

Vega Risk in RiskManager

Model Overview

Lisa R. Goldberg

Michael Y. Hayes

Ola Mahmoud (Marie Cure Fellow)¹

Tamas Matrai

May 2011

¹ Ola Mahmoud was a Marie Curie Fellows at MSCI. The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7-PEOPLE-ITN-2008 under grant agreement number PITN-GA-2009-237984 (project name: RISK). The funding is gratefully acknowledged.

Contents

Contents	2
1 Introduction	3
2 Background	3
2.1 <i>Implied Volatility Concepts</i>	3
2.2 <i>Vega Risk</i>	7
3 Understanding the Smile.....	8
3.1 <i>Fixed-Delta and Fixed-Strike Volatility Models</i>	8
3.2 <i>Smile Dynamics</i>	11
4 Vega Risk in RiskManager	14
5 The New Vega Risk Model.....	16
6 Backtesting the Model	16

1 Introduction

Vega risk is the risk due to variations in volatility, or the *volatility of volatility*. An account of Vega risk is important in any portfolio containing options, as it can generate exposure to the option's volatility, even if all the other Greeks are hedged away. The key idea to the implementation of Vega risk in RiskManager is to treat implied volatility as any other risk factor that portfolios are exposed to. The main difference, however, lies in the existence of the volatility smile, that is the different values of implied volatility at a given date depending on the moneyness of the option.

This paper gives an overview of the Vega risk implementation in RiskManager via risk factors which account for smile risk. We start by reviewing some basic concepts on option implied volatility and its risk. We then carry out some diagnostic studies to help us gain a better understanding of the implied volatility smile and its risk. Two enhancements to the existing RiskManager model are introduced and tested.

2 Background

Remark on why one needs to study and account for implied volatility and its risk:

- Gives market view of risk-neutral volatility
- Determines the volatility risk premium
- Measures the forward-looking leverage effect
- Important risk factor for options

2.1 Implied Volatility Concepts

The Black-Scholes Framework

An implied volatility is linked to a particular option valuation model, the most common one being that of Black-Scholes. This is the model adopted in RiskManager, under which the price of a call option on an underlying asset is given by

$$C = Se^{-q\tau}N(d_1) - Ke^{-r\tau}N(d_2),$$

where S is the spot price, q is the dividend yield, τ is the time to maturity, K is the strike, r is the risk-free interest rate, N is the standard normal cumulative distribution function, and

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

The price P of a put option on the same underlying (and with same strike and maturity) can be immediately derived via Put-Call-Parity:

$$C + Ke^{-r\tau} = P + Se^{-q\tau}.$$

As a foreign currency is analogous to a stock paying a known dividend yield, namely the foreign risk-free interest rate, and FX option can be priced using the same formula above, where the spot is the value of one unit of foreign currency in the domestic currency.

The Black-Scholes valuation formula implies a positive relationship between the price of an option and the volatility parameter, assuming everything else is kept constant. Figure 1 illustrates that as the volatility increases, so does the premium.

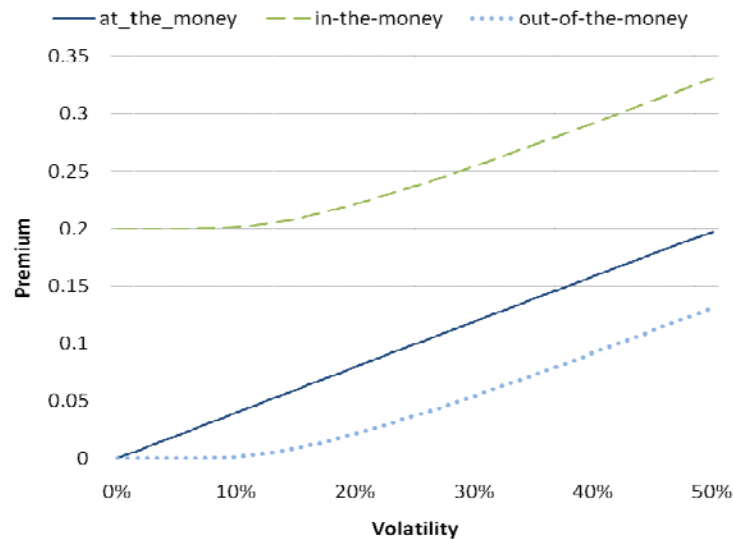


Figure 1: Value of a generic 1-year-to-maturity call option as a function of volatility. Risk-free rate and yield are set to zero, spot is taken to equal 1, in- and out-of-the-money calls have strikes equal to spot minus/plus 0.2, respectively.

Moneyiness

The value of option *moneyiness* can be taken in absolute terms as strike or in relative terms as strike divided by the forward price of the underlying asset (for a forward contract maturity equal to that of the option). The RiskManager convention is to take Black-Scholes *Delta* δ as a measure of moneyiness. Delta is the first-order derivative of the Black-Scholes option price with respect to the underlying. For a given spot S and volatility σ , the call delta δ_c and put delta δ_p are given by

$$\delta_c = e^{-qT} N(d_1) \text{ and } \delta_p = \delta_c - e^{-qT} .$$

Under this approach to moneyiness, an at-the-money volatility is taken to be that implied by a call option with a delta close to 0.5 (or 50). We will simply refer to at-the-money implieds as *50-Delta*, hence assuming that the ATM implied volatility quoted by the market refers to a 50-Delta call implied volatility. This is not exactly true, but it is a good approximation.

