

Advancing the universality of quadrature methods to any underlying process for option pricing

Article (Published Version)

Chen, Ding, Härkönen, Hannu J and Newton, David P (2014) Advancing the universality of quadrature methods to any underlying process for option pricing. *Journal of Financial Economics*, 114 (3). pp. 600-612. ISSN 0304-405X

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/id/eprint/49885/>

This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

Copyright and reuse:

Sussex Research Online is a digital repository of the research output of the University.

Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.



Advancing the universality of quadrature methods to any underlying process for option pricing[☆]

Ding Chen, Hannu J. Härkönen, David P. Newton^{*}

Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, UK

ARTICLE INFO

Article history:

Received 25 April 2013

Received in revised form

14 October 2013

Accepted 3 November 2013

Available online 2 August 2014

JEL classification:

G13

C63

Keywords:

Universal quadrature

QUAD

Option pricing

Numerical techniques

Transition density function

ABSTRACT

Exceptional accuracy and speed for option pricing are available via quadrature (Andricopoulos, Widdicks, Duck, and Newton, 2003), extending into multiple dimensions with complex path-dependency and early exercise (Andricopoulos, Widdicks, Newton, and Duck, 2007). However, the exposition is incomplete, leaving many modelling processes outside the Black-Scholes-Merton framework unattainable. We show how to remove the remaining major block to universal application. Although this had appeared highly problematic, the solution turns out to be conceptually simple and implementation is straightforward (we provide code on the Journal of Financial Economics website at <http://jfe.rochester.edu>). Crucially, the method retains its speed and flexibility across complex combinations of option features but is now applicable across other underlying processes.

© 2014 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/3.0/>).

1. Introduction

Numerical techniques are widely required in derivatives pricing, since it is often the case that no analytic equation has been found for the valuation of a particular class of option. Ideally, in place of numerical methods, we would eventually have a suite of analytic solutions to cover all derivatives pricing situations or, failing that, analytic approximations of sufficient accuracy and utility for all practical cases. For example, the work of Kristensen and Mele (2011) is highly encouraging, yet we remain a long way from generality along this route. Beyond the solutions of Black and Scholes (1973) and Merton (1973) and a limited set of other cases (generally those with no early exercise),

numerical techniques are frequently required. The available numerical techniques are classified as trees (Cox, Ross, and Rubinstein, 1979), solution of partial differential equations usually by finite difference methods starting with the most basic explicit method (Brennan and Schwartz, 1977), Monte Carlo simulation (Boyle, 1977) and quadrature in the form of the QUAD technique (Andricopoulos, Widdicks, Duck, and Newton, 2003).

Each of these has been the subject of modification and refinement, especially in relation to handling early exercise with Monte Carlo (Longstaff and Schwartz, 2001) and path-dependent features with the other techniques. Andricopoulos, Widdicks, Newton, and Duck (2007) further developed QUAD into a flexible, robust option pricing tool of wide applicability, covering multiple dimensions, early exercise and heavy path-dependence in complex combinations of exercise features. QUAD is usually overwhelmingly fast, making it especially useful in those cases in which computation with other methods is inconveniently slow. However, it has largely been limited to the Black-Scholes-Merton framework. Overcoming

[☆] David P. Newton thanks the Cambridge Judge Business School, University of Cambridge, UK, for facilities provided during a visiting fellowship in 2013.

^{*} Corresponding author. Tel.: +44 115 8467414.

E-mail address: David.Newton@nottingham.ac.uk (D.P. Newton).

this limitation is the subject of our paper, which completes the exposition of the method.

Just as the mathematics of trees, finite difference and Monte Carlo approaches were all known and used in the natural sciences and engineering long before their introduction into finance, basic quadrature goes back centuries. In essence, it is the calculation of an area under a graph via an approximation, splitting the area into a series of shapes, such as rectangles, and summing their individual areas. Taking smaller shapes produces more accurate results, converging on the correct one. Well-known methods for doing this are the Trapezium Rule, Simpson's Method and Gaussian quadrature, and there are others. Each has differing properties and is more or less easy to program, but of particular interest is the rate of convergence to a correct solution as the number of calculations is increased in progressively finer approximations.

A key concept in the financial application of quadrature, sometimes not appreciated, is that the mathematical quadrature component is merely a computational engine to be chosen appropriately to fit into the wider calculations of the particular options problem (Andricopoulos, Widdicks, Duck, and Newton, 2003; Andricopoulos, Widdicks, Newton, and Duck, 2007). Thus, even the very simple Trapezium Rule can be adequate when elements in the wider calculations are less refined. Similarly, Gaussian quadrature, though in itself a very fast scheme, only provides useful extra speed over what may be the best practical engine, Simpson's Method, where unusually heavy calculational demands are made on the quadrature component versus the rest of the computational scheme. We shall return to the engine analogy later, when we show how previously intractable problems in applying quadrature can be circumvented by including a second type of numerical "engine".

The foundation work was presented in the Black-Scholes-Merton framework but (as explained in Section 2.2) the technique applies whenever the conditional probability density function is known. This restricts the immediate use of the technique to the Black-Scholes-Merton setup, to Merton's jump-diffusion model (Merton, 1976) and to certain interest rate models such as those of Vasicek (1977) and Cox, Ingersoll, and Ross (1985). Extension to Merton's process is straightforward. The interest rate models are more subtle, though Heap (2008) has successfully extended the coverage to some (but not all) interest rate derivatives with mean-reverting underlying processes.

A notable advance was made by O'Sullivan (2005), who used the observation that many useful processes without a well-known density function do, nonetheless, have a well understood characteristic function. The density function, as the inverse Fourier transform (FFT) of the characteristic function, can be computed using fast Fourier transform and the output may then be inserted in the QUAD scheme to price derivatives. We refer to this method as FFT-QUAD. O'Sullivan's method applies in particular to exponential Levy processes. This made FFT-QUAD an important advance but it does suffer several drawbacks. First, it requires two integrations even for a derivative on a single underlying process. This brings the complexity of the algorithm to at least $O(N^2)$, where N is the number of grid points used in the numerical integrations; by comparison, the original QUAD has a much better complexity of just $O(N)$ for vanilla

options. Second, it does not cover every option type; for example, the single-variable FFT-QUAD cannot be used to price heavily path-dependent options in stochastic volatility frameworks, since it does not keep track of the evolution of the volatility process in moving from one observation point to the next.

O'Sullivan's FFT-QUAD was improved considerably by the CONV technique of Lord, Fang, Bervoets, and Oosterlee (2007). We refer to this method as CONV-QUAD (Staunton, 2007). This excellent method uses the observation that the fundamental pricing integral may usually be regarded as the convolution (strictly speaking, the cross-correlation) of the payoff and the density function. The beauty of this insight is that the two integrals of FFT-QUAD may then be replaced by two fast Fourier transforms. This brings the complexity of the algorithm down to $O(N \log(N))$ and, for example, for Bermudan options (on M observation points), the complexity remains at $O(MN \log(N))$, which beats even QUAD's $O(MN^2)$. The CONV-QUAD method applies to exponential Levy processes and, hence, in particular to the Black-Scholes-Merton model, thereby improving on the speed and accuracy of the plain QUAD technique of Andricopoulos, Widdicks, Duck, and Newton (2003) and Andricopoulos, Widdicks, Newton, and Duck (2007). Due to its nearly linear speed, it clearly replaces plain QUAD as the fastest method for a great many cases.

Useful as these developments were, the road to full universality for underlying processes remains blocked. The CONV method cannot be applied to, for example, the CEV or the Heston processes with early exercise and, while a single-variable characteristic function for the latter has been used in O'Sullivan (2005) and Fang and Oosterlee (2008) to price European options, a universal QUAD-style treatment of these processes is still lacking.

In this paper we return to the methods of Andricopoulos, Widdicks, Duck, and Newton (2003) and Andricopoulos, Widdicks, Newton, and Duck (2007) and provide option pricing techniques for the missing underlying processes. At the core of this extension is the use of closed-form approximations for the appropriate single- or two-variable transition density functions. By using these approximations we can price complex combinations of option features precisely as if we were working in the Black-Scholes-Merton framework. Thus, we advance the range of the earlier papers without losing their generality; the universality promised in the title of the first paper (Andricopoulos, Widdicks, Duck, and Newton, 2003) is finally arrived at.

2. Basics

Descriptions of the QUAD method can be found in Andricopoulos, Widdicks, Duck, and Newton (2003), Andricopoulos, Widdicks, Newton, and Duck (2007) and Chen (2013). We also provide a detailed appendix on the Journal of Financial Economics website (<http://jfe.rochester.edu>).

