volatilities for the S&P 500 and EURUSD. We use a 3-volatility model which accounts for non-normal skewness and kurtosis. The volatilities control the level, slope and curvature of the Black-Scholes implied volatility smile; accordingly we term them "base", "skew" and "smile" volatility. We define instantaneous skewness and kurtosis as simple ratios of the volatilities, and show that when these metrics are held constant, the model is relative sticky-delta.

For the S&P 500 in 2008, skew and smile volatility are highly correlated with base volatility. Instantaneous skewness and kurtosis are remarkably stable, including over the market dislocation of the last four months of the year. Daily changes in all three volatilities are correlated with daily returns. For EURUSD in 2006, base and smile volatility are closely correlated, but in contrast to the equity case, skew volatility is independent and changes sign. This change in sign appears to provide advance warning of the two major market moves of the year. However, daily changes in the volatilities are uncorrelated with daily returns.

1. THE MODEL

1.1 Basic definitions and results

Here we recall some basic definitions and results from Carey (2005, 2006). Let (X_t) be a positive-valued adapted stochastic process on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), Q), t \geq 0$, satisfying the usual conditions. Here (X_t) is interpreted as the market price of a financial security, (\mathcal{F}_t) as the market information structure, and Q as an equivalent martingale measure. Let $\delta t > 0$ denote a finite period of time, and define $\delta X_t = X_{t+\delta t} - X_t$, so that the security return over the interval t to $t + \delta t$ reads $\delta X_t / X_t$. $E_t \cdot$ denotes expectation conditional on \mathcal{F}_t , and j is a generic positive integer.

We define j -th order finite-period volatility $\Sigma_{j,t}$ via the identity:

$$\Sigma_{j,t}^j \equiv E_t(\delta X_t / X_t)^j / \delta t ,$$

with the convention that $\Sigma_{j,t}$ equals the non-negative root when j is even, and we let $\sigma_{j,t} = \lim_{\delta t \searrow 0} \Sigma_{j,t}$ denote j -th order instantaneous volatility. In Carey (2005) a simple thought experiment shows how to interpret $\Sigma_{j,t}$ and its limit $\sigma_{j,t}$ as standard changes.

Mathematically, the quantities $(\sigma_{j,t} X_t)^j$, $j=1,2,\dots$ can be characterised as *infinitesimal moments*. They are defined in Gikhman and Skorokhod (1972) in the context of Markov processes, and have begun to appear in the finance literature — see e.g. Johannes (2004).

We assume that the security has *market-independent returns* under the equivalent martingale measure Q . We define this to mean that for any t and δt , the security return $\delta X_t / X_t$ is independent of \mathcal{F}_t , the information available to investors at time t . If the information available to investors is limited to that represented by the security price history, or if all extraneous information is irrelevant to investors, this reduces to the more common assumption of independent (relative) increments.¹ It also implies that the instantaneous volatility processes $\sigma_{j,t}, j=1,2,\dots$ are deterministic. We make the further assumption that over arbitrarily small intervals of time, the security almost surely does not double in price.

Under these premises, the characteristic function of the log return $\ln(X_{t+\tau}/X_t)$ reads:

$$\phi(v) = \exp \sum_{j=1}^{\infty} \binom{iv}{j} \int_t^{t+\tau} \sigma_{j,u}^j du ,$$

with $\sigma_{1,t} = b_t$, the instantaneous cost of carry of the security. See Appendix 1 for a derivation. This formula and a related one were originally obtained in Carey (2005, 2006) under the assumption that the finite-period volatilities are deterministic for very small time horizons δt . While not as clean as the assumption we make above, it perhaps carries more intuition. When higher-order volatilities are zero, this characteristic function reduces to that of the Black and Scholes (1973) and Merton (1973) model.

The market price C of a European-exercise, path-independent option satisfies the partial differential equation (PDE):

$$\partial_t C(X,t) + \sum_{j=1}^{\infty} \frac{1}{j!} \sigma_{j,t}^j X^j \partial_X^j C(X,t) = r_t C(X,t) ,$$

where r_t is the riskless rate of interest at time t . This can be viewed as an infinite-order version of the (second-order) Black-Scholes PDE. As indicated, the sequence of volatilities σ_j , $j=1,2,\dots$ can be truncated at $j=2$ to yield the Black-Scholes model. It cannot however be truncated at any higher order, and must remain infinite.²

¹ Often the market information structure is exogenously limited to the security price history, albeit implicitly (and perhaps, sometimes, unwittingly). There is however no good basis for this assumption, as pointed out by e.g. Babbs and Selby (1998) and Brody et al (2007).

² Truncation by way of approximation can be highly unstable, no matter how high the truncation point.

We define the average volatility σ_j via:

$$\sigma_j^j \equiv \frac{1}{\tau} \int_t^{t+\tau} \sigma_{j,u}^j du ,$$

again with the convention that σ_j be non-negative when j is even. The (Fisher) skewness and kurtosis of a random variable are given by:

$$\gamma_1 = \frac{c_3}{c_2^{3/2}} , \quad \gamma_2 = \frac{c_4}{c_2^2}$$

where c_j is the j -th central cumulant — see e.g. Abramowitz and Stegun (1972), p.928. It can be shown that when all infinitesimal moments of a stochastic process exist, so do the infinitesimal cumulants, and further that the two sequences are identically equal (see Appendix 2). It follows that, to first order in δt :

$$c_j = \sigma_j^j \delta t , \quad j = 1, 2, \dots$$

and the skewness and kurtosis corresponding to this approximation read:

$$\gamma_1 = \left(\frac{\sigma_{3,t}}{\sigma_{2,t}} \right)^3 \frac{1}{\sqrt{\delta t}} , \quad \gamma_2 = \left(\frac{\sigma_{4,t}}{\sigma_{2,t}} \right)^4 \frac{1}{\delta t} . \quad (1)$$

These quantities are not standardised for time, and become infinite as δt tends to zero. This is not a reflection of any misbehaving on the part of the stochastic process, but rather that these coefficients were not designed for this setting. We therefore suggest defining *instantaneous skewness* χ_t and *instantaneous kurtosis* κ_t as:

$$\chi_t = \frac{\sigma_{3,t}}{\sigma_{2,t}} , \quad \kappa_t = \frac{\sigma_{4,t}}{\sigma_{2,t}}$$

where the exponents in (1), serving no purpose, have been dropped. Note that these coefficients are defined for any stochastic process for which they exist, and not simply for a process with \mathcal{F}_t -independent increments. We further define *average instantaneous skewness* and *kurtosis* as:

$$\chi = \frac{\sigma_3}{\sigma_2} , \quad \kappa = \frac{\sigma_4}{\sigma_2} .$$

This is something of a misnomer, because $\chi \neq \int_t^{t+\tau} \chi_u du$ and $\kappa \neq \int_t^{t+\tau} \kappa_u du$ in general. However, it is always the case that $\chi_t = \lim_{\tau \rightarrow t} \chi$ and $\kappa_t = \lim_{\tau \rightarrow t} \kappa$. The quantities χ, κ will provide the basis for an alternative parameterisation of the model presented in the subsection 1.2.

1.2 The reduced Merton jump-diffusion implementation

Here we make use of the Merton (1976) jump-diffusion model to obtain a pricing formula that depends on the volatilities of order 2, 3 and 4. The full Merton model with lognormal jumps has four free risk parameters: the volatility σ , the jump rate λ , the jump size mean γ and its standard deviation δ . We only require three free parameters, and so reduce the parameter set by exogenously setting $\lambda T = 1$. This choice is motivated both by the analytical simplification that results, and by the empirical behaviour of the marked-implied λ which, in the equity index setting at least, tends to decrease as time to expiration increases (see Carey (2005)).

The option pricing formula then simplifies to:

$$\sum_{n=0}^{\infty} \frac{BSM_n}{n!}$$

where BSM_n is the Black-Scholes-Merton formula with cost of carry and volatility given by:

$$b_n = b + \frac{n\gamma - m_1}{T}, \quad v_n = \sqrt{\sigma^2 + \frac{n\delta^2}{T}}$$

respectively. As noted in Carey (2005), we do not assume that the risk-neutral process is a Merton lognormal jump-diffusion (with or without $\lambda T = 1$); the formula is simply borrowed for reasons of analytical tractability.

A general formula giving j -th order volatility in terms of the Merton jump-diffusion parameters can be found in Carey (2005). With this formula, j -th order volatility is readily computed from the jump-diffusion parameters. However, because the formula is highly nonlinear, obtaining the jump-diffusion parameters from a set of volatilities is more involved. Under the simplification set out above, the parameter equations reduce to:

$$\begin{aligned} \sigma_2^2 T &= \sigma^2 T + A^2 B - 2A + 1, \\ \sigma_3^3 T &= A^3 B^3 - 3A^2 B + 3A - 1, \\ \sigma_4^4 T &= A^4 B^6 - 4A^3 B^3 + 6A^2 B - 4A + 1 \end{aligned}$$


subject to the constraints $A > 0, B \geq 1$, where $A = e^\gamma, B = e^{\delta^2}$. The last two equations can be solved analytically for A, B using symbolic mathematics software (we use MuPad, packaged with Matlab). The first equation then yields σ .

In this way, we obtain an option pricing formula that depends on the three risk parameters $\sigma_2 - \sigma_3 - \sigma_4$. However, this formula can be somewhat unstable, because there are many combinations of these three parameters that are not admissible. A combination of parameters that is admissible can abruptly become inadmissible under a small change in one of the parameters. In the next section we will see evidence that the ratios $\chi = \sigma_3 / \sigma_2$ and $\kappa = \sigma_4 / \sigma_2$, defined above as average skewness and kurtosis, are substantially more stable over time than their respective numerators. It will be useful then to consider the alternative parameter interface $\sigma_2 - \chi - \kappa$, in which σ_2 can be varied over a wide range without leaving the domain of the pricing formula.

Below is an example of a volatility surface with constant volatilities. The volatilities chosen here correspond to typical values found in S&P 500 options market in the course of the year 2008, at roughly the 3-month expiration mark. The input volatilities are $\sigma_2 = 22\%, \sigma_3 = -17\%, \sigma_4 = 18\%$. The corresponding jump-diffusion

parameters are $\sigma = 8.68\%$, $(\lambda = 4)$, $\gamma = -5.91\%$, $\delta = 8.81\%$. The spot underlying is 100, and both the cost of carry and the rate of interest are zero.

