Vega risk and the smile

Allan M. Malz

The RiskMetrics Group, 44 Wall Street, New York, New York 10005, USA

Vega risk is analytically easy to "nest" into the standard risk management framework, but is complicated by the prevalence of volatility smiles and term structures in most option markets. Volatility smiles, in spite of their occasionally treacherous effects on option books, are often neglected by risk managers. This paper provides a guide to incorporating vega risk into a "classical" value-at-risk (VaR) model. The paper includes a tractable approach to capturing the effects of the volatility smile and term structure on vega risk and their interaction with other risk factors. The author also presents a summary of the statistical behavior of implied volatility and relates it to observed differences in the pattern of implied volatilities from that predicted by the Black–Scholes model. The VaR computation strategy employs the smile to compensate for the shortcomings of the Black–Scholes model and provide improved VaR forecasts without a specific alternative model.

1. INTRODUCTION

Vega risk, the risk arising from fluctuations in implied volatility, can be a large part of the risk of a portfolio containing options. Vega risk is analytically easy to "nest" into the standard risk management framework. The treatment of vega risk in portfolios is, however, complicated by the prevalence of volatility smiles and term structures in most option markets. These option market phenomena run counter to the benchmark Black–Scholes option pricing model, in which implied and historical volatility are constant over time and invariant with respect to option maturity and likelihood of exercise. Researchers have developed models incorporating stochastic volatility, jump-diffusions, and other alternatives to the pure constant-volatility diffusion environment of the Black–Scholes model. Some option dealers and risk managers use these models to refine their option prices and return forecasts.

However, volatility smiles, in spite of their occasionally treacherous effects on option books, are often neglected by risk managers. A common metric for market risk for a portfolio of financial assets is value-at-risk (VaR), a measure of the maximum loss to the portfolio that can occur over a specified time horizon with a specified confidence level. This paper provides a guide to incorporating vega risk into a "classical" VaR model. The paper suggests a practical approach to capturing the effects of the volatility smile and term structure on vega risk and their interaction with other risk factors. In our discussion, we will present several examples using a high-quality database of foreign exchange implied volatilities.[1]

[1] The data, sourced from JP Morgan's foreign exchange desk and available from RiskMetrics, includes implied volatilities for a wide range of currency pairs and maturities, extensive coverage of volatility smiles, and term structures.

While our approach does not incorporate an alternative option pricing model to Black–Scholes, it does take account of the observed departures of real-world option pricing from the standard model by incorporating the volatility smile and term structure. If option markets are reasonably efficient, the smile and term structure will reflect the same deviations of volatility behavior from the random walk that the alternative models attempt to capture. Using market prices in this way avoids committing the risk manager to a particular forecasting model and exploits instead the readily available information embedded in market prices.

Section 2 describes how to compute a VaR for a portfolio containing options that takes vega risk into account in the simplest possible way, disregarding the discrepancies between the Black–Scholes model and real-world option prices. This captures exposure to the general level of implied volatility and the correlations of implied volatility with returns on the underlying market factors. The following sections describe the discrepancies between the model's predictions and actual option price behavior, and propose an adaptation of the standard VaR model that takes these discrepancies into account without requiring specification of an alternative model.

2. VEGA RISK IN THE BLACK–SCHOLES MODEL

Option positions are exposed to a range of market risks. Delta and gamma risk are the exposures of an option position to changes in the prices of the underlying assets. Vega, denoted κ, is the exposure of an option position to changes in the implied volatility of the option. The change in the option value is defined as a partial derivative, i.e., it assumes all other factors determining the option value are held constant:

$$\kappa = \frac{\partial(\text{option value})}{\partial(\text{implied volatility})}.$$

Vega is measured in dollars or other base currency units, implied volatility in percent per annum.

An implied volatility is linked to a particular option valuation model. The most common is the Black–Scholes model. Because implied volatility is defined only in the context of a particular model, an option pricing model is required to measure vega. Vega is then defined as the partial derivative of the call or put pricing formula with respect to the implied volatility. The Black–Scholes model vega of a call or put option on a dividend-paying asset is

$$\kappa = \sqrt{\tau}\, S_t \mathrm{e}^{-d\tau} \phi\left(\frac{\ln(S_t/X) + (r - d + \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}\right), \tag{1}$$

where σ is the implied volatility, S_t the current underlying price, X the exercise price, τ the remaining time to maturity, r the τ-year financing or money market

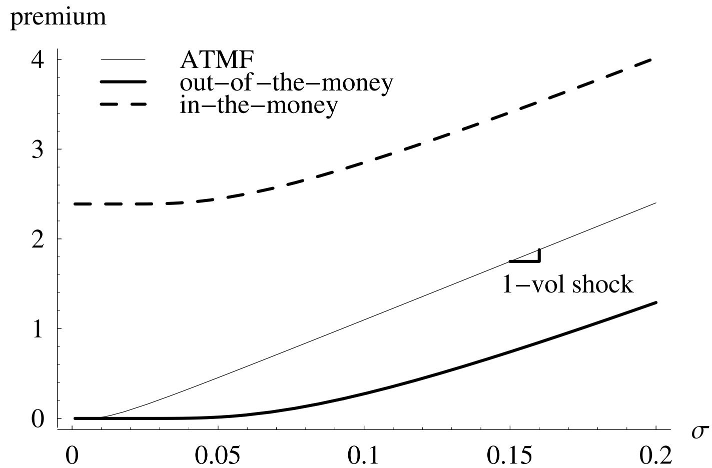

FIGURE 1. Option value as a function of implied volatility. One-month dollar–yen calls, premium in JPY. Out-of-the-money exercise price 2.5% above forward rate; in-the-money exercise price 2.5% below forward rate.

rate, d the dividend, interest, or other cashflow the underlying asset is expected to return over the term of the option, and ϕ the standard normal density.

Some properties of vega in the Black–Scholes model are illustrated in Figures 1 and 2:

- The value of a put or call increases monotonically with the implied volatility. The vega of a long option position is therefore always positive.
- Vega is at a maximum for an at-the-money option. The value of an option is therefore almost linear in the implied volatility for an at-the-money option.
- The vega of a longer-term option is greater than that of a shorter-term option.

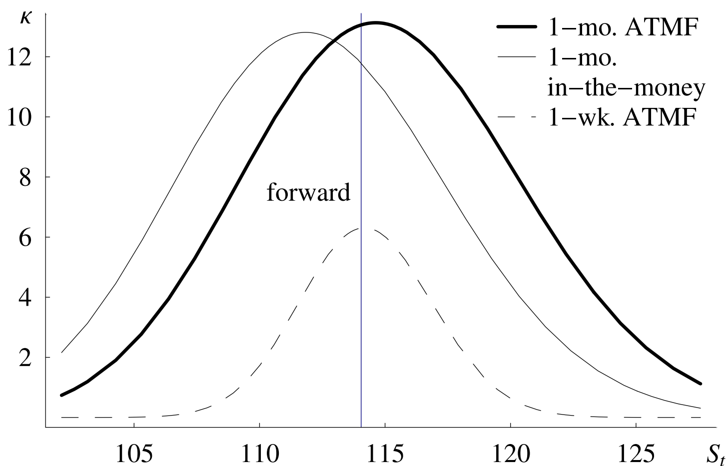

FIGURE 2. Option vega. Dollar–yen calls, premium in JPY, implied volatility 16.7%. In-the-money exercise price 2.5% below forward rate.

The vega exposure of an option position is defined as the change in position value if the implied volatility rises and is calculated as the option vega times the underlying amount of the option (the "number of options").

