Slope Matters to Land on the Right Price

Uwe Wystup, Sebastiaan Van Mulken
MathFinance AG, Frankfurt am Main

November 2023

Traders and Sales often argue that if you already have a volatility surface and can hence price all vanilla options, for exotics, e.g., for barrier options, all you need to do is to add a barrier to the contract description and extract entry on the pricing screen. We will illustrate in this column why, unfortunately, it is not so easy, even for simple European digitals.
In the last FX column we discussed the observed price discrepancy on different pricing platforms for the Europen digital. It is clear that the volatility surface plays a huge role in the pricing of this contract. In this column we would like to expand a bit more on how to obtain such a volatility surface and the impact of the shape (that is the level *and slope*) on the price of exotics, where we use the digital contract as a basic example.
Similar as in the previous FX column, we want to price an European USD-JPY digital put with strike $K = 115.00$ and a payoff of 5 million USD, with the premium given in USD. The horizon date is March 29, 2023, the spot reference is $S_0 = 132.87$ JPY per USD and the time to maturity is 84 days (June 21, 2023) with delivery date June 23, 2023. A snippet of the market on the horizon date is shown in Table 1.

T	ATM	25RR	10RR	25BF	10BF	Fwd	Rate(USD)
1W	11.340	-1.0825	-1.855	0.2675	0.7050	132.848	5.000
2W	12.255	-1.2000	-2.210	0.2850	0.7755	132.494	5.070
1M	11.695	-1.3350	-2.580	0.3000	0.8400	132.317	5.095
2M	12.485	-1.7700	-3.385	0.3650	1.0200	131.467	5.175
3M	12.500	-1.8100	-3.580	0.3900	1.1450	130.708	5.225
6M	11.735	-1.1532	-3.070	0.4125	1.3150	128.455	5.245

Table 1: *Snippet of the market data on March 29, 2023. T is the tenor, ATM is the at the money volatility, 25RR and 10RR are risk reversal volatility at the 25 and 10 delta respectively; BF is the butterfly. Fwd is the forward rate. Rate is the USD money market interest rate for deposits, all quotes in %.*

With the above USD-JPY market, we showed that the payoff and hence the value of the USD-paying digital is the sum of a vanilla and the JPY-paying digital. Furthermore, the JPY-paying digital can be approximated by a tight vanilla call- or put-spread with strikes around the given center strike of the digital.

Volatility in the delta space

As can be seen in the market snippet in Table 1, the volatility quoting mechanisms are FX specific and differ significantly, for example, from the equity option market where prices are quoted for certain strikes. The volatilities quoted are usually for the 10% and 25% delta of the butterfly (BF) and the risk reversal (RR) and the at-the-money (ATM) volatility. Since the Black-Scholes formula needs the volatility for a given strike, we need to transform the volatilities from the delta space to the strike space. This is not as trivial as it seems, for example one has to be very cautious which at-the-money convention is used and which delta type is used.
The butterfly structure consists of a short ATM-straddle and long strangle with 25-delta strikes (or 10-delta strikes). Therefore, the volatility quote is the average of the 25-delta call and 25-delta put minus the at-the-money volatility,

$$\mathsf{BF}_T(\Delta) = \frac{\sigma_T^{\mathsf{call}}(\Delta) + \sigma_T^{\mathsf{put}}(\Delta)}{2} - \mathsf{ATM}_T \tag{0.1}$$

for a given time tenor T and delta Δ. The risk reversal (RR) is the difference between the implied volatilities of a 25-delta call and 25-delta put,

$$\text{RR}_T(\Delta) = \sigma_T^{\text{call}}(\Delta) - \sigma_T^{\text{put}}(\Delta). \tag{0.2}$$

The above two formulas yield the following simplified expressions[1] for the delta call and delta put volatilities,

$$\begin{aligned} \sigma_T^{\text{call}}(\Delta) =& \text{ATM}_T + \text{BF}_T(\Delta) + 0.5\text{RR}_T(\Delta), \\ \sigma_T^{\text{put}}(\Delta) =& \text{ATM}_T + \text{BF}_T(\Delta) - 0.5\text{RR}_T(\Delta), \end{aligned} \tag{0.3}$$

which we can use based on the market at time of maturity (June 21, 2023). Talking about inter- and extrapolation, the 84 days to expiry is somewhere at the far right end between the 2M and 3M tenor. We challenge you to think about the best way to interpolate the volatilities and rates, which has quite some freedom, as the interpolation method doesn't necessary have to be the same for the volatilities and the forward or USD rate.

