Model-free valuation of barrier options

Peter Austing and Yuan Li provide an analytic formula for valuing continuous barrier options. The model exactly fits the implied volatility smile in a manifestly arbitrage-free way and produces prices consistent with an underlying stochastic volatility dynamic

The standard approach to valuing continuous barrier options, developed in the literature by Jex (1999) and Lipton (2002), is to use a local stochastic volatility (LSV) model. Such models begin with an underlying stochastic volatility (SV) process, for example, stochastic alpha, beta, rho (SABR), Heston or a number of other choices. The SV is first calibrated to give vanilla option prices that closely match the market-implied volatility smile at one or more expiries. Then, the volatility of the SV (the volatility of volatility) is reduced by a proportion known as the mixing fraction, so the smiles generated are flatter than the market. Finally, a local volatility (LV) correction factor is introduced and calibrated to bring vanillas back to the correct market prices.

The mixing fraction measures how far the model is between SV (mixing zero) and LV (mixing one). An added bonus is that, in the mixing zero case, the LV correction can be used to fill in any mismatch between the pure SV smile and the true market smile. In practice, the mixing is tuned to match the price of a one-touch or double no-touch option, and it is often around one half.

The underlying SV process chosen varies considerably, yet market participants agree closely on barrier option prices. Thus, although the true underlying model is not known, this is not relevant as long as everyone believes in local stochastic volatility. Such a remarkable fact makes barrier options safe to trade in large volumes, and they are often treated almost as flow products.

Consider the reason for the approximate model invariance. It has been argued by Austing (2014) that this can be explained by the reflection principle. If we consider a pure normal SV model with constant rates, and without spot-volatility correlation, then the reflection principle holds; thus, valuation depends only on the terminal spot distribution, which is determined by the implied volatility smile. This remains approximately true in a lognormal model, and two SV models that generate similar terminal smile convexities will generate similar barrier option prices. Adding in a local volatility correction smoothly deforms the price towards the Dupire (1994) valuation, and so the impact is similar across the underlying SV models.

Such an argument is heuristic, and we shall not dwell upon it here. Rather, we take as our starting point the empirical observation that barrier prices are approximately model-independent within the class of SV models. Then, if prices do not actually depend on the complex stochastic spot/volatility dynamics, we are naturally curious to discover if we can remove it completely and provide an analytic formula for valuation.

The speed of modern computers, together with the realisation that a simple underlying SV is sufficient, means LSV models are now almost universally used for quoting barrier prices. Sometimes they are even used for trading desk risk management. However, enhancements in regulatory requirements mean trading books must be valued in simulations of increasing size and complexity, and accurate pricing within these simulations becomes more important. As a result, analytic approaches to valuation are as interesting as ever.

An analytic approach introduced by Lipton & McGhee (2002) and Wystup (2003), usually known as the Vanna-Volga method, works by thoughtfully adding the smile cost of second-order Greeks to the theoretical (flat smile) value. It was the workhorse for foreign exchange trading desks for a number of years, and particularly appreciated for the intuitive hedges it provided to traders. However, as the smile adjustment is built from a small number of vanillas, there are inevitable pricing anomalies (see Bossens *et al* 2010; Moni 2011).

Our approach will be to write down a joint probability distribution for the terminal spot S_T and the maximum value of spot in the interval $t \in [0, T]$, which we denote M_T. One starting point would be to use a simple copula, with marginals for S_T and M_T determined from the volatility smile and one-touch prices, respectively. This is problematic, because the copula would need to maintain the crucial relation $M_T \geqslant S_T$ to avoid arbitrage. Furthermore, as we know empirically that one-touch prices in SV are determined from the volatility smile, a solution built from a full marginal for M_T is likely to be over-determined.

Instead, we will write down a formula that determines both one-touch prices and more general barrier option prices directly from the volatility smile.

Constructing the distribution

We wish to construct a joint probability distribution for S_T and M_T with the property that $M_T \geqslant S_T$, with probability 1. As usual, we would like our prices to reduce to Black-Scholes when the volatility smile is flat; this motivates us to begin with a simple Brownian motion, W_t. We define the maximum to be $A_T = \sup_{t\in[0,T]} W_t$, and then the reflection principle tells us that:

$$P(W_T < x, A_T < y) = \begin{cases} 0, & y < 0 \\ N(x) + N(2y - x) - 1, & y > 0,\ y > x \\ 2N(y) - 1, & y > 0,\ y < x \end{cases} \tag{1}$$

where $N(x)$ is the cumulative normal function and $T = 1$.