2.1. QUAD in the Black-Scholes-Merton framework

Start with the well-known Black-Scholes-Merton partial differential equation for an option with an underlying

asset following geometric Brownian motion:

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + (r - D_c)S \frac{\partial V}{\partial S} - rV = 0, \quad (1)$$

where V is the price of the derivative product and is a function of S , the value of the underlying asset, and time, t . The risk-free interest rate is r , the volatility of the underlying asset is σ , and continuous dividend yield is D_c . In the earlier versions of QUAD, we found it convenient to take the log transform of the underlying asset (the method works equally well using the actual asset price). Suppose y is the corresponding value of the transform of the underlying at time $t + \Delta t$, where Δt is a time step:

$$x = \log S \quad (2)$$

and

$$y = \log S_{+\Delta t} \quad (3)$$

It is important to note that Δt is not restricted to small time periods; for example, were QUAD to be applied to a plain European call option (no need; we have the analytic solution!) then the complete time to expiry would be taken in a single time step, Δt . At expiry, the final condition (payoff) becomes $\max(e^y - X, 0)$, where X is the strike (exercise price). The solution for the value of the option at time t on an underlying asset S is then

$$V(x, t) = A(x) \int_{-\infty}^{\infty} B(x, y) V(y, t + \Delta t) dy, \quad (4)$$

where

$$A(x) = \frac{1}{\sigma\sqrt{2\pi\Delta t}} \exp\left(-\frac{kx}{2} - \frac{\sigma^2 k^2 \Delta t}{8} - r\Delta t\right) \quad (5)$$

and

$$B(x, y) = \exp\left(\frac{-(x-y)^2}{2\sigma^2\Delta t} + \frac{ky}{2}\right), \quad (6)$$

where, in turn,

$$k = \frac{2(r - D_c)}{\sigma^2} - 1. \quad (7)$$

For a plain vanilla call option, for example, the integrand becomes

$$f(x, y) = B(x, y) \cdot \max(e^y - X, 0). \quad (8)$$

This integral is key. Next, any of the many quadrature methods can be employed as a valuation engine for what is a European option with known payoff. The integration covers a single time step, Δt . The (considerable) advantage comes from treating more complex and interesting options problems as equivalent to series of European options. Regions where there is no boundary condition to deal with can each be jumped in a single step. Although the range of integration is infinite, this is easily handled by truncation of the range, provided the integrand outside the truncated range is suitably small. Highly accurate calculations are then possible to any level required. American exercise is readily handled by extrapolation from the Bermudan case.

If the QUAD grid is constructed to coincide with the discontinuities such as a strike price, barrier or exercise

boundary, then convergence is perfectly smooth and suitable for improvement through extrapolation via a simple Richardson-type procedure, because only distribution error remains. It is so smooth that extrapolated results can often be further extrapolated themselves. This extrapolation is applicable in all cases, including those with early exercise (but see [Andricopoulos, Widdicks, Newton, and Duck, 2007](#), for adaptive quadrature). Assuming the discontinuities are correctly located, for Simpson's rule the extrapolated results converge as $(\Delta y)^8$, which is $1/N^8$. For comparison, a trinomial tree converges merely with $1/N$ or in some cases with only $1/\sqrt{N}$, and finite difference methods at best converge at the rate of $(\Delta S^2, \Delta t^2)$, where ΔS and Δt are the step sizes in the S and t directions respectively.

Impressive as the outcomes are, they are limited in scope for the underlying processes. In order to address this, we will need to reformulate the QUAD method.

2.2. QUAD and transition densities

In [Section 2.1](#), we took as our starting point the Black-Scholes-Merton partial differential equation and formulated the price of the option in a particular fashion in order to pave the way for numerical integration by quadrature. Mathematically, this was a use of the corresponding Green's function. As we move to settings where Green's functions are not available, it will become necessary to reformulate this in terms of probability density functions. For this we simply model the asset price directly.

In this vein, the starting point is

$$V(x, t) = e^{-r\tau} \int_{-\infty}^{\infty} V(y, T) f_{\tau}(y|x) dy, \quad (9)$$

where V remains the value of the option (this time as a function of the asset price), and $f_{\tau}(y|x)$ is the risk-neutral conditional density function of the asset price over time-step $\tau = T - t$.

To see the equivalence with the previous approach, we note that in the Black-Scholes-Merton setup the density reads as

$$f_{\tau}(y|x) = \frac{1}{y\sigma\sqrt{2\pi\tau}} \exp\left(-\frac{1}{2}\left(\frac{\ln y - \ln x + (r - \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}\right)^2\right). \quad (10)$$

To get from Eqs. (9) to (4), we merely need to carry out a change of variables from y to $\log y$. In the original equations, factors $A(x)$ and $B(x, y)$ are isolated for computational purposes only – collecting to $A(x)$ the factors that are independent of y speeds up the computations, but conceptually Eq. (9) is the key.

In particular, we see that we can, in principle, use quadrature techniques as long as we have some method of computing the values $f_{\tau}(y|x)$. One particularly natural way to achieve this is through characteristic functions. This is the route taken in [O'Sullivan \(2005\)](#) and [Lord, Fang, Bervoets, and Oosterlee, \(2007\)](#), and indeed progress can be made with this approach: it applies to a great many processes and the use of fast Fourier transform (see [Carr and Madan, 1999](#)) allows speedy implementation. As

explained in the Introduction, the CONV-QUAD method of Lord, Fang, Bervoets, and Oosterlee (2007) is particularly fast. This method hinges on one absolutely crucial condition: the transition density $f_\tau(y|x)$ for the log-asset price must depend only on $y - x$. This condition is satisfied (from their very definition) by exponential Levy processes, which makes CONV-QUAD the method of choice for this large class of processes. However, this condition is not satisfied in many interesting cases, including several local and stochastic volatility models, and we are left with a roadblock to the generality of the methods under the umbrella term QUAD.

We resolve this final difficulty to completion of the method by using approximations of the transition density. This opens up the remaining previously unattainable areas for pricing. Just as a fundamental feature of QUAD, in earlier papers, is interchangeability of quadrature methods as calculational engines, here, in the same spirit, any sufficiently accurate approximation to a transition density function can be employed. In some cases, even where a density function is available, it transpires that the approximation route is superior. To illustrate this, we use two sources: Ait-Sahalia (2008) and Henry-Labordère (2009), though we favour the former as the best all-purpose engine. In the process of formulating QUAD into this more generally applicable version, all the capabilities to handle complex combinations of features in Andricopoulos, Widdicks, Duck, and Newton (2003) and Andricopoulos, Widdicks, Newton, and Duck (2007) are retained.

In Section 3, we begin with a discussion of local volatility models, using this as a way of introducing approximations of the density function and covering several practically useful models. We then move to spot interest rate models, for which approximation techniques can be used to improve the techniques of Heap (2008) alluded to in the Introduction. Finally, we explain how approximations can be used to price options under stochastic volatility models such as Heston's and SABR (both with early exercise) that could not previously be priced via QUAD-style techniques. We also choose these models to illustrate the universality now achieved for QUAD methods because they are stochastically two dimensional and, thus, may be regarded as somewhat more difficult to implement than single dimensional processes.

3. Local volatility models

Local volatility models have been a popular route for practitioners to fit the day's volatility smile, from plain vanilla options, prior to pricing more complex options. Discrete and continuous derivations are due to Derman and Kani (1994) and Dupire (1994). Although calibrated to a current smile, these models generally fail to reproduce future volatility smiles; in other words, they are not well suited to modelling the dynamics of smiles (for early empirical work on the dynamics of implied volatility surfaces, see Dumas, Fleming, and Whaley, 1998). Therefore, in Section 5.2 we show how to handle fully stochastic volatility models while here we focus on (time-homogeneous)

local volatility models, of the form

$$dS_t = rS_t dt + \sigma(S_t) dW_t, \tag{11}$$

where $\sigma(S_t)$ is a deterministic function (the local volatility function) of the asset price S_t , generalizing the constant volatility under Black-Scholes-Merton. For this family of models, we use an approximation of the appropriate density function. There are two reasons for this: first, the density function need not, in general, be known in a closed or even semi-closed form and, second, the approximation is sometimes easier to work with than the actual density function, even when this is known.

We next describe a general method of approximating the density function due to Ait-Sahalia. We then consider specific examples: the constant elasticity of variance (CEV) process of Cox (1996) and quadratic local volatility. Once the relevant approximations have been found, using them is as routine as using the density function for the Normal Distribution in the well-known Black-Scholes equations for calculating plain vanilla European option values.

