Calibrating Heston for credit risk

Marco de Innocentis and Sergei Levendorskiĭ describe a new method for market-implied calibration of the Heston model for equity, based on an improved version of the parabolic pricing algorithm. This pricing method, when used in the calibration, is much faster and more accurate, and better reproduces the implied volatilities, than popular alternatives. As such, it is suitable for use in a typical internal model method counterparty risk engine

One of the main reasons for the popularity of the Heston model (Heston 1993) is the availability of several algorithms for the fast pricing of European options (Carr & Madan 1999; Kahl & Jäckel 2005; Levendorskiĭ 2012). This allows calibration to be performed relatively quickly, and the model can then be used, for example, to price exotic contracts. In addition to its use in pricing, the Heston model has become a popular choice in counterparty credit risk (CCR) modelling.

Under the Basel III internal model method (IMM) approach, banks can compute their exposure at default (EAD), one of the key ingredients in the calculation of economic capital, using the output of the bank's internal Monte Carlo engine for calculating future exposures, subject to regulatory approval. The Heston model is particularly attractive in CCR modelling as it captures some distinctive features of the historical distribution of returns and performs well in historical backtesting. Since a bank may need to simulate portfolios with hundreds of distinct underlying assets over tens of thousands of scenarios for maturities of up to several years, a fast pricing algorithm is essential. The pricing method must also be sufficiently accurate to cope with long and very short maturities as well as options that are very far in- and out-of-the-money. Any situations with which the method struggles may well occur when hundreds or thousands of assets are simulated across tens of thousands of Monte Carlo paths. Should the method become very slow in such cases, then the simulation may not be completed in a reasonable time. If accuracy is sacrificed for speed, then incorrect prices can manifest themselves as spikes in exposure profiles, which in turn can result in increased risk charges.

Most banks use market-implied (risk neutral) calibration for the calculation of credit value adjustment (CVA), ie, the adjustment to the fair price of over-the-counter derivatives instruments to account for CCR. Risk-neutral calibration is also becoming increasingly popular for the calculation of exposure, especially since the publication of the Basel Committee's review of the credit valuation adjustment risk framework (Basel Committee on Banking Supervision 2015). This proposes banks use market-implied models for the calculation of the Basel III CVA capital charge, which aims to capitalise the risk of future changes in CVA.[1]

We develop a new market-implied calibration methodology based on an improved version of the (fractional) parabolic deformation method; this was introduced by Boyarchenko & Levendorskiĭ (2014) for pricing European options in Lévy models and was adjusted to the Heston model, as well as more general affine models, by Levendorskiĭ (2012, 2016). We show this method is much faster and more accurate than alternative methodologies, and the calibrated model performs well under the kind of backtesting required by regulators. Our benchmarking tests are designed based on recommendations in the Federal Reserve's Supervisory Guidance Letter SR-11-7 (Board of Governors of the Federal Reserve System 2011).

[1] *This is in order to recognise the hedges banks use to target the exposure component of CVA variability as well as ensure consistency with accounting standards.*

Heston model

Model specification and review of pricing methods. In the Heston model (Heston 1993), the stock price follows the dynamics:

$$\mathrm{d}S_t = (r-\delta)S_t\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}\tilde{W}_t^{(1)} \tag{1}$$

$$\mathrm{d}v_t = \kappa(m - v_t)\,\mathrm{d}t + \sigma_0\sqrt{v_t}\,\mathrm{d}W_t^{(2)} \tag{2}$$

under the risk-neutral measure $\mathbb{Q}$ used for pricing, where $\tilde{W}_t^{(1)}$ and $W_t^{(2)}$ are Brownian motions with correlation ρ, ie, $[\tilde{W}^{(1)}, W^{(2)}]_t = \rho_t$, and κ (the reversion speed of variance), m (the long-term mean) and σ_0 (the volatility of the variance) are positive. We will also assume $\rho \in (-1, 1)$.[2] In the Heston model, the characteristic function of S_t given S_0 and v_0 can be calculated explicitly by solving the associated system of Riccati equations. The analyticity of the characteristic exponent was used in the context of exponential Lévy models by Boyarchenko & Levendorskiĭ (2014), who derived a general formula for the price of a European option as an integral over a line parallel to the real axis (see Boyarchenko & Levendorskiĭ (2014) and references therein). This feature is extremely useful, since the discretisation error of the infinite trapezoid rule decays as $\exp[-2\pi d/\zeta]$, where d is the half-width of the strip of analyticity of the integrand, and ζ is the mesh. This result was used by Feng & Linetsky (2008), who gave approximate error bounds and recommendations for the choice of mesh for several classes of problems arising in finance. More accurate bounds and prescriptions for optimal choices of both line of integration and mesh were derived in Boyarchenko & Levendorskiĭ (2014) and Levendorskiĭ (2012). They also introduced the (fractional) parabolic deformation of the line of integration, which allows the number of terms $N(\varepsilon)$ sufficient to satisfy an error tolerance ε to be much smaller.

