Expanded forward volatility

*Using a one time-step finite difference implementation, **Jesper Andreasen** and **Brian Huge** eliminate the arbitrages in the wings of the volatility smile that result from most expansion techniques for local stochastic volatility models, including the widely used SABR model*

Modelling the implied volatility smile using local and stochastic volatility has been the subject of much research over the past 20 years (see, for example, Dupire, 1996, Hagan *et al*, 2002, Heston, 1993, Jex, Henderson & Wang, 1999, Lewis, 2000, and Lipton, 2002). Interest rate option desks typically need to maintain very large amounts of interlinked volatility data. For each currency, there might be 20 expiries and 20 tenors, that is, 400 volatility smiles. Furthermore, the smiles might be linked across different currencies. Interpolation of observed discrete quotes to a continuous curve is needed for the pricing of general caps and swaptions. At the same time, extrapolation of option quotes are needed for constant maturity swap (CMS) pricing. For these purposes, the industry standard has been the SABR model using expansions as in Hagan *et al* (2002). The implied volatility expansions have the advantages that they are fast and simple to code.

But the expansions are not very accurate, particularly not for long maturities or low strikes. Numerical examples of this can be found in Paulot (2009). This is, however, largely irrelevant as the SABR expansions are generally only used for the pricing of European options, and not, for example, for calibration of full dynamic models.

The main practical problem with the expansions is that they imply negative densities for low strikes and occasionally also for high strikes. With the low rates we have today, this problem is more acute than ever. Furthermore, the SABR model only has four parameters to handle the above-mentioned tasks, which is not enough flexibility to exactly fit all option quotes.

As in Balland (2006) and Lewis (2007), we extend the stochastic volatility process to include a constant elasticity of variance (CEV) skew on the volatility of volatility. The CEV volatility process allows us to have more explicit control of the extrapolated high-strike volatilities, which in turn allows better control of CMS prices. Further, we will use a non-parametric volatility function for the spot process, which enables us to have an exact fit to all the observed quotes and gives us the ability to model negative option strikes.

Rather than going through heat kernel expansions as in Hagan *et al* (2002), we follow Balland (2006) and use a short-maturity expansion for the implied volatility of the option. The short-maturity expansion also yields results for the short-maturity limit of the Dupire (1996) forward volatility, that is, the short-maturity limit of the conditional expected local variance:

$$\vartheta(k)^2 = \lim_{t\downarrow 0} E\left[ds(t)^2 / dt \,\middle|\, s(t) = k \right]$$

The forward volatility is used in a single time step implicit finite difference discretisation of the Dupire (1994) forward partial differential equation (PDE). This precludes arbitrage and so avoids negative densities. We also derive an adjustment of the forward volatility function to compensate for the pricing in one, rather than multiple, time steps. The single-step finite difference grid generates all prices in one go and this can in turn be used for calibrating the model directly to observed CMS prices.

We provide two calibration procedures: an implicit method that works by iteration of the connection from parameters to prices in a non-linear solver, and a direct method that given an arbitrage-free continuous curve of option prices directly infers the parameters of the model. We show that the implicit calibration method can be used to fit the model to 10 discrete strikes in approximately one millisecond of CPU time.

Finally, we show how we are able to control the wings of the smile by varying the functional form of the diffusion of the spot and volatility processes, and the impact this has on CMS pricing.

Short maturity expansion

First, we will outline the short-maturity expansion. Our approach is similar to that used in Balland (2006). We consider the model:

$$\begin{aligned} ds &= z\sigma(s)dW \\ dz &= \varepsilon(z)dZ \end{aligned} \qquad (1)$$

where W and Z are Brownian motions with correlation ρ.

The non-parametric form of the volatility function $\sigma(s)$ allows us to have a perfect fit to any discrete or continuous set of observed arbitrage-free option quotes.

We can write the price of a European call option on a fixing $s(T)$ as:

$$c(t) = E_t\left[\left(s(T) - k \right)^+ \right] = g\left(t, s(t), \nu(t)\right)$$

where $\nu(t)$ is the implied normal volatility and g is the normal (Bachelier) option pricing formula:

$$g(\tau, s, \nu) = (s-k)\Phi\left(\frac{s-k}{\nu\sqrt{\tau}}\right) + \nu\sqrt{\tau}\phi\left(\frac{s-k}{\nu\sqrt{\tau}}\right), \quad \tau = T - t \qquad (2)$$

Applying Itô's lemma to (2) yields:

$$dc = -g_\tau dt + g_s ds + \frac{1}{2} g_{ss} ds^2 + g_\nu d\nu + \frac{1}{2} g_{\nu\nu} d\nu^2 + g_{s\nu} ds \cdot d\nu \qquad (3)$$

where subscripts denote partial derivatives. In the following, we assume $\nu > 0$.