Another common convention used in practice and in the RiskManager Vega risk model is to take the Delta of a call option to be 1 minus the Delta of a put (with same maturity and strike). This is again not strictly true, since the Delta of a call and the Delta of a put do not sum strictly to one, but to one multiplied by the discount factor of the dividend yield (or the foreign risk-free interest rate for an FX option).

Under these two assumptions, one may map option prices and their implied volatilities (as defined next) as a function of call Delta, with Delta running from 0 to 100 (or simply from 0 to 1), thereby spanning the entire range of all possible values.

Implied Volatilities

All inputs to the Black-Scholes option pricing formula are observable, except for the volatility parameter. Given a market price for an option, one may back out a volatility value which equates the formula and

market prices. This is precisely the *implied volatility*, namely the volatility implied by option prices observed in the market under a particular valuation model (in our case that of Black and Scholes).

Implied volatilities are a forward looking measure of the market’s perception of volatility. Indeed, traders use implied volatilities (of the most actively traded options) to monitor the market’s opinion about the expected volatility level. In the context of the Black-Scholes model, an option implied volatility can therefore be interpreted as the risk neutral estimate of the volatility of the returns to the underlying asset over the life of the option.

Under the assumptions of the Black-Scholes model, asset return volatility is constant, which is rarely supported by empirical research. Realized as well as option implied volatilities fluctuate over any given period of time – hence the importance of Vega. Moreover, empirically observed option implied volatilities differ for options with different strikes and different maturities. Any portfolio of options is therefore not only affected by changes in the level volatility over time, but also by changes in the *volatility smile* and the *term structure of volatility*, as described next.

Volatility Smiles

For a given analysis date and option maturity, a plot of implied volatility as a function of moneyness is known as the *volatility smile*. Depending on the notion of moneyness one works with, the smile can be parameterized in different ways, as shown in Figure 2. The term ‘smile’ describes the characteristic shape of the volatility/moneyness plot: at-the-money options tend to have the lowest implied volatilities. The volatility smile for FX options has the general shape shown in Figure 3. The likelihood of extreme moves in exchange rates and the demand for insurance is reflected in the shape of the smile: the option price (and hence its implied volatility) increases as it moves away from being at-the-money in either direction.




Figure 2: Parameterizations of the volatility smile: implied volatility plotted as a function of strike K (left), moneyness x (middle), and Black-Scholes delta δ (right).

The volatility smile for equities is often referred to as *volatility skew*, as illustrated in Figure 3. The implied volatility decreases for options that have a larger upside, i.e. for deep in-the-money puts or out-of-the-money calls. Because equity options traded in the USA only revealed a skew-like behaviour after the market crash of 1987, one reason for the equity volatility skew is traders’ concern about the occurrence of another crash. Options are then priced accordingly, resulting in deep out-of-the-money puts and deep in-the-money calls being more expensive. [Another reason: Black’s leverage effect?]

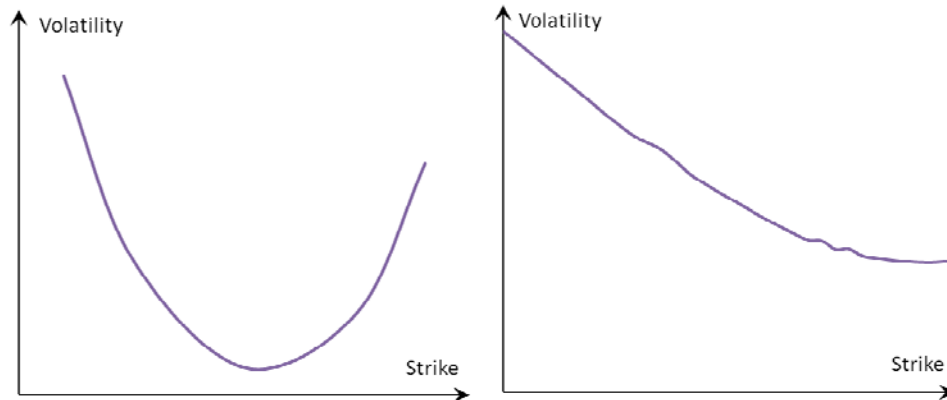


Figure 3: The volatility smile for options on GBP/USD (left) and on the STOXX50 (right). Implied volatility is plotted as a function of increasing strike levels, for the same spot and same time to maturity.

Term Structure of Volatility

The term structure of implied volatility describes, for a given moneyness or strike at a given point in time, the relationship between implied volatility and option maturity.

Volatility Surface

The *volatility surface* is simply a combination of the volatility smile and its term structure. It defines a relationship between implied volatility and the two dimensions of moneyness and maturity. Figure 4 shows the implied volatility surface on September 9th 2009 for options on the STOXX50.




Figure 4: The implied volatility surface for options on the STOXX50 on September 9th 2009.

2.2 Vega Risk

Definition of Vega

Portfolios containing options are exposed to a range of risks, the so-called *Greeks*, each of which measures the sensitivity of a given portfolio to some variable (spot, time, etc.). The constant volatility assumption of the Black-Scholes model implies no sensitivity with respect to volatility. However, as in practice implied volatilities change over time, the option value Π is also liable to changes in the movements of volatility.

