The vanna-volga method for implied volatilities

The vanna-volga method is a popular approach for constructing implied volatility curves in the options market. In this article, ***Antonio Castagna*** *and* ***Fabio Mercurio*** *give it both theoretical and practical support by showing its tractability and robustness*

The vanna-volga (VV) method is an empirical procedure that can be used to infer an implied-volatility smile from three available quotes for a given maturity.[1] It is based on the construction of locally replicating portfolios whose associated hedging costs are added to corresponding Black-Scholes prices to produce smile-consistent values. Besides being intuitive and easy to implement, this procedure has a clear financial interpretation, which further supports its use in practice.

The VV method is commonly used in foreign exchange options markets, where three main volatility quotes are typically available for a given market maturity: the delta-neutral straddle, referred to as at-the-money (ATM); the risk reversal for 25Δ call and put; and the (vega-weighted) butterfly with 25Δ wings.[2] The application of VV allows us to derive implied volatilities for any option's delta, in particular for those outside the basic range set by the 25Δ put and call quotes.

In the financial literature, the VV approach was introduced by Lipton & McGhee (2002), who compare different approaches to the pricing of double-no-touch options, and by Wystup (2003), who describes its application to the valuation of one-touch options. However, their analyses are rather informal and mostly based on numerical examples. In this article, instead, we will review the VV procedure in more detail and derive some important results concerning the tractability of the method and its robustness.

We start by describing the replication argument the VV procedure is based on and derive closed-form formulas for the weights in the hedging portfolio to render the smile construction more explicit. We then show that the VV price functional satisfies typical no-arbitrage conditions and test the robustness of the resulting smile by showing that: changing the three initial pairs of strike and volatility consistently eventually produces the same implied-volatility curve; and the VV method, if readapted to price European-style claims, is consistent with static-replication arguments. Finally, we derive first- and second-order approximations for the implied volatilities induced by the VV option price.

Since the VV method provides an efficient tool for interpolating or extrapolating implied volatilities, we also compare it with other popular functional forms, like that of Malz (1997) and that of Hagan *et al* (2002).

All the proofs of the propositions in this article are omitted for brevity. However, they can be found in Castagna & Mercurio (2005).

The VV method: the replicating portfolio

We consider a forex option market where, for a given maturity T, three basic options are quoted: the 25Δ put, the ATM and the 25Δ call. We denote the corresponding strikes by K_i, $i = 1, 2, 3$, $K_1 < K_2 < K_3$, and set $\mathcal{K} := \{K_1, K_2, K_3\}$.[3] The market-implied volatility associated with K_1 is denoted by σ_i, $i = 1, 2, 3$.

The VV method serves the purpose of defining an implied-volatility smile that is consistent with the basic volatilities σ_i. The rationale behind it stems from a replication argument in a flat-smile world where the constant level of implied volatility varies stochastically over time. This argument is presented hereafter, where for simplicity we consider the same type of options, namely calls.

It is well known that in the Black-Scholes (BS) model, the payout of a European-style call with maturity T and strike K can be replicated by a dynamic Δ-hedging strategy, whose value (including the bank account part) matches, at every time t, the option price $C^{BS}(t; K)$ given by:

$$
\begin{aligned}
C^{BS}(t;K) &= S_t e^{-r^f \tau} \Phi\left(\frac{\ln \frac{S_t}{K} + \left(r^d - r^f + \frac{1}{2}\sigma^2 \right)\tau}{\sigma\sqrt{\tau}} \right) \\
&\quad - K e^{-r^d \tau} \Phi\left(\frac{\ln \frac{S_t}{K} + \left(r^d - r^f - \frac{1}{2}\sigma^2 \right)\tau}{\sigma\sqrt{\tau}} \right)
\end{aligned}
\tag{1}
$$

where S_t denotes the exchange rate at time t, $\tau := T - t$, r^d and r^f denote, respectively, the domestic and foreign risk-free rates, and

[1] *The terms vanna and volga are commonly used by practitioners to denote the partial derivatives ∂Vega/∂Spot and ∂Vega/∂Vol of an option's vega with respect to the underlying asset and its volatility, respectively. The reason for naming the procedure this way will be clear below*

[2] *We drop the '%' sign after the level of the Δ, in accordance with market jargon. Therefore, a 25Δ call is a call whose delta is 0.25. Analogously, a 25Δ put is one whose delta is –0.25*

[3] *For the exact definition of strikes K_1, K_2 and K_3, we refer to Bisesti, Castagna & Mercurio (2005), where a thorough description of the main quotes in a forex option market is also provided*

[4] *The option price when the underlying asset is an exchange rate was in fact derived by Garman & Kohlhagen (1983). Since their formula follows from the BS assumption, we prefer to state we are using the BS model, also because the VV method can in principle be applied to other underlyings*

σ is the constant BS implied volatility.[4] In real financial markets, however, volatility is stochastic and traders hedge the associated risk by constructing portfolios that are vega-neutral in a BS (flat-smile) world.

