NBIM DISCUSSION NOTE

#15 - 2012

Modelling the implied tail risk of foreign-exchange positions

Executive Summary

- Standard risk models based on historical data have deficiencies: the results of these models depend heavily on the period used for calibration. In addition, the models make unrealistic ex-ante assumptions (such as on the distribution that price changes follow) that may lead to underestimation of tail risks.
- In this note, we walk through a risk-modelling approach that, as far as possible, uses forward-looking, market-implied information from the currency option market to shed light on the risk of the Fund's currency position. This model is used at NBIM as a complement to standard risk models.
- Analysing the Fund's new benchmark for fixed-income investments, we find that increasing the allocation to emerging-market currencies adds to short-term currency tail risk, when measured in a common reference currency. The market-implied tail-risk model predicts expected losses in adverse years to be 30-60 percent higher than the standard models.

Motivation and background

From a risk management perspective, tail risks and return distribution asymmetries of investments are important to analyse. In this note, we describe a modelling approach that addresses some of the weaknesses of standard risk models. We use the model internally as a complement to standard models to evaluate tail risk in foreign-exchange (FX) positions. Examples of tail events for FX positions could be single-currency devaluations or more widespread flight-to-quality/carry-trade-unwind episodes. Formal tail-risk modelling involves a lot of mathematics, and while we have left out some details, we have kept relevant technicalities to motivate the approach we have chosen.

Most risk models in practical use at asset managers assume that FX rate changes follow a normal distribution. This is clearly not the case in the real world. Figure 1 shows the empirical weekly NZD/JPY FX rate return from January 1997 against a theoretical normal distribution having the same volatility as the data set. The NZD/JPY return on the x-axis will be negative if the NZD weakens against the JPY. It is clearly seen that the left tail of the fitted normal distribution does not reflect the fat tail in the empirical distribution. In other words, if a normal distribution is used when modelling the FX risk, we would underestimate the tail risk of a long NZD, short JPY position. This position is an example of a typical carry strategy, where one borrows in lower-yielding currencies such as JPY and invests in higher-yielding currencies such as NZD.

Figure 1. Empirical NZD/JPY distribution versus a fitted normal distribution.

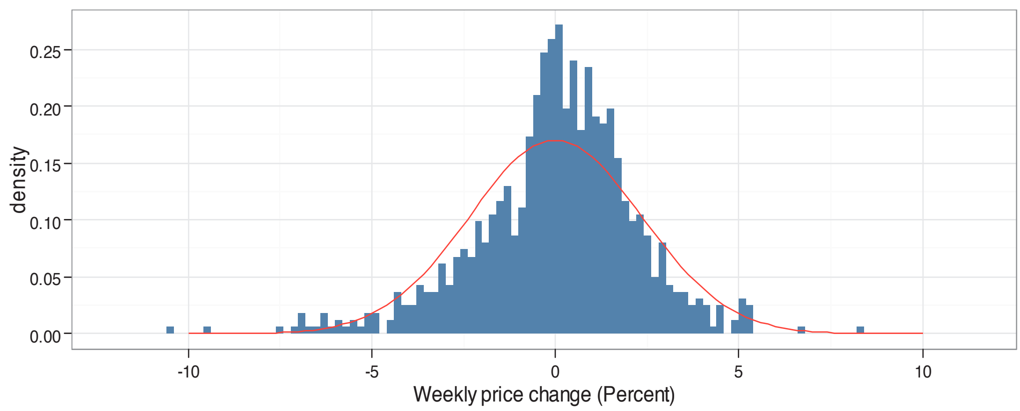

Dynamic volatility models, such as GARCH models, could potentially model fat-tailed distributions for a single financial time series such as NZD/JPY. However, these models do not easily extend to the multivariate case that is needed for portfolio risk modelling. Another approach to mitigate this deficiency is to use a parametric distribution that allows fat tails (Student-t, skewed Student-t and normal inverse Gaussian (NIG) are examples). Estimation of the parameters for the normal distribution is straightforward, which is probably why it is often used as the default assumption. Estimating the parameters for fat-tailed distributions can be much more complex, and using these will not mitigate another important issue: the selection of historical data used for estimation.

Measurement of risk usually involves using price history as input for estimating the models used for forward projections of risk. To make the risk estimates relevant for the current period, a rolling historical window is often used, either explicitly or implicitly by applying a decay factor to return series. If the resulting data set is used to calibrate a parametric model and it does not contain any large moves, a parametric model will have difficulties in predicting large moves in the future. One of the challenges then is in anticipating that large moves are possible when prices and rates have been relatively stable in the estimation period. Figure 2 illustrates this issue by looking at four previous emerging-market (EM) devaluation episodes (against USD) for four managed floating currencies. Measured volatility would be relatively low leading up to the devaluation, as price movements were small before. Volatility measurement for floating exchange rates has the same problem in periods of apparent stability.

Figure 2. Four currency episodes. Indexed to value before drop.

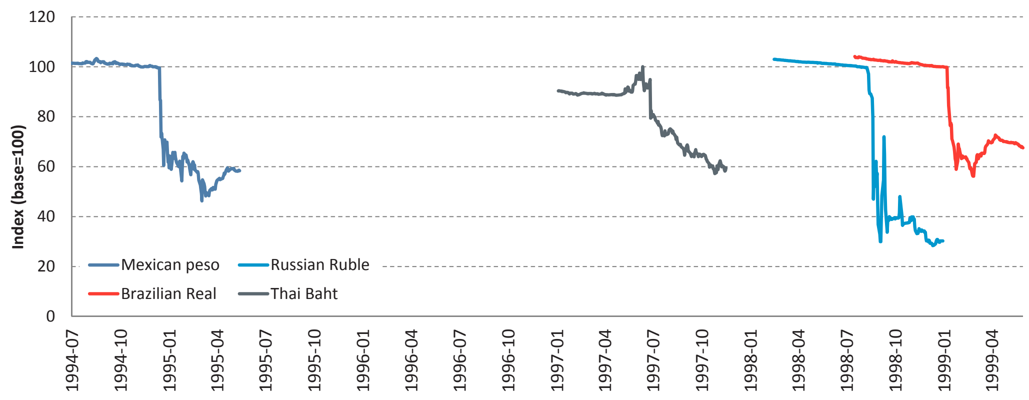

One of our goals is to develop a risk model that captures the fat-tailed feature of the empirical distributions of FX rates. In addition, we would like it to be independent of the sampling period and update its estimates quickly to new information.