It is worth taking a moment to examine (1). First, we have zero probability when $y < 0$. This is because the Brownian motion starts at $W_0 = 0$, so the maximum must always be greater. Later, we will have a similar requirement on the maximum versus initial spot, $M_T > S_0$. Second, when $y < x$ the cumulative probability is independent of x, so when we differentiate we find there is zero probability density in the region $W_T > A_T$.

To convert (1) into a copula, we first obtain the marginal distributions u_1 and u_2 by taking the large y and large x limits of (1). Inverting these gives us:

$$x = N^{-1}(u_1) \tag{2}$$

$$y = N^{-1}(\tfrac{1}{2}(u_2 + 1)) \tag{3}$$

leading to the copula:

$$C(u_1, u_2) = \begin{cases} u_1 + N[2N^{-1}(\frac{1}{2}(u_2+1)) - N^{-1}(u_1)] - 1, & \frac{1}{2}(u_2+1) > u_1 \\ u_2, & \frac{1}{2}(u_2+1) < u_1 \end{cases} \tag{4}$$

Let us suppose we are provided with marginal distributions $P(S_T < K) = F_1(K)$, derived from the implied volatility smile, and $P(M_T < B) = F_2(B)$ from one-touch prices. Then, we can attempt to create a joint distribution by setting:

$$u_1 = F_1(K) \tag{5}$$

$$u_2 = F_2(B) \tag{6}$$

and substituting these into (4).

If we now look at the region $\frac{1}{2}(u_2 + 1) < u_1$ in which there is zero density, this corresponds to:

$$\tfrac{1}{2}(F_2(B) + 1) < F_1(K) \tag{7}$$

which needs to be true if and only if $B < K$. We know from the reflection principle that F_1 places a heavy constraint on F_2, and, assuming F_1 and F_2 are continuous, the constraint in (7) tells us the marginal distribution of the maximum has to be given by:

$$F_2(B) = \begin{cases} 2F_1(B) - 1, & F_1(B) > \frac{1}{2} \\ 0, & \text{otherwise} \end{cases} \tag{8}$$

Thus, we have constructed a joint probability distribution for S_T, M_T, with the property that $M_T \geqslant S_T$, with probability 1. At this stage, the marginal for the spot S_T matches the input smile, while the marginal for the maximum M_T is a consequence of the construction and does not necessarily match one-touch prices. We have not yet succeeded in meeting the other constraint, $M_T > S_0$. Instead, the marginal distribution (8) tells us that $M_T > F_1^{-1}(\frac{1}{2})$.

To fix this, we will introduce two parameters, a drift and a volatility, into the Brownian motion from which we generated our copula. If we consider a new process:

$$\mathrm{d}X_t = \nu\,\mathrm{d}t + \sigma\,\mathrm{d}W_t \tag{9}$$

and set $A_T = \sup_{t\in[0,T]} X_t$, then standard methods show that the joint distribution (1) becomes:

$$P(x,y) = \begin{cases} 0, & y < 0 \\ N\left(\dfrac{x-\nu}{\sigma}\right) - \mathrm{e}^{2y\nu/\sigma^2} N\left(\dfrac{x-2y-\nu}{\sigma}\right), & y > 0,\ y > x \\ N\left(\dfrac{y-\nu}{\sigma}\right) - \mathrm{e}^{2y\nu/\sigma^2} N\left(\dfrac{-y-\nu}{\sigma}\right), & y > 0,\ y < x \end{cases} \tag{10}$$

The scaling parameter σ is irrelevant for constructing a copula and can be set to 1. Nevertheless, we shall retain it for reasons that will become apparent.