Figure 1
Example volatility surface




The input volatilities are $\sigma_2 = 22\%$, $\sigma_3 = -17\%$, $\sigma_4 = 18\%$. The spot and forward prices both equal 100 units.

1.3 Volatility smile dynamics and price sensitivities

First we examine the behaviour of the volatility smile under changes in the spot price of the underlying security. In figure 2, the spot price is increased and decreased by a small amount. Evidently, the volatility smile shifts right and left respectively, that is, the smile follows the spot price. This behaviour is known as sticky moneyness (see e.g. Daglish et al (2007)), and is a hallmark of jump-diffusion/Lévy processes.

Next we examine the effect of changes in the volatilities on the Black-Scholes implied volatility smile and the option prices it represents. Beginning with the graph for third-order volatility in figure 3 (labelled HOV[3]) we can observe an increase (decrease) in smile asymmetry as σ_3 increases (decreases) in magnitude; thus σ_3 predominantly affects the skew (or smirk) aspect of the volatility smile. This is confirmed by the graph of option price sensitivities to σ_3 in figure 4. An increase in magnitude increases the price of below-the-money option and lowers the price of above-the-money options (note that in our example, σ_3 is negative). We shall refer to σ_3 as *skew volatility*.

Figure 2
The Black-Scholes implied volatility smile
under changes in spot price

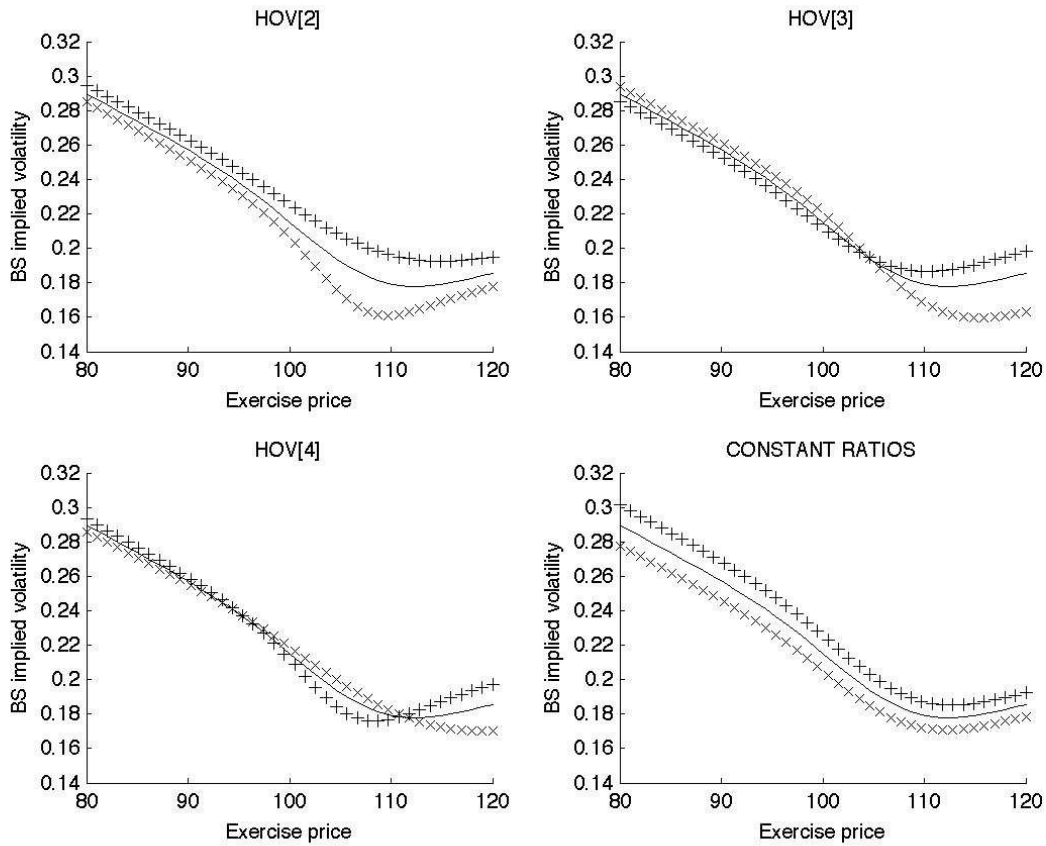


The solid line is the initial volatility smile. The pluses (crosses) are the volatility smile resulting from an increase (decrease) of 2 units in the underlying security price. The initial spot price is 100 units. Volatilities remain constant.

With fourth-order volatility (labelled HOV[4]), we notice from figure 3 that the implied volatility smile curvature increases (decreases) as σ_4 increases (decreases); thus σ_4 predominantly affects the smile shape in implied volatilities. This is confirmed by the graph of option price sensitivities to σ_4 in figure 4. An increase in magnitude mildly increases the price of away-from-the-money options, but sharply lowers the prices of at-the-money options. We shall term σ_4 *smile volatility*.

Turning now to second-volatility, we see in figure 3 (labelled HOV[2]) that its influence on the implied volatility smile is more clear-cut: unlike with skew or smile volatility, an increase (decrease) increases (decreases) implied volatilities across all exercise prices. This effect is most pronounced in a region some way to the right of the spot price. Indeed the option price sensitivity to σ_2 , plotted in figure 4, is positive everywhere and has a bell shape similar to that of the Black-Scholes vega.

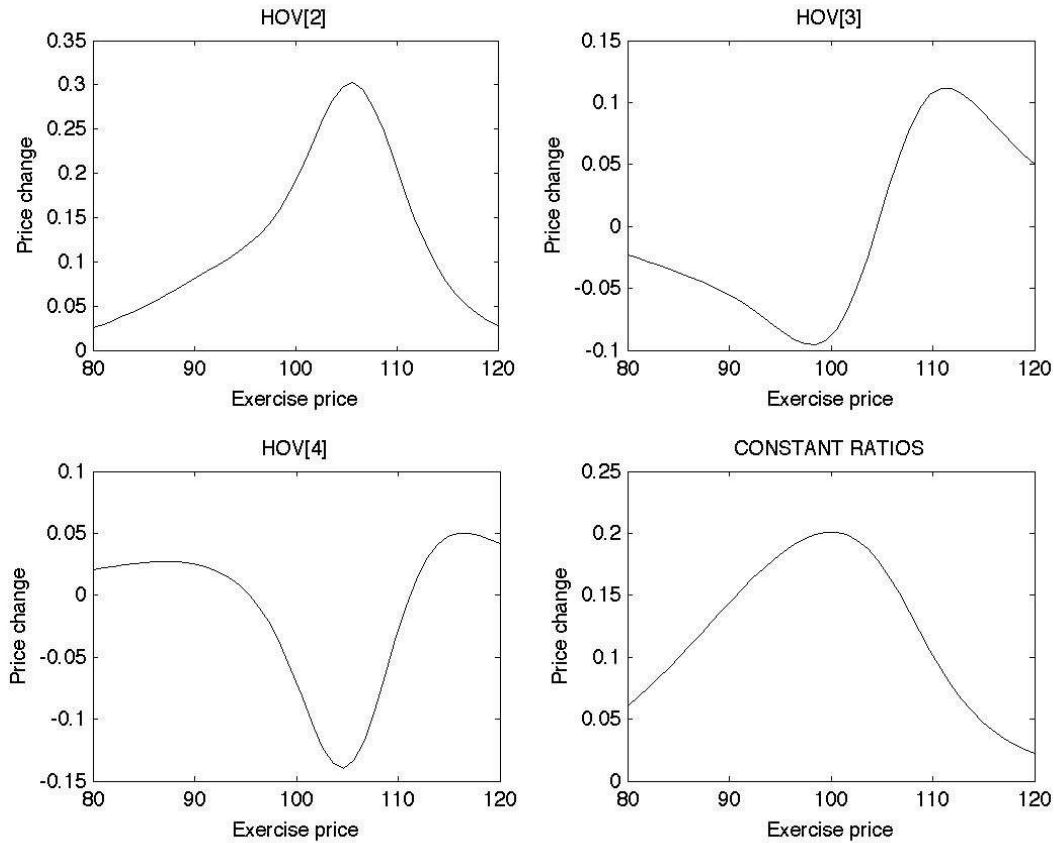
Figure 3
 Black-Scholes implied volatility smile
 sensitivities to higher-order volatilities



The solid line is the initial volatility smile. The pluses (crosses) are the volatility smile resulting from an increase (decrease) of 0.01 in higher-order volatility. The spot underlying is at 100.


The above analysis holds under the $\sigma_2 - \sigma_3 - \sigma_4$ parameterisation, and each volatility was varied in turn while holding the others constant. Under the alternative $\sigma_2 - \chi - \kappa$ parameterisation, introduced in subsection 1.2, the picture for the last two parameters of the interface is very similar. Indeed, varying the skewness metric $\chi = \sigma_3 / \sigma_2$ by itself is the same as varying skew volatility σ_3 independently of the other two volatilities, and likewise varying the kurtosis $\kappa = \sigma_4 / \sigma_2$ by itself is the same as varying smile volatility σ_4 , also independently.

Figure 4
Call or put option price sensitivities to higher-order volatilities



However, varying second-order volatility σ_2 in the $\sigma_2 - \chi - \kappa$ interface has a different result than in the $\sigma_2 - \sigma_3 - \sigma_4$ interface. The plot in figure 3 labelled "CONSTANT RATIOS" shows that an increase (decrease) in σ_2 , while holding skewness χ and kurtosis κ constant, results in a parallel shift upwards (downwards) of the entire implied volatility curve. The corresponding option price sensitivities are shown in figure 4, with a positive-valued bell-shaped curve that is somewhat wider and flatter than for the $\sigma_2 - \sigma_3 - \sigma_4$ parameterisation.

Figure 5
 Black-Scholes implied volatility smile sensitivity
 to base volatility in the $\sigma_2 - \chi - \kappa$ parameterisation



The solid line is the initial volatility smile. The spot underlying is at 100.