3.1. Ait-Sahalia's algorithm

We now describe briefly Ait-Sahalia's method, which is explained with numerous examples worked in his paper (Ait-Sahalia, 1999). For theoretical justification of the computations, see his later paper (Ait-Sahalia, 2002). The algorithm begins from normalization of the process by modifying it to one with unit diffusion. From Ito's Lemma, this is achieved by taking $Y_t = \gamma(X_t)$, where

$$\gamma(x) = \int^x \frac{du}{\sigma(u)} \tag{12}$$

and where an alternate σ is $\sigma(S_t)$. The rest of the algorithm involves working out an approximation for the transition density $p_Y(t, y|y_0)$ of Y_t . This gives us the approximation we are looking for, as

$$p_X(t, x|x_0) = \frac{p_Y(t, \gamma(x))\gamma'(x_0)}{\sigma(x)}. \tag{13}$$

To find an approximation for $p_Y(t, y|y_0)$, we first define

$$\mu_Y(y) = \frac{\mu(\gamma^{-1}(y))}{\sigma(\gamma^{-1}(y))} - \frac{1}{2} \frac{\partial \sigma}{\partial x}(\gamma^{-1}(y)). \tag{14}$$

The significance of this function is that Y_t satisfies the stochastic differential equation $dY_t = \mu_Y(Y_t) dt + dW_t$. Our approximation to order L for $p_Y(t, y|y_0)$ is then

$$\tilde{p}_Y^L(t, y|y_0) = \frac{1}{\sqrt{t}} \phi\left(\frac{y-y_0}{\sqrt{t}}\right) \exp\left(\int_{y_0}^y \mu_Y(w) dw\right) \sum_{l=0}^L c_l(y|y_0) \frac{t^l}{l!}, \tag{15}$$

where ϕ is the density of the normal distribution, $\phi(z) = e^{-z^2/2} / \sqrt{2\pi}$, and where the coefficients $c_l(y|y_0)$ are defined recursively as follows: set $c_0(y|y_0) = 1$ and define, for $j \geq 1$,

$$c_j(y|y_0) = \frac{j}{(y-y_0)^j} \int_{y_0}^y (w-y_0)^{j-1} \left(\lambda(w)c_{j-1}(w|y_0) + \frac{\partial^2 c_{j-1}}{\partial w^2}(w|y_0) \right) dw, \tag{16}$$

where $\lambda = -\frac{1}{2}(\mu_Y^2(y) + (\partial\mu/\partial y)(y))$.

While the expressions may look somewhat complicated, the algorithm is very easy to implement. All the steps can be carried out on a symbolic algebra package such as Mathematica or Matlab/MuPAD. Our computations were carried out using a ten to fifteen line script in which the user can specify the order of the expansion (number L in the above equations), and the program carries out the recursion in a matter of seconds. The resulting approximation can then be ported to whichever programming language one wishes to use for more numerically intensive algorithms.

An alternate method is to work out γ , γ^{-1} and μ_Y with pen and paper, leaving the recursion to the computer. For QUAD applications, $\gamma(S)$ and $\gamma(S_0)$ are then calculated within the routine computing the density function and plugged into the approximation for p_Y . The advantage of this approach is that the density calculation becomes slightly faster because $\gamma(S)$ is precomputed and the approximation for p_Y tends to be slightly easier to work with.

Finally, we note that with $t=0.1$, first- and second-order expansions are typically sufficient for achieving root mean square errors of order 10^{-5} and 10^{-8} , respectively. Similar accuracy was reported by Ait-Sahalia (1999).

3.2. CEV process

We apply the procedure of Section 3.1 first to Cox's CEV process (Cox, 1996). We begin here for two reasons. First, for this model the density function exists in a (semi-) closed form, so that it is easy to compare the approximation $p_t^0(f|f_0)$ with the actual value. Second, we demonstrate the highly convenient feature that use of approximations to the density function, far from being inferior, can in some cases be a superior route.

3.2.1. The density function and related issues

We use the following formulation for the CEV model:

$$dS_t = rS_t dt + \sigma S_t^\beta dW_t. \tag{17}$$

Moreover, we require that the process is absorbed by zero: if $S_t=0$ for some $t \geq 0$, then $S_\tau = 0$ for all $\tau \geq t$ as well.

The evolution of this process is fairly well understood. In particular, the transition density is known in a semi-closed form. The prefix “semi-” is explained by the presence of Bessel functions in the expression for the density function. In fact, if $r=0$, we have (for $S > 0$)

$$p_t^0(S|S_0) = \frac{S^{1/2-2\beta}}{\sigma^2|1-\beta|t} \sqrt{S_0} \exp\left(-\frac{S^{2(1-\beta)} + S_0^{2(1-\beta)}}{2\sigma^2(1-\beta)^2t}\right) I_\nu\left(\frac{(SS_0)^{1-\beta}}{\sigma^2(1-\beta)^2t}\right) \tag{18}$$

where $\nu = 1/2(1-\beta)$ and I_ν denotes the modified Bessel function of the first kind:

$$I_\nu(z) = \left(\frac{z}{2}\right)^\nu \sum_{k=0}^\infty \frac{\left(\frac{z}{2}\right)^{2k}}{k! \Gamma(\nu+k+1)}. \tag{19}$$

A process S_t with a positive drift $r > 0$ can be obtained from one with zero drift; call it X_t , as follows: $S_t = e^{rt} X_{\tau(t)}$,

where

$$\tau(t) = \frac{1}{2r(\beta-1)} \left(e^{2rt(\beta-1)} - 1 \right). \tag{20}$$

The density of S_t is then related to p_t^0 as follows:

$$p_t(S_t|S_0) = e^{-rt} p_{\tau(t)}^0(e^{-rt} S_t|S_0). \tag{21}$$

Whether $r=0$ or not, we also have a positive (though frequently negligible) probability of default

$$p(0|S_0) = 1 - \int_0^\infty p_t(x|S_0) dx \tag{22}$$

Finally, by our assumption that the process absorbs at zero, we have $p_t(S|S_0) = 0$ for $S < 0$.

In principle, we can now price exotic and path-dependent options using the algorithms of Andricopoulos, Widdicks, Duck, and Newton (2003) and Andricopoulos, Widdicks, Newton, and Duck (2007) by replacing the Green's function by Eq. (18). The only difference is that the CEV process admits a non-zero probability of S_t hitting zero and this must in some way be incorporated into the QUAD scheme. We will return to this point shortly.

The CONV-QUAD method of Lord, Fang, Bervoets, and Oosterleem (2007) cannot be used in the CEV context. First of all, CONV requires that we work with log-asset price $s_t = \log S_t$, and it is unclear how this should be done when S_t has a non-negative (and sometimes significant) chance of hitting zero. More seriously, the CEV process does not satisfy the CONV condition that the density function can be written as

$$f(s_t|s_0) = f(s_t - s_0). \tag{23}$$

Notice that this would imply in particular that $s_{t+\Delta t} - s_t$ should be independent of s_t , and this fails for the CEV process. To see why, we use Ito's Lemma to find that

$$ds_t = \left(r - \frac{1}{2} \sigma^2 e^{2(\beta-1)s_t} \right) dt + \sigma e^{(\beta-1)s_t} dW_t. \tag{24}$$

We, therefore, have

$$s_{t+\Delta t} = s_t + \int_t^{t+\Delta} \left(r - \frac{1}{2} \sigma^2 e^{2(\beta-1)s_u} \right) du + \int_t^{t+\Delta} \sigma e^{(\beta-1)s_u} dW_u, \tag{25}$$

so that $s_{t+\Delta t} - s_t$ is independent of the initial state s_t if and only if $\beta=1$, i.e. if we have geometric Brownian motion.

3.2.2. Implementation

We can now use the QUAD methods of Andricopoulos, Widdicks, Duck, and Newton (2003) and Andricopoulos, Widdicks, Newton, and Duck (2007) to price exotic options: we merely have to replace the Black-Scholes density function with the CEV density function. There are, however, two points of note.

First, with certain parameters the CEV has a significant probability of hitting $S_T=0$ over a time step from t to $T = t + \tau$, and this probability cannot be read directly from the density function. However, we have

$$\mathbb{P}(S_T = 0 | S_t = x) = 1 - \int_0^\infty p_\tau(y|x) dy. \tag{26}$$

If this p_0 is non-negligible, we should replace Eq. (9) by

$$V(x, t) = e^{-rt} \left(V(0, T) \mathbb{P}(S_T = 0 | S_t = x) + \int_0^\infty V(y, T) f_\tau(y|x) dy \right). \tag{27}$$

In practice, it could make sense to compute whether $S_t=0$ is significant over the entire lifespan of the option. If not, we can ignore this point and use precisely the same QUAD routines as with geometric Brownian motion; however, modifying the algorithm is a trivial matter. This absorbing barrier at zero is well known. [Rebonato \(2004\)](#), for example, provides a financial commentary and a review of technical solutions for alternate numerical methods.

