



TI 2006-066/2

Tinbergen Institute Discussion Paper

Optimal Fourier Inversion in Semi-analytical Option Pricing

Roger Lord¹

Christian Kahl²

¹ *Rabobank International, Utrecht, The Netherlands; Erasmus Universiteit Rotterdam; and Tinbergen Institute;*

² *University of Wuppertal, Wuppertal, Germany, and ABN AMRO, London, UK.*

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31

1018 WB Amsterdam

The Netherlands

Tel.: +31(0)20 551 3500

Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50

3062 PA Rotterdam

The Netherlands

Tel.: +31(0)10 408 8900

Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at
<http://www.tinbergen.nl>.

Optimal Fourier inversion in semi-analytical option pricing*

Roger Lord[†]

Christian Kahl[‡]

First version 9th March 2006

This version: 18th July 2006

Abstract

At the time of writing this article, Fourier inversion is the computational method of choice for a fast and accurate calculation of plain vanilla option prices in models with an analytically available characteristic function. Shifting the contour of integration along the complex plane allows for different representations of the inverse Fourier integral. In this article, we present the optimal contour of the Fourier integral, taking into account numerical issues such as cancellation and explosion of the characteristic function. This allows for robust and fast option pricing for almost all levels of strikes and maturities.

1 Introduction

In recent years the family of models which allow for semi-analytical solutions by the aid of Fourier inversion has been growing considerably. Let us therefore introduce φ as the characteristic function of the log-underlying S

$$\varphi(u) := \mathbb{E}[e^{iu \ln S_\tau}] . \quad (1)$$

Knowing the characteristic function allows us to express the forward price of a European call with strike K and maturity τ can be expressed very similarly to the Black-Scholes price as

$$C(S, K, \tau) = F\Pi_1 - K\Pi_2 , \quad (2)$$

with F being the forward value of the underlying and

$$\Pi_1 := \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left(\frac{e^{-iuk} \varphi(u-i)}{iu\varphi(-i)} \right) du = \mathbb{S}(S_\tau > K) , \quad (3)$$

where $k = \ln(K)$ and \mathbb{S} is the stock-price measure, i.e the measure induced by using $S(\tau)$ as the numeraire asset. Moreover we have

$$\Pi_2 := \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left(\frac{e^{-iuk} \varphi(u)}{iu} \right) du = \mathbb{P}(S_\tau > K) , \quad (4)$$

*The authors would like to express their gratitude to Peter Jäckel and Antoon Pelsser for helpful comments and fruitful discussions. The first author would like to thank ManWo Ng for an inspiring first look at related problems in his Master's thesis. Also, Jeroen Kerkhof is thanked for his comments on an early draft of this paper.

[†]Modelling and Research, Rabobank International, P.O. Box 17100, 3500 HG Utrecht, The Netherlands, and Erasmus University of Rotterdam, Econometric and Tinbergen Institute, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

[‡]Faculty of Mathematics and Science, Department of Mathematics, Chair of Applied Mathematics / Numerical Analysis, University of Wuppertal, Gaußstraße 20, Wuppertal, D-42119, Germany, and Quantitative Analytics Group, ABN AMRO, 250 Bishopsgate, London EC2M 4AA, UK

where \mathbb{P} is the martingale measure corresponding to the zero coupon bond expiring at τ . The probabilities are found by inverting the characteristic function, an approach which dates back to Lévy [Lév25]. His inversion theorem is restricted to cases where the random variable is strictly positive. Gurland [Gur48] and Gil-Pelaez [GP51] derived this more general inversion theorem. We refer the interested reader to Lukacs [Luk70] for a detailed account of the history of these approaches. The formulae (2)-(4) were first used in the context of option pricing by Heston [Hes93].

To allow for greater flexibility Carr and Madan [CM99] found an alternative representation for the European call price, derived by taking the Fourier transform of the damped option price $\hat{C}(\alpha) = C \exp(\alpha k)$ with respect to the log-strike price, and inverting it

$$C(S, K, \tau, \alpha) = \frac{e^{-\alpha k}}{\pi} \int_0^{\infty} \operatorname{Re} \left(\frac{e^{-ivk} \varphi(v - i(\alpha + 1))}{-(v - i\alpha)(v - i(\alpha + 1))} \right) dv. \quad (5)$$

The damping parameter α ensures that the damped call price is L^1 -integrable, a sufficient condition for its Fourier transform to exist. To simplify further notation we define

$$\psi(v, \alpha) := \operatorname{Re} \left(\frac{e^{-ivk} \varphi(v - i(\alpha + 1))}{-(v - i\alpha)(v - i(\alpha + 1))} \right). \quad (6)$$

The Fourier transform was taken with respect to the log-strike price in view of utilising the fast Fourier transform (FFT) to retrieve option values for a whole grid of strikes with just one evaluation of the FFT. The focus of Carr and Madan was purely on call options, which can be retrieved by using $\alpha > 0$. Though this approach was new to the area of option pricing, the idea of damping functions on the positive real line in order to be able to find their Fourier transform is an idea that goes back to at least Dubner and Abate [DA68].

Raible [Rai00] and Lewis [Lew01]¹ considered the same type of approach, except that here the transform was taken with respect to the log-forward and log-spot price respectively. Both authors demonstrate that the formula in (5) is quite general in that it can be adapted to a wide variety of European payoff functions, provided we can analytically calculate the Fourier transform of the damped payoff function. In Lewis [Lew01], however, an important step was made by considering the resulting integral as a contour integral in the complex plane. By shifting the contour (effectively changing α in (5)), various parity relations are obtained. As we will stick with the formula obtained by taking the Fourier transform of the log-strike price, the result of Lewis here comes down to:

$$C(S, K, \tau, \alpha) = R(S, K, \alpha) + \frac{1}{2\pi} \int_{-\infty - i\alpha}^{\infty - i\alpha} e^{-izk} \frac{\varphi(z - i)}{-z(z - i)} dz, \quad (7)$$

where the residue term equals

$$R(S, K, \alpha(k)) = F \cdot 1_{\{\alpha \leq 0\}} - K \cdot 1_{\{\alpha \leq -1\}} - \frac{1}{2} (F \cdot 1_{\{\alpha = 0\}} - K \cdot 1_{\{\alpha = -1\}}). \quad (8)$$

Carr and Madan already noticed in their seminal article that for strike prices far from the at-the-money level (ATM), and for short maturities, the integrand in (5) becomes highly oscillatory, and difficult to integrate numerically. As an alternative they considered taking a Fourier transform of out-of-the-money option prices, though this is also prone to numerical difficulties for short maturities. Andersen and Andreasen [AA02] suggested another approach to stabilise the numerical Fourier inversion. They used the Black-Scholes model as a control variate by subtracting the Black-Scholes characteristic function from the integrand and adding the Black-Scholes price. This approach would work perfectly well if the characteristic functions of both models would be close and if we would know an appropriate volatility level for the Black-Scholes model. Nonetheless, this approach may yield better results than using (7) with

¹We thank Ariel Almendral Vázquez for pointing out the equivalence of both approaches to us.

a default value of α . Lee [Lee04] intensively discussed Carr and Madan’s approach, and proposed an algorithm to arrive at an optimal α in the situation where the Fourier integral in (7) is approximated by the discrete Fourier transform (DFT). The algorithm consists of maximising the sum of the truncation error and the discretisation error with respect to the parameters of the discretisation, as well as α . Although this approach seems to work quite well for the examples Lee considered, it is quite specifically tailored towards the use of the DFT. Secondly, for most characteristic functions the estimated truncation error is typically a very conservative estimate of the true truncation error, which will certainly affect the resulting α .

The truncation error can actually be completely avoided by transforming the infinite integral to a finite domain using the limiting behaviour of the characteristic function as shown by Kahl and Jäckel [KJ05] for the Heston model. This reduces the sources of error from two to one, just leaving the discretisation error. Transforming the range of integration in this way precludes the use of the FFT algorithm. Therefore we are not able to use the power of the FFT, but this is not a real issue here. Firstly, using a different value of α for each strike/maturity pair would also already preclude the use of the FFT. Secondly, when calibrating a model to quoted option prices one typically has quotes for just a couple of strikes and maturities. Using the FFT would require a uniform grid in the log-strike direction in (2). The strikes of the options to which we calibrate will typically not lie on this grid, so that an additional source of error is introduced when using the FFT: interpolation error. Combined with the fact that the FFT binds us to a uniform grid usually makes it favourable to use a direct integration of (5) and (7). In terms of practical reliability and robustness it therefore seems more appropriate to follow the lead by Kahl and Jäckel [KJ05] and use an adaptive numerical integration scheme such as the adaptive Gauss-Lobatto scheme developed by Gander and Gautschi [GG00]. This does not necessarily minimise the overall computational workload, but it certainly ensures that the results are sufficiently accurate.

The outline of this article is as follows. In the next section 2, we summarise the characteristic functions of the models we want to discuss in the following. Their analytical features are analysed in order to be able to transform the integration domain to a finite interval. The very heart of this article follows, namely the appropriate choice of α in section 3. It is shown that the optimal choice of α is related to the contour shift used in saddle-point approximations. The penultimate section gives some numerical results underlining the previous results. Finally, we conclude.