Define $x = (s - k)/\nu$. Using Itô's lemma yields:

$$\begin{aligned} dx &= \frac{1}{\nu}ds - \frac{s-k}{\nu^2}d\nu - \frac{1}{\nu^2}ds \cdot d\nu + \frac{s-k}{\nu^3}d\nu^2 \\ &= \frac{1}{\nu}\left(ds - x d\nu\right) + O(dt) \\ dx^2 &= \frac{1}{\nu^2}\left(ds^2 + x^2 d\nu^2 - 2x ds \cdot d\nu\right) \end{aligned} \quad (4)$$

The normal option pricing function, g, has the following properties:

$$\begin{aligned} g_\nu &= \nu\tau g_{ss} \\ g_{\nu\nu} &= \left(\frac{s-k}{\nu}\right)^2 g_{ss} \\ g_{s\nu} &= -\frac{s-k}{\nu} g_{ss} \\ 0 &= -g_\tau + \frac{1}{2}\nu^2 g_{ss} \end{aligned} \quad (5)$$

Using the properties in (5) we can transform (3) to:

$$dc - g_s ds = \frac{1}{2} g_{ss}\left[\nu^2\left(dx^2 - dt\right) + 2\tau\nu d\nu\right] \quad (6)$$

The left hand side of (6) is the change in value of a hedged portfolio. Taking conditional expectations yields:

$$0 = \frac{1}{2} g_{ss}\nu^2 E_t\left[dx^2 - dt\right] + g_{ss}\tau\nu E_t\left[d\nu\right] \quad (7)$$

As $g_{ss} > 0$ for $\nu > 0$, and for any diffusion, $E_t[dx^2 - dt] = 0$ is equivalent to $dx^2 = dt$, we obtain the condition:

$$0 = \left(dx^2 - dt\right) + 2\frac{\tau}{\nu}E_t\left[d\nu\right] \quad (8)$$

For small maturities, $\tau \to 0$, we arrive at the arbitrage condition:

$$\sigma_x^2 \equiv \frac{dx^2}{dt} = 1 \quad (9)$$

Note that this is a diffusion condition rather than the drift conditions that we normally see in financial mathematics.

As x must be a function of the state variables (s, z), the diffusion condition (9) leads to the differential equation:

$$\begin{aligned} 1 = \sigma_x^2 &= \left(x_s ds + x_z dz\right)^2 / dt \\ &= z^2\sigma(s)^2 x_s^2 + \varepsilon(z)^2 x_z^2 + 2\rho z\sigma(s)\varepsilon(z) x_s x_z \end{aligned} \quad (10)$$

Given the functions σ, ε, we need to solve this non-linear first-order differential equation subject to the boundary condition $x(s = k, z) = 0$. Once we have the solution $x(s, z)$, we can find the implied normal volatility as:

$$\nu = \frac{s-k}{x(s,z)} \quad (11)$$

We note that the error of the implied volatility is $O(\tau)$.

The result implies that for any choice of $\sigma(s)$, $\varepsilon(z)$ any function $x = x(s, z)$ that satisfies $dx^2 = dt$ leads to an implied volatility given by $\nu = (s - k)/x$.

We could have chosen to derive the short-maturity expansion in implied Black-Scholes (lognormal) volatility $\bar{\nu}$ instead of implied normal volatility. Instead of x we should then have chosen the transformation $\bar{x} = \ln(s/k)/\bar{\nu}$. The diffusion condition would be the same so $\bar{x} = x$. This relates short-maturity implied lognormal and normal volatilities by the simple relationship:

$$\frac{\bar{\nu}}{\nu} = \frac{\ln s/k}{s-k} \quad (12)$$

The expansion results that we present in the following can easily be switched between use in implied normal and implied lognormal volatility form by use of the relation (12).

Deterministic volatility

First, we will consider the case with $\varepsilon(z) = 0$. In this case, $z \equiv 1$, and the differential equation (10) reduces to the ordinary differential equation (ODE):

$$x_s^2\sigma(s)^2 = 1 \quad (13)$$

Using the boundary condition $x(s = k) = 0$, we find the solution:

$$x = \int_k^s \frac{1}{\sigma(u)} du \quad (14)$$

with corresponding implied normal and Black volatilities given by:

$$\begin{aligned} \nu &= \frac{s-k}{\int_k^s \sigma(u)^{-1} du} \\ \bar{\nu} &= \frac{\ln(s/k)}{\int_k^s \sigma(u)^{-1} du} \end{aligned} \quad (15)$$

These results appears in many places, for example in Andersen & Ratcliffe (2002).