The simplest approach to computing a VaR for a portfolio containing options is to treat implied volatility analogously to other market risk factors, such as equity, foreign exchange, and interest rates of different maturities. Logarithmic changes—"returns"—in implied volatility are assumed to be jointly normally distributed with the remaining risk factors in the portfolio with mean zero and a constant variance-covariance matrix.[2] An analogous assumption of lognormality is frequently made for interest rates in risk models and avoids assigning a positive probability to negative implied volatilities. Under the model assumptions, the logarithm of tomorrow's implied volatility is normally distributed with a mean of zero and a standard deviation $\nu_{\mathrm{vol}}\sqrt{\tau}$, where ν_{vol} is the standard deviation of the implied volatility, or "vol-of-vol", with τ the VaR horizon in days, and the 90% confidence interval for next-day implied volatility is $1.65\sigma e^{\pm\nu_{\mathrm{vol}}}$.[3]

If the only market risk of an option were vega, the VaR of a portfolio containing a single option would be

$$\text{underlying amount} \cdot \kappa \cdot 1.65\sigma(e^{\nu_{\mathrm{vol}}} - 1) \approx \text{underlying amount} \cdot \kappa \cdot 1.65\sigma \cdot \nu_{\mathrm{vol}}.$$

An option position is exposed to delta and gamma risks as well as vega, and these must also be accounted for in VaR. In a simple parametric VaR framework incorporating exposure to both the underlying price and to implied volatility, the VaR of a single option with a 1-day horizon and a 95% confidence level is

$$1.65\left|\text{underlying amount}\right|\sqrt{\begin{bmatrix}\delta S\nu_{\mathrm{spot}} & \kappa\sigma\nu_{\mathrm{vol}}\end{bmatrix}\begin{bmatrix}1 & \rho_{\mathrm{spot,vol}}\\ \rho_{\mathrm{spot,vol}} & 1\end{bmatrix}\begin{bmatrix}\delta S\nu_{\mathrm{spot}}\\ \kappa\sigma\nu_{\mathrm{vol}}\end{bmatrix}},$$

where δ denotes the option delta and ν_{spot} the historical volatility of the underlying asset.[4]

Example 2.1 Consider a short position in a 1-month at-the-money forward (ATMF) dollar–yen put, i.e, a put on the dollar, with premium and exercise price denominated in yen and with an exercise price equal to the current forward price of the asset. The underlying amount is \$1 000 000, the current spot rate is

[2] The standard approach to VaR computation is described by Morgan Guaranty Trust Company (1996).

[3] More precisely, the expected value of the logarithm of the next-day implied volatility is $\sigma - \frac{1}{2}\nu_{\mathrm{vol}}^2$, but we ignore the negligible second term. Note that at the risk of some confusion, we express implied volatility at an annual rate and the vol-of-vol at a daily rate.

[4] In principle, exposure to volatility is a nonlinear risk, making parametric VaR inaccurate, though, for at-the-money options, vega risk is almost exactly linear, as seen in Figure 1. Even for at-the-money options, parametric VaR is less accurate than Monte Carlo VaR, as the exposure to the underlying is highly nonlinear. See Mina and Ulmer (1999) for more detail.

¥120, the financing rate, i.e., the yen-denominated money-market rate, is 50 basis points, the US money-market rate is 5%, and the implied volatility is 15%. The initial option vega is $\kappa = 13.76$. The vega of the short option position is $-1\,000\,000\kappa = -¥13\,759\,100$, or $-\$114\,659$.

The standard deviation of 1-day logarithmic changes in 1-month implied volatility is $\nu_{\text{vol}} = 0.0567$ (5.67%). The VaR arising from vega risk, disregarding other market risks, is $1.65 \cdot \frac{1\,000\,000}{120} \cdot 13.76 \cdot 0.15 \cdot 0.0567 = \$1609$. The dollar–yen put option delta is $-0.489\,32$ and the delta equivalent is a long dollar position of $1\,000\,000 \cdot 0.48932 = \$489\,320$. The daily spot rate volatility is $\nu_{\text{spot}} = 0.0097$. The VaR of the delta equivalent, ignoring vega as well as the gamma and higher-order nonlinear exposures to the underlying spot rate, is $1.65 \cdot 1\,000\,000 \cdot 0.489\,32 \cdot 0.0097 = \$7794$.[5] The correlation of dollar–yen returns with logarithmic changes in the 1-month implied volatility is -0.395. The tendency for a weakening dollar to coincide with a rise in option prices might be due to an increased desire of Japanese exporters to hedge their foreign exchange earnings if a weaker dollar appears likelier.

The parametric VaR is

$$1.65 \cdot 1\,000\,000 \cdot \sqrt{\begin{bmatrix} 0.489 \cdot 0.0097 \\ -\frac{1}{120} \cdot 13.7591 \cdot 0.15 \cdot 0.0567 \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} 1 & -0.395 \\ -0.395 & 1 \end{bmatrix} \begin{bmatrix} 0.489 \cdot 0.0097 \\ -\frac{1}{120} \cdot 13.759 \cdot 0.15 \cdot 0.0567 \end{bmatrix}},$$

or \$8558. The correlation benefit reduces VaR by $7794 + 1609 - 8558 = \$845$.

In this example, delta risk—exposure to the price of the underlying asset—predominates. In some option portfolios, vega risk can be much greater than delta risk, e.g., in a delta-hedged option, a straddle (a call and put with the same exercise price), and a strangle (a call and a put with different exercise prices).

Example 2.2 We now add a delta hedge to the short dollar–yen put of Example 2.1, so the portfolio consists of a short put position together with a short dollar position of \$489 320. This portfolio will generally have much smaller fluctuations in value than the naked option, but is exposed to gamma and vega risk.

To illustrate, consider the Monte Carlo results for the VaR of the hedged option portfolio. We generate n shocks $(u_{i,\text{spot}}, u_{i,\text{vol}})$ from a bivariate normal distribution with a mean of zero and a variance-covariance matrix based on ν_{spot}, ν_{vol}, and $\rho_{\text{spot,vol}}$. The portfolio is revalued using the Black–Scholes model.[6]

[5] Historical volatilities and correlations are those prevailing on 8 February 1999.

[6] We have ignored the aging of the option and the forward points paid to carry the long dollar position overnight. Even with a large interest rate differential such as 5% favoring the yen, the points add about 2 basis points to the daily cost of holding yen against dollars, negligible compared with the standard deviation of 1%. Note also that the VaR on the delta hedge is not centered exactly at zero. Rather, because we convert to dollars by dividing by the exchange rate, there is a slight downward bias: if the average shock u_i is zero, the average of the e^{-u_i} is slightly less than zero.

Table 1. Returns on dollar–yen option portfolios.

	0.05 quantile	0.95 quantile	Mean	Median
Examples 2.1 and 2.2. Hedged short dollar–yen put option portfolio: naive approach				
Unhedged option, return to spot fluctuations	−9158	6837	−386	72
Delta hedge	−7698	7749	−38	−69
Hedged option, return to spot fluctuations	−1659	14	−424	−186
Option, return to vol fluctuations	−1689	1528	−35	9
Unhedged option, return to spot and vol fluctuations	−9880	7599	−424	54
Hedged option, return to spot and vol fluctuations	−2589	1240	−462	−338
Example 4.1. Hedged short dollar–yen put option portfolio: fixed smile technique				
Hedged option with dollar-bearish smile	−2474	1078	−543	−460
Hedged option with dollar-bullish smile	−3062	1451	−468	−288
Example 4.2. Hedged dollar-bearish risk reversal: fixed smile technique				
Risk reversal with dollar-bearish smile	−712	406	−128	−124
Risk reversal, ignoring smile	−216	551	84	24
Example 4.3. Hedged short dollar–yen put option portfolio: random smile technique				
Hedged option with dollar-bearish smile	−2384	1082	−523	−447

The data in the table are 1-day changes in portfolio value. The underlying amount is \$1 000 000 and all options have a tenor of 1 month. The initial data are: spot rate ¥120, USD risk-free rate 5.5%, JPY risk-free rate 0.5%, implied volatility 15%. The risk reversal portfolio contains a long 25-delta yen call and a short 25-delta yen put. The 0.05 quantile $\equiv$ VaR $\times$ (-1) at a 95% confidence level.