For demonstration purpose, we will use a basic linear interpolation for the BF, RR and ATM volatilities, the forward rate and also for the USD rate. The call and put volatilities corresponding to the given deltas are summarized in Table 2.

	10 delta	25 delta	ATM delta
Risk Reversal	-3.541%	-1.802%	12.497%
Butterfly	1.12%	0.385%	12.497%
σ^{call}	11.8465%	11.981%	12.497%
σ^{put}	15.3875%	13.783%	12.497%

Table 2: *Call and put volatilities corresponding to the 10- and 25-delta, obtained from equation (0.3). The RR, BF and ATM volatilities are based on a linear interpolation.*

Using the above formulas and the (linearly) interpolated market volatilities for $T = 84$, the implied call and put volatilities in the delta space for the market snippet given in Table 1 are shown in Figure 1. Given the implied volatility for these five points, we can now construct a volatility smile in the delta space. One of the simplest ways to construct the volatility smile is to fit following quadratic function to the five points in the delta space:

$$\sigma_\Delta = \sigma_{\text{ATM}} + \alpha(\Delta - \Delta_{\text{ATM}}) + \beta(\Delta - \Delta_{\text{ATM}})^2. \tag{0.4}$$

Malz proposed a simplification of this quadratic approach by making the assumptions that $\Delta_{\text{ATM}} = 0.5$ and 25-delta call is equal to the 75-delta put [1]. This would be correct for volatilities quoted on the forward-delta space and for premium-excluded volatility quotes. If this was the case, using only the 25-delta points, one can solve the general quadratic formula above explicitly and obtain

$$\sigma_\Delta = \sigma_{\text{ATM}} - 2\text{RR}_{25\text{Delta}}(\Delta - 0.5) + 16\text{BF}_{25\text{Delta}}(\Delta - 0.5)^2. \tag{0.5}$$

This function is plotted on top of the implied volatilities in Figure 1. It is clear by construction that this function fits the ± 25 delta perfectly, but it is a bad fit for the 10 deltas. This is

[1] Smile construction is generally more complex, taking all the conventions into account, but we kept it simple and rather educational in this column.

