

Foreign Exchange options and the Volatility Smile¹

Reimer Bener and Marije Elkenbracht-Huizing

Companies and institutions increasingly use options to reduce their currency risk. The volatility smile is a crucial phenomenon in the valuation of these options. For banks the volatility smile is an important research topic. Here we give some insight in what it is, where it stems from and how to cope with it.

1. Introduction to the Foreign Exchange vanilla options market and the volatility smile

Volatility smile

Foreign Exchange (FX) European vanilla options are valued with the well-known Black Scholes model. The only unobserved input to this model is the volatility. We can also invert the relation and calculate which so-called implied volatility should be used to result in a certain price. If all Black-Scholes assumptions would hold the implied volatility would be the same for all European vanilla options on a specific underlying FX rate. In reality we will find different implied volatilities for different strikes and maturities. In fact, all assumptions of the standard Black-Scholes model that do not hold express themselves in the so-called implied volatility surface. Thus, the Black-Scholes model effectively acts as a quotation convention. An example of an implied volatility surface is given in figure 1.

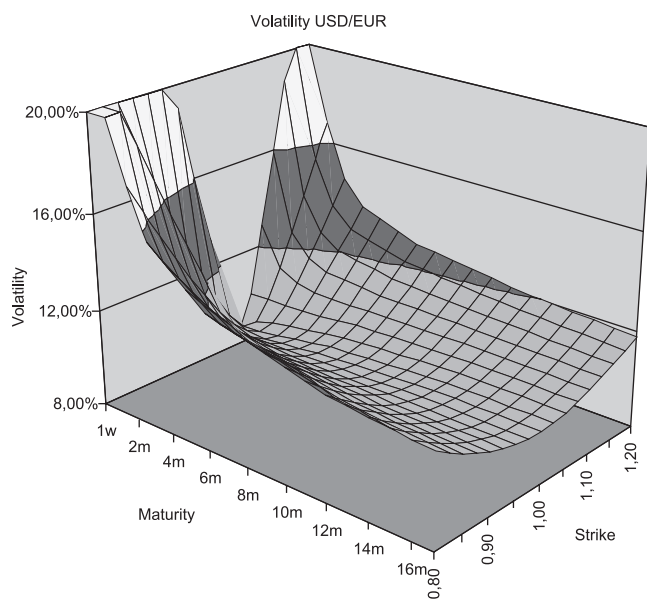


Figure 1: EUR/USD Implied Volatility Surface 16-jan-03.

When we regard implied volatilities for a specific maturity only, one generally encounters shapes as plotted in figure 2. These are known as volatility smile, skew and frown respectively.

In stock markets one generally encounters the volatility skew. The main reasons given for this are the expectation that the volatility will go up in a downward market and vice versa, and the concern for market crashes. In FX markets one mainly encounters the volatility smile. Here the smile can be seen as the general believe that returns are not expected to be lognormal, but will be more extreme in reality. The skew reflects that one currency is expected to appreciate against the other currency. The volatility frown is a very unusual shape. It corresponds with an expected distribution of returns with less weight in the tails, i.e. less kurtosis, than the lognormal distribution. One can encounter this shape with pegged currencies.

Volatility quotation

In figure 4 we give an example of how in the FX market implied volatilities are quoted².

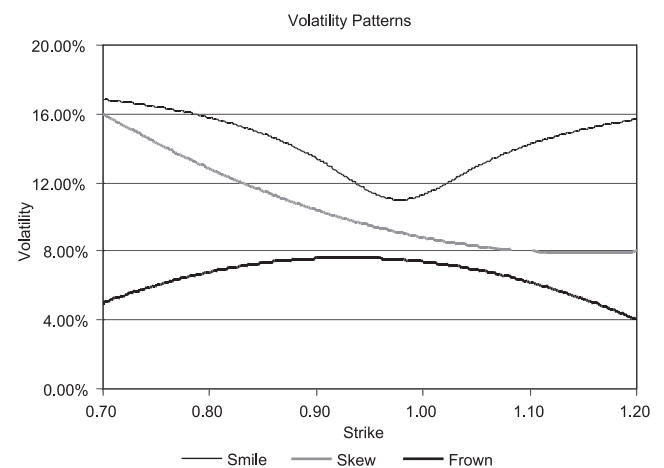


Figure 2: Implied volatility smile, skew and frown.

