Out of the past
Into the future

Mathematical issues within the Heston stochastic volatility model

Roger Lord
http://people.few.eur.nl/lord/

March 27th, 2006

Contents

- The Heston model
- The complex logarithm

  Joint work with Christian Kahl (University of Wuppertal, ABN·AMRO), forthcoming.

- Euler discretisations

  Joint work with Remmert Koekkoek (Robeco Alternative Investments) and Dick van Dijk (Erasmus University Rotterdam), "A comparison of biased simulation schemes for stochastic volatility models", available online.

The Heston model

In the Heston [1993] model the underlying asset and the variance evolve according to the following SDEs:

$$dS(t) = \mu S(t)dt + \sqrt{V(t)}S(t)dW_S(t)$$

$$dV(t) = -\kappa\left(V(t) - \theta\right)dt + \omega\sqrt{V(t)}dW_V(t)$$

where $dW_S(t) \cdot dW_V(t) = \rho\, dt$.

Mainly used in equity/FX, but the mean-reverting square root process is also often used as the stochastic volatility driver within interest rate models.

The Heston model (2)

The mean-reverting square root process used for the stochastic variance was first used in finance for the short interest rate by Cox, Ingersoll and Ross [1985] model. The corresponding parabolic equation was already studied by Feller [1951].

Important pro's of the Heston model are:

- For $V(0), \kappa, \theta \geq 0$, we have that $V(t) \geq 0$ a.s.;
- Semi-analytical pricing of European options;
- It fits the market prices quite reasonably;

The Heston model (3)