Pricing methodology used in the paper. The pricing algorithm used in our calibration is an improved version of the method introduced in Levendorskiĭ (2012) (described in the appendix), which is aimed at reducing arbitrage. Especially for very short or long maturities, the approximate bounds in Levendorskiĭ (2012), on which the numerical prescriptions for the discretisation mesh ζ, offset ω and number of terms N rely (see the appendix), may fail to hold, which can occasionally result in incorrect prices. Hence, our version of the algorithm works as follows:

(1) We use the (fractional) parabolic method of order $\alpha = 1.7$, as described in Levendorskiĭ (2012). In particular, we use the explicit procedure to determine the maximal strip of analyticity $\mathrm{Im}\,\xi \in (\lambda_-(\tau), \lambda_+(\tau))$ of the characteristic function. The procedure for the choice of the deformed contour and the pricing formula (for call, put or covered call) is formulated in terms of substrips $\mathrm{Im}\,\xi \in (\lambda_-(\tau), -1)$, $\mathrm{Im}\,\xi \in (-1, 0)$ and

[2] *In the equity market, usually* $\rho < 0$.

$\mathrm{Im}\,\xi \in (0, \lambda_+(\tau))$ of the analyticity of the integrand; the widest substrip is recommended in Levendorskiĭ (2012), and explicit recommendations for the choice of the parameters of the deformation and simplified trapezoid rule are given (see the appendix).

■ (2) If the vanilla price breaches the no-arbitrage bound,[3] then we use the parabolic method again, with the following modifications.

■ (2a) Instead of using the largest strip for $\mathrm{Im}\,\xi$ in the pricing formula, between $(\lambda_-(\tau), -1)$, $(-1, 0)$ and $(0, \lambda_+(\tau))$, we use the second largest.

■ (2b) If the price still breaches the no-arbitrage bound, then we use the third strip.

Henceforth, we will refer to this as the 'multi-strip' (MS) version of the parabolic method (ParMS).

■ **Alternative pricing methods.** The following alternative methodologies were used in order to benchmark the algorithm described above.

■ **Carr-Madan (CM) method.** This is based on the algorithm employed in Carr & Madan (1999), where it was used for a specific example rather than presented as a universal prescription. We use it in the form described in Albrecher *et al* (2007), under which it has become popular with practitioners. This is equivalent to the choice of $\omega = -1.75$, $\zeta = 0.25$ and $N = 4096$ points in the notation of the appendix. The price is calculated using the fast Fourier transform and linearly interpolating over the results in log-strike space.

■ **Adaptive quadrature.** This is based on applying a high-precision adaptive quadrature scheme to the 'semi-analytical' pricing formula in Heston (1993), together with the change of variables in Kahl & Jäckel (2005), in order to reduce the integral from $(0, \infty)$ to $(0, 1)$. We use the NAG function `d01sjc`, which implements an algorithm based on Gauss 10-point and Kronrod 21-point rules.

■ **Lewis-Lipton formula.** This approach is equivalent to the reduction to a covered call (ie, the strip $\mathrm{Im}\,\xi \in (0, 1)$) and using $\omega = -1/2$ in the pricing formula. The strip $\mathrm{Im}\,\xi \in (-1, 0)$ is convenient for very large maturities, when typically $\lambda_+(\tau)$ and $-1 - \lambda_-(\tau)$, which determine the width of the strips corresponding to the put and call option cases (see Levendorskiĭ 2012), are both very small.

■ **COS method.** This method, introduced in Fang & Oosterlee (2008) for European vanillas and later extended to American and discrete barrier options, is based on a cosine expansion of the pricing kernel. As shown in Boyarchenko & Levendorskiĭ (2014), COS is equivalent to the standard inverse Fourier transform algorithm applied to the option with a modified payoff. Although increasingly popular with practitioners, as shown in Boyarchenko & Levendorskiĭ (2014) and de Innocentis & Levendorskiĭ (2014), it can produce large errors that are difficult to estimate. Fang & Oosterlee (2008) only provide an *ad hoc* prescription for the choice of numerical parameters in their scheme, rather than recommendations based on error estimates. In our numerical examples, we use the prescription given in Fang & Oosterlee (2008, section 5.3) for the choice of the truncation parameters, and $N = 160$ terms (the largest number used in the numerical examples in Fang & Oosterlee (2008)).

[3] *Assuming a constant interest rate $r > 0$, this takes the form:*

$$\max\{0, \mathrm{e}^{-r(T-t)}(F(t,T) - K)\} \leqslant V_{\mathrm{call}}(t, T; K) \leqslant \mathrm{e}^{-r(T-t)} F(t, T)$$

where $V_{\mathrm{call}}(t, T; K)$ denotes the time t price of a call option with strike K and maturity T, and $F(t, T)$ is the corresponding forward price of the underlying.