In fact, this pattern extends over a surprisingly wide interval around the original σ_2 . Figure 5 shows the result of increasing (decreasing) σ_2 at intervals of 5%. Again we see almost perfectly parallel shifts in implied volatilities. Together with the sticky-moneyness nature of the model, in which the implied volatility curve shifts left or right with the spot price, the sensitivity with respect to σ_2 within the $\sigma_2 - \chi - \kappa$ parameterisation is strongly consistent with the application of the popular trader's rule known as "relative sticky moneyness," in which the volatility curve shifts both sideways with the spot price, and up and down with ATM volatility — see Daglish et al (2007). We shall refer to σ_2 as *base volatility*.

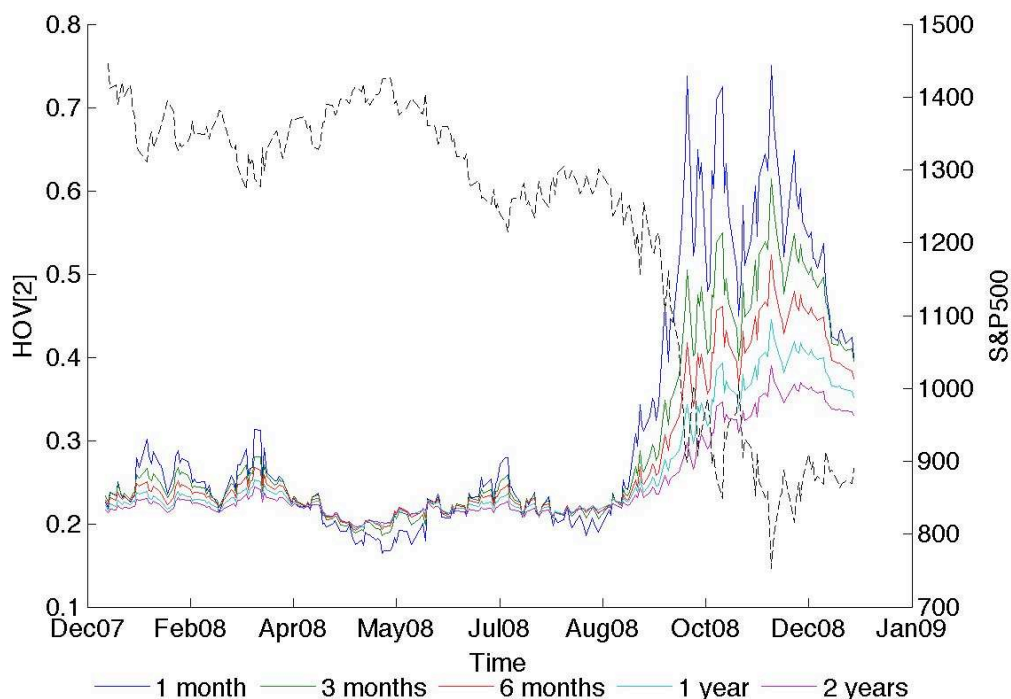
2. S&P 500 2008 TIME SERIES

The data consists of 101402 pairs of price quotes for call and put options price. The data was culled of both call and put quotes showing a bid or ask equal to zero, as well as a mid-quotes less than 0.5 points. Applicable rates of interest and index forward prices were implied from option prices via linear regression. A separate model fitting was carried out for each subset of quotes sharing the same trade and expiration dates, for a total of 2277 separate calibrations. Each calibration consisted of the minimisation of the sum of squared differences between model and market prices, weighted by the inverse of their bid-ask spreads. The jump-diffusion parameters were fitted in preference to the volatilities, due to the domain-related issues discussed in section 1.2. The fit to market prices can be measured via the absolute deviation between the market implied volatility (averaged over the call and the put) and the model implied volatility. We have a mean and interquartile range for this metric of respectively 0.64 and 0.61 percentage points.

The volatilities of order 2, 3 and 4 are plotted in figures 6(a)-(c). These time series are for the five fixed expirations of 1-3-6 months and 1-2 years. The fixed expiration data was obtained by cubic spline interpolation along the expiration axis. The closing S&P 500 index is plotted for context. The severe market turmoil of September through December stands in stark contrast to the first part of the year. For the sake of expediency we will refer to the last four months of 2008 as "the financial crisis."

We make a total of eight observations. The first four mirror the points made in Carey (2005) on the basis of snapshot parameter estimates, also for the S&P 500.

Figure 6(a)
Implied base volatility for five fixed expirations
S&P 500



The dashed line is the S&P 500 index.

Figure 6(b)
 Implied skew volatility for five fixed expirations
 S&P 500


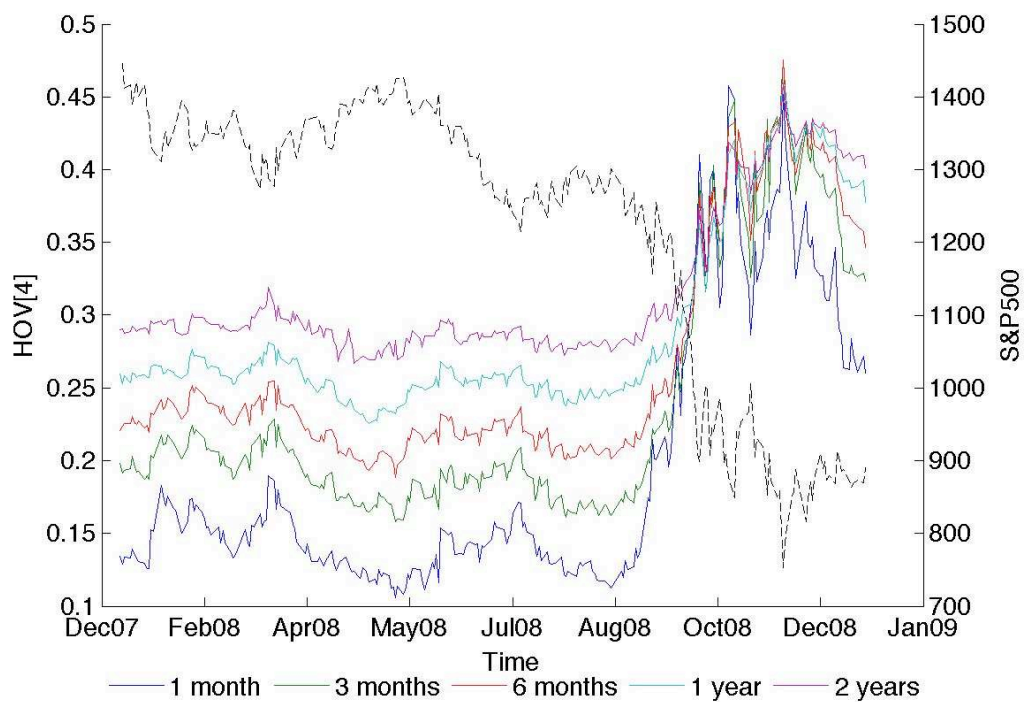


Figure 6(c)
 Implied smile volatility for five fixed expirations
 S&P 500




The dashed lines are the S&P 500 index.

i. **Base volatility closely tracks at-the-money volatility.** The latter is plotted in figure 7(a), and is visually indistinguishable from the base volatility time series. To get a clearer picture, consider the ratio of base volatility to at-the-money (ATM) volatility, plotted in figure 7(b). Base volatility is slightly greater for the 1-month expiration, fluctuating around 1.02 times at-the-money volatility; but the situation reverses at about the 3-month mark, whereupon the ratio steadily decreases to about 0.95 at the 2-year mark, with a further dip to below 0.9 during the financial crisis. These conclusions are confirmed by the base and ATM volatility averages, plotted in figure 8.

In addition, the standard deviations of both quantities, also plotted in figure 8, are also broadly similar. The standard deviation of ATM volatility is marginally greater than that of base volatility, a difference that is more pronounced at longer expirations as well as during the financial crisis.

Figure 7(a)
At-the-money Black-Scholes implied volatility for five fixed expirations
S&P 500




The dashed line is the S&P 500 index.

ii. **Skew and smile volatility have the same order of magnitude and standard deviation as base volatility.** This is revealed by a cursory examination of figures 6(a)-(c), in which base volatility ranges from 10% to 80% across all expirations, while skew and smile volatility both range (in magnitude) from 10% to 50%. The mean and standard deviations, plotted in figure 8, confirm this. The three volatilities are closest in magnitude between 6 and 12 months to expiration. The proximity of skew and smile volatility standard deviations is especially striking: they are essentially equal from the 3-month expiration mark onward; and for the period prior to the financial crisis, are joined by the base volatility standard deviation from the 6-month mark onward.

Figure 7(b)

Base volatility to at-the-money volatility ratio for five fixed expirations
S&P 500