We note that (14) implies the following relationship between x and the forward volatility:

$$\frac{\partial x}{\partial k} = -\frac{1}{\sigma(k)} \quad (16)$$

Suppose we have x from a stochastic volatility model like (1), that is, given as the solution to (10) for some volatility functions $\sigma(s)$, $\varepsilon(z)$ and correlation ρ. Define the function ϑ by:

$$\vartheta(k) = -\left(\frac{\partial x}{\partial k}\right)^{-1} \quad (17)$$

and consider the deterministic volatility model:

$$ds = \vartheta(s) dW \quad (18)$$

It now follows that:

$$x^{LV} \equiv \int_k^s \vartheta(u)^{-1} du = x \quad (19)$$

So the stochastic volatility model (1) and the local volatility model (18) will produce the same short-maturity expansion option prices.

The above is a short-maturity limit version of the general result by Gyongy (1986) and Dupire (1996) that the model:

$$d\bar{s} = a(t, \bar{s}) dW, \quad \bar{s}(0) = s(0) \quad (20)$$

1 Short maturity expansion implied volatility

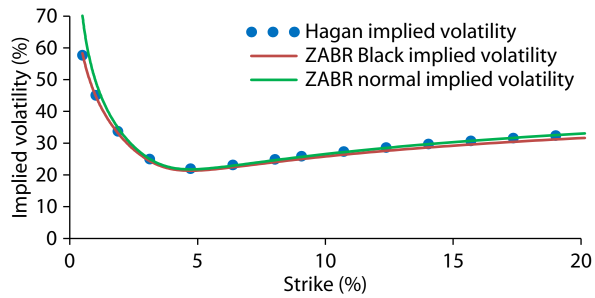

Note: implied Black volatility for our short maturity expansions in Black (red) or normal volatility (green) compared with Hagan's formula (dotted blue)

2 Comparison of σ(·) and ϑ(·)

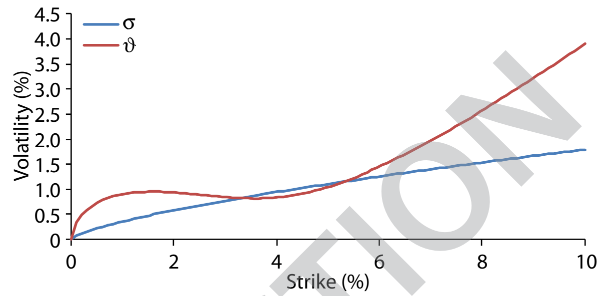

Note: volatility $\sigma(s) = 0.0873 \cdot s^{0.7}$ (blue) compared with the forward volatility $\vartheta(k) = J(y)z\sigma(k)$ (red) for the case of $\varepsilon = 0.47$, $\rho = -0.48$

produces the same option prices as the model (1) if $a(\cdot)$ is chosen to be:

$$a(t,k)^2 = E\left[ds(t)^2 / dt \,\middle|\, s(t) = k\right] \tag{21}$$

We conclude that in the short-maturity limit, the conditional expected variance of the underlying is related to the transformed variable x by:

$$\vartheta(k)^2 = \lim_{t \downarrow 0} E\left[ds(t)^2 / dt \,\middle|\, s(t) = k\right] = \left(\frac{\partial x}{\partial k}\right)^{-2} \tag{22}$$

This constitutes a way of relating the two dimensional pricing problem (1) to the simpler one-dimensional pricing problem (18). We will make use of this relationship to generate arbitrage-free prices later in this article.

The SABR model

Here, we will rederive the main result of Hagan *et al* (2002) by solving the diffusion condition for the lognormal volatility process case, $\varepsilon(z) = \varepsilon \cdot z$.

First, we use the transformation:

$$y = \frac{1}{z}\int_k^s \frac{1}{\sigma(u)} du \tag{23}$$

to get:

$$\begin{aligned} dy &= dW - \varepsilon y dZ + O(dt) \\ &= \left[1 + \varepsilon^2 y^2 - 2\rho\varepsilon y\right]^{1/2} dB + O(dt) \\ &\equiv J(y) dB + O(dt) \end{aligned} \tag{24}$$

where B is a new Brownian motion. As $y(s = k) = 0$, we can now get x by normalising the volatility of y, hence:

$$\begin{aligned} x &= \int_0^y J(u)^{-1} du = \frac{1}{\varepsilon} \ln \frac{J(y) - \rho + \varepsilon y}{1 - \rho} \\ \nu &= \frac{s - k}{x} \\ \bar{\nu} &= \frac{\ln(s/k)}{x} \end{aligned} \tag{25}$$

For the CEV case $\sigma(s) = \sigma_0 \cdot s^\beta$, we have:

$$y = \frac{1}{z\sigma_0} \frac{s^{1-\beta} - k^{1-\beta}}{1 - \beta} \tag{26}$$

These formulas are basically the result of Hagan *et al* (2002). This is extended to include maturity and various refinements for the CEV case. The Hagan result does, however, produce implied volatility smiles that are virtually identical to those produced with formula (25). In figure 1, we compare the Hagan expansion with (25). To be precise, the Hagan expansion used here and in the following is (2.17) in Hagan *et al* (2002).

