

A Finite Volume - Alternating Direction Implicit Approach for the Calibration of Stochastic Local Volatility Models

Maarten Wyns^{*1} and Jacques Du Toit²

¹Department of Mathematics and Computer Science, University of Antwerp,
Middelheimlaan 1, B-2020 Antwerp, Belgium. Email:

`maarten.wyns@uantwerpen.be`;

²Numerical Algorithms Group, Peter House, Oxford street, Manchester, M1 5AN,
United Kingdom. Email: `jacques@nag.co.uk`.

February 10, 2017

Abstract

Calibration of stochastic local volatility (SLV) models to their underlying local volatility model is often performed by numerically solving a two-dimensional non-linear forward Kolmogorov equation. We propose a novel finite volume (FV) discretization in the numerical solution of general 1D and 2D forward Kolmogorov equations. The FV method does not require a transformation of the PDE. This constitutes a main advantage in the calibration of SLV models as the pertinent PDE coefficients are often nonsmooth. Moreover, the FV discretization has the crucial property that the total numerical mass is conserved. Applying the FV discretization in the calibration of SLV models yields a non-linear system of ODEs. Numerical time stepping is performed by the Hundsdorfer–Verwer ADI scheme to increase the computational efficiency. The non-linearity in the system of ODEs is handled by introducing an inner iteration. Ample numerical experiments are presented that illustrate the effectiveness of the calibration procedure.

Key words: Stochastic local volatility; Forward Kolmogorov equation; Finite volume discretization; ADI methods; Calibration.

2010 AMS Subject Classification: 35K55, 97M30.

1 Introduction

In contemporary financial mathematics, *stochastic local volatility* (SLV) models constitute state-of-the-art models to describe asset price processes, notably foreign exchange (FX) rates. Let S_τ represent the exchange rate at time $\tau \geq 0$ and let the spot value S_0 be given. For modelling the exchange rate we consider the transformation $X_\tau = \log(S_\tau/S_0)$ since the transformed variable X_τ reflects better the duality between the exchange rate S_τ and $1/S_\tau$, where the latter one is the exchange rate when the role of the domestic currency and the foreign currency are swapped. In this article we consider SLV models of the type

$$\begin{cases} dX_\tau = (r_d - r_f - \frac{1}{2}\sigma_{SLV}^2(X_\tau, \tau)\psi^2(V_\tau))d\tau + \sigma_{SLV}(X_\tau, \tau)\psi(V_\tau)dW_\tau^{(1)}, \\ dV_\tau = \kappa(\eta - V_\tau)d\tau + \xi V_\tau^\alpha dW_\tau^{(2)}, \end{cases} \quad (1.1)$$

^{*}Corresponding author.

with ψ a non-negative function on \mathbb{R}^+ such that $\psi(0) = 0$, $\alpha, \kappa, \eta, \xi$ strictly positive parameters, $dW_\tau^{(1)} \cdot dW_\tau^{(2)} = \rho d\tau$, $-1 \leq \rho \leq 1$ and given spot values $X_0 = 0, V_0$. The non-negative function $\sigma_{SLV}(x, \tau)$ is called the *leverage function* and the constant r_d , respectively r_f , denotes the risk-free interest rate in the domestic currency, respectively in the foreign currency. It is readily seen that the process V_τ is always non-negative. In [1] it is shown that the boundary $V_\tau = 0$ is attainable for $0 < \alpha < 1/2$ and for $\alpha = 1/2$ if $2\kappa\eta < \xi^2$. For $\alpha > 1/2$ it holds that $V_\tau = 0$ is an unattainable boundary. Furthermore, $V_\tau = \infty$ is an unattainable boundary for all values of $\alpha > 0$. The choice $\psi(v) = \sqrt{v}$, $\alpha = 1/2$ corresponds to the Heston-based SLV model and the choice $\psi(v) = v$, $\alpha = 1$ corresponds to the SLV model described in [28].

SLV models constitute a natural combination of *local volatility* (LV) models and *stochastic volatility* (SV) models, [21, 28]. The former type of models can be described by the stochastic differential equation (SDE)

$$dX_{LV,\tau} = (r_d - r_f - \frac{1}{2}\sigma_{LV}^2(X_{LV,\tau}, \tau))d\tau + \sigma_{LV}(X_{LV,\tau}, \tau)dW_\tau. \quad (1.2)$$

The function $\sigma_{LV}(x, \tau)$ is called the *local volatility function* and can be determined by the Dupire formula [4] such that the LV model reproduces the known market prices for European call and put options. Since the LV model is completely determined by the market prices of European call and put options, it offers no flexibility in matching the market dynamics. The underlying pure SV model corresponding with (1.1) is given by the system of SDEs

$$\begin{cases} dX_{SV,\tau} = (r_d - r_f - \frac{1}{2}\psi^2(V_{SV,\tau}))d\tau + \psi(V_{SV,\tau})dW_\tau^{(1)}, \\ dV_{SV,\tau} = \kappa(\eta - V_{SV,\tau})d\tau + \xi V_{SV,\tau}^\alpha dW_\tau^{(2)}, \end{cases} \quad (1.3)$$

where the function ψ and the parameters need to satisfy the same restrictions as above. SV models are typically well-suited to reflect forward volatilities, but they are often unable to capture the volatility smile exactly, see e.g. [30]. By combining features of the LV model with features of the SV models, SLV models are able to match the market dynamics and to reproduce the market prices for European call and put options. If the leverage function is identically equal to one, then the SLV model reduces to the SV model (1.3). If the parameter ξ equals zero, then the SLV model reduces to a purely LV model.