The *Vega* V of an option Π is its sensitivity to the volatility σ of the underlying asset. It is formally defined as the partial derivative

$$V = \frac{\partial \Pi}{\partial \sigma},$$

which measures the rate of change of the value of the option with respect to changes in its implied volatility, while assuming that all other variables determining the option value (spot, strike, time to maturity, etc.) are held constant. In the Black-Scholes model, the above expression for Vega V is thereby given as

$$V = \sqrt{\tau} S e^{-q\tau} \phi(d_1),$$

where all the variables are defined as in Section 2.1, and ϕ is the standard normal density function. More intuitively, Vega expresses the change of the option value for every 1% change of its implied volatility. In practice, it is measured and quoted in units of the base currency of the option.

Properties of Vega

Under the Black-Scholes model and its related formulas above, some fundamental properties of Vega and its behaviour with respect to the other parameters can be inferred. First, since the value of an option increases for higher volatility (recall Figure 1), the Vega of an option is always positive. Moreover, when plotted against moneyness (and assuming all other variables of the Black-Scholes valuation function are held constant), Vega is at its maximum for an at-the-money option, as illustrated in Figure 5. This implies that the value of an at-the-money option is approximately linearly (and positively) proportional to its implied volatility. Finally, Vega increases with longer time to maturity, also shown in Figure 5.




Figure 5: Vega as a function of log moneyness, for strike K and forward F . Each curve corresponds to a different time-to-maturity and depicted in decreasing maturity order.

3 Understanding the Smile

Understanding the shape of the smile and how it evolves over time is an essential prerequisite to developing the Vega risk model. In this Section, we review fundamental properties of the volatility smile in the equity and FX space. We also perform some exploratory analyses to help us obtain a better picture of smile dynamics.

3.1 Fixed-Delta and Fixed-Strike Volatility Models

Derman [?] introduced models that parameterize the volatility smile in equity index markets. This development is motivated by the empirical observation that the way in which the volatility smile evolves in the future depends on the regime of the underlying, for example whether it is trending, jumpy or constrained within a specific range. Derman’s so-called *sticky* models aim to describe the aspects of the smile that remain invariant (or ‘sticky’) as the underlying moves. As these were developed for equity index models, the assumption is an implied volatility skew which varies linearly when the strike moves away from its at-the-money value. We recall Derman’s *sticky-strike* and *sticky-delta* models, which are relevant to the RiskManager Vega risk model.

Sticky-Strike

The assumption in a sticky-strike model is that the volatility of an option with a particular strike is independent of the level of the underlying: if the spot moves, the fixed-strike volatility, which we denote by σ_k throughout, remains unchanged. In equity markets, and under the assumption of a monotonically decreasing volatility skew as a function of strike over spot, the at-the-money volatility σ_{ATM} decreases as the spot increases when fixing the strike. The change in the option value in this model is affected only by its moneyness, and the total delta of the option therefore equals the Black-Scholes delta.

The sticky-strike rule is often the underlying model used when the movement of the underlying is bound to a specific range. In a stable range-bounded market, the volatility is likely to maintain its current level without significant increases or decreases, and so the implied volatility of an option held at a given strike remains unchanged.

Sticky-Delta

In a sticky-delta model, one assumes that the volatility of an option depends only on its delta (thought of as its moneyness), and its dependence on the underlying movement comes implicitly from its dependence on delta, which is defined in terms of strike and spot. The idea in this model is that as the spot moves, the fixed-delta volatility σ_δ , in particular the at-the-money volatility $\sigma_{ATM} \approx \sigma_{50\delta}$, remains unchanged. To keep a constant moneyness, the volatility of a given fixed strike increases with the spot. Straightforward calculus shows that the total delta of the option is greater than the Black-Scholes delta:

$$\Delta \equiv \frac{d\Pi}{dS} = \frac{\partial\Pi}{\partial S} + \frac{\partial\Pi}{\partial\sigma} \frac{\partial\sigma}{\partial S} = \delta + V \frac{\partial\sigma}{\partial S},$$

with *smile correction* term given by

$$\frac{\partial\sigma}{\partial S} = \frac{\partial\sigma}{\partial\delta} \frac{\partial\delta}{\partial S} = \gamma \frac{\partial\sigma}{\partial\delta} = \gamma \frac{\partial}{\partial\delta} (a\delta^2 + b\delta + c) = \gamma(2a\delta + b).$$

The sticky-delta model is usually used in an upward trending market, where the underlying is changing positively in its level, but its realized volatility is not undergoing significant variation. Under this stable-trending assumption, the implied volatility of an at-the-money option can be estimated via the realized volatility. Therefore, since the realized volatility remains unchanged, so does the at-the-money volatility.

Sticky Models in FX space

Derman’s sticky models have been developed for options on equity indices. The underlying assumption is therefore a negatively sloping volatility smile. The shape of the smile for options on currencies is not necessarily downward sloping, but is generally parabola-shaped or sometimes also upward sloping. In that case, the way fixed-delta and fixed-strike volatilities vary with a movement in the underlying is altered as well.

Comparing Fixed-Delta and Fixed-Strike Risk Factors

By *fixed-delta risk factors* we mean changes (or returns) to implied volatility for a fixed value of delta. Fixed-delta volatilities remain constant under the assumptions of Derman’s sticky-delta model. Similarly, *fixed-strike risk factors* are changes to implied volatility for a fixed strike, which remains unchanged in a sticky-strike model

RiskManager currently employs fixed-delta implied volatilities in its Black-Scholes valuation function, the implication of which we analyse in our backtesting analyses in Section 6. Here, we aim to understand the general behaviour of fixed-delta risk factors and their fixed-strike counterparts.