Maintaining the assumption of flat but stochastic implied volatilities, the presence of three basic options in the market makes it possible to build a portfolio that zeros out partial derivatives up to the second order. In fact, denoting respectively by Δ_t and x_i the units of the underlying asset and options with strikes K_i held at time t and setting $C_i^{BS}(t) = C^{BS}(t; K_i)$, under diffusion dynamics for both S_t and $\sigma = \sigma_t$, we have by Itô's lemma:

$$
\begin{aligned}
& dC^{BS}(t;K) - \Delta_t dS_t - \sum_{i=1}^{3} x_i dC_i^{BS}(t) \\
&= \left[\frac{\partial C^{BS}(t;K)}{\partial t} - \sum_{i=1}^{3} x_i \frac{\partial C_i^{BS}(t)}{\partial t}\right] dt \\
&+ \left[\frac{\partial C^{BS}(t;K)}{\partial S} - \Delta_t - \sum_{i=1}^{3} x_i \frac{\partial C_i^{BS}(t)}{\partial S}\right] dS_t \\
&+ \left[\frac{\partial C^{BS}(t;K)}{\partial \sigma} - \sum_{i=1}^{3} x_i \frac{\partial C_i^{BS}(t)}{\partial \sigma}\right] d\sigma_t \\
&+ \frac{1}{2}\left[\frac{\partial^2 C^{BS}(t;K)}{\partial S^2} - \sum_{i=1}^{3} x_i \frac{\partial^2 C_i^{BS}(t)}{\partial S^2}\right] (dS_t)^2 \\
&+ \frac{1}{2}\left[\frac{\partial^2 C^{BS}(t;K)}{\partial \sigma^2} - \sum_{i=1}^{3} x_i \frac{\partial^2 C_i^{BS}(t)}{\partial \sigma^2}\right] (d\sigma_t)^2 \\
&+ \left[\frac{\partial^2 C^{BS}(t;K)}{\partial S \partial \sigma} - \sum_{i=1}^{3} x_i \frac{\partial^2 C_i^{BS}(t)}{\partial S \partial \sigma}\right] dS_t d\sigma_t
\end{aligned}
\tag{2}
$$

Choosing Δ_t and x_i so as to zero out the coefficients of dS_t, $d\sigma_t$, $(d\sigma_t)^2$ and $dS_t d\sigma_t$,[5] the portfolio comprises a long position in the call with strike K, and short positions in x_i calls with strike K_i and short the amount Δ_t of the underlying, and is locally risk-free at time t, in that no stochastic terms are involved in its differential[6]:

$$
\begin{aligned}
& dC^{BS}(t;K) - \Delta_t dS_t - \sum_{i=1}^{3} x_i dC_i^{BS}(t) \\
&= r^d \left[C^{BS}(t;K) - \Delta_t S_t - \sum_{i=1}^{3} x_i C_i^{BS}(t)\right] dt
\end{aligned}
\tag{3}
$$

Therefore, when volatility is stochastic and options are valued with the BS formula, we can still have a (locally) perfect hedge, provided that we hold suitable amounts of three more options to rule out the model risk. (The hedging strategy is irrespective of the true asset and volatility dynamics, under the assumption of no jumps.)

■ **Remark 1.** The validity of the previous replication argument may be questioned because no stochastic-volatility model can produce implied volatilities that are flat and stochastic at the same time. The simultaneous presence of these features, though inconsistent from a theoretical point of view, can however be justified on empirical grounds. In fact, the practical advantages of the BS paradigm are so clear that many forex option traders run their books by revaluing and hedging according to a BS flat-smile model, with the ATM volatility being continuously updated to the actual market level.[7]

The first step in the VV procedure is the construction of the above hedging portfolio, whose weights x_i are explicitly calculated in the following section.

