

The Implied Volatility Surfaces*

August 13, 2007

1 Introduction

The volatility surface, or matrix (we will use without any distinction the two terms), is the map of the implied volatilities quoted by the market for plain vanilla options struck at different levels and expiring at different dates. Implied volatility is the parameter σ to plug into the Black-Scholes formula to calculate the price of an option.

In practice, the matrix is built according to three main conventions, each prevailing as a standard in the market according to the traded underlying: the *sticky strike*, the *sticky delta*, and finally the *sticky absolute*. These are simple rules used to conveniently quote and trade options written on different assets and, as such, are not intended to model the evolution of the volatility surface.

When the *sticky strike* rule is effective, implied volatilities are mapped, for each expiry, with respect to the strike prices; this is the rule usually adopted in official markets (e.g.: equity options and futures options). The name *sticky strike* is referred to the fact that implied volatilities do not to change if the underlying asset's price changes. Clearly, that almost never happens since the volatility matrix is not at all constant in reality. Nevertheless, the assumption is believed to be in force for small movements of the underlying asset. Accordingly, traders quote option prices for specific strikes, and usually in terms of premiums, so that one has to back out the implied volatility from those. In table 1 an example of *sticky strike* volatility matrix is presented

If the *sticky delta* rule is adopted, implied volatilities are mapped, for each expiry, with respect to the Delta¹ of the option; this rule is usually used in over-the-counter (OTC) markets (e.g.: FX options). The underpinning assumption is that options are priced depending on their Delta, so that when the underlying asset's price moves and the Delta of an option changes accordingly, a different implied volatility has to be plugged into the formula. A stylized representation of the *sticky delta* volatility matrix is in table 2

Lastly, the *sticky absolute* rule produces matrices with implied volatilities mapped, for each expiry, in terms of absolute distance, measured in some units of price, from the at-the-money strike². This rule, which is in some way a mix of the two described above, prevails in some over the counter markets, such as that for swaptions and for bond options. It implies that the implied volatility for a given strike changes along movements in the underlying asset's price, since also the absolute distance from the at-the-money is different. A *sticky absolute* volatility matrix is shown in table 3.³

Whichever rule is prevailing the main problem a market maker has to cope with, is the building of a consistent volatility surface for a wide range (in terms of expiries and strikes) of options,

*Jason 2007. All rights reserved

¹Though very likely superfluous, we just specify that Delta is the partial derivative of the option price with respect to the underlying asset's price.

²In most of cases, under this rule, the at-the-money strike is set equal to the forward price of the underlying asset

³It can be shown that the sticky strike, delta and absolute rules all produce arbitrage opportunities, should the surface behave as predicted by them. This is the reason why they are mainly regarded as quoting mechanisms and not expressions of actual behaviors of volatility surfaces.

	60.0	70.0	80.0	90.0	100.0	110.0	120.0	130.0	140.0
1m									
2m									
3m									

Table 1: Sticky Strike Matrix

	Δ Put					Δ Call			
	10.0%	20.0%	30.0%	40.0%	ATM	40%	30%	20%	10%
1m									
2m									
3m									

Table 2: Sticky Delta Matrix

provided the knowledge of a few prices. The problem is twofold. First, she needs a tool to interpolate/extrapolate implied volatilities amongst strikes for a given expiry. Second, she faces the problem of the interpolation amongst available expiries. To address the former issue, different interpolation schemes, based in models or simply on interpolating functions, must be introduced. The latter issue is instead tackled by proposing a weighting scheme to consistently include working days, eventful days and holidays.

Before analyzing both issues, we first have to choose how to represent and handle a volatility surface in an efficient, intuitive and convenient way, capable to satisfy different instances presented by a financial institution. This is explained in the following section.

2 Criteria for an Efficient and Convenient Representation of the Volatility Surface

The representation of the volatility surface is not directly related to the specific conventions of the reference market. Actually, not all the rules we have described in the previous section, grant and imply a convenient way to handle the volatility matrix. We here list some critical features that the representation should have:

- **Parsimony:** the representation contains the smallest amount of information needed to retrieve the entire volatility surface for all strikes and expiries.
- **Consistency:** the information contained in the representation is consistently organized along the expiries and strikes, so as to make the integration of missing points, either by interpolation or extrapolation, easily possible.
- **Intuitiveness:** the information provides the user with a clear picture about the shape of the volatility surface, and each piece of the information distinctly affects one specific trait of the volatility surface.