Table 1 shows results of the Monte Carlo simulation. Incorporating vega risk gives a more accurate picture of the risks of the hedged position, raising the VaR from \$1659 to \$2589.

3. LIMITATIONS OF THE BLACK–SCHOLES MODEL

3.1 *Anomalies in the Statistical Behavior of Implied Volatility*

The Black–Scholes model assumes that asset returns follow a random walk with a constant volatility. The Black–Scholes implied volatility is the market estimate of that volatility. Since volatility is a constant, vega risk does not exist in the Black–Scholes model. In real-life markets, implied volatility fluctuates widely. There is therefore a contradiction, overlooked in the previous section, in using the Black–Scholes implied volatility and the Black–Scholes vega to measure vega risk. The Black–Scholes model is only a useful first approximation to the "true" model governing volatility and the Black–Scholes implied volatility is not necessarily a correct measure of anticipated volatility. It is rather a marketadjusted parameter in the Black–Scholes option pricing formulas which is related to, but not identical to, anticipated volatility.

To illustrate variations in implied volatility, we consider 11 assets, described in Table 2. Table 3 displays descriptive statistics on implied volatility, computed for daily and weekly observations over long sample intervals. The range of variation of implied volatility is wide. For example, annualized 1-month ATMF dollar–yen vols have ranged from 6.2% to 40%. As Examples 2.1 and 2.2 make clear, vega risk can add considerably to the VaR of a portfolio containing options. One reason for this is the typically high daily standard deviation of log changes in implied volatility, or vol-of-vol. For dollar–yen, it has averaged about 5.3%, or 80% per annum.

The assumption made in the previous section, that log implied volatility returns follow a normal distribution, is also only a crude approximation. In fact, implied volatility returns are highly kurtotic and generally skewed, particularly in daily observations. Implied volatility is strongly mean-reverting, as shown in

TABLE 2. Reference information for implied volatility data.

Asset	Description	Sample period	Sample size Weekly	Daily
CME S&P500	S&P500 index, 3-month futures	28Jan83–08Oct99	4063	863
CME Eurodollar	USD 3-month Libor, 3-month futures	21Mar85–08Oct99	3633	758
Liffe short sterling	GBP 3-month Libor, 3-month futures	03Apr90–12Dec99	2406	503
CBOT 30-year bond	US Treasury bond, $4\frac{1}{2}$-month futures	24Feb86–12Dec99	3466	721
Nymex crude light oil	Crude oil, 1-month futures	21Mar85–08Oct99	1781	374
Gold	Spot New York gold, USD per ounce	05Jan85–04Jan00	1267	259
USD–JPY	Spot dollar–yen exchange rate	31Mar92–31Aug99	1897	387
USD–EUR	Spot dollar–mark exchange rate	31Mar92–13Sep99	1907	389
GBP–EUR	Spot sterling–mark exchange rate	31May90–14Dec99	2455	500
USD–MXP	Spot dollar–Mexican peso exchange rate	02Feb97–14Dec99	724	148
USD–THB	Spot dollar–Thai baht exchange rate	20Jan95–31Aug99	1176	240

At-the-money volatilities for futures options, computed from puts and calls with exercise prices adjacent to the futures settlement price. Constant maturity implied volatilities by linear interpolation of major-contract implied volatilities with adjacent maturities. One-month at-the-money forward volatilities for currencies and gold.

TABLE 3. Full sample descriptive statistics on implied volatilities.

	Implied volatility levels			Implied volatility logarithmic returns			
	Mean	Min.	Max.	Mean	SD	Skewness	Kurtosis
A. Daily observations							
CME S&P500	16.9	9.6	81.1	−0.00	4.38	4.42*	106.20*
CME Eurodollar	15.5	5.4	40.6	−0.02	4.71	0.42*	7.09*
Liffe short sterling	13.2	5.2	36.0	0.01	5.06	−0.24*	8.88*
CBOT 30-year bond	10.4	6.9	23.5	−0.01	2.70	−0.57*	17.69*
Nymex crude light oil	29.0	14.0	59.6	0.05	4.02	0.52*	11.41*
Gold	11.4	4.0	53.5	0.08	6.94	1.93*	15.47*
USD–JPY	11.8	6.2	40.0	0.02	5.27	1.01*	5.36*
USD–EUR	10.8	6.2	24.0	−0.01	4.14	1.85*	14.95*
GBP–EUR	7.5	2.5	18.0	0.01	5.66	0.26*	51.05*
USD–MXP	12.2	5.7	43.0	0.03	9.89	6.54*	99.09*
USD–THB	12.1	1.4	60.0	0.02	6.62	2.78*	31.07*
B. Weekly observations							
CME S&P500	16.8	9.7	58.0	0.02	7.84	2.14*	21.37*
CME Eurodollar	15.5	5.5	39.7	−0.07	9.86	0.70*	5.78*
Liffe short sterling	13.2	5.3	36.0	0.06	10.70	0.04	6.20*
CBOT 30-year bond	10.4	7.0	20.2	−0.04	5.54	0.33*	7.56*
Nymex crude light oil	29.0	14.3	56.5	0.24	8.37	0.72*	3.05*
Gold	11.4	4.0	53.0	0.33	16.00	1.57*	10.31*
USD–JPY	11.8	6.4	28.5	0.10	10.90	1.13*	4.59*
USD–EUR	10.8	6.5	21.5	−0.03	8.41	1.28*	5.06*
GBP–EUR	7.5	2.6	18.0	0.06	11.40	0.53*	9.03*
USD–MXP	11.9	5.7	32.0	0.15	23.00	3.62*	22.53*
USD–THB	12.1	1.4	60.0	0.16	17.10	2.34*	14.12*

Implied volatility levels: percent per annum. Implied volatility returns: daily logarithmic changes, in percent. Mean and SD: mean and standard deviation of daily vol returns in percent. Kurtosis: coefficient of kurtosis excess. Skewness: skewness coefficient. Asterisk denotes test statistics significantly different from zero (test of normal distribution) at 99% confidence level.

Tables 4 and 5. Table 4 displays autocorrelation coefficients for the 11 assets studied. In most cases, negative and statistically significant coefficients are predominant among the daily autocorrelations, and the first-order weekly autocorrelation coefficient is negative and statistically significant.

Table 5 displays variance ratio tests. Under the hypothesis that implied volatility follows a random walk, the vol-of-vol over a longer time interval is equal to the vol-of-vol over a shorter period multiplied by the square root of the ratios of the time intervals. The vol-of-vol over a trading week, say, is expected to be about $\sqrt{5}$ times the daily vol-of-vol. This property can be tested using variance ratios, which measure the ratio between long- and short-period return variances, normalized by the ratios of the time intervals. If implied volatility follows a random walk, its variance ratios are expected to be close to 1. If instead implied volatility is mean-reverting, the per-period volatility of longer-term implied volatility returns will be smaller than the volatility of one-period returns and its variance ratios will be smaller than 1. The variance ratios of implied volatilities provide additional evidence of mean reversion. In almost all cases, the variance ratios are less than unity, and decline for longer horizons.