Calibration methodology

The objective of the calibration is to determine the set of Heston parameters $\Theta = \{\rho, \sigma_0, \kappa, m, v_0\}$ that, for a given stock or equity index on a given calibration date, best reproduces the corresponding implied volatility surface. Since the Heston model parameters are constant through time, an exact calibration to the surface cannot be achieved. Instead, an 'optimal' parameter set Θ is determined by minimising the fitting error using a least-squares optimisation procedure as follows.[4] We calibrate to a set of times to maturity $\tau_k, k = 1, 2 \ldots, M$. For each maturity τ_k, we select a grid of strikes $K_n^-(\tau_k), K_{n-1}^-(\tau_k), \ldots, K_1^-(\tau_k), K_0(\tau_k), K_1^+(\tau_k), \ldots, K_n^+(\tau_k)$, where $K_0(T) = F(T)$, $T = \tau_1, \tau_2, \ldots, \tau_M$ and:

$$K_j^{\pm}(T) = F(T) + j\frac{K^{\pm}(T) - F(T)}{n}, \quad j = 1, 2, \ldots, n \qquad (3)$$

where in turn $F(T)$ denotes the underlying forward price for maturity T; $K^{\pm}(T) = S\exp(\pm N\sigma(T)\sqrt{T} - 0.5\sigma^2(T)T)$; S is the underlying spot price;[5] $\sigma(T)$ denotes the at-the-money-forward (ATMF) implied volatility at maturity T; and N is a positive integer, ie, for each maturity, we try to match the market volatilities lying approximately within N standard deviations of the ATMF point. The calibration minimises the sum of squares of differences between model and market-implied volatilities of ATMF or out-of-the-money-forward options with the given strikes and maturities. An advantage of fitting to implied volatilities rather than prices is the former tend to be of a similar order of magnitude across different strikes and maturities, resulting in a more global fit. This is especially useful when pricing exotics, which can be replicated as portfolios of vanillas. To reflect the maturities of OTC equity options traded in the market, we use times to maturity corresponding to 30 days, 90 days, 180 days, one year and three years, along with $n = 4$ (that is, nine strikes per maturity) and $N = 3$ standard deviations. For the minimisation, we use the Levenberg-Marquardt method for constrained optimisation, with the following admissible region for the model parameters in (1) and (2):[6]

$$\rho \in (-0.99, -0.1), \quad \sigma_0 \in (0.001, 4), \quad \kappa \in (0.001, 4)$$
$$m \in (0.005, 4) \quad \text{and} \quad v_0 \in (0.005, 4)$$

The latter was determined from extensive empirical analysis after fitting to options on several liquid underlyings between January 2008 and December 2016.

Numerical results

■ **Pricer benchmarking.** We use a set of 10,000 five-dimensional Halton points to cover the admissible region used for the calibration of the Heston parameters. Assuming zero interest and dividend rates, at each point we calculate the price of an in-the-money option with moneyness

[4] *The procedure described in this section was developed in collaboration with Bechir Saidi.*

[5] *For ease of notation, we suppress the dependency on the spot date t.*

[6] *A single starting point in the centre of the region is used. The authors experimented extensively using multiple starting points determined by a (quasi) random sequence; however, with the Levenberg-Marquardt algorithm, a single starting point usually suffices to find a near-global minimum. The authors would like to thank the anonymous referee who recommended further investigation of the calibration with multiple starting points.*

A. Average times and percentage of no-arbitrage breaches in the prices of vanillas, butterfly spreads and calendar spreads observed with each pricing method with respect to the benchmark

	ParMS	Par	KJ	CM	COS
Average pricing time (ms)	0.13	0.11	4.51	14.63	3.31
% vanilla breaches	0	0.33	0.93	20.26	31.21
% BS breaches	1.30	1.33	1.75	12.43	9.21
% CS breaches	0.04	0.08	0.45	11.01	8.74
% average calibration time (s)	1.25	N/A	13.04	20.94	7.12

The calibration times refer to the procedure described in the numerical results section

1 Error distribution relative to the benchmark for MS, Par, KJ, CM and COS methods, for strike/spot between 1/5 and 5, and maturity between 1W and 5Y