We can use (25) to retrieve the forward volatility function of SABR from:

$$\frac{\partial x}{\partial k} = \frac{\partial x}{\partial y}\frac{\partial y}{\partial k} = \frac{1}{J(y)}\left(-\frac{1}{z}\frac{1}{\sigma(k)}\right) \tag{27}$$

Hence:

$$\vartheta(k) = J(y) z \sigma(k) \tag{28}$$

This result can also be deduced from results in Doust (2010).

Figure 2 compares the function $\sigma(s)$ to the forward volatility function $\vartheta(k)$.

The ZABR model

Next, we consider the extended SABR model where the volatility process is of the CEV type $\varepsilon(z) = \varepsilon \cdot z^\gamma$.

Again, we will introduce an intermediate variable:

$$y = z^{\gamma-2} \int_k^s \frac{1}{\sigma(u)} du \tag{29}$$

For which Itô expansion yields:

$$dy = z^{\gamma-1}\left(dW + (\gamma - 2)\varepsilon y dZ\right) + O(dt) \tag{30}$$

Define $x = z^{1-\gamma} f(y)$, for some function $f(y)$, and we get:

$$\begin{aligned} dx &= z^{1-\gamma} f'(y) dy + (1-\gamma)\varepsilon f(y) dZ + O(dt) \\ &= f'(y) dW + \left[(\gamma - 2)\varepsilon y f'(y) + (1-\gamma)\varepsilon f(y)\right] dZ + O(dt) \end{aligned} \tag{31}$$

We conclude that the diffusion condition (8) is satisfied if f solves the ODE:

3 Control of high-strike behaviour

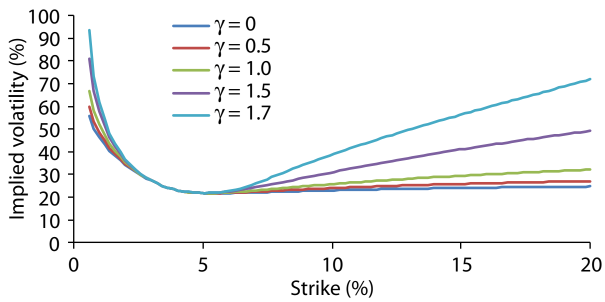

Note: implied Black volatility smiles for different choices of γ. Other parameters the same as in figure 1

4 Problem of negative density

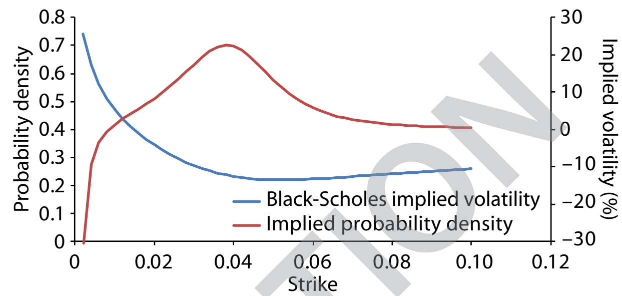

Note: density (red) and implied Black-Scholes volatilities (blue) from the Hagan expansion. Parameters as in figure 1

$$\begin{aligned}
1 &= A(y)f'(y)^2 + B(y)f(y)f'(y) + Cf(y)^2 \\
A(y) &= 1 + (\gamma-2)^2\varepsilon^2y^2 + 2\rho(\gamma-2)\varepsilon y \\
B(y) &= 2\rho(1-\gamma)\varepsilon + 2(1-\gamma)(\gamma-2)\varepsilon^2 y \\
C &= (1-\gamma)^2\varepsilon^2 \\
f(0) &= 0
\end{aligned} \tag{32}$$

The ODE (32) can be rearranged as:

$$f'(y) = \frac{-B(y)f + \sqrt{B(y)^2f^2 - 4A(y)\left(Cf^2-1\right)}}{2A(y)} \equiv F(y,f) \tag{33}$$

which can be solved by standard techniques for the integration of ODEs. We can evaluate the solution for all strikes in one sweep by:

$$\begin{aligned}
\frac{\partial y}{\partial k} &= -z^{\gamma-2}\sigma(k)^{-1} \\
\frac{\partial x}{\partial k} &= z^{1-\gamma}\frac{\partial f}{\partial k} = z^{1-\gamma}f'(y)\frac{\partial y}{\partial k} = -z^{-1}F\left(y,z^{\gamma-1}x\right)\sigma(k)^{-1} \\
x(k=s) &= y(k=s) = 0
\end{aligned} \tag{34}$$

Again, we can find the forward volatility function as:

$$\vartheta(k) = -\left(\frac{\partial x}{\partial k}\right)^{-1} = z\sigma(k)f'(y)^{-1} = z\sigma(k)F\left(y,z^{\gamma-1}x\right)^{-1} \tag{35}$$

Equations (34) and (35) will typically be evaluated at $z = z(0) = 1$. Rather than numerically solving the two ODEs in (34) separately, we favour solving (32) as a joint system.