In financial practice, see e.g. [2, 28], one will first determine the SV parameters κ, η, ξ, ρ such that the underlying SV model matches the current market dynamics. Afterwards, the leverage function $\sigma_{SLV}(x, \tau)$ is calibrated such that the SLV model and the underlying LV model define the same fair value for non-path-dependent European options. Then, since LV models are calibrated exactly to the known market prices for vanilla options, the SLV model also reproduces the known market prices for European call and put options. Let $p(x, v, \tau; X_0, V_0)$ denote the joint density of (X_τ, V_τ) under the SLV model (1.1) and let $p_{LV}(x, \tau; X_0)$ denote the density of $X_{LV,\tau}$ under (1.2). It is well-known, see e.g. [9, 27], that both the SLV model (1.1) and the LV model (1.2) define the same marginal distribution for S_τ , i.e.

$$p_{LV}(x, \tau; X_0) = \int_0^\infty p(x, v, \tau; X_0, V_0)dv, \quad (1.4)$$

for $\tau > 0$, if

$$\sigma_{LV}^2(x, \tau) = \mathbb{E}[\sigma_{SLV}^2(X_\tau, \tau)\psi^2(V_\tau)|X_\tau = x] = \sigma_{SLV}^2(x, \tau)\mathbb{E}[\psi^2(V_\tau)|X_\tau = x]. \quad (1.5)$$

If one can determine the conditional expectation above, and if the leverage function is defined by (1.5), then both models yield the same fair value for non-path-dependent European options and hence the SLV model is calibrated exactly to the market prices for vanilla options. This is, however, highly non-trivial since the conditional expectation itself depends on the leverage function.

In the past years, a variety of numerical techniques, see e.g. [2, 5, 11, 24, 30], has been proposed in order to approximate the conditional expectation from (1.5) and to approximate the appropriate

leverage function. The authors in [11, 30] make use of Monte Carlo techniques, whereas in [2, 5, 24] partial differential equation (PDE) methods are applied.

In this article, the PDE approach is considered. For the effective calibration, the conditional expectation is then often rewritten as, cf. [2, 5, 24],

$$\mathbb{E}[\psi^2(V_\tau)|X_\tau = x] = \frac{\int_0^\infty \psi^2(v)p(x, v, \tau; X_0, V_0)dv}{\int_0^\infty p(x, v, \tau; X_0, V_0)dv}. \quad (1.6)$$

It can be shown, see e.g. [25], that the joint density function satisfies the *forward Kolmogorov equation*

$$\begin{aligned} \frac{\partial}{\partial \tau} p &= \frac{\partial^2}{\partial x^2} \left(\frac{1}{2} \sigma_{SLV}^2 \psi^2(v) p \right) + \frac{\partial^2}{\partial x \partial v} (\rho \xi \sigma_{SLV} \psi(v) v^\alpha p) + \frac{\partial^2}{\partial v^2} \left(\frac{1}{2} \xi^2 v^{2\alpha} p \right) \\ &- \frac{\partial}{\partial x} \left((r_d - r_f - \frac{1}{2} \sigma_{SLV}^2 \psi^2(v)) p \right) - \frac{\partial}{\partial v} (\kappa(\eta - v) p), \end{aligned} \quad (1.7)$$

for $x \in \mathbb{R}, v > 0, \tau > 0$ and with initial condition $p(x, v, 0; X_0, V_0) = \delta(x)\delta(v - V_0)$ where δ denotes the Dirac delta function. Once the joint density p is known, one can easily determine the leverage function by computing the integrals in (1.6) and the SLV model is calibrated exactly to the LV model. By combining (1.5), (1.6) and (1.7) it is readily seen that one has to solve a highly non-linear PDE in order to perform the calibration.

In financial mathematics, convection-diffusion equations of the type (1.7) are often discretized by means of finite difference methods, see e.g. [2, 24]. If the parameter α is less or equal to $1/2$, however, it holds that $V_\tau = 0$ is attainable and defining a proper boundary condition at $v = 0$ is a non-trivial task. Moreover, finite difference methods are often not mass-conservative whereas conservation of mass is a key property of forward Kolmogorov equations. The finite volume method proposed in [5] manages to deal with the issues above. The latter method, however, makes use of a transformation of the original PDE (1.7) which incorporates derivatives of the leverage function. As the leverage function is often nonsmooth and only known at a finite number of points, this could lead to undesirable (erratic) behaviour.

In this paper we will introduce a finite volume - alternating direction implicit discretization for the numerical solution of general, non-transformed forward Kolmogorov equations of the type (1.7). The discretization makes use of the general *method of lines* (MOL), cf. [20]. The PDE is first discretized in the spatial variables, yielding large systems of stiff ordinary differential equations (ODEs). These so-called *semidiscrete systems* are subsequently solved by applying a suitable implicit time stepping method. The spatial discretization is performed by *finite volume methods* to keep the total numerical mass equal to one and to handle the boundary conditions in a natural way. Since the PDE (1.7) is multidimensional, the temporal discretization is performed by an *alternating direction implicit* (ADI) scheme, more precisely the *Hundsdorfer-Verwer* (HV) scheme. This can yield a large computational advantage in comparison with standard (non-split) implicit time stepping methods. Finally, for the calibration of the SLV model to the LV model, an inner iteration is introduced in order to handle the non-linearity from inserting (1.6) into (1.7).

An outline of the rest of our paper is as follows.

In Section 2 a finite volume discretization is introduced for the spatial discretization of general 1D and 2D forward Kolmogorov equations. The performance of the finite volume discretization is illustrated by ample numerical experiments.

The spatial discretization results in a large system of ODEs. In Section 3 an ADI temporal discretization scheme is applied to increase the computational efficiency in the numerical solution of this ODE system.