To simplify notation, we have denoted the joint distribution $P(X_T < x, A_T < y)$ by $P(x, y)$, and we denote its two marginals by:

$$P_X(x) = N\left(\frac{x-\nu}{\sigma}\right) \tag{11}$$

$$P_A(y) = N\left(\frac{y-\nu}{\sigma}\right) - \mathrm{e}^{2y\nu/\sigma^2} N\left(\frac{-y-\nu}{\sigma}\right) \tag{12}$$

Then, our joint distribution for S_T and M_T is obtained by substituting:

$$x = P_X^{-1}(F_1(K)) \tag{13}$$

$$y = P_A^{-1}(F_2(B)) \tag{14}$$

into (10). As before, the marginal distribution for the maximum is determined from the marginal for the terminal spot by imposing $x = y$ when $K = B$. This gives us:

$$F_2(B) = P_A P_X^{-1} F_1(B) \tag{15}$$

for the cumulative probability distribution of the maximum M_T, or equivalently:

$$y = P_X^{-1}(F_1(B)) \tag{16}$$

We need to choose ν so that the maximum cannot be lower than the initial spot, that is, $P(M_T < S_0) = 0$. Since the probability is zero when $y < 0$, this means we need to arrange things so that $y = 0$ when $B = S_0$. This is achieved by choosing:

$$\nu = -\sigma N^{-1} F_1(S_0) \tag{17}$$

To summarise, we have constructed a joint probability distribution between the terminal spot S_T and the maximum M_T. The distribution is given by (10), with x and y given by (13) and (16). The parameter σ scales and can be set to 1. The parameter ν is given by (17).

Matching stochastic volatility prices

So far, we have constructed a joint probability distribution between the spot at expiry, S_T, and the maximum of the spot between start time and expiry time, M_T. We can use this distribution to value options that have a continuous upper barrier and a European payout. The probability distribution exactly matches the market prices of vanilla options, since the marginal for S_T is an input that can be derived from the market smile. However, the values of one-touch options are given by the other marginal distribution, which is an output from the model and not necessarily expected to match market prices. Our aim in what follows is to extend the model so that one-touch prices become consistent with an underlying stochastic volatility model.

Let us denote the distribution we have constructed so far by $P_{\nu,\sigma}$. We recall that it is given by:

$$P_{\nu,\sigma}(K, B) = P(P_X^{-1}(F_1(K)), P_X^{-1}(F_1(B))) \tag{18}$$

where $P(x, y)$, $P_X(x)$ and F_1 are as defined above. As we have noted, the parameter σ scales and the parameter ν is determined by (17), so this distribution is unique.