	- 150.0	- 100.0	- 75.0	- 50.0	ATM	+ 50	+75	+100	+150
1m									
2m									
3m									

Table 3: Sticky Absolute Matrix

The representation is *parsimonious* if one can devise a suitable interpolation-extrapolation scheme amongst strikes and expiries, which requires only a few points as input. In principle, this seems a hard task to achieve, since for each given expiry, a volatility smile has as many degrees of freedom as considered strikes. However, from an empirical point of view, volatilities do not move independently from one another, and one may reasonably assume that the degrees of freedom are only three: 1) level, ii) slope and iii) convexity. In fact, as a principal component analysis can show, most of shape variations can be explained either by a parallel shift of the smile or by a tilt to the right or to the left or by a relative change of the wings with respect to the central strike. Therefore, three is the minimum number of points, for each expiry, needed to represent these stylized movements: the volatility for the at-the-money strike, that for an out-of-the-money call and that for out-of-the-money put.⁴

These strike triplets, one for each expiry, must also be chosen in such a way that the resulting representation is *consistent*. To make things clear let us think of a very simple volatility surface with only two expiries: one week and ten years. For both expiries, one of the three strikes to choose may be set equal to the current price of the underlying asset (at-the-money spot). This choice is reasonable but not necessarily the best one. In fact, it would be better to replace the two at-the-money spot values with the forward prices at the two expiries, which can be viewed as expected values of the future underlying asset under suitable measures (the corresponding forward risk adjusted measures).

Things can even be worse for the other two points, since a meaningful selection criterion likely leads to different values for the two expiries: two chosen strikes may convey a good amount of information regarding the smile for the one week expiry, but may be not so informative for the ten year expiry. In fact, what matters (under a probabilistic point of view) is the relative distance of a strike from the central one, possibly expressed in volatility units, which makes the chosen strikes, and their corresponding implied volatilities, comparable throughout the entire range of expiries. A meaningful distance measure, familiar to practitioners, is provided by the Delta of an option (in absolute terms), since it is a common indicator used in the market and it has the same signaling power as the relative distance from the at-the-money (in units of total standard deviation). For this reason, we will select, for each expiry, the volatility for the at-the-money and the 25Δ call and the 25Δ put.⁵ These two Delta levels are introduced because they are almost midway between the center of the smile and the extreme wings (0Δ put and 0Δ call) and also because they are the strikes associated with the highest Volga, thus containing the maximum information on the underlying asset's fourth moment, and hence on the curvature of the smile.

Finally, the representation is *intuitive* if it is directly expressed in terms of three qualitative features of the surface, instead of three implied volatilities. These features, already mentioned above, are the level, the steepness and the convexity of the smile for each maturity. The level is correctly measured by the at-the money volatility. As for the steepness, we can use the difference between the 25Δ call and the 25Δ put, which in the market lore is called *risk reversal* (or *collar*). A good indicator for the convexity is the average level of the volatility for the two 25Δ wings, with respect to the at-the-money level: in the market jargon it is referred to as the *butterfly*. In a representation like this one, a user is able to change the shape of the volatility surface by simply changing these three indicators.

As for the set of expiries, a fixed number of maturities expressed as a fraction or multiple of years (and not not as a fixed date) is the most intuitive and consistent choice to represent the volatility surface. This makes it easier to compare times in the matrix and is more respondent to the requirements of intuitiveness and consistency.

To sum up the considerations above, a convenient and efficient way to represent the volatility surface can be obtained by organizing the information as follows: for each expiry (expressed as time to maturity and in year units) store the at-the-money volatility, the *risk reversal* and the *butterfly* for the 25Δ call and put. The at-the-money is referred to a strike set equal to the

⁴Three points for each expiry can be interpolated by a stochastic volatility model (e.g.: Heston's (1993) model), though they will typically be not enough to ensure a stable calibration.

⁵We drop the “%” sign after the level of the Δ , in accordance to the 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.