Calculating the VV weights

We assume hereafter that the constant BS volatility is the ATM one, thus setting $\sigma = \sigma_2$ $(= \sigma_{ATM})$. We also assume that $t = 0$, dropping accordingly the argument t in the call prices. From equation (2), we see that the weights $x_1 = x_1(K)$, $x_2 = x_2(K)$ and $x_3 = x_3(K)$, for which the resulting portfolio of European-style calls with maturity T and strikes K_1, K_2 and K_3 has the same vega, ∂Vega/∂Vol and ∂Vega/∂Spot as the call with strike K,[8] can be found by solving the following system:

$$
\begin{aligned}
\frac{\partial C^{BS}}{\partial \sigma}(K) &= \sum_{i=1}^{3} x_i(K) \frac{\partial C^{BS}}{\partial \sigma}(K_i) \\
\frac{\partial^2 C^{BS}}{\partial \sigma^2}(K) &= \sum_{i=1}^{3} x_i(K) \frac{\partial^2 C^{BS}}{\partial \sigma^2}(K_i) \\
\frac{\partial^2 C^{BS}}{\partial \sigma \partial S_0}(K) &= \sum_{i=1}^{3} x_i(K) \frac{\partial^2 C^{BS}}{\partial \sigma \partial S_0}(K_i)
\end{aligned}
\tag{4}
$$

Denoting by $\mathcal{V}(K)$ the vega of a European-style option with maturity T and strike K:

$$
\begin{aligned}
\mathcal{V}(K) &= \frac{\partial C^{BS}}{\partial \sigma}(K) = S_0 e^{-r^f T} \sqrt{T} \varphi(d_1(K)) \\
d_1(K) &= \frac{\ln \frac{S_0}{K} + \left(r^d - r^f + \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}
\end{aligned}
\tag{5}
$$

where $\varphi(x) = \Phi'(x)$ is the normal density function, and calculating the second-order derivatives:

$$
\begin{aligned}
\frac{\partial^2 C^{BS}}{\partial \sigma^2}(K) &= \frac{\mathcal{V}(K)}{\sigma} d_1(K) d_2(K) \\
\frac{\partial^2 C^{BS}}{\partial \sigma \partial S_0}(K) &= -\frac{\mathcal{V}(K)}{S_0 \sigma \sqrt{T}} d_2(K) \\
d_2(K) &= d_1(K) - \sigma\sqrt{T}
\end{aligned}
$$

we can prove the following.

■ **Proposition 1.** The system (4) admits a unique solution, which is given by:

$$
\begin{aligned}
x_1(K) &= \frac{\mathcal{V}(K)}{\mathcal{V}(K_1)} \frac{\ln \frac{K_2}{K} \ln \frac{K_3}{K}}{\ln \frac{K_2}{K_1} \ln \frac{K_3}{K_1}} \\
x_2(K) &= \frac{\mathcal{V}(K)}{\mathcal{V}(K_2)} \frac{\ln \frac{K}{K_1} \ln \frac{K_3}{K}}{\ln \frac{K_2}{K_1} \ln \frac{K_3}{K_2}} \\
x_3(K) &= \frac{\mathcal{V}(K)}{\mathcal{V}(K_3)} \frac{\ln \frac{K}{K_1} \ln \frac{K}{K_2}}{\ln \frac{K_3}{K_1} \ln \frac{K_3}{K_2}}
\end{aligned}
\tag{6}
$$

In particular, if $K = K_j$ then $x_i(K) = 1$ for $i = j$ and zero otherwise.

The VV option price

We can now proceed to the definition of an option price that is consistent with the market prices of the basic options.

[5] The coefficient of $(dS_t)^2$ will be zeroed accordingly, due to the relation linking an option's gamma and vega in the BS world

[6] We also use the BS partial differential equation

[7] 'Continuously' typically means a daily or slightly more frequent update

[8] This explains the name assigned to the smile-construction procedure, given the meaning of the terms vanna and volga

1 Implied volatilities plotted against strikes and put deltas (in absolute value), calibrated to the three basic euro/dollar quotes and compared with the 10Δ (put and call) volatilities