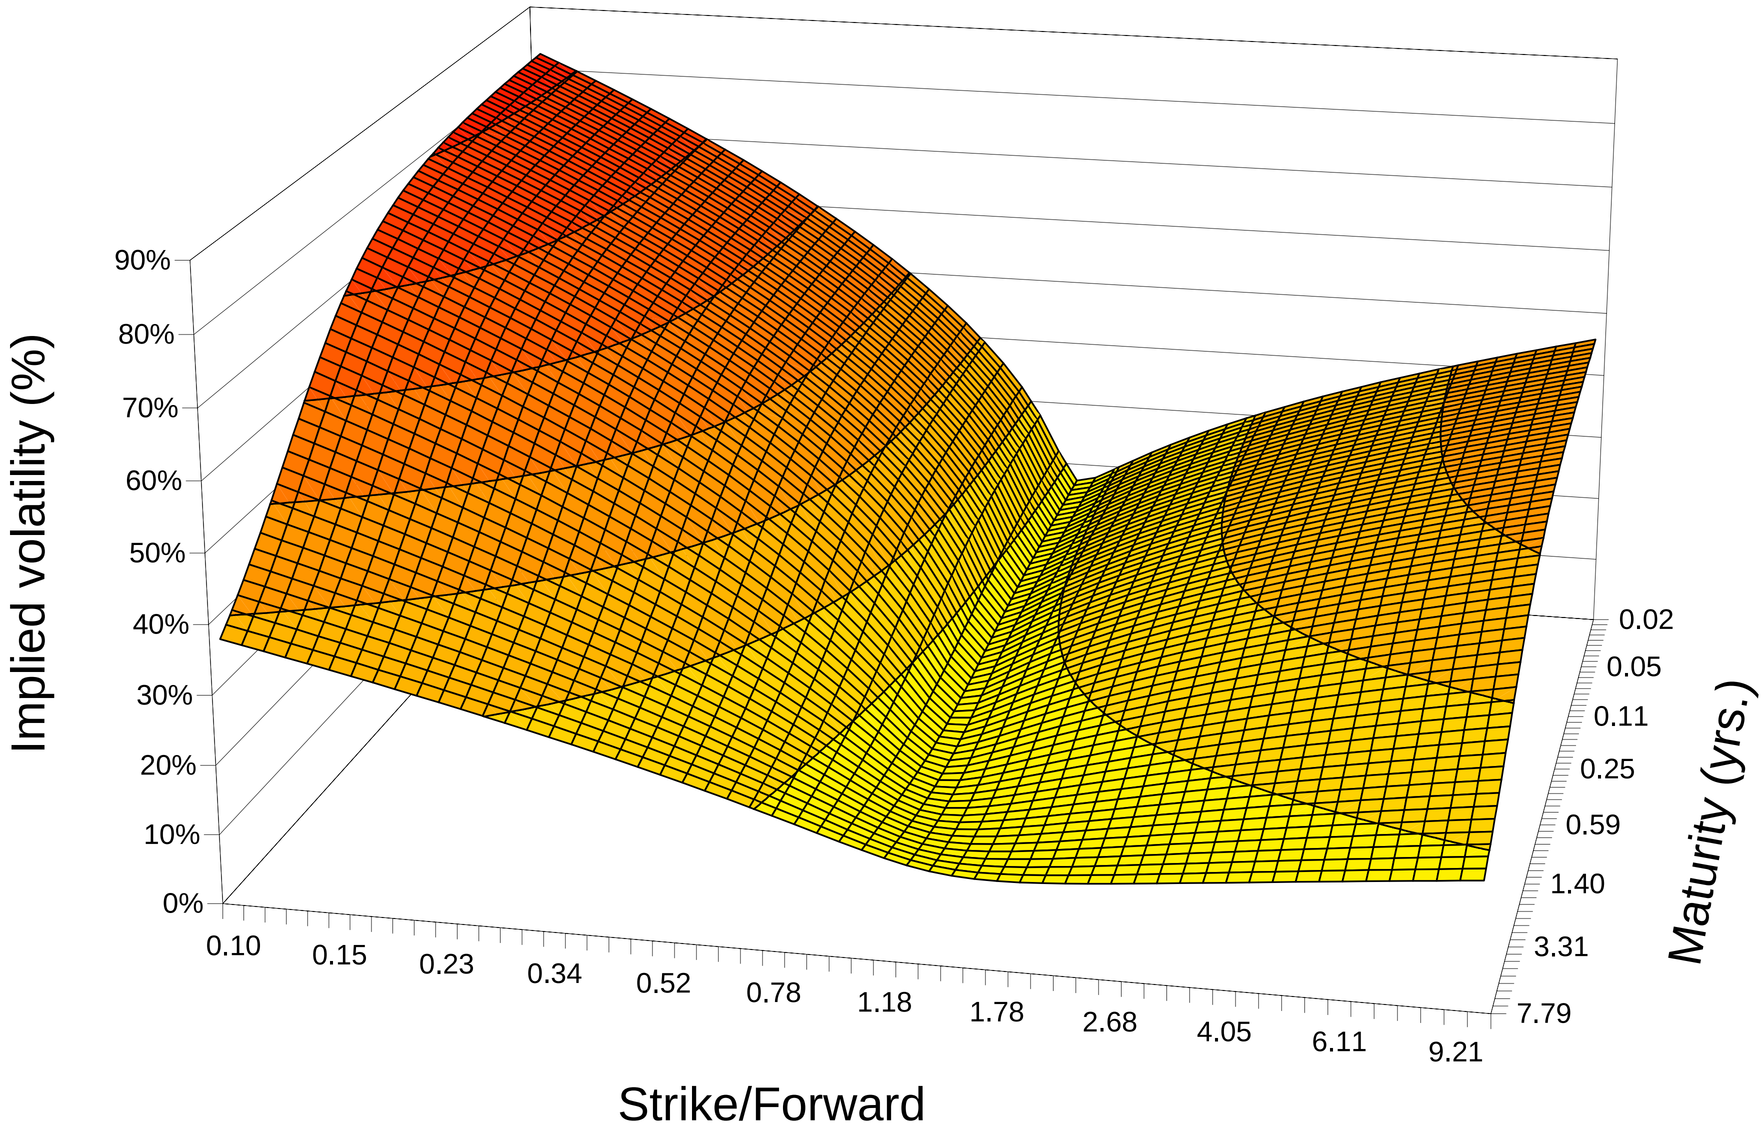

The Heston model (4)

If we have the characteristic function $\phi(u)$, $u \in \mathbb{C}$, which is equal to $\mathbb{E}_t\left[\exp(iu \ln S(T))\right]$, option prices are given by a formula similar in structure to the Black-Scholes [1973] formula:

$$\frac{C(t,T,S(t),K)}{P(t,T)} = \mathbb{E}_t\left[\left(S(T)-K\right)^+\right] = F(t,T)\cdot\Pi_1 - K\cdot\Pi_0$$

The probabilities Π_0 and Π_1 follow from the Gil-Pelaez inversion formula [1951]:

$$\Pi_j = \tfrac{1}{2} + \tfrac{1}{\pi}\int_0^\infty \mathrm{Re}\left(\frac{e^{-ix\ln K}\cdot\phi(x+yi)}{ix}\right)dx \qquad y=-1,0$$

The Heston model (5)

Hence, we need to be able to evaluate $\phi(x+yi)$ for:

- $x \geq 0$ and $y \in \{-1,0\}$

Carr and Madan [1999] have a more efficient representation that actually extends the requirement to:

- $x \geq 0$ and $y \leq -1$

In related work we show that the Carr-Madan formulation allows us to push Heston to the limit (i.e. value options with extreme strikes and maturities).