Expiry	ATM	25 Δ RiskRev	25 Δ BFly
τ_1	σ_1	rr_1	fly_1
τ_2	σ_2	rr_2	fly_2
τ_3	σ_3	rr_3	fly_3
τ_4	σ_4	rr_4	fly_4
τ_5	σ_5	rr_5	fly_5
τ_6	σ_6	rr_6	fly_6

Table 4: Stylized representation of a volatility surface in the compact form

the forward price for each expiry. One could also choose the at-the-money 0 Δ straddle, whose definition will be given below. Practitioners accustomed with FX options will easily realize that this representation coincides with the standard way to handle a volatility matrix in the FX market. An example of such representation, in a stylized form, is provided in Table 4.

3 Interpolation Among Strikes

In this section we will review some of the most commonly adopted approach to interpolate/extrapolate among strikes the volatility surfaces.

3.1 Interpolating with Polynomials and Splines

If only few volatilities are available (i.e.: 3 or 5) for each expiry, then one can resort to some polynomial interpolation. In case we have 3 volatilities available (typically that referred to an almost ATM strike σ_{K_A} , and other two referred to an OTM Put and an OTM Call, σ_{K_P} and σ_{K_C}), then we can use a 2nd degree polynomial.

Let's assume we have a *sticky strike* matrix.

$$\sigma_K = a \cdot \sigma_{K_A} + b \cdot r(K - K_A) + c \cdot f(K - K_A)^2 \quad (1)$$

where $r = \sigma_{K_C} - \sigma_{K_P}$ and $f = (\sigma_{K_C} + \sigma_{K_P})/2 - \sigma_{K_A}$. When $K = K_A$, we have:

$$\sigma_K = a \cdot \sigma_{K_A} + b \cdot r0 + c \cdot f0 \quad (2)$$

and we set $a = 1$.

By substitution of 1 in the definition of r , we have:

$$\sigma_{K_C} - \sigma_{K_P} = \frac{b}{2} \cdot f \quad (3)$$

implying $b = 2$.

Finally, plugging (1) in the definition of f yields:

$$f = 0.5 \cdot (\sigma_{K_C} + \sigma_{K_P}) - \sigma_{K_A} = 0.25^2 \cdot c \cdot f \quad (4)$$

and we can set the coefficient $c = 16$.

Hence equation (1) can be written as:

$$\sigma_K = \sigma_{K_A} + 2 \cdot r(K - K_A) + 16 \cdot f(K - K_A)^2 \quad (5)$$

In case of a *sticky delta* matrix, the interpolation modifies as follows: Let's assume that at time t we know the risk reversal, defined as:

$$rr(t, T; 25) = \sigma_{25call}(t, T) - \sigma_{25put}(t, T) \quad (6)$$

and the butterfly defined as:

$$bfly(t, T; 25) = 0.5(\sigma_{25call}(t, T) + \sigma_{25put}(t, T)) - \sigma_{ATM}(t, T) \quad (7)$$

. We assume we have this quantities for the 25 Δ for a given expiry T . This is usually the case, though any other two Δ can be chosen (though on opposite sides with respect the ATM level). Define the ATM Δ of the Put as the

$$\Delta_{ATM} \equiv \Delta(S, t, T, K_{ATM}, \sigma_{ATM}, r_d, r_e, -1) \quad (8)$$

where K_{ATM} is the ATM strike (according to the definition, it can be equal to the ATM spot, ATM forward, or the ATM zero-delta).

Besides, we recognize that the relationship between the Δ of a call and the Δ of a put both with strike K and expiring in T is

$$\Delta_{put}(S, t, T, K, \sigma_K, r_d, r_e) = Df^e(t, T) - \Delta_{call}(S, t, T, K, \sigma_K, r_d, r_e) \quad (9)$$

where

$$\Delta_{call}(S, t, T, K, \sigma_K, r_d, r_e) \equiv \Delta(S, t, T, K, \sigma_K, r_d, r_e, 1)$$

In order to set up a proper interpolation function passing through the three points given by the main strikes, we choose once again a 2nd order polynomial of the following kind:

$$\begin{aligned} \sigma_{\Delta Put}(t, T) &= a \cdot \sigma_{ATM}(t, T) + b \cdot rr(t, T; 25)(\Delta_{put}(K_{\Delta Put}) \\ &\quad - \Delta_{put}(K_{ATM})) + c \cdot bfly(t, T; 25)(\Delta_{put}(K_{\Delta Put}) - \Delta_{put}(K_{ATM}))^2 \end{aligned} \quad (10)$$

We are using a slightly loose notation, where the only argument of the function Δ is the strike and the others are omitted. When $\Delta_{put}(K_{\Delta Put}) = \Delta_{put}(K_{ATM})$, we have $\sigma_{\Delta Put}(t, T) = a \cdot \sigma_{ATM}(t, T)$, so we can immediately set $a = 1$.