Increasing γ lifts the wings of the implied volatility smile whereas the implied volatility smile for strikes close to at-the-money are virtually unaffected. This is illustrated in figure 3. This can in turn be used to give us better control over the CMS prices.

It should here be noted that the ODE representation (32) has previously been obtained by Balland (2006) for the lognormal case $\sigma(s) = \sigma \cdot s$. Further, it should be noted that Henry-Labordère (2008) has a treatment of the general non-CEV case $\varepsilon = \varepsilon(z)$.

For quick identification of the model parameters, the following second-order Taylor expansion is convenient:

$$\begin{aligned}
\nu(k) &= \nu(s) + \nu'(s)(k-s) + \frac{1}{2}\nu''(s)(k-s)^2 + O\left((k-s)^3\right) \\
\nu(s) &= z\sigma(s) \\
\nu'(s) &= \frac{1}{2}\left[z^{\gamma-1}\rho\varepsilon + z\sigma'(s)\right] \\
\nu''(s) &= \frac{1}{6z\sigma(s)}\Big[z^{2\gamma-2}\left((-5+2\gamma)\rho^2+2\right)\varepsilon^2 \\
&\qquad + z^2\left(2\sigma(s)\sigma''(s) - \sigma'(s)^2\right)\Big]
\end{aligned} \tag{36}$$

For the CEV case $\sigma(k) = \omega \cdot (k - \underline{s})^\beta/(s - \underline{s})^\beta$ and $z = 1$ we have:

$$\begin{aligned}
\nu(s) &= \omega \\
\nu'(s) &= \frac{1}{2}\left[\rho\varepsilon + \omega(s-\underline{s})^{-1}\beta\right] \\
\nu''(s) &= \frac{1}{6\omega}\left[\left((-5+2\gamma)\rho^2+2\right)\varepsilon^2 + \omega^2(s-\underline{s})^{-2}\beta(\beta-2)\right]
\end{aligned} \tag{37}$$

For a given set of discrete quotes $\hat{\nu}(k_1), \ldots, \hat{\nu}(k_n)$, the Taylor expansion (37) can be used for regressing the triple $\nu(s)$, $\nu'(s)$, $\nu''(s)$. One can in turn solve (37) to get parameter estimates for β, ρ, ε.

Finite difference volatility

Using the implied volatility coming from the short-maturity expansions, (15), (25) and (34), directly for pricing using (2) will not give arbitrage-free option prices. Our short-maturity expansions suffer from the same problem of potential negative implied densities for low strikes as the original Hagan expansion. In figure 4, we have plotted an example of the implied volatilities and the implied density coming from the Hagan expansion.

To avoid this problem, we will instead use the forward volatilities derived in (28) and (35) as the basis for our pricing.

The forward volatility $\vartheta(k)$ can be used to generate option prices as the solution to the Dupire (1994) forward PDE:

$$c_t(t,k) = \frac{1}{2}\vartheta(k)^2 c_{kk}(t,k)$$
$$c(0,k) = (s-k)^+ \tag{38}$$

The usual way of solving this numerically is to set up a time discretisation with multiple time steps and then use a finite difference solver. However, to gain speed we will instead use the single time step implicit finite difference approach introduced in Andreasen & Huge (2011). Here we need to solve the ODE:

$$c(t,k) - \frac{1}{2}t\theta(k)^2 c_{kk}(t,k) = (s-k)^+ \tag{39}$$

In Andreasen & Huge (2011) it is shown that this approach generates a set of arbitrage-free call prices for any choice of θ. It is also shown that the one-step finite difference price is the Laplace transform of the solution to (38). The Laplace transform of the Gaussian distribution is the Laplace distribution:

$$\int_0^\infty e^{-t/T}\frac{1}{\nu\sqrt{t}}\phi\left(\frac{s-k}{\nu\sqrt{t}}\right)dt = \sqrt{\frac{T}{2\nu^2}}e^{-|s-k|\sqrt{\frac{T}{2\nu^2}}} \tag{40}$$

which is peaked at $k = s$. Therefore if we choose $\theta = \vartheta$ we will also get a peak in the densities.

Instead, we will find an adjustment for the forward volatility function based on our expansion results. As option prices generated by (38) and (39) should be the same, we can substitute $c_{kk} = 2c_t/\vartheta^2$ from (38) into (39) and rearrange to find:

$$\begin{aligned}\theta(k)^2 &= \vartheta(k)^2\frac{c(t,k)-(s-k)^+}{tc_t(t,k)}\\ &\approx \vartheta(k)^2\frac{\left(g(t,s,\nu)-(s-k)^+\right)/t}{\partial g(t,s,\nu)/\partial t}\\ &= \vartheta(k)^2\cdot 2\left(1-\xi\frac{\Phi(-\xi)}{\phi(\xi)}\right),\quad \xi = |x|t^{-1/2}\\ &\equiv \vartheta(k)^2\cdot P(x)^2\end{aligned} \tag{41}$$

where the second (approximate) equality involves approximating the option prices by our expansion result.