In Section 4 the finite volume discretization is used for the calibration of SLV models, yielding a large non-linear system of ODEs. The ADI scheme is applied for the numerical solution of this system of ODEs and an iteration procedure is described for handling the non-linearity.

In Section 5 numerical experiments are presented to illustrate the performance of the obtained calibration procedure and the final Section 6 concludes.

2 Spatial discretization of forward Kolmogorov equations

In the general MOL approach the PDE is first discretized in the spatial variables by for example finite difference (FD) or finite volume (FV) methods. In this section a spatial discretization is proposed for a general two-dimensional forward Kolmogorov equation of the type

$$\frac{\partial}{\partial \tau} p + \frac{\partial}{\partial x} (\mu_1 p) + \frac{\partial}{\partial y} (\mu_2 p) = \frac{\partial^2}{\partial x^2} \left(\frac{1}{2} \sigma_1^2 p \right) + \frac{\partial^2}{\partial x \partial y} (\rho \sigma_1 \sigma_2 p) + \frac{\partial^2}{\partial y^2} \left(\frac{1}{2} \sigma_2^2 p \right), \quad (2.1)$$

with $x, y \in \mathbb{R}$, $\tau > 0$, and where $\sigma_1, \sigma_2, \mu_1, \mu_2$ are real coefficient functions of x, y, τ . Moreover, the functions σ_1, σ_2 are required to be non-negative and it is assumed that there exist values X_0, Y_0 such that the initial function is given by $p(x, y, 0) = \delta(x - X_0)\delta(y - Y_0)$. Due to the form of the coefficients it is possible that the spatial domain is naturally restricted. For example, if $\mu_2(x, y, \tau) = \kappa(\eta - y)$ and $\sigma_2(x, y, \tau) = \xi y^\alpha$ with $\kappa, \eta, \xi, \alpha$ strictly positive constants, then the domain in the y -direction is naturally restricted to $y \geq 0$, cf. PDE (1.7).

Since the solution of forward Kolmogorov equations represents the density of an underlying stochastic process, conservation of mass is a fundamental property and the use of FV schemes is appropriate. While FD methods are well known in finance, FV methods are less common in financial applications and we briefly recall the basic idea.

2.1 Introduction to finite volume discretizations

Finite volume methods were originally developed to solve conservation laws, or more generally to solve PDEs in conservative form. For example, consider the one-dimensional conservative PDE

$$\frac{\partial}{\partial \tau} p + \frac{\partial}{\partial x} (a(p, x, \tau)p) = \frac{\partial}{\partial x} (b(p, x, \tau) \frac{\partial}{\partial x} p), \quad (2.2)$$

for $x \in \Omega$, $\tau > 0$, where Ω is an interval in \mathbb{R} . Both sides of equation (2.2) can be integrated in x over an interval (more generally, a cell) $[x_l, x_u]$ in order to get

$$\frac{\partial}{\partial \tau} \int_{x_l}^{x_u} p dx = f(p, x_l, \tau) - f(p, x_u, \tau), \quad (2.3)$$

where $f(p, x, \tau)$ is a function given by

$$f(p, x, \tau) = a(p, x, \tau)p - b(p, x, \tau) \frac{\partial}{\partial x} p.$$

The function f is typically called the *flux* of p and $f(p, x, \tau)|_{x=x_l}, f(p, x, \tau)|_{x=x_u}$ represent the fluxes at the left and right boundaries of the cell $[x_l, x_u]$. Relationship (2.3) shows that the total integral of p , which typically represents a mass, momentum or some similar quantity, changes only as a result of the flux difference over the cell. If equation (2.2) is considered over a bounded domain $[x_{\min}, x_{\max}]$ and we assume that $f(p, x_{\min}, \tau) = f(p, x_{\max}, \tau)$ for all τ , i.e. the flux at the left boundary is exactly matched by the flux at the right boundary, then the space integral of p over $[x_{\min}, x_{\max}]$ is constant in time. This means that the total mass or momentum is conserved. If the spatial domain Ω of the PDE is unbounded, and if the interval $[x_{\min}, x_{\max}]$ is wide enough, then one will often have that $f(p, x_{\min}, \tau) \approx f(p, x_{\max}, \tau) \approx 0$ for all $\tau > 0$.

To construct a numerical FV scheme we start with a discretization of the spatial domain. If the spatial domain is unbounded, it needs to be truncated to a wide, finite interval $[x_{\min}, x_{\max}]$. Then, consider the discretization

$$x_{\min} = x_1 < x_2 < \dots < x_m = x_{\max},$$

of the domain of interest and denote

$$\Delta x_i = x_i - x_{i-1}, \quad \text{for } 2 \leq i \leq m,$$

and $\Delta x_1 = \Delta x_{m+1} = 0$. Define mid-points

$$x_{i-0.5} = x_i - \frac{1}{2} \Delta x_i = \frac{x_{i-1} + x_i}{2} \quad \text{for } 2 \leq i \leq m,$$

and let $\Omega_i = [x_{i-0.5}, x_{i+0.5}]$ be cells for $i = 1, 2, \dots, m$ where we define $x_{0.5} := x_1$ and $x_{m+0.5} := x_m$. This yields a *vertex centred grid* with cell vertices $x_{i-0.5}$. We can now consider the cell average $\bar{p}_i(\tau)$ which is defined by

$$\bar{p}_i(\tau) = \frac{1}{x_{i+0.5} - x_{i-0.5}} \int_{\Omega_i} p(x, \tau) dx,$$

and which is typically the quantity that FV schemes approximate. If we assume that the grid is *smooth* in the sense that $\Delta x_{i+1} - \Delta x_i = \mathcal{O}(\Delta x^2)$ where Δx denotes a maximal mesh width, then the cell average $\bar{p}_i(\tau)$ is a second-order approximation to $p(x_i, \tau)$. Differentiating $\bar{p}_i(\tau)$ in τ and using (2.3) gives

$$\bar{p}'_i(\tau) = \frac{f(p, x_{i-0.5}, \tau) - f(p, x_{i+0.5}, \tau)}{x_{i+0.5} - x_{i-0.5}} \quad (2.4)$$

which is just another way of stating the conservation property since if we sum over all cells and pull out the derivative in time we find that

$$\frac{\partial}{\partial \tau} \sum_{i=1}^m \bar{p}_i(\tau) (x_{i+0.5} - x_{i-0.5}) = 0,$$

provided $f(p, x_{\min}, \tau) = f(p, x_{\max}, \tau)$.