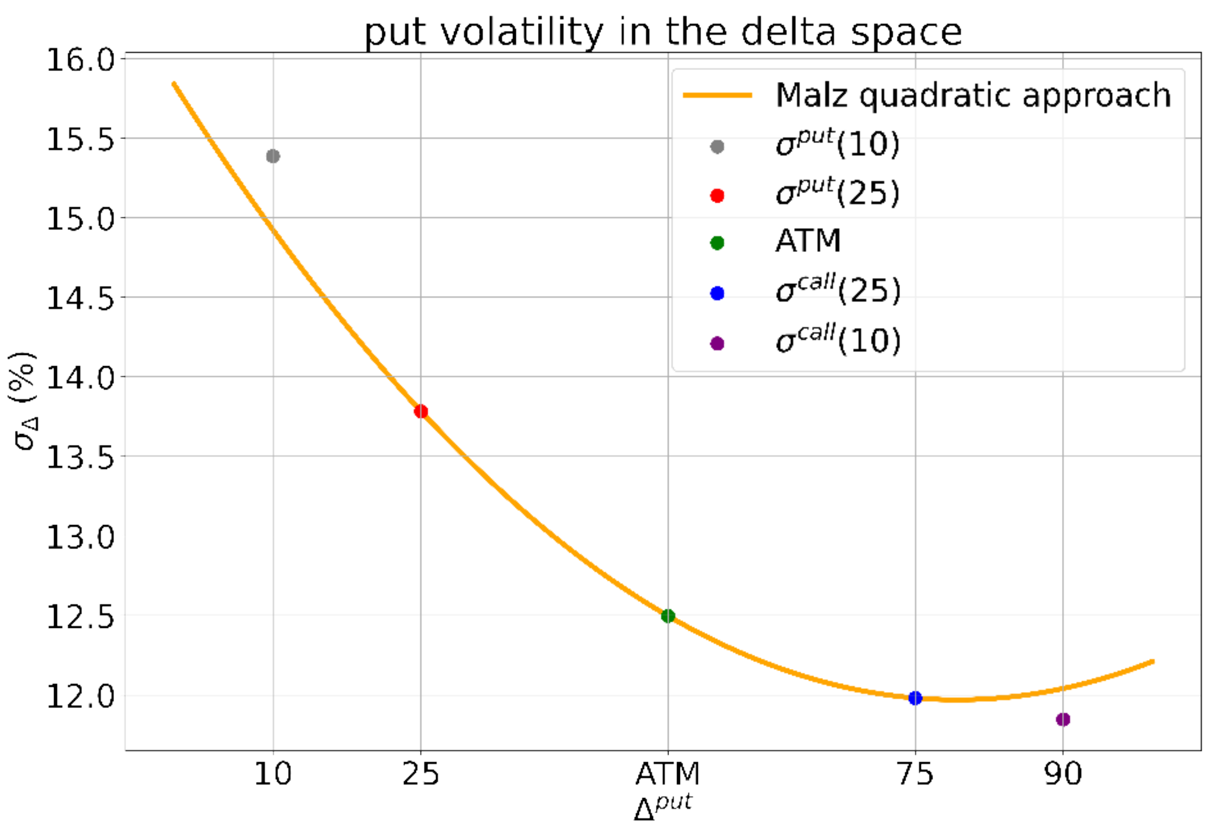

Figure 1: *(Put) volatility smile in the delta space. Here the relationship $\Delta_T^{\text{call}} - \Delta_T^{\text{put}} = 1$ is used, meaning the 25 delta call is equal to the 75 delta put for the same time to maturity T. The interpolation is obtained using the Malz quadratic approach.*

normally not its biggest problem, as the 10 delta volatility may itself be a number calculated by a market data provider via (a typically unknown) extrapolation. However, as pointed out by Wystup & Reiswich, the use of this functional relationship has a number of other problems [2].

From the delta space to the strike space

Since the Black-Scholes equation requires strike-volatility pairs, we need to convert volatilities quoted on the delta space to volatilities quoted on the strike space. But one has to know how the given delta is quoted, since the formula to calculate the strike for given delta and implied volatility depends on the delta conventions (and on the volatility we are trying to find). In our example it is straight forward: the currency pair USD-JPY contains only currencies from the Organization for Economic Co-operation and Development (OECD) economies and its premium currency is in FOR (USD); therefore the delta convention is a *premium-adjusted spot delta* $\Delta_{\text{S, pa}}(K, \sigma, \phi)$. The premium-adjusted spot delta is given by

$$\Delta_{\text{S, pa}}(K, \sigma, \phi) = \phi \frac{K}{F} e^{-r_f T} \mathcal{N}(\phi d_-) \tag{0.6}$$

where $\phi = \pm 1$ for a call (+1) or put (-1) delta and d_-, in terms of the forward F, is given by

$$d_- = \frac{\ln\left(\frac{F}{K}\right) - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}.$$

The right-hand side of Equation (0.6) is equal to the amount of FOR to buy per 1 FOR. Equation (0.6) can only be solved numerically using a root finding procedure, which also means that for a given delta it is possible to get more than one corresponding strike. In such cases it is common to search for strikes corresponding to deltas which are on the right hand side of the delta maximum [2]. Solving equation (0.6) yields the values shown in Table (3).

strike	10 delta put	25 delta put	ATM delta	75 delta put	90 delta put
K	119.26	125.25	130.75	135.41	138.86

Table 3: *Strikes corresponding to the put deltas using equation (0.6).*

To finally get the volatility smile in the strike space, these five points need to be inter- and extrapolated. For this interpolation we will use the raw Stochastic Volatility Inspired (SVI) parameterization [4], which models the total implied variance $w^{\text{SVI}}_{\text{imp}}(k)$ via the hyperbola (another conic section)

$$w^{\text{SVI}}_{\text{imp}}(k) = a + b\left(\rho(k-m) + \sqrt{(k-m)^2 + \sigma^2}\right) \tag{0.7}$$

on the log-moneyness space $k = \ln \frac{K}{F}$, where the total implied variance is defined as $w_{\text{imp}} = \sigma^2_{\text{imp}} T$. In short, the a (think of ATM) gives the overall level of variance, b (think of Butterfly = convexity) the angle between the left and right asymptotes, the curvature σ the smoothness of the vertex (a kink for $\sigma = 0$), ρ (think of Risk Reversal) the orientation and finally m (the minimum of the curve, not necessarily ATM) translates the graph. In this formula, σ is not the volatility of the underlying's price process. The calibrated SVI curve is shown in Figure (2) as the green line.