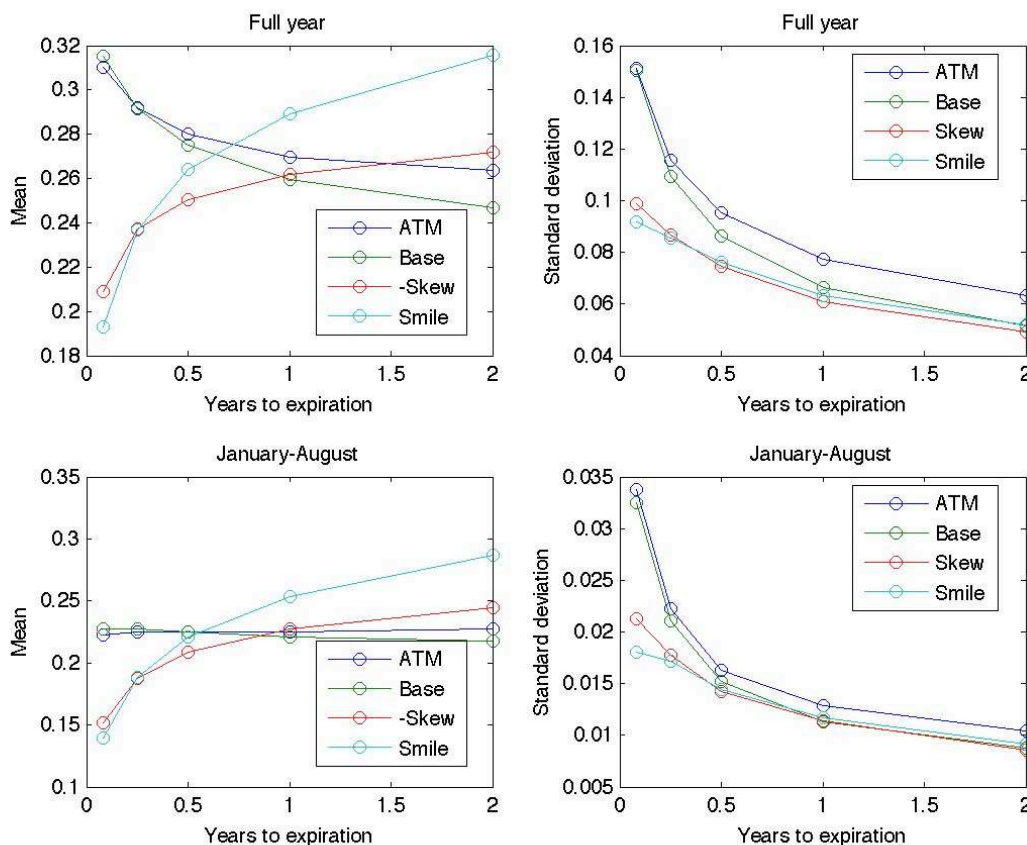
iii. **Skew volatility is consistently negative.** This is evident from the time series in figure 6(b). A negative skew volatility implies negative instantaneous skewness. This feature is not surprising, given that negative skewness in the risk-neutral distribution of returns is a well-documented fixture of major equity index markets.

iv. **The term structure of skew and smile volatilities is very pronounced.** Considering the first eight months of 2008, this is seen in the time-series in figures 6(b)-(c), as well as in the averages plotted in figure 8. In contrast, the term structure for the base volatilities is essentially flat. Time-inhomogeneity of the parameters is a common feature of option-pricing models featuring a process with independent increments: specifically, the implied skewness and kurtosis increase with expiration. Often this is viewed as evidence that such processes are misspecified, on the grounds that the skewness and kurtosis cannot be expected to increase all the time. However, the presence of parameter bias in a model may or may not be an issue, depending on the specific usage. The Black-Scholes model for instance is still widely used in many markets for trading first-generation derivatives, including the vanilla options considered here.

Interestingly, during the financial crisis the situation is reversed: the term structure of base volatility becomes sharply sloped (figure 6(a)), dropping off as time to expiration increases, while the term structure of skew and smile volatility flattens out, at some points completely (figures 6(b) and (c)). As the crisis abates, towards the end of the year, the term structures begin to return to their original state.

Figure 8

Mean and standard deviation for at-the-money, base, skew and smile volatilities S&P 500

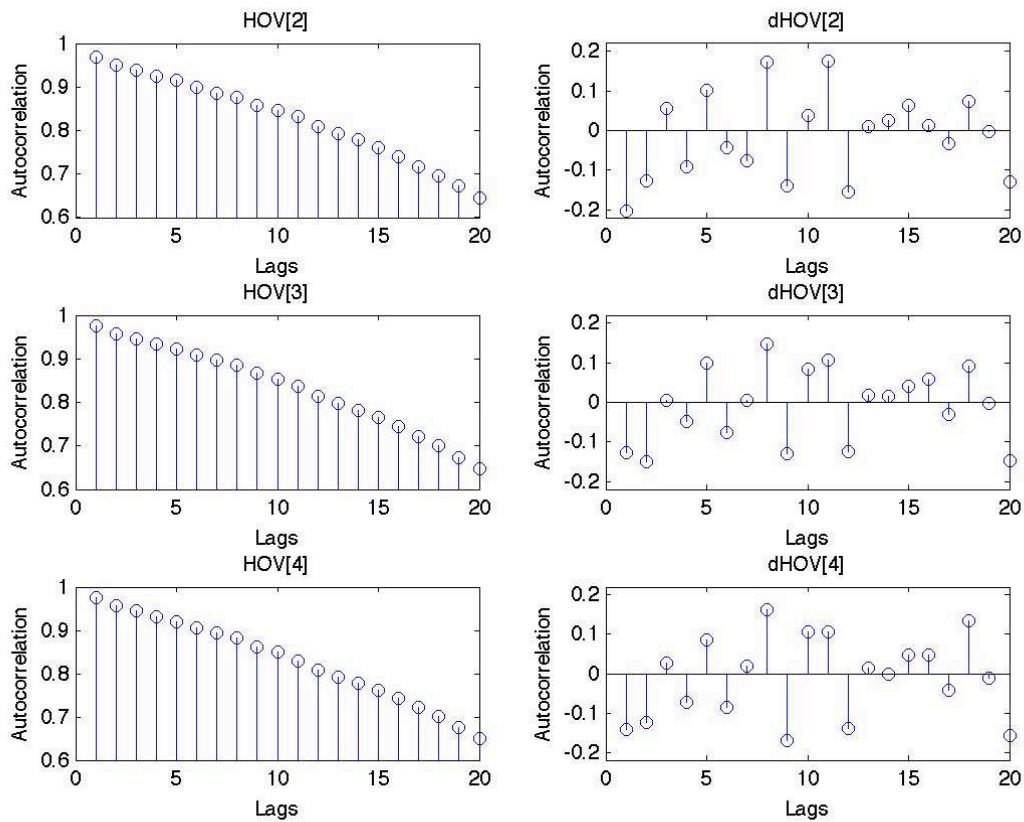


The bottom figures do not cover the market turmoil of the last four months of 2008, and are thus more representative of normal market conditions.

v. Daily changes in the volatilities show mild negative serial correlation.

Figure 9 plots the autocorrelations for the volatility time series, as well as for the daily changes (i.e. the first differences). The correlograms suggest that the volatilities are not stationary (the correlations show a slow, linear decrease), but that they are first-difference stationary. For all three volatilities there is evidence of modest negative autocorrelation in daily changes over the first two lags, indicating a mild propensity for self-correction over the two days following a change. The correlograms do not change markedly when the financial crisis is left out.

Figure 9
Correlograms for base, skew and smile volatilities
S&P 500

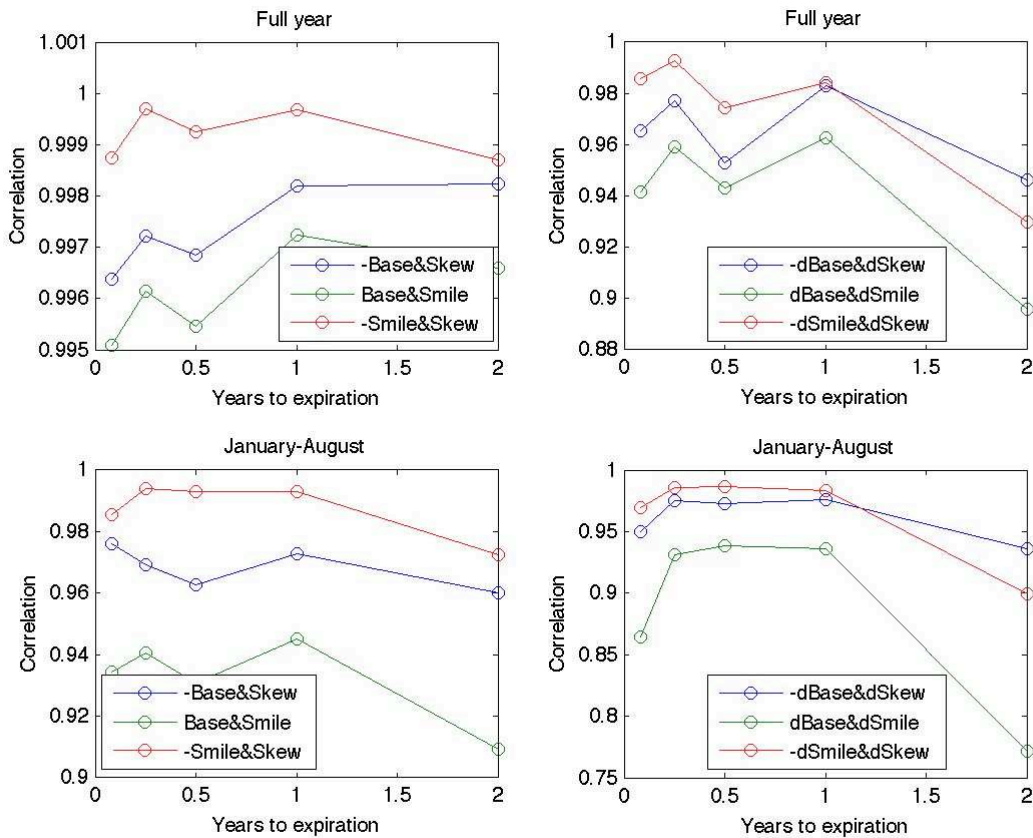


The left column are the ACFs for the volatility time series with fixed expiration 30 days. The right column are the ACFs for the first differences of the same series.

vi. **The volatilities are highly cross-correlated.** This is true for the absolute quantities as well as, to a slightly lesser extent, their daily changes. The correlations are plotted in figure 10. For the period prior to the financial crisis, the skew-to-base correlation ranges from 0.96 to 0.98 depending on time to expiration, while the smile-to-base correlation ranges from 0.91 to 0.94. The highest correlation however is between skew and smile volatilities, ranging from 0.97 to 0.99. Cross-correlations for the daily changes follow the same pattern, but with somewhat lower coefficients, and with a marked drop-off at longer expirations: the correlation coefficients are greatest for 3- to 12-month expirations, ranging from 0.93 to 0.98.

When the financial crisis is included, the general pattern is preserved, but the correlation coefficients shift upwards. The absolute quantities, in particular, range above 0.995, indicating almost perfect correlation between the volatilities during the financial crisis.