For the Risk Reversal we use the following relationship:

$$\begin{aligned} rr(t, T; 25) &= \sigma_{25Call}(t, T) - \sigma_{25Put}(t, T) \\ &= a \cdot \sigma_{ATM}(t, T) + b \cdot rr(t, T; 25)(\Delta_{put}(K_{25Call}) - \Delta_{put}(K_{ATM})) \\ &\quad + c \cdot bfly(t, T; 25)(\Delta_{put}(K_{25Call}) - \Delta_{put}(K_{ATM}))^2 \\ &\quad - a \cdot \sigma_{ATM}(t, T) - b \cdot rr(t, T; 25)(\Delta_{put}(K_{25Put}) - \Delta_{put}(K_{ATM})) \\ &\quad - c \cdot bfly(t, T; 25)(\Delta_{put}(K_{25Put}) - \Delta_{put}(K_{ATM}))^2 \end{aligned} \quad (11)$$

$$\begin{aligned} bfly(t, T; 25) &= \frac{\sigma_{25Call}(t, T) + \sigma_{25Put}(t, T)}{2} - \sigma_{ATM}(t, T) \\ &= \frac{1}{2}[a \cdot \sigma_{ATM}(t, T) + b \cdot rr(t, T; 25)(\Delta_{put}(K_{25Call}) - \Delta_{put}(K_{ATM})) \\ &\quad + c \cdot bfly(t, T; 25)(\Delta_{put}(K_{25Call}) - \Delta_{put}(K_{ATM}))^2 \\ &\quad + a \cdot \sigma_{ATM}(t, T) - b \cdot rr(t, T; 25)(\Delta_{put}(K_{25Put}) - \Delta_{put}(K_{ATM})) \\ &\quad + c \cdot bfly(t, T; 25)(\Delta_{put}(K_{25Put}) - \Delta_{put}(K_{ATM}))^2] \\ &\quad - \sigma_{ATM}(t, T) \end{aligned} \quad (12)$$

From the two equations above we can set up a system to calculate r and f .

In case of a *sticky absolute* matrix, the interpolation can be performed similarly to the n case of a *sticky strike case*, by setting the $K_A = 0$ and interpolating among the available volatilities referred to the corresponding absolute spreads s_i (e.g.: $s_1 = -100, s_2 = 100$). Once the parameters f and r are calculated, the volatility for a given strike K is derived from the polynomial calculated in $s_K = K - K_A$.

In case more than 3 volatilities are available for each expiry, one may resort to a *cubic splines* interpolation. Assume we are in a *sticky strike* environment and we are provided with a set of strikes K_i and corresponding volatilities σ_i , $i = [1, \dots, n]$. If $K_i < K < K_{i+1}$, a linear interpolation yields:

$$\sigma_K = A\sigma_i + B\sigma_{i+1} \quad (13)$$

where $A = \frac{K_{i+1}-K}{K_{i+1}-K_i}$ and $B = 1 - A$. Adding the two second derivatives for the volatilities, we get:

$$\sigma_K = A\sigma_i + B\sigma_{i+1} + C\sigma_i'' + D\sigma_{i+1}'' \quad (14)$$

where $C = \frac{1}{6}(A^3 - A)(K_{i+1} - K_i)^2$ and $D = \frac{1}{6}(B^3 - B)(K_{i+1} - K_i)^2$. The two derivatives are unknown, but by imposing the first derivatives to be continuous:

$$\frac{d\sigma_K}{dK} = \frac{\sigma_{i+1} - \sigma_i}{K_{i+1} - K_i} - \frac{3A^2 - 1}{6}(K_{i+1} - K_i)\sigma_i'' + \frac{3B^2 - 1}{6}(K_{i+1} - K_i)\sigma_{i+1}'' \quad (15)$$

we obtain the following recursive system:

$$\frac{K_i - K_{i-1}}{6}\sigma_{i-1}'' + \frac{K_{i+1} - K_{i-1}}{3}\sigma_i'' + \frac{K_{i+1} - K_i}{6}\sigma_{i+1}'' = \frac{\sigma_{i+1} - \sigma_i}{K_{i+1} - K_i} - \frac{\sigma_i - \sigma_{i-1}}{K_i - K_{i-1}} \quad (16)$$