One notable exception to this pattern is the USD–THB series. Statistically and economically significant positive autocorrelations and variance ratios in excess of

TABLE 4. Autocorrelation coefficients of implied volatility returns.

	Autocorrelation coefficients				Box–Pierce statistics	
	1	2	3	4	$\hat{Q}_5$	$\hat{Q}_{10}$
A. Daily observations						
CME S&P500	−8.2**	−6.8**	−2.2	0.6	48.5**	62.1**
CME Eurodollar	0.8	−8.5**	−2.1	2.2	31.2**	48.0**
Liffe short sterling	−3.1	−1.0	−1.4	−3.8*	10.1	27.7**
CBOT 30-year bond	−2.6	−5.0**	−0.9	−1.8	13.0**	20.9**
Nymex crude light oil	−12.7**	2.5	−2.2	−0.1	30.9**	38.8**
Gold	10.6**	−5.8**	−9.1**	−5.1*	34.2**	45.2**
USD–JPY	−0.8	−6.5**	−8.4**	−2.2	22.6**	29.5**
USD–EUR	−4.1*	−0.8	−9.6**	−2.6	23.4**	37.8**
GBP–EUR	−15.0**	−0.5	3.0	−5.2**	63.8**	68.5**
USD–MXP	3.6	−2.5	−1.1	−0.6	1.5	24.1**
USD–THB	11.1**	10.4**	8.6**	2.2	37.2**	40.5**
B. Weekly observations						
CME S&P500	−5.9*	−9.7**	−3.2	−3.1	23.8**	35.4**
CME Eurodollar	−9.7**	−5.0	−1.6	−4.7	10.9	13.2
Liffe short sterling	−5.5	−12.6**	8.0*	−4.5	14.2**	16.2
CBOT 30-year bond	−9.1**	−2.6	−1.7	−2.8	9.1	11.8
Nymex crude light oil	−8.4	−6.0	−4.5	1.5	5.6	9.2
Gold	−11.6*	−5.7	−9.0	6.8	8.9	12.9
USD–JPY	−14.0**	−2.6	−9.6*	−0.5	11.7*	15.3
USD–EUR	−15.2**	2.8	−3.7	−6.4	17.0**	19.9*
GBP–EUR	−14.9**	−2.2	9.9**	−5.1	20.4**	30.1**
USD–MXP	−19.5**	9.6	−6.2	−8.0	9.8	14.1
USD–THB	12.1*	4.1	−5.8	−2.3	6.7	16.1

Autocorrelation coefficients in percent. Asterisk (double asterisk) denotes test statistics significantly different from zero at 95% (99%) confidence level (test against normal distribution).

Box–Pierce statistics in percent. Asterisk (double asterisk) denotes test statistics significantly greater than zero at 95% (99%) confidence level (test against chi-square distribution).

TABLE 5. Variance ratios of implied volatility returns.

	$q = 2$	$q = 4$	$q = 8$	$q = 16$
A. Daily observations				
CME S&P500	0.92	0.80	0.72	0.61
CME Eurodollar	1.01	0.92	0.84	0.76
Liffe short sterling	0.97	0.94	0.89	0.76
CBOT 30-year bond	0.98	0.91	0.81	0.75
Nymex crude light oil	0.87	0.83	0.78	0.70
Gold	1.11*	1.06	0.95	0.82
USD–JPY	0.99	0.88	0.76	0.64
USD–EUR	0.96	0.89	0.77	0.68
GBP–EUR	0.85	0.79	0.71	0.65
USD–MXP	1.04	1.03	0.95	0.95
USD–THB	1.11	1.32*	1.51*	1.64*
B. Weekly observations				
CME S&P500	0.94	0.80	0.62	0.50
CME Eurodollar	0.91	0.80	0.68	0.57
Liffe short sterling	0.95	0.84	0.76	0.69
CBOT 30-year bond	0.91	0.83	0.69	0.49
Nymex crude light oil	0.90	0.78	0.68	0.55
Gold	0.89	0.73	0.63	0.42
USD–JPY	0.86	0.72	0.55	0.43
USD–EUR	0.85	0.79	0.62	0.52
GBP–EUR	0.85	0.80	0.76	0.68
USD–MXP	0.82	0.82	0.62	0.38
USD–THB	1.13	1.22	1.28	1.02

Asterisk denotes test statistics significantly greater than zero at 95% confidence level (test against chi-square distribution). The variance ratios and test statistics are $\overline{\mathrm{VR}}(q)$ and $\psi^*(q)$ as described by Campbell, Lo, MacKinlay (1997, pp. 48ff.). The statistics are computed assuming a mean implied volatility return equal to zero.

unity are evident even at lags of many weeks. This can be explained by the abandonment of the fixed exchange rate system on 2 July 1997: in the course of a currency crisis that results in the breaking of a peg, the currency often depreciates and implied volatilities rise over a protracted period of weeks or even months.

3.2 *Anomalies in the Cross-Section Pattern of Implied Volatility*

Apart from the time-variation and statistical behavior of implied volatility, the cross-sectional pattern of observed implied volatilities also differs from that predicted by the Black–Scholes model. In real-life markets, the implied volatility for a given asset at a given time is not an identical constant for all options on a given asset observed at a given time, but differs for options with different exercise prices and different times to maturity. These patterns of implied volatility are

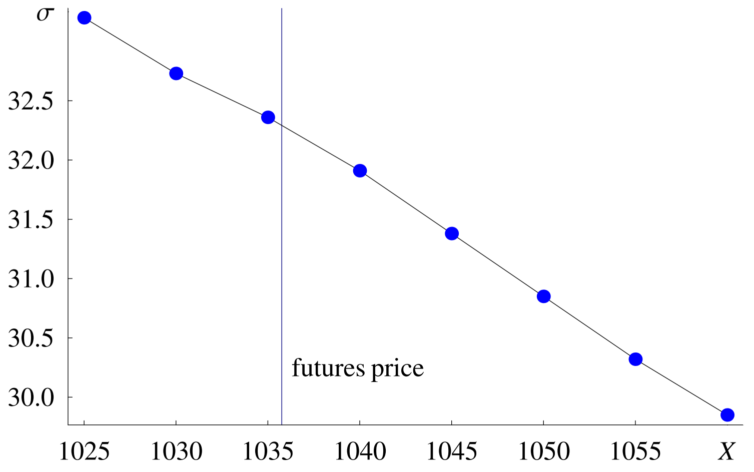

FIGURE 3. The volatility smile in the S&P index market. Implied volatilities on 11 September 1998 of December 1998 options on December 1998 futures.

known as the volatility smile and the term structure of volatility and have been observed for options on most assets. The volatility surface, the plot of the implied volatilities of a particular asset as a function of exercise price and time to maturity, combines these phenomena:[7]

Smile The volatility smile describes the characteristic shape of the plot of implied volatilities of options of a given time to expiry against the delta or against the exercise price: out-of-the-money options often have higher implied volatilities than at-the-money options. Typically, the curvature is less for longer-term options than for short-term options.

Smirk and sneer The volatility smile is often skewed, so that out-of-the-money call options have implied volatilities which differ from those of equally out-of-the-money put. The lopsided pattern of implied volatility across exercise prices is often called the "smirk" or "sneer" by market participants, as in Figure 3.