GBP/USD Spot 1.6459

	Vols	25 Delta Strangles	25 Delta Risk Revs
1W	8.08		
1M	8.18	0.19	0.59
3M	8.26	0.21	0.52
6M	8.38		
1Y	8.48	0.22	0.41
2Y	8.67		

Figure 4: Example of volatility quotation.

We find below “Vols” the volatilities to be used for At-the-Money³ (ATM) options of various maturities. Furthermore, we encounter in this quotation Strangles (STR) and Risk Reversals (RR). A strangle is a long position in an Out-of-the-Money (OTM) call and an OTM put. A strangle is a bet on a large move of the underlying either upwards or downwards. Note that where the ATM indicates the level of the smile, the STR can be regarded as a measure of the curvature or convexity of the volatility smile. A risk reversal is a combination of a long OTM call and a short OTM put. A RR can be seen as a measure of skewness, i.e. the slope of the smile. When RR’s are positive, the market favors the foreign currency.

The implied volatilities correspond to 25-delta and ATM options. Delta is the sensitivity of the option to the spot FX rate and is always between 0% and 100% of the notional. It can be shown that an ATM option has a delta around 50%. A 25-delta call (put) corresponds to option with a strike above (below) the strike of an ATM option. For a precise description of the calculation of delta we refer to the separate section on “Market conventions for calculation of delta”.

A 25-delta RR quote is the difference between the volatility of a 25-delta call and a 25-delta put. A 25-delta STR is equal to the average volatility of a 25-delta call and put minus the ATM volatility. Therefore, the volatility of a 25-delta call and put can be obtained from these quotes as follows

$$v_{C,25} = v_{ATM} + STR_{25} + \frac{1}{2} RR_{25}$$

$$v_{P,25} = v_{ATM} + STR_{25} - \frac{1}{2} RR_{25}$$

Usually quotes also exist for 10-delta RR’s and STR’s, although these options are not as liquid.

2. Volatility interpolation

Clearly to derive valuations for European vanilla options for other delta’s one needs to interpolate between and extrapolate outside the available quotes. But inter- and extrapolation is also required for the derivation of prices for European style derived products, like European digitals. European digitals pay out a fixed amount if the spot at

maturity ends above (or below) the strike and otherwise nothing. Finally inter- and extrapolation is needed when one would like to apply certain methods that try to account for the volatility smile in the valuation of non-European style exotic options. Some of these methods need a continuum of traded strikes and their implied volatilities in order to convert these volatilities to so-called local volatilities ([D], [DK], [Ru], see also Section 4).

2.1 Various interpolation methods

Jackwerth [J] classifies three approaches to inter- and extrapolate volatility quotes. First, fit a function through the option prices. Second, fit a function through the implied volatility quotes. And third, assume a process for the underlying FX rate, fit its parameters to market data and as a result one can calculate all option prices (sometimes analytically but always with Monte Carlo).

Fitting a function through options prices is rarely used. The shape of the function is more difficult to build with standard functions than the implied volatility surface [Ro, S]. Also, one has to take care of certain no-arbitrage conditions.

Of fitting a function through implied volatilities we will describe and compare three methods that are often used: linear, cubic spline and a second-degree polynomial in delta. Furthermore, we will describe in Section 3 how one can also use a mixture of lognormal densities. This can be seen as an example of Jackwerth’s third approach. Note that there are numerous other possibilities, some of which will be mentioned in Section 4.

The call and put delta quotes have to be combined in order to do the interpolation. Therefore, we will construct a horizontal grid that measures the distance to the ATM delta:

$$x = \Delta - \Delta_{ATM}$$

This transformation is illustrated using an example in the table below.

Note that the non-uniqueness of the premium-included call delta (see e.g. figure 3) does not cause any problems for the interpolation, since we will only consider deltas between zero and the At-The-Money (ATM) delta.

Linear

An interpolation technique has to be employed to obtain the volatility corresponding to an arbitrary x , i.e. to

	Δ	x	v
Call	0.10	-0.38	14.00%
Call	0.25	-0.23	13.00%
ATM	0.48	0.00	11.00%
ATM	-0.48	0.00	11.00%
Put	-0.25	0.23	12.00%
Put	-0.10	0.38	13.00%

Table 1: Transformation of put and call delta’s.