The complex logarithm

The characteristic function in the Heston model is:

$$\phi(u) = \mathbb{E}_t\left[\exp(iu \ln S(T))\right] = e^{C(T-t,u)+D(T-t,u)\cdot V(0)+iu\ln F(t,T)}$$

Since D does not pose a problem, we focus on C:

$$C(\tau,u) = \frac{\kappa\theta}{\omega^2}\left(\left(\beta(u)+d(u)\right)\tau - 2\ln\left(\frac{c(u)e^{d(u)\tau}-1}{c(u)-1}\right)\right)$$

$$\beta(u) = \kappa - \rho\omega u i$$

$$d(u) = \sqrt{\beta(u)^2 + \omega^2 u(u+i)}$$

$$c(u) = \frac{\beta(u)+d(u)}{\beta(u)-d(u)}$$

The complex logarithm (2)

Although the square root is also multivalued, one can easily show that this does not pose a problem here. The only other multivalued function we encounter is the complex logarithm. If $z = re^{i\theta}$ for $r \in (0,\infty)$ and $\theta \in \mathbb{R}$, then the complex logarithm is given by:

$$\ln z = \ln|r| + (t + 2\pi n)i$$

where $t \in [-\pi, \pi)$ is the principal argument, and $n \in \mathbb{Z}$. If we, as many, implement the Heston model without noticing this issue, we will have typically restricted the complex logarithm to its principal branch ($n = 0$).

The complex logarithm (3)

This approach will produce discontinuities in the option price for moderate to long maturities, and large values of κ. Depicted is the integrand of Π_1 as a function of v, taken from Kahl and Jäckel [2005]:

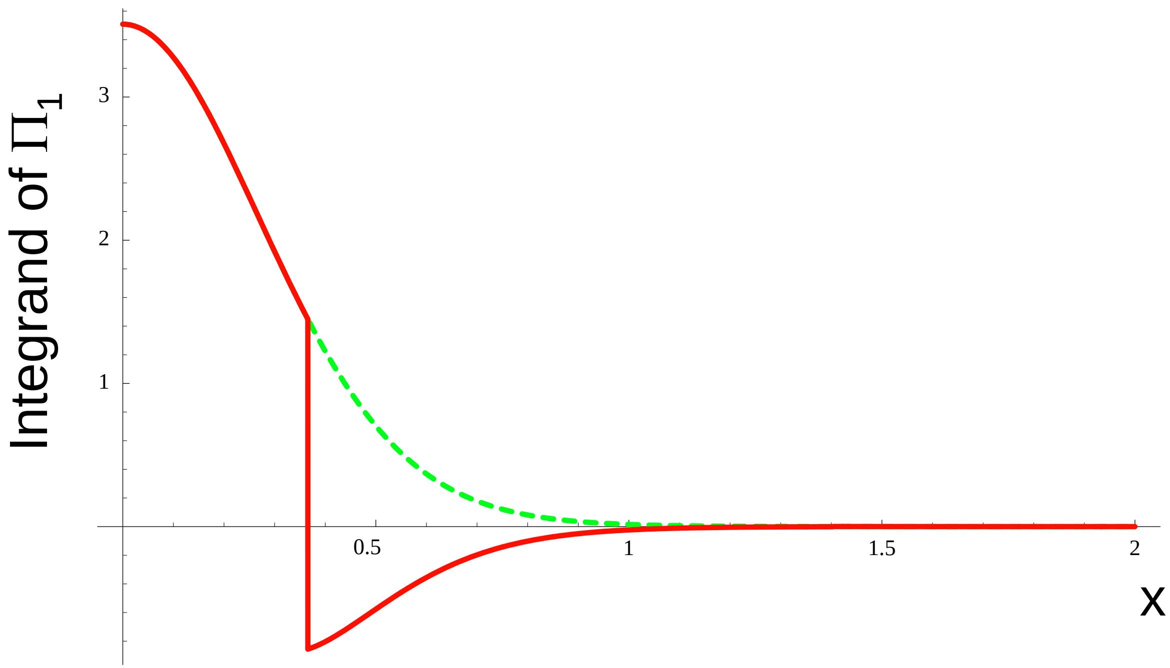

The complex logarithm (4)

Many authors (e.g. Schöbel and Zhu [1999], Lee [2004]) have suggested to keep track of the number of jumps by comparing to neighbouring points in the numerical integration scheme. This will be hard:

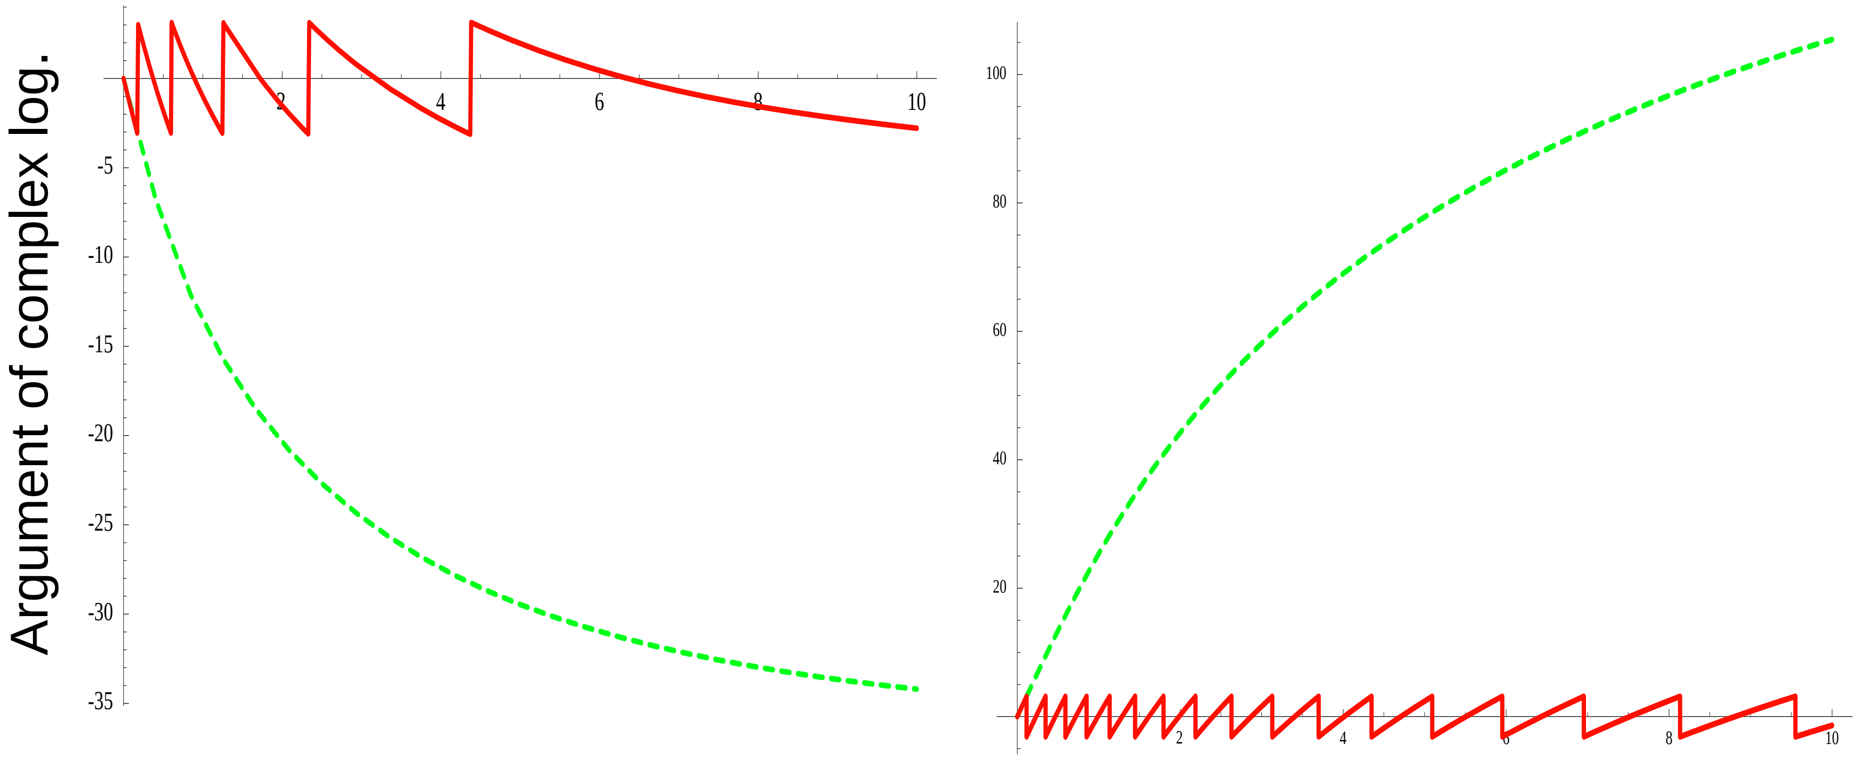

The complex logarithm (5)

Luckily, Kahl and Jäckel [2005] recently came up with an algorithm that seems to work. The main idea is that in the calculation of terms like:

$$z - 1 = re^{i\theta} - 1$$

we lose information on the exact phase of the number. Let us denote the phase interval by:

$$p(\theta) = \mathrm{int}\left[\frac{\theta + \pi}{2\pi}\right]$$

which is 0 if $\theta \in [-\pi,\pi)$, 1 if it is in $[\pi,3\pi)$, etc.

The complex logarithm (6)

If subtracting 1 does not change the phase interval of the complex number, then we can write:

$$z-1=re^{i\theta}-1=|z-1|\,e^{i\left(\arg(z-1)+2\pi p(\theta)\right)}$$

This is the algorithm Kahl and Jäckel dub the rotation count algorithm. The algorithm seems to work perfectly fine in the evaluation of $\phi(u)$ for any allowed value of $u \in \mathbb{C}$, but this is far from a formal proof. Clearly we do not want to have a counterexample pop up in the week before bonuses are announced, so a formal proof is required.