A forward-looking risk model

Foreign-exchange option markets trade on expectations of *future volatility*. The next sections will describe our approach to constructing a currency-risk model based on price data from these markets. A final section will present a short portfolio analysis by applying the model to the Fund's new fixed-income benchmark and comparing the results to standard models. The approach as presented is equally applicable to the currency exposure of the equity benchmark. However, the FX risk for an unhedged equity position (where the business risk of the firm will dominate) accounts for a smaller part of its total risk than for a fixed-income position. It is important to note that we do not believe that one model can have all the answers. However, the aim is to improve on basic risk models, and we believe the option-based approach can provide several key benefits, as detailed below. Figure 3 shows the outline of the steps we follow to tie what we call a "forward-looking" risk model for foreign-exchange positions together.

Figure 3. FX modelling steps, full market-implied model.

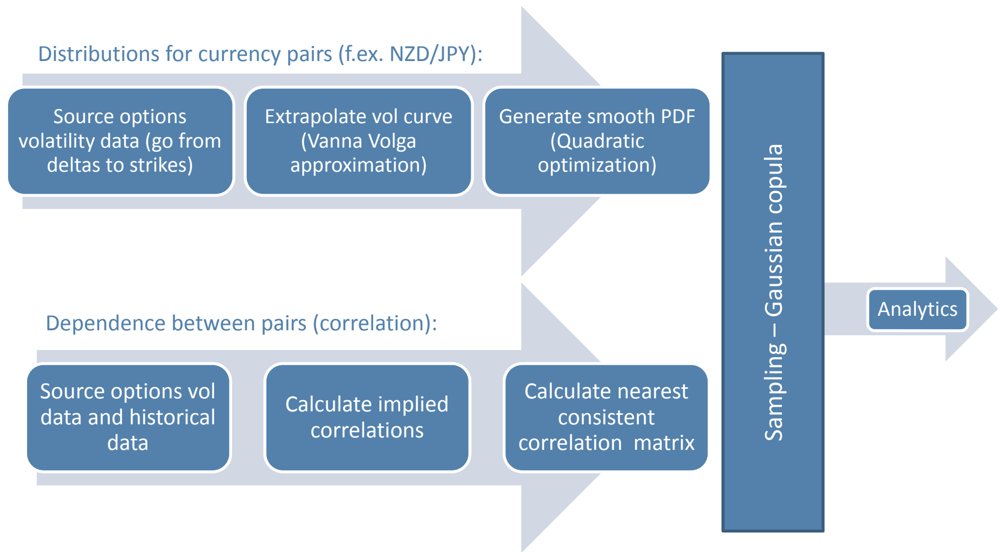

Briefly put, the risk model consists of three steps. In the first step, we calculate the market-implied distribution for individual currency pairs. In the second step, we calculate a correlation matrix to describe the dependencies between all FX pairs, again using data from option markets as far as possible. In the third step, we employ a simulation procedure using the parameters from the first two steps to arrive at a distribution for the portfolio we are analysing. From this distribution we can calculate all the risk statistics we are interested in.

A summary of the key differences between the risk model described in this note and a standard (lognormal) risk model is provided in Table 1.

Table 1. Comparison of standard models and option-based model.

	Standard lognormal risk model	**NBIM option-implied method**
	Model assumptions and requirements	
Asset distribution	Log returns follow a normal distribution	Assumption-free
Dependence structure	Linear (multivariate lognormal)	Linear (Gaussian copula)
Data requirements	Historical price data	Option volatilities
Inference	Assumes that future distribution can be inferred from historical price movements.	Assumes that risk-neutral implied probabilities are good proxies for real probabilities, i.e. it assumes that option prices contain aggregated information on the market view of the probability distributions of future prices.
	Model limitations and advantages	
Single-asset tails	Does not model fat tails.	Allows fat tails since it is distribution-free.
Correlated tail events	Does not model clustering of tail events.	Does not model clustering of tail events.
Data-quality sensitivity	Quality of price quotes less of an issue, but results are highly dependent on historical period used being representative.	Dependent on good-quality data and a sufficient number of data points (strikes) available.
Implementation effort	Low	Medium

Implied volatility

An option contract gives its owner the right, but not the obligation, to buy (call option) or sell (put option) an asset at a predetermined price (strike) at some date in the future[1]. This right to buy or sell has value and will therefore have a price. The convention in the FX option market is to quote prices at "sticky deltas", with the actual strike being implied by the Black-Scholes option-pricing model. This means that quotes will be provided in volatility terms for a given delta value. The delta of an option contract represents its sensitivity to price change. More formally, the delta is the derivative of the option-pricing formula with respect to a price change. The strike for the 25 delta call is the strike that makes this derivative equal to 0.25, and for the 25 delta put it is -0.25. This implies that, according to the option-pricing model, a 1 unit increase in the FX rate will increase the value of the call option by 0.25 units and the put option will lose 0.25 units in value.

Informally, we can say that the implied volatility of an option measures the probability of reaching the strike. Figure 4 shows the implied volatility (y-axis) versus deltas for puts and calls (x-axis) on the NZD/JPY pair for the next year. The Black-Scholes model used to price options assumes that prices follow geometric Brownian motions where volatility is flat and constant across strikes. The empirical data show that the NZD has a tendency to suddenly weaken, in a non-continuous way, against the JPY. Option sellers therefore demand a higher price for selling out-of-the-money puts than out-of-the-money calls on the NZD/JPY currency pair. When the option-pricing model is inverted to solve for the volatility

1 A so-called European option gives the holder the right to exercise on a specific date, while the holder of an American option would have the right to exercise the contract on all days up to an agreed date.

that is consistent with the prices in the market, we observe a *volatility skew* across strikes: the 10 delta put is quoted at around 24 percent volatility and the 10 delta call at around 14 percent volatility.

Figure 4. Implied volatility one year ahead for NZD/JPY. P: puts, C: calls (25/07/2012).

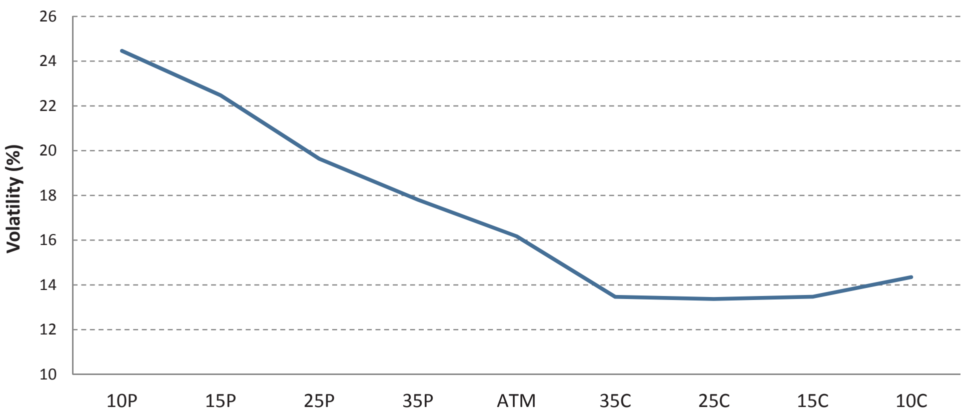

It has been documented that, for assets with active option markets, traded volatilities have a tendency to be higher over time than the volatility that is subsequently realised[2]. This would lead us to overestimate risks and is a potential weakness of the implied-volatility approach. Our working assumption is that the pricing in option markets provides useful additional information. Another compelling argument for using the method is that the underlying options can actually be traded if one wants to manage the risk exposure.