1 Numerical demonstration of the model

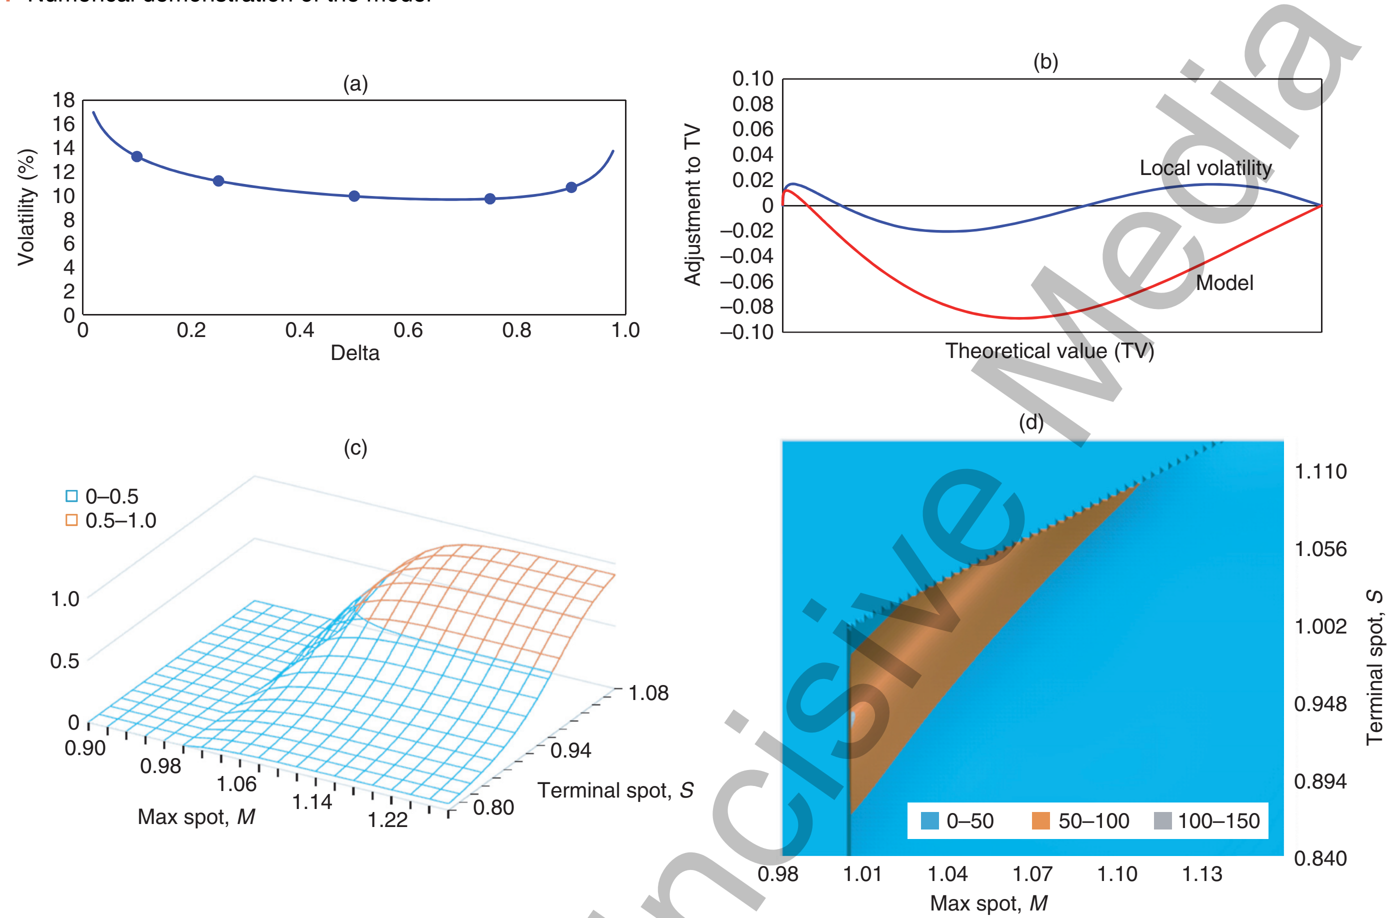

(a) Implied volatility smile. (b) One-touch plot. (c) Joint cumulative distribution. (d) Joint probability density

However, we can easily generate families of distributions by defining, for example:

$$Q_{\nu_1,\sigma_1;\nu_2,\sigma_2}(K,B) = \tfrac{1}{2}P_{\nu_1,\sigma_1}(K,B) + \tfrac{1}{2}P_{\nu_2,\sigma_2}(K,B) \tag{19}$$

We have four parameters: ν_1, σ_1, ν_2 and σ_2. As before, one parameter is redundant through scaling, and one is fixed by the requirement $P(M_T < S_0) = 0$. A helpful approach is to replace σ_1 and σ_2 with a single parameter α:

$$\sigma_1 = (1-\alpha)\sigma_{\text{atm}} \tag{20}$$

$$\sigma_2 = (1+\alpha)\sigma_{\text{atm}} \tag{21}$$

to remove the scaling, where σ_{atm} is the at-the-money volatility.

Returning to our earlier observation that LSV barrier prices are approximately independent of the details of the underlying SV model, we mention that this continues to work even if the SV is reduced to a simple mixture model. Here, the SV 'dynamic' is to select from two constant volatilities at time zero with equal probability. While Piterbarg (2003) elegantly demonstrates that this model is pathological for general contracts, it has long been observed that it generates correct barrier prices when incorporated into LSV and, indeed, as attested by Brigo *et al* (2015), it is used in production in a number of major financial institutions. As before, this is a consequence of the reflection principle.

In order to get realistic SV prices, we are going to extend our original probability distribution to that generated by a mixture model. Reminiscent of Brigo & Mercurio (2000), we employ this approach not to use a mixture model directly, but rather to help us generate an appropriate and tractable model. We note at the outset that the following approach continues to work if we extend to a continuum of volatilities. This corresponds to integrating over the terminal volatility distribution of an arbitrary (correlation-free) SV model, with an approximation, as the reflection principle is not exact in the lognormal case.