The complex logarithm (7)

Can the subtraction of 1 change the phase interval? Yes it can, consider the following complex function:

$$f(x) = (1 - x) - (\tfrac{1}{2} - x)i$$

f(x)+1, f(x), f(x)-1 and f(x)-2
in the complex plane

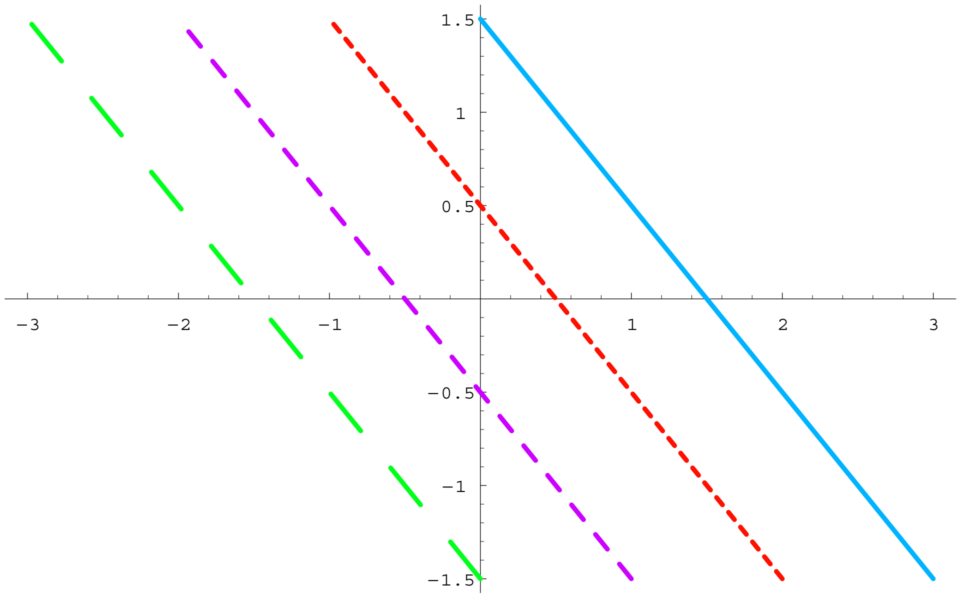

Argument as a function of x

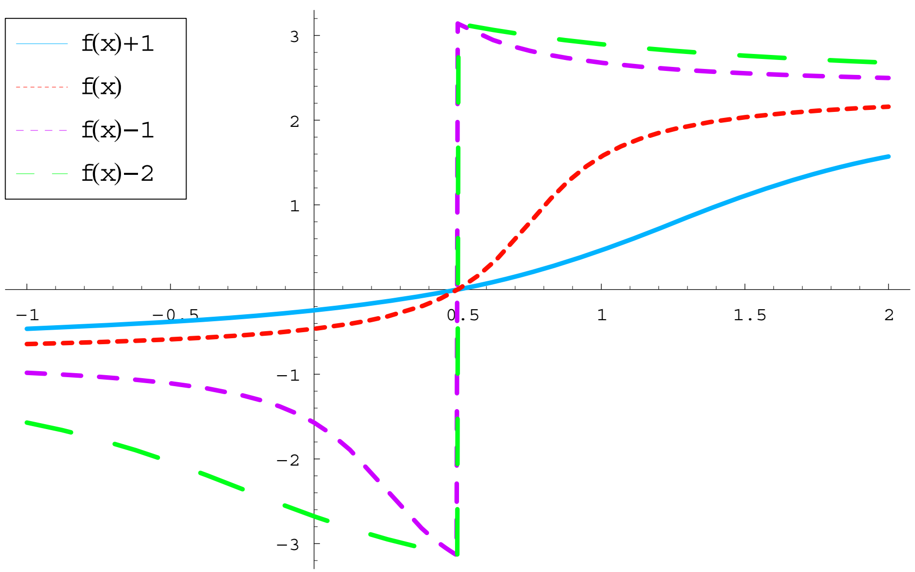

The complex logarithm (8)

If the rotation count algorithm does indeed work, this must mean that in the evaluation of:

$$\ln\left(\frac{c(u)e^{d(u)\tau}-1}{c(u)-1}\right)$$

both $c(u)$ and $c(u)e^{d(u)\tau}$ never cross the real axis in the interval [0,1). Indeed, we are able to prove that:

$$|c(u)| \geq 1 \quad \text{for } -1 \leq \rho \leq \tfrac{\kappa}{\omega}$$

Also, $\mathrm{Im}(c(u)) > 0$ for $\rho > \kappa/\omega$, so the denominator never poses a problem.

The complex logarithm (9)

Remembering that:

$$d(u) = \sqrt{\beta(u)^2 + \omega^2 u(u+i)}$$

we can ensure that $\mathrm{Re}(d(u)) \geq 0$ by always taking the nonnegative solution to the square root. Hence, if $|c(u)| \geq 1$, we immediately have:

$$|c(u)e^{d(u)\tau}| = |c(u)|\, e^{\mathrm{Re}(d(u))\cdot\tau} \geq 1$$

so that the numerator is also solved for $-1 \leq \rho < \kappa/\omega$. In fact, we typically have $\rho << 0$, $\kappa >> 1$ and $\omega \leq 1$, so the algorithm works for all realistic parameter values.