The function $P(x)^2$ can conveniently be approximated with a third- or fifth-order polynomial. Specifically:

$$\Phi(x)/\phi(x) \approx \sum_n a_n u^n,\quad u = 1/(1+px) \tag{42}$$

where the constants p, a_1, a_2, ... can be found in (26.2.16) and (26.2.17) of Abramowitz & Stegun (1972), .

The finite difference discretisation of (39) is:

$$\left[1-\frac{1}{2}t\theta(k)^2\delta_{kk}\right]c(t,k) = (s-k)^+ \tag{43}$$

where δ_{kk} is the second-order difference operator. This equation can be represented as a tridiagonal matrix equation on the grid $\{k_0, k_1, \ldots, k_n\}$, which in turn can be solved for $\{c(t, k_i)\}$ in linear CPU time using the tridag() algorithm in Press *et al* (1992).

5 Probability density and forward volatility

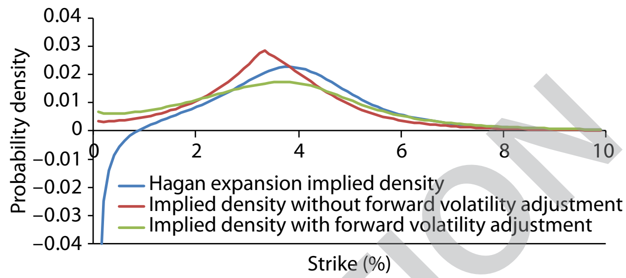

Note: the density with (green) and without (red) the forward volatility adjustment compared with the implied density of the Hagan

As an alternative to the finite difference solution (43), one could use the exact solution methodology for the ODEs of the type (39) described in Lipton & Sepp (2011). However, for this methodology to be computationally effective the forward volatility function θ(k) needs to be well approximated by a piecewise linear function with few knot points over the full domain of the solution. This is generally not the case here, as can be seen in figure 2. We have therefore chosen to base our solution on (43).

In figure 5, we have plotted the density both with and without the forward volatility adjustment. For reference, we have also plotted the implied density from the Hagan expansion. We see that the finite difference generated option prices have corresponding implied densities that are positive, that is, arbitrage is precluded. We also see that using our forward volatility result, $\vartheta(k)$, directly in the single time step finite difference solver produces a density that is peaked around at-the-money. This, however, is eliminated when using the adjusted forward volatility $\theta(k)$.

Calibrating the volatility function

First consider the case where we have a continuous curve of arbitrage-free option prices. This could for example be produced by the Andreasen & Huge (2011) interpolation scheme or come from another ZABR model. Calculate the forward volatility function by the discrete Dupire equation:

$$\theta(k)^2 = 2\frac{c(t,k)-(s-k)^+}{t\delta_{kk}c(t,k)} \tag{44}$$

Using (35) we can calibrate the volatility function:

$$\sigma(k) = \frac{F\left(y, z^{\gamma-1}x\right)\theta(k)}{zP(x)} \tag{45}$$

where x, y are found from (34) as the solution to the ODE system:

$$\begin{aligned}\frac{\partial y}{\partial k} &= -\frac{z^{\gamma-1}P(x)}{\theta(k)F\left(y,z^{\gamma-1}x\right)}\\ \frac{\partial x}{\partial k} &= -\frac{P(x)}{\theta(k)}\\ y(k=s) &= x(k=s) = 0\end{aligned} \tag{46}$$

The ODE system (46) can be solved for all strikes in one sweep.

6 Low-strike behaviour

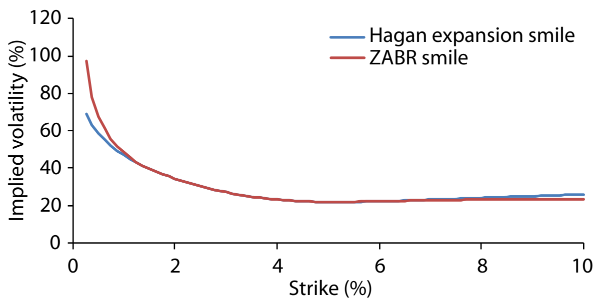

Note: implied Black volatilities from the ZABR model (red) that is calibrated to the Hagan expansion smile (blue) for strikes in [0.02, 0.06]. Hagan parameters as in figure 1