We can set up our distribution to mimic the mixture model and, therefore, be sure to generate realistic SV barrier prices. To do so, we choose α so that a mixture model with two volatility states σ_1 and σ_2 generates smile convexity to match the market-implied volatility smile. For our numerical results, we choose to define smile convexity using market 25-delta and 75-delta strikes, K_{25}, K_{75}, by calculating:

$$\text{ST}_{25} = \tfrac{1}{2}(\sigma(K_{25}) + \sigma(K_{75})) - \sigma_{\text{atm}}$$

where σ is the market- or model-implied volatility. We then solve numerically for α so that:

$$\text{ST}_{25}^{\text{model}}(\alpha) = \text{ST}_{25}^{\text{market}} \tag{22}$$

Since valuing vanilla options in the mixture model is a simple matter of averaging the Black-Scholes formula over the two volatilities, this is a simple procedure.

2 One-touch prices and densities for the model and two additional values of α

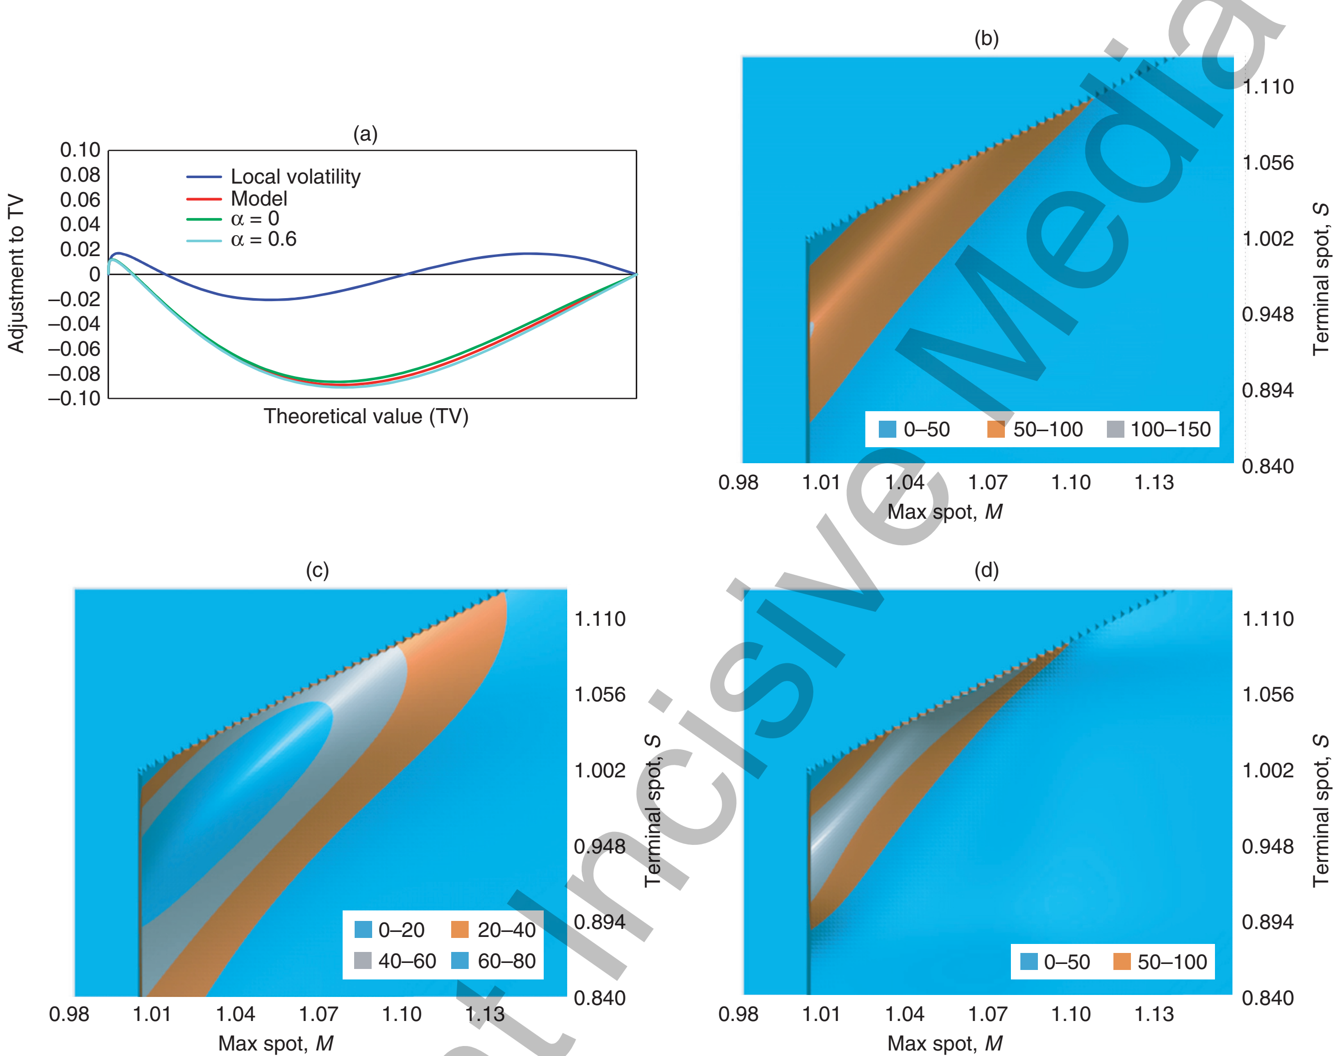

(a) One-touch plot. (b) Joint probability density. (c) Joint probability density: $\alpha = 0$. (d) Joint probability density: $\alpha = 0.6$

Having fixed α and, as a result, σ_1 and σ_2, we choose:

$$\nu_1 = \beta - \tfrac{1}{2}\sigma_1^2 \tag{23}$$

$$\nu_2 = \beta - \tfrac{1}{2}\sigma_2^2 \tag{24}$$

and solve numerically for β to achieve $P(M_T < S_0) = 0$. This fully determines the parameters in the joint probability distribution (19), and so allows us to price barrier options. If the true underlying market were a mixture model, then this procedure ensures our constructed probability distribution would exactly match that market. As a result, since a mixture market adequately generates SV-style barrier prices, our model will do the same.

In order to demonstrate the model numerically, we set up a simple volatility surface in a spreadsheet and generated the classic one-touch plots of Jex (1999). For simplicity, we defined our implied volatility at strike K and expiry time T using three parameters, a, b, c:

$$\sigma^2_{\text{implied}}(K, T) = \sqrt{ay^2 + by + c} \tag{25}$$