Foreign-exchange market data

Liquidity is an important consideration when implementing a model driven by market data. If trading volumes are too thin, liquidity effects rather than information content will dominate changes in prices, making the results of the model less relevant.

Due to the global nature of the foreign-exchange market, and to the fact that a lot of trading takes place over the counter (OTC), putting an exact number on trading activity is difficult. The Bank of International Settlements (BIS 2010) estimates the average daily turnover traded OTC for all foreign-exchange derivatives products (forwards, swaps, options etc.) combined to be USD 2.5 trillion, USD 2.1 trillion of which in advanced economies and USD 0.4 trillion in emerging-market economies. Turnover in "currency options and other products" contributes 9 percent and 6 percent of the total in advanced and emerging countries respectively. Although much lower in emerging markets, trading activity there is rapidly expanding – turnover was up 140 percent from 2004 to 2010. USD plays an important role in FX trading, particularly in emerging markets: it was one of the currencies in 95 percent of all FX derivatives transactions in 2010.

Figure 5 shows how far away from the current price the implied strikes are for 10 and 25 delta puts and calls maturing in one year, as quoted by Bloomberg. Official statistics for the distribution of trading volumes along the volatility curve are not available. Discussions with trading counterparties indicate, however, that there is some liquidity in most emerging-market currency pairs versus USD one year ahead at strikes around the 25 delta level. Trading activity around the 10 delta level is much lower. We use at-the-money (ATM) and 25 delta quotes and acknowledge that we might overshoot on some of the risk statistics calculated from the smallest emerging-market currencies, where liquidity might be lower and market-makers may adjust their volatility quotes upwards.

2 See Carr and Wu (2009).

Figure 5. Option deltas translated to price changes against USD (20/07/2012).

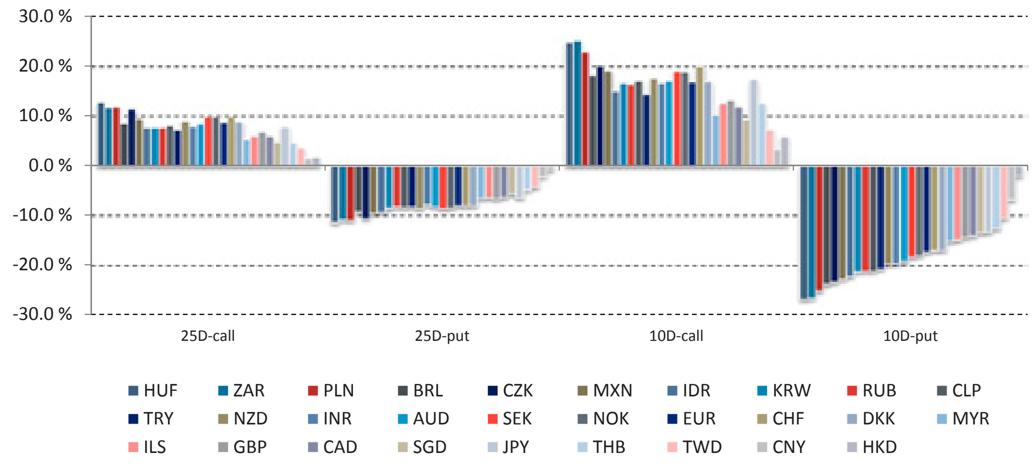

Market-implied distributions from implied volatilities

In the following, we rely heavily on a result that follows from the framework of arbitrage pricing theory: it is possible to mathematically deduce the risk-neutral market-implied probability distribution from the implied volatilities of traded options. We describe the result briefly below; see Breeden and Litzenberger (1978) for a more thorough description. If we let the price of a call option $c(x)$ be given by the risk-neutral expectation:

$$c(x) = \int_x^\infty (s - x) f(s) ds \tag{1}$$

where x is the strike price and $f(s)$ is the implied probability density for the currency pair, we can show that the second derivative of $c(x)$ with respect to x is $f(x)$:

$$\begin{aligned} \frac{\partial c}{\partial x} &= \frac{\partial}{\partial x} \int_x^\infty s f(s) ds - \frac{\partial}{\partial x} x \int_x^\infty f(s) ds \\ \frac{\partial c}{\partial x} = -x f(x) - \int_x^\infty f(s) ds + x f(x) &= -\int_x^\infty f(s) ds \\ \frac{\partial^2 c}{\partial x^2} &= f(x) \end{aligned} \tag{2}$$

Hence it is possible to extract the implied probability density function $f(x)$ for the currency pair given the implied volatility curve: if we have the volatility curve, we have $c(x)$ and its derivatives, holding other inputs to the option-pricing model constant.

It is worth noting that although we treat the risk-neutral probabilities above as "real probabilities", they are synthetic probabilities valid within the standard derivatives-pricing framework. To get from risk-neutral to real probabilities, we would have to calculate the premium investors require as compensation for taking on the risk. If we assume the risk premium is a constant[3], the main effect would be to shift the probability distribution by the premium required. Since our interest is in the shape of the probability distribution, not the absolute location, we chose to ignore this effect. For an accessible introduction to the relation between real and risk-neutral probabilities, see Bank of England (2012).

The vanna-volga method

As shown above, to calculate implied probability densities we need to take second derivatives of the option price with regard to the strike. To cover all possible outcomes, this means that we need call prices for all possible strikes. However, as quotes can only be reliably sourced for the 25 delta (25D) puts and calls and the ATM option, we need to approximate values of implied volatilities for missing strikes. The vanna-volga method is commonly used to accomplish this task. See Mercurio and Castagna (2005) and Mercurio and Castagna (2007) for a more detailed description than provided below.