Figure 10
Higher-order volatility cross-correlations
S&P 500



The left column are the correlations for the base, skew and smile volatilities, as functions of time to expiration. The right column are the correlations for daily changes in base, skew and smile volatilities, as functions of time to expiration. The bottom figures do not cover the financial crisis of the last four month of 2008, and are more representative of normal market conditions.

vii. **Instantaneous skewness and kurtosis are surprisingly stable.** This point is essentially a corollary to point (vi). Figure 11(a) plots the skew-to-base ratio, which we defined as average instantaneous skewness. This ratio fluctuate within well-defined ranges, which show a strong expiration dependency. Skew volatility ranges in magnitude from 60-70% of base volatility at the 1-month expiration mark, to 105-115% of base volatility at the 2-year mark. The financial crisis appears to result in a slight increase in the standard deviation of the ratio, but overall does not appear to push it in one direction or the other.

Figure 11(b) plots the smile-to-base ratio, which we defined as average instantaneous kurtosis, and shows a similar picture. Smile volatility ranges in magnitude from 55-65% of base volatility at the 1-month expiration mark, to 120-140% of base volatility at the 2-year mark. Again the financial crisis seems to slightly increase the standard deviation of the series, but less so than for skewness; also, the ratio is not influenced either way.

Figure 11(a)
 Instantaneous skewness for five fixed expirations
 S&P 500

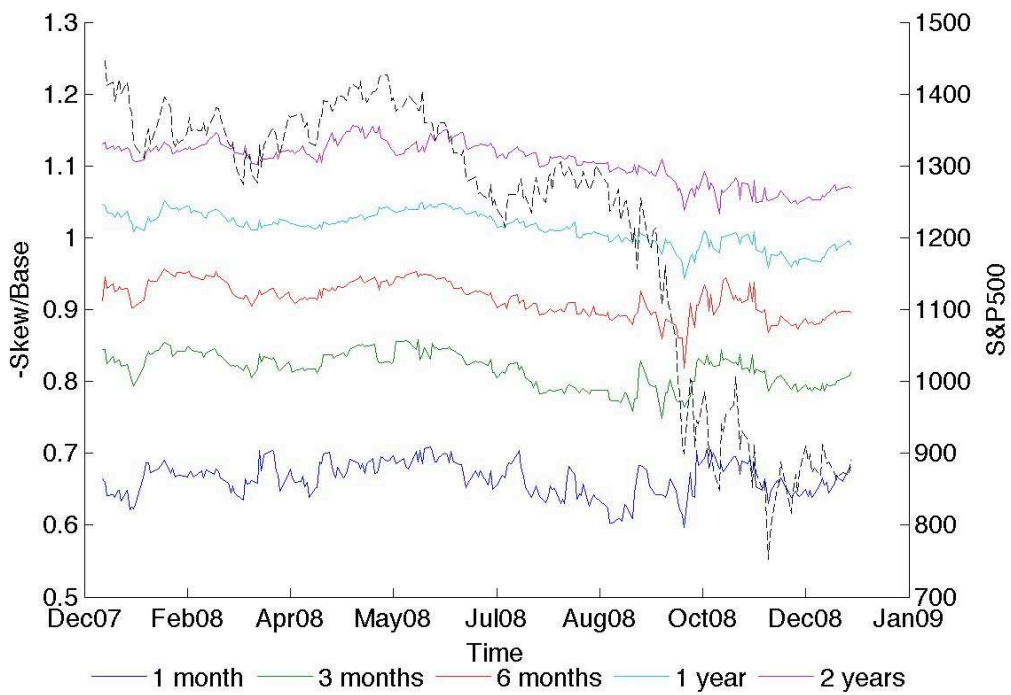
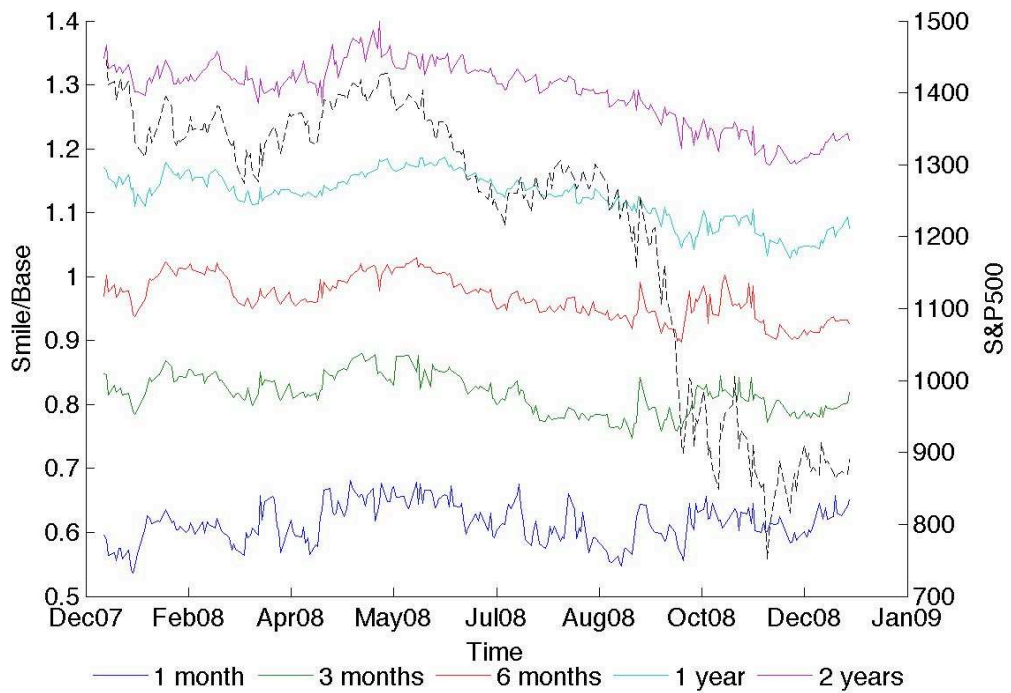


Figure 11(b)
 Instantaneous kurtosis for five fixed expirations
 S&P 500



The dashed lines are the S&P 500 index.

viii. **Changes in the volatilities are correlated with daily returns.** It is a stylised fact of equity index markets that changes in at-the-money volatility are negatively correlated with index returns: positive (negative) returns tend to be associated with decreases (increases) in implied volatility. Table 1 reports the correlation coefficients between daily returns and daily changes in base, skew and smile volatilities, for five fixed expirations.

Given the similarity between ATM volatility and base volatility, it is not surprising that the latter should display a high negative correlation with daily returns, ranging from -0.91 to -0.94. This association is confirmed by figure 12(a), a scatter plot of daily returns³ versus daily changes in 1-month base volatility. Based on the best-fit line, an index return of +/- 1% is associated with an (absolute) change in base volatility of -/+ 1.23%. The R² for this regression is 0.83.

Changes in skew and smile volatilities are also correlated with returns. This was expected since, as noted above, daily changes in the volatilities are cross-correlated. The correlation between index returns and changes in skew volatility ranges from +0.83 to +0.90 (table 1), and the best fit line in figure 12(b) suggests an (absolute) change in skew volatility of +/- 0.66% for every +/- 1% index return, with an R² of 0.73. The correlation between index returns and changes in smile volatility ranges from -0.77 to -0.88 (table 1), and the best fit line in figure 12(c) suggests an (absolute) change in smile volatility of -/+ 0.60% for every +/- 1% index return, with an R² of 0.69.


Table 1
Correlation between daily returns and daily changes in the volatilities,
for five fixed expirations.
S&P 500

Time to expiration	Correlation between daily returns and daily changes in ...		
	Base volatility	Skew volatility	Smile volatility
1 month	-0.9130	0.8553	-0.8333
3 months	-0.9325	0.8877	-0.8574
6 months	-0.9316	0.8707	-0.8471
1 year	-0.9351	0.9028	-0.8849
2 years	-0.9141	0.8297	-0.7748

³ The returns we use here are synthetic, and are computed based on the 1-month option-implied forward index. This was necessary because the options market closes later than the equities market, so that the index closing price for the day is stale relative to the closing option prices used here.

Figure 12(a)


Daily changes in base volatility versus daily returns, S&P 500



The best fit line has slope=-1.23, intercept=0.00, and $R^2=0.83$.
Note: the returns used here are synthetic—see footnote 1.

Figure 12(b)


Daily changes in skew volatility versus daily returns, S&P 500



The best fit line has slope=0.66, intercept=0.00, and $R^2=0.73$. Note:
the returns used here are synthetic—see footnote 1.

Figure 12(c)

Daily changes in smile volatility versus daily returns, S&P 500



*The best fit line has slope=-0.60, intercept=0.00, and $R^2=0.69$.
Note: the returns used here are synthetic—see footnote 1.*

3. EURUSD 2006 TIME SERIES

The data, obtained from the British Bankers Association (BBA), consists of a set of quotes for each of 253 business days in 2006. Each set comprises an at-the-money, a 25-delta risk-reversal, and a 25-delta strangle volatility, for fixed expirations of 1, 3 and 12 months. These volatility quotes were converted to option prices as detailed in Appendix 3. We used USD and EUR Libor as interest rates. As with the S&P 500 data, the model was calibrated to option prices using squared pricing errors. However, since the number of free parameters in our formula (base, skew and smile volatility) matches the number of inputs (the three volatility quotes), the fit to market prices is exact.

The higher-order volatilities for base, skew and smile volatilities are plotted in figures 13(a)-(c). The contemporaneous EURUSD rate is plotted for context. We make a total of eight observations, in parallel with section 2.