where $y = \log(K/F_T)$ is the log moneyness with respect to the forward F_T.[1] Our plots were generated with expiry $T = 1$ year, spot $S_0 = 1$, risk-free rate 3% and dividend yield 5%. The volatility parameters were $a = 0.005$, $b = -0.0005$ and $c = 0.0001$, giving an at-the-money volatility of 10%, a 25-delta strangle of 0.54% and a 25-delta risk-reversal of −1.5%. The results are shown in figure 1. With the specified market data, we calculated $\alpha = 0.44$ and $\beta = -0.014$.

Looking at the one-touch plot in figure 1, we see the familiar moustache shape, with the model price lower than the local volatility price; this is to be expected, since our model is a proxy for stochastic volatility. We have also plotted the joint cumulative distribution between M_T and S_T and the joint density function. It is interesting to note that there is a sharp change in density from zero along the lines $M_T = S_0 \equiv 1$ and $M_T = S_T$. This

[1] *This implied volatility parameterisation was chosen to be as simple as possible while satisfying the Lee bounds.*

makes sense, as, for example, the line $M_T = S_0$ represents those paths that spend a brief time around S_0 before heading downwards.

The parameter α is fixed to generate stochastic volatility-style prices by imposing (22). While the ability to obtain SV barrier prices analytically is a significant step forward, the true market price of a barrier option usually lies between the stochastic volatility price and the local volatility price. The most sophisticated market participants use local-stochastic volatility models to achieve a mixing between SV and LV. This invites the question of whether we can tune α in order to find a proxy for the LSV price.

Figure 2 shows what happens when we vary α. We note that one-touch prices do not change much. In particular, even when α is zero, so that our underlying mixture model has no stochasticity, the one-touch prices do not come close to the local volatility prices. This shows that the parameter α does not act like a mixing between SV and LV. This should not be surprising, as there is no reason to believe the original copula we constructed earlier in this article would provide a proxy for LV prices.