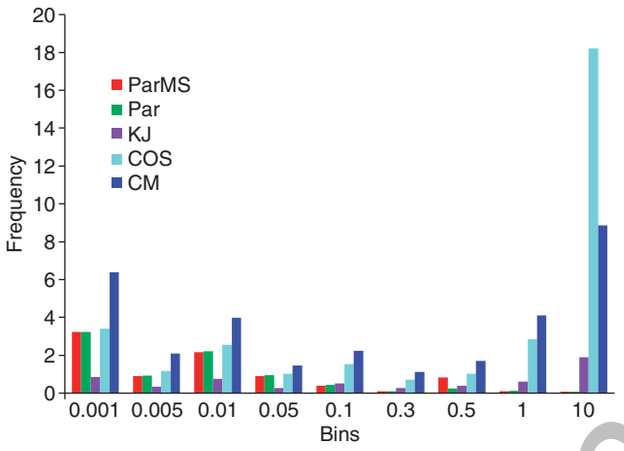

A value of, for example, 3% above the bin labelled 0.001 indicates that 3% of the relative errors fall between 0.001 and the boundary of the next bin (0.005). The first bin for relative errors between 0 and 0.001 was removed due to its large values. The frequencies for ParMS and Par were about 90% in this first bin, while KJ, COS and CM had frequencies just below 70%

K/S_t between 1/5 and 5, and maturities between 1W and 7Y, using our version of the parabolic method (ParMS), the original parabolic method as described in Levendorskiĭ (2012) (Par), the Kahl-Jäckel adaptive quadrature method (KJ) and the CM and COS methods.[7] For this benchmarking exercise only, we apply CM separately at each strike, ie, without interpolation. Our benchmark is calculated using the Lewis-Lipton formula and an iterative procedure for ζ, $\Lambda = N\zeta$: starting with initial values $\zeta_0 = 0.02$ and $\Lambda_0 = 500$, we double Λ in a loop until the results of two successive evaluations have a relative difference under 10^{-3}. Then, using this Λ, we decrease the mesh ζ by a factor of 1.5 each time, until the results of two successive evaluations have a relative difference under 10^{-3}.

For each method, we look at

- the percentage of call option prices V_t that breach the no-arbitrage pricing bound within an absolute tolerance of 0.0001, eg, for a call option on a non-dividend-paying asset S_t:

$$\max\{0, S_t - Ke^{-r(T-t)}\} \leq V_{\text{call}}(t, T; K) \leq S_t$$

where T is the maturity of the option;

- the percentage of prices resulting in negative prices for butterfly spreads;
- the percentage of prices resulting in negative prices for calendar spreads; and
- the total calculation times.

Table A shows the percentage of no-arbitrage condition breaches and the average calculation time per option (in milliseconds). The parabolic method is by far the fastest, with an average time of 0.13 ms per option for the MS version vs 4.51 ms for KJ. ParMS has by far the lowest percentage of no-arbitrage breaches, followed closely by the original parabolic method. For the KJ method, the percentages of breaches are around 1%; however, the calculation time is much slower than ParMS at 4.51 ms, which is likely to be too slow for use in a counterparty risk engine. For both COS and CM, the breach percentages are much higher (over 20% for vanilla prices calculated with CM), and the calculation times are comparable with that of KJ. Figure 1 shows the distribution of errors relative to the benchmark for the different pricing methods. From this comparison, it can be seen that in more than 18% of cases, COS results in relative errors of 1,000% and above. This is due to the fact that with COS the errors for maturities larger than 3Y can be extremely high (a similar effect for barrier options was documented in de Innocentis & Levendorskiĭ (2014)). Errors in the same range also occur for CM and KJ in about 8% and 2% of cases, respectively. As discussed in the introduction, such large mispricings can have serious repercussions in economic capital calculations.

Calibration. To examine the robustness of the calibration, we take the parameters calibrated to the Standard & Poor's 500 (S&P 500) volatility surface for May 30, 2008 and compare

- the market-implied volatilities;[8]
- the implied volatilities calculated using the ParMS with parameters calibrated using the ParMS method;
- the implied volatilities calculated using the ParMS method with parameters calibrated using CM; and
- the implied volatilities calculated using CM with parameters calibrated using CM.

Figure 2 shows the volatility smiles corresponding to maturities of one and nine months. The following conclusions can be drawn:

- The results shown demonstrate the ParMS calibration produces a good fit to the market smile, even for maturities and moneyness levels not included in the calibration, ie, good out-of-sample performance (note the very close match between the red and blue lines in figure 2(b) for the 9M maturity, which was not included in the calibration).
- The smiles calibrated using CM show much larger deviations from the market smile.
- The CM algorithm does not produce reliable results for short maturities, even with the parameter set found with its own calibration (cf, the differences between the green and dashed black curves). Note such large errors for ITM options approaching expiry can lead to miscalculation of exposure in a counterparty risk engine.
- CM is also slower than ParMS, taking almost 21 seconds to calibrate, on average, versus ParMS's 1.25 seconds for the 2010–16 period.[9]

[7] *Specifically, we use 1W, 1M, 3M, 6M, 9M, 1Y, 2Y, 3.5Y and 5Y, strikes* $20, 30, \ldots, 500$ *and spot price* $S_t = 100$.