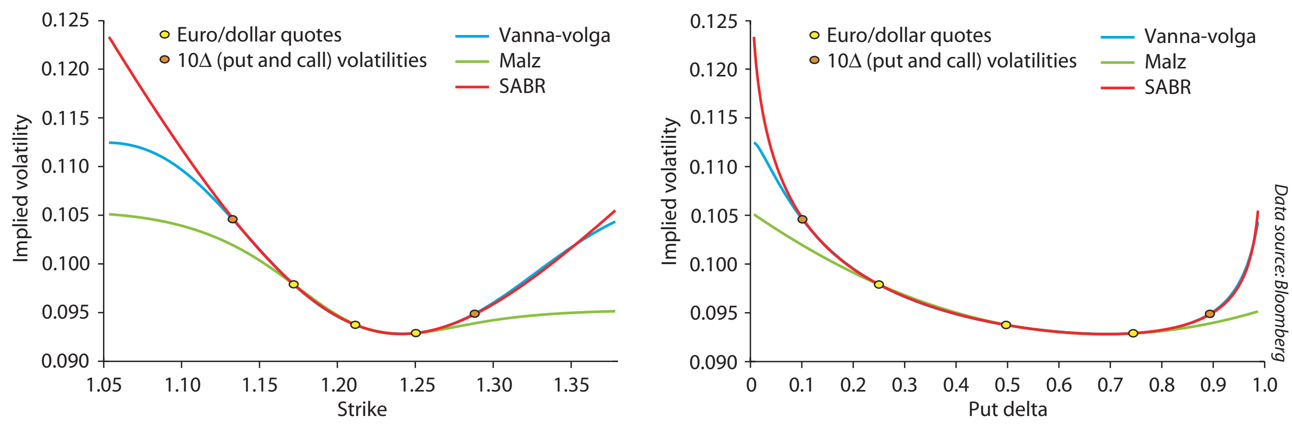

The above replication argument shows that a portfolio comprising $x_i(K)$ units of the option with strike K_i (and Δ_0 units of the underlying asset) gives a local perfect hedge in a BS world. The hedging strategy, however, has to be implemented at prevailing market prices, which generates an extra cost with respect to the BS value of the options portfolio. Such a cost has to be added to the BS price (1), with $t = 0$, to produce an arbitrage-free price that is consistent with the quoted option prices $C^{MKT}(K_1)$, $C^{MKT}(K_2)$ and $C^{MKT}(K_3)$.

In fact, in case of a short maturity, that is, for a small T, equation (3) can be approximated as:

$$\begin{aligned}&(S_T - K)^+ - C^{BS}(K) - \Delta_0[S_T - S_0]\\&-\sum_{i=1}^{3} x_i\left[(S_T - K_i)^+ - C^{BS}(K_i)\right]\\&\approx r^d\left[C^{BS}(K) - \Delta_0 S_0 - \sum_{i=1}^{3} x_i C^{BS}(K_i)\right]T\end{aligned}$$

so that setting:

$$C(K) = C^{BS}(K) + \sum_{i=1}^{3} x_i(K)\left[C^{MKT}(K_i) - C^{BS}(K_i)\right] \tag{7}$$

we have:

$$\begin{aligned}&(S_T - K)^+ \approx C(K) + \Delta_0[S_T - S_0]\\&+\sum_{i=1}^{3} x_i\left[(S_T - K_i)^+ - C^{MKT}(K_i)\right]\\&+r^d\left[C(K) - \Delta_0 S_0 - \sum_{i=1}^{3} x_i C^{MKT}(K_i)\right]T\end{aligned}$$

Therefore, when actual market prices are considered, the option payout $(S_T - K)^+$ can still be replicated by buying Δ_0 units of the underlying asset and x_i options with strike K_i (investing the resulting cash at rate r^d), provided one starts from the initial endowment $C(K)$.

The quantity $C(K)$ in (7) is thus defined as the VV option's premium, implicitly assuming that the replication error is also negligible for longer maturities. Such a premium equals the BS price $C^{BS}(K)$ plus the cost difference of the hedging portfolio induced by the market implied volatilities with respect to the constant volatility σ. Since we set $\sigma = \sigma_2$, the market volatility for strike K_2, (7) can be simplified to:

$$\begin{aligned}C(K) &= C^{BS}(K) + x_1(K)\left[C^{MKT}(K_1) - C^{BS}(K_1)\right]\\&+x_3(K)\left[C^{MKT}(K_3) - C^{BS}(K_3)\right]\end{aligned}$$

Remark 2. Expressing the system (4) in the form $b = Ax$ and setting $c = (c_1, c_2, c_3)'$, where $c_i := C^{MKT}(K_i) - C^{BS}(K_i)$, and $y = (y_1, y_2, y_3)' := (A')^{-1}c$, we can also write:

$$\begin{aligned}C(K) &= C^{BS}(K) + y_1\frac{\partial C^{BS}}{\partial \sigma}(K)\\&+y_2\frac{\partial^2 C^{BS}}{\partial \sigma^2}(K) + y_3\frac{\partial^2 C^{BS}}{\partial \sigma \partial S_0}(K)\end{aligned}$$

The difference between the VV and BS prices can thus be interpreted as the sum of the option's vega, ∂Vega/∂Vol and ∂Vega/∂Spot, weighted by their respective hedging cost y.[9] Besides being quite intuitive, this representation also has the advantage that the weights y are independent of the strike K and, as such, can be calculated once for all. However, we prefer to stick to the definition (7), since it allows an easier derivation of our approximations below.