The second and more serious point is that the presence of the Bessel function in the density function equation (18) is troublesome. Function $I_\nu(z)$ eventually explodes and while the growth is, in theory, absorbed by the exponential term preceding it in Eq. (18), this can lead to complications as the implementation tries to evaluate something of the form $0 \cdot \infty$. This is particularly prone to occur when pricing options involving very small time steps, as the t in the denominator then pushes the argument of I_ν toward hazardingly large values.

Ait-Sahalia's approximation for the density function solves the latter problem completely. The computations required for the approximation (see [Section 3.1](#)) are easy to carry out. We find that

$$\gamma(x) = \frac{x^{1-\beta}}{\sigma(1-\beta)}, \tag{28}$$

$$\gamma^{-1}(y) = (\sigma(1-\beta)y)^{1/(1-\beta)}, \tag{29}$$

$$\mu_Y = r(1-\beta)Y_t - \frac{\beta}{2(1-\beta)} Y_t^{-1}. \tag{30}$$

The first-order approximation for $p_X(t, x|x_0)$ is then

$$\tilde{p}_X^L(t, x|x_0) = \frac{\tilde{p}_Y^L \left(t, \frac{x^{1-\beta}}{\sigma(1-\beta)} \middle| \frac{x_0^{1-\beta}}{\sigma(1-\beta)} \right)}{\sigma x^\beta}, \tag{31}$$

where

$$\tilde{p}_Y^L(t, y|y_0) = \frac{1}{\sqrt{t}} \phi \left(\frac{y-y_0}{\sqrt{t}} \right) \exp(r(1-\beta)(y^2 - y_0^2)) + \frac{\beta}{2(\beta-1)} \ln \left(\frac{y}{y_0} \right) \sum_{i=0}^2 c_i(y|y_0) \frac{t^i}{i!}, \tag{32}$$

where in turn $c_0(y|y_0) = 1$ and

$$c_1(y|y_0) = -\frac{r^2(1-\beta)^2(y^3 + y^2y_0 + yy_0^2 + y_0^3)}{6} - \frac{(20r\beta^2 - 16r\beta - 8r\beta^3 + \beta^2 + 4r)(y + y_0)}{8(\beta-1)^2} \frac{\beta}{4(\beta-1)^2yy_0}. \tag{33}$$

The next term is about five times the size of this one, so we have opted not to reproduce it here.

3.3. Quadratic local volatility model

Compare an alternate local volatility form (see also [Dumas, Fleming, and Whaley, 1998](#)) given by

$$dS_t = rS_t dt + \sigma S_t(S_t + b) dW_t, \tag{34}$$

where $\sigma > 0$ and b are constants. We can find an approximation for the density of this process in the same way as before and obtain

$$\gamma(x) = \frac{1}{b\sigma} \ln \left(\frac{x}{x+b} \right), \tag{35}$$

$$\gamma^{-1}(y) = \frac{b}{e^{-b\sigma y} - 1}, \tag{36}$$

$$\mu_Y = \frac{b^2\sigma^2(e^{bY_t\sigma} + 1) - 2r(e^{bY_t\sigma} - 1)^2}{2b\sigma(e^{bY_t\sigma} - 1)}. \tag{37}$$

The first-order approximation for $p_X(t, x|x_0)$ is then

$$\tilde{p}_X^L(t, x|x_0) = \frac{\tilde{p}_Y^L \left(t, \frac{1}{b\sigma} \ln \left(\frac{x}{x+b} \right) \middle| \frac{1}{b\sigma} \ln \left(\frac{x_0}{x_0+b} \right) \right)}{\sigma x(x+b)} \tag{38}$$

where

$$\tilde{p}_Y^L(t, y|y_0) = \exp \left(\frac{(-e^{by\sigma} + e^{by_0\sigma})r}{b^2\sigma^2} + \frac{(y-y_0)(2r-b^2\sigma^2)}{b\sigma} \right) + \ln \left(\frac{1-e^{by\sigma}}{1-e^{by_0\sigma}} \right) \frac{1}{\sqrt{t}} \phi \left(\frac{y-y_0}{\sqrt{t}} \right) \sum_{i=0}^2 c_i(y|y_0) \frac{t^i}{i!}, \tag{39}$$

where in turn $c_0(y|y_0) = 1$ and

$$c_1(y|y_0) = -\frac{r^2(e^{2by\sigma} - e^{2by_0\sigma})}{4b^3\sigma^3(y-y_0)} + \frac{r(r+b^2\sigma^2)(e^{by\sigma} - e^{by_0\sigma})}{b^3\sigma^3(y-y_0)} - \frac{(-2r+b^2\sigma^2)^2}{8b^2\sigma^2} \tag{40}$$

Here, we also opt not to reproduce the much more complicated $c_2(y|y_0)$ term.

3.4. Numerical results

QUAD, in its original and modified forms, is an exceptionally fast method. Our purpose next is to demonstrate that the final modification presented in this paper does not degrade performance to any significant degree.

Computation times depend, of course, on both computer and software. The computers used in this work changed between [Andricopoulos, Widdicks, Duck, and Newton \(2003\)](#) and parts of [Andricopoulos, Widdicks, Newton, and Duck \(2007\)](#), from 550 MHz to 2.4 GHz Pentium-based university cluster computers using an optimized Fortran compiler. For the present paper, work was carried out on a university high performance computing service, using one node of eight cores at 3.0 GHz. Parallel codes were compiled using an Intel C++ compiler and OpenMPI. As with the earlier papers, we seek high intrinsic speed and fast convergence for test cases, such that more complex cases requiring intensive calculations are completed in reasonable time periods. Therefore, our interest is in performances relative to the Black-Scholes-Merton QUAD of [Andricopoulos, Widdicks, Duck, and Newton \(2003\)](#) and

Andricopoulos, Widdicks, Newton, and Duck (2007). We use parameters as follows:

Each of the processes has $S_0 = 10$ and $r = 0.05$. The Black-Scholes-Merton process has $\sigma = 0.2$; the CEV process has $\sigma = 0.6$ and $\beta = 0.5$; and the quadratic process has $b = 1$ and $\sigma = 0.02$ (in the last two cases, σ is chosen so as to make the initial instantaneous volatility $\sigma(S_0)/S_0$ approximately equal to the Black-Scholes volatility 0.2).

Our proxy for accuracy of these computations, denoted by K , is related to the step size as follows:

$$\Delta t = \frac{7.5 \cdot T}{K \log(M)} \quad (41)$$

where T denotes the maturity of the option and M is the number of observation points. The thinking behind this somewhat bizarre looking equation is as follows. First, we want our algorithms to have a constant step size in all grids; second, as in Andricopoulos, Widdicks, Duck, and Newton (2003), we want an integer proxy that is inversely related to the step size; and, finally, we want the achieved accuracy to be about the same across the range of options we are interested in.

We initially took Δt to be a constant times T/KM , but this produced disproportionately dense grids for frequently observed options: the prices of these options were often an order of magnitude more accurate than those of less frequently observed options (and, of course, the computational times were longer). After experimenting with Δt equal to a constant times T/\sqrt{MK} , we finally settled for the equation above, the constant 7.5 being to some extent arbitrary, decided by trial and error. However, we include a proper mathematical treatment of K in the online appendix on the QUAD method (Journal of Financial Economics website at <http://jfe.rochester.edu>).

We record for each of these processes the root mean square error and the total time taken for computation of a package of eight options. For Bermudans, this package consists of options with strikes 10.0 and 10.5, maturities 0.5 and 1.0, giving four combinations; then each with numbers of observations 6 and 30, totalling eight combinations. For barrier options, we consider strikes 10.0 and 10.5, barriers 9.5 and 9.9; the numbers of observations 6 and 30 and the maturity constant at 0.5, totalling eight

combinations. In both cases, the error is computed with respect to the QUAD price obtained using $K=512$. In the case of the CEV model, the reference is computed using the second-order approximation (see below for further comments on the choice of order). We record the root mean square errors of the prices obtained by Richardson extrapolation.