The complex logarithm (10)

The rotation count algorithm greatly simplifies an implementation of the Heston model, and has now been verified for all relevant parameter values. In our paper we show that it can be extended to:

- All extensions of the Heston model with jumps in stock and/or volatility (SVJ, SVCJ);
- The Stein-Stein [1991] and Schöbel-Zhu [1999] stochastic volatility models;
- Broadie and Kaya's exact simulation algorithm for the Heston model.

Euler discretisations

The next step is of course to price exotic options. If we need to price both path-dependent and early exercise features, Monte Carlo simulation is the most flexible method. Focusing on the stochastic variance equation, a naïve Euler discretisation yields:

$$V(t+\Delta t) = \left(1-\kappa\Delta t\right)V(t) + \kappa\theta\Delta t + \omega\sqrt{V(t)}\cdot\Delta W_V(t)$$

This has a positive probability of becoming negative:

$$\mathbb{P}\left(V(t+\Delta t)<0\right) = N\left(\frac{-\left(1-\kappa\,\Delta t\right)V(t)-\kappa\theta\Delta t}{\eta\sqrt{V(t)\Delta t}}\right)$$

Euler discretisations (2)

Secondly, the usual convergence (weak or strong) theorems require the drift and diffusion to:

- Satisfy a linear growth condition;
- Be globally Lipschitz.

Global Lipschitzianity of the diffusion term means that there is a constant C, such that for all x, y:

$$\left| \omega\sqrt{x} - \omega\sqrt{y} \right| \leq C \cdot |x - y|$$

but clearly this is not the case here!

Euler discretisations (3)

Two issues:

- How to “fix” negative values of the variance?
- Prove that the scheme converges.

In our paper we unify all known “fixes” as:

$$V(t+\Delta t) = f_1\big(V(t)\big) - \kappa\Delta t \cdot \big(f_2\big(V(t)\big) - \theta\big) + \omega\sqrt{f_3\big(V(t)\big)} \cdot \Delta W_V(t)$$

where f_1, f_2, f_3 satisfy:

- $f_i(x) = x$ for $x \geq 0$;
- $f_3(x) \geq 0$ for all x.

Euler discretisations (4)

Two of these "fixes" seem common practice:

- Absorption: set V(t) = 0, and continue;
- Reflection: set V(t) = -V(t), and continue.

From a practical point of view, we want that "fix" that creates the smallest amount of bias. In our paper we make a strong case for *full truncation*:

$$V(t+\Delta t) = V(t) - \kappa\Delta t \cdot \left(V(t)^{+} - \theta\right) + \omega\sqrt{V(t)^{+}} \cdot \Delta W_V(t)$$

i.e. we let V(t) drift deterministically once it becomes negative.

Euler discretisations (5)

Let us compare the following schemes:

- Euler schemes (Absorption, reflection, full truncation, and schemes from Deelstra and Delbaen [1998], Higham and Mao [2005]);
- Exact scheme by Broadie and Kaya [2005];
- “2^{nd}” order scheme by Kahl and Jäckel [2005];

in terms of bias and root-mean-squared (RMS) error, when pricing a 5y European call option:

$$\text{RMS} = \sqrt{\text{bias}^2 + \text{stderr}^2}$$

Euler discretisations (6)