2 Characteristic functions and integration domain transformation

In this section we will introduce the models we will investigate in the remainder of this paper. We will provide their characteristic functions and will analyse their limiting behaviour, in order to be able to transform the integration domain to a finite one. To keep things general we will analyse an affine jump-diffusion stochastic volatility model, as well as a model of the exponential Lévy class: the Variance Gamma (VG) model. The affine-jump diffusion model encompasses as special cases the models of Black-Scholes, Heston [Hes93], Stein and Stein [SS91] and Schöbel and Zhu [SZ99], as well as their respective extensions to include jumps in the asset price.

2.1 Affine diffusion stochastic volatility model

The affine diffusion stochastic volatility model we consider here is characterised by the following 2-dimensional system of stochastic differential equations

$$dS(t) = rS(t)dt + \hat{\eta} \cdot \sigma(t)^p S(t) dW_S(t), \quad (9)$$

$$d\sigma(t) = \kappa(\theta - \sigma(t))dt + \omega\sigma(t)^{1-p} dW_V(t) \quad (10)$$

with correlated Brownian motions $dW_S(t)dW_V(t) = \rho dt$. For $p = 0$ we have the standard Black-Scholes model, whilst for $p = 1/2$ we obtain the Heston stochastic volatility model. In a model with a single underlying asset the effect of $\hat{\eta}$ can be fully subsumed by the parameters of the variance process, so that it is safe to assume that $\hat{\eta} = 1$. If the same stochastic volatility driver is used for multiple assets,

$\hat{\eta}$ serves as a relative scaling to indicate how volatile each asset is compared to the other. An example of such a model is the stochastic volatility extension of the BGM/J market models due to Andersen and Andreasen [AA02]. To simplify the notation in the following we assume $\hat{\eta} = 1$ for $p \neq 0$, though of course all results remain valid if $\hat{\eta}$ is unequal to 1. Finally, for $p = 1$ the model is equivalent to the Schöbel-Zhu model, itself a generalisation of the Stein and Stein model to allow for non-zero correlation between the volatility and the spot price. Though the Schöbel-Zhu model is often referred to as being a non-affine stochastic volatility model, it is actually affine, as recently became evident from the study of quadratic models by Gaspar [Gas04] and Cheng and Scaillet [CS05]. The latter paper showed that a linear-quadratic jump-diffusion (LQJD) model is equivalent to an affine jump-diffusion (AJD) model with an augmented state vector, which in the case of the Schöbel-Zhu model is $\ln S(t)$, $\sigma(t)$ and $\sigma(t)^2$. All affine-diffusion models have in common that the characteristic function is exponentially affine in the state variables. In the model we consider here, we have:

$$\varphi_{\text{Affine}}(u) = e^{iuf + A(u, \tau) + B_\sigma(u, \tau)\sigma_0 + B_v(u, \tau)v_0}, \quad (11)$$

where $f = \ln F$, the logarithm of the forward price of the underlying asset. The functions A , B_σ and B_v satisfy the usual system of Riccati equations. In case of the Schöbel-Zhu model, we use the formulation of Lord and Kahl [LK06] to specify the characteristic function as:

$$A(u, \tau) = A_\sigma(\tau) + \frac{1}{4}(\beta - D)\tau - \frac{1}{2} \ln \left(\frac{Ge^{-D\tau} - 1}{G - 1} \right), \quad (12)$$

$$B_\sigma(u, \tau) = 2\kappa\theta \left(\frac{(\beta - D)(1 - e^{-\frac{1}{2}D\tau})^2}{D\gamma(1 - Ge^{-D\tau})} \right), \quad (13)$$

$$B_v(u, \tau) = \frac{(\beta - D)(1 - e^{-D\tau})}{2\gamma(1 - Ge^{-D\tau})}, \quad (14)$$

where we need some additional variables

$$\hat{\alpha} = -\frac{1}{2}u(u + i), \quad \beta = 2(\kappa - i\omega\rho u), \quad \gamma = 2\omega^2, \quad (15)$$

as well as

$$D = \sqrt{\beta^2 - 4\hat{\alpha}\gamma}, \quad G = \frac{\beta - D}{\beta + D}. \quad (16)$$

Finally some involved calculations lead to

$$A_\sigma(\tau) = \left(\frac{(\beta - D)\kappa^2\theta^2}{2D^3\omega^2} \right) \left(\beta(D\tau - 4) + D(D\tau - 2) + \frac{4e^{-\frac{1}{2}D\tau} \left(\frac{D^2 - 2\beta^2}{\beta + D} e^{-\frac{1}{2}D\tau} + 2\beta \right)}{1 - Ge^{-D\tau}} \right).$$

Transformation of the integration domain requires the asymptotics of the different components which are given in the following proposition.

Proposition 2.1 Assuming that $\kappa, \theta, \omega, \tau > 0$ and $\rho \in (-1, 1)$ we obtain the following asymptotics for the integrand of the Schöbel-Zhu model:

$$\lim_{u \rightarrow \infty} \psi(u, \alpha) \approx \psi(0, \alpha) \cdot e^{-uC_\infty} \cdot \text{Re} \left(\frac{e^{iut_\infty}}{-u^2} \right) = \psi(0, \alpha) \cdot e^{-uC_\infty} \cdot \frac{\cos(ut_\infty)}{-u^2} \quad (17)$$

with

$$C_\infty = D_\infty (\tau + V_0/\omega^2) / 4, \quad (18)$$

as well as

$$t_\infty = \beta_\infty (\tau + V_0/\omega^2) / 4 + \ln(F/K), \quad (19)$$

where the auxiliary variable is given by $D_\infty = 2\omega\sqrt{1 - \rho^2}$ and $\beta_\infty = -2\rho\omega$.

Proof: The proof can be found in the appendix A. □

Remark 2.1 Using the limiting behaviour of the Schöbel-Zhu model we can transform the integration domain using the transformation function $g(x) = -\frac{\ln x}{C_\infty}$.

The characteristic function of the Heston model is given by

$$A(\tau, u) = \frac{\kappa\theta}{\omega^2} \left((\beta - D)\tau - 2 \ln \left(\frac{Ge^{-D\tau} - 1}{G - 1} \right) \right), \quad (20)$$

$$B_v(\tau, u) = \frac{\beta - D}{\omega^2} \left(\frac{1 - e^{-D\tau}}{1 - Ge^{-D\tau}} \right), \quad (21)$$

where G and D are as before, with the exception that the auxiliary variables $\hat{\alpha}$, β and γ are redefined as:

$$\hat{\alpha} = -\frac{1}{2}u(u + i), \quad \beta = \kappa - i\omega\rho u, \quad \gamma = \frac{1}{2}\omega^2. \quad (22)$$

The asymptotics are due to Kahl and Jäckel [KJ05] [Prop. 3.1]:

Proposition 2.2 Assuming that $\kappa, \theta, \omega, \tau > 0$ and $\rho \in (-1, 1)$ we obtain the following asymptotics, for the integrand of the Heston model:

$$\lim_{u \rightarrow \infty} \psi(u, \alpha) \approx \psi(0, \alpha) \cdot e^{-uC_\infty} \cdot \operatorname{Re} \left(\frac{e^{iut_\infty}}{-u^2} \right) = \psi(0, \alpha) \cdot e^{-uC_\infty} \cdot \frac{\cos(ut_\infty)}{-u^2} \quad (23)$$

with

$$C_\infty = \left(\alpha d_\infty \tau + \frac{\sqrt{1 - \rho^2}}{\omega} V_0 \right) = \frac{\sqrt{1 - \rho^2}}{\omega} (V_0 + \kappa\theta\tau), \quad (24)$$

as well as

$$t_\infty = -\frac{\rho(\kappa\theta\tau + V_0)}{\omega} + \ln(F/K), \quad (25)$$

where $d_\infty = \omega\sqrt{1 - \rho^2}$.

Remark 2.2 The implementation of the characteristic function in the Heston and Schöbel-Zhu model requires the computation of the complex logarithm involved in $A(\tau, u)$ in equation (12) and equation (20). The complex logarithm is a multivalued function so that we need to choose the right branch to guarantee continuity. Kahl and Jäckel developed the so called *rotation-count-algorithm* [KJ05][Alg.1] to accomplish this task and Lord and Kahl [LK06] finally verified this algorithm for market relevant parameter configurations. In addition to this, it was found in [LK06] that the formulation used above is guaranteed to be continuous for the same parameter configurations if we restrict the complex logarithm to its principal branch.

We conclude this section with an analysis of the easiest affine diffusion: the Black-Scholes model. Its characteristic function is given by

$$\varphi(u) = \exp \left(iuf - \frac{1}{2}\hat{\eta}^2 u(u + i)\tau \right). \quad (26)$$

Clearly the characteristic function of the Black-Scholes model decays faster than the characteristic function of the stochastic volatility models.