Term structure The term structure of implied volatility describes the pattern of options with the same exercise price but different maturities, which generally have different implied volatilities.

Portfolios containing options are exposed not only to changes in the level of implied volatility but also to:

- changes in the curvature and skewness of the smile;
- changes in implied volatility *along* the smile;
- changes in the slope of the term structure of volatility;

[7] On the volatility smile, see, among many other examples, Heynen (1994) and Peña, Rubio, and Serna (1999) on equity indexes, Melick and Thomas (1997) for crude oil futures, and Xu and Taylor (1994) and Malz (1996) for currencies. On the term structure of volatility, see, e.g., Campa and Chang (1995) on currencies and Stein (1989) on equity indexes.

- changes along the term structure of volatility as the option "ages".

In Section 4, we will examine how these anomalies affect VaR estimation, focusing on foreign exchange markets. Before doing so, it may be useful to introduce some institutional features and quotation conventions of the foreign exchange option market.[8] In foreign exchange option markets, dealers use the Black–Scholes option delta as a metric for exercise price and the Black–Scholes implied volatility as a metric for price. The delta and implied volatility are related to current market data via the Black–Scholes pricing functions. The volatility smile can then be represented as a schedule of implied volatilities of options on the same underlying and with the same maturity but different deltas. This convention is unambiguous, since a unique exercise price corresponds to each call or put delta and a unique option premium in currency units corresponds to each implied volatility. The most liquid option markets are for ATMF (50-delta options) and for 25-delta calls and puts. There are also actively traded markets in 10-delta calls and puts.

The liquidity of out-of-the-money options in foreign exchange markets has led to the widespread trading of option combinations, in particular the straddle, the strangle, and the risk reversal. The straddle combines a call and a put with the same exercise price, usually the current forward outright rate. The strangle and the risk reversal combine an out-of-the-money call and an out-of-the-money put with the same delta, usually 25% (less often 35, 30, or 10%), and the same maturity. The exercise price of the call component is higher than the current forward outright rate, and the exercise price of the put is lower by approximately the same proportional amount. The strangle consists of a long out-of-the-money put and call. The risk reversal consists of a long out-of-the-money call and a short out-of-the-money put.

The prices of these option combinations are expressed in vols rather than currency units in over-the-counter market parlance. The implied volatility of the put component of an at-the-money forward straddle is identical to that of the call (by virtue of put–call parity) and is referred to as the straddle or at-the-money forward volatility and denoted here by atm_t. The straddle price, denoted str_t, is generally quoted as the spread between the average of the out-of-the-money implied volatilities and the ATMF volatility. The risk reversal price, denoted rr_t, is generally quoted as the spread between the out-of-the-money base currency call and put volatilities. It can be positive or negative.

In Figure 4, a quadratic polynomial is fitted to the observed ATMF and 25- and 75-delta implied volatilities. In the quadratic case, the smile has an intuitive parametrization in terms of the quoted option prices:

$$\sigma(\delta) = \mathrm{atm}_t - 2\mathrm{rr}_t(\delta - 0.50) + 16\mathrm{str}_t(\delta - 0.50)^2. \tag{2}$$

This closed form for the smile provides intuition for the currency option price conventions: the strangle volatility is a "measure of location" of the smile, the

[8] See Malz (1997) for detail on these conventions.

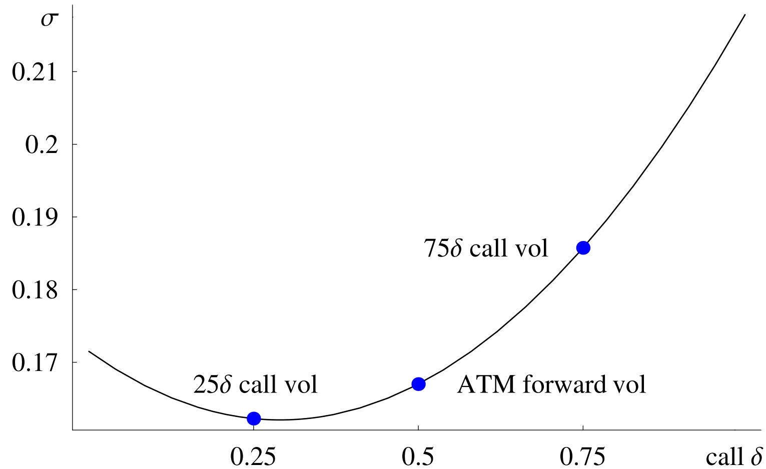

FIGURE 4. The volatility smile in the foreign exchange market. One-month dollar–yen options, 8 February 1999. Source: DataMetrics.

strangle price indicates the degree of curvature, and the risk reversal price indicates the degree of skewness of the smile.

The smile displayed in Figure 3 is skewed toward lower values of the underlying asset. This is commonly called a "put skew," since out-of-the-money puts have higher vols than equally out-of-the-money calls. This is something of a misnomer, since the smiles of puts and calls are identical when graphed against exercise price. If the put and call smiles differed, it would violate put–call parity and open up arbitrage opportunities.

3.3 *Alternative Option Pricing Models to Black–Scholes*

The cross-sectional and statistical anomalies in option prices have given rise to a large body of research developing alternative models to Black–Scholes. The alternative models seek consistency with the anomalies and with the statistical behavior of the underlying asset prices and implied volatilities.

The most widely accepted of these modeling approaches represents the underlying asset volatility as stochastic rather than constant. In this approach, an additional stochastic process is introduced to describe volatility behavior.[9] Recent work has focused on the autoregressive conditional heteroscedasticity (ARCH) family of models of stochastic volatility, with which the researcher can attempt to capture both the volatility smile and the observed dynamics of asset prices and volatility.[10]

[9] An example of this approach is given by Hull and White (1987). Several papers, including Stein and Stein (1991) and Heston (1993), propose a mean-reverting specification for the logarithm of volatility, which is quite consistent with the empirical observations presented here.

[10] Duan (1999), for example, presents a generalized ARCH (GARCH) model that captures additional features of the observed statistical properties of underlying asset volatility, in particular the fact that returns appear to follow a fat-tailed distribution even after volatility innovations are accounted for.

Another approach, less widely applied, is to model the asset price as following a jump-diffusion process. As in the stochastic volatility class of models, an additional stochastic process is introduced, here describing the size and timing of discrete changes in the underlying asset price. This approach is intuitively appealing, but regarded as somewhat less successful than the stochastic volatility family of models in replicating the statistical behavior of asset volatility.[11]

The stochastic volatility and jump-diffusion approaches have several drawbacks. First, they introduce additional sources of risk, jump risk and volatility risk, that cannot be perfectly hedged in the standard continuous-trading model. As a result, the jump and volatility risks must be assumed to be nonsystematic, or their returns must be known, in order to derive the option price by no-arbitrage arguments. Otherwise, the option price must be derived in the context of a particular equilibrium model of the economy in which tastes and technology are specified, or assumptions such as diversifiability must be made that set the price of these additional risk factors to zero. Second, the parameters of the resulting option models must generally be estimated statistically rather than being implied by the set of observed option prices. The models do not have satisfactory forecasting power for asset and option returns, so, to increase accuracy, practitioners update the parameters of the models frequently in order to increase their accuracy. This is problematic in precisely the same way that time-varying volatility is for the Black–Scholes model.