The data used are implied volatilities and implied strikes for fixed deltas lying in the range between 20 and 80 for equities, and 10 to 90 for FX options. The number of fixed delta nodes we are given by the vendor (?) are 23 and 5 for equities and FX, respectively. Fixed-strike risk factors are obtained from these fixed-delta nodes as follows:

Volatility of risk factors

We find that generally fixed-strike implied volatility returns are less volatile than their fixed-delta counterparts, across all maturities and option underlyings. Figure 6 displays this property for 1-month-to-maturity options on the S&P 500.

[Possible explanations?]




Figure 6: 21-day exponentially weighted moving average of the volatilities of fixed-delta and fixed-strike risk factors (at Delta = 50) for 1-year-to-maturity options on the S&P 500.

Correlation with the underlying

Assuming a downward sloping volatility smile as in the equity options market, the sticky-strike model implies negative correlation between the fixed-delta implied volatilities and the underlying. In comparison, fixed-strike implied volatilities show less negative correlation. Figure 7 shows the correlation with the underlying for fixed-delta and fixed-strike risk factors.

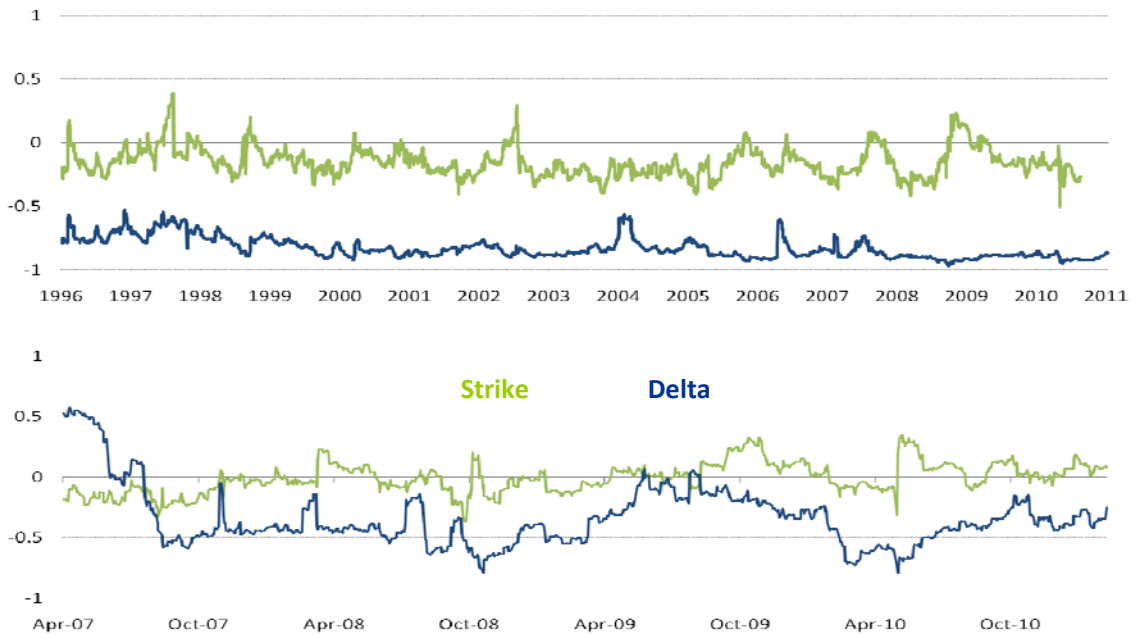


Figure 7: 21-day exponentially weighted moving average of the correlations of fixed-delta and fixed-strike risk factors (at Delta = 50) with the underlying. Top Figure: 1-year-to-maturity options on the S&P 500. Bottom Figure: 1-month-to-maturity options on the GBP/USD exchange rate.

Forecasting implied volatility returns

One main question is when one should use one or the other risk factor in practice. As RiskManager currently only provides fixed-delta risk factors, we investigate the value added by complementing (or replacing) these with fixed-strike risk factors. To this end, we compare the forecasting accuracy of implied volatility based on the two risk factors.

Assume we aim to forecast fixed-strike implied volatility changes. These define the price changes of an actual options contract. One important aspect of this forecast is the sensitivity of the risk factors to the underlying returns. We model the returns $r(\sigma_K)$ to the fixed-strike implied volatilities σ_K as

$$r(\sigma_K) = \beta r_S + \varepsilon$$

where β is the sensitivity of the given risk factor to the spot, r_S is the return to the spot, and ε is the residual. When using the fixed-strike risk factor to forecast fixed-strike implied volatility changes, our model is therefore simply

$$r(\sigma_K) = \beta(\sigma_K) r_S + \varepsilon_K$$

where $\beta(\sigma_K)$ is the sensitivity of fixed-strike volatility σ_K to the spot. For the fixed-delta risk factors, one needs to account for smile correction:

$$r(\sigma_K) = \left[\beta(\sigma_\delta) + \frac{d\sigma}{dS} \right] r_S + \varepsilon_\delta$$

where $\beta(\sigma_\delta)$ is the sensitivity of fixed-strike volatility σ_δ to the spot, and $d\sigma/dS$ corresponds to *riding the smile*.

We back-test the above two models for 1-month-to-maturity options on the S&P 500 and on GBP/USD currency exchange rates. A 50-day moving average of daily fixed-strike and fixed-delta squared residuals, ϵ_K and ϵ_δ respectively, is shown in Figure 8.