Market conventions for calculation of delta

In the FX market implied volatilities are quoted in terms of delta. There are various definitions of delta. Hence, for the correct interpretation of the implied volatility quotes it is important to know what definition is used.

Let us start with stating the Black-Scholes formula for FX European vanilla options:

$$V_0 = \varepsilon(S_0 e^{-r_f T_d} N(\varepsilon d_1) - K e^{-r_d T_d} N(\varepsilon d_2))$$

where,

$$d_1 = \frac{\ln(S_0 / K) + r_d T_d - r_f T_d + v^2 T_e / 2}{v \sqrt{T_e}}, d_2 = d_1 - v \sqrt{T_e}$$

and

S_0 spot FX rate denoted in domestic units per unit of foreign currency

K strike using the same quotation as the spot rate

T_e time from today until expiry of the option

T_d time from spot until delivery of the option

r_d domestic interest rate corresponding with period T_d

r_f foreign interest rate corresponding with period T_d

v volatility corresponding with strike K and period T_e

ε 1 for a call, -1 for a put

$N(\cdot)$ cumulative normal distribution

Hence V_0 is the value of the option expressed in domestic currency on a notional of one unit of foreign currency.

The Black-Scholes delta of the option is equal to

$$\Delta_{BS} = \frac{\partial V_0}{\partial S_0} = \varepsilon \cdot e^{-r_f T_d} N(\varepsilon d_1)$$

In all currency markets, except the eurodollar market, the premium in the foreign currency is included in the delta. This “premium-included” delta has to be calculated as follows

$$\Delta_p = \Delta_{BS} - \frac{V_0}{S_0} = \varepsilon \cdot \frac{K}{S_0} e^{-r_d T_d} N(\varepsilon d_2)$$

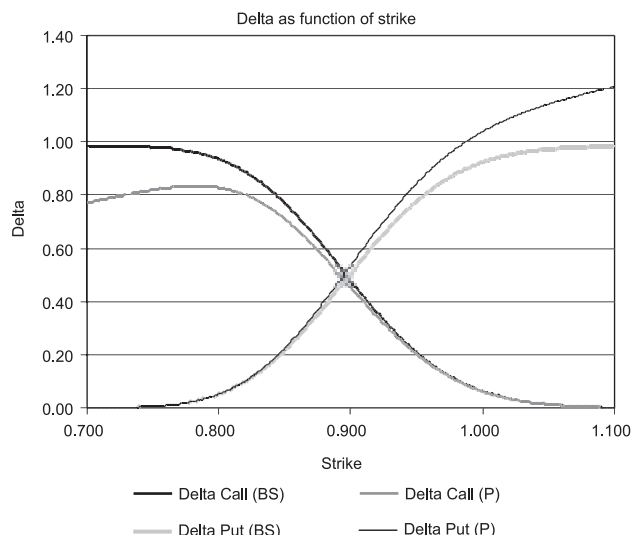


Figure 3: Black-Scholes and premium-included delta as function of strike.

The logic of this premium-included delta can be illustrated with an example. Consider a bank that sells a call on the foreign currency. This option can be delta hedged with an amount of delta of the foreign currency. However, the bank will only have to buy an amount equal to the premium-included delta when it receives the premium in foreign currency.

It can be observed from the above formula that the premium-included delta for a call is not strictly decreasing in strike like the Black-Scholes call delta. Therefore, a premium-included call delta can correspond to two possible strike prices (see figure 3).

For emerging markets (EM) and for maturities of more than two years, it is usual for forward delta's to be quoted. These are defined as follows

$$\Delta_{BS}^F = e^{r_f T_d} \Delta_{BS} \text{ and } \Delta_P^F = e^{r_f T_d} \Delta_P$$

The At-the-Money (ATM) strike refers to the strike of a zero delta straddle, i.e. the strike for which the call delta is equal to the put delta. This strike can be calculated analytically. The table below shows the ATM delta and the ATM strike for each market.