Lemma 2.1 Assuming that the volatility $\hat{\eta} > 0$, then we obtain the following asymptotics, for the Black-Scholes model:

$$\lim_{u \rightarrow \infty} \psi(u, \alpha) \approx \psi(0, \alpha) \cdot e^{-u^2 C_\infty} \cdot \frac{\cos(ut_\infty)}{-u^2}, \quad (27)$$

with

$$C_\infty = \frac{\hat{\eta}^2 \tau}{2}, \quad (28)$$

as well as

$$t_\infty = \ln(F/K) + \hat{\eta}^2 \tau / 2. \quad (29)$$

Hence we can choose $g(x) = \sqrt{-\frac{\ln(x)}{C_\infty}}$ as the interval transforming function. From the tail behaviour of the stochastic volatility models it is already clear that the rate of decay of the characteristic function will increase when ω is decreased. In fact, for $\omega \rightarrow 0$ the affine diffusion model degenerates to the Black-Scholes model and we obtain:

Lemma 2.2 Let $\varphi(u)$ be the characteristic function of the Heston or Schöbel-Zhu model with limiting behaviour given by equation (18) of Prop. 2.1 and equation (24) of Prop. 2.2 then we obtain

$$\lim_{\omega \rightarrow 0} e^{-uC_\infty} \cdot \frac{\cos(ut_\infty)}{-u^2} = 0. \quad (30)$$

2.2 Affine jump-diffusion stochastic volatility model

To show how the analysis changes when jumps are included, we will add jumps to the underlying asset. The analysis is much the same when jumps are added to the stochastic volatility driver. When jumps are added the SDE in (9) changes to:

$$dS(t) = (r - \lambda\mu) S(t)dt + \hat{\eta} \cdot \sigma_t^p S(t)dW_S(t) + J_{N(t)} S(t)dN(t), \quad (31)$$

where N is a Poisson process independent of the Wiener process W_S with intensity parameter λ such that $E[N_t] = \lambda t$. The random variable J_i describes the size of the i^{th} jump. We will here only consider lognormally distributed jumps, though of course any jump size distribution could be used. With lognormal jump sizes

$$\ln(1 + J_t) \sim \mathcal{N}\left(\ln(1 + \mu) - \frac{\eta^2}{2}, \eta\right), \quad (32)$$

the model collapses to the Merton jump-diffusion model when $p = 0$, and to the Bates model [Bat96] when $p = 1/2$. As the Poisson process and the jump sizes are independent of the Brownian motions in the model, the characteristic function of the full model $\varphi_{\text{AJD}}(u)$ with jumps can be found as the product of the characteristic function of the affine diffusion part and $\varphi_{\text{Affine}}(u)$ and the characteristic function of the jump part $\varphi_{\text{Jump}}(u)$:

$$\varphi_{\text{AJD}}(u) = \varphi_{\text{Affine}}(u) \cdot \varphi_{\text{Jump}}(u). \quad (33)$$

For lognormal jump size $\varphi_{\text{Jump}}(u)$ is equal to

$$\varphi_{\text{Jump}}(u) = \exp\left(-\lambda\mu i u \tau + \lambda\tau \left((1 + \mu)^{iu} \exp(\eta^2(iu/2)(iu - 1)) - 1\right)\right), \quad (34)$$

It is straightforward to deduce that for $u = x + iy$, we have:

$$\lim_{x \rightarrow \infty} \varphi_{\text{Jump}}(x + iy) \approx \exp(-\lambda\mu i x \tau), \quad (35)$$

which does not influence the asymptotic behaviour of $\varphi_{\text{AJD}}(x + iy)$. Thus all affine jump-diffusion models considered allow us to transform the inverse Fourier integral to a finite domain which strongly simplifies the numerical computation of the semi-analytical option price.

2.3 Variance Gamma model

In the Variance Gamma model the stock price process is modelled as:

$$S(t) = F(t) \exp(\omega t + \theta G(t) + \sigma W(G(t))) \quad (36)$$

where $W(t)$ is a standard Brownian motion, and $G(t)$ is a Gamma process with parameter ν . The forward price of the underlying asset at time t is denoted by $F(t)$. Finally, the parameter ω ensures that the exponential term has a mean equal to 1:

$$\omega = \frac{1}{\nu} \ln \left(1 - \theta\nu - \frac{1}{2}\sigma^2\nu \right) \quad (37)$$

As before we let $f(t) = \ln F(t)$ and introduce $\tilde{f}(t) = f(t) + \omega t$. The conditional characteristic function is then specified as:

$$\varphi_{\text{VG}}(u) = \frac{\exp\left(iu\tilde{f}(T)\right)}{\left(1 - iu\left(\theta + \frac{1}{2}i\sigma^2u\right)\nu\right)^{\tau/\nu}} \quad (38)$$

In contrast to the affine jump-diffusion stochastic volatility models we analysed till thus far, this characteristic function decays only polynomially, also leading to a polynomial decay for the integrand in (5). For all transformation functions we tried, this polynomial decay, combined with the oscillatory nature of the characteristic function, causes the oscillations to bunch up at one end of the finite interval. Finding an appropriate transformation of the integration domain is therefore still an unsolved problem at the time of writing this article. Fortunately the easy analytical structure of (38) allows us to bound the characteristic function quite sharply, as has been done in [Lee05]. Combining this approach with a suitable choice of α will still reduce the numerical difficulties of the Fourier inversion significantly.

3 On the choice of alpha

Now we come to the very heart of this article - the optimal choice of α . Taking a closer look at the representation of the option price in (2) we recognise that when using this approach the range of option prices that can be calculated numerically is limited by cancellation errors. Indeed the crucial point for cancellation is the addition of $\frac{1}{2}(F - K)$ to the integral. Assuming that we use a highly sophisticated numerical integration scheme such as the suggested adaptive Gauss-Lobatto method, we can approximate the integrand up to a relative accuracy of 16 digits², which is the machine precision of a double number. Without loss of generality we set the forward value to $F(T) = 1$. Calculating out-of-the money (OTM) options with small times to maturity thus requires the subtraction of two values of almost the same size to obtain a much smaller option price. Cancellation is the consequence. Hence within the original Heston parametrisation we are not able to compute option prices below the machine size precision, independent of the integration scheme that is used. Clearly this problem is not specific to the Heston model when using equations (2)-(4) to recover option prices. Moreover, this problem will occur in almost every situation where a Fourier transform has to be numerically inverted.

The Carr-Madan representation, or in fact its more general formulation in (7), has several advantages. First of all, the number of numerical integrals is reduced from two to one. As an added benefit the denominator of the integrand is now a quadratic function in the integrating variable v , and as such decays faster than the integrands in (3)-(4). Finally, and most importantly for our purposes in this paper, the Carr-Madan representation (5) allows us to split the problem of tiny option prices from the problem of restricted machine size precision since $\exp(-\alpha k)$ serves as a scaling factor. An appropriate choice of α enables us to find a scaling which allows us to calculate arbitrarily small option prices. Unfortunately, the situation is not as simple as that, since by changing α , the integrand can become either

²The relative accuracy is provided by the macro `DBL_EPSILON` in the C header `<float.h>` which is the smallest positive number x such that $1 + x \neq 1$ in the computer's floating point number representation.

strongly peaked when getting close to the poles of the integrand, or highly oscillatory when reaching the maximum allowed α . It is therefore of substantial importance to have an appropriate choice of α .

Though the literature on mathematical finance has recognised the need for an appropriate choice of α , there has not been much work in this direction. Before discussing some recent studies, we mention a related article by Levendorskiĭ and Zherder [LZ02], which considers a different contour shift than the one considered here. Though they successfully apply their technique to the KoBoL/CGMY model of order smaller than 1, it turns out that their technique is not directly applicable to the stochastic volatility models considered here. We return to this in the beginning of section 3. As far as the damping factor α is concerned, Carr and Madan suggested to use one quarter of the maximal allowed α whilst Raible recommended that, for the models he considered, choosing α equal to 25 works best. Schoutens, Simons and Tistaert [SST04] suggested using an α equal to 0.75. Clearly, only the suggestion of Carr and Madan will work for any given model, as the option pricing formula in (5)-(7) is only well-defined when the $(\alpha + 1)$ -th moment exists. For the ad-hoc choices of $\alpha = 25$ or $\alpha = 0.75$ this need not be the case. The only rigorous study into an optimal choice of α has been performed by Lee [Lee04]. Lee suggests to minimise the sum of the discretisation error and the truncation error with respect to α and the quadrature parameters. In our approach, we transform the integration interval to a finite one, thus avoiding truncation error. Furthermore, by using adaptive quadrature we virtually avoid any discretisation error, so that there is nothing really left to minimise.