[8] *The data for this exercise was provided by a Tier 1 bank in London.*

[9] *This is on a PC with Intel(R) Xeon(R) CPU, 2.27 GHz (2 processors), with 16 cores and 16 GB RAM and 64-bit, running Windows 7 Enterprise, using Visual Studio 2010.*

2 S&P 500: May 30, 2008

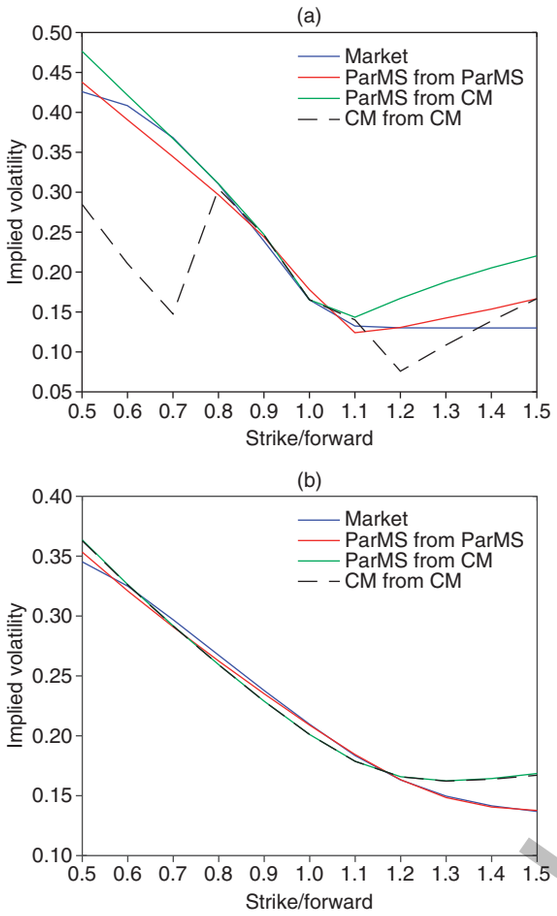

Volatility smile comparison at (a) 1M and (b) 9M maturities: market smile (blue line), smile reconstructed from ParMS calibration using ParMS pricer (red line), smile from CM calibration using ParMS pricer (green line) and smile from CM calibration using CM pricer (dashed black line)

The results for COS (not shown) are qualitatively similar to those for CM. COS also struggles to reproduce the market smile accurately, albeit to a smaller extent than CM. This is due to the fact that COS tends to produce large errors for times to maturity larger than 3Y.

Backtesting

Backtesting attempts to assess the predictive power of a risk-factor evolution (RFE) model by comparing its past predictions with the distribution of realised outcomes. For a set of historical dates, the RFE model is calibrated to the implied volatility surface as of that day. For each calibration date – in our case, each month end date t between January 29, 2010 and December 30, 2016 – we calculate the cumulative distribution function (CDF) of the corresponding realised value of the stock price at time $t + \Delta$, where the horizon Δ varies between 1W and 1Y, using the Heston parameters calibrated at t, with zero drift. The assumption of zero drift is based both on the difficulty of estimating the future drift and on grounds of conservatism. Since equity prices are positively correlated with the state of the financial system, possible underestimations of the exposures in scenarios of bull markets are less of a concern, especially for regulators, than underestimating exposures under conditions of market distress. Typically, a mixture of exception counting (EC) and distributional tests is used (see, for example, Anfuso *et al* 2014). For each test, we calculate the p-value at a given time horizon and assign a red, amber or green colour to it according to whether it lies in $[0, 0.01)$, $[0.01, 0.05)$ or $[0.05, 1]$, respectively.[10] For left and right EC at a given time horizon Δ, the calculation of the p-value works as follows. First, one counts the number of exceptions, ie, the number E of CDFs of realised values at horizon Δ that lies in the corresponding tail. The p-value is then given by:

$$p = 1 - \sum_{k=0}^{E-1} \binom{N}{k} \beta^k (1-\beta)^{N-k}$$

where β is the percentile level of the test (eg, $\beta = 0.01$) and $N \geqslant 1$ is the total number of observations (eg, for monthly calibrations over seven years, $N = 7 \times 12 = 84$). If no exceptions are observed, then we set $p = 1$.

Figure 3 shows the results obtained with the S&P 500 as underlying,[11] for horizons ranging between one week and one year, left and right EC with levels $\beta = 0.01, 0.05, 0.1, 0.15, 0.20, 0.25, 0.30$ (corresponding to left tail exceptions, EXCL), and levels $\beta = 0.70, 0.75, 0.80, 0.85, 0.90, 0.95, 0.99$ (corresponding to right tail exceptions, EXCR), as well as the Anderson-Darling (AD) and Cramér-von Mises (CvM) distributional tests. These tests are passed in almost all cases the authors have come across. Figure 3(a) shows the results obtained with ParMS calibration, while figure 3(b) shows the corresponding results for CM calibration. Note the CM calibration produces two red and amber results for EC with a 1W horizon, which do not occur with ParMS.[12] This is likely a consequence of the fact that CM struggles for low maturities, as noted above (cf, the blue and dashed black lines in figure 2(a)).