Market		ATM Delta	ATM Strike
EUR/USD maturities < 2y	Δ_{BS}	$\varepsilon \cdot e^{-r_f T_d} / 2$	$S_0 e^{(r_d - r_f) T_d - v^2_{ATM} T_e / 2}$
EUR/USD maturities \geq 2y	Δ_{BS}^F	$\varepsilon / 2$	$S_0 e^{(r_d - r_f) T_d - v^2_{ATM} T_e / 2}$
maturities < 2y	Δ_P	$\varepsilon \cdot e^{-r_d T_d - v^2_{ATM} T_e / 2} / 2$	$S_0 e^{(r_d - r_f) T_d - v^2_{ATM} T_e / 2}$
maturities \geq 2y, EM	Δ_P^F	$\varepsilon \cdot e^{-v^2_{ATM} T_e / 2} / 2$	$S_0 e^{(r_d - r_f) T_d - v^2_{ATM} T_e / 2}$

Table 2: ATM Delta and Strike for different delta definitions.

an arbitrary delta. The easiest method is a linear interpolation between the points. A disadvantage of this method is the sharp unrealistic notch that results at the ATM point. Further, to obtain the volatilities for high absolute values of x (or for low deltas) extrapolation is required. We choose to do this linearly as well.

Cubic Spline

Somewhat more sophisticated is the use of a cubic spline interpolation with linear extrapolation at the bounds. The cubic spline interpolation links the points (x, v) with third degree polynomial functions, such that the first and second derivative at the interior points are continuous. Furthermore, the first derivative at the left (right) bound is chosen to be equal to the slope between the first (last) two points. For an algorithm of this method, we refer to chapter 3 in Numerical Recipes [NR].

Second-degree polynomial in delta

The third method that is considered fits a second-degree polynomial through the three (most liquid) quotes. When the ATM delta is equal to 0.5, this results in the following formula

$$v(\Delta) = v_{ATM} - 2RR_{25}x + 16STR_{25}x^2$$

This formula is derived in [Ma] and is known to be common market practice according to [BMR].

2.2 Converting deltas into strikes

In the previous Section we have described three methods that can be used to calculate the volatility for an arbitrary delta between zero and the ATM with an interpolation method. The volatility corresponding to a certain strike can be computed as follows. First, the strike is compared with the strike of an ATM option. If the strike exceeds the ATM strike a call delta will be used, otherwise a put delta. After that, delta is chosen in a way that it matches the analytic delta calculated with the volatility obtained with interpolation. Thus,

$$h(\Delta) = f(g(\Delta)) - \Delta = 0$$

where,

$f(v)$ The analytic delta corresponding to the volatility v .

$g(\Delta)$ The interpolated volatility corresponding to Δ .

The derivative of $h(\Delta)$ is required to be able to use the Newton-Raphson method. This derivative can be calculated using the chain rule

$$\frac{\partial h(u)}{\partial u} = \frac{\partial f(v)}{\partial v} \bigg|_{v=g(u)} \frac{\partial g(u)}{\partial u} - 1$$

The derivative of the analytic delta with respect to the volatility depends on the definition of delta. When the Black-Scholes forward delta is used, the derivative is equal to

$$\frac{\partial f(v)}{\partial v} = n(d_1) \frac{-\ln(S_0 / K) - (r_d - r_f)T_d + v^2 T_e / 2}{v^2 \sqrt{T_e}}$$

and if the premium is included in delta it is equal to

$$\frac{\partial f(v)}{\partial v} = \frac{K}{S_0} e^{-r_d T_d} n(d_2) \frac{-\ln(S_0 / K) - (r_d - r_f)T_d - v^2 T_e / 2}{v^2 \sqrt{T_e}}$$

Note that when spot deltas are used, the derivative has to be discounted by $e^{-r_f T_d}$.

The derivative of $g(u)$ depends on the interpolation method. For all the discussed interpolation methods, this derivative can easily be derived.

3. Implied density and lognormal mixtures

Implied density

It is possible to derive the so-called implied risk-neutral density function (RND) from vanilla option prices for a given maturity and a continuum of strikes. The implied RND is the probability distribution of the FX spot rate at time T : S_T under the risk neutral probability measure. The implied RND is equal to:

$$\phi(S_T) = e^{r_d T_d} \frac{\partial^2 V_0}{\partial K^2} \bigg|_{K=S_T}$$

The derivation, due to Breeden & Lizenberger [BL], can be found in the appendix. The derivative $\partial^2 V_0 / \partial K^2$ can be calculated with finite differences. When all Black-Scholes assumptions would hold and volatility would be constant over all strikes, the implied RND would look like the “No smile” graph as plotted in figure 5: a perfect log-normal distribution. However, in practice one encounters the “Smile” graph: the implied RND has so-called fat tails and extra density in the center. The higher the volatility in the wings compared with the ATM volatility, the more mass in the tails.