We first assume that volatility is flat and stochastic. Then, by using the standard approach in arbitrage pricing theory, we set up a hedged portfolio that is long the option with strike K and expiry t that we are trying to price, short an amount of the underlying and short some amounts x_i of the calls we know the prices of (ATM, 25D put, 25D call), where $i \in 1,2,3$ indexes the three options. The x_i weights chosen are such that the short portfolio will have the same sensitivities to changes in Black-Scholes parameters as the option we are trying to price. More formally, the portfolio is locally hedged to all partial derivatives up to the second order. The sensitivity to price change is hedged by the short position in the underlying; it can also be shown that if two options have the same vega (derivative of the option price with respect to volatility) and same maturity, they will have the same gamma (second derivative with respect to price). The weights x_i can then be found by matching three sensitivities: the vega, the second derivative with respect to volatility (volga), and the partial derivative of the option price first with respect to volatility then to the FX rate (vanna):

$$\begin{aligned}
\text{vega} &= \frac{\partial C^{BS}}{\partial \sigma}(t;K) &= \sum_{i=1}^{3} x_i(t;K)\frac{\partial C^{BS}}{\partial \sigma}(t;K_i) \\
\text{volga} &= \frac{\partial^2 C^{BS}}{\partial \sigma^2}(t;K) &= \sum_{i=1}^{3} x_i(t;K)\frac{\partial^2 C^{BS}}{\partial \sigma^2}(t;K_i) \\
\text{vanna} &= \frac{\partial^2 C^{BS}}{\partial \sigma \partial S_t}(t;K) &= \sum_{i=1}^{3} x_i(t;K)\frac{\partial^2 C^{BS}}{\partial \sigma \partial S_t}(t;K_i)
\end{aligned} \tag{3-5}$$

Mercurio and Castagna (2007) show that the weights can be uniquely determined analytically. They further show that, given the weights, the price of the option can be approximated by adjusting the flat-volatility price C^{BS} by the premium differences of actual market prices C^{MKT} (where volatility is not flat) to a flat-volatility price for the short options in the hedging portfolio:

$$C(K) = C^{BS}(K) + \sum_{i=1}^{3} x_i(K)\left[C^{MKT}(K_i) - C^{BS}(K_i)\right] \tag{6}$$

An approximation of this expression can be inverted to find the implied volatility at the unquoted strike K. We include a more detailed development of the expression in the appendix to this note. Figure 6 shows two volatility curves, for BRL and CHF against USD, resulting from applying the vanna-volga approximation to the volatility mid-quotes for ATM, 25D put and 25D call options. The resulting curve overlaps quite well with quotes at other strikes in the market.

3 This is a strong assumption. The consensus in modern financial economics is that risk premiums are time-varying; for a review, see NBIM Discussion Note 1-2011. Further, Ilmanen (2011) argues that tail risks are particularly well-rewarded (i.e. overpriced) following tail-risk events. If that is the case, a model based on option volatilities would overestimate the "true" tail risk in a position. A counter to such an argument would be that option volatilities are tradable; if tail risk is clearly mis-priced, this should eventually attract profit-seeking volatility sellers.

Figure 6. Extrapolation, CHF/USD and BRL/USD volatility curves (13/06/2012).

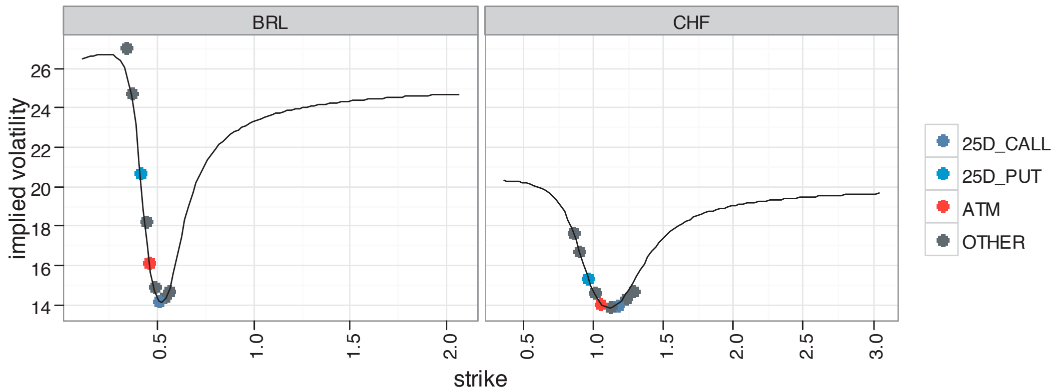

Model implementation

Single FX rate distributions

For the liquidity and data availability reasons discussed previously, we use option market quotes from Bloomberg for all currencies against USD to estimate implied probability densities. We have found that, for the three strikes needed for extrapolation, the Bloomberg quotes are very much in agreement with quotes received directly from our counterparties.

The vanna-volga extrapolation scheme produces relatively smooth curves outside the traded range, see Figure 7. It is seen that the EM currencies are more dispersed in terms of level of volatility and the volatility skew is higher overall compared to developed-market currencies, when measured against USD.

Figure 7. Volatility extrapolation. Developed versus emerging markets (20/07/2012).

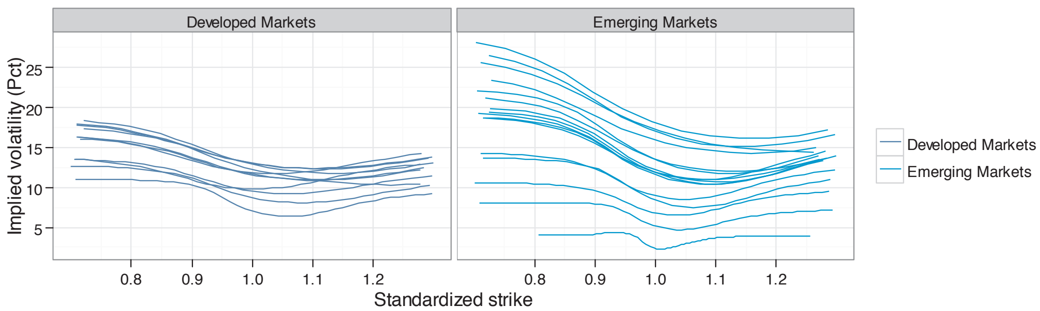

For some currencies, if a call price curve is created using the extrapolated volatility curves and second derivatives taken directly, implied probability density will go negative. This is most problematic for currency pairs with steep volatility skews. To circumvent this issue, we adopt the novel approach described in Fengler (2005). This scheme converts the differentiation problem into a spline-fitting problem. In short, call prices calculated from our first-pass extrapolation of the volatility curve are fed into a procedure that fits a smoothing spline that penalises deviations from an arbitrage constraint: the call price must be a decreasing and convex function of the option's strike price.

As a spline is defined by knots connected by first and second derivatives, the procedure allows us to get rid of negative probabilities and calculate implied probability densities in the same operation. The spline-fitting procedure is a quadratic minimisation problem with constraints.

Figure 8 demonstrates the effect of the smoothing spline applied to the call price. A minor adjustment to the call price function created by the extrapolation procedure creates a smooth probability density function for the BRL. In this example, the higher option premium demanded by put sellers in the option market manifests itself as a clearly defined heavy downside tail for the implied BRL/USD rate distribution.