Figure 13(a)
 Implied base volatility for three fixed expirations
 EURUSD


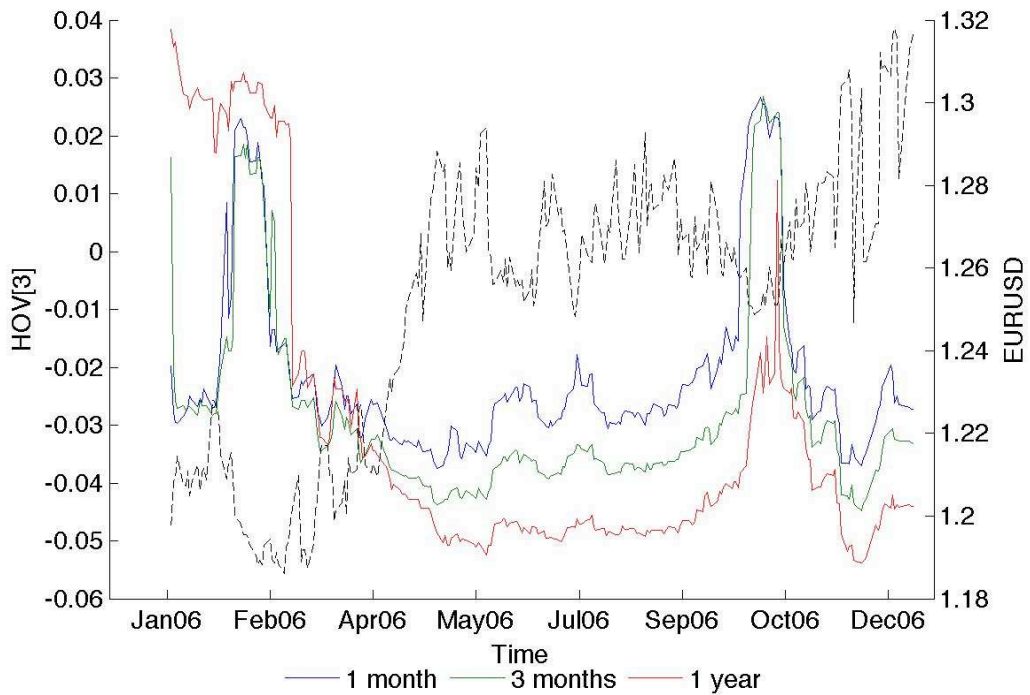



Figure 13(b)
 Implied skew volatility for three fixed expirations
 EURUSD



The dashed lines are the EURUSD rate.

Figure 13(c)
Implied smile volatility for three fixed expirations
EURUSD



The dashed line is the EURUSD rate.


i. **Base volatility closely tracks ATM volatility.** As with the S&P 500, base volatility is visually indistinguishable from ATM volatility, plotted in figure 14(a). The ratio of base volatility to ATM volatility, plotted in figure 14(b), fluctuates between 1.025 and 1.04 for much of 2006, before increasing steeply to the 1.05-1.06 range at the end of the year. The averages for base and at-the-money volatilities, plotted in figure 15, bear this out. As for their standard deviations, they can be observed in figure 15 to be essentially identical.

ii. **Skew and smile volatilities have the same order of magnitude and standard deviation as base volatility.** Smile volatility ranges from a low of 3%, reached at the 1-month expiration, to just under 10%, a high reached at the 1-year expiration, as seen in figure 13(c). This compares to the 6%-10% range achieved by base volatility in figure 13(a). In figure 13(b), skew volatility ranges from -1% to -5% when negative (i.e. most of the time) and from 1% to 2% when positive. Interestingly, when skew volatility changes sign, it does so abruptly and without pausing in the [-1%, +1%] interval.

Figure 15 plots the means and standard deviations. Skew and smile volatility are seen to also be comparable to base volatility in standard deviation. Smile volatility has about half the standard deviation of base volatility, while skew volatility is greater and rises to three times as much at the 12-month expiration.

Figure 14(a)

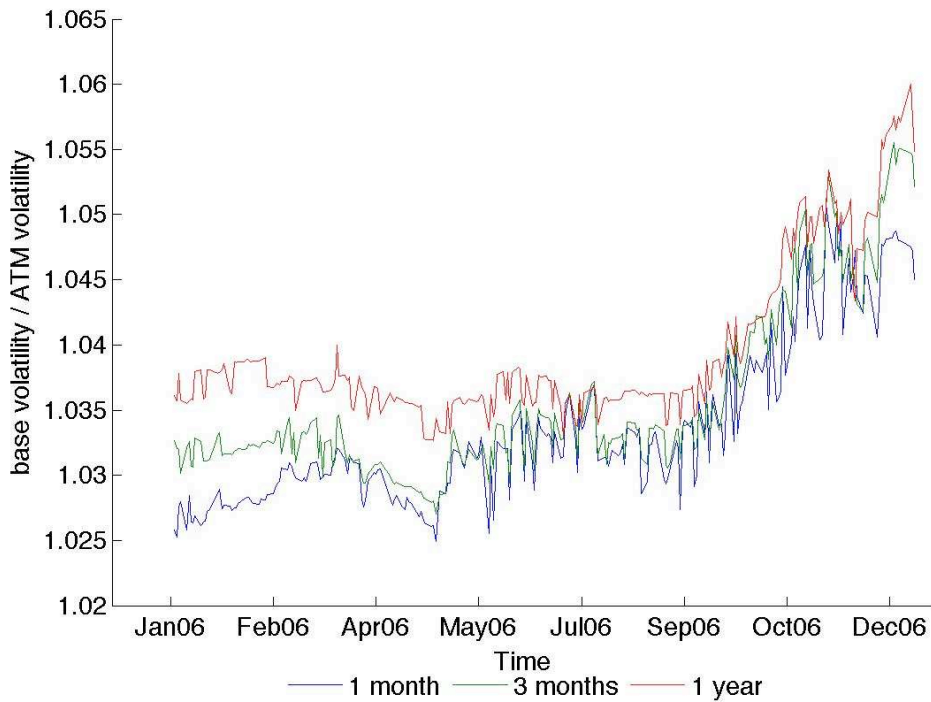
At-the-money Black-Scholes implied volatility for three fixed expirations
EURUSD



The dashed line is the EURUSD rate.

Figure 14(b)

Base volatility to at-the-money volatility ratio for three fixed expirations
EURUSD



iii. **Skew volatility alternates between positive and negative.** While skew volatility is negative during most of 2006, indicating negative skewness in EURUSD, two short episodes in which it turns positive can be observed in figure 13(b). These episodes appear to last about a month each, and seem to signal the two major upward moves in EURUSD that occur during the year: the first from 1.20 to 1.28 (March-May), and the second from 1.26 to 1.30 (October-December). It may be that they correspond to the building up of speculative positions in anticipation of the price moves to follow. We observe that Bates (1991) also notes a marked change in skewness prior to the 1987 stockmarket crash.

iv. **The term structure of skew and smile volatilities is very pronounced.** For smile volatility, the term structure is very steep, as can be seen both in the time-series in figures 13(b)-(c), and in the averages plotted in figure 15. From a mean value of 4% at the 1-month expiration, it rises to almost 9% at the 1-year mark. In contrast, base volatility has a more modest upward slope, starting at 8% and ending at (coincidentally) 9%.

Skew volatility also has a term structure that increases with time to expiration, albeit to a lesser extent. Because of the intermittent changes of sign, this is not so apparent in figure 15, which plots the mean. In the time-series chart in figure 13(b) however, we see that during periods when EURUSD is range-bound (e.g. May-September), 1-month skew volatility is about 1% less (in magnitude) than that at the 3-month mark, which in turn is 1% less than 1-year skew volatility.

Figure 15

Mean and standard deviation for at-the-money, base, skew and smile volatilities EURUSD

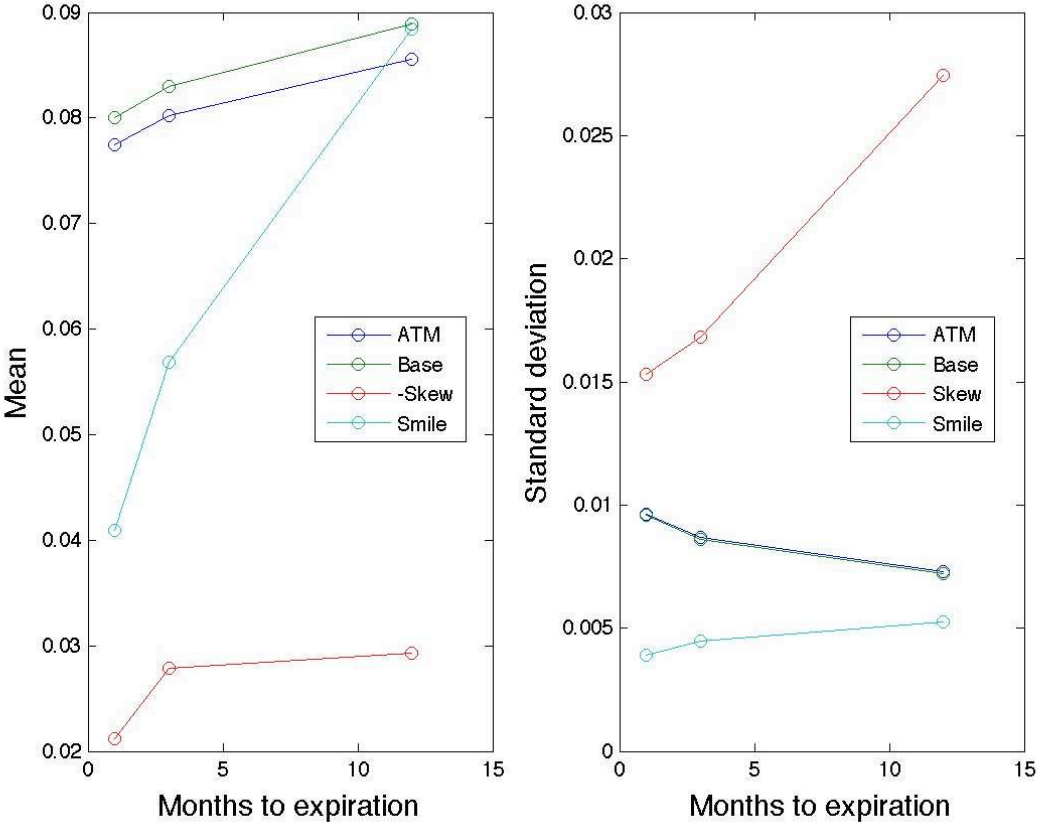
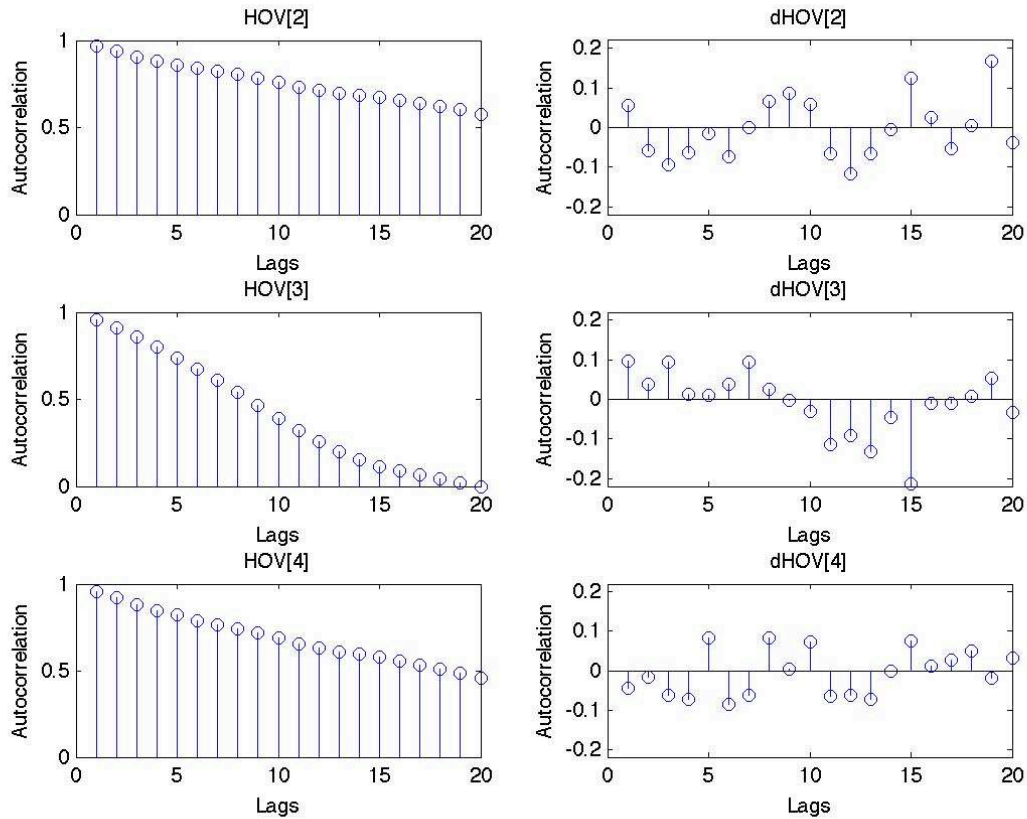