8 Variation with γ

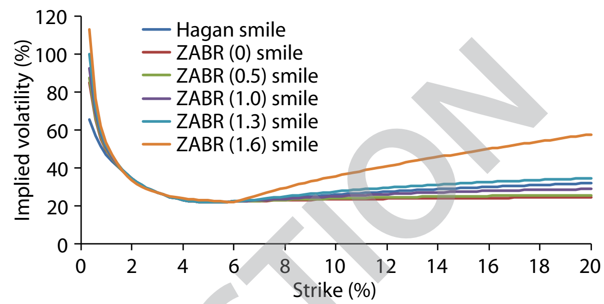

Note: calibration result for different choices of γ. All models calibrated to Hagan expansion prices for strikes in [0.02, 0.06]. Parameters as in figure 1

7 Positive probability of negative rates

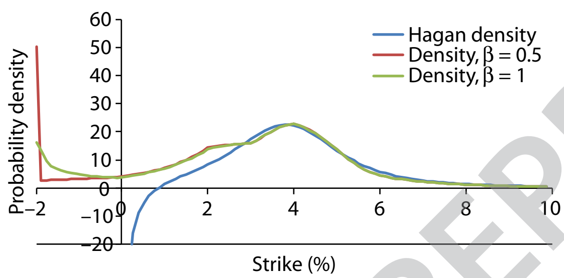

Note: densities for Hagan (blue) and two ZABR models with $\beta = 0.5$ (red) and $\beta = 1$ (green). The ZABR models are calibrated to the Hagan expansion prices for strikes in [0.02, 0.06]. Parameters as in figure 1

9 Convexity adjustment dependence on γ

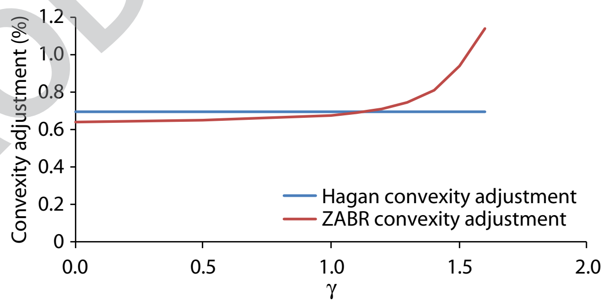

Note: convexity adjustment for a CMS forward 10×10 for different choices of γ in ZABR model (red) compared with Hagan (blue). ZABR models calibrated to Hagan expansion option prices for strikes in [0.02, 0.06]

However, typically, we prefer to calibrate directly to the observed discrete quotes. This is done by solving the ODEs in (34) and (35) and including the one-step finite difference adjustment (41):

$$\begin{aligned}\frac{\partial y}{\partial k} &= -\frac{z^{\gamma-2}}{\sigma(k)}\\ \frac{\partial x}{\partial k} &= -\frac{F\left(y, z^{\gamma-1}x\right)}{z\sigma(k)}\\ \theta(k) &= \frac{P(x)z\sigma(k)}{F\left(y, z^{\gamma-1}x\right)}\\ x(k=s) &= y(k=s) = 0\end{aligned} \qquad (47)$$

After numerical solution of (47), we find the option prices using the one-step finite difference algorithm in (43). On top of this, we use a non-linear solver to calibrate the volatility function $\sigma(k)$ to observed discrete option quotes. As we get all option prices in one sweep, we can include CMS forwards and option quotes in the calibration without additional computational costs.

Even though non-linear iteration is involved, this procedure is very fast. Typically, we can calibrate a non-parametric volatility function with 10 knot points to a given smile in roughly 50 iterations, which takes approximately one millisecond of CPU time.

When it comes to outright pricing speed, the ZABR model is capable of generating 100,000 smiles each consisting of 256 strikes in approximately seven seconds. It should be stressed that this includes both numerical ODE and finite difference solutions. This is actually faster than direct use of Hagan's SABR expansion, which takes 10 seconds for the same task. The reason for this difference is mainly that one time-step finite difference is faster at producing prices than the Black formula. An alternative to the ZABR model for producing arbitrage-free option prices is the Fourier-based models found in, for example, Lipton (2002). For a displaced Heston (1993) model, numerical solution for 100,000 smiles consisting of 256 strikes via the fast Fourier transform with the Black-Scholes formula used as a control variate takes around 18 seconds (see Andreasen & Andersen, 2002). It should be noted that this type of model is considerably less flexible with respect to fitting discrete quotes and more difficult to implement.

Though we generally use (47) in conjunction with a non-linear solver for the calibration, the direct calibration methodology (46) is relevant as it admits direct calibration of one ZABR model to another.

The stochastic process x has unit diffusion and thus, in the sense of the short-maturity limit, is normally distributed. So it is natural to use a uniform spacing in x and a non-uniform spacing in k. For this, the ODE system (47) can conveniently be transformed to:

$$\frac{\partial y}{\partial x}=\frac{z^{\gamma-1}}{F\left(y,z^{\gamma-1}x\right)}$$
$$\frac{\partial k}{\partial x}=-\frac{z\sigma(k)}{F\left(y,z^{\gamma-1}x\right)} \tag{48}$$
$$\theta(k)=\frac{P(x)z\sigma(k)}{F\left(y,z^{\gamma-1}x\right)}$$
$$y(x=0)=0,\quad k(x=0)=s$$

In our implementation, we solve (48) on a uniform x grid to generate and fix a non-uniform strike grid $\{k_0, k_1, \ldots, k_n\}$ that is used in the numerical solution of (47) during calibration and pricing. As a final remark, we note that ODEs in this section typically will be solved at $z = z(0) = 1$.