A third approach focuses on replicating as precisely as possible the cross-sectional anomalies in option prices. It represents volatility as nonconstant but deterministic, varying in accordance with a local volatility function with the future date and the then-prevailing level of the underlying asset. In the deterministic volatility approach, the volatility surface for European options is used to estimate the local volatility function.[12]

The deterministic volatility model has the advantage of preserving the no-arbitrage derivation of option prices, even for exotic options, without being tied to the Black–Scholes hypothesis of a constant volatility. As a result, the stochastic process of the underlying asset is implicit in the volatility function and can also be retrieved.[13] However, it appears that the local volatility function is not stable. As a result, the deterministic volatility approach forecasts future option prices less well than the internally inconsistent but simpler approach of using the volatility surface to forecast future implied volatilities and thus option prices.[14] This approach is adopted in the next section to incorporate the statistical and cross-sectional anomalies in implied volatility into VaR estimation.

[11] See, e.g., Bates (1991). Bates (1996) combines the jump-diffusion and stochastic volatility approaches to explain currency option behavior. Bakshi, Cao, and Chen (1997) provide an empirical assessment of a range of option models.

[12] Examples of this approach include those of Derman and Kani (1994), Dupire (1997), and Demeterfi *et al.* (1999).

[13] See Rubinstein (1994).

[14] See the extensive empirical study by Dumas, Fleming, and Whaley (1998).

4. VEGA RISK IN THE PRESENCE OF A VOLATILITY SMILE AND TERM STRUCTURE

In obtaining more accurate VaR estimates than with the "naive" Black–Scholes approach of Section 2, one can apply an alternative model of the joint behavior of asset price returns and volatility, such as stochastically or deterministically variable volatility, or jump-diffusion. Several drawbacks of these approaches were noted in the previous section. None of the models current in the literature has been able to accurately replicate and forecast the joint behavior of the entire volatility surface and the spot market over time.

Rather, we draw a forecast of the joint behavior of the underlying price and implied volatility directly from the volatility surface. The rationale for this approach, like that of the deterministic volatility approach, is that the volatility surface contains an implicit risk-neutral estimate of the distribution of future asset returns and volatility. Both approaches take into account the phenomena that stochastic volatility and jump-diffusion models attempt to capture, as well as other phenomena such as shifting correlations between the level of the underlying asset price and the shape and level of the volatility surface, insofar as these phenomena have become embedded in the market-assessed constellation of implied volatilities.

The main advantage of the approach here is that it does not rely on a particular model of the process but draws on the more complicated model expressed in the market. However, by relying on current implied volatilities rather than historical prices, it can be influenced by shifts in risk premiums as well as changes in implicit market forecasts. These shifts in risk premiums cannot, in the current state of the art, be accurately distinguished from changes in expectations.[15] We build this approach in steps, first assuming the shape of the volatility smile to be constant over the forecast horizon, and then permitting the shape to vary over time.

4.1 *Fixed Volatility Smile*

In the simpler approach to incorporating the smile, we ignore changes in the shape of the volatility smile and take into account only parallel shifts in the volatility smile and changes in implied volatility along the volatility smile as the spot rate changes. The only new wrinkle here is that implied volatility scenarios are determined not only by the normally distributed shocks $u_{i,\mathrm{vol}}$ to the general level of implied volatility, but by a displacement along the volatility smile. The displacement along the smile is determined by the shock to the cash price, as shown in Figure 5.

In the fixed-smile approach, we assume the shock to implied volatility shifts the entire volatility smile up or down, but does not change the *shape* of the

[15] Recently, Aït-Sahalia and Lo (2000) has attempted to distinguish empirically between the risk premiums and expectations embedded in risk-neutral probability distributions.

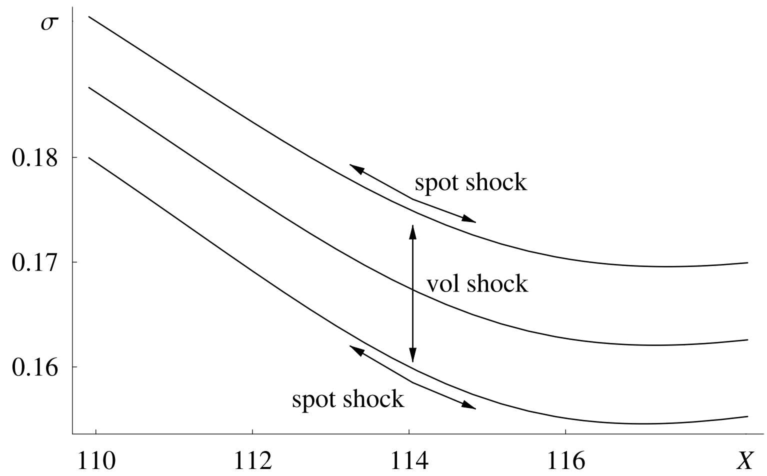

FIGURE 5. Perturbation of spot and vol under the fixed smile approach.
Source: DataMetrics.

smile. The smile shifts in exercise price–vol space, but not in delta–vol space.[16] The fixed smile approach partially compensates for one of the shortcomings of conventional VaR, the fact that the assumption of normal returns is only an approximate description of the actual behavior of cash prices and implied volatilities. The volatility smile is one manifestation of nonnormal returns, and the fixed smile approach lets us take it into account using standard VaR computations.

Example 4.1 We return to Example 2.2 of a delta-hedged short at-the-money dollar–yen option and continue to assume that the at-the-money volatility is 15%. In addition, we assume that the risk reversal and strangle are −2.5% and 0.5 vols. We use Monte Carlo computation, as in Example 2.2 above.[17]

A positive shock to the underlying shifts the smile to the right, so that the at-the-money implied volatility now applies to a higher exercise price. An existing put (call) option which yesterday was at-the-money is now slightly out-of-the-money (in-the-money). It is revalued at the implied volatility corresponding to its new, lower (higher) delta. For the dollar-bearish volatility smile of our example, this means that a positive shock to the exchange rate (dollar up) leads to a rise in the implied volatility of an at-the-money put and a decline in the implied volatility of an at-the-money call along the volatility smile.

In the Monte Carlo computation of VaR, we combine correlated shocks to spot and implied volatility. To understand the net effect of the smile on VaR, we

[16] Derman (1999) labels this approach the "sticky smile" approach to implied volatility.

[17] These data are close to those observed on 8 February 1999 and displayed in Figure 4, but are rounded off in the example for ease of exposition. The volatility skew in the dollar–yen market has typically, over the decade or so for which we have data, favored the yen, i.e., the average risk reversal is negative. A risk reversal of −0.025 would be considered high, but not unusually so. Similarly, a curvature of 0.005 is high, but not extreme, for dollar–yen.

need to analyze the effect on portfolio value of the four possible combinations of positive and negative shocks to these two factors:

Spot up, vol up This is the worst combination for the hedged short at-the-money put. The rise in the dollar creates a "bad gamma" trading loss (as will be the case for a drop in spot). The dollar appreciation also drives the volatility of the put higher along the smile, as it becomes, say, a 40- rather than 50-delta option. Thus volatility rises both as a result of the shock to volatility and as a result of the change in moneyness.

Spot up, vol down The rise in the dollar's value creates a bad gamma trading loss and drives the volatility of the put higher along the smile. In this case, the rise in vol along the smile is offset by its decline due to the vol shock.

Spot down, vol up The fall in the dollar's value also creates a gamma loss. The vol falls along the volatility smile as the put becomes, say, a 60-delta option, offsetting the rise due to the vol shock.

Spot down, vol down This is the best combination for the hedged short at-the-money put. The vol falls both along the smile and because of the vol shock. The result is a gain offset only by the gamma loss.

A similar analysis can be applied to hedged long put, and long and short call, portfolios. The long put portfolio behaves precisely as does the short put portfolio, but with signs reversed. The hedged call portfolio behaves in precisely the same way as the hedged put portfolio.