The VV option price has several interesting features that we analyse in the following.

When $K = K_j$, $C(K_j) = C^{MKT}(K_j)$, since $x_i(K_j) = 1$ for $i = j$ and zero otherwise. Therefore, (7) defines a rule for either interpolating or extrapolating prices from the three option quotes $C^{MKT}(K_1)$, $C^{MKT}(K_2)$ and $C^{MKT}(K_3)$.

The option price $C(K)$, as a function of the strike K, is twice differentiable and satisfies the following (no-arbitrage) conditions:

$$\begin{aligned}&\lim_{K\to 0^+} C(K) = S_0 e^{-r^f T} \text{ and } \lim_{K\to +\infty} C(K) = 0\\&\lim_{K\to 0^+} \tfrac{dC}{dK}(K) = -e^{-r^d T} \text{ and } \lim_{K\to +\infty} K\tfrac{dC}{dK}(K) = 0\end{aligned}$$

These properties, which are trivially satisfied by $C^{BS}(K)$, follow from the fact that, for each i, both $x_i(K)$ and $dx_i(K)/dK$ go to zero for $K \to 0^+$ or $K \to +\infty$.

To avoid arbitrage opportunities, the option price $C(K)$ should also be a convex function of the strike K, that is, $(d^2C)/(dK^2)(K)$

[9] *We thank one of the referees for suggesting this alternative formulation*

> 0 for each $K > 0$. This property, which is not true in general[10], holds however for typical market parameters, so that (7) leads to prices that are arbitrage-free in practice.

The VV implied-volatility curve $K \to \varsigma(K)$ can be obtained by inverting (7), for each considered K, through the BS formula. An example of such a curve is shown in figure 1. Since, by construction, $\varsigma(K_i) = \sigma_i$, the function $\varsigma(K)$ yields an interpolation-extrapolation tool for the market implied volatilities.

Comparison with other interpolation rules

Contrary to other interpolation schemes proposed in the financial literature, the VV pricing formula (7) has several advantages: it has a clear financial rationale supporting it, based on the hedging argument leading to its definition; it allows for an automatic calibration to the main volatility data, being an explicit function of σ_1, σ_2, σ_3; and it can be extended to any European-style derivative (see our second consistency result below). To our knowledge, no other functional form enjoys the same features.

Compared, for example, with the second-order polynomial function (in Δ) proposed by Malz (1997), the interpolation (7) equally perfectly fits the three points provided, but, in accordance with typical market quotes, boosts the volatility value both for low- and high-put deltas. A graphical comparison, based on market data, between the two functional forms is shown in figure 1, where their difference at extreme strikes is clearly highlighted.

The interpolation (7) also yields a very good approximation of the smile induced, after calibration to strikes K_i, by the most renowned stochastic-volatility models in the financial literature, especially within the range $[K_1, K_3]$. This is not surprising, since the three strikes provide information on the second, third and fourth moments of the marginal distribution of the underlying asset, so that models agreeing on these three points are likely to produce very similar smiles. As a confirmation of this statement, in figure 1, we also consider the example of the SABR functional form of Hagan *et al* (2002), which has become a standard in the market as far as the modelling of implied volatilities is concerned. The SABR and VV curves tend to agree quite well in the range set by the two 10Δ options (in the given example they almost overlap), typically departing from each other only for illiquid strikes. The advantage of using the VV interpolation is that no calibration procedure is involved, since σ_1, σ_2, σ_3 are direct inputs of formula (7).

A. Discount factors: July 1, 2005

Expiry	Dollar	Euro
3M: October 3, 2005	0.9902752	0.9945049

B. Strikes and volatilities corresponding to the three main deltas: July 1, 2005

Delta	Strike	Volatility
25D put	1.1720	9.79%
At-the-money	1.2115	9.375%
25D call	1.2504	9.29%

In figure 1, we compare the volatility smiles yielded by the VV price (7), the Malz (1997) quadratic interpolation and the SABR functional form[11], plotting the respective implied volatilities both against strikes and put deltas. The three plots are obtained after calibration to the three basic quotes σ_1, σ_2, σ_3, using the following euro/dollar data as of July 1, 2005 (provided by Bloomberg): $T = 3M$,[12] $S_0 = 1.205$, $\sigma_1 = 9.79\%$, $\sigma_2 = 9.375\%$, $\sigma_3 = 9.29\%$, $K_1 = 1.1720$, $K_2 = 1.2115$ and $K_3 = 1.2504$ (see also tables A and B).