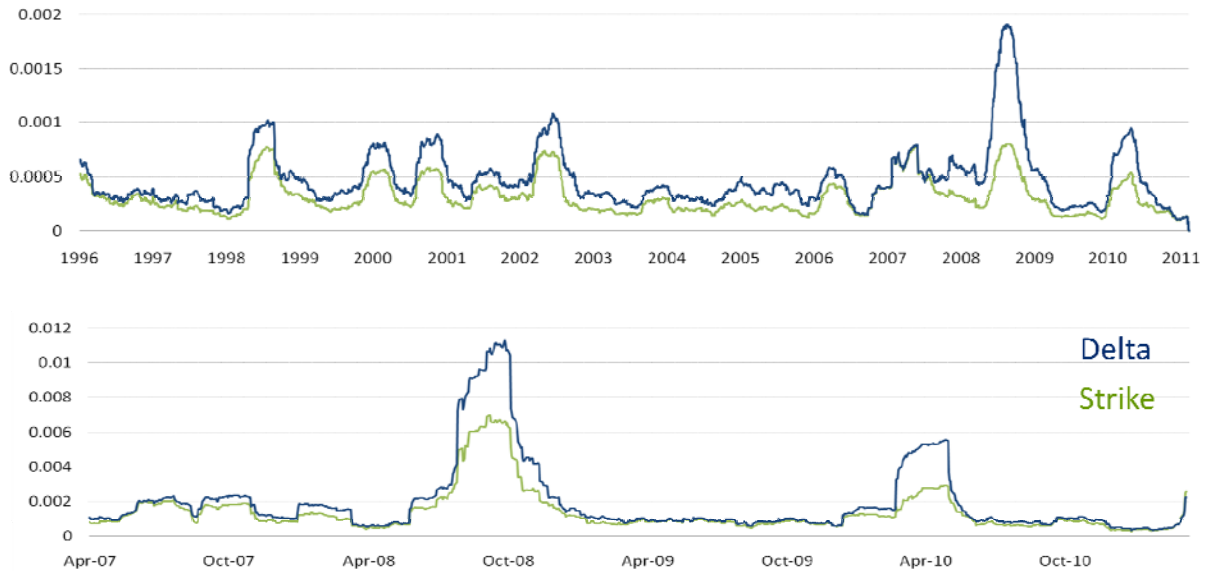


Figure 8: 50-day moving average of fixed-delta and fixed-strike squared residuals for 1-month-to-maturity options on the S&P 500 (top graph) and GBP/USD (bottom graph).

Clearly, fixed-strike implied volatility returns consistently yield smaller residuals over time, especially during some of the turbulent market times, hence providing a better risk factor for forecasting purposes. We finally point out that we have performed an analogous analysis for other times to maturity and other equity and FX option underlyings, all of which produced similar results.

3.2 Smile Dynamics

Principal Component Analysis

We use the standard methodology for extracting the most important uncorrelated sources of variation in the volatility smile via PCA, for both equity and currency markets. An analysis of the principal components of daily changes in the fixed-delta and fixed-strike volatilities for 1-month-to-maturity options on the S&P 500 and GBP/USD currencies are shown in Figures 9 and 10 with call implied volatilities for 13 and 5 different delta-nodes, respectively. Fixed-strike returns are obtained using linear interpolation as described in Section (?).

Fixed-Delta			Fixed-Strike		
Component	Eigenvalue	Cumulative R ²	Component	Eigenvalue	Cumulative R ²
P1	0.08337	0.9910	P1	0.11004	0.8883
P2	0.00069	0.9992	P2	0.01264	0.9905
P3	3.85E-05	0.9997	P3	7.65E-04	0.9966

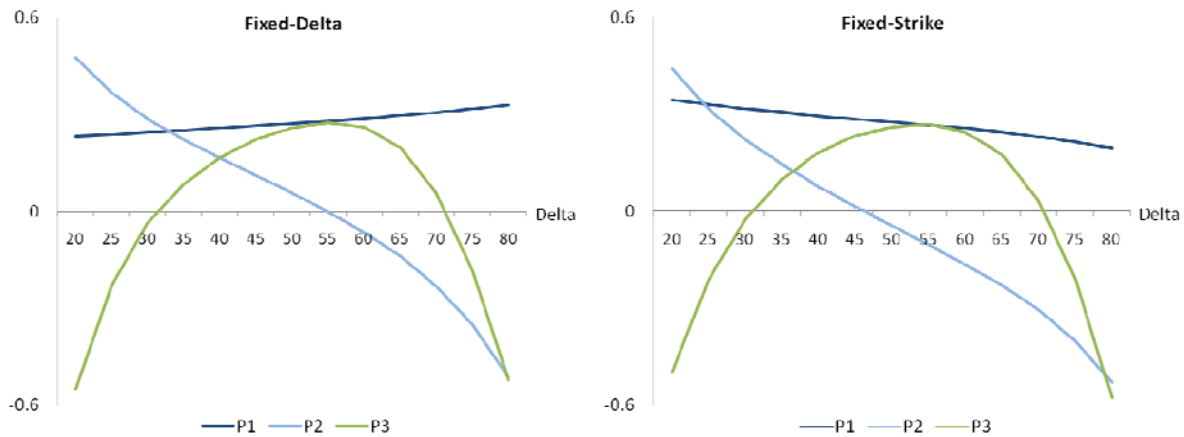


Figure 9:

Fixed-Delta			Fixed-Strike		
Component	Eigenvalue	Cumulative R ²	Component	Eigenvalue	Cumulative R ²
P1	0.00726	0.9978	P1	0.01912	0.9019
P2	0.00016	0.9991	P2	0.00181	0.9874
P3	6.21E-06	0.9999	P3	2.59E-04	0.9996




Figure 10:

Almost 100% of the total variation of both fixed-delta and fixed-strike implied volatilities can be explained by just three risk factors: parallel shift, tilt and curvature. The parallel shift component of the fixed-delta returns accounts for almost all of this variation (99%), whereas it contributes with 88%-90% to the variation of the fixed-strike implied volatility changes. The standard shift-twist-butterfly trend

revealed from this analysis leads us to consider parameterizing the smile in terms of three risk factors which capture almost all of the variation in the smile dynamics.