Figure 8. Smooth spline interpolation of call price and resulting PDF (BRL/USD on 13/06/2012).

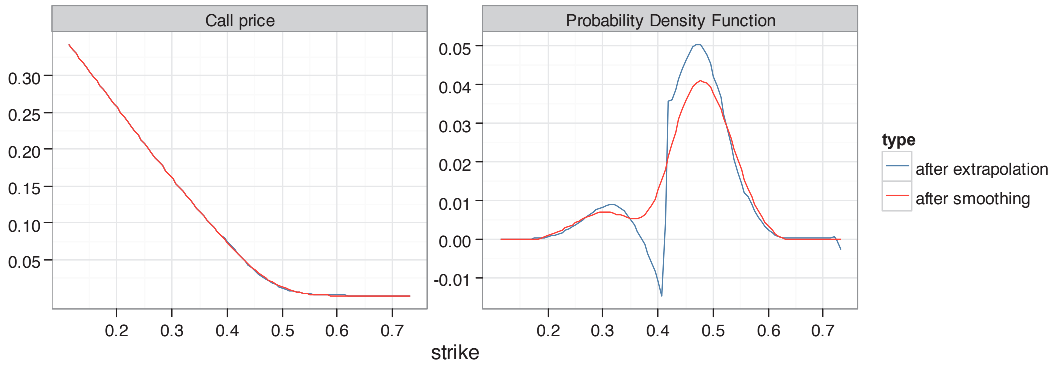

Given full probability distributions for single currencies, we can calculate some preliminary analytics. In Figure 9, we show the probability of three sets of currencies depreciating more than 15 percent against USD, as calculated using four different assumptions and models in mid-June 2012. Two of them are option-based; one uses a flat volatility curve equal to the ATM volatility, and the other uses the volatility curves created as described above. The last two both use a normal distribution assumption and different historical data windows of 60 days and 3 years.

The option-based methods predict higher probabilities of tail events versus USD than the methods based on historical data for all currencies. This is due to two factors: traded implied volatilities are currently higher than realised volatilities, and, as expected, tails are heavier when calculated using information from the full volatility curve than from the flat ATM volatility estimates. Emerging markets generally have heavier tails. It is also interesting to see that the option-based model calculates a non-zero probability of managed floating currencies such as TWD and THB moving by more than 15 percent against USD in the next year. Speculation of any breaks would feed into our estimates as soon as the market starts to trade on it.

Figure 9. Probability of 15 percent depreciation against USD (13/06/2012).

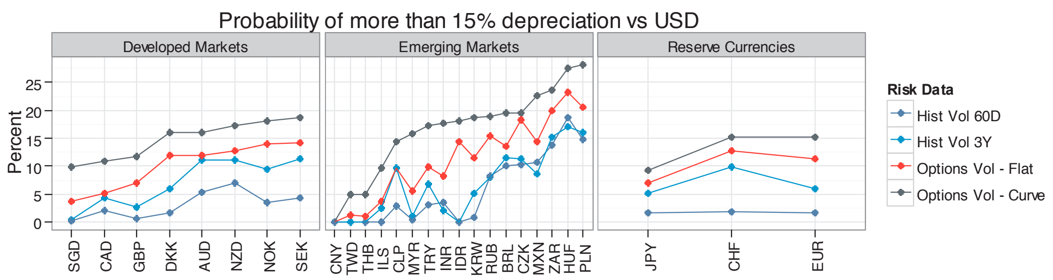

Tail-risk estimates for traditional "reserve currencies" such as JPY, CHF and EUR (which would be expected ex ante to have a lower and possibly negative correlation to equity markets and risky assets in a downturn) show a marginally lower tail risk than for other currencies. The impact of the recent European situation, in which EUR and CHF (pegged to EUR) did not trade like reserve currencies, can be seen in Figure 10. The example also highlights a key benefit of using an option-based approach: increased implied skewness in the probability distribution of the EUR/USD rate is quickly taken into the model as the market perception of EUR changes (the upper tail, which is not shown, does not move as much). A parametric historical model would only allow for this if skewness was explicitly modelled and large moves had already occurred in the period used to calibrate the model.

Figure 10. Reserve-currency tail risk estimated using option volatility curves.

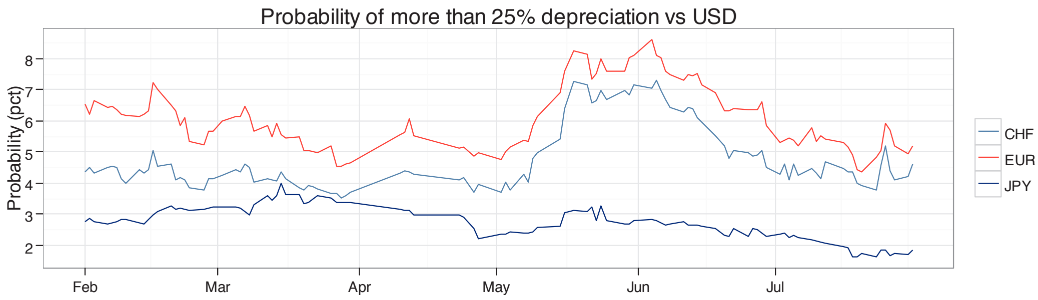

Dependence measure – implied correlations

Since we use market information directly in our modelling of single-currency-pair distributions, and do not make ex-ante assumptions on the (marginal) distributions that FX rates follow, a copula-based approach to dependence modelling is well-suited. This allows us to model single currencies separate from dependencies across rates. In our simulation model, we use a so-called Gaussian copula to draw correlated uniform variables, corresponding to points on the marginal cumulative density functions of currencies against USD.

A case for using other copula models with stronger tail dependencies than the Gaussian can certainly be made, as synchronised large moves in currencies against USD are plausible in a flight-to-safety scenario. However, in selecting a Gaussian copula as our dependence generator, we can compare standard multivariate lognormal risk models (implemented in standard risk software) against market-implied models in the same simulation setup.

The global currency market is one of few markets where the correlation can be inferred from traded prices. As long as all three currencies in a currency triangle are traded and liquid, a no-arbitrage argument requires the following to hold for the correlation between USD/AUD and USD/JPY rates – see, for example, Haug (1996) or Lopez and Walter (2000):

$$\rho_{\frac{USD}{AUD},\frac{USD}{JPY}} = \frac{\sigma^2_{\frac{USD}{AUD}} + \sigma^2_{\frac{USD}{JPY}} - \sigma^2_{\frac{AUD}{JPY}}}{2\sigma_{\frac{USD}{AUD}}\sigma_{\frac{USD}{JPY}}} \tag{7}$$

All variables on the right-hand side of the equation can be inferred from the option market. We discard any correlation skew and focus on the ATM options in our application. To construct a full correlation matrix that also includes pairs of currencies with illiquid option markets, we collect a blend of market-implied and historical correlations into a candidate correlation matrix. This matrix does not necessarily have the consistency properties (i.e. invertibility) required to be used in simulation. We therefore use a procedure from Higham (2002) to find a "nearest consistent correlation matrix" which can be used for our simulation procedure.