Computational time is for the entire bundle of options. The computational time for a single 6-step or 30-step option is typically about 1/60th or just under 1/4th of the total computational time. Also, 0.01% relative error for a 6-step option can be achieved in about one second for the Black-Scholes-Merton (BSM) process and in less than five seconds for the other processes on a common laptop equipped with Intel Core 2 Duo P7450 at 2.13 GHz.

Tables 1 and 2 illustrate for Bermudan options and down-and-out barrier options. The first comment we make on these results is that the QUAD prices for CEV and quadratic processes converge to the correct price at the same rate as those for the Black-Scholes-Merton process. Secondly, the pricing of each option under the CEV or the quadratic process takes three to five times as long as for the Black-Scholes-Merton process. This reflects the fact that our approximations are slightly more complicated than the usual Black-Scholes-Merton density and, therefore, take longer to compute. Given the enormous computational speed advantage of QUAD, this is a perfectly acceptable result.

The first conclusion was merely a consistency check: equal rate of convergence is what we expected. The second conclusion means that many exotic and path-dependent options can now be priced in local volatility situations essentially as quickly as if we were working in the Black-Scholes-Merton framework. In particular, there is no need to resort to Monte Carlo or solution of partial differential equations via finite difference methods.

Finally, we note the remarkable speed at which the extrapolated prices converge toward the correct value. The Bermudan prices are immediately within about 10^{-8} of the correct price. Bearing in mind that the first-order approximations are themselves accurate to order 10^{-5} , this means that these prices, obtained in less than a second for the Black-Scholes model and in about 1.5 and 3 seconds

Table 1

Bermudan options. A comparison of root mean square errors and computation times (in seconds) for the entire bundle of eight options, as described in the text, with different underlying processes. Extrapolated results are denoted by (e). Conclusion: QUAD calculation times are not seriously degraded by moving to alternate processes.

K	BSM	BSM(e)	CEV	CEV(e)	Quadratic	Quadratic(e)
4	0.001603 [0.0064]		0.001633 [0.0229]		0.0016996 [0.0321]	
6	0.000735 [0.0132]		0.000731 [0.0492]		0.0007336 [0.0705]	
8	0.000429 [0.0226]	4.90×10^{-08} [0.0289]	0.000424 [0.0835]	3.01×10^{-08} [0.1063]	0.0004475 [0.1224]	9.40×10^{-08} [0.1545]
12	0.000188 [0.0487]	4.96×10^{-10} [0.0619]	0.000196 [0.1811]	1.61×10^{-09} [0.2303]	0.0001979 [0.2732]	3.65×10^{-09} [0.3437]
16	1.08×10^{-04} [0.0857]	1.36×10^{-10} [0.1082]	2.07×10^{-04} [0.3125]	1.68×10^{-07} [0.3960]	0.0001104 [0.4793]	5.38×10^{-11} [0.6017]
32	3.07×10^{-05} [0.3301]	1.37×10^{-09} [0.4158]	5.49×10^{-05} [1.2314]	1.44×10^{-09} [1.5439]	5.03×10^{-05} [1.8833]	1.05×10^{-08} [2.3626]

Table 2

Down-and-out barrier options. A comparison of root-mean-square errors and computation times (in seconds) for the entire bundle of eight options, as described in the text, with different underlying processes. Extrapolated results are denoted by (e). Again, there is no serious degradation and QUAD incorporating the extra approximation engine for the density function retains its extreme speed across different underlying processes.

K	BSM	BSM(e)	CEV	CEV(e)	Quadratic	Quadratic(e)
4	0.01101 [0.0059]		0.01099 [0.0205]		0.01151 [0.0313]	
6	0.00684 [0.0127]		0.00683 [0.0446]		0.00719 [0.0697]	
8	0.00494 [0.0224]	1.38×10^{-05} [0.0282]	0.00491 [0.0791]	1.45×10^{-05} [0.0996]	0.00518 [0.1230]	1.37×10^{-05} [0.1543]
12	0.00314 [0.0495]	4.04×10^{-06} [0.0623]	0.00311 [0.1757]	4.82×10^{-06} [0.2203]	0.00330 [0.2809]	4.54×10^{-06} [0.3506]
16	0.00228 [0.0869]	2.12×10^{-06} [0.1092]	0.00226 [0.3110]	2.3×106^{-06} [0.3901]	0.00240 [0.4881]	2.26×10^{-06} [0.6111]
32	0.00107 [0.3391]	6.52×10^{-07} [0.4260]	0.00105 [1.2273]	6.95×10^{-06} [1.5383]	0.00112 [1.9464]	7.2×102^{-07} [2.4345]

Table 3

Computational time of QUAD without density functions (in milliseconds). This illustrates the computational cost of setting up grids and computing numerical integrals (see text).

Option	4	6	8	10	16	32
Bermudan	0.9	1.67	2.67	3.87	8.90	32.3
Barrier	0.79	1.55	2.56	3.80	9.00	33.6

for CEV and quadratic models, respectively, are as accurate as the method allows. If even greater accuracy is required, second-order approximations must be used and then, including extrapolations, the computations can still be carried out in a matter of seconds.

In Table 3, we illustrate the computational cost of a complicated density function by “pricing” an option using a density function that is set to some constant. The option “price” is, of course, meaningless but the computational time indicates how much time is needed merely for setting up the grids and computing the numerical integrals. Even in Black-Scholes QUAD, 89% of the computational effort is used in the computation of grids and so on but for other processes the difference is greater. In practical implementations, therefore, it pays to optimize the computation of the density function!

3.5. Henry-Labordère's algorithm

Just as there is more than one engine of approximation via quadrature, there is more than one for approximating the density function. To demonstrate this, we note that an alternate method of approximating the density function has been found by Henry-Labordère (2009). His method applies nicely to processes that can be transformed into a process X_t satisfying a stochastic differential equation of the form

$$dX_t = \sigma(X_t) dW_t; \tag{42}$$

and so, in particular, the CEV model is covered. Moreover, it is easy to modify QUAD routines so that they use a forward rate as the random variable and, in this way, Henry-Labordère's method becomes as widely applicable as Ait-Sahalia's.

Here, we merely note that in the second-order his approximation is very similar in speed and accuracy to Ait-Sahalia's approximation in the first-order. In contrast with Ait-Sahalia's method, it is not easy to extend Henry-Labordère's expansions to higher order, since this requires knowledge of higher order heat kernel coefficients which cannot be deduced easily from the lower order terms. These expansions are, however, fast and accurate to evaluate (once established) so that they provide a viable alternative to the methods we have described so far.

4. Interest rate models

At this point, it is useful to note problems previously encountered with interest rate models used either with bonds or as a second underlying in more general derivatives pricing. Heap (2008) had some success in applying QUAD to these processes, especially Vasicek (1977), but for the CIR model (Cox, Ingersoll, and Ross, 1985), limitations remain. We show briefly how to circumvent these.

The CIR model is given by

$$dr_t = a(b - r_t) dt + \sigma\sqrt{r_t} dW_t. \tag{43}$$

This process has a known transition density and QUAD type option pricing can be carried with certain restrictions, as before. To see how these restrictions arise, recall that the density function for this process is

$$p_t(r|r_0) = \alpha \exp\left(-\frac{\gamma - \kappa}{\sigma^2}(r_0 - r + \kappa\theta\gamma)\right) \left(\frac{re^{\gamma t}}{r_0}\right)^{q/2} I_q\left(2\sqrt{\alpha^2 r r_0 e^{-\gamma t}}\right), \tag{44}$$

where $\gamma = \sqrt{\kappa^2 + 2\sigma^2}$, $\alpha = 2\gamma/\sigma^2(1 - e^{-\gamma t})$, $q = 2\kappa\theta/\sigma^2 - 1$ and I_q denotes the modified Bessel function of the first kind.

The difficulties with this expression stem from the presence of the Bessel functions. The reasons are precisely as we noted in the context of the CEV model. As with the CEV model, the actual breakdown point depends on the parameters but, as observed by Heap (2008), financially reasonable parameters can be chosen for which the density cannot be evaluated for $t < 0.05$. This makes the pricing of a one-year option with more than 20 observation points impossible.

As with the CEV process, Ait-Sahalia's approximations can be used to rescue the situation. Again, the approximation is the best compromise between speed and accuracy. In fact, the first-order approximation becomes more and more accurate when t approaches zero, so the exact formula and the approximation may be seen to complement one another. QUAD style techniques can then be implemented easily, as with Vasicek's model.

5. Stochastic volatility models

We now turn to stochastic volatility models. In these, the asset price, S_t , follows a geometric Brownian motion within which the asset's variance, v_t , itself follows a stochastic process.