Quadratic Smiles

Standard PCA disentangled the three types of sources affecting the shape of the smile. RiskManager uses a quadratic model in terms of these three risk factors to describe smile dynamics. Using this model, the implied volatilities for any other level of delta can be easily inferred. The risk factors used are inspired by the following three most liquid FX option structures:

- *ATM straddle* ('ATM'), which is given by the sum of a call and a put struck at the at-the-money level. Its implied volatility σ_{ATM} is approximated by the volatility at 50-Delta.
- *25-Delta risk reversal* ('RR'), which is given by going long a 25-Delta call and short a 25-Delta put (or vice versa). Its volatility σ_{RR} is given by $\sigma_{RR} = \sigma_{25} - \sigma_{75}$.
- *25-Delta Vega-weighted butterfly* ('VWB'), which is built by selling an ATM straddle and buying a symmetric *25-Delta strangle*. The latter is simply the sum of a base currency call and put, both struck at 25 Delta. Its volatility is given by $\sigma_{VWB} = 0.5 (\sigma_{25} + \sigma_{75}) - \sigma_{ATM}$.

Although designed by Malz [?] for the FX options market, these risk factors can be applied in equity space as well. The 50-Delta volatility can be considered the central point of the smile, and the 25- and 75-Delta volatilities are roughly the *wings* of the smile. Malz's framework for studying the shape and movement of the smile uses the moneyness conventions described in Section 2.1. His model uses a second-order polynomial in Delta to fit the above three implied volatility nodes. This way, an implied volatility σ_δ for a given Delta δ (and a fixed time to maturity) can be interpolated from the smile described by

$$\sigma_\delta = a \cdot \sigma_{ATM} + b \cdot \sigma_{RR} \cdot (\delta - 0.5) + c \cdot \sigma_{VWB} \cdot (\delta - 0.5)^2.$$

Under the assumption that 50-Delta is simply ATM, we can set the coefficient a to be 1. Similarly, using the same assumptions and substituting the expressions for σ_{RR} and σ_{VWB} explicitly, we can derive that $b = 2$, and $c = 16$. The volatility smile can therefore be approximated as

$$\sigma_\delta = \sigma_{ATM} + 2 \cdot \sigma_{RR} \cdot (\delta - 0.5) + 16 \cdot \sigma_{VWB} \cdot (\delta - 0.5)^2.$$

This model will be used to understand how the smile changes. Recalling the PCA study performed in the previous section, one can in fact disentangle from this quadratic model the same three basic movements of the smile: σ_{ATM} is the parallel up and down movement of the smile; σ_{RR} describes the slope movement (large values imply steeper smiles), and σ_{VWB} is the curvature movement (large values imply a more convex smile). Therefore, this quadratic model is essentially a function in the principal components of the smile. Indeed, these three risk factors coincide with the first three principal components extracted earlier, in terms of their volatilities and correlations with the underlying spot, details for which will be skipped here.

A second-order polynomial fit to three nodes generally performs quite well in the range within the two outermost nodes. However, it is unclear whether implied volatilities are over- or underestimated for high (> 75) or low (< 25) Deltas. To understand how well Malz's quadratic parameterization explains implied volatilities at large and small Deltas, we compare empirical volatility returns at the outermost nodes to those implied by the quadratic model. We also perform the same analysis, but using the fixed-strike risk factors [EXPLAIN]. The data we use are for 1-month-to-maturity options on the S&P 500, with smallest and largest option implied volatilities available at 20- and 80-Delta, respectively, and for 1-month-to-maturity currency options on GBP/USD, with smallest and largest volatilities given at the 10- and 90-Delta nodes, respectively. Table 1 shows the percent unexplained variance when regressing the

empirical volatility nodes against their quadratic estimators. Note that these numbers imply values of R^2 in the range of approximately 97%–100%. Similar analyses performed on other equity and FX underlyings and other option maturities consistently show similar R^2 numbers.

	S&P 500		GBP/USD	
	Delta = 20	Delta = 80	Delta = 10	Delta = 90
Fixed-Delta	3.14%	2.44%	0.17%	0.25%
Fixed-Strike	1.54%	2.33%	0.24%	0.48%

Table 1: Average unexplained variance in the volatility smile wings for 1-month-to-maturity options on the S&P 500 and on the GBP/USD exchange rate. The average is taken over 1996 – 2011 for S&P 500 and 2007 – 2011 for GBP/USD.

4 Vega Risk in RiskManager

The diagnostic studies carried out in Section 3 give us a clearer picture of the shape of the smile and how to parameterize it. The underlying idea in accounting for the risk of implied volatility and the smile in RiskManager is to treat it analogously to any other market risk factor that a portfolio is exposed to.

The Implied Volatility Risk Factor

Implied volatility risk factors are defined in RiskManager as logarithmic returns to fixed-delta implied volatilities. More intuitively, these are the proportional changes over time to implied volatility at a given delta. RiskManager assumes that these changes follow a random walk with zero mean.