Equation (2.4) is typically taken as the starting point for the numerical discretization. Denote by

$$P_i(\tau) \approx \bar{p}_i(\tau), \quad \text{for } 1 \leq i \leq m,$$

the numerical approximations for the cell averages and let P be the vector that contains these approximations. The numerical discretization is then defined by

$$P'_i(\tau) = \frac{f_{i-0.5}(P, \tau) - f_{i+0.5}(P, \tau)}{x_{i+0.5} - x_{i-0.5}} = [f_{i-0.5}(P, \tau) - f_{i+0.5}(P, \tau)] \frac{2}{\Delta x_i + \Delta x_{i+1}}, \quad (2.5)$$

where the $f_{i\pm 0.5}$ are numerical fluxes that approximate the exact fluxes $f(p, x_{i\pm 0.5}, \tau)$. By defining a discretization of the type (2.5), it readily follows that the total numerical integral (mass)

$$\sum_{i=1}^{m_1} P_i(\tau) (x_{i+0.5} - x_{i-0.5}) \quad (2.6)$$

stays constant in time provided that $f_{0.5}(P, \tau) = f_{m+0.5}(P, \tau)$. It is clear that the exact fluxes from (2.4) involve the unknown function p at the cell boundaries $x_{i\pm 0.5}$. Therefore, we define the numerical fluxes by

$$f_{i\pm 0.5}(P, \tau) = a(P_{i\pm 0.5}, x_{i\pm 0.5}, \tau) P_{i\pm 0.5} - b(P_{i\pm 0.5}, x_{i\pm 0.5}, \tau) P_{x, i\pm 0.5}, \quad (2.7)$$

where the $P_{i\pm 0.5}(\tau)$ form approximations to the exact values $p(x_{i\pm 0.5}, \tau)$ and the $P_{x, i\pm 0.5}(\tau)$ form approximations to $\frac{\partial}{\partial x} p(x, \tau)|_{x=x_{i\pm 0.5}}$. Since the cell average is a second-order approximation to p , the approximations P_i can be used to define $P_{i\pm 0.5}$ and $P_{x, i\pm 0.5}$. The manner in which the latter values are computed at the cell boundaries from the surrounding P_i plays a large part in defining the characteristics of the numerical scheme. Lastly, inserting (2.7) into (2.5) yields a system of (possibly non-linear) ODEs which is solved with a suitable time integration procedure.

Recall that conservation of mass is a fundamental property of forward Kolmogorov equations and the use of FV schemes is appropriate. Forward Kolmogorov equations of the type (2.1) are, however, not in conservative form and hence straightforward application of standard FV schemes is not possible. Moreover, rewriting PDE (2.1) in conservative form would involve derivatives of the coefficient functions, which are not known in general practical applications. In the remainder of this section, a FV-based discretization of the spatial derivatives in the non-transformed PDE (2.1) is introduced such that conservation of total mass is guaranteed. We start by explaining the discretization for a general one-dimensional forward Kolmogorov equation, and then generalise it to the two-dimensional case.

2.2 One-dimensional forward Kolmogorov equations

Standard one-dimensional forward Kolmogorov equations are also not written in conservative form and their solutions represent density functions of underlying stochastic processes. In this subsection a FV-based discretization is introduced for the general one-dimensional equation

$$\frac{\partial}{\partial \tau} p + \frac{\partial}{\partial x} (\mu p) = \frac{\partial^2}{\partial x^2} \left(\frac{1}{2} \sigma^2 p \right), \quad (2.8)$$

for $x \in \mathbb{R}$, $\tau > 0$, where σ, μ are real functions of x and τ , with σ non-negative and with initial function given by $p(x, 0) = \delta(x - X_0)$ for some real X_0 . Spatial discretization by FD or FV methods is often applied on a finite grid. By consequence, the spatial domain has to be truncated to $[x_{\min}, x_{\max}]$, where the boundaries are chosen sufficiently far away from X_0 such that the truncation error is negligible. Recall that the form of σ, μ can naturally restrict the spatial domain of the PDE to for example $x \geq 0$. In the latter case, the lower boundary is naturally defined as $x_{\min} = 0$.