Beside the SVI parameterization, one could also choose a second order polynomial (SOP) fit, which has the general form $y(x) = ax^2 + bx + c$, through the five implied volatilities. This is shown in Figure (2) as the blue line. Besides that the fit in general is not as good as the SVI, especially around the 25 delta put at $K = 125.25$, it also deviates quite a bit in the extrapolated tail. For example, at $K = 115.00$, the SVI parameterization has a volatility around 16.5%, where the second order polynomial fit is very close to a volatility of 17%. As we will discuss later in this column, the level and slope of the volatility smile will have huge impact on the price of the digital put.

Finally we also show the Malz quadratic fit that we used in the delta space, which is shown as the orange line in Figure (2). This shows very nicely that mapping between the delta and the strike space is not linear, as the quadratic fit in the delta space is not a quadratic equation anymore in the strike space.

Absence of static arbitrage

Besides taking term structure into account, the volatility smile, which prices the quoted market instruments correctly, must satisfy a number of conditions. Absence of static arbitrage implies that the corresponding total variance must be an increasing function of the maturity (absence of calendar spread arbitrage). Furthermore, a negative implied probability density implies potential butterfly arbitrage.

The reason is that the risk-neutral density $\phi(K, T)$ can be extracted from option prices, see

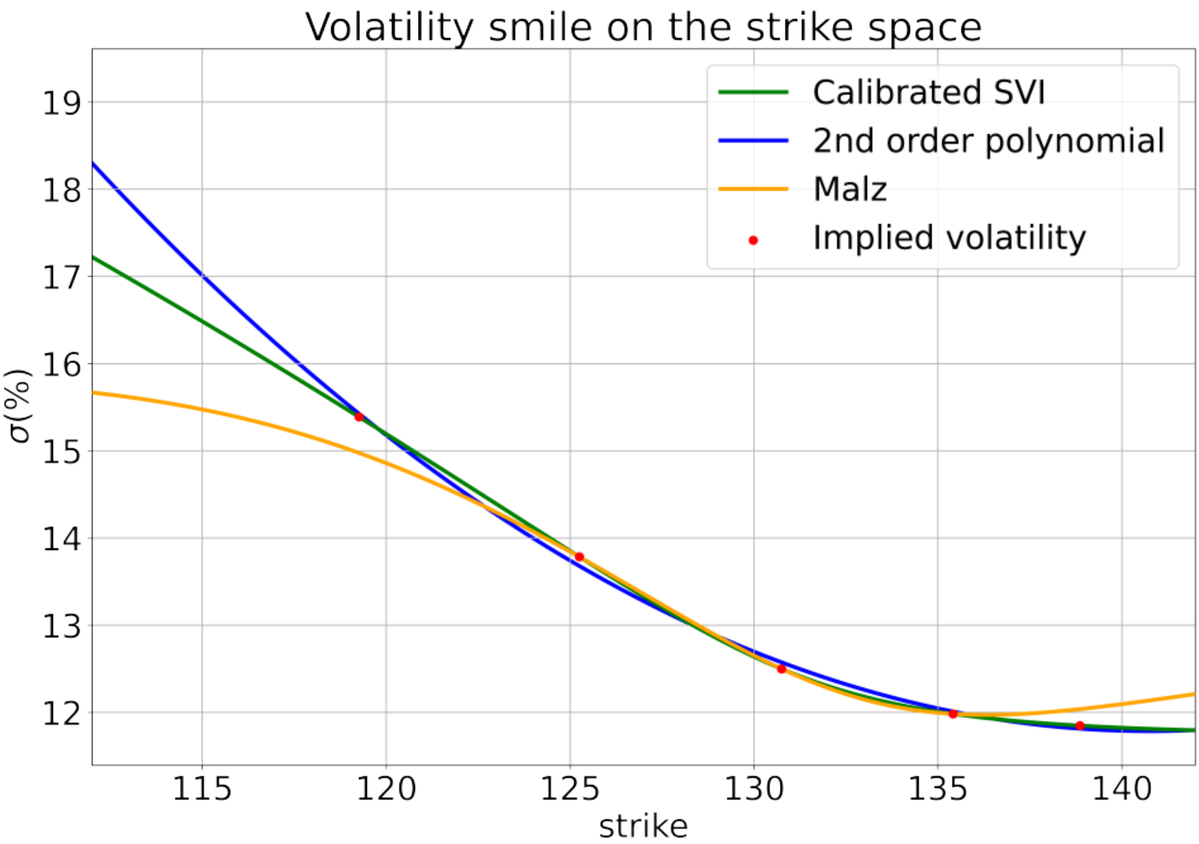

Figure 2: *USD-JPY implied volatility surface on the strike space. The green line is the SVI parameterization, the blue line is the fit of a second order polynomial and the orange line is the Malz volatility smile converted to the strike space ignoring the 10-delta quotes by construction.*