A reason for the good backtesting performance of the Heston model is its ability to capture the main features of the historical distribution of returns. Figure 4 shows the Heston probability density function (PDF) calibrated to fit the implied volatility surface of the S&P 500 index as of January 30, 2015, compared with that estimated from the monthly returns of the index between January 31, 2008 and January 30, 2015. The fat left tail and the thin right tail are why the Heston model is historically realistic and successful in risk-factor backtesting.

Conclusions

The parabolic method, even without the 'multi-strip' feature, is both faster and more accurate than popular alternatives, to such an extent that none of the other methods used in our comparisons are suitable for a typical IMM counterparty risk engine, either due to insufficient accuracy or speed. As discussed in the introduction to this article, a fast algorithm is essential when simulating portfolios with hundreds or thousands of assets across

[10] *For a procedure to deal with overlapping backtesting horizons, see Anfuso et al (2014).*

[11] *Similar results (all or mostly green) were also obtained for several liquid stocks and indexes.*

[12] *Although the results look similar for this underlying in terms of colours, there are large differences in some of the p-values. The p-values were omitted at the request of the bank that provided our market data.*

3 Backtesting results for the S&P 500 index over the period 2010–16, using (a) ParaMS calibration and (b) CM calibration

(a)

Name	Quant.	1W	2W	1M	3M	6M	1Y
EXCL	0.01						
EXCR	0.99						
EXCL	0.05						
EXCR	0.95						
EXCL	0.10						
EXCR	0.90						
EXCL	0.15						
EXCR	0.85						
EXCL	0.20						
EXCR	0.80						
EXCL	0.25						
EXCR	0.75						
EXCL	0.30						
EXCR	0.70						
CvM							
AD							

(b)

Name	Quant.	1W	2W	1M	3M	6M	1Y
EXCL	0.01						
EXCR	0.99						
EXCL	0.05						
EXCR	0.95						
EXCL	0.10						
EXCR	0.90						
EXCL	0.15						
EXCR	0.85						
EXCL	0.20						
EXCR	0.80						
EXCL	0.25						
EXCR	0.75						
EXCL	0.30						
EXCR	0.70						
CvM							
AD							

'Quant.' denotes quantile. Note that CM results in some spurious exceptions for short horizons (cf, figure 2(a))

4 Comparison of the empirical PDF of historical S&P 500 monthly returns in January 2008–January 2015 and the Heston PDF calibrated as of January 30, 2015 to European options

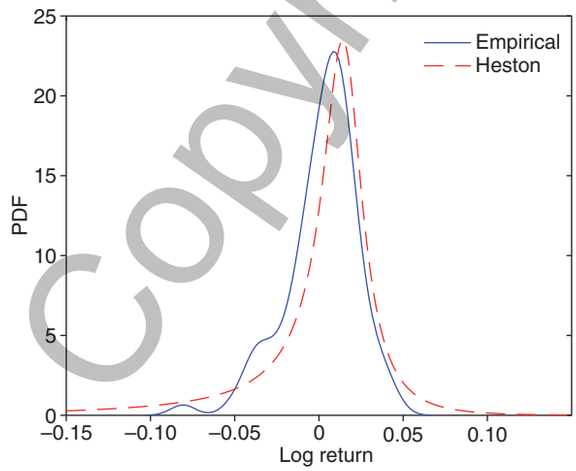

tens of thousands of scenarios. In light of this, the advantage of the multi-strip version of the method becomes apparent. Although, as shown in table A, the original method gives prices outside the bounds only about 0.33% of the time, even a single exploding price can cause significant difficulties in a counterparty risk engine. With the multi-strip version, this problem is never seen to occur.

We have also found the Heston model, with market-implied calibration to equity and equity index options, backtests well. Figure 3 shows both EC and distributional tests pass at several horizons. Moreover, the EC tests were carried out for several left and right quantiles, which is itself akin to a distributional test. This result is especially surprising since a market-implied, rather than historical, calibration was used.