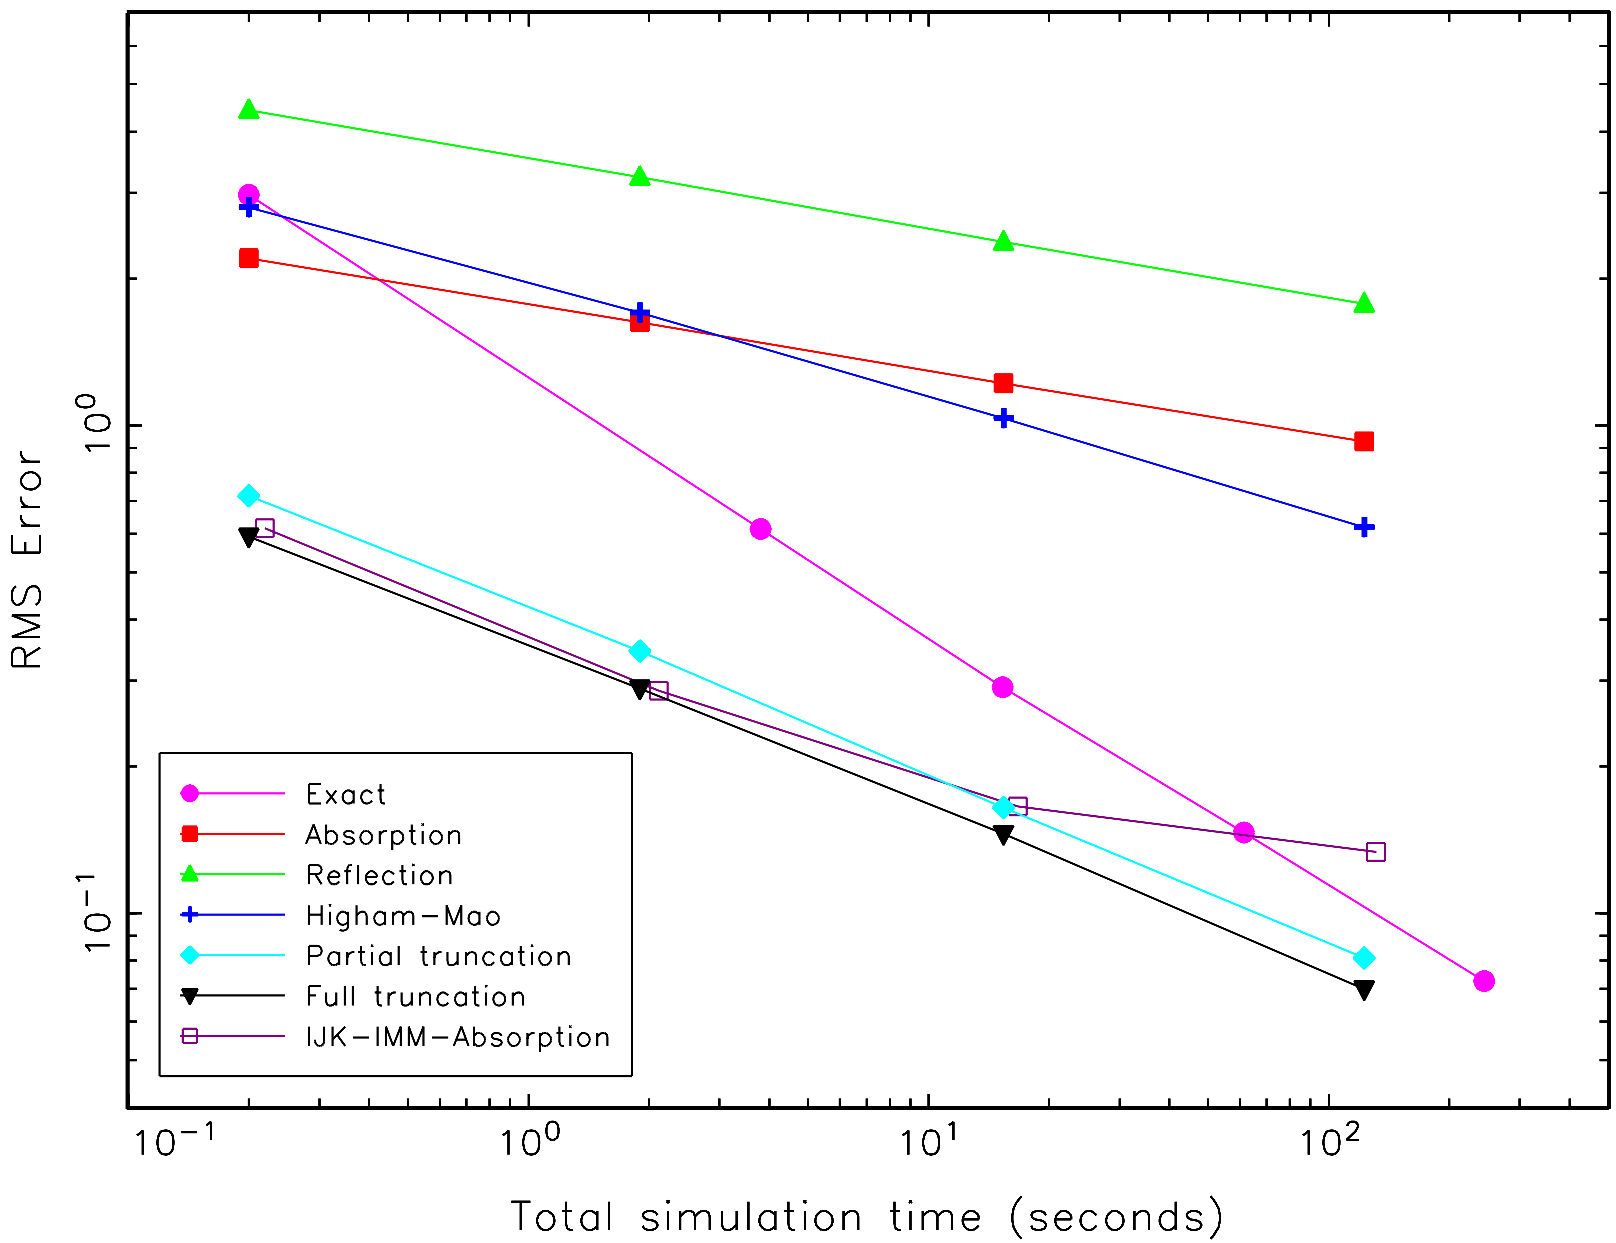

Euler discretisations (7)