Breeden and Litzenberger [5], using the formula

$$\phi(K,T) = e^{+r_d T}\frac{\partial C^2(K,T)}{\partial K^2} \approx e^{+r_d T}\left(\frac{C(K+\Delta K,T) - 2C(K,T) + C(K-\Delta K,T)}{(\Delta K)^2}\right). \tag{0.8}$$

The right-hand side of equation (0.8) is the compounded second-order central difference approximation of the second derivative. Basically the second derivative can be viewed as a series of very tight overlapping butterfly spreads. Hence the name butterfly arbitrage. More details can be found in a previous FX column "Negative Butterflies and Why We Check Butterfly Arbitrage by a Non-Negative Probability Density" [3].
The probability density function for both the calibrated SVI and SOP are shown in Figure (3). It is clear that the probability density for both the SVI calibration and the second order polynomial fit are strictly positive, meaning that there is no indication of butterfly arbitrage. Several other conditions that have to be considered to obtain the perfect volatility surface are summarized in one of the previous columns [6].

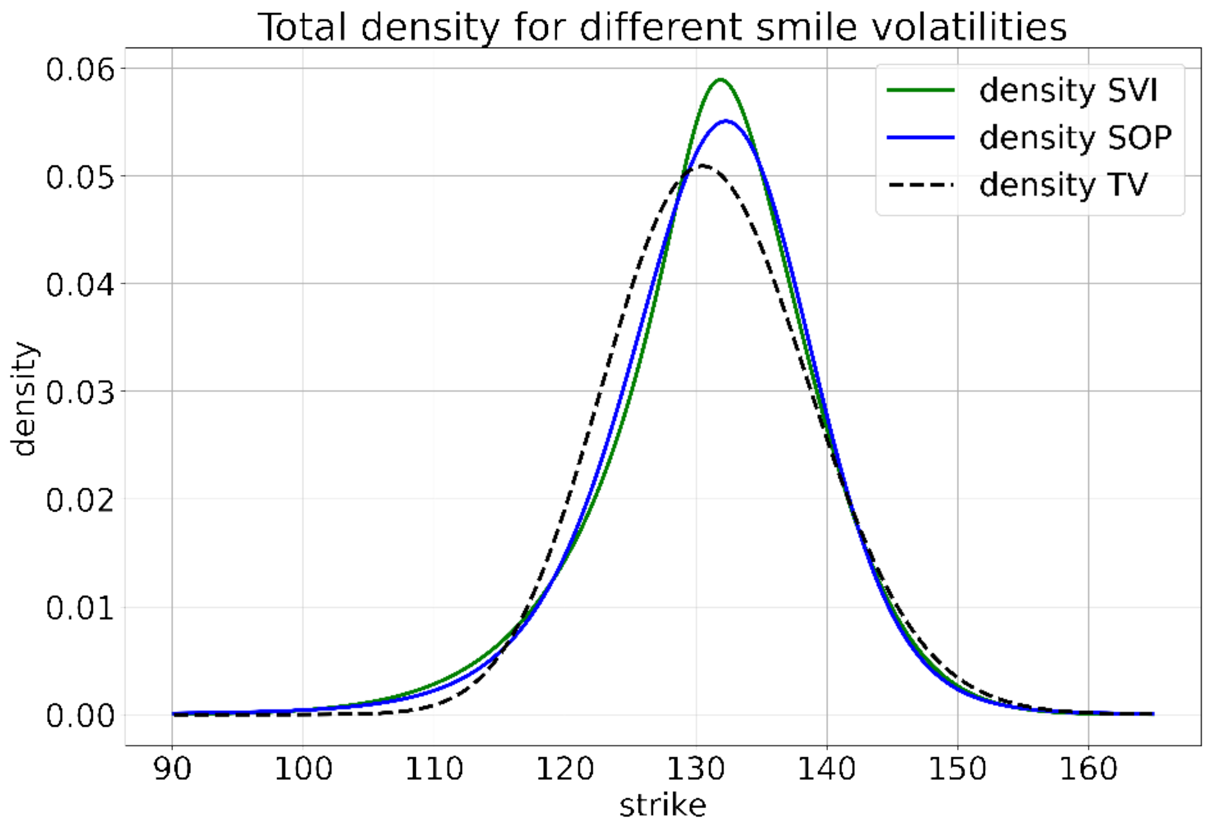

Figure 3: *Plot of the density function $\phi(K)$ for the market calibrated SVI curve, for the second order polynomial (SOP) and the theoretical value (TV) when only a flat ATM volatility is used. The graphs are strictly positive, meaning that for both fits there is no indication of butterfly arbitrage.*

Gatheral et al. have shown that using specific boundaries for the SVI-parameterization, the resulting volatility surface will be free of butterfly arbitrage [4].