Once the three functional forms are calibrated to the liquid quotes σ_1, σ_2, σ_3, one may then compare their values at extreme strikes with the corresponding quotes that may be provided by brokers or market-makers. To this end, in figure 1, we also report the implied volatilities of the 10Δ put and call options (respectively equal to 10.46% and 9.49%, again provided by Bloomberg) to show that the Malz (1997) quadratic function is typically not consistent with the quotes for strikes outside the basic interval $[K_1, K_3]$.

[10] *One can actually find cases where the inequality is violated for some strike K*
[11] *We fix the SABR β parameter to 0.6. Other values of β produce, anyway, quite similar calibrated volatilities*
[12] *To be precise, on that date the three-month expiry counted 94 days*

| Derivatives

A ticking
time-bomb?

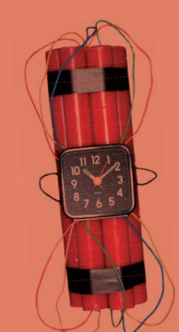

Market leaders
read Risk
Do you?

Risk
risk.net

2 Euro/dollar implied volatilities and their approximations plotted against strikes and put deltas (in absolute value)

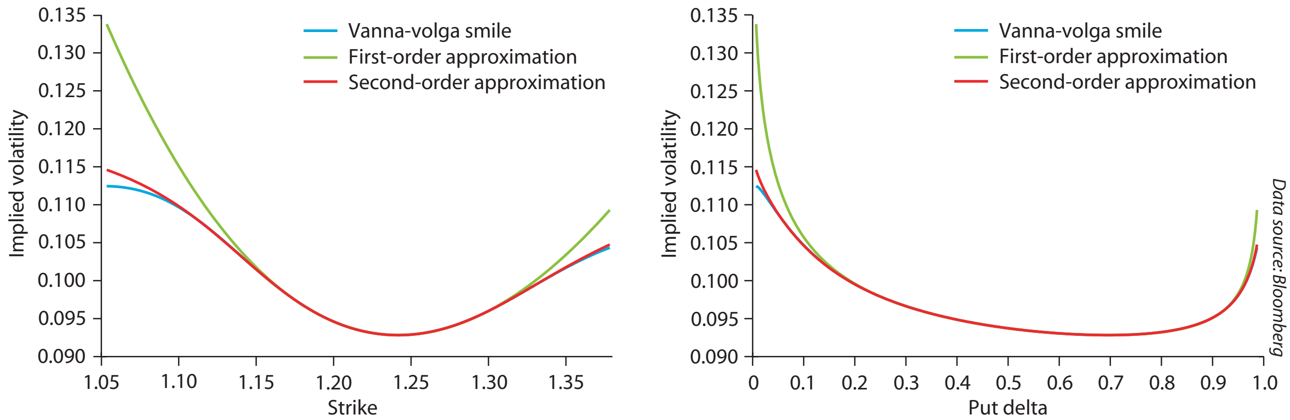

Two consistency results for the VV price

We now state two important consistency results that hold for the option price (7) and that give further support to the VV procedure.

The first result is as follows. One may wonder what happens if we apply the VV curve construction method when starting from three other strikes whose associated prices coincide with those coming from formula (7). Clearly, for the procedure to be robust, we would want the two curves to exactly coincide. This is indeed the case.

In fact, consider a new set of strikes $\mathcal{H} := \{H_1, H_2, H_3\}$, for which we set:

$$\begin{aligned} C^{\mathcal{H}}(H_i) &= C^{\mathcal{K}}(H_i) \\ &= C^{BS}(H_i) + \sum_{j=1}^{3} x_j(H_i)\left[C^{MKT}(K_j) - C^{BS}(K_j)\right] \end{aligned} \tag{8}$$

where the superscripts $\mathcal{H}$ and $\mathcal{K}$ highlight the set of strikes the pricing procedure is based on, and weights x_j are obtained from $\mathcal{K}$ with formulas (6). For a generic strike K, denoting by $x_i(K; \mathcal{H})$ the weights for K that are derived starting from the set $\mathcal{H}$, the option price associated to $\mathcal{H}$ is defined, analogously to (7), by:

$$C^{\mathcal{H}}(K) = C^{BS}(K) + \sum_{j=1}^{3} x_j(K;\mathcal{H})\left[C^{\mathcal{H}}(H_j) - C^{BS}(H_j)\right]$$

where the second term in the sum is now not necessarily zero since H_2 is in general different from K_2. The following proposition states the desired consistency result.