Bias as a function of the strike of the option, expressed in implied volatility (in bp, i.e. 1% = 100 bp):

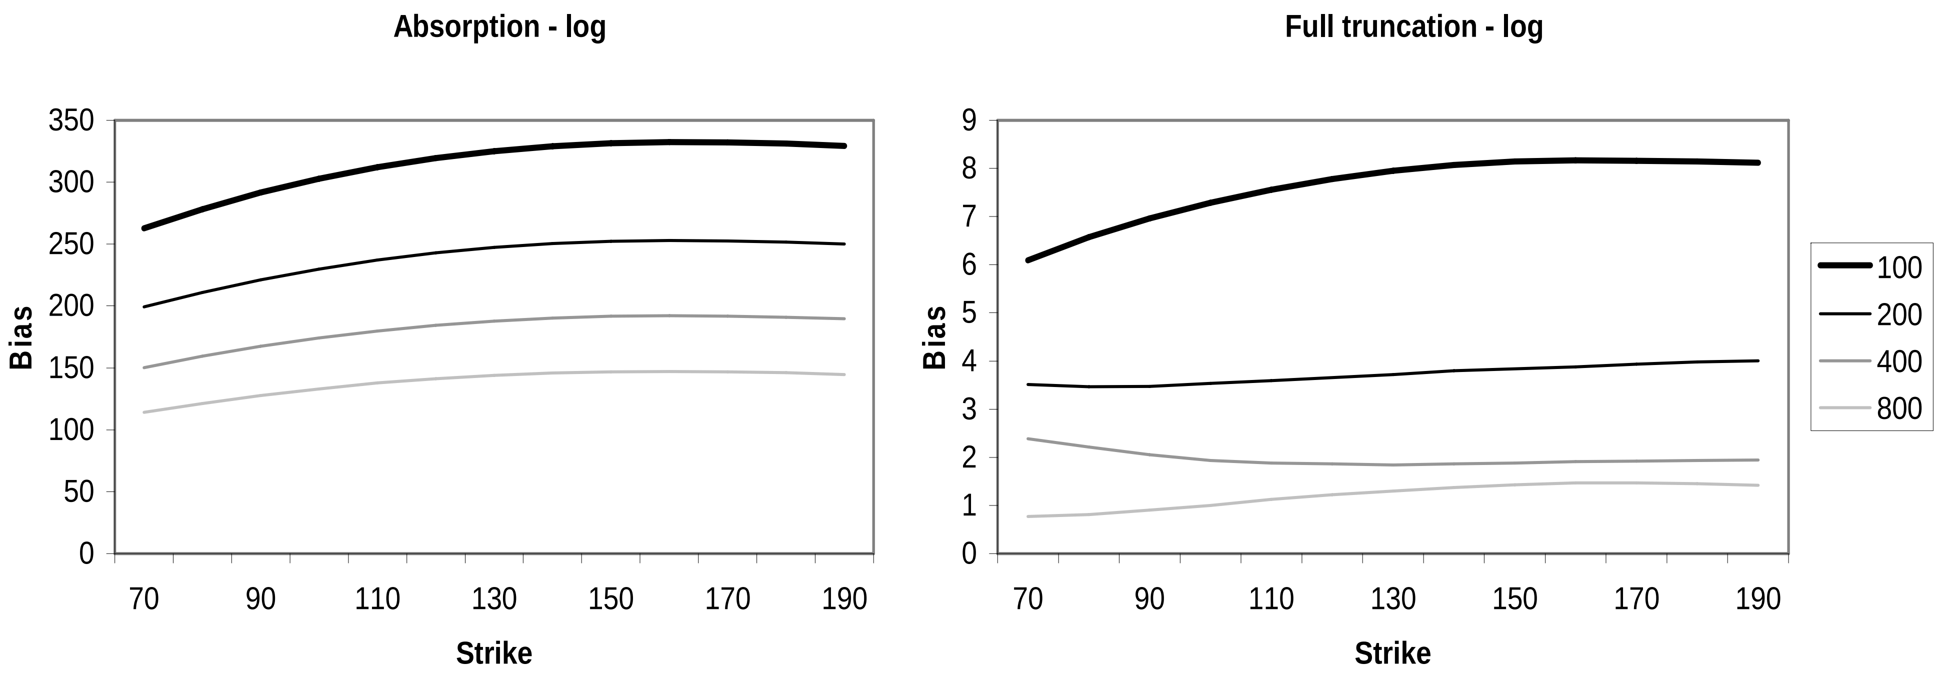

Clearly using a good “fix” has a huge impact!

Conclusions

- Touched upon the problem of the complex logarithm in the Heston model;
- Demonstrated that the rotation count algorithm works for all plausible parameter values, thus allowing for a very efficient implementation;
- Touched upon the problem of Euler discretisations of the mean-reverting square-root process;
- Demonstrated that thinking about the boundary behaviour of these discretisations is very important and can lead to great efficiency improvements;