To understand the net effect of the smile on VaR, we need also to consider the correlation between the at-the-money implied volatility and the spot rate. Because of the pronounced negative correlation between spot and vol, there is a preponderance of scenarios in which spot and vol move in opposite directions, which leads to offsetting effects of the smile and the vol shock. As a result, the smile lowers VaR from $2589 to $2474. In general, one can expect a bearish volatility smile, i.e., one skewed to lower strikes, when there is a negative correlation between implied volatility and the underlying price. The smile is thus under normal circumstances somewhat "self-insuring". With a dollar-bullish smile, in contrast, VaR is increased to $3062.

Example 4.2 The smile can have a large an impact on the VaR of a risk reversal. A risk reversal is vega neutral: its delta is identically equal to 0.50 at the time it is initiated, since it combines a short and long position with deltas of 25%, while the vega is close to zero at initiation, as seen in Figure 6. As is the case for any portfolio of options, the delta and vega must be actively managed as the spot rate changes. At the time of initiation, in the absence of a volatility smile, the return profile of the risk reversal is nearly symmetrical. In the presence of a smile, the return distribution is no longer symmetrical, but can be heavily skewed.

Consider a short position in a yen-bullish risk reversal consisting of a long position in a 25-delta dollar put and a 25-delta dollar call in a market in which the volatility smile is as shown in Figure 4. Dealers might have such a position

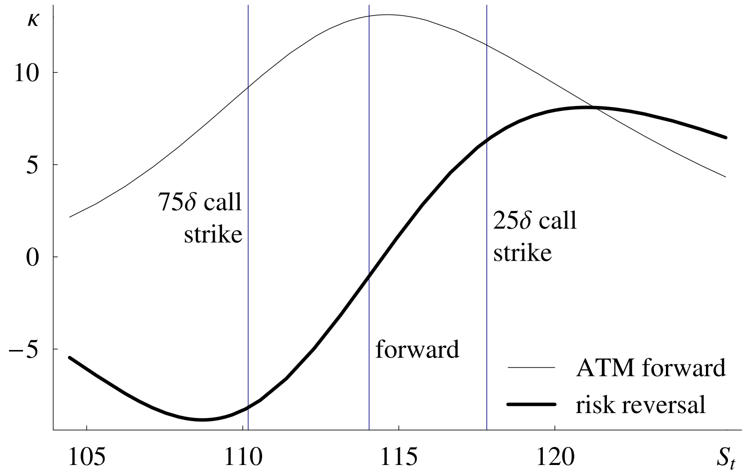

FIGURE 6. Vega of a risk reversal. One-month dollar–yen options, premium in JPY, 8 February 1999. Black–Scholes vega incorporating the smile by evaluating equation (1), for each value of the spot rate, at the implied volatility associated with the delta for that spot rate. Source: DataMetrics.

because of customer demand for protection against a weaker dollar. Because a risk reversal is vega-neutral and has low gamma risk, the VaR of a delta-hedged risk reversal is very low in the absence of the smile. In the presence of a smile, however, the VaR increases dramatically. In the case at hand, VaR more than triples, from \$216 to \$712, as seen in Table 1.

4.2 *Random Volatility Smile*

There is a more complex approach to incorporating smile effects on VaR, in which the 25-, 50-, and 75-delta implied volatilities are permitted to vary in correlated fashion. For exchange-traded options, one can permit the at-the-money implied volatility and the vols corresponding to one or two strikes on each side of the current futures price to vary together. The procedure is mechanically analogous to time bucketing along a maturity spectrum. To obtain an implied volatility for the shocked value of the underlying price, we interpolate between the shocked smile points using the interpolation scheme described above.

Example 4.3 We first calculate correlated shocks to spot and the 25-, 50-, and 75-delta implied volatilities. The vector of standard deviations is (0.0097, 0.0558, 0.0567, 0.0554) and the correlation matrix is

$$\begin{bmatrix} 1 & -0.300 & -0.395 & -0.483 \\ & 1 & 0.986 & 0.946 \\ & & 1 & 0.986 \\ & & & 1 \end{bmatrix}.$$

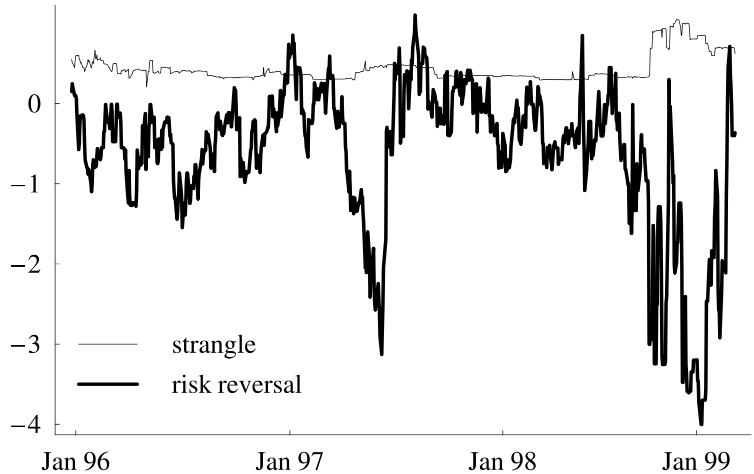

FIGURE 7. Dollar–yen 1-month 25-delta risk reversals and strangles. Source: DataMetrics.

The 25-, 50-, and 75-delta implied volatilities are highly correlated with one another. The 75-delta call vol has a markedly higher negative correlation with the spot rate, i.e., when the dollar weakens the skewness against the dollar becomes more pronounced and the risk reversal price becomes more negative. This may amplify some of the effects of the smile on VaR outlined in the previous section. In the case of the short at-the-money put portfolio, the effects are small, as seen in Table 1. The VaR is now $2384.

Both the term structure and the shape of the volatility smile generally change slowly over time compared with the level of implied volatility. However, when the levels of spot and implied volatility change abruptly, it is often accompanied by a sharp change in the slope of the term structure and the curvature of the smile. We are unlikely to pick up these relationships with our lognormality assumption regarding implied volatilities, and it is probably best to stress-test for their effects. A table of changes in term spreads and measures of smile curvature accompanying large changes in cash prices and implied volatility can provide scenarios for such stress tests. Figure 7 shows how abruptly the smile can change shape.

4.3 *The Term Structure of Volatility and Vega Risk*

One effect of the term structure of volatility is that as an option "ages", its implied volatility changes along the term structure. The mechanics of the aging of options in moving their implied volatilities along the prevailing term structure are analogous to the mechanics of cash price shocks moving implied volatilities of options along the smile. The difference is that the effect of "aging" is deterministic rather than random.

Except for very short term options (less than 1 month), the VaR horizon of 1 day is typically short compared with the spread between the vols of options with different maturities. The shortening of the option maturity by 1 day, or even 10 days, is therefore unlikely to change the implied volatility materially along the term structure. For example, if the spread between the 1-month implied volatility and the 3-month implied volatility is a relatively high 1.20 vols, the aging of a 3-month option will change its volatility by only 15 basis points (0.0015) per day. This is typically small compared with the vol-of-vol, which in Example 2.1 was 0.0567. We will therefore ignore changes in the term of the option.

Just as risk managers must address changes in the shape of the volatility smile as well changes along the smile due to variations in cash prices, they must take into account changes in the shape of the term structure of volatility. Term structure risks can be addressed by dividing the maturity spectrum into buckets. For each underlying asset, options in each maturity bucket has a vega risk correlated with that of other maturity buckets, with the cash asset prices, and with other assets. In practice, the availability of implied volatility data may limit the fineness of the maturity grid.