■ **Proposition 2.** The call prices based on $\mathcal{H}$ coincide with those based on $\mathcal{K}$, namely, for each strike K:

$$C^{\mathcal{H}}(K) = C^{\mathcal{K}}(K) \tag{9}$$

A second consistency result that can be proven for the option price (7) concerns the pricing of European-style derivatives and their static replication. To this end, assume that $h(x)$ is a real function that is defined for $x \in [0, \infty)$, is well behaved at infinity and is twice differentiable. Given the simple claim with payout $h(S_T)$ at time T, we denote by V its price at time zero, when taking into account the whole smile of the underlying at time T. By Carr & Madan (1998), we have:

$$V = e^{-r^d T} h(0) + S_0 e^{-r^f T} h'(0) + \int_0^{+\infty} h''(K) C(K)\, dK$$

The same reasoning adopted above (see 'The VV method: the replicating portfolio') with regard to the local hedge of the call with strike K can also be applied to the general payout $h(S_T)$. We can thus construct a portfolio of European-style calls with maturity T and strikes K_1, K_2 and K_3, such that the portfolio has the same vega, ∂Vega/∂Vol and ∂Vega/∂Spot as the given derivative. Denoting by V^{BS} the claim price under the BS model, this is achieved by finding the corresponding portfolio weights x_1^h, x_2^h and x_3^h, which are always unique (see Proposition 1). We can then define a new (smile-consistent) price for our derivative as:

$$\bar{V} = V^{BS} + \sum_{i=1}^{3} x_i^h\left[C^{MKT}(K_i) - C^{BS}(K_i)\right] \tag{10}$$

which is the obvious generalisation of (7). Our second consistency result is stated in the following.

■ **Proposition 3.** The claim price that is consistent with the option prices (7) is equal to the claim price that is obtained by adjusting its BS price by the cost difference of the hedging portfolio when using market prices $C^{MKT}(K_i)$ instead of the constant-volatility prices $C^{BS}(K_i)$. In formulas:

$$V = \bar{V}$$

Therefore, if we calculate the hedging portfolio for the claim under flat volatility and add to the BS claim price the cost difference of the hedging portfolio (market price minus constant-volatility price), obtaining $\bar{V}$, we exactly retrieve the claim price V as obtained through the risk-neutral density implied by the call option prices that are consistent with the market smile.[13]

As an example of a possible application of this result, Castagna & Mercurio (2005) consider the specific case of a quanto option, showing that its pricing can be achieved by using the three basic options only and not the virtually infinite range that is necessary when using static replication arguments.

An approximation for implied volatilities

The specific expression of the VV option price, combined with our analytical formula (6) for the weights, allows for the derivation of a straightforward approximation for the VV implied volatility $\varsigma(K)$, by expanding both members of (7) at first order

[13] *Different but equivalent expressions for such a density can be found in Castagna & Mercurio (2005) and Beneder & Baker (2005)*

in $\sigma = \sigma_2$. We obtain:

$$\varsigma(K) \approx \varsigma_1(K) := \frac{\ln\frac{K_2}{K}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_1}}\sigma_1 + \frac{\ln\frac{K}{K_1}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_2}}\sigma_2 + \frac{\ln\frac{K}{K_1}\ln\frac{K}{K_2}}{\ln\frac{K_3}{K_1}\ln\frac{K_3}{K_2}}\sigma_3 \qquad (11)$$

The implied volatility $\varsigma(K)$ can thus be approximated by a linear combination of the basic volatilities σ_i, with coefficients that add up to one (as tedious but straightforward algebra shows). It is also easily seen that the approximation is a quadratic function of $\ln K$, so that one can resort to a simple parabolic interpolation when log co-ordinates are used.