Lognormal mixture

The lognormal mixture approach approximates the terminal probability density of the spot exchange rate with a mixture of lognormal variables. The parameters of the lognormal variables are chosen such that the resulting distribution implies the prices observed in the market. The motiva-




Figure 5: Implied risk-neutral density function with and without smile.




Figure 6: Volatility smile with different interpolation methods.

tion for using lognormal variables is twofold: the distribution stays close to the lognormal one and well-known analytic results can be used. In fact one can see this method as one of the simplest alternative spot processes of which the parameters are calibrated to the available quotes.

As discussed in Section 1 the market of FX options is mainly characterized by three quotes per maturity. A fourth constraint is given by the fact that in the risk-neutral world the expected spot price should be equal to the forward price. Hence, when two lognormals variables with equal probability are used, we need to obtain four parameters. These parameters can be calculated exactly by solving four (non-linear) equations using a numerical scheme (e.g. Broyden's method [NR]). Strictly speaking, there is no guarantee that a solution will be found, but we observed that we were able to solve the equations in realistic market situations.

Comparison

For an interpolation method to be fully satisfactory there are five requirements:

1. The resulting volatilities and prices should fit the market as close as possible.
2. The price function V_0 should be twice differentiable with respect to strike K .
3. The implied RND should be nonnegative.
4. The implied RND should integrate to one.
5. The expected spot price should be equal to forward price with respect to the implied RND.

All our methods satisfy by construction the last two requirements [BH].

In figure 6, we show the results of a comparison between the different methods. In this (fictitious) example the 6 months ATM, RR and STR are 11.0%, 1.0% and 1.5%, respectively.

It can be observed that the smile is more pronounced with the second-degree polynomial and the mixture method, i.e. higher volatilities for low delta's and lower volatilities around the ATM point. The relatively flat linear / spline graph can be explained by the fact that the first derivative at the bounds is equal to the slope between the first / last two points.

Note that the spline interpolation could result in a downward sloping curve, for example, in case the 25-delta put volatility is below the ATM volatility (which can be realistic). To obtain higher volatilities for low delta's with the spline method, one has to include 10-delta quotes. However, these quotes are less reliable because OTM options are not as liquid.

In figure 7, we show the smile in terms of strike. Also, in this graph it can be observed that for very low/high strikes the volatility can be very different depending on the interpolation method. However, to put the results in perspective,

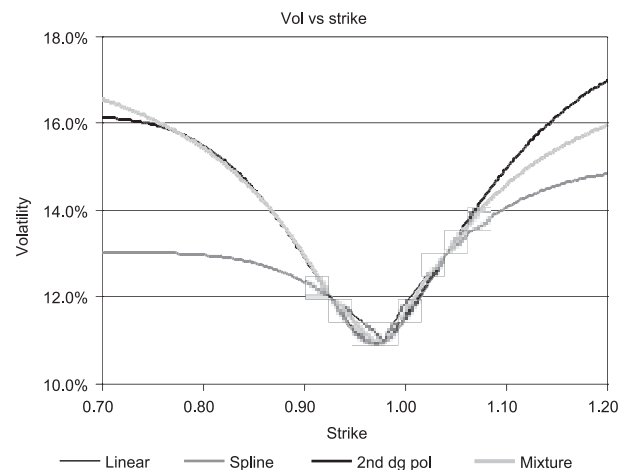


Figure 7: Volatility smile in terms of strike.




Figure 8: Implied density functions.

we note that for deeply OTM or ITM European vanilla options the impact of volatility on the price is limited. That is because low delta options are not very sensitive to the volatility (i.e. these options have low vega). For exotic options however, the differences can be more important.