Our aim is to show that Bermudan (leading to American, by extrapolation) and path-dependent options can now be priced via QUAD when the asset price follows such a process. This is necessarily slightly more involved than what has gone before, since we need to keep track of the variance process as well. We do this by considering the joint density function of the pair (S_t, v_t) . To see why this is necessary, consider a Bermudan put option with only two observation points T_1 and T_2 . To price this in a QUAD-like fashion, we start by computing the payoffs of the option for some terminal grid of S_{T_2} values. This presents no problems. We then move to time T_1 . At each grid point at time T_1 we can either exercise (in which case we can readily compute the pay-off) or stick with the option, in which case we have to price a European put option. The question then is what the volatility should be. The variance is stochastic, and the price of the put option is different for different values of v_{T_1} , so clearly we need to use a two-dimensional grid and compute the value of the put option for each point (S_{T_1}, v_{T_1}) . After choosing the maximum between the exercise and non-exercise prices, we obtain a two-dimensional grid of values. Finally, we compute the price of the option at the initial time by summing over the two-dimensional grid. Here, we must weight each of the values by the two-dimensional density function $f_{T_1}(S_{T_1}, v_{T_1}|S_0, v_0)$ (also by other weights, depending our choice of quadrature scheme).

5.1. Heston's model

The Heston (1993) model has proven highly popular, doubtless due to its relative tractability with European style options, so we use it here to demonstrate how the final modification of QUAD presented in this paper can handle such models, including American and other more complex features.

5.1.1. Definition and maximum likelihood expansion

In Heston's model, the asset price follows geometric Brownian motion

$$dS_t = rS_t dt + \sqrt{v_t}S_t dW_t^{(1)}, \tag{45}$$

where the variance v_t is assumed to follow another, mean-reverting process, of the CIR type

$$dv_t = \kappa(\theta - v_t) dt + \sigma_v \sqrt{v_t} dW_t^{(2)}, \tag{46}$$

and the underlying Brownian motions are correlated, that is $dW_t^{(1)}dW_t^{(2)} = \rho dt$ for some $-1 < \rho < 1$. The

interpretation of this model is well documented. We merely remark that the variance reverts toward level θ at rate κ .

The characteristic function of $s_t = \log S_t$ is known in a closed form (see Heston, 1993). This means that vanilla options can be priced quickly using Fourier inversion and FFT. Moreover, even the joint characteristic function is known (see Zhylyevskyy, 2010, 2012). Pricing, therefore, becomes theoretically straightforward: we compute the joint density function from the characteristic function and use two-dimensional QUAD in the spirit of Andricopoulos, Widdicks, Newton, and Duck (2007). This part of the process, though, is quite slow: Zhylyevskyy (2010) reports computation times of about 44 seconds for the Fourier inversion on a computer equipped with Intel Core-2 2.83 GHz and 4 gigabytes of RAM.

It is natural to ask whether some CONV-style idea could be used for computational efficiency but this, too, is unlikely. In fact, it is easy to see that the Heston process, regarded as a two-dimensional process, does not satisfy the (two-dimensional version of) the condition in Eq. (23). To see why, simply consider two scenarios for $v_{t+\Delta t} - v_t$, where v_t is at or far away from the mean-reversion level θ . In the first case, $v_{t+\Delta t}$ is likely to remain near θ , so that the increment is concentrated near zero. In the second case, $v_{t+\Delta t}$ is likely to be inverting closer to θ , leading to a non-zero increment.

Once again, a convenient approximation can be used with QUAD. The (logarithm of the) joint density function has been found by Ait-Sahalia and Kimmel (2007). The idea of the method is to seek an approximation of the functional form

$$L_t(s, v|s_0, v_0) = \ln\left(\frac{1}{2\pi t\sqrt{\Omega}}\right) + \frac{C^{-1}(s, v|s_0, v_0)}{t} + \sum_{j=0}^J C^j(s, v|s_0, v_0) \frac{t^j}{j!}, \tag{47}$$

where the $C^j(s, v|s_0, v_0)$ are functions yet to be defined and Ω is the determinant of the diffusion matrix with respect to a pair of uncorrelated Brownian motions, i.e.

$$\Omega = (1 - \rho^2)\sigma_v v \tag{48}$$

in the case of the Heston model. We view this essentially as a power series with respect to the time variable. The coefficient functions $C^j(s, v|s_0, v_0)$, which are independent of t , are then expanded as Taylor series about (s_0, v_0) and the coefficients of these expansions are worked out recursively using Fokker-Planck equations and, in practice, a symbolic algebra program such as Mathematica. Finally, we obtain an approximation for the density function simply by exponentiating $L_t(s, v|s_0, v_0)$. More details can be found in Ait-Sahalia and Kimmel (2007, pp. 419–421).

In the notation of the present paper, t is the time step in the QUAD procedure, Δt . As with local volatility models, the approximation is asymptotic in time. The approximations, moreover, are sensitive: we find that even with a second-order approximation (i.e. $J=2$) we should not expect the approximation to be accurate for larger steps than 0.1 (see Section 5.1.2 for a solution to this). We also

find that the approximation becomes less accurate as ν_0 approaches zero and that further care needs to be taken with low-likelihood events. As explained above, we are working with a Taylor expansion on (s, ν) and so, if we venture too far away from (s_0, ν_0) , the series will blow up. In practice, this problem can be avoided by setting the density to zero when s is 5 to 6 standard deviations away from s_0 (though the blow ups, if left to occur, are very easy to detect through absurd output values, such as 10^{80}).

5.1.2. Implementation

Once a suitable likelihood approximation has been ported to a preferred programming language, such as Fortran or C++, implementation of QUAD schemes is straightforward. The multi-dimensional QUAD techniques of [Andricopoulos, Widdicks, Newton, and Duck \(2007\)](#) are used, with one of the assets replaced by ν_0 .

The asymptotic nature of the density approximation brings with it some problems, the solution of which is perhaps best explained simply by pricing a European call option under Heston's model. In theory, this can be done in one step by evaluating the expectation

$$\mathbb{E}[(S_t - K)^+ | \mathcal{F}_0] = \int_0^\infty \int_0^\infty (S_t - K)^+ f(S_t, \nu_t | s_0, \nu_0) dS_t d\nu_t. \tag{49}$$

As we have noted, we cannot expect our approximation of the density to be accurate beyond $t=0.1$. To circumvent this problem, we evaluate the expectation in steps. Mathematically, this amounts to repeated use of the tower property

$$\mathbb{E}[- | \mathcal{F}_0] = \mathbb{E}[\mathbb{E}[- | \mathcal{F}_s] | \mathcal{F}_0] \tag{50}$$

where $0 \leq s \leq t$ while in our implementation this means inclusion of dummy layers, where we store to each grid point (s_1, ν_1) the simple expectation worked out from the previous grid.

The same principle applies to the pricing of more complicated options. In terms of implementation, the best solution is to store a constant indicating how large a time step is tolerated, Tol (e.g. 0.1). We then work backward from one observation point to the next as usual, except that if the time step t exceeds this constant then we divide the step into

$$N = \frac{\Delta t}{\text{Tol}} \tag{51}$$

equally spaced steps and add a total of $N - 1$ dummy layers in between the observations.

The need for the dummy layers, of course, adds to the computational cost. To gain an idea of how this works, consider European options with varying maturities $N\text{Tol}$ while keeping the step sizes constant in both s and ν directions. The variance process is mean reverting, so we work with the same number of grid points in this direction for all the layers. As for the (log-)asset price, the i th has $C\sqrt{i/\text{Tol}}$ grid points for some constant C . In total, we end up with

$$C\sqrt{\text{Tol}} \sum_{i=1}^N \sqrt{i} = O(N^{3/2}) \tag{52}$$

computations. This compares unfavourably with $O(\sqrt{N})$ that the plain Black-Scholes QUAD would achieve but is far from disastrous.

5.2. The SABR model

The SABR model ([Hagan, Kumar, Lesniewski, and Woodward, 2002](#)) has proven popular with practitioners and can be linked to the LIBOR market model (e.g. see [Rebonato, McKay, and White, 2009](#)). This model addresses the problem with local volatility models that, although they are constructed to fit the day's smile, they do not deal correctly with the dynamics of volatility surfaces.

The SABR model may be regarded as a stochastic volatility variant of the CEV process. More precisely, we require that

$$dS_t = rS_t dt + \nu_t S_t^\beta dW_t^{(1)}, \tag{53}$$

where $0 \leq \beta \leq 1$ and the volatility satisfies

$$d\nu_t = \sigma_\nu \nu_t dW_t^{(2)}, \tag{54}$$

where $\sigma_\nu \geq 0$ and the Brownian motions are correlated as before.