However, the joint density function, also shown in figure 2, varies dramatically with α. This is an important point. It means we can construct families of models that price vanilla options correctly, and match one-touch prices fairly closely, yet in which other barrier contracts vary significantly. Thus, the choice of α to satisfy (22) is a crucial part of our model. It is this choice that allows us to rely on the argument of the underlying mixture model to obtain meaningful barrier option prices.

An alternative way of generating prices close to LSV is to create a distribution that mixes true local volatility (calculated by a numerical partial differential equation solution) with the SV proxy from our model. In practice, this is as simple as calculating the price in LV and mixing with the price from our model. While this may leave some feeling uneasy due to the Piterbarg (2003) hangover, there is probably no need for such qualms. It is a legitimate means of creating a joint terminal probability distribution, and it does not suffer from the temporal arbitrage problem.

Conclusion

We have developed an analytic model that generates barrier option prices that approximate stochastic volatility models. The model provides the joint cumulative probability distribution of S_T and M_T, the terminal spot and maximum spot, respectively. Thus, certain contracts, such as one-touches and digital payouts with a continuous barrier, can be valued analytically. Vanilla payouts with continuous barriers can be valued fast with a single numerical integration.

The model takes the market-implied volatility smile as an input and ensures the vanilla options are exactly repriced. This is reminiscent of local stochastic volatility, in which the local volatility correction adjusts for any mismatch between the underlying SV smile and the true market smile.

We have emphasised two main assumptions in our model. First, we rely on the reflection principle and, therefore, have implicitly approximated the drift in any underlying SV model by a constant. This means we are assuming interest rates are not strongly varying over time. Similarly, since we assume a lognormal underlying spot process, volatility contributes to the drift and this means we are assuming individual volatility paths generated by the SV are not too strongly varying in time.

The second assumption is that we used an implicit underlying mixture model to generate the smile convexity. We used the fact that this underlying model is known to generate SV-style barrier prices to argue that our model is sensible. However, as the mixture model does not have spot-volatility correlation, it does not generate any skew. Therefore, the skew fit is achieved entirely through the copula, without reference to any underlying dynamic model. The full impact of this remains open, and it would be interesting to extend to SV models with correlation using the methods of McGhee & Trabalzini (2014).

In addition to providing an arbitrage-free alternative to Vanna-Volga and a simple means of generating LSV-style prices by mixing with local volatility, it is pleasing to see a simple analytic explanation for the classic one-touch moustaches. Indeed, those who are energetic enough could use the formula to generate continuous vanilla hedge portfolios to improve on the classic trader rules. ■

Peter Austing is a quantitative researcher at Citadel in London. Yuan Li is an associate director within the FX quant team at ANZ in London. **The authors are grateful to Anqi Zhou and Minying Lin for indispensable discussions. The opinions expressed in this paper are those of the authors, and do not necessarily represent the views of their employers.**
Emails: peter.austing@citadel.com, yuan.li2@anz.com.

REFERENCES

Austing P, 2014
Smile Pricing Explained
Palgrave Macmillan

Bossens F, R Grégory, NS Skantzos and G Deelstra, 2010
Vanna-Volga methods applied to FX derivatives: from theory to market practice
International Journal of Theoretical and Applied Finance 13(8), pages 1293–1324

Brigo D and F Mercurio, 2000
A mixed-up smile
Risk September, pages 123–126

Brigo D, C Pisani and F Rapisarda 2015
The multivariate mixture dynamics model: shifted dynamics and correlation skew
Working Paper, SSRN

Dupire B, 1994
Pricing with a smile
Risk January, pages 18–20

Jex M, 1999
Pricing exotics under the smile
Risk November, pages 72–75

Lipton A, 2002
The vol smile problem
Risk February, pages 61–65

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

McGhee WA and R Trabalzini, 2014
Deconstructing the volatility smile
Working Paper, SSRN

Moni C, 2011
Two little Vannas go to market
Preprint, SSRN

Piterbarg VV, 2003
Mixture of models: a simple recipe for a hangover
Wilmott Magazine, pages 72–77

Wystup U, 2003
The market price of one-touch options in foreign exchange markets
Derivatives Week 12(13), pages 8–9