for $i = 2, \dots, n - 1$. To solve the system, two conditions at the boundaries of the first derivatives must be set. One natural choice is to set $\sigma'_1 = \sigma'_n = 0$. A second possible choice is to set them equal to $\sigma'_1 = \frac{\sigma_2 - \sigma_1}{K_2 - K_1}$ and $\sigma'_n = \frac{\sigma_n - \sigma_{n-1}}{K_n - K_{n-1}}$, so that the extrapolation beyond the available points is not a constant level.

The *sticky absolute* case can be handled in the same way explained above as for the polynomial interpolation. As for the *sticky delta* case, use the available Δ 's instead of the strikes in the cubic spline, by translating them all in terms of put's Δ 's, as explained above.

3.2 Interpolating by a Stochastic Volatility Model

An alternative methodology to interpolate among strike is using a stochastic volatility model. For example, one may calibrate an Heston's model to available volatilities (actually, to prices implied by them): some problems arise as stability is concerned, but the model is reach enough to satisfactorily fit market prices in most of cases.

Another popular model is the SABR model, by Hagan et al. (2002), which is used to generate volatility smiles for any expiry and tenor. This model is quite appealing since in its framework the authors derive an explicit function for equivalent Black-Scholes implied volatilities:

$$\sigma(K) \simeq \frac{\alpha}{(S_{a,b}(0)K)^{\frac{1-\beta}{2}} \left[1 + \frac{(1-\beta)^2}{24} \ln^2 \left(\frac{S_{a,b}(0)}{K} \right) + \frac{(1-\beta)^4}{1920} \ln^4 \left(\frac{S_{a,b}(0)}{K} \right) \right]} x(z) \Gamma(K) \quad (17)$$

where

$$\begin{aligned} \Gamma(K) &:= 1 + \left[\frac{(1-\beta)^2 \alpha^2}{24(S_{a,b}(0)K)^{1-\beta}} + \frac{\rho\beta\epsilon\alpha}{4(S_{a,b}(0)K)^{\frac{1-\beta}{2}}} + \epsilon^2 \frac{2-3\rho^2}{24} \right] T_a \\ z &:= \frac{\epsilon}{\alpha} (S_{a,b}(0)K)^{\frac{1-\beta}{2}} \ln \left(\frac{S_{a,b}}{K} \right) \\ x(z) &:= \ln \left[\frac{\sqrt{1-2\rho z + z^2} + z - \rho}{1-\rho} \right] \end{aligned}$$

with the particular case

$$\sigma(K = S_{a,b}(0)) \simeq \frac{\alpha}{S_{a,b}^{1-\beta}(0)} \Gamma(S_{a,b}(0))$$

for the ATM volatility. Calibration can be performed directly on the market volatility quotes. The model based interpolation is in practice referred to a *sticky strike* environment, though it can be applied also to the other contexts, though the calibration requires the strikes corresponding to the available volatilities and the output is naturally provided for levels of strikes. Besides, each model implies a smile dynamics as the underlying asset's price moves. If one wants to be stuck to one of the described rules, then she has to recalibrate with the new underlying price and with inputs re-arranged accordingly to the adopted rule. That is to say: if the *sticky strike* rule is in force, then one has simply to recalibrate with the new underlying asset's price and the same inputs; if the *sticky delta* or *sticky absolute* rule is adopted, then one has first to recalculate the strikes corresponding to the different levels of Δ 's, or absolute distance from the ATM, given the new underlying asset's price, then she calibrates.

4 Smile Interpolation among Expiries: Implied Volatility's Term Structure

In the options market a set of maturities are actively traded and provide the market makers with a guide to price options with any expiry. One could first try and use some function fitting, more or less accurately, the given expiries, and then interpolate/extrapolate by means of it. For example, a form of the Heston type for the instantaneous variance of the underlying asset:

$$\eta^2(t) = \sigma_\infty^2 + (\eta_0^2 - \sigma_\infty^2)e^{-\kappa t} \quad (18)$$

has been proposed to fit the FX options volatility term structure of at-the-money volatilities. This expression could also be used in other markets, such as the equity or commodity options markets. In (18), the instantaneous variance $\eta^2(t)$ evolves toward a long-term average level σ_∞^2 by a mean-reversion speed measured by the parameter κ , starting from an initial level η_0^2 . In order to retrieve the implied volatility for a given expiry, one has to integrate equation (18), divide the result by the time to the expiry T and take the square root:

$$\varsigma(T) = \sqrt{\frac{\int_0^T \eta^2(t) dt}{T}} = \sqrt{\sigma_\infty^2 + (\eta_0^2 - \sigma_\infty^2) \frac{1 - e^{-\kappa T}}{\kappa T}} \quad (19)$$