Example 4.4 (Term structure of dollar–yen volatilities) Returning to the example of dollar–yen, the vector of standard deviations is (0.0097, 0.0567, 0.0226) and the correlation matrix of the spot rate and the 1- and 12-month implied volatilities is

$$\begin{bmatrix} 1 & -0.395 & -0.085 \\ & 1 & 0.769 \\ & & 1 \end{bmatrix}.$$

In this example, shorter- and longer-term implied volatilities are not as highly correlated with one another as points on the volatility smile. The correlations of implied volatility with the spot rate decay rapidly as the time to maturity

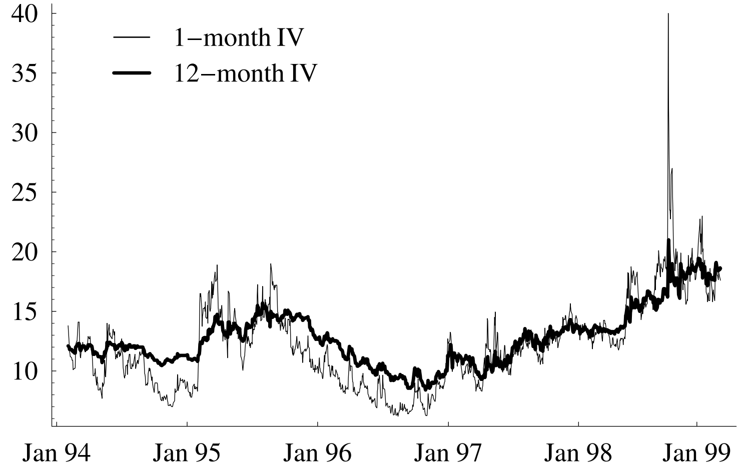

FIGURE 8. Dollar–yen 1-month and 12-month at-the-money forward implied volatility. Source: DataMetrics.

increases. As one would expect in light of its mean-reversion properties, longer-term implied volatility is less variable than shorter-term implied volatility. These important difference between shorter- and longer-term volatility need to be captured in an accurate VaR analysis of a portfolio containing both.

As noted at the end of the last section, and can be seen in Figure 8, the term structure can change shape abruptly, generally at the same time that cash prices and the level and shape of volatility smile are shifting. Stress testing is therefore a required supplement to standard VaR analysis.

5. CONCLUSIONS

We conclude with some generalizations and suggestions for risk managers on dealing with vega risk, which can be responsible for a large part or even the bulk of the VaR of a portfolio containing options.

The volatility smile can have a large impact on VaR. However, the smile may also reduce vega risk because large volatility changes along the smile are offset by volatility changes correlated with spot movements.

A vega-neutral portfolio has only a small exposure to the general level of implied volatility, but may have a very large exposure to the volatility smile. Risk managers encountering vega-neutral trades should be skeptical of VaR reports that do not take account of the smile.

In order to lay out the structure of the risk measurement techniques proposed here as clearly as possible, we have focused largely on VaR computation. Recently, however, risk management research and practice have become more aware of the inadequacy of VaR as a central risk measure in view of the evident nonnormality of asset price behavior.[18] The techniques presented in this paper could readily be adapted to provide basic stress-test measures such as the loss from an adverse 5- or 10-vol move in a subset of markets.

For example, the techniques presented in Section 4 assume that implied volatility is jointly lognormally distributed with other asset prices and rely on the volatility smile to correct this assumption with the market's expected return distribution (in the risk-neutral sense). Rapid changes in the shape of the smile are likely to accompany sharp changes in the levels of implied volatility and cash assets. Stress testing can reveal the vulnerability of a portfolio to large smile changes, especially for a vega-neutral or low-vega portfolio. Extensions of the work presented in this paper and further research on vega risk management generally should take account of work currently being done on stress testing, extreme value theory, and other approaches to measuring the risk of large losses.

[18] Many recent examples of this work can be found on the *All About Value-at-Risk* web site at `www.gloriamundi.org/var/stress.html` and in the April 2000 issue of the *RiskMetrics Journal*, available `www.riskmetrics.com/research/journals/rmj2q00.pdf`.

REFERENCES

Aït-Sahalia, Y., and Lo, A. W. (2000). Nonparametric risk management and implied risk aversion. *Journal Of Econometrics*, **94**(1/2), 9–51.

Bakshi, G., Cao, C., and Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance*, **52**(5), 2003–2049.

Bates, D. S. (1991). The Crash of '87: was it expected? The evidence from options markets. *Journal of Finance*, **46**(3), 1009–1044.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Journal of Financial Studies*, **9**(1), 69–107.

Campa, J. M., and Chang, P. K. (1995). Testing the expectations hypothesis on the term structure of volatilities in foreign exchange options. *Journal of Finance*, **50**(2), 529–547.

Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*. Princeton University Press.

Demeterfi, K., Derman, E., Kamal, M., and Zou, J. (1999). A guide to volatility and variance swaps. *Journal of Derivatives*, **6**(4), 9–32.

Derman, E. (1999). Regimes of volatility. Quantitative Strategies Research Notes, Goldman Sachs. Available at `www.gs.com/qs/doc/regimes.pdf`.

Derman, E., and Kani, I. (1994). Riding on a smile. *Risk*, **7**(2), 32–39.

Duan, J.-C. (1999). Conditionally fat-tailed distributions and the volatility smile in options. Mimeo. Available at `www.bm.ust.hk/~jcduan/ft_opt.pdf`.

Dumas, B., Fleming, J., and Whaley, R. E. (1998). Implied volatility functions: Empirical tests. *Journal of Finance*, **53**(6), 2059–2106.

Dupire, B. (1997). Pricing and hedging with smiles. In: *Mathematics of Derivative Securities* (ed. M. A. Dempster and S. R. Pliska), pp. 103–111. Cambridge University Press.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, **6**(2), 327–343.

Heynen, R. (1994). An empirical investigation of observed smile patterns. *Review of Futures Markets*, **13**(2), 317–354.

Hull, J., and White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, **47**(2), 281–300.

Malz, A. M. (1996). Using option prices to estimate realignment probabilities in the European Monetary System: The case of sterling-mark. *Journal of International Money and Finance*, **15**(5), 717–748.

Malz, A. M. (1997). Estimating the probability distributions of the future exchange rate from option prices. *Journal of Derivatives*, **5**(2), 18–36.

Melick, W. R., and Thomas, C. P. (1997). Recovering an asset's implied PDF from option prices: An application to oil prices during the Gulf Crisis. *Journal of Financial and Quantitative Analysis*, **32**(1), 91–115.

Mina, J., and Ulmer, A. (1999). Delta-gamma four ways. Working Paper, RiskMetrics Group.

Morgan Guaranty Trust Company (1996). *RiskMetrics Technical Document*, 4th edn.

Peña, I., Rubio, G., and Serna, G. (1999). Why do we smile? On the determinants of the implied volatility function. *Journal of Banking and Finance*, **23**(8), 1151–1179.

Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, **49**(3), 771–818.

Stein, E. M., and Stein, J. C. (1991). Stock price distributions with stochastic volatility: An analytic approach. *Review of Financial Studies*, **4**(4), 727–752.

Stein, J. (1989). Overreactions in the option markets. *Journal of Finance*, **44**(4), 1011–1023.

Xu, X., and Taylor, S. J. (1994). The term structure of volatility implied by foreign exchange options. *Journal of Financial and Quantitative Analysis*, **29**(1), 57–74.