There does not seem to be a strong academic consensus on whether traded FX correlations actually contain information beyond what is present in the price history. Lopez and Walter (2000) find implied correlations useful in forecasting observed correlations, but also find that it does not fully incorporate all information in the historical data. We find it interesting to track a model based on market-implied correlations alongside models based on historical correlations.

Simulation setup

We work with option volatilities on a one-year horizon and set up the simulation so that each FX rate has an expected value equal to the one-year forward rate. The Monte Carlo simulation is implemented as follows: we repeat the instructions below n = 100,000 times (as mentioned, for this study we use USD as the base currency):

- Draw a vector X from a multivariate Gaussian distribution with our estimated correlation matrices. Transform this vector to correlated uniforms by applying the univariate normal cumulative density function (CDF).
- For each FX rate, transform the uniform to a sampled outcome Y by applying an inverse CDF:
  - For the forward-looking method: the market-implied empirical inverse CDF
  - For a standard model: the inverse CDF of the normal distribution
- Revalue the portfolio (in the base currency) with the new FX rates Y and calculate profits and losses (P&L) versus the forward value of today's portfolio.

Portfolio analytics

At the end of June 2012, the Fund's strategic benchmark for fixed income, set by the Ministry of Finance, was expanded to include government bonds for all currencies in the Barclays Global Aggregate index. Ten emerging-market currencies were added to the benchmark as a result, having a combined weight of 5.9 percent at the time of implementation. Inclusion of EM bonds may lead to increased short-term currency risks when a developed-market currency is used as the reference currency. The reference currency used here is USD.

In this section, we look at some analytics for the currency composition of the Fund's new fixed-income benchmark as of 2 July 2012. The total value was NOK 1,428 billion at the end of the second quarter of 2012. It is important to note that this analysis does not take into account any credit or government-bond risks; it looks at FX risks in isolation. The currency composition of the new benchmark is shown in Figure 11.

Figure 11. Currency composition of the fixed-income benchmark (02/07/2012).

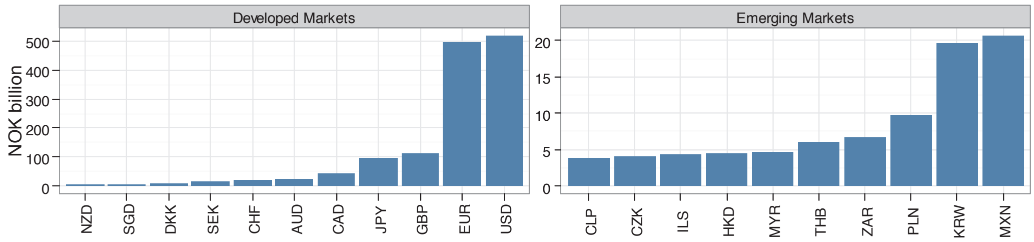

With this simulation setup, we can compare risk models. In Figure 12, we show the distributions one year ahead produced by four different simulation runs. In two of them, we use option-market information for the single-currency distributions. In the first, historical correlations over a short window are used, while the second uses correlations as implied by the market. Two normally-distributed models, one using a short, daily-frequency data window and the other using a longer, weekly-frequency data window, are run as benchmarks.

Figure 12. Risk model comparison (02/07/2012).

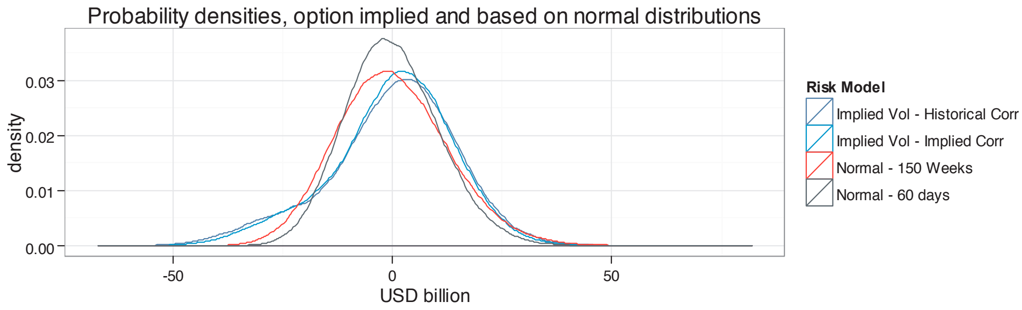

Two observations are apparent from the figure. First, the option market is pricing a much heavier downside tail for the portfolio than the normal models calculate. Second, for this snapshot there is not a large difference between a model using implied versus historical correlations. Standard risk statistics run on the positions are shown in Table 1. VaR84 indicates Value-at-Risk, the negative of the value at the 16th percentile of the P&L distribution; ES84 is taken to be the expected shortfall over the 16th percentile, the expected value of losses should losses below the VaR84 occur.

Table 1. Standard risk statistics, USD billion.

	VaR84	ES84	VaR95	ES95	VaR99	ES99
Implied Vol – Historical Corr	14.3	25.2	28.4	35.6	40.3	44.9
Implied Vol – Implied Corr	13.7	23.6	26.4	33.3	37.7	42.3
Normal – 150 Weeks	12.6	18.4	19.8	24.2	27.0	30.1
Normal – 60 days	10.6	15.5	16.8	20.5	22.9	25.8

For the percentile-based risk statistics close to levels usually reported by fund managers (standard deviation is equal to the VaR84 statistic when a normal model is used), the option-implied risk statistics are 10-40 percent higher. The differences increase as we go further out into the tail. The much higher expected shortfall statistics reflect the heavier tails modelled by the option-implied method: informally, the ES84 statistic shows that, in the worst year out of six, the drawdown in portfolio value is expected to be 30-60 percent more than standard models would indicate.

Figure 13. Implied distribution, with and without EUR (02/07/2012).