Usually, though (18) is qualitatively appealing, the fitting to the market data is not completely satisfactory, and one must resort to some time-dependency of the parameters (or to the introduction of a free extra-parameter) in order to improve the performance.

As far as interest rates are concerned, formula (18) is typically not flexible enough to recover the specific shape of the volatility term structure observed in the market, so that one has to devise a more suitable formulation. A possible example is given by the following function, which has been proposed for modelling the instantaneous volatility of forward (LIBOR) rates expiring at time T_i (see, for example, Brigo and Mercurio, 2006):

$$\eta_i(t) = (a(T_i - t) + b)e^{-b(T_i - t)} + d,$$

leading to the implied volatility:

$$\varsigma_i(T) = \sqrt{\frac{\int_0^T ((a(T_i - t) + b)e^{-b(T_i - t)} + d)^2 dt}{T}} \quad (20)$$

Also this function, though richer than the first one above, is seldom able to perfectly fit the available market quotes simultaneously.

Functions (19) and (20) are examples of parametric forms that can be chosen to fit the at-the-money volatilities in a given options market. However, the problem of finding a suitable parametrization is not fundamental by itself, since one may reasonably accept to take the term

structures provided by the market as a matter of fact (unless they engender an arbitrage) and simply interpolate between consecutive maturities.⁶

The volatility interpolation between two given expiries needs to be correctly implemented by taking into account the daily variations, and weighing the days with sensible factors. Focussing on the interpolation issue, we would like to stress that the interpolation of the term structure refers only to the at-the-money volatility, thus only concerning the *level* of the surface.

Let's start with a set of M standard traded expiries $\mathcal{T} = \{T_1, T_2, \dots, T_M\}$. Our problem is to devise a method to consistently interpolate between two contiguous dates. Let T be an expiry date between two dates $T_i \leq T \leq T_{i+1}$ and $\eta(t)$ be the instantaneous volatility of the spot process. The squared implied volatility for expiry T , in a BS world, is given by

$$\varsigma^2(T) = \frac{1}{T} \int_0^T \eta^2(t) dt$$

Such a volatility can be obtained by linearly interpolating the mean variance

$$V(T) = \int_0^T \eta^2(t) dt = T\varsigma^2(T)$$

between its time T_i and time T_{i+1} values. We get:

$$\begin{aligned} \int_0^T \eta^2(t) dt &= \frac{V(T_{i+1}) - V(T_i)}{T_{i+1} - T_i} (T - T_i) + V(T_i) \\ &= \frac{T_{i+1}\varsigma^2(T_{i+1}) - T_i\varsigma^2(T_i)}{T_{i+1} - T_i} (T - T_i) + T_i\varsigma^2(T_i) \\ &= \frac{T - T_i}{T_{i+1} - T_i} T_{i+1}\varsigma^2(T_{i+1}) + \frac{T_{i+1} - T}{T_{i+1} - T_i} T_i\varsigma^2(T_i) \end{aligned}$$

and hence

$$\varsigma(T) = \sqrt{\frac{T - T_i}{T_{i+1} - T_i} \frac{T_{i+1}}{T} \varsigma^2(T_{i+1}) + \frac{T_{i+1} - T}{T_{i+1} - T_i} \frac{T_i}{T} \varsigma^2(T_i)}. \quad (21)$$

This is what we call the *total variance interpolation* method.

As a subsequent step, we then take into account the weighting of the days entering in the interpolation, since holidays and other events may affect heavily the daily volatility. One simple way to account for different weighting of the days is as follows:

- Calculate the number of days N_1 occurring between T_i and T and the number of days N_2 occurring between T_i and T_{i+1} .
- Associate each day with a proper weight w_i :
 - $w_i = 1$ for a normal business day;
 - $w_i < 1$ for holidays in the underlying markets (e.g., in the FX market, it can be set equal to 0.5 if the day is holiday in one of the two countries involved in the exchange rate, or it can set at 0 for weekends);
 - $w_i > 1$ for days when special events are expected (e.g.: key economic figures are released).