Nonetheless, choosing the right α can be very important, in particular for short maturities and/or strikes that are away from the at-the-money level. Figure 1 shows the relative error for different values of α . To generate this figure we used the adaptive Gauss-Lobatto scheme on the finite interval for different relative and absolute tolerance levels. It is obvious from Figure 1 that for certain contract




Figure 1: Relative pricing errors for the Heston model using the adaptive Gauss-Lobatto scheme for different absolute and relative tolerance levels over varying values of α . Underlying: $dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_S(t)$ with $S = F = 1$ and $\mu = 0$. Variance: $dV_t = \kappa(\theta - V_t)dt + \omega\sqrt{V_t}dW_V(t)$ with $V_0 = \theta = 0.1$, $\kappa = 1$, $\omega = 1$ and $\rho = -0.9$. (A): $\tau = 1/52$ and $K = 2$. Call value = $3.25E-126$, optimal $\alpha = 541.93$, (B) $\tau = 1/12$ and $K = 1.5$ Call value = $1.1802E-17$, optimal $\alpha = 121.24$.

parameters the outcome of the numerical integration is highly sensitive to the choice of α . Luckily there is a range of α 's which leads to the correct value, although the width of this range becomes increasingly smaller as the maturity of the option decreases and the strikes move away from the at-the-money level. Figure 2 compares the integrand of (5), i.e. $e^{-\alpha k} \psi(v, \alpha)$, in the Heston model for various values of α . The integration domain $[0, \infty)$ has been transformed to the unit interval using the transformation advocated in section 2. Furthermore, all integrands have been scaled by a constant such that their values in $x = 0$ coincide. Whilst choosing an α below the optimal one leads to oscillation at the left end of the interval the overestimation close to the maximum allowed α makes the characteristic function oscillatory at the other end. Vice versa the optimal choice of α leads to an integrand which is neither peaked nor

oscillating at all. Since the remaining shape of the optimal characteristic function differs depending on the parameter configuration, a standard saddle-point approximation would be too inaccurate. For that reason we suggest to estimate the optimal α and in addition use an adaptive quadrature scheme to obtain robust and accurate option prices. When possible, we also advocate transforming the integration domain to a finite one so that an analytical estimation of an appropriate upper limit of integration in (5) can be avoided.




Figure 2: Integrand in 5 for different values of α , where the transformation from (24) has been used. Underlying: $dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_S(t)$ with $S = F = 1$ and $\mu = 0$. Variance: $dV_t = \kappa(\theta - V_t)dt + \omega\sqrt{V_t}dW_V(t)$ with $V_0 = \theta = 0.1$, $\kappa = 1$, $\omega = 1$. (A): $\tau = 1/52$, $\rho = -0.9$ and $K = 2$. (B) $\tau = 1/360$, $\rho = -0.2$ and $K = 2.4$.

3.1 Minimum and maximum allowed alpha

This section examines how to determine the strip of regularity for the characteristic function, a problem which has already been investigated at great lengths in Andersen and Piterbarg [AP04]. Whilst Andersen and Piterbarg concentrate on the existence of ξ -th moment

$$\mu(\zeta, T) = \mathbb{E}[S(T)^\zeta] , \quad \text{for } \zeta > 1 , \quad (39)$$

we want to estimate the whole strip of regularity, thus choosing $\zeta \in \mathbb{R}$ such that $\mu(\zeta, T) < \infty$. Clearly this range will be of the form (ζ_-, ζ_+) . Let us therefore define:

$$\Lambda_x = \{u = x + iy \in \mathbb{C} \mid \zeta_- < -y < \zeta_+\} . \quad (40)$$

For all $u \in \Lambda_x$ the existence of the characteristic function is thus guaranteed, since:

$$|\varphi(u)| = |\mathbb{E}[e^{iu \ln S_\tau}]| \leq \mathbb{E}[|e^{iu \ln S_\tau}|] = \varphi(iy) . \quad (41)$$

The range (α_-, α_+) which contains all allowed α 's clearly corresponds to $(\zeta_- - 1, \zeta_+ - 1)$ as $\psi(v, \alpha)$ requires an evaluation of $\varphi(v - i(\alpha + 1))$, see (6). In the remaining subsections we give the strip of regularity for the models considered in section 2.

3.1.1 Affine diffusion stochastic volatility model

The crucial point for explosions in the characteristic function of an affine diffusion stochastic volatility model is the Riccati equation for $B_v(\tau, u)$ given by

$$\frac{\partial B_v}{\partial \tau} = \hat{\alpha}(u) - \beta(u)B_v + \gamma B_v^2 . \quad (42)$$

This equation holds for both the Heston and the Schöbel-Zhu model, with $\hat{\alpha}, \beta$ and γ defined in equation (15) in the Schöbel-Zhu model, whereas for the Heston model they are defined in (22). When considering the moment stability of the Schöbel-Zhu model, we can therefore immediately use the results of the Heston model, as long as we double κ and ω . Following the lead by Andersen and Piterbarg [AP04], who analysed moment explosions within the Heston model³, one can estimate the critical value of ζ with almost the same argumentation they use to determine the critical time to maturity. It can be shown that $\mu(\zeta, T)$ given by equation (39) is finite if $T < T^*$ and infinite for $T \geq T^*$, where T^* is given by one of three possibilities:

1. $D(-\zeta i)^2 \geq 0$, $\beta(-\zeta i) \geq 0$ or $\zeta \in [0, 1]$

$$T^* = \infty \quad (43)$$

2. $D(-\zeta i)^2 \geq 0$, $\beta(-\zeta i) < 0$

$$T^* = \frac{1}{c} \ln \left(\frac{\beta - c}{\beta + c} \right) \quad (44)$$

3. $D(-\zeta i)^2 < 0$

$$T^* = \frac{2}{c} \left(1_{\{\beta > 0\}} \pi + \arctan \left(-\frac{c}{\beta} \right) \right). \quad (45)$$

where β is shorthand notation for $\beta(-\zeta i)$ and we introduced:

$$c = \frac{1}{2} |D(-\zeta i)|. \quad (46)$$

As mentioned, Andersen and Piterbarg analysed the case where $\zeta > 1$. For $\zeta \in [0, 1]$ it is clear that all moments will be finite, as we can use the inequality $S(T)^\zeta < 1 + S(T)$, so that the moment will be bounded from above by the finite first moment, plus 1. It can be shown that for $\zeta < 0$ the results of Andersen and Piterbarg for $\zeta > 1$ still remain valid. The strip of regularity Λ_x can then be calculated by fixing the critical time $T^* = \tau$ and solving the nonlinear equation for α .

The efficiency of a numerical solution for a nonlinear equation strongly depends on a good approximation for the starting value. Here we only discuss the approximation of the starting values for some practically relevant parameter configurations such as $\rho \approx -1$ for an equity model and $\rho = 0$ in the interest rate case. The critical time in the notation of Andersen and Piterbarg is closely related to the sign of $D(-\zeta i)^2$ which is a second order polynomial in ζ

$$D(-\zeta i)^2 = \kappa^2 + \omega(\omega - 2\kappa\rho)\zeta - \omega^2(1 - \rho^2)\zeta^2, \quad (47)$$

with roots equal to

$$\zeta_{D\pm} = \frac{\omega - 2\kappa\rho \pm \sqrt{(\omega - 2\kappa\rho)^2 + 4(1 - \rho^2)\kappa^2}}{2\omega(1 - \rho^2)}. \quad (48)$$

Lemma 3.1 We have $\zeta_+ \geq \zeta_{D+} \geq 1$ and $\zeta_- \leq \zeta_{D-} < 0$.

Proof: For $\rho \leq 0$ the zeroes of $D(-\zeta i)^2$ are not the critical moments, as is evident from the analysis of Andersen and Piterbarg and Lord and Kahl [LK06]. Furthermore we have $D(-i)^2 = (\kappa - \omega\rho)^2 \geq 0$ and $D(0)^2 = \kappa^2 > 0$. \square

Proposition 3.1 The maximum allowed ζ for a highly negative correlation $\rho \approx -1$ can be approximated by

$$\zeta_+ \approx \frac{\omega - 2\kappa\rho \pm \sqrt{(\omega - 2\kappa\rho)^2 + 4(1 - \rho^2)(\kappa^2 + 4\pi^2\tau^{-2})}}{2\omega(1 - \rho^2)}. \quad (49)$$

³In fact, they analysed the moment stability of stochastic volatility models to be able to indicate when the values of certain contracts requiring the existence of higher moments are finite. Examples of such contracts are Eurodollar futures, LIBOR-in-arrears and CMS payments.

Proof: For $\rho \rightarrow -1$, ζ_{D_+} tends to infinity so that we can approximate (45) as $\tau = 2\pi/c$, leading to the analytical solution stated in equation (49). \square

As a final remark, when we have $\rho = 0$, as in an interest rate setting, symmetry is introduced into the problem: the critical time of the $(\zeta_{D_+} + z)^{\text{th}}$ moment for $z > 0$ is equal to the critical time of the $(\zeta_{D_-} - z)^{\text{th}}$ moment. Furthermore $\zeta_{D_+} + \zeta_{D_-} = 1$ here so that $\zeta_+ = 1 - \zeta_-$.