The level *and* slope of the volatility smile matter

With the volatility smiles obtained in the previous section for both the SVI and SOP, we can now calculate the prices of the digital put using the vanilla put spread. In Table (4) we summarize the results for $K = 115.00$, $K = 125.00$ and $K = 130.00$.
The values of the digital put contracts for both the calibrated SVI and SOP at strikes closer to the ATM ($K = 125.00$ and $K = 130.00$) seem to diverge towards the same value. As one would maybe expect from Figure 2, we see that for $K = 115.00$ the price difference is relatively large, since the difference in level of the SVI volatility smile and SOP volatility smile is also larger in the wings. The extreme, and obvious wrong, case would be to take the volatility obtained from the Malz curve. Here, at $K = 114.50$ and $K = 115.50$, the volatilities are 15.52% and 15.43% respectively. The resulting put price are 38.87 USD and 51.67 USD.
However, one should keep in mind that the notional of the plain vanilla puts in the put spread is 5M JPY. Concrete, this means for example that the price $\mathsf{P}^{\mathsf{Put}}_{\mathsf{SVI}}(K = 114.50) = 57.27$ USD per 5M JPY is roughly $1.145 \cdot 10^{-5} = 0.00001145$ USD per JPY. The difference between the digital values $\mathsf{P}^{\mathsf{dig}}_{\mathsf{SVI}}$ and $\mathsf{P}^{\mathsf{dig}}_{\mathsf{SOP}}$ may seem relatively large (approximately 24,000 USD), but taking the 5M notional into consideration this difference would be smaller than the bid-offer spread. Additionally, when traded a sales margin is added and it is not strange to encounter relatively wide bid-offer spreads for these far out of the money digital contracts.

strike	σ(SVI)	σ(SOP)	P^{Put}_{SVI}	P^{Put}_{SOP}	P^{dig}_{SVI}	P^{dig}_{SOP}	P^{dig}_{MF}
$\boldsymbol{K = 115.00}$					153,410	129,236	153,306
$K_1 = 114.50$	16.61%	17.22 %	57.27	68.87			
$K_2 = 115.50$	16.36%	16.81 %	68.87	78.47			
$\boldsymbol{K = 125.00}$					864,595	874,005	867,546
$K_1 = 124.50$	13.98%	13.86 %	321.12	314.72			
$K_2 = 125.50$	13.72%	13.62 %	376.46	370.66			
$\boldsymbol{K = 130.00}$					1,872,779	1,906,804	1,863,207
$K_1 = 129.50$	12.74%	12.78 %	696.09	699.14			
$K_2 = 130.50$	12.61%	12.67 %	806.91	811.97			

Table 4: *Volatilities from both the SVI and SOP for a given strike, together with the vanilla put option price. The digital put contract price is based on the given vanilla put spread. For comparison, we also show the digital put contract price obtained from the MathFinance AG library (MF) by replication. All the prices in this table are in US dollar.*

Another important observation to make is that despite the level of the SOP volatility smile is *higher* than the SVI volatility smile for $K = 115.00$, the price of the digital put contract based on the SOP volatilities is actually *lower* than the SVI digital put contract value. It is clear that the *slope* is playing an important role here: both the slope and level of the inter- and extrapolated volatility smile at the strike enter the value of a digital.