- Set

$$\tau_1 = \frac{\sum_{i=1}^{N_1} w_i}{N_y}, \quad \tau_2 = \frac{\sum_{i=1}^{N_2} w_i}{N_y},$$

where N_y is the total number of days in the year ($N_y = 365$ or $N_y = 366$);

⁶Extrapolation outside the available range of expiries is a minor issue, normally managed by traders by adding some spread over the last quoted option's maturity.

- Replace $T - T_i$ with τ_1 and $T_{i+1} - T_i$ with τ_2 in the interpolation formula (21):

$$\varsigma(T) = \sqrt{\frac{\tau_1}{\tau_2} \frac{T_{i+1}}{T} \varsigma^2(T_{i+1}) + \frac{\tau_2 - \tau_1}{\tau_2} \frac{T_i}{T} \varsigma^2(T_i)}. \quad (22)$$

The weighting has tangible effects for expiries up to 1 year. After that, it is quite immaterial to use either (21) or (22).

As mentioned before, the interpolation is used only for the at-the-money volatilities. If the implied volatility of an out-of-the-money option has to be recovered, then first one interpolates by means of (22) the at-the-money volatility, then she may apply the same adjustment over a simple linear interpolation also on the other strikes (in whatever form they are expressed, i.e.: values, Delta's or absolute distance from the at-the-money).

5 Consistency check

Every volatility surface must satisfy some conditions in order to rule out any arbitrage opportunities exploited by means of static positions set up at time t : in this sense they are necessary and sufficient conditions to declare the volatility surface admissible

Let $\mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, \omega)$ the price of an option at time t , expiring in T and struck at K . The other input of the *B&S* price are the implied volatility σ related to the strike K and expiry T , r_d is the risk free rate and δ is the asset's flows (both continuously compounded). The parameter ω indicates the price of a call (put) option when it is set at 1 (-1).

The first condition is a couple of constraints which prevents the possibility to establish an arbitrage by trading a call spread or a put spread:

$$\frac{\partial \mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, \omega)}{\partial K} \Big|_{\omega=1} \leq 0 \quad (23)$$

$$\frac{\partial \mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, \omega)}{\partial K} \Big|_{\omega=-1} \geq 0 \quad (24)$$

Basically, call options struck at a given level must be worth more than an otherwise identical call struck at a higher price; specularly a put option must be worth more than an otherwise identical put struck at a lower price.

More conditions on the slope of the smile are provided by the following constraints, for extreme strikes. The option price $\mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, 1)$, as a function of the strike K , is twice differentiable and satisfies the following (no-arbitrage) conditions:

- i) $\lim_{K \rightarrow 0^+} \mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, 1) = S_0 e^{-\delta T}$ and
- ii) $\lim_{K \rightarrow +\infty} \mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, 1) = 0$;
- iii) $\lim_{K \rightarrow 0^+} \frac{d\mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, 1)}{dK} = -e^{-r_d T}$ and $\lim_{K \rightarrow +\infty} K \frac{d\mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, 1)}{dK} = 0$.

The second condition set a constraint on the convexity of the surface at a given maturity, ruling out the possibility to buy a butterfly for nothing or even with a positive cash-flow:

$$\frac{\partial^2 \mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, \omega)}{\partial K^2} > 0 \quad (25)$$

The third and final condition is on the slope with respect the time to maturity of the surface: it should be not be possible to buy an option and sell an otherwise identical options with a shorter time to maturity without paying a premium.

$$\frac{\partial \mathcal{O}(S, t, T, K, \sigma_K(t, T), r_d, \delta, \omega)}{\partial T} > 0 \quad (26)$$

It is worth spending some words for the last condition. The Θ of an option is normally a positive function of the time to maturity, and that is exactly what has been expressed in the condition above. Nevertheless, sometimes for deep in the money options, and when there is a huge difference between the domestic and foreign interest rate, the Θ could be negatively related to the time to maturity. To avoid this kind of problem, one should check the admissibility of the volatility surface with out of the money options, for which condition (26) is always a bounding constraint.