The implied density functions based on the same example are shown in figure 8. One clearly sees what happens if the price function V_0 is not twice differentiable with respect to strike K : the linear interpolation implies a density with a discontinuity around the ATM point. This is not surprising given the sharp notch that we observe in figure 7 for linear interpolation. The density function with the cubic spline method has a notch at the strikes that correspond to the 25-delta point. For these two methods the implied RND can also become negative. The second-degree polynomial interpolation and the lognormal mixture method both imply a smooth density function that can not become negative. Overall these latter two approaches perform very similar, although there is some empirical evidence that the second-degree polynomial interpolation approach outperforms the lognormal mixture approach ([BP], [C]). For an alternative method and further discussion of this topic we refer to [BH].

4. Volatility smile and exotics

The value of most exotic options is not only determined by the terminal distribution of the spot FX rate, but also by the transition densities. There are various classes of models that all start off with making an assumption of the stochastic process followed by the spot FX rate. Sometimes the volatility is modeled as a stochastic process as well.

Local volatility

One of the earliest models for taking the smile into account, is the local volatility model as developed by [D], [DK] and [Ru]. The tree implementation builds on the original Cox, Ross, and Rubinstein binomial tree [CRR]. The standard spot process is extended as:

$$\frac{dS_t}{S_t} = (r_d - r_f)dt + \sigma(S_t, t)dW$$

where $\sigma(S_t, t)$ is a deterministic process that is fully characterized once the complete continuum of traded strikes and their implied volatilities is available. Main advantage of this model is that it fits per definition the European vanilla prices perfectly. The calibration boils down to an analytic formula for $\sigma(S_t, t)$, where one only has to care about the existence and preferably smoothness of again the second derivative. Main drawback of this method is the unrealistic implied forecast of the volatility surface. E.g. the further one puts oneself in time ahead, the flatter the forecasted volatility surface. This can result in bad hedge performance, sometimes even worse than if one would have used Black Scholes models [HKLW].

Stochastic volatility

Another popular approach is assuming that the volatility is a stochastic process. This kind of modeling is more in line with the real world, where volatility is generally seen as being stochastic. A key result for this type of models is due to Heston [H], who derived a semi-analytic formula for valuing European vanilla options. Such an analytic expression is especially important for calibration purposes.

Stochastic volatility models are known to fit the long end reasonably well, but are less able to fit to the sometimes steep short-end volatilities. Therefore, it is less suitable for short term options. Furthermore, calibration difficulties can result in unstable parameters and hence in hedge problems for longer term options.

Other approaches

There are several other approaches to model the volatility smile, like incorporating jumps (e.g. Bates [B]) or combining local and stochastic volatility [LM]. For a model to be useful one would like:

- the postulated process to match as close as possible the characteristics of the real world, not only for the underlying FX rate, but also for the volatility surface,
- an (semi)-analytic formula for calls and puts for fast calibration,
- good fitting to European vanilla prices; as these are used as hedges they should be priced consistently with the exotic,
- that it is possible to hedge the product throughout its life, so among others the parameters of the model have to be stable.

None of the described approaches (yet) meets all these criteria. Therefore, modeling the volatility smile in the valuation of exotic options is still a subject of ongoing research.

5. Conclusion

The volatility smile is an important phenomenon in the FX options market. In this article we have described the smile, how it is quoted and where it stems from. Next we have described some methods of how one can derive from a few quotes a whole range of option prices and the resulting risk neutral density function. Finally, we have discussed some of the most important models that attempt to account for the volatility smile in the valuation of exotics. It is clear that the valuation of exotics is very much dependent on the implied risk neutral density function. However, unfortunately there is no model yet without any drawback. As differences between various models for the same product can be significant, these models should be treated with utmost care.

Appendix

In this Appendix we prove the result of Breeden & Litzenberger [BL], who have shown that the implied RND is equal to:

$$\phi(S_T) = e^{r_d T_d} \left. \frac{\partial^2 C_0}{\partial K^2} \right|_{K=S_T}$$

We start by expressing the value of European call option in terms of the risk-neutral density function

$$C_0 = e^{-r_d T_d} \int_K^{\infty} (S_T - K) \phi(S_T) dS_T$$

Differentiating this expression once with respect to the strike price gives

$$\frac{\partial C_0}{\partial K} = e^{-r_d T_d} \left(-K \phi(K) - \int_K^{\infty} \phi(S_T) dS_T + K \phi(K) \right) = -e^{-r_d T_d} \int_K^{\infty} \phi(S_T) dS_T$$

Now, differentiating again gives

$$\frac{\partial^2 C_0}{\partial K^2} = e^{-r_d T_d} \phi(K)$$

which proves the results. The result can also be derived for a put option. It is assumed that it is impossible that none of the currencies ever becomes worthless, i.e. $\lim_{S_T \rightarrow 0} \phi(S_T) = \lim_{S_T \rightarrow \infty} \phi(S_T) = 0$.