Ignoring the effect of the slope means to keep the volatility within the vanilla put spread constant, which will result in a wrong valuation. In the previous FX column, we showed that the value of the digital contract within the Black-Scholes (BS) framework is given by

$$\begin{aligned} P^{\text{digital}} &= -\frac{\partial \mathsf{P}^{\mathsf{BS}}(K,T,\sigma(K))}{\partial K} \\ &= \underbrace{-p_K(\sigma(K))}_{\text{flat digital}} \underbrace{-p_\sigma(\sigma(K))\sigma'(K)}_{\text{windmill adjustment}}, \end{aligned} \tag{0.9}$$

where $p_K(\sigma(K)) := \frac{\partial P^{\mathsf{BS}}(K,T,\sigma(K))}{\partial K}$ is the dual delta of the vanilla option, $p_\sigma(\sigma(K)) := \frac{\partial P^{\mathsf{BS}}(K,T,\sigma(K))}{\partial \sigma}$ the vanilla vega and $\sigma'(K)$ is the slope of the volatility smile on the strike space. The first part of the second line in Equation (0.9) is the value of the digital with a flat volatility. The second part is the windmill adjustment. This comes from the fact that the volatility is also a function of the strike, resulting in an extra term when taking the derivative with respect to the strike. The vanilla vega is always positive, the slope of the smile on the strike space can be positive or negative. The windmill effect, for a negative slope, will lower the value of the digital put and increase the value of the digital call. For the digital put contract with strike $K = 115.00$, this result is summarized in Table (5).

For $K = 115.00$ the digital contract with a flat volatility (232,813 USD) is overestimating the value of the digital that takes the smile in consideration (153,410 USD). The windmill adjustment accounts for the smile impact on the digital contract (-78,671 USD), which results in more or less the same valuation as the replication (154,142 USD).

As we showed in the previous column, the price mismatch in the flat volatility replication is due to the fact that on one side the flat volatility is lower than the smile volatility, while on the other side the flat volatility is higher than the smile volatility. This results in different valuations

$P(K = 115.00)$	Formula	Value
Replication	$\frac{P^{\mathsf{BS}}(\frac{1}{K_1},\sigma(K_1))-P^{\mathsf{BS}}(\frac{1}{K_2},\sigma(K_2))}{\frac{1}{K_1}-\frac{1}{K_2}}$	153,410
Theoretical digital	$e^{-r_f T}\mathcal{N}(-d_+)$	232,813
Windmill Adjustment	$S_0 p_\sigma(\sigma(K))\sigma'(K)$	-78,671
Total Value		154,142

Table 5: *Comparison between the replicated digital contract (SVI) and the theoretical digital with the windmill adjustment. Prices are in US Dollar.*

of the put spread, and hence is the replication value different. In numbers, specifically for the calibrated SVI curve: $\sigma(K = 114.5) = 16.61\%$ which is higher than $\sigma(K = 115.00) = 16.48\%$, while $\sigma(K = 115.5) = 16.36\%$ is actually lower than the mid-point volatility.

The impact of the selection of market data input

As a case study from one of our industry projects, we show here how much care should be addressed to choosing the input market data. A data vendor advertised to provide volatility quotes for 5-delta points. One can then try to fit volatility curves through 7 points, see Figure 4.