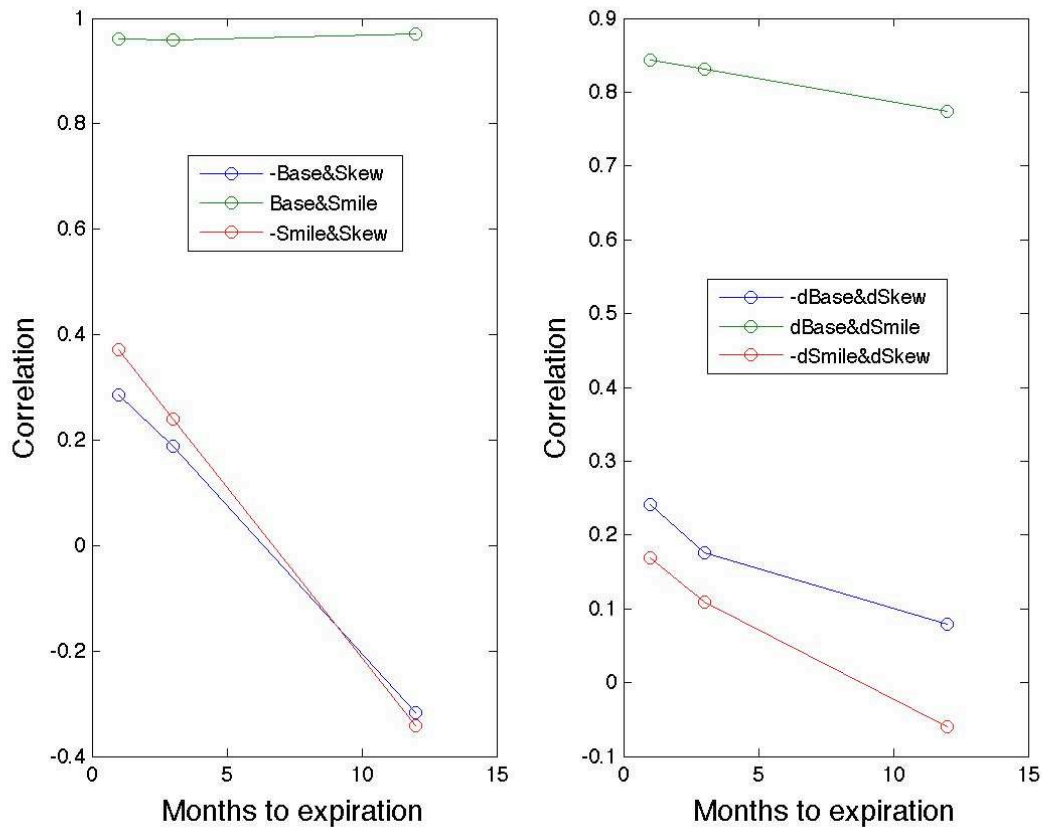
Figure 16
Correlograms for base, skew and smile volatilities
EURUSD



The left column are the ACFs for the volatility time series with fixed expiration 1 month. The right column are the ACFs for the first differences of the same series.

v. Daily changes in skew volatility show mild serial correlation, but base and smile volatilities show none. Figure 16 plots the autocorrelations for the volatility time series, as well as for the daily changes (i.e. the first differences). As for the S&P 500, the correlograms suggest that none of the three volatilities are stationary (the correlations show a linear decrease), but that they are first-difference stationary. For skew volatility there is evidence of modest positive autocorrelation in daily changes over the first few lags, indicating a mild tendency for moves to self-reinforce. The other two volatilities however show no serial correlation in daily changes.

Figure 17
Higher-order volatility cross-correlations
EURUSD



The left column are the correlations for the base, skew and smile volatilities, as functions of time to expiration. The right column are the correlations for daily changes in base, skew and smile volatilities, as functions of time to expiration.

vi. **Base and smile volatilities are highly correlated, while skew volatility is largely independent.** This holds both for the absolute quantities as well as for the daily changes. The correlations are plotted in Figure 17. The base-to-smile volatility correlation is above 0.95 for all three expirations, and the correlation for daily changes ranges from 0.84 (at 1-month expirations) to 0.88 (at the 1-year mark). In contrast, the skew-to-base and skew-to-smile volatility correlations range between -0.4 (1-month expirations) and +0.4 (1-year expirations); correlations for the daily changes are similarly low in magnitude.

Figure 18(a)
 Instantaneous skewness for five fixed expirations
 EURUSD





Figure 18(b)
 Instantaneous kurtosis for five fixed expirations
 EURUSD



The dashed lines are the EURUSD rate.

vii. **Instantaneous skewness is variable and changes sign, while instantaneous kurtosis is stable.** Figure 18(a) plots the skew-to-base ratio. Similarly to the skew volatility time-series, instantaneous skewness is negative over much of the year, except for a positive episode at the beginning (especially over February) and over the month of October. These positive episodes coincide with troughs in the time-series for EURUSD, subsequent to which there is an abrupt rise in the exchange rate. Over the rest of the year, instantaneous skewness appears to have a term structure that increases with time to expiration. In contrast, instantaneous kurtosis in figure 18(b) is stable over the entire year, albeit with a mild upward drift. The term structure is very pronounced, rising from about 0.5 at the 1-month mark to 1 at the 12-month mark.


viii. **Daily changes in the volatilities are largely uncorrelated with daily returns.** Table 2 reports the correlation coefficients between daily returns and daily changes in base, skew and smile volatilities, for three fixed times to expiration. In contrast to the S&P 500, the correlation coefficients are small and no greater than 0.25 in magnitude. For the sake of symmetry with our treatment of the equity index, we plot daily returns versus daily changes in 1-month base, skew and smile volatility, in figures 19(a)-(c). The weakness of the association is confirmed by an R^2 statistic that is close to zero in all three cases.

Table 2
Correlation between daily returns and daily changes in the volatilities,
for three fixed expirations.
EURUSD

Time to expiration	Correlation between daily returns and daily changes in ...		
	Base volatility	Skew volatility	Smile volatility
1 month	0.2484	-0.1986	0.2400
3 months	0.1989	-0.1553	0.1202
1 year	0.1415	-0.0708	0.1692

Figure 19(a)


Daily changes in base volatility versus daily returns, EURUSD



The best fit line has slope=0.06, intercept=0.00, and $R^2=0.06$. The volatility series used is for 1-month expirations.


Figure 19(b)

Daily changes in skew volatility versus daily returns, EURUSD



The best fit line has slope=-0.11, intercept=0.00, and $R^2=0.04$.

Figure 19(c)
Daily changes in smile volatility versus daily returns
EURUSD



The best fit line has slope=0.03, intercept=0.00, and $R^2=0.06$.