Controlling the wings

Here, we will give a few examples to illustrate how we are able to control the behaviour of the wings of the smile. Consider a model with:

$$\sigma(s)=\omega(s)\left(s-\underline{s}\right)^{\beta} \tag{49}$$

where ω is a non-parametric curve and $\underline{s}$ is the lower bound of the spot process. For $\beta < 1$, we have absorption at $\underline{s}$ and for $\beta \geq 1$ the barrier is unattainable. Our finite difference solution imposes absorption for the cases where the barrier is attainable. In figure 6, we fit this model to Hagan prices for $\beta = 0.5$ and $\beta = 1$ and $\underline{s} = -0.02$ and we see that the fit is good for positive strikes.

In figure 7, we have plotted the resulting densities. As before, the Hagan expansion produces negative densities for low positive strikes. For $\beta = 0.5$, we have absorption at the barrier and for $\beta = 1$ we see that the density below zero is spread out.

We now use the model with $\beta = 1$ to illustrate the effect of γ. For different levels of γ, we have calibrated the model to the Hagan expansion prices for strikes between 0.02 and 0.06. In figure 8, we see that all models are well calibrated in the sense that the models all produce the same smiles for strikes between 2% and 6%. We can also see the biggest impact of varying γ is for high strikes.

One way of fixing γ is to choose it to match CMS forwards or option quotes. In figure 9, we have shown the impact on a CMS convexity adjustment.

Conclusion

We have used a simple method to derive short-maturity expansions for forward volatilities from stochastic volatility models. The solution is an ODE that can be solved numerically for all strikes in one sweep including adjustment of the forward volatility function to compensate for the one-step finite difference option pricing. Finally, we use the one-step finite difference scheme to generate option prices. The approach is very fast and it generates arbitrage-free option prices. We have added flexibility to the original SABR model to get an exact fit of all quoted option prices and better control of the wings of the smile for improved CMS pricing. Also we can add CMS prices to the calibration without additional computational costs. ■

Jesper Andreasen is the head of and Brian Huge is chief analyst in the quantitative research department at Danske Markets in Copenhagen. They would like to thank colleagues Morten Karlsmark and Jesper Ferkinghoff-Borg for assistance with Taylor expansions. Email: kwant.daddy@danskebank.com, brian.huge@danskebank.com

References

Abramowitz M and I Stegun, 1972
Handbook of mathematical functions
Dover Publications, New York

Andersen L and R Ratcliffe, 2002
Extended Libor market models with stochastic volatility
Working paper, General Re Financial Products

Andreasen J and L Andersen, 2002
Volatile volatilities
Risk December, pages 163–168, available at *www.risk.net/1530294*

Andreasen J and B Huge, 2011
Volatility interpolation
Risk March, pages 86–89, available at *www.risk.net/2029822*

Balland P, 2006
Forward smile
Presentation, ICBI Global Derivatives

Doust P, 2010
No arbitrage SABR
Working paper, Royal Bank of Scotland

Dupire B, 1994
Pricing with a smile
Risk January, pages 18–20, available at *www.risk.net/1500211*

Dupire B, 1996
A unified theory of volatility
Working paper, Banque Paribas

Fiedler M, 1986
Special matrices and their application in numerical mathematics
Martinus Nijhoff, Dordrecht

Gyongy I, 1986
Mimicking the one-dimensional marginal distributions of processes having an Ito differential
Probability Theory and Related Fields 71, pages 501–516

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

Henry-Labordère P, 2008
Analysis, geometry, and modeling in finance
Chapman & Hall

Heston S, 1993
A closed-form solution for options with stochastic volatility with applications to bond and currency options
Review of Financial Studies 6, pages 327–343

Jex M, R Henderson and D Wang, 1999
Pricing exotics under the smile
Working paper, JP Morgan

Lewis A, 2000
Option valuation under stochastic volatility
Finance Press

Lewis A, 2007
Geometries and smile asymptotics for a class of stochastic volatility models
Presentation, UC Santa Barbara

Lipton A, 2002
The vol smile problem
Risk February, pages 61–65, available at *www.risk.net/1530435*

Lipton A and A Sepp, 2011
Filling the gaps
Risk September, pages 78–83, available at *www.risk.net/2111211*

Paulot L, 2009
Asymptotic implied volatility at the second order with application to the SABR model
Working paper, Sophis Technology

Press W, B Flannery, S Teukolsky and W Vetterling, 1992
Numerical recipes in C
Cambridge University Press