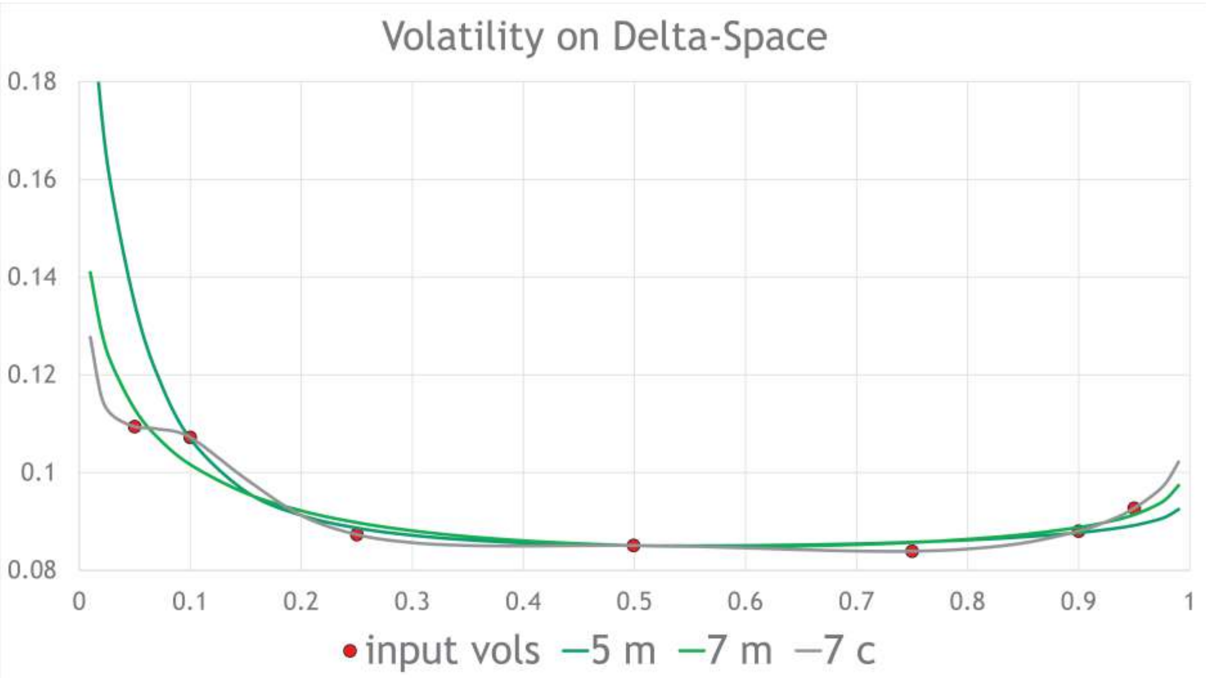

Figure 4: *Interpolation of volatility input (red dots). Put delta on the x-axis, implied volatility on the y-axis, 5m (the SVI version of the MathFinance library) takes on the inner 5 volatility data points into account and ignores the 5-delta volatility; 7m does the best fit using all volatility data points; 7c uses also all points, but cubic spline interpolation.*

On the left hand side, the 5-delta quote seems lower than expected. It doesn't quite seem to fit in the picture. In Sesame street we would have gotten the task to identify the point that does not fit in the series. A volatility smile construction should work for all cases, and as you can see here, the combination of the interpolation and input data can make so much of difference on the level and on the slope of the smile. It is clear that cubic splines create waves in the curve and can lead to windmill-effect terms riding a roller coaster. Which SVI to take instead depends on how serious the 5-delta put quote is. Solve the riddle: why is the point so low? (Answer - to be placed somewhere else in the journal - : The data vendor had used three contributors for 5-delta quotes. In the specific example, one of the three did not contribute a number, and the missing number was interpreted as a zero, which lowers the average).

Conclusion

The key to exotics pricing is the exact shape of the vanilla volatility smile surface, at the very least. While for vanilla options the level of volatility is the only thing that matters. For exotics, even simple digitals, the slope of the smile on the strike space contributes non-negligible amounts to the digital price. Therefore, is it crucial to interpolate and extrapolate carefully and be aware that all choices matter. Another interesting way of obtaining the volatility smile and pricing first-generation exotic options in the FX market would be by using the Vanna-Volga pricing [7].

We illustrated the impact on digitals. However, this translates to the entire range of exotics. A one-touch can be thought of as two digitals, and a reverse-knock-out can be replicated by a spread of regular knock-out options plus a no-touch. As a consequence the windmill-effect is one of the big contributions to pricing exotics with smile. We will cover this in detail in the next FX column.

Bibliography

[1] Malz, Allen. Estimating the probability distribution of the future exchange rate from option prices, Journal of Derivatives, Winter 1997, pp. 18–36.

[2] Wystup, Uwe and Reiswich, Dimitri. FX Volatility Smile Construction, Willmot, Vol. 2012, Issue 60, July 2012, pp. 58–69.

[3] Wystup, U. Negative Butterflies and Why We Check Butterfly Arbitrage by a Non-Negative Probability Density, Wilmott, volume 2022, issue 121, September 2022, pp. 48–49.

[4] Gatheral, Jim and Jacquier, Antoine, Arbitrage-Free SVI Volatility Surfaces, Quantitative Finance, Vol. 14, No. 1, pp. 59–71.

[5] Breeden, Douglas T. and Litzenberger, Robert H., Prices of State-Contingent Claims Implicit in Option Prices, The Journal of Business, Vol. 51, No. 4, October 1978, pp. 621–651.

[6] Wystup, U., Arbitrage in the Perfect Volatility Surface, Wilmott, issue 97, September 2018, pp. 16–17.

[7] Bossens, Frederic and Rayée, Grégory and Skantzos, Nikos S. and Deelstra, Griselda, Vanna-Volga Methods Applied to FX Derivatives: From Theory to Market Practice, International Journal of Theoretical and Applied Finance, Vol. 13(8), April 2010, pp. 1293–1324.

[8] Wystup, U. Vanna-Volga Pricing, June 2008, https://www.mathfinance.com/wp-content/uploads/2020/09/wystup_vannavolga_eqf.pdf