About the authors

Reimer Beder has studied econometrics at Erasmus University Rotterdam. At present he works part time at ABN AMRO while for the remaining time he is working on his PhD at Erasmus University.

Marije Elkenbracht-Huizing completed a PhD in mathematics at Leiden University. Afterwards, in 1997, she started working for ABN AMRO, where she currently heads the FX and Equity group within the Market Risk Modeling & Product Analysis department.

Notes

- [1] The opinions expressed in this article are those of the authors, and do not necessarily reflect the views of ABN Amro Bank.
- [2] This example is taken from the Reuters pages BBAVOLFIX in which daily volatility fixings are quoted that are composed of input from a number of banks.
- [3] An At-the-Money option has approximately a strike equal to the current forward. For a precise definition of At-the-Money, see the separate section on "Market conventions for calculation of delta".

References

- [BH] Brunner, B., and Hafner, R., *Arbitrage-Free Estimation of the Risk-Neutral Density from the Implied Volatility Smile*, Working Paper, University of Augsburg, June 2002.
- [BL] Breeden, D.T. and R.H. Litzenberger (1978), *Prices of State-contingent Claims Implicit in Option Prices*, Journal of Business 51, No. 4, pp. 621-651, October 1978.
- [BMR] Brigo, D., F. Mercurio and F. Rapisarda, *Lognormal-mixture dynamics and calibration to market volatility smiles*, working paper, Banca IMI.
- [BP] Bliss, R.R., Panigirtzoglou, N., *Testing the stability of implied probability density functions*, Journal of Banking & finance 26 (2002), pp. 381-422.
- [C] Cooper, N., *Testing techniques for estimating implied RNDs from the prices of european-style options*, Proceedings of the BIS workshop on implied PDFs, Bank of International Settlements, Basle, 1999.
- [CRR] Cox, J., Ross, S., and Rubinstein, M., *Option Pricing: A Simplified Approach*, Journal of Financial Economics, 7, No. 3 (1979), pp. 229-263.
- [DK] Derman, E. and Kani, I., *Riding on a smile*, Risk 7, No. 2 (1994), pp. 32-39.
- [D] Dupire, B., *Pricing with a smile*, Risk 7, No. 1 (1994), pp 18-20.
- [J] Jackwerth, J.C., *Option-Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review*, Journal of Derivatives 7(2) (1999), pp 83-100.
- [H] Heston, S.L., *A Closed Form Solution for Options with Stochastic Volatility with Applications to Bonds and Currency Options*, Review of Financial Studies, 6, 2 (1993), pp. 327 - 343.
- [HKLW] Hagan, P.S., D. Kumar, A.S. Lesniewski and D.E. Woodward, *Managing Smile Risk*, July 2002, Wilmott Magazine, pp 84-108.
- [NR] Press, W.H., S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C++*, *The Art of Scientific Computing*, second edition, 2002.
- [LM] Lipton, A., W. McGhee, *Universal Barriers*, Risk May 2002.
- [Ma] Maltz, A.M., *Estimating the probability distribution of the future exchange rate from option prices*, Journal of derivatives 5(2) (1997), pp. 18-36.
- [Me] Merton, R.C., *Option pricing when underlying stock returns are discontinuous*, Journal of Financial Economics, 3 (1976), pp. 125 - 144.
- [Ru] Rubinstein, M., *Implied binomial trees*, Journal of Finance, 49, No. 3 (1994), pp. 771-818.
- [Ro] Rosenberg, J.V., *Implied volatility functions: A reprise*, Journal of derivatives 7(3) (2000), pp. 51-64.
- [S] Shimko, D., *Bounds of probability*, Risk 6, No 4 (1993), pp. 33-37.