Appendix: parabolic pricing method

For the application of the (fractional) parabolic deformations, it is convenient to use the following change of variables. The underlying stock is represented as $S_t = \exp[X_t + by_t]$, where:

$$\mathrm{d}X_t = (\mu_0 + \mu_1 y_t)\,\mathrm{d}t + \sqrt{y_t}\,\mathrm{d}W_t^{(1)} \tag{4}$$

$$\mathrm{d}y_t = \kappa(\theta - y_t)\,\mathrm{d}t + \sigma\sqrt{y_t}\,\mathrm{d}W_t^{(2)} \tag{5}$$

Here, $W_t^{(1)}$ and $W_t^{(2)}$ are independent Brownian motions, and:

$$y_t = (1-\rho^2)v_t, \qquad X_t = \ln S_t - \frac{\rho}{\sigma_0}v_t$$

as well as:

$$b = \frac{\rho}{\sigma_0(1-\rho^2)}, \qquad \theta = m(1-\rho^2), \qquad \sigma = \sigma_0(1-\rho^2)^{1/2}$$

and the parameters μ_0 and μ_1 are given by:

$$\mu_0 = r - \delta - \frac{\rho\kappa\theta}{\sigma(1-\rho^2)^{1/2}} = r - \delta - b\kappa\theta = r - \delta - \kappa m\rho/\sigma_0 \tag{6}$$

$$\mu_1 = \frac{\rho\kappa}{\sigma(1-\rho^2)^{1/2}} - \frac{1}{2(1-\rho^2)} = b\kappa - \frac{1}{2(1-\rho^2)} = \frac{\rho\kappa/\sigma_0 - 1/2}{1-\rho^2} \tag{7}$$

Consider a European call or put option on the stock S_t with strike K and maturity date $T > 0$, and denote its terminal payoff by $G(\ln S_t)$. Let $V(t,x,y)$ be the option price at $t \leqslant T$ under $\mathbb{Q}$, such that $X_t = x$ and $v_t = y$. The option price is given by:

$$V(t,x,y) = \mathbb{E}^{\mathbb{Q}}_{t,x,y}[\mathrm{e}^{-r\tau}G(X_T + by_T)]$$

We expand the payoff G into the Fourier integral:

$$G(\hat{x}) = -\frac{K}{2\pi}\int_{\mathrm{Im}\,\xi=\omega} \frac{\mathrm{e}^{\mathrm{i}(\hat{x}-\ln K)\xi}}{\xi(\xi+\mathrm{i})}\,\mathrm{d}\xi \tag{8}$$

where we need $\omega > 0$ for the put and $\omega < -1$ for the call payoff. If we denote time to maturity by $\tau = T - t$, it can be shown there exist $\lambda_-(\tau) < -1 < \lambda_+(\tau)$ such that, if $\omega \in (0, \lambda_+(\tau))$ for the put, or $\omega \in (\lambda_-(\tau), -1)$ for the call, then one can substitute (8) into the expression for $V(t,x,y)$ and change the order of taking expectation and integration using Fubini's theorem. Taking the symmetry of the integrand into account, the result is:

$$V(t,x,y) = -\frac{K\mathrm{e}^{-r\tau}}{\pi}\,\mathrm{Re}\int_{\mathrm{i}\omega}^{\mathrm{i}\omega+\infty} \frac{\mathrm{e}^{\mathrm{i}x'\xi + yB(\tau,\xi) + C_1(\tau,\xi)}}{\xi(\xi+\mathrm{i})}\,\mathrm{d}\xi \tag{9}$$

where $x' = x - \ln K + \mu_0 \tau$ and:

$$B(\tau,\xi) = \sigma^{-2}(\kappa - R(\xi))\frac{1 - D_1(\xi)\mathrm{e}^{-\tau R(\xi)}}{1 - D(\xi)\mathrm{e}^{-\tau R(\xi)}} \tag{10}$$

$$C(\tau,\xi) = -r\tau + \mathrm{i}\mu_0\tau\xi + C_1(\tau,\xi) \tag{11}$$

$$C_1(\tau,\xi) = \frac{\kappa\theta}{\sigma^2}\left((\kappa - R(\xi))\tau - 2\ln\frac{1 - D(\xi)\mathrm{e}^{-\tau R(\xi)}}{1 - D(\xi)}\right) \tag{12}$$

and:

$$D(\xi) = \frac{\mathrm{i}b\xi - \beta_-(\xi)}{\mathrm{i}b\xi - \beta_+(\xi)} = \frac{\mathrm{i}b\sigma^2\xi - \kappa + R(\xi)}{\mathrm{i}b\sigma^2\xi - \kappa - R(\xi)} \tag{13}$$

$$D_1(\xi) = \frac{\beta_+(\xi)}{\beta_-(\xi)}D(\xi) = \frac{\kappa + R(\xi)}{\kappa - R(\xi)}D(\xi) \tag{14}$$

where in turn:

$$R(\xi) = \sqrt{\kappa^2 + \sigma^2(-2\mathrm{i}\mu_1\xi + \xi^2)}$$

The integral on the right-hand side of (9) is discretised using the simplified trapezoid rule. The universal recommendations for the choice of N and ζ make adaptive schemes unnecessary.

The pricing algorithm is based on the following ingredients:

- a procedure for the numerical calculation of $\lambda_-(\tau)$ and $\lambda_+(\tau)$ (see Levendorskiĭ 2012);
- a deformation of the contour of integration in (9) using a conformal map, such that along the deformed contour the integrand decays much more rapidly; and
- prescriptions for the choice of ω, ζ and Λ, and the other numerical parameters involved, according to the user's error tolerance (see Levendorskiĭ 2012).