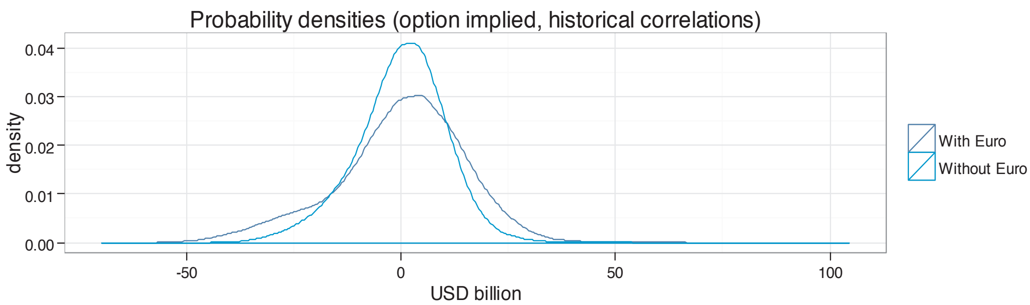

It is worth highlighting that the tail risk in Figure 12 is higher than what we would expect in a less turbulent period for the markets. The market is currently pessimistic about EUR, which shows up as a higher implied tail risk. This is illustrated in Figure 13, where we compare the distribution of the current benchmark currency composition to a benchmark where EM currencies and the USD position are held as today and the EUR holdings are distributed to the other developed-market currencies pro rata according to their current weight in the benchmark.

Figure 14. Tails of P&L distribution, varying amounts of EM (02/07/2012).

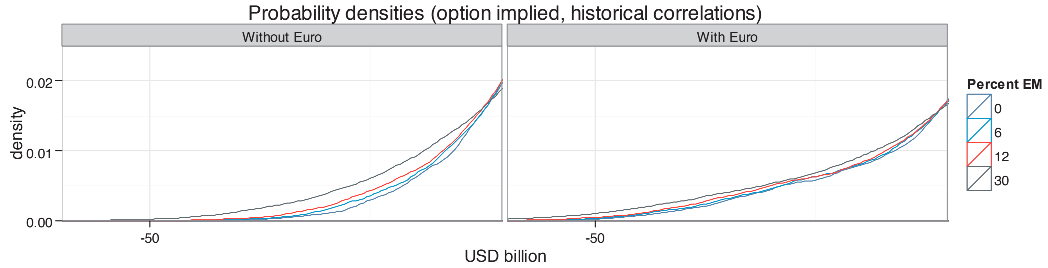

Finally, we illustrate the effect on currency risk of varying the fraction of the benchmark invested in emerging-market bonds. We consider two base portfolios without EM bonds: one including EUR and one without (following the same procedure as before). EM currency exposure is then increased pro rata according to the weights in Figure 11. The tails of the resulting distributions are shown in Figure 14. Increasing the EM allocation increases the tail risk for both portfolios, but the impact is larger for the portfolio not holding EUR. The impact on the non-EUR holding portfolio is in line with what we would expect from adding EM currencies to a developed-markets portfolio in a more normal market environment. It should also be noted that the EM portfolio does not include RUB or BRL, which might be included in the Fund's benchmark at a later stage. These currencies are among the EM currencies that have the highest volatility skews. Adding these currencies to the EM portfolio will further increase the downside tail.

FX drawdowns and other risky assets

Analysis of tail risks for currency positions is particularly important in a total portfolio context, as episodes of simultaneous weakening of higher-yielding currencies have a tendency to occur at the same time that risky assets such as equities or corporate bonds sell off (see Ilmanen (2011) for more detail). This is illustrated in Figure 15 where we show the performance of the S&P 500 and the Barclays US Corporate High Yield index against the USD performance of an equal-weighted basket of four EM currencies and their measured average pairwise correlation.

The chart shows that, as the 2008 selloff takes hold, EM currency correlations rise and the value of the currency basket is reduced in tandem with the credit and equity indices. The implication is that, although the improvement in expected long-term risk/reward might be positive from adding investments in emerging-market bonds and equities, one should expect sharper short-term drawdowns in the USD (and presumably NOK) value of the portfolio in periods of distress due to the added tail risks.

Figure 15. EM basket performance and average pairwise correlation.

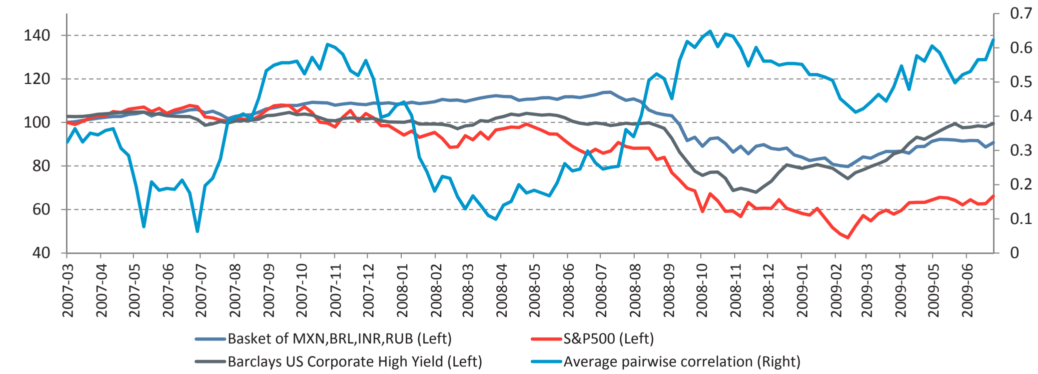

Extension and further work

We do not provide a historical test for whether the implied tail-risk estimates correspond better than the results of a standard model to what is subsequently realised. Further work in this direction might reveal whether we need to include explicit modelling of time-varying risk premiums to get better risk estimates from an option-implied risk model.

The model as described allows for heavy tails for single FX rates. However, as the dependence structure is modelled using a Gaussian copula, strong tail dependencies will not be modelled. Refining the dependence modelling would be an obvious extension to the method.

The generic process described in this note, where we extract implied probability distributions for assets from actively traded option markets and connect them via a measure of dependence, is applicable to other asset classes as well. We are currently working on a similar model based on prices in the option markets for single equities. Option-implied modelling for an equities portfolio introduces an additional set of challenges, such as thinner data coverage (many stocks do not have active option markets) and dimensionality (the Fund's equity portfolio and benchmark combined count over 8,000 stocks). Additionally, there is empirically a strong downside tail dependence within the equity asset class that should be considered.

We have found the risk-modelling insights gained from working with option data to provide additional useful information. Future initiatives may extend our efforts into interest rates, credit and cross-asset modelling as well.