A graphical representation of the accuracy of the approximation (11) is shown in figure 2, where we use the same euro/dollar data as for figure 1. The approximation (11) is extremely accurate inside the interval $[K_1, K_3]$. The wings, however, tend to be overvalued. In fact, being the quadratic functional form in the log-strike, the no-arbitrage conditions derived by Lee (2004) for the asymptotic value of implied volatilities are violated. This drawback is addressed by a second, more precise, approximation, which is asymptotically constant at extreme strikes, and is obtained by expanding both members of (7) at second order in $\sigma = \sigma_2$:

$$\varsigma(K) \approx \varsigma_2(K) := \sigma_2 + \frac{-\sigma_2 + \sqrt{\sigma_2^2 + d_1(K)d_2(K)\left(2\sigma_2 D_1(K) + D_2(K)\right)}}{d_1(K)d_2(K)} \qquad (12)$$

where:

$$D_1(K) := \varsigma_1(K) - \sigma_2$$

$$D_2(K) := \frac{\ln\frac{K_2}{K}\ln\frac{K_3}{K}}{\ln\frac{K_2}{K_1}\ln\frac{K_3}{K_1}} d_1(K_1)d_2(K_1)(\sigma_1 - \sigma_2)^2 + \frac{\ln\frac{K}{K_1}\ln\frac{K}{K_2}}{\ln\frac{K_3}{K_1}\ln\frac{K_3}{K_2}} d_1(K_3)d_2(K_3)(\sigma_3 - \sigma_2)^2$$

As we can see from figure 2, the approximation (12) is also extremely accurate in the wings, even for extreme values of put deltas. Its only drawback is that it may not be defined, due to the presence of a square-root term. The radicand, however, is positive in most practical applications.

Conclusions

We have described the VV approach, an empirical procedure to construct implied volatility curves in the forex market. We have seen that the procedure leads to a smile-consistent pricing formula for any European-style contingent claim. We have also compared the VV option prices with those coming from other functional forms known in the financial literature. We have then shown consistency results and proposed efficient approximations for the VV implied volatilities.

The VV smile-construction procedure and the related pricing formula are rather general. In fact, even though they have been developed for forex options, they can be applied in any market where at least three reliable volatility quotes are available for a given maturity. The application also seems quite promising in other markets, where European-style exotic payouts are more common than in the forex market. Another possibility is the interest rate market, where CMS convexity adjustments can be calculated by combining the VV price functional with the replication argument in Mercurio & Pallavicini (2006).

A last, unsolved issue concerns the valuation of path-dependent exotic options by means of a generalisation of the empirical procedure that we have illustrated in this article. This is, in general, a quite complex issue to deal with, considering also that the quoted implied volatilities only contain information on marginal densities, which is of course not sufficient for valuing path-dependent derivatives. For exotic claims, *ad hoc* procedures are usually used. For instance, barrier option prices can be obtained by weighing the cost difference of the 'replicating' strategy by the (risk-neutral) probability of not crossing the barrier before maturity (see Lipton & McGhee (2002) and Wystup (2003) for a description of the procedure for one-touch and double-no-touch options, respectively). However, not only are such adjustments harder to justify theoretically than those in the plain vanilla case, but, from a practical point of view, they can even have the opposite sign with respect to that implied in market prices (when very steep and convex smiles occur). We leave the analysis of this issue to future research. ■

Antonio Castagna is an equity derivatives trader at Banca Profilo, Milan. Fabio Mercurio is head of financial modelling at Banca IMI, Milan. They would like to thank Giulio Sartorelli and three anonymous referees for their comments and helpful suggestions. Email: antonio.castagna@bancaprofilo.it, fabio.mercurio@bancaimi.it

References

Beneder R and G Baker, 2005
Pricing multi-currency options with smile
Internal report, ABN Amro

Bisesti L, A Castagna and F Mercurio, 2005
Consistent pricing and hedging of an FX options book
Kyoto Economic Review 74(1), pages 65–83

Carr P and D Madan, 1998
Towards a theory of volatility trading
In Volatility, edited by R Jarrow, Risk Books, pages 417–427

Castagna A and F Mercurio, 2005
Consistent pricing of FX options
Internal report, Banca IMI, available at *www.fabiomercurio.it/consistentfxsmile2b.pdf*

Garman B and S Kohlhagen, 1983
Foreign currency option values
Journal of International Money and Finance 2, pages 231–237

Hagan P, D Kumar, A Lesniewski and D Woodward, 2002
Managing smile risk
Wilmott, September, pages 84–108

Lee R, 2004
The moment formula for implied volatility at extreme strikes
Mathematical Finance 14(3), pages 469–480

Lipton A and W McGhee, 2002
Universal barriers
Risk May, pages 81–85

Malz A, 1997
Estimating the probability distribution of the future exchange rate from option prices
Journal of Derivatives, winter, pages 18–36

Mercurio F and A Pallavicini, 2006
Smiling at convexity
Risk August, pages 64–69

Wystup U, 2003
The market price of one-touch options in foreign exchange markets
Derivatives Week 12(13), pages 1–4