The conformal maps $\chi^+_{\alpha;\omega;a}$ and $\chi^-_{\alpha;\omega;a}$ can be used for the contour deformation when $x' \geqslant 0$ and $x' \leqslant 0$, respectively. These are defined by:

$$\begin{aligned}\xi &= \chi^{\pm}_{\alpha;\omega;a}(\eta) = \mathrm{i}(\pm a + \omega) \mp \mathrm{i}a(1 \pm \mathrm{i}\eta)^{\alpha} \\ &:= \mathrm{i}(\pm a + \omega) \mp \mathrm{i}a\exp[\alpha\ln(1 \pm \mathrm{i}\eta)]\end{aligned} \tag{15}$$

Here, $\alpha \in [1,2)$ and $a > 0$. The deformation is justified by Cauchy's theorem. The formula for the new contour resembles that of a fractional parabola, hence the name of the method (parabolic iFT). If the 'plus' map is used, this gives the following approximation for the option price:

$$\begin{aligned}V(t,x,y) &= -\frac{K\mathrm{e}^{-r\tau}}{\pi}\zeta \\ &\times \mathrm{Re}\sum_{j=0}^{N-1}\frac{\mathrm{e}^{\mathrm{i}x'\chi^+_{\alpha;\omega;a}(\eta_j) + yB(\tau,\chi^+_{\alpha;\omega;a}(\eta_j)) + C_1(\tau,\chi^+_{\alpha;\omega;a}(\eta_j))}}{\chi^+_{\alpha;\omega;a}(\eta_j)(\chi^+_{\alpha;\omega;a}(\eta_j) + \mathrm{i})} \\ &\times \alpha a(1 + \mathrm{i}\eta_j)^{\alpha-1}(1 - \delta_{j0}/2)\end{aligned} \tag{16}$$

The expression for the option price if the map $\chi^-_{\alpha;\omega;a}$ is used to deform the integration contour is similar. ■

Marco de Innocentis is a senior quantitative analyst at Credit Suisse, London, and an honorary fellow in the Department of Mathematics at the University of Leicester. Sergei Levendorskiĭ is founder and partner at Calico Science Consulting in Austin, Texas. The views expressed herein are those of the authors only. No other representation should be attributed.
Marco de Innocentis would like to thank Bechir Saidi and Lorenzo Pitotti for useful discussions.
Email: marcdein@gmail.com (corresponding author), levendorskii@gmail.com.

REFERENCES

Albrecher H, P Mayer, W Schoutens and J Tistaert, 2007
The little Heston trap
Wilmott January, pages 83–92

Anfuso F, D Karyampas and A Nawroth, 2014
Credit exposure models backtesting under Basel III
Risk August, pages 82–87

Basel Committee on Banking Supervision, 2015
Review of the credit valuation adjustment risk framework
Consultative Document, July, Bank for International Settlements
Available at *www.bis.org/bcbs/publ/d325.htm*

Board of Governors of the Federal Reserve System, 2011
Guidance on model risk management
Supervisory Guidance Letter SR-11-7, April 4, Federal Reserve
Available at *www.federalreserve.gov/bankinforeg/srletters/sr1107.htm*

Boyarchenko S and S Levendorskiĭ, 2014
Efficient variations of Fourier transform in applications to option pricing
The Journal of Computational Finance 18(2), pages 57–90

Carr P and DB Madan, 1999
Option valuation using the fast Fourier transform
The Journal of Computational Finance 2(4), pages 61–73

de Innocentis M and SZ Levendorskiĭ, 2014
Pricing discrete barrier options and credit default swaps under Lévy processes
Quantitative Finance 14(8), pages 1337–1365

Fang F and CW Oosterlee, 2008
A novel pricing method for European options based on Fourier-cosine series expansions
SIAM Journal on Scientific Computing 31(2), pages 826–848

Feng L and V Linetsky, 2008
Pricing discretely monitored barrier options and defaultable bonds in Lévy process models: a fast Hilbert transform approach
Mathematical Finance 18(3), pages 337–384

Heston SL, 1993
A closed-form solution for options with stochastic volatility with applications to bond and currency options
Review of Financial Studies 6(2), pages 327–343

Kahl C and P Jäckel, 2005
Not so complex logarithms in the Heston model
Wilmott September, pages 94–103

Levendorskiĭ S, 2012
Efficient pricing and reliable calibration in the Heston model
International Journal of Theoretical and Applied Finance 15(7), pages 1–44

Levendorskiĭ S, 2016
Pitfalls of the Fourier transform method in affine models, and remedies
Applied Mathematical Finance 23(2), pages 81–134