As before, define a spatial mesh $x_{\min} = x_1 < x_2 < \dots < x_m = x_{\max}$, let $\Delta x_i = x_i - x_{i-1}$ be the mesh widths, with $\Delta x_1 = \Delta x_{m+1} = 0$, and define

$$x_{i-0.5} = x_i - \frac{1}{2} \Delta x_i = \frac{x_{i-1} + x_i}{2} \quad \text{for } 2 \leq i \leq m,$$

with $x_{0.5} = x_1$ and $x_{m+0.5} = x_m$. This yields a vertex centred grid with cells $\Omega_i = [x_{i-0.5}, x_{i+0.5}]$. Let the $P_i(\tau)$ denote approximations to the exact cell averages

$$\bar{p}_i(\tau) = \frac{1}{x_{i+0.5} - x_{i-0.5}} \int_{\Omega_i} p(x, \tau) dx,$$

and let P be the vector containing these approximations. Analogously to the previous section (see equations (2.4) and (2.5), as well as [20]) we define discretizations of the form

$$P'_i(\tau) = [f_{i-0.5}(P, \tau) - f_{i+0.5}(P, \tau)] \frac{2}{\Delta x_i + \Delta x_{i+1}} \quad (2.9)$$

where the numerical fluxes are given by

$$f_{i\pm 0.5}(P, \tau) = f_{a,i\pm 0.5}(P, \tau) + f_{d,i\pm 0.5}(P, \tau),$$

with

$$f_{a,i\pm 0.5}(P, \tau) \approx \mu(x_{i\pm 0.5}, \tau) p(x_{i\pm 0.5}, \tau), \quad (2.10)$$

and

$$f_{d,i\pm 0.5}(P, \tau) \approx -\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma^2(x, \tau) p(x, \tau) \right) \Big|_{x=x_{i\pm 0.5}}. \quad (2.11)$$

For the ease of presentation, from now on we omit the dependence of the parameters on τ and set $\mu_{i\pm 0.5} = \mu(x_{i\pm 0.5}, \tau)$ and $\sigma_i = \sigma(x_i, \tau)$. Note that $f_{0.5}(P, \tau)$, respectively $f_{m+0.5}(P, \tau)$, corresponds with the flux at the boundary $x_{\min} = x_1$, respectively $x_{\max} = x_m$.

The advection part of the PDE (2.8) is written in conservative form. For the inner cell boundaries, i.e. for $x_{i-0.5}$ with $2 \leq i \leq m$, we consider the second-order central FV scheme, cf. [20], and define $f_{a,i-0.5}(P, \tau)$ in (2.10) as

$$f_{a,i-0.5}(P, \tau) = \mu_{i-0.5} \frac{P_{i-1}(\tau) + P_i(\tau)}{2}.$$

The diffusion part is not written in conservative form and hence it is not possible to apply standard FV schemes to this term directly. The idea of the second-order FV scheme for (2.11), see e.g. [20], is generalised by defining

$$f_{d,i-0.5}(P, \tau) = -\left(\frac{1}{2} \sigma_i^2 P_i(\tau) - \frac{1}{2} \sigma_{i-1}^2 P_{i-1}(\tau) \right) \frac{1}{\Delta x_i}$$

for $2 \leq i \leq m$. It is readily seen that

$$\left(\frac{1}{2}\sigma_i^2 p(x_i, \tau) - \frac{1}{2}\sigma_{i-1}^2 p(x_{i-1}, \tau)\right) \frac{1}{\Delta x_i}$$

is a second-order approximation of $\frac{\partial}{\partial x}(\frac{1}{2}\sigma^2 p)$ at the point $x_{i-0.5}$ which explains the choice for this discretization. Inserting these expressions back into (2.9) we get

$$\begin{aligned} P'_i(\tau) &= \frac{\sigma_{i-1}^2 P_{i-1}(\tau)}{\Delta x_i(\Delta x_i + \Delta x_{i+1})} - \frac{\sigma_i^2 P_i(\tau)}{\Delta x_i \Delta x_{i+1}} + \frac{\sigma_{i+1}^2 P_{i+1}(\tau)}{\Delta x_{i+1}(\Delta x_i + \Delta x_{i+1})} \\ &+ \left[\mu_{i-0.5} \frac{P_{i-1}(\tau) + P_i(\tau)}{2} - \mu_{i+0.5} \frac{P_i(\tau) + P_{i+1}(\tau)}{2} \right] \frac{2}{\Delta x_i + \Delta x_{i+1}}, \end{aligned} \quad (2.12)$$

for $2 \leq i \leq m-1$. Note that by applying the second-order central FD scheme for diffusion on non-uniform spatial grids, cf. [14], on the term $\frac{\partial^2}{\partial x^2}(\frac{1}{2}\sigma^2 p)$, one would end up with the same discretization for the diffusion term.

To complete this semidiscretization, it also has to be defined at the boundaries of the truncated domain such that conservation of the total mass is guaranteed. Given that p represents a density function, it follows that

$$\int_{-\infty}^{\infty} p(x, \tau) dx = 1, \quad \forall \tau > 0,$$

and hence

$$\int_{-\infty}^{\infty} \left[\frac{\partial}{\partial \tau} p \right] dx = \int_{-\infty}^{\infty} \left[\frac{1}{2} \frac{\partial^2}{\partial x^2} (\sigma^2 p) - \frac{\partial}{\partial x} (\mu p) \right] dx = 0.$$

Assuming that $[x_{\min}, x_{\max}]$ is chosen sufficiently wide, the condition above can be approximated by

$$\left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma^2 p \right) - (\mu p) \right] \Big|_{x=x_{\min}}^{x=x_{\max}} = 0,$$

reflecting the fact that the total flux over the interval $[x_{\min}, x_{\max}]$ is zero. As stated above, for some choices of coefficient functions the spatial domain is naturally restricted. For example, if the PDE (2.8) stems from a non-negative process, x_{\min} can be set equal to zero. The no-flux boundary condition above then still holds on the naturally restricted domain.