3.1.2 Affine jump-diffusion stochastic volatility model

When independent jumps are added to the underlying asset we have two possible sources of explosions: the affine part of the model requires $\varphi_{\text{Affine}}(x+iy)$ to be finite, whereas the jump part requires $\varphi_{\text{Jump}}(x+iy)$ to be finite. The second is only a practical problem since the jump component is always theoretically finite, though it may exceed the largest number representable on a finite computer system. Since we already discussed the affine part in the last section we concentrate on the latter here. As before we will focus on lognormally distributed jumps which are added to the underlying asset. For that reason we introduce

$$f_{\text{Jump}}(\zeta) = \ln(\varphi_{\text{Jump}}(-i\zeta)) = -\lambda\mu\tau\zeta + \lambda\tau \left((1 + \mu)^\zeta \exp(\eta^2 (\zeta/2) (\zeta - 1)) - 1 \right). \quad (50)$$

Our aim is to find a range for ζ such that the characteristic function can be calculated using standard double precision. We therefore impose an upper bound on (50) equal to $\frac{1}{4} \ln(\text{DBL_MAX}) \approx 177 = d_{\text{Max}}$. The motivation for this choice is that we still want to leave

τ	$\hat{\zeta}_{\text{Max}}$	$\zeta_{\text{Jump,Max}}$	steps	$\hat{\zeta}_{\text{Min}}$	$\zeta_{\text{Jump,Min}}$	steps
1/52	39.6266	39.6265	2	-57.6884	-57.6886	2
1	30.6901	30.6857	3	-48.7408	-48.7478	3
10	24.445	24.4042	3	-42.3941	-42.4663	3

Table 1: Jump diffusion: $\mu = \lambda = \eta = 0.1$. Number of Newton-Raphson iterations steps to obtain an accuracy of 10^{-8} .

enough room for the remaining terms of the characteristic function. In view of the optimal choice of α discussed in section 3 this upper bound is large enough to guarantee that the critical α of the jump part is greater than the optimal α . The next step is to calculate the critical values of ζ . This requires the solution of a nonlinear equation which can be done by using the standard Newton-Raphson method. To find a good starting point, let us define the function g_{Jump} as a slight modification of f_{Jump}

$$g_{\text{Jump}}(\zeta) = \ln(\varphi_{\text{Jump}}(-i\zeta)) + \lambda\mu\tau\zeta = \lambda\tau \left((1 + \mu)^\zeta \exp(\eta^2 (\zeta/2) (\zeta - 1)) - 1 \right). \quad (51)$$

This enables us to compute solution of the nonlinear equation $g_{\text{Jump,Max}}(\hat{\zeta}) = d_{\text{Max}}$ analytically

$$\hat{\zeta}_{\text{Max}} = \frac{\eta^2 - 2 \ln(\mu + 1) + \sqrt{8 \ln\left(\frac{d_{\text{Max}} + \lambda\tau}{\lambda\tau}\right) \eta^2 + (2 \ln(\mu + 1) - \eta^2)^2}}{2\eta^2}, \quad (52)$$

$$\hat{\zeta}_{\text{Min}} = \frac{\eta^2 - 2 \ln(\mu + 1) - \sqrt{8 \ln\left(\frac{d_{\text{Max}} + \lambda\tau}{\lambda\tau}\right) \eta^2 + (2 \ln(\mu + 1) - \eta^2)^2}}{2\eta^2}. \quad (53)$$

Now one can use this ζ as a starting point for the Newton-Raphson method on $f_{\text{Jump}}(\zeta_{\text{Jump}}) = d_{\text{Max}}$. Table 1 shows that the initial guess $\hat{\zeta}$ is a good approximation for the solution of the nonlinear equation. Concluding, although the strip of regularity of the affine jump-diffusion stochastic volatility model is equivalent to that of the Heston model, numerical explosions of the characteristic function suggest that we should restrict the strip of regularity (ζ_-, ζ_+) to a strip of numerical regularity:

$$\zeta_{\text{Min}} = \text{Max}(\zeta_{\text{Jump,Min}}, \zeta_-), \quad \zeta_{\text{Max}} = \text{Min}(\zeta_{\text{Jump,Max}}, \zeta_+). \quad (54)$$

Though we have focused on lognormally distributed jumps here, the same caveats apply to models with other jump size distributions.

3.1.3 Variance Gamma model

Due to the comparative simplicity of the characteristic function for the VG model given by equation (38), we can straightforwardly deduce the strip of regularity (ζ_-, ζ_+) as

$$\zeta_{\pm} = -\frac{\theta}{\sigma^2} \pm \sqrt{\frac{\theta^2}{\sigma^4} + \frac{2}{\nu\sigma^2}}. \quad (55)$$

Since the variance gamma model is fully time-homogeneous, the maximum and minimum allowed ζ , and therefore also α , do not depend on the time to maturity, contrary to what we find in the affine jump-diffusion stochastic volatility models.

3.2 Optimal alpha

Before we embark upon our quest for the optimal α , we return to the method considered by Levendorskiĭ and Zherder [LZ02], which we mentioned in the beginning of this section. Their integration-along-cut (IAC) method is specifically tailored towards KoBoL/CGMY processes of order smaller than 1. For these processes it turns out that the characteristic function $\varphi(u)$ is analytic for $u \in \mathbb{C}$, with cuts $[i\zeta_+, i\infty)$ and $(-i\infty, i\zeta_-]$. In words, this means that the characteristic function is well-defined on the whole complex plane, apart from being discontinuous along the mentioned cuts. Instead of integrating along $(-\infty - i\alpha, \infty - i\alpha)$ in the complex plane to obtain the option price via equation (7), they integrate along one of the two aforementioned cuts. Unfortunately this technique is not generally applicable to any process. For example, it is shown in Lord and Kahl [LK06] that for the stochastic volatility models we consider here, the characteristic function has an infinite number of singularities along the imaginary axis. As we are seeking a method that is fully general and applicable to any model, we will therefore not consider this method any further in this article.

The literature on option pricing has recognised the need for a good choice of α . Choosing α too small or too big leads to problems with either cancellation errors or highly oscillating integrands as shown in figures 1 and 2. The best α ensures that the absolute value of the integrand is as constant as possible over the whole integration area. The oscillation of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ on the finite interval $[-1, 1]$ can be measured by the total variation

$$\text{TV}(f) = \int_{-1}^1 \left| \frac{\partial f}{\partial x}(x) \right| dx. \quad (56)$$

As shown by Förster and Petras [FP91] a small total variation reduces the approximation error of the applied numerical integration scheme. Ideally one should choose α such that the total variation of the integrand is minimised⁴

$$\alpha^* = \underset{\alpha \in \{\alpha_{\text{Min}}, \alpha_{\text{Max}}\}}{\text{argmin}} e^{-\alpha k} \int_0^{\infty} \left| \frac{\partial \psi(v, \alpha)}{\partial v \psi(0, \alpha)} \right| dv. \quad (57)$$

Clearly this optimization problem is of a rather theoretical nature as it has to be solved again for each option price and would therefore require many more function evaluations than the original problem. A related metric can be found by considering Figure 1, whence it is clear that in the optimal range of α 's the derivative of the error of the numerical scheme with respect to α will roughly be equal to zero. If we choose the squared difference between the true price and the value resulting from our numerical scheme as a relevant criterion, the first order condition for optimality becomes:

$$\alpha^* = \underset{\alpha \in \{\alpha_{\text{Min}}, \alpha_{\text{Max}}\}}{\text{argmin}} \frac{\partial}{\partial \alpha} e^{-\alpha k} \int_0^{\infty} \psi(v, \alpha) dv. \quad (58)$$

⁴Note that one should rescale the integrand to make the total variation comparable for different values of α .

where the integral on the right-hand side should be interpreted as the numerical approximation of this integral. A simple but highly effective idea to simplify these optimisation problems is to minimise only the maximal value of the integrand, which occurs in $v = 0$. We therefore suggest to choose α according to

$$\alpha^* = \operatorname{argmin}_{\alpha \in \{\alpha_{\text{Min}}, \alpha_{\text{Max}}\}} |e^{-\alpha k} \psi(-(\alpha + 1)i)| . \quad (59)$$

or equivalently

$$\alpha^* = \operatorname{argmin}_{\alpha \in \{\alpha_{\text{Min}}, \alpha_{\text{Max}}\}} \left[-\alpha k + \frac{1}{2} \ln(\psi(-(\alpha + 1)i)^2) \right] =: \Psi(\alpha, k), \quad (60)$$

rendering the calculation more stable since the function ψ is not necessarily positive. Finding the optimal α thus requires to find the minimum of Ψ . The optimal α resulting from this optimisation problem will be referred to as the payoff-dependent α , as ψ depends on the particular payoff function we are considering. In the following we will also consider a payoff-independent alternative:

$$\alpha^* = \operatorname{argmin}_{\alpha \in \{\alpha_{\text{Min}}, \alpha_{\text{Max}}\}} [-\alpha k + \ln(\varphi(-(\alpha + 1)i))] =: \Phi(\alpha, k) , \quad (61)$$

Coincidentally, this payoff-independent way of choosing α has a close link to how the optimal contour is chosen in the area of saddlepoint approximations, something we discuss in the following section. A related article by Choudhury and Whitt [CW97] considers the contour shift following from (60) to avoid numerical problems in the numerical inversion of non-probability transforms. To the best of our knowledge this choice of α has not been picked up in the area of option pricing.