APPENDIX 1: CHARACTERISTIC FUNCTION OF LOG RETURNS

Here we derive the characteristic function of log returns from Section 1. To avoid repeated technical digressions, we group some auxiliary results below, which can be skipped with no loss of intuition.

$$\text{There exists a } \lambda > 0 \text{ such that } |\delta X_u / X_u| < 1 \text{ whenever } \delta u < \lambda. \quad (\text{R1})$$

$$\left| (\delta X_u / X_u)^{j+1} \right| < \left| (\delta X_u / X_u)^j \right| < 1 \text{ and } \lim_{j \rightarrow \infty} (\delta X_u / X_u)^j = 0, \delta u < \lambda. \quad (\text{R2})$$

$$\left| \Sigma_{j+1,u}^{j+1} \right| < \left| \Sigma_{j,u}^j \right| < 1 \text{ and } \lim_{j \rightarrow \infty} \Sigma_{j,u}^j = 0, \delta u < \lambda. \quad (\text{R3})$$

$$\left| \sigma_{j+1,u}^{j+1} \right| < \left| \sigma_{j,u}^j \right| < 1 \text{ and } \lim_{j \rightarrow \infty} \sigma_{j,u}^j = 0. \quad (\text{R4})$$

$$\left| \int_t^T \sigma_{j,u}^j du \right| < \int_t^T \left| \sigma_{j,u}^j \right| du < 1, j = 1, 2, \dots, \delta u < \lambda. \quad (\text{R5})$$

$$\int_t^T \sigma_{j,u}^j du = \lim_{\delta u \rightarrow 0} \sum_u \Sigma_{j,u}^j \delta u. \quad (\text{R6})$$

$$\text{The series } \sum_{j=1}^n \binom{z}{j} (\delta X_u / X_u)^j, n = 1, 2, \dots \text{ converges absolutely, } \delta u < \lambda, z \in \mathbb{C}. \quad (\text{R7})$$

$$\text{The series } \sum_{j=1}^n \binom{z}{j} \Sigma_{j,u,\delta u}^j \delta u, n = 1, 2, \dots \text{ converges uniformly in } \delta u < \min(\lambda, 1), z \in \mathbb{C}. \quad (\text{R8})$$

$$\text{The series } \sum_{j=1}^n \binom{z}{j} \int_t^T \sigma_{j,u}^j du, n = 1, 2, \dots \text{ converges absolutely, } z \in \mathbb{C}. \quad (\text{R9})$$

(R1) restates an assumption from Section 1. (R2)-(R4) each follow from the preceding result with no difficulty. (R5) holds by the monotonicity property of integrals. (R6) holds since $\Sigma_{j,u}^j \delta u = \sigma_{j,u}^j \delta u + o(\delta u)$ and by (R4). (R7) holds by the ratio test using (R1). (R8) holds by the Weierstrass M-test since:

$$\left| \binom{z}{j} \Sigma_{j,u,\delta u}^j \delta u \right| < \left| \binom{z}{j} \right| = M_j$$

and the partial sums $\sum_j M_j$ converge (absolutely) by the ratio test. (R9) holds by the ratio test using (R5).

Moving on to the derivation of the c.f., we assume that $\delta u < \min(\lambda, 1)$. Note first that:

$$E_t \left[(X_T / X_t)^z \right] = E_t \left[\prod_u (X_{u+\delta u} / X_u)^z \right] = \prod_u E_u \left[(X_{u+\delta u} / X_u)^z \right], \quad (\text{A1.1})$$

where the product is over the time indices $u = t, t + \delta u, \dots, t + n\delta u = T$. Here we used the law of iterated expectations, together with the assumption that the relative increments are market-independent. Now:

$$(X_{u+\delta u}/X_u)^z = (1 + \delta X_u/X_u)^z = 1 + \sum_{j=1}^{\infty} \binom{z}{j} (\delta X_u/X_u)^j,$$

where we have used the identity:

$$(1 + \beta)^z = 1 + \sum_{j=1}^{\infty} \binom{z}{j} \beta^j, \quad |\beta| < 1.$$

Then by (R7):

$$E_u (X_{u+\delta u}/X_u)^z = 1 + \sum_{j=1}^{\infty} E_u \left[\binom{z}{j} (\delta X_u/X_u)^j \right] = 1 + \sum_{j=1}^{\infty} \binom{z}{j} \sigma_{j,u,\delta u}^j \delta u. \quad (\text{A1.2})$$

Combining (A1.1) and (A1.2) gives:

$$E_t \left[(X_T/X_t)^z \right] = \prod_u \left[1 + \sum_{j=1}^{\infty} \binom{z}{j} \Sigma_{j,u}^j \delta u \right]. \quad (\text{A1.3})$$

Observe next that:

$$1 + \sum_{j=1}^{\infty} \binom{z}{j} \Sigma_{j,u}^j \delta u = \exp \left[o(\delta u) + \sum_{j=1}^{\infty} \binom{z}{j} \Sigma_{j,u}^j \delta u \right], \quad (\text{A1.4})$$

where $o(\delta u)$ represents terms which vanish with δu faster than δu (that is, $o(\delta u)/\delta u \rightarrow 0$ as $\delta u \rightarrow 0$). Replacing (A1.4) into (A1.3) and simplifying yields:

$$E_t \left[(X_T/X_t)^z \right] = \exp \left[\frac{o(\delta u)}{\delta u} (T-t) + \sum_u \sum_{j=1}^{\infty} \binom{z}{j} \Sigma_{j,u}^j \delta u \right]. \quad (\text{A1.5})$$

By the definition of instantaneous volatility, along with (R3), (R6), (R8) and (R9):

$$\lim_{\delta u \rightarrow 0} \sum_u \sum_{j=1}^{\infty} \binom{z}{j} \Sigma_{j,u}^j \delta u = \sum_{j=1}^{\infty} \binom{z}{j} \lim_{\delta u \rightarrow 0} \sum_u \Sigma_{j,u}^j \delta u = \sum_{j=1}^{\infty} \binom{z}{j} \int_t^T \sigma_{j,u}^j du.$$

Thus, taking the limit $\delta u \rightarrow 0$ in (A1.5) gives:

$$E_t \left[(X_T/X_t)^z \right] = \exp \sum_{j=1}^{\infty} \binom{z}{j} \int_t^T \sigma_{j,u}^j du.$$

Setting $z = iv$ yields the c.f. $\phi(v) = E_t \left[\exp \{ iv \ln(X_T/X_t) \} \right]$ as required.

APPENDIX 2: INSTANTANEOUS MOMENTS AND CUMULANTS

The raw moments μ_j and cumulants c_j of a random variable satisfy the relation:

$$\mu_j = c_j + \sum_{k=1}^{j-1} \binom{j-1}{k-1} \mu_{j-k} c_k \quad (\text{A2.1})$$

(Stuart and Ord (1994), p.119). Let m_j denote the j -th infinitesimal moment of a stochastic process; that is, the j -th raw moment of the process increment over an interval of time δt has the form $\mu_j = m_j \delta t + o(\delta t)$. To prove that m_j is also the j -th infinitesimal cumulant we need to show that:

$$c_n = m_n \delta t + o(\delta t) \quad (\text{A2.2})$$

for every $n=1,2,\dots$. We proceed by induction. Note first that equation (A2.2) holds for $n=1$ by virtue of the identity $\mu_1 = c_1$. Next, suppose that (A2.2) holds for $n = j-1$. Then (A2.1) reads:

$$c_j = m_j \delta t + o(\delta t) - \sum_{k=1}^{j-1} \binom{j-1}{k-1} [m_k \delta t + o(\delta t)] [m_{j-k} \delta t + o(\delta t)].$$

Since the sum on the right is $o(\delta t)$, it follows that (A2.2) holds for $n = j$, and the result follows.

APPENDIX 3: FX MARKET QUOTES

The FX option quotes are handled as follows. Let ATM, RR and STR denote respectively the at-the-money, risk-reversal and strangle volatility quotes. The implied volatilities for the 25 delta call and (-)25 delta put are computed as:

$$\begin{aligned}\sigma_C &= ATM + STR + \frac{1}{2} RR \\ \sigma_P &= ATM + STR - \frac{1}{2} RR\end{aligned}$$

as per the BBA documentation, Wystup (2006) and the treatment in Carr and Wu (200X).⁴

For EURUSD, the premium currency is USD, and the deltas are of the regular spot kind. To determine the strike prices corresponding to a given delta level, we use the following formula:

$$K = F \exp\left[-\varepsilon\sigma\sqrt{\tau}N^{-1}\left(\varepsilon e^{r_f\tau}\Delta\right) + \frac{1}{2}\sigma^2\tau\right],$$

where $\varepsilon=1$ for a call option, -1 for a put. For example, the ATM strike is $K = F \exp\left[\frac{1}{2}\sigma^2\tau\right]$, corresponding to the ATM $\Delta = \frac{1}{2}\varepsilon e^{r_f\tau}$. For more details see Reiswich and Wystup (2009).

⁴ This may not correspond to market practice. There are sources which detail a different interpretation of market quotes, but they are both sparse and conflicting—these include Hakala et al (2008) and Reiswich and Wystup (2009).

REFERENCES

- Abramovitz, M. and I. Stegun (1972) *Handbook of Mathematical Functions*, 10th printing, National Bureau of Standards, United States Department of Commerce.
- Babbs, S. and M. Selby (1998) Pricing by arbitrage under arbitrary information, *Mathematical Finance* **8**(2), 163-168.
- Bates, D. (1991) The crash of '87: Was it expected? Evidence from options markets, *Journal of Finance* **46**(3), 1009-1044.
- Black, F. and M. Scholes (1973) The pricing of options and corporate liabilities, *Journal of Political Economy* **81**(3), 637-654.
- Brody, D., L. Hughston and A. Macrina (2007) Information-based asset pricing, *International Journal of Theoretical and Applied Finance* **11**(1), 107-142.
- Carey, A. (2005) Higher-order volatility, working paper, SSRN e-library.
- Carey, A. (2006) Higher-order volatility: dynamics and sensitivities, working paper, SSRN e-library.
- Carr, P. and L. Wu (2007) Stochastic skew in currency options, *Journal of Financial Economics* **86**, 213-247.
- Daglish, T., J. Hull and W. Suo (2007) Volatility surfaces: theory, rules of thumb, and empirical evidence, *Quantitative Finance* **7**(5), 507-524.
- Gikhman, I. and A. Skorokhod (1972) *Stochastic Differential Equations*, Springer-Verlag.
- Hakala, J., R. Stuart and T. Hamosfakidis (2008) FX derivatives, presentation slides, Standard Chartered Bank.
<http://www.mathfinance.com/workshop/2008/papers/hakala/slides.pdf>
- Johannes, M. (2004) The statistical and economic role of jumps in continuous-time interest rate models, *Journal of Finance* **59**(1), 227-260.
- Merton, R. (1973) Theory of rational option pricing, *Bell Journal of Economics and Management Science* **4**(1), 141-183.
- Merton, R. (1976) Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* **3**(1-2), 125-144.
- Reiswich, D. and U. Wystup (2009) FX volatility smile construction, working paper, Frankfurt School of Finance and Management.
- Stuart, A. and K. Ord (1994) *Kendall's Advanced Theory of Statistics (Volume 1)*, 6th edition, Arnold.
- Wystup, U. (2006) *FX Options and Structured Products*, Wiley.