The SABR model does not appear to have a characteristic function in a closed form and so the approach in the present paper is again required. Thus, in the work of [Henry-Labordère \(2009\)](#), approximations for the density function are found and these can be used for QUAD type pricing as an alternative to [Ait-Sahalia \(2008\)](#). Formally, pricing becomes identical to the two-variable QUAD of [Andricopoulos, Widdicks, Newton, and Duck \(2007\)](#).

5.3. Numerical results

We are now in a position to calculate option prices with the same flexibility and variety of features as in [Andricopoulos, Widdicks, Duck, and Newton \(2003\)](#) and [Andricopoulos, Widdicks, Newton, and Duck \(2007\)](#) but with other underlying processes.

We illustrate via the Heston and SABR models (two-dimensional processes). The method starts from an extremely fast base, with European options generally giving error terms less than order of 10^{-10} within seconds. Consequently, just as in the earlier papers, although the computational burden increases exponentially with the number of observation points, Bermudan and (by extrapolation) American options are still calculated with great speed and precision.

We calculate the root mean square errors of a bundle of eight options using second-order Ait-Sahalia approximation with the same setup as for local volatility models in [Section 3.4](#), with $\sigma_0 = 0.2$. Benchmarking is done using $K=128$. For Heston, $r=0.05$, $\kappa=2$, $\theta=0.3$, $\sigma=0.6$ and $\rho=-0.75$; SABR has $r=0.05$, $\sigma=0.6$, $\beta=0.5$ and $\rho=-0.75$. We again use K as a proxy for accuracy. Sufficient accuracy is attainable with low values of K but we tabulate to higher values to show the slowdown at higher values (we also use high K to obtain benchmark prices).

In Tables 4 and 5, we report errors and computation times for Bermudan options and for down-and-out barrier options. In Tables 6 and 7, we demonstrate convergence of Bermudan prices toward American prices for put and call options with strike 10.5 and 9.8 and maturity 0.5, calculated against number of observations, with number of

Table 4

Bermudan options. Error versus speed for Heston and SABR models with root mean square errors and computation times (in seconds).

K	Heston	SABR
8	0.382701 [16.4]	0.354873 [87.3]
12	0.027617 [78.8]	2.63×10^{-3} [413]
16	4.09×10^{-3} [243]	1.56×10^{-4} [1.25 × 10 ³]
28	1.41×10^{-3} [2.20 × 10 ³]	1.71×10^{-5} [1.12 × 10 ⁴]
36	6.17×10^{-4} [5.99 × 10 ³]	1.14×10^{-5} [3.03 × 10 ⁴]
52	2.02×10^{-5} [2.59 × 10 ⁴]	1.11×10^{-6} [1.30 × 10 ⁵]

Table 5

Down-and-out barrier options. Error versus speed for Heston and SABR models with root mean square errors and computation times (in seconds).

K	Heston	SABR
8	0.69527 [0.79]	0.367201 [7.25]
12	3.70×10^{-3} [10.46]	0.033376 [94.5]
16	2.10×10^{-3} [49.0]	4.42×10^{-3} [447]
24	6.65×10^{-4} [358]	4.56×10^{-4} [3.26 × 10 ³]
32	4.72×10^{-4} [1.34 × 10 ³]	3.9×10^{-6} [1.23 × 10 ⁴]
36	3.07×10^{-5} [2.27 × 10 ³]	3.19×10^{-7} [2.07 × 10 ⁴]

Table 6

Convergence of Bermudan put against number of observation points and computation times [in seconds]. Extrapolated results are denoted by (e).

Observation	Heston	Heston(e)	SABR	SABR(e)
8	0.014623 [1.41 × 10 ³]	0.014932 [2.57 × 10 ³]	0.059798 [7.07 × 10 ³]	3.51×10^{-05} [1.32 × 10 ⁴]
16	0.010713 [1.95 × 10 ³]	0.003410 [3.79 × 10 ³]	0.029903 [9.81 × 10 ³]	3.11×10^{-06} [1.98 × 10 ⁴]
24	0.007543 [2.29 × 10 ³]	0.000479 [4.51 × 10 ³]	0.019935 [1.15 × 10 ⁴]	6.61×10^{-07} [2.25 × 10 ⁴]
32	0.005692 [2.54 × 10 ³]	5.09×10^{-05} [5.03 × 10 ³]	0.014952 [1.28 × 10 ⁴]	6.64×10^{-07} [2.51 × 10 ⁴]
40	0.004557 [2.79 × 10 ³]	5.71×10^{-06} [5.54 × 10 ³]	0.011961 [1.40 × 10 ⁴]	1.52×10^{-06} [2.79 × 10 ⁴]
48	0.003798 [2.85 × 10 ³]	1.35×10^{-06} [5.66 × 10 ³]	0.009967 [1.43 × 10 ⁴]	6.55×10^{-06} [2.83 × 10 ⁴]

Table 7

Convergence of Bermudan call against number of observation points and computation times [in seconds]. Extrapolated results are denoted by (e).

Observation	Heston	Heston(e)	SABR	SABR(e)
8	0.020118 [1.37 × 10 ³]	0.012881 [2.51 × 10 ³]	0.135567 [6.87 × 10 ³]	0.026572 [1.26 × 10 ⁴]
16	0.008138 [1.94 × 10 ³]	0.004317 [3.77 × 10 ³]	0.072595 [9.74 × 10 ³]	0.004407 [1.89 × 10 ⁴]
24	0.003353 [2.28 × 10 ³]	0.001498 [4.49 × 10 ³]	0.048968 [1.15 × 10 ⁴]	0.000957 [2.25 × 10 ⁴]
32	0.001303 [2.52 × 10 ³]	0.000416 [4.99 × 10 ³]	0.036832 [1.27 × 10 ⁴]	0.000225 [2.51 × 10 ⁴]
40	0.000417 [2.71 × 10 ³]	2.25×10^{-05} [5.37 × 10 ³]	0.029484 [1.36 × 10 ⁴]	4.26×10^{-05} [2.70 × 10 ⁴]
48	0.000103 [3.02 × 10 ³]	9.07×10^{-06} [5.73 × 10 ³]	0.024572 [1.52 × 10 ⁴]	3.34×10^{-06} [2.88 × 10 ⁴]

QUAD steps set at $K=52$ (the benchmark is calculated by Bermudan put option price extrapolated with 63 and 64 observation points using $K=128$ QUAD steps). Extrapolation against number of observation points is extremely effective in this case – an error term of order 10^{-5} can be achieved within 40 observation points (the Bermudan prices obtained with $K=52$ themselves contain error term of order 10^{-5} for the Heston model and 10^{-6} for the SABR model).

6. Conclusion

The numerical techniques of quadrature introduced previously are applicable whenever the conditional probability density function is known, restricting the immediate use of the method to the Black-Scholes-Merton world, Merton's jump-diffusion model and certain interest rate models. The method can be extended via Fourier transform for processes without a known density function but with known characteristic functions but it then cannot handle the full range of option features. Worse, the road to full universality for underlying processes is blocked in this direction.

The distinguishing quality of QUAD has been its exceptional speed combined with flexibility in handling any option feature or combinations of features, and it is important while adding a new set of capabilities not to lose those previously developed. The universality of the method, promised in the title of the first paper (Andricopoulos, Widdicks, Duck, and Newton, 2003), was partly delivered by Andricopoulos, Widdicks, Newton, and Duck (2007) through extension to options involving simultaneously complex path dependency, early exercise features and multiple underlyings but remained severely limited in the range of underlying processes that could be handled. With the removal of this limitation, the exposition of the method is now complete. QUAD implementation is still straightforward (we provide code in an online appendix) and remains applicable to cases with complex combinations of features.

Gaining the capability to handle previously intractable underlying processes, while maintaining the full range of application previously established, involves returning to

the methods of Andricopoulos, Widdicks, Duck, and Newton (2003) and Andricopoulos, Widdicks, Newton, and Duck (2007) and finding techniques for the missing processes. At the core of this extension is the use of closed-form approximations for the appropriate single or two-variable transition density functions. Just as plain QUAD works with any one of several quadrature methods as calculational engines, any sufficiently accurate approximation to a transition density function could, in principle, be used and we illustrate the method via the approximations of Aït-Sahalia (2008) and Henry-Labordère (2009), the former proving the better all-purpose choice. In this ultimate version of QUAD it can be the case that, even when a density function is available, the approximation route is superior.