Let us for a moment consider whether solving the optimal α from (60) is a well-posed problem. From the theory of saddlepoint approximations (see [Dan54]) we know that the minimum of the damped characteristic exponent in (61) is unique under relatively mild conditions on the first derivative of the characteristic exponent (cf. equation (6.1) in [Dan54]), which are satisfied here. This analysis leads us to conclude that the payoff-dependent function Ψ will have a local minimum in three ranges: $\alpha \in (\alpha_{\text{Min}}, -1)$, $\alpha \in (-1, 0)$ and $\alpha \in (0, \alpha_{\text{Max}})$, since Ψ is infinite when $\alpha = -1$ or 0 . To find the optimal payoff-dependent α^* we have to solve the nonlinear equation:

$$\frac{\partial \Psi(\alpha, k)}{\partial \alpha} = 0 . \quad (62)$$

Its typical shape can be found in figure 3.




Figure 3: Function Ψ given by equation (60) and its first derivative for the Heston model w.r.t. α . Underlying: $dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_S(t)$ with $S = F = 1$ and $\mu = 0$. Variance: $dV_t = \kappa(\theta - V_t)dt + \omega\sqrt{V_t}dW_V(t)$ with $V_0 = \theta = 0.1$, $\kappa = 1$, $\omega = 1$ and $\rho = -0.7$. (A): $\tau = 1$ and $K = 1.2$. (B): $\tau = 1/2$ and $K = 0.7$.

Figure 3 indicates that for $F < K$, $\alpha^* \in (0, \alpha_{\text{Max}})$ and for $F > K$ we will have $\alpha^* \in (\alpha_{\text{Min}}, -1)$. Though this is no rigorous proof, more often than not we find that this rule of thumb holds true. This allows us to restrict the search for the optimal α to just one of the three aforementioned ranges, if computational time is of the essence. In words the rule of thumb states that it is always preferable to price out-of-the-money options. The following two subsections explore the payoff-independent choice of α for two models in which it can be calculated analytically: the the Black-Scholes model and the VG model.

3.2.1 Black-Scholes model

As the characteristic function of the Black-Scholes formula has a convenient analytical structure, the optimal payoff-independent α we propose in equation (61) can be calculated analytically. We find:

$$\alpha^* = \operatorname{argmin}_{\alpha \in \mathbb{R}} \Phi(\alpha, k) = -\frac{f - k + \frac{1}{2}\widehat{\eta}^2\tau}{\widehat{\eta}^2\tau} = -\frac{d_1}{\widehat{\eta}\sqrt{\tau}}, \quad (63)$$

where d_1 is a well-known part of the Black-Scholes option pricing formula. Note that the payoff-independent α also roughly obeys the rule of thumb we stated earlier. The payoff-dependent α^* can also be solved in closed-form, though this does not lead to nice analytical expressions as (62) becomes a fourth-order polynomial in α .

There is one more thing to be said about the payoff-independent α^* in the Black-Scholes model. Suppose we write:

$$\varphi(v - (\alpha + 1)i) = \exp(\varphi_r(v, \alpha) + i\varphi_i(v, \alpha)), \quad (64)$$

with both φ_r and φ_i being real-valued functions. The Carr-Madan representation of the option price in (5) can then be written as:

$$C(\alpha) = \frac{e^{-\alpha k}}{\pi} \int_0^\infty e^{\varphi_r(v, \alpha)} \frac{(\alpha(\alpha + 1) - v^2) \cos(kv - \varphi_i(v, \alpha)) - (1 + 2\alpha)v \sin(kv - \varphi_i(v, \alpha))}{(v - i(\alpha + 1))(v + i(\alpha + 1))(v - i\alpha)(v + i\alpha)} dv,$$

which is a real-valued integral. So far this discussion has been model-independent. Upon inspection it seems logical to choose α such that all oscillations due to the sine and the cosine are removed completely, by setting:

$$kv - \varphi_i(v, \alpha) = 0. \quad (65)$$

Typically, this equation will not have a unique solution, though it does in the Black-Scholes case. The solution for α in (65) is exactly the payoff-independent α^* we derived earlier in (63).

3.2.2 Variance Gamma model

The Variance Gamma model and the Black-Scholes model are two of the few models in which the optimal payoff-independent α can be found analytically. Standard yet tedious calculations yield:

$$\alpha^* = \operatorname{argmin}_{\alpha \in \mathbb{R}} \Phi(\alpha, k) = -\frac{\theta}{\sigma^2} - 1 + \frac{\tau}{\nu\widetilde{m}} - \operatorname{sgn}(\widetilde{m}) \sqrt{\frac{\theta^2}{\sigma^4} + \frac{2}{\nu\sigma^2} + \frac{\tau^2}{\nu^2\widetilde{m}^2}}, \quad (66)$$

where we introduced $\widetilde{m} = \widetilde{f} - k = f - k + \omega\tau$, a quantity related to the log-moneyness of the option. One can easily check that $\alpha^* \in (\alpha_-, \alpha_+)$, as should be the case. We mention that Ait-Sahalia and Yu [ASY06] provide the saddlepoint of the VG model and several other models where it can be calculated analytically. As we will show in the next section, the saddlepoint, minus one, coincides exactly with our payoff-independent α^* .

3.3 Saddlepoint approximations

The concept of saddlepoint approximations dates back to Daniels [Dan54] in the context of calculating probability densities via Fourier inversion

$$p(x) = \frac{1}{2\pi} \int_{-\infty-i\alpha}^{\infty-i\alpha} e^{-izx} \varphi_p(z) dz, \quad (67)$$

where φ_p is the characteristic function associated with the density p . Though the standard inversion formula would set α equal to zero, we can shift the contour of integration by choosing α different from zero. Saddlepoint approximations typically continue along the following lines. Let $M(z) = \ln \varphi_p(z)$, the characteristic exponent, then Daniels advocated to choose α as the minimum of the damped characteristic exponent (61):

$$\alpha^* = \underset{\alpha \in \{\alpha_{\text{Min}}, \alpha_{\text{Max}}\}}{\text{argmin}} -\alpha x + M(-i\alpha). \quad (68)$$

Thus we obtain $M'(-i\alpha^*) = ix$. Note that the payoff-independent α coincides with the saddlepoint minus one. Applying a Taylor expansion around its minimum leads to

$$M(z) - izx = M(-i\alpha^*) - \alpha^* x + \frac{1}{2} M''(-i\alpha^*) (z + i\alpha^*)^2 + \mathcal{O}(z^3), \quad (69)$$

so that the density can be approximated via

$$p(x) = \frac{1}{2\pi} \int_{-\infty-i\alpha^*}^{\infty-i\alpha^*} e^{-izx} \varphi_p(z) dz \approx \frac{\varphi_p(-i\alpha^*) e^{-\alpha^* x}}{2\pi} \int_{-\infty-i\alpha^*}^{\infty-i\alpha^*} e^{\frac{1}{2} M''(-i\alpha^*) z^2} dz \quad (70)$$

$$= \frac{\varphi_p(-i\alpha^*) e^{-\alpha^* x}}{\sqrt{-2\pi M''(-i\alpha^*)}} =: p_{\text{simple}}(x). \quad (71)$$

Instead of numerically evaluating an integral over an infinite domain, this "simple" saddlepoint approximation requires only a few function evaluations⁵. As Ait-Sahalia and Yu [ASY06] mention, the name "saddlepoint" stems from the shape of the right-hand side of (69) in a neighbourhood of its minimum, which can be seen as a saddle. Based on this saddlepoint approximation to the density function, an approximation can be derived for the cumulative density function. The most famous of these is the Lugannani-Rice formula, see Lugannani and Rice [LR80].

All saddlepoint approximations are found to work remarkably well in the tails of the distribution. As such, saddlepoint approximations in finance have mainly been used for Value at Risk and expected loss calculations. Though option pricing mainly deals with the bulk of the distribution, Rogers and Zane [RZ98] applied the Lugannani-Rice formula to compute the probabilities Π_1 (3) and Π_2 (4). As expected they obtained accurate results for options close to maturity and away from the at-the-money level, though the accuracy was lower around the at-the-money level. For small option prices this approach will lead to cancellation errors, so that we suggest to apply the saddlepoint approximation directly to the Carr-Madan representation (5). With abuse of notation, let us define $\psi(z) = \psi(v, \alpha)$ when $z = v - i\alpha$ for $v, \alpha \in \mathbb{R}$, i.e. $\psi(z) = \varphi^{(z-i)}/(-z(z-i))$. Instead of expanding the characteristic exponent around its minimum, M here has to be taken equal to:

$$M(z) = \frac{1}{2} \ln \psi(z)^2, \quad (72)$$

leading to the "simple" saddlepoint approximation of the option price

$$C(S, K, \tau, \alpha^*) = R(S, K, \alpha^*) + \frac{1}{2\pi} \int_{-\infty-i\alpha^*}^{\infty-i\alpha^*} e^{-izx} \psi(z) dz \approx \frac{\psi(-i\alpha^*) e^{-\alpha^* x}}{\sqrt{-2\pi M''(-i\alpha^*)}} =: C_{\text{simple}}(x), \quad (73)$$

⁵The number of function evaluations depends on the estimation of the second order derivative.