In theory it is possible that there is a positive flux at one of the boundaries, and exactly the same negative flux at the other boundary. However, since the solution represents a density function, it is more realistic to impose that no mass is coming in or going out at each of the boundaries. In light of this, we assume that the following boundary conditions hold:

$$\left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma^2 p \right) - (\mu p) \right] \Big|_{x=x_{\min}} = 0, \quad (2.13)$$

$$\left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma^2 p \right) - (\mu p) \right] \Big|_{x=x_{\max}} = 0.$$

The numerical equivalent of the first condition is to say that the flux at $x_1 = x_{\min}$ is zero, i.e. $f_1(P, \tau) \equiv f_{0.5}(P, \tau) = 0$. This can be achieved by creating a ghost point $x_0 = x_1 - \Delta x_2$ and using (2.13) to define the value of $\frac{1}{2}\sigma_0^2 P_0(\tau)$ at the ghost point by

$$\frac{\frac{1}{2}\sigma_2^2 P_2(\tau) - \frac{1}{2}\sigma_0^2 P_0(\tau)}{2\Delta x_2} - \mu_1 P_1(\tau) = 0,$$

where we make use of the fact that the cell averages form second-order approximations to the point values. Turning to (2.11) we define the diffusive flux $f_{d,0.5}$ at x_1 as

$$f_{d,0.5}(P, \tau) = -\frac{\frac{1}{2}\sigma_2^2 P_2(\tau) - \frac{1}{2}\sigma_0^2 P_0(\tau)}{2\Delta x_2} = -\mu_1 P_1(\tau).$$

Since x_1 is the left boundary of the first cell, the flux on the boundary x_{\min} stemming from the advection part (see (2.10)) can be approximated as $f_{a,0.5} = \mu_1 P_1(\tau)$. Inserting these expressions into (2.9) we obtain

$$P'_1(\tau) = -f_{1.5}(P, \tau) \frac{2}{\Delta x_1 + \Delta x_2} = -f_{1.5}(P, \tau) \frac{2}{\Delta x_2}. \quad (2.14)$$

The boundary condition at x_{\max} can be handled analogously in order to get

$$P'_m(\tau) = f_{m-0.5}(P, \tau) \frac{2}{\Delta x_m}. \quad (2.15)$$

By performing the discretization of the boundary conditions in this way, it follows that $f_{0.5}(P, \tau) = f_{m+0.5}(P, \tau)$ and we ensure that mass is conserved in the numerical scheme.

Combining (2.12), (2.14) and (2.15) we see that the total discretization can be written as a system of ODEs

$$P'(\tau) = A(\tau)P(\tau) \quad (2.16)$$

for $\tau > 0$, with given matrix $A(\tau)$. Since $P_i(\tau)$ represent cell averages it is natural to define the initial vector as

$$P_i(0) = \begin{cases} \frac{2}{\Delta x_i + \Delta x_{i+1}} & \text{if } X_0 \in [x_{i-0.5}, x_{i+0.5}], \\ 0 & \text{otherwise.} \end{cases}$$

In general, the exact solution of the system of ODEs (2.16) can not be computed analytically and one relies on numerical methods in order to approximate it. Since the discretization above often leads to stiff semidiscrete systems, suitable implicit time stepping schemes such as the *Crank–Nicolson scheme* are widely considered, see e.g. [29].

2.3 Numerical experiments for 1D forward Kolmogorov equations

In this subsection the performance of the FV discretization is tested by considering two practical examples. As a first example, consider the SDE

$$dS_\tau = (r_d - r_f)S_\tau d\tau + \sigma_{BS}S_\tau dW_\tau,$$

with $\sigma_{BS} > 0$, corresponding to the classical Black–Scholes model. Then, the underlying density is known exactly and given by

$$p(s, \tau) = \frac{1}{\sigma_{BS}\sqrt{\tau}} \phi\left(\frac{\log(s/S_0) - (r_d - r_f - \frac{1}{2}\sigma_{BS}^2)\tau}{\sigma_{BS}\sqrt{\tau}}\right) \frac{1}{s}, \quad \text{for } s > 0, \tau > 0, \quad (2.17)$$

where $\phi(x)$ is the density function of a standard normally distributed random variable. The density function $p(s, \tau)$ from (2.17) satisfies the PDE

$$\frac{\partial}{\partial \tau} p = \frac{\partial^2}{\partial s^2} \left(\frac{1}{2}\sigma_{BS}^2 s^2 p\right) - \frac{\partial}{\partial s} ((r_d - r_f)sp),$$

for $s, \tau > 0$, with $p(s, 0) = \delta(s - S_0)$. This PDE is of the form (2.8) with a natural restriction of the spatial domain. Note that for the numerical experiment we don't apply the log-transformation from Section 1 in order to have non-constant coefficients which makes the problem more challenging.

Firstly, the spatial domain is truncated to $[S_{\min}, S_{\max}] = [0, 30S_0]$ and we define a non-uniform grid $S_{\min} = s_1 < s_2 < \dots < s_m = S_{\max}$ by making use of a sinh transformation of a uniform underlying grid, cf. [14]. In Figure 1 the spatial grid is shown for the values $S_0 = 100, m = 50$ and from $s = 0$ to $s = 5S_0$ to illustrate the smaller mesh widths around the point $s = S_0$. Applying the FV discretization from Subsection 2.2 then yields approximations $P_i(\tau)$ of the exact values $p(s_i, \tau)$.

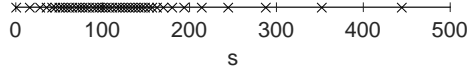


Figure 1: Illustration of the non-uniform grid around $s = S_0$ for the Black-Scholes example and the actual values $S_0 = 100, m = 50$.