Implied volatilities are naturally correlated with those of other options, with the underlying, and with other risk factors. The assumption in RiskManager is that all of these risk factors, which are represented in terms of logarithmic returns as defined above, are jointly normally distributed with mean zero and constant standard deviation. The basic approach is then to use Monte Carlo to simulate changes in implied volatilities, apply normally distributed shocks to the implied volatility, reprice the option, and finally computing risk forecasts (in particular VaR) from these simulated scenarios. Detailed explanation of each of these steps follows next.

4.1 A Simple No-Smile Case

Assume for the sake of illustration that volatilities were constant across all strikes for a given maturity. This is in agreement with the simple Black-Scholes model assumptions. The aim is to forecast 1-day VaR at some confidence level, say 95%.

Monte Carlo simulations under the joint-normality assumption are generated to calculate VaR numbers. The first step is therefore the estimation of the volatilities of the risk factors (volatility of implied volatility and spot) and their correlations. These can be equal-weighted historical estimates or exponentially weighted moving averages with a specified half-life. Multivariate normal scenarios are generated from these volatilities and correlations, where the mean is assumed to be zero. These scenarios are the shocks we use to perturb, say, today’s spot rate and implied volatility by applying the logarithmic changes to them. The position is next revalued and for each of these perturbed scenarios, and any risk measure, in particular VaR, can be ‘empirically’ calculated from the distribution of the scenario position changes.

4.2 Incorporating the Smile

We now incorporate the effect of the volatility smile in forecasting implied volatility risk. This methodology captures the differences in Vega risk among options on the same underlying, but with different maturities and different strikes.

The Base Scenario

The volatility smile on any given day is specified by the implied volatility risk factors of the ATM straddle, the risk reversal, and the Vega-weighted butterfly (recall Section 3.2), all of which are easily derived from the implied volatilities at the 25, 50, and 75 delta nodes only. The quadratic function

$$\sigma(\delta) = \sigma_{ATM} + 2 \cdot \sigma_{RR} \cdot (\delta - 0.5) + 16 \cdot \sigma_{VWB} \cdot (\delta - 0.5)^2$$

in terms of delta δ is then used to interpolate implied volatilities for any delta. We have seen in Section 3 that the quadratic fit to the implied volatility provides good explanatory power for the smile, even at its wings outside the 25-delta and 75-delta range. This model is referred to as the *quadratic base scenario*, because it models the volatility that will be perturbed by the Monte Carlo scenarios, i.e. we shock the base scenario to generate Monte Carlo scenarios.

Monte Carlo VaR

Perturbing the Smile

Because of the increased dimensionality that the smile introduces, implied volatility scenarios are now not only determined by the normally distributed shocks to the general level of implied volatility (as in Section 4.1), but also by a displacement *along the smile*. This displacement is determined by a shock to the underlying, as illustrated in Figure 11.

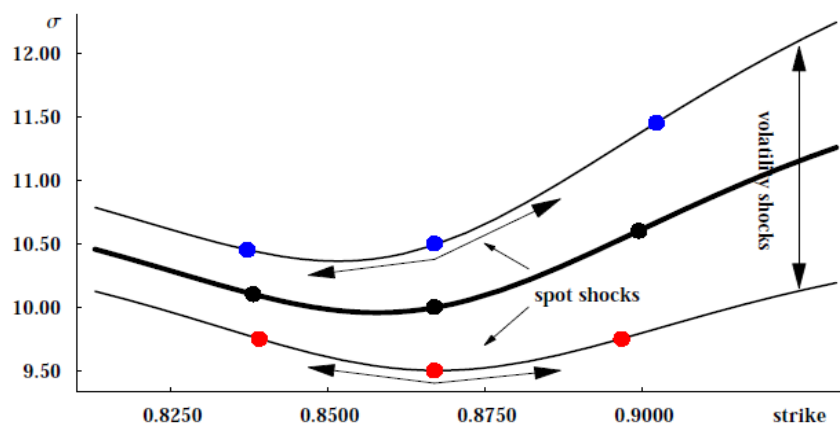


Figure 11: Perturbation of the implied volatility smile via shocks to the volatility and shocks to the spot.

Note also from this Figure that volatility shocks are assumed to induce only parallel shifts to the volatility smile, and not any changes to the shape of the smile. In the Monte Carlo scenario generation, we therefore combine correlated shocks to implied volatility and the underlying.

4.3 Summary of Methodology

5 The New Vega Risk Model

- Empirical base scenario
- Fixed-strike risk factors (explain how they are interpolated)

6 Backtesting the Model

- Overview of tests: volatility forecasts, delta forecasts, VaR/ES violations, etc
- Scope of backtests: STOXX 50, FX options, daily data, history, etc.
- A simple illustrative portfolio: single option portfolios (in FX)
- Explanation of the more complex portfolios we are testing (market-maker-type hedged portfolios)
- Explanation of backtesting methodology
- Results!

Acknowledgements

Angelo Barbieri, Thomas Ta, Carlo Acerbi, ...