Tolerance	Heston			optimal- α		
	avg. FE	max Error	avg. Error	avg. FE	max Error	avg. Error
Min.	24	173.63%	4.8071%	12	3.17%	0.048%
10^{-3}	212.9	32.91%	0.2724%	73.9	0.47%	0.0038%
10^{-5}	540.4	2.5%	0.0044%	179.8	0.03%	0.0001%
10^{-7}	1097.4	0.09%	9.7E-5%	373.5	0.0001%	7.7E-7%

Table 2: Average number of function evaluations (FE) and error measured in implied volatility for the Heston model calculated over a whole range of strikes and maturities. Underlying: $dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_S(t)$ with $S = F = 1$ and $\mu = 0$. Variance: $dV_t = \kappa(\theta - V_t)dt + \omega\sqrt{V_t}dW_V(t)$ with $V_0 = \theta = 0.16$, $\kappa = 1$, $\omega = 2$, $\rho = -0.8$, $\tau \in [1, 5/4, 6/4, \dots, 15]$, and $K \in [1/10, 2/10, \dots, 4]$.

where α^* here is the payoff-dependent optimal α .

Naturally, higher-order expansions of the density can also be derived. Aït-Sahalia and Yu give a higher-order expansion derived by expanding the left-hand side of (69) to fourth order in z :

$$p(x) = p_{\text{simple}}(x) \left(1 + \frac{1}{8} \frac{M^{(4)}(-i\alpha^*)}{(M^{(2)}(-i\alpha^*))^2} - \frac{5}{24} \frac{(M^{(3)}(-i\alpha^*))^2}{(M^{(2)}(-i\alpha^*))^3} \right) =: p_{\text{sophisticated}}(x). \quad (74)$$

We will refer to the analogue of (74) in the context of option pricing as the "sophisticated" saddlepoint approximation of the option price. The numerical effort for this approximation is increased considerably since we have to approximate the fourth order derivative numerically⁶.

As a final note, we mention that the saddlepoint approximations considered here can be seen as an expansion around a Gaussian base (the term $\frac{1}{2}z^2$ in (69)). Aït-Sahalia and Yu show how to derive saddlepoint approximations which expand the characteristic exponent around a non-Gaussian base, and demonstrate that this improves the accuracy of the approximations considerably for jump-diffusion and Lévy models. Though this route is viable, the problem of determining which base is most suitable for a given model is obviously a very model-dependent problem. As we would like to keep our pricing methods as model-independent as possible, the only model-dependent part being of course the specification of the characteristic function and a transformation function⁷, we only consider the saddlepoint approximations around a Gaussian base here.

4 Numerical results

In this final section we present some numerical tests comparing the different numerical methods for semi-analytical option pricing with regard to their approximation quality and the computational effort. In table 2 we compare the number of function evaluations and the computational error for the standard Heston formulation (cf. (2)) to the optimal- α method (cf. (7)) using the optimal payoff-dependent α from (60)).

One can recognize that the optimal choice of α decreases both the necessary number of function evaluations and the computational error. The parameter configuration was taken from Kahl and Jäckel [KJ05][Fig.7] and is such that the standard Heston implementation leads to a smooth implied volatility surface when using a relative tolerance of 10^{-12} in the adaptive Gauss-Lobatto scheme. Another great advantage is that the optimal choice of α allows to calculate option prices for all ranges of strikes and maturities. Figure 4 shows the implied volatility surface for the Heston model calculated with the optimal α . This surface is perfectly smooth. In the upper right corner we have options which are close to maturity and decisively out-of-the-money. In table 3 we present the corresponding option prices. Using the optimal α in conjunction with the adaptive Gauss-Lobatto scheme, and transforming the integration domain to a finite interval, allows us to compute option prices down to the lowest number⁸

⁶As we have closed-form expressions for all characteristic functions considered in this article, one could theoretically derive all the required derivatives in closed-form.

⁷In case the infinite domain of integration cannot be transformed to a finite one we require an appropriate upper limit of integration.

⁸This information can be found again in the C header `<float.h>` in the macro `DBL_MIN`.




Figure 4: Black implied volatilities from the Heston model. Underlying: $dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_S(t)$ with $S = F = 1$ and $\mu = 0$. Variance: $dV_t = \kappa(\theta - V_t)dt + \omega\sqrt{V_t}dW_V(t)$ with $V_0 = \theta = 0.1$, $\kappa = 1$, $\omega = 1$, $\rho = -0.7$, $\tau \in [1/52, 2/52, \dots, 3/2]$, and $K \in [1/10, 2/10, \dots, 10]$.

$\tau - K$	9.5	9.6	9.7	9.8	9.9	10.0
$1/52$	6.4232E-260	2.6773E-261	1.1522E-262	5.1158E-264	2.3423E-265	1.1052E-266
$2/52$	3.4710E-133	6.9920E-134	1.4313E-134	2.9768E-135	6.2873E-136	1.3483E-136
$3/52$	76979E-91	26221E-91	9.0293E-92	3.1424E-92	1.1051E-92	3.9263E-93
$4/52$	1.2869E-69	5.7020E-70	2.5472E-70	1.1471E-70	5.2069E-71	2.3818E-71

Table 3: Option prices within the Heston model where the parameter configuration is the same as in the figure 4.

representable in floating point precision. Moreover this method can be used as a Black-Box algorithm to semi-analytical option pricing since we avoid numerical instabilities for small option prices.

As a final assessment of how well our method performs, we will compare the optimal- α method to the saddlepoint approximations discussed in section 3.3. In figure 5 we present the error made in




Figure 5: Error in the implied volatility surface for the saddlepoint approximation in the Heston model. Underlying: $dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_S(t)$ with $S = F = 1$ and $\mu = 0$. Variance: $dV_t = \kappa(\theta - V_t)dt + \omega\sqrt{V_t}dW_V(t)$ with $V_0 = \theta = 0.1$, $\kappa = 1$, $\omega = 1$ and $\rho = -0.5$. (A): Simple saddlepoint-approximation (70). (B): Sophisticated saddlepoint-approximation (74).

the implied volatility surface using the saddlepoint approximation to approximate the inverse Fourier integral (7). As mentioned these saddlepoint approximations are different to the Lugannani-Rice based approximation Rogers and Zane [RZ98] considered in their article, though we find their accuracy to be comparable to ours. For this reason we focus on the simple and sophisticated saddlepoint approximation we suggested in section 3.3. Both saddlepoint approximations work extremely well in the tails of the probability density but fail for strikes that are around the at-the-money level. The method is most efficient for small times to maturity. Comparing the different saddlepoint methods we see that including higher order terms of the Taylor expansion certainly leads to an improved accuracy, at the cost of having to compute 4-th order derivatives numerically.

As a fair comparison of the saddlepoint methods to the optimal α method we finally present the error in the implied volatility surface for the optimal α method, using a maximum of 12 function evaluations per option price. This is more than competitive to the sophisticated saddlepoint approximation, as just using central-difference formulae of order $O(h^2)$, with h being the stepsize used for the numerical differentiation, would already require 5 function evaluations. Calculating the derivatives numerically with two different values of h and subsequently applying Richardson extrapolation would require 10 function evaluations. Clearly, using a more sophisticated scheme to calculate the required derivatives numerically will require more than 12 function evaluations. We find that the stripped-down optimal α method produces an error of less than half a percent over the whole surface, which is more than sufficient for all practical purposes. Furthermore the stripped-down optimal α method only seems to produce marginally worse results than the sophisticated saddlepoint method for options with low strikes and high maturities. From a computational point of view the optimal α method should therefore certainly be preferred, as numerical integration of a function is a much more stable operation than the numerical differentiation required in the calculation of the saddlepoint approximations.