Appendix – Implied volatility from vanna-volga approximation

Recall that the x's are the weights of the three strikes we have implied volatility data for, where the sum of the weights is 1. The C^{BS} function is the standard formula for calculating the value of a currency option. The analytical solution in terms of the weights for the equation set in the main text is:

$$x_1(t;K) = \frac{\nu(t;K)}{\nu(t;K_1)} \frac{\ln\frac{K_2}{K}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_1}}$$
$$x_2(t;K) = \frac{\nu(t;K)}{\nu(t;K_2)} \frac{\ln\frac{K}{K_1}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_2}}$$
$$x_3(t;K) = \frac{\nu(t;K)}{\nu(t;K_3)} \frac{\ln\frac{K}{K_1}\ln\frac{K}{K_2}}{\ln\frac{K_3}{K_2}\ln\frac{K_3}{K_1}}$$

The greeks (the partial derivatives) are as follows

$$\begin{aligned}
\frac{\partial^2 C^{BS}}{\partial\sigma^2}(t;K) &= \frac{\nu(t;K)}{\sigma} d_1(t;K) d_2(t;K) \\
\frac{\partial^2 C^{BS}}{\partial\sigma\partial S_t}(t;K) &= -\frac{\nu(t;K)}{S_t\sigma\sqrt{\tau}} d_2(t;K) \\
d_2(t;K) &= d_1(t;K) - \sigma\sqrt{\tau}
\end{aligned}$$

where

$$\begin{aligned}
\nu(t;K) &= \frac{\partial C^{BS}}{\partial\sigma}(t;K) = S_t \exp^{-r^f\tau}\sqrt{\tau}\varphi(d_1(t;K)) \\
d_1(t;K) &= \frac{\ln\frac{S_t}{K} + (r^d - r^f + \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}} \\
\varphi(x) &= \Phi'(x) = \frac{1}{\sqrt{2\pi}}\exp^{-\frac{1}{2}x^2}
\end{aligned}$$

We can use the fact that we know the implied volatility for the options in the replicating portfolio. This means that the following relation will hold:

$$C(K) = C^{BS}(K) + \sum_{i=1}^{3} x_i(K)\left[C^{MKT}(K_i) - C^{BS}(K_i)\right]$$

C^{MKT} is the market value of the options in the replicating portfolio. C^{BS} is the value of the options calculated with a flat volatility. When we have the calculated the price C(K) it is possible to calculate an approximate implied volatility at the unquoted strike K. The approximation is done using a second-order Taylor expansion, and is done around σ. First we develop the options in the replicating portfolio around σ. Then we develop the unquoted option price $C(K)$ around σ:

$$
\begin{aligned}
&C(K) \approx C^{BS}(K) + \sum_{i=1}^{3} x_i(K)\left[\nu(K_i)(\sigma(K_i)-\sigma) + \frac{1}{2}\frac{\partial^2 C^{BS}}{\partial\sigma^2}(K_i)(\sigma(K_i)-\sigma)^2\right] \\
&C(K) - C^{BS}(K) \approx \nu(K)(\sigma(K)-\sigma) + \frac{1}{2}\frac{\partial^2 C^{BS}}{\partial\sigma^2}(K)(\sigma(K)-\sigma)^2 \\
&\nu(K)(\sigma(K)-\sigma) + \frac{1}{2}\frac{\partial^2 C^{BS}}{\partial\sigma^2}(K)(\sigma(K)-\sigma)^2 \approx \sum_{i=1}^{3} x_i(K)\left[\nu(K_i)(\sigma(K_i)-\sigma) + \frac{1}{2}\frac{\partial^2 C^{BS}}{\partial\sigma^2}(K_i)(\sigma(K_i)-\sigma)^2\right]
\end{aligned}
$$

The last equation is a second-degree equation in $\sigma(K)$. The solution of this equation is:

$$
\begin{aligned}
\sigma(K) &\approx \sigma + \frac{-\sigma + \sqrt{\sigma^2 + d_1(K)d_2(K)(2\sigma D_1(K) + D_2(K))}}{d_1(K)d_2(K)} \\
D_1(K) &= \frac{\ln\frac{K_2}{K}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_1}}\sigma_{25\Delta p} + \frac{\ln\frac{K}{K_1}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_2}}\sigma_{ATM} + \frac{\ln\frac{K}{K_1}\ln\frac{K}{K_2}}{\ln\frac{K_3}{K_2}\ln\frac{K_3}{K_1}}\sigma_{25\Delta c} \\
D_2(K) &= \frac{\ln\frac{K_2}{K}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_1}}d_1(K)d_2(K)(\sigma_{25\Delta p}-\sigma)^2 + \frac{\ln\frac{K}{K_1}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_2}}d_1(K)d_2(K)(\sigma_{ATM}-\sigma)^2 \\
&+ \frac{\ln\frac{K}{K_1}\ln\frac{K}{K_2}}{\ln\frac{K_3}{K_2}\ln\frac{K_3}{K_1}}d_1(K)d_2(K)(\sigma_{25\Delta c}-\sigma)^2
\end{aligned}
$$

which is the vanna-volga expression used.

References

Bank for International Settlements (2010): "Derivatives in Emerging Markets", *BIS Quarterly Review*, December.

Bank of England (2012): "Notes on the Bank of England Option-Implied Probability Density Functions", http://www.bankofengland.co.uk/statistics/Documents/impliedpdfs/BofEOptionImpliedProbabilityDensityFunctions.pdf

Breeden, D. and R. Litzenberger (1978): "Prices of State-Contingent Claims Implicit in Option Prices", *Journal of Business*, 51, 4.

Carr, P. and L. Wu (2009): "Variance Risk Premiums", *Review of Financial Studies*, 22, 3.

Fengler, M. (2005): "Arbitrage-Free Smoothing of the Implied Volatility Surface", SFB 649 Discussion Paper 2005-019, http://sfb649.wiwi.hu-berlin.de/papers/pdf/SFB649DP2005-019.pdf.

Haug, E. (1996): "Implisitt korrelasjon i valutamarkedet", BETA 1/96, espenhaug.com.

Higham, N. (2002): "Computing the Nearest Correlation Matrix – A Problem From Finance", *IMA Journal of Numerical Analysis*, 22, 3, 329-343.

Ilmanen, A. (2011): *Expected returns: An Investor's Guide To Harvesting Market Rewards*, John Wiley.

Mercurio, F. and A. Castagna (2005): "Consistent Pricing of FX Options", Banca IMI Working Paper, http://www.fabiomercurio.it/consistentfxsmile.pdf.

Mercurio, F. and A. Castagna (2007): "The vanna-volga method for implied volatilities", *Risk*, 20, 1, 106-111.

NBIM (2011): "On risk premium variation", NBIM Discussion Note 1-2011.

Walter, C. and L. Lopez (2000): "Is Implied Correlation Worth Calculating? Evidence from Foreign Exchange Options and Historical Data", Federal Reserve Bank of San Francisco Working Paper Series 2000-02, http://www.frbsf.org/econrsrch/workingp/2000/wp00-02.pdf.

Norges Bank Investment Management (NBIM)
Bankplassen 2
Postboks 1179 Sentrum
N-0107 Oslo

Tel.: +47 24 07 30 00
www.nbim.no