Appendix A. Heston probability density, convergence and extrapolation

As initial variance ν_0 is set progressively closer to zero, the estimate of the probability density function for the Heston model eventually blows up (the extreme tail of the density surface explodes toward infinity). This is a well-known property and a practical, if inelegant, solution is simply to set a lower bound for variance, so that the part of the density that has blown up is cut off in the lower tail of the density function. The cutoff value depends on how other parameters of the Heston model are set but is easily deduced using either C++ or Mathematica. Percentage errors of the order of 10^{-5} are found for $\nu_0=0.02$, for example, and so density estimates are sufficiently accurate. The only requirement for extra care would be with many observation points combined with particularly low ν_0 , to check compounded cutoff errors, but in practice this can be avoided.

Turning to convergence and extrapolation, the total error under stochastic volatility models consists of two parts: quadrature error and density estimation error. Quadrature error depends upon the scheme chosen for the valuation (Trapezium, Simpson, Gaussian, etc.), with the rate of convergence varying in some order of the quadrature step size. Density estimation error is produced by Aït-Sahalia's algorithm (or a substitute, such as Henry-Labordère's) and, although its accuracy increases as step size decreases, it does not have a uniform convergence rate and increasingly becomes the dominant source of error.

Initially, convergence is very fast but slows down as the error term becomes almost entirely dominated by estimation error of the density function. Here, estimation error is of second-order in Aït-Sahalia's algorithm. Fig. A1 shows quadrature step range 350–400 where quadrature error can be neglected compared with density estimation error. The density estimation error shows convergence, though that convergence is not smooth. In this case, extrapolation performs poorly. In other words, the error term can be extrapolated with bigger step sizes when quadrature error is still dominant but, as step size becomes smaller, quadrature error becomes insignificant and extrapolation no longer works. Fig. A2 illustrates.

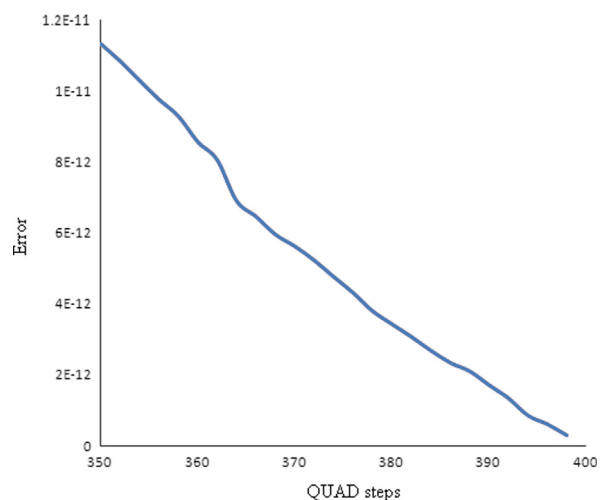


Fig. A1. Density estimation error (quadrature error is negligible in this range).

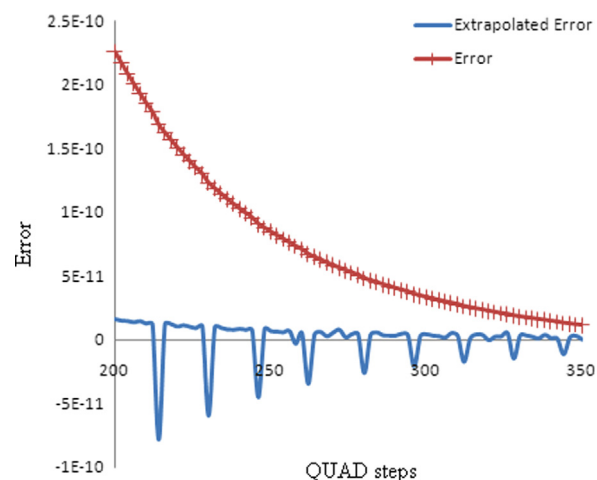


Fig. A2. Extrapolation error.

References

- Aït-Sahalia, Y., 1999. Transition densities for interest rate and other nonlinear diffusions. *Journal of Finance* 54, 1361–1395.
- Aït-Sahalia, Y., 2002. Maximum likelihood estimation of discretely sampled diffusions: a closed-form approximation approach. *Econometrica* 70, 223–262.
- Aït-Sahalia, Y., 2008. Closed-form likelihood expansions for multivariate diffusions. *Annals of Statistics* 36, 906–937.
- Aït-Sahalia, Y., Kimmel, R., 2007. Maximum likelihood estimation of stochastic volatility models. *Journal of Financial Economics* 83, 413–452.
- Andricopoulos, A.D., Widdicks, M., Duck, P.W., Newton, D.P., 2003. Universal option valuation using quadrature methods. *Journal of Financial Economics* 67, 447–471. (see also Corrigendum, *Journal of Financial Economics* 73, 603 (2004)).
- Andricopoulos, A.D., Widdicks, M., Newton, D.P., Duck, P.W., 2007. Extending quadrature methods to value multi-asset and complex path-dependent options. *Journal of Financial Economics* 83, 471–499.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–659.
- Boyle, P.P., 1977. Options: a Monte Carlo approach. *Journal of Financial Economics* 4, 323–338.

- Brennan, M.J., Schwartz, E.S., 1977. Convertible bonds: valuation and optimal strategies for call and conversion. *Journal of Finance* 32, 1699–1715.
- Carr, P., Madan, D., 1999. Option pricing and the fast Fourier transform. *Journal of Computational Finance* 2, 61–73.
- Chen, D., 2013. On Option Pricing and Equity Volatility Dynamics (Ph.D. dissertation). Nottingham University, Nottingham, UK.
- Cox, J.C., 1996. The constant elasticity of variance option pricing model. *Journal of Portfolio Management* 23, 15–17.
- Cox, J.C., Ingersoll, J.E., Ross, S.A., 1985. A theory of the term structure of interest rates. *Econometrica* 53, 385–408.
- Cox, J.C., Ross, S.A., Rubinstein, M., 1979. Option pricing: a simplified approach. *Journal of Financial Economics* 7, 229–264.
- Derman, E., Kani, I., 1994. Riding on a smile. *Risk* 7, 32–39.
- Dumas, B., Fleming, J., Whaley, R.E., 1998. Implied volatility functions: empirical tests. *Journal of Finance* 53, 2059–2106.
- Dupire, B., 1994. Pricing with a smile. *Risk* 7, 18–20.
- Fang, F., Oosterlee, C.W., 2008. A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing* 31, 826–848.
- Hagan, P.S., Kumar, D., Lesniewski, A., Woodward, D.E., 2002. Managing smile risk. *Wilmott Magazine* (September), 84–108.
- Heap, J., 2008. Enhanced Techniques for Complex Interest Rate Derivatives (Ph.D. dissertation). Manchester University, Manchester, UK.
- Henry-Labordère, P., 2009. *Analysis, Geometry, and Modeling in Finance: Advanced Methods in Option Pricing*. Chapman and Hall, London, UK.
- Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–343.
- Kristensen, D., Mele, A., 2011. Adding and subtracting Black-Scholes: a new approach to approximating derivative prices in continuous-time models. *Journal of Financial Economics* 102, 390–415.
- Longstaff, F.A., Schwartz, E.S., 2001. Valuing American options by simulation: a simple least squares approach. *Review of Financial Studies* 14, 113–147.
- Lord, R., Fang, F., Bervoets, F., Oosterlee, K., 2007. A fast and accurate FFT-based method for pricing early-exercise options under Levy processes. *SIAM Journal on Scientific Computing* 30; *Social Science Research Network* (<http://ssrn.com/abstract=966046>).
- Merton, R.C., 1973. Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4, 141–183.
- Merton, R.C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- O'Sullivan, C., 2005. Path dependent option pricing under Levy processes. In: *European Finance Association Moscow Meeting*, *Social Science Research Network* (<http://ssrn.com/abstract=673424>).
- Rebonato, R., 2004. *Volatility and Correlation: The Perfect Hedger and the Fox*, second ed. Wiley, West Sussex, UK.
- Rebonato, R., McKay, K., White, R., 2009. *The SABR/LIBOR Market Model: Pricing, Calibration and Hedging for Complex Interest Rate Derivatives*. Wiley, West Sussex, UK.
- Staunton, M., 2007. Convolution for Levy. *Wilmott Magazine* (September), 62–63.
- Vasicek, O., 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5, 177–188.
- Zhylyevskyy, O., 2010. A fast Fourier transform technique for pricing American options under stochastic volatility. *Review of Derivatives Research* 13, 1–24.
- Zhylyevskyy, O., 2012. Efficient pricing of European style options under Heston's stochastic volatility model. *Theoretical Economics Letters* 2, 16–20.