References

Appendix A: Data Considerations

Client Service Information is Available 24 Hours a Day

clientservice@msci.com

Americas

Americas	1.888.588.4567 (toll free)
Atlanta	+ 1.404.551.3212
Boston	+ 1.617.532.0920
Chicago	+ 1.312.675.0545
Montreal	+ 1.514.847.7506
Monterrey	+ 52.81.1253.4020
New York	+ 1.212.804.3901
San Francisco	+ 1.415.836.8800
Sao Paulo	+ 55.11.3706.1360
Stamford	+1.203.325.5630
Toronto	+ 1.416.628.1007

Europe, Middle East & Africa

Amsterdam	+ 31.20.462.1382
Cape Town	+ 27.21.673.0100
Frankfurt	+ 49.69.133.859.00
Geneva	+ 41.22.817.9777
London	+ 44.20.7618.2222
Madrid	+ 34.91.700.7275
Milan	+ 39.02.5849.0415
Paris	0800.91.59.17 (toll free)
Zurich	+ 41.44.220.9300

Asia Pacific

China North	10800.852.1032 (toll free)
China South	10800.152.1032 (toll free)
Hong Kong	+ 852.2844.9333
Seoul	+827.0768.88984
Singapore	800.852.3749 (toll free)
Sydney	+ 61.2.9033.9333
Tokyo	+ 81.3.5226.8222

Notice and Disclaimer

- This document and all of the information contained in it, including without limitation all text, data, graphs, charts (collectively, the "Information") is the property of MSCI Inc. or its subsidiaries (collectively, "MSCI"), or MSCI's licensors, direct or indirect suppliers or any third party involved in making or compiling any Information (collectively, with MSCI, the "Information Providers") and is provided for informational purposes only. The Information may not be reproduced or disseminated in whole or in part without prior written permission from MSCI.
- The Information may not be used to create derivative works or to verify or correct other data or information. For example (but without limitation), the Information may not be used to create indices, databases, risk models, analytics, software, or in connection with the issuing, offering, sponsoring, managing or marketing of any securities, portfolios, financial products or other investment vehicles utilizing or based on, linked to, tracking or otherwise derived from the Information or any other MSCI data, information, products or services.
- The user of the Information assumes the entire risk of any use it may make or permit to be made of the Information. NONE OF THE INFORMATION PROVIDERS MAKES ANY EXPRESS OR IMPLIED WARRANTIES OR REPRESENTATIONS WITH RESPECT TO THE INFORMATION (OR THE RESULTS TO BE OBTAINED BY THE USE THEREOF), AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, EACH INFORMATION PROVIDER EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES (INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF ORIGINALITY, ACCURACY, TIMELINESS, NON-INFRINGEMENT, COMPLETENESS, MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE) WITH RESPECT TO ANY OF THE INFORMATION.
- Without limiting any of the foregoing and to the maximum extent permitted by applicable law, in no event shall any Information Provider have any liability regarding any of the Information for any direct, indirect, special, punitive, consequential (including lost profits) or any other damages even if notified of the possibility of such damages. The foregoing shall not exclude or limit any liability that may not by applicable law be excluded or limited, including without limitation (as applicable), any liability for death or personal injury to the extent that such injury results from the negligence or wilful default of itself, its servants, agents or sub-contractors.
- Information containing any historical information, data or analysis should not be taken as an indication or guarantee of any future performance, analysis, forecast or prediction. Past performance does not guarantee future results.
- None of the Information constitutes an offer to sell (or a solicitation of an offer to buy), any security, financial product or other investment vehicle or any trading strategy.
- MSCI's indirect wholly-owned subsidiary Institutional Shareholder Services, Inc. ("ISS") is a Registered Investment Adviser under the Investment Advisers Act of 1940. Except with respect to any applicable products or services from ISS (including applicable products or services from MSCI ESG Research Information, which are provided by ISS), none of MSCI's products or services recommends, endorses, approves or otherwise expresses any opinion regarding any issuer, securities, financial products or instruments or trading strategies and none of MSCI's products or services is intended to constitute investment advice or a recommendation to make (or refrain from making) any kind of investment decision and may not be relied on as such.
- The MSCI ESG Indices use ratings and other data, analysis and information from MSCI ESG Research. MSCI ESG Research is produced by ISS or its subsidiaries. Issuers mentioned or included in any MSCI ESG Research materials may be a client of MSCI, ISS, or another MSCI subsidiary, or the parent of, or affiliated with, a client of MSCI, ISS, or another MSCI subsidiary, including ISS Corporate Services, Inc., which provides tools and services to issuers. MSCI ESG Research materials, including materials utilized in any MSCI ESG Indices or other products, have not been submitted to, nor received approval from, the United States Securities and Exchange Commission or any other regulatory body.
- Any use of or access to products, services or information of MSCI requires a license from MSCI. MSCI, Barra, RiskMetrics, ISS, CFRA, FEA, and other MSCI brands and product names are the trademarks, service marks, or registered trademarks of MSCI or its subsidiaries in the United States and other jurisdictions. The Global Industry Classification Standard (GICS) was developed by and is the exclusive property of MSCI and Standard & Poor's. "Global Industry Classification Standard (GICS)" is a service mark of MSCI and Standard & Poor's.

About MSCI

MSCI Inc. is a leading provider of investment decision support tools to investors globally, including asset managers, banks, hedge funds and pension funds. MSCI products and services include indices, portfolio risk and performance analytics, and governance tools.

The company's flagship product offerings are: the MSCI indices which include over 148,000 daily indices covering more than 70 countries; Barra portfolio risk and performance analytics covering global equity and fixed income markets; RiskMetrics market and credit risk analytics; ISS governance research and outsourced proxy voting and reporting services; FEA valuation models and risk management software for the energy and commodities markets; and CFRA forensic accounting risk research, legal/regulatory risk assessment, and due-diligence. MSCI is headquartered in New York, with research and commercial offices around the world.