When trying to determine the performance of a numerical method with respect to a reference solution, it is important to take note of the computational environment in which values are calculated, and to understand the impact that has on the comparison. Our calculations take place in 64 bit IEEE floating point arithmetic. Since the solution of forward Kolmogorov equations represents a probability density function, the magnitude of the solution varies dramatically over the computational domain. This is especially true of the initial condition (a Dirac delta), and also more generally with naturally bounded stochastic processes with an attainable boundary, c.f. the SLV model (1.1) with $0 < \alpha \leq 1/2$. Since IEEE floating point has a fixed-length mantissa, the density function cannot be represented to a high *absolute* accuracy uniformly over the domain. In areas where the density function is large, only high *relative* accuracy (correct number of digits) can be achieved. In addition, since the numerical solution is obtained by using implicit time stepping, we should not expect high relative accuracy of the numerical density in regions where the exact solution is small. This is because when solving the linear systems we combine many terms with very different magnitudes and then sum them up, which in IEEE arithmetic will lead to a loss of relative accuracy. Therefore when comparing the numerical solution and the reference solution we adopt a mixed absolute-relative error metric: we use relative error when the reference solution is larger than 1, absolute error if the reference solution is less than 1, and we take the maximal error value over the whole domain. More precisely, let

$$\epsilon_i(m) = \begin{cases} \left| \frac{p(s_i, T) - P_i(T)}{p(s_i, T)} \right| & \text{if } p(s_i, T) > 1, \\ |p(s_i, T) - P_i(T)| & \text{else.} \end{cases}$$

The total mixed spatial error is then defined by

$$\epsilon(m) = \max_{1 \leq i \leq m} \epsilon_i(m).$$

The value of 1 is somewhat arbitrary. The results, however, are not that sensitive to the crossover value as long as it is not too small. For the actual experiments, the values $P_i(T)$ are approximated by applying the Crank–Nicolson time stepping scheme with a large number of steps such that the temporal discretization error is negligible. In the left plot of Figure 2 the total mixed spatial error is shown for the relevant situation where $r_d = 0.03$, $r_f = 0.01$, $\sigma_{BS} = 0.2$, $T = 1$ and for the number of spatial grid points $m = \{50, 100, \dots, 1000\}$. The corresponding numerical solution for $m = 200$ is shown in the right plot of Figure 2. The convergence plot clearly indicates that the FV discretization is second-order convergent with respect to the current initial-boundary value problem.

As a second example we consider the Cox–Ingersoll–Ross (CIR) process, cf. [3],

$$dV_\tau = \kappa(\eta - V_\tau)d\tau + \xi\sqrt{V_\tau}dW_\tau,$$

where κ, η, ξ are strictly positive parameters. The corresponding density function is given by, see e.g. [3],

$$p(v, \tau) = ce^{-u_0 - u_1 \left(\frac{v_1}{u_0}\right)^{q/2}} I_q(2\sqrt{u_0 u_1}), \quad (2.18)$$

where

$$c = \frac{2\kappa}{\xi^2(1 - e^{-\kappa\tau})}, \quad u_0 = cV_0 e^{-\kappa\tau}, \quad u_1 = cv, \quad q = \frac{2\kappa\eta}{\xi^2} - 1,$$

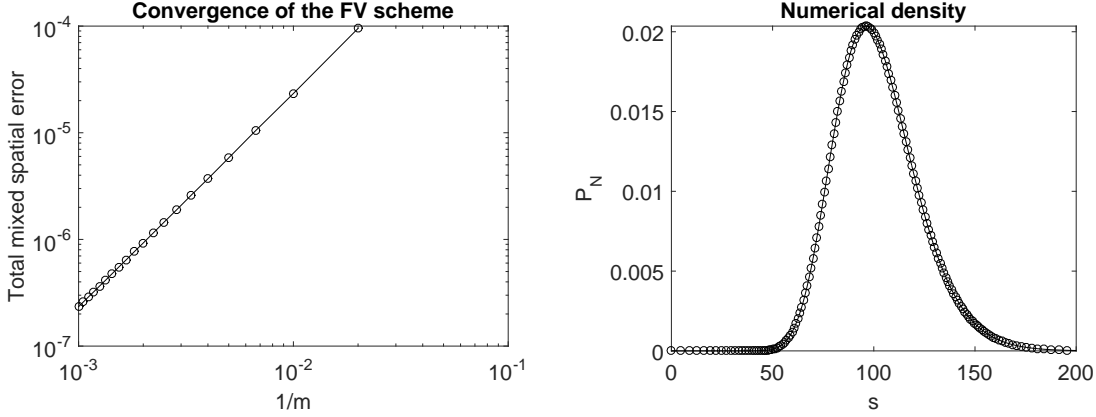


Figure 2: Convergence results within the 1D Black–Scholes model. The parameter values are $r_d = 0.03, r_f = 0.01, \sigma_{BS} = 0.2, T = 1$.

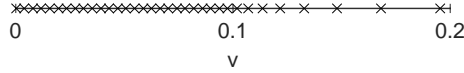


Figure 3: Illustration of the non-uniform grid around $v = 0$ and $v = V_0$ for the CIR example and the actual values $V_0 = 0.0625, m = 50$.

and $I_q(\cdot)$ is the modified Bessel function of the first kind of order q . Note that the value of q is directly related with the so-called Feller condition, i.e. with the possibility that $V_\tau = 0$ is attainable. The density function (2.18) satisfies the forward Kolmogorov equation

$$\frac{\partial}{\partial \tau} p = \frac{\partial^2}{\partial v^2} \left(\frac{1}{2} \xi^2 v p \right) - \frac{\partial}{\partial v} (\kappa(\eta - v)p), \quad (2.19)$$

for $v, \tau > 0$, with $p(v, 0) = \delta(v - V_0)$. It is readily seen that if the Feller condition is violated, i.e. if $q < 0$, then the density from (2.18) is not defined at $v = 0$ and the density function tends to infinity as v tends to zero. In addition, around $v = 0$ the PDE (2.19) is strongly convection dominated which is very challenging for numerical discretization methods.

The domain is truncated to $[V_{\min}, V_{\max}] = [0, 15]$ and a non-uniform grid $0 = v_1 < v_2 < \dots < v_m = V_{\max}$ is defined by making use of a sinh transformation of a uniform underlying grid. The spatial grid is shown in Figure 3 for the values $V_0 = 0.0625, m = 50$ and from $v = 0$ to $v = 0.2$ to illustrate the smaller mesh widths around $v = 0$ and $v = V_0$. Afterwards the FV discretization is applied which leads to approximations $P_j(T)$, ($1 \leq j \leq m$). Recall that if $q < 0$, then the density function tends to infinity as v tends to zero and at $v = 0$ the exact density function is not defined. By increasing the number of spatial grid points m , the value of the second grid point v_2 tends to zero and adequately comparing the difference between $p(v_2, T)$ and $P_2(T)$ becomes difficult. In view of this, we opt to compute the error on similar spatial domains. Let v_{low} be the smallest non-zero grid point, i.e. the point v_2 , if the total number of spatial grid points m is 50. We then define the total mixed spatial error by

$$\max_{j_1 \leq j \leq m} \epsilon_j(m),$$

where, for given m , j_1 is the lowest index such that $v_{j_1} \geq v_{low}$ and

$$\epsilon_j(m) = \begin{cases} \left| \frac{p(v_j, T) - P_j(T)}{p(v_j, T)} \right| & \text{if } p(v_j, T) > 1, \\ |p(v_j, T) - P_j(T)| & \text{else.} \end{cases}$$

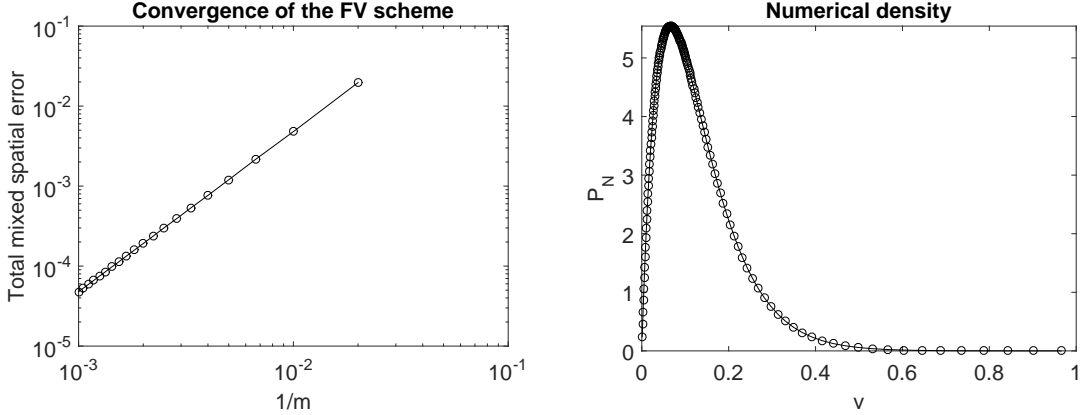


Figure 4: Convergence results within the CIR model. The parameters are given by Set A.

The approximations $P_j(T)$ are determined by considering a large number of time steps such that the temporal discretization error is negligible. Please note that the choice $m = 50$ for defining v_{low} is not crucial. The conclusions of the numerical experiments are essentially unchanged as long as v_{low} is defined via one of the coarsest grids considered in the experiment.

For the actual experiment we consider two sets of parameters:

	κ	η	ξ	V_0	T
Set A	5	0.16	0.9	0.0625	0.25
Set B	1.15	0.0348	0.39	0.0348	0.25

These sets are taken from [7] and were also used in [26]. For Set A we have $q = 0.98$ and the variance process remains strictly positive. For Set B we have $q = -0.47$ and $V_\tau = 0$ is attainable. In the left plot of Figure 4, respectively Figure 5, the total mixed spatial error is shown for the parameters of Set A, respectively Set B, and for the number of spatial grid points $m = \{50, 100, \dots, 1000\}$. In the right plots, the corresponding numerical solutions are shown for $m = 200$. The convergence plots indicate that the FV discretization is convergent with respect to the current initial-boundary value problems. Additional experiments suggest that FV discretization is second-order convergent if the Feller condition is satisfied. If $q < 0$ the order of convergence can drop to one. In addition, all the experiments confirm that the total numerical mass (2.6) stays constantly equal to one, even if the Feller condition is strongly violated.

2.4 Two-Dimensional forward Kolmogorov equations

In this subsection, the FV discretization from the one-dimensional case is used to define a spatial discretization for the general two-dimensional forward Kolmogorov equation (2.1). Suppose the spatial domain is truncated to $[x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$, and the spatial grid points in the x -direction, respectively y -direction, are given by

$$x_{\min} = x_1 < x_2 < \dots < x_{m_1} = x_{\max},$$

respectively

$$y_{\min} = y_1 < y_2 < \dots < y_{m_2} = y_{\max}.$$

Let $\Delta x_i = x_i - x_{i-1}$ and $\Delta y_j = y_j - y_{j-1}$ be the spatial mesh widths, where $\Delta x_1 = \Delta x_{m_1+1} = \Delta y_1 = \Delta y_{m_2+1} = 0$, and define volumes

$$\Omega_{i,j} := [x_{i-0.5}, x_{i+0.5}] \times [y_{j-0.5}, y_{j+0.5}],$$

where $x_{i\pm 0.5}$, respectively $y_{j\pm 0.5}$, are defined analogously as in the one-dimensional case.