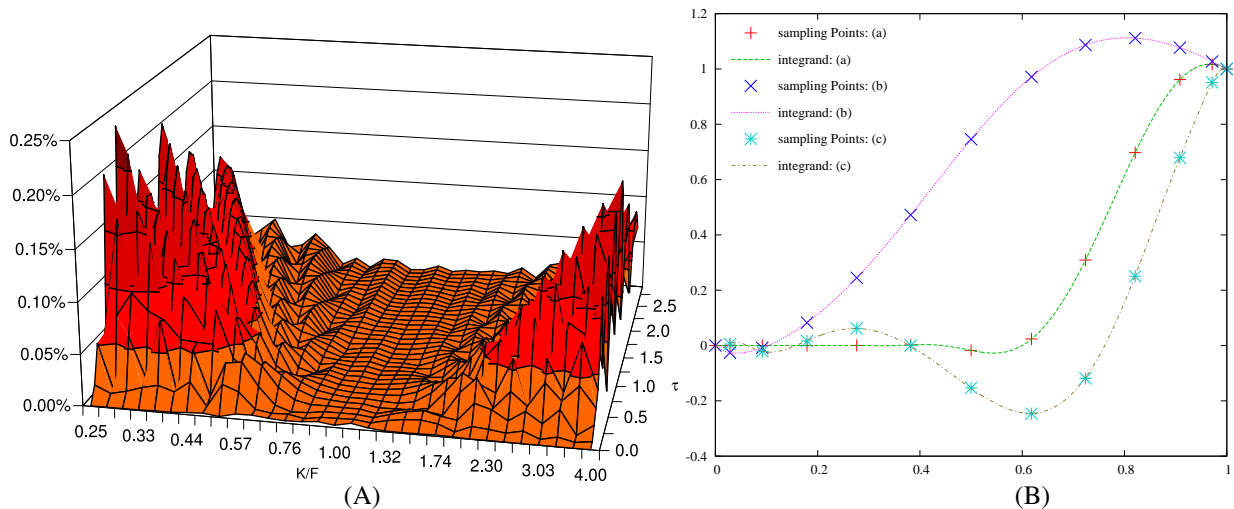


Figure 6: (A): Error in the implied volatility surface for the optimal α method in the Heston model using 12 function evaluations in the adaptive Gauss-Lobatto method for each price. Underlying: $dS_t = \mu S_t dt + \sqrt{V_t} S_t dW_S(t)$ with $S = F = 1$ and $\mu = 0$. Variance: $dV_t = \kappa(\theta - V_t)dt + \omega\sqrt{V_t}dW_V(t)$ with $V_0 = \theta = 0.1$, $\kappa = 1$, $\omega = 1$ and $\rho = -0.5$. (B): Sampling points for the adaptive Gauss-Lobatto quadrature scheme for (a) $K = 1$ and $\tau = 1/2$, (b) $K = 4$ and $\tau = 1/2$ and (c) $K = 0.4$ and $\tau = 3$.

5 Conclusion

In this article we investigated the problem of handling the Fourier inversion required in semi-analytical option pricing as best as possible. Using the Carr-Madan representation of the option price and shifting the contour of integration along the complex plane allows for different representations of the European call price. The key point is to choose the contour of integration best suited for numerical quadrature, which can be changed by varying the damping parameter α . Changing the contour of integration alters the behaviour of the integrand considerably, so that we have to ensure that the integrand does not become

too peaked or too oscillatory. Furthermore, cancellation errors also have to be avoided at all cost. By transforming the integration domain as discussed in section 2, as well as choosing the optimal α as in section 3, we obtain a call price formula which allows for a robust implementation enabling us to price options down to machine size precision. The optimal choice of the damping parameter α is the only way to overcome any numerical instabilities and guarantee accurate results.

A Proofs

In this section we present the missing proofs. **Proof:** [Prop. 2.1] One can successively deduce that

$$\begin{aligned}\lim_{u \rightarrow \infty} \frac{\beta(u)}{u} &= -2i\rho\omega =: i \cdot \beta_\infty \\ \lim_{u \rightarrow \infty} \frac{D(u)}{u} &= 2\omega\sqrt{1-\rho^2} =: D_\infty \\ \lim_{u \rightarrow \infty} G(u) &= \text{const.}\end{aligned}$$

Using the limiting behaviour of D we can directly deduce that

$$\lim_{u \rightarrow \infty} A_\sigma(u) = 0 ,$$

as well as

$$\lim_{u \rightarrow \infty} B_\sigma(u) = 0 ,$$

which finally leads to

$$\begin{aligned}\lim_{u \rightarrow \infty} \frac{A(u)}{u} &= \frac{1}{4}(\beta_\infty - D_\infty)\tau , \\ \lim_{u \rightarrow \infty} \frac{B_v(u)}{u} &= \frac{\beta_\infty - D_\infty}{2\gamma} .\end{aligned}$$

Combining all result we obtain

$$C_\infty = D_\infty (\tau + V_0/\omega^2) / 4 .$$

which completes the proof. □

References

- [AA02] L.B.G. Andersen and J. Andreasen. Volatile volatilities. *Risk*, 15(12):163–168, 2002.
- [AP04] L. Andersen and V. Piterbarg. Moment Explosions in Stochastic Volatility Models. Technical report, Bank of America, 2004. ssrn.com/abstract=559481.
- [ASY06] Y. Aït-Sahalia and Jialin. Yu. Saddlepoint Approximations for Continuous-Time Markov Processes. *Working paper*, 2006. <http://www.princeton.edu/~yacine/saddlepoint.pdf>.
- [Bat96] D.S. Bates. Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *The Review of Financial Studies*, 9(1):69–107, 1996.
- [CM99] P. Carr and D. Madan. Option valuation using the Fast Fourier Transform. *Journal of Computational Finance*, 2(4):61–73, 1999.

- [CS05] P. Cheng and O. Scaillett. Linear-quadratic jump-diffusion modelling with application to stochastic volatility. *Working paper, Crédit Suisse, FAME and HEC Genève*, 2005.
- [CW97] G.L. Choudhury and W. Whitt. Probabilistic scaling for the numerical inversion of non-probability transforms. *INFORMS J. Computing*, 9:175–184, 1997.
- [DA68] H. Dubner and J. Abate. Numerical inversion of Laplace transforms by relating them to the finite Fourier cosine transform. *Journal of the ACM*, 15(1):115–123, 1968.
- [Dan54] H.E. Daniels. Saddlepoint Approximations in Statistics. *The Annals of Mathematical Statistics*, 25(4):631–650, 1954.
- [FP91] K.-J. Förster and K. Petras. Error estimates in Gaussian quadrature for functions of bounded variations. *SIAM Journal of Numerical Analysis*, 28(3):880–889, 1991.
- [Gas04] R. Gaspar. General quadratic term structures of bond, futures and forward prices. *SSE/EFI Working paper Series in Economics and Finance*, (559), 2004.
- [GG00] W. Gander and W. Gautschi. Adaptive Quadrature — Revisited. *BIT*, 40(1):84–101, March 2000. CS technical report: <ftp://inf.ethz.ch/pub/publications/tech-reports/3xx/306.ps.gz>.
- [GP51] J. Gil-Pelaez. Note on the inversion theorem. *Biometrika*, 37:481–482, 1951.
- [Gur48] J. Gurland. Inversion formulae for the distribution of ratios. *Annals of Mathematical Statistics*, 19:228–237, 1948.
- [Hes93] S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6:327–343, 1993.
- [KJ05] C. Kahl and P. Jäckel. Not-so-complex logarithms in the Heston model. *Wilmott*, September, 2005. <http://www.math.uni-wuppertal.de/~kahl/publications.html>.
- [Lee04] R. Lee. The Moment Formula for Implied Volatility at Extreme Strikes. *Mathematical Finance*, 14:469–480, 2004.
- [Lee05] R. Lee. Option Pricing by Transform Methods: Extensions, Unification, and Error Control. *Journal of Computational Finance*, 7(3):51–86, 2005.
- [Lév25] P. Lévy. *Calcul des probabilités*. Gauthier-Villars, Paris, 1925.
- [Lew01] A. Lewis. A Simple Option Formula for General Jump-Diffusion and Other Exponential Levy Processes. 2001. <http://ssrn.com/abstract=282110>.
- [LK06] R. Lord and C. Kahl. Why the rotation count algorithm works. *Working paper*, 2006. <http://www.math.uni-wuppertal.de/~kahl/publications/WhyTheRotationCountAlgorithmWorks.pdf>.
- [LR80] R. Lugannani and S. Rice. Saddlepoint approximation for the distribution of the sum of independent random variables. *Advances in Applied Probability*, 12:475–490, 1980.
- [Luk70] E. Lukacs. *Characteristic functions*. Griffin, London, 1970.
- [LZ02] S. Levendorskiĭ and V.M. Zherder. Fast option pricing under regular Lévy processes of exponential type. working paper, University of Texas at Austin and Rostov State University of Economics, 2002.

- [Rai00] S. Raible. *Levy Processes in Finance: Theory, Numerics and Empirical Facts*. PhD thesis, Albert-Ludwigs-Universität Freiburg, Germany, 2000.
- [RZ98] L.C.G. Rogers and O. Zane. Saddlepoint approximations to option pricing. *The annals of applied probability*, 9(2), 1998.
- [SS91] E. Stein and J. Stein. Stock-Price Distributions with Stochastic Volatility - An Analytic Approach. *Review of Financial Studies*, 4:727–752, 1991.
- [SST04] W. Schoutens, E. Simons, and J. Tistaert. A perfect calibration! Now what? *Wilmott Magazine*, March 2004.
- [SZ99] R. Schöbel and J. Zhu. Stochastic Volatility with an Ornstein-Uhlenbeck Process: An Extension. *European Finance Review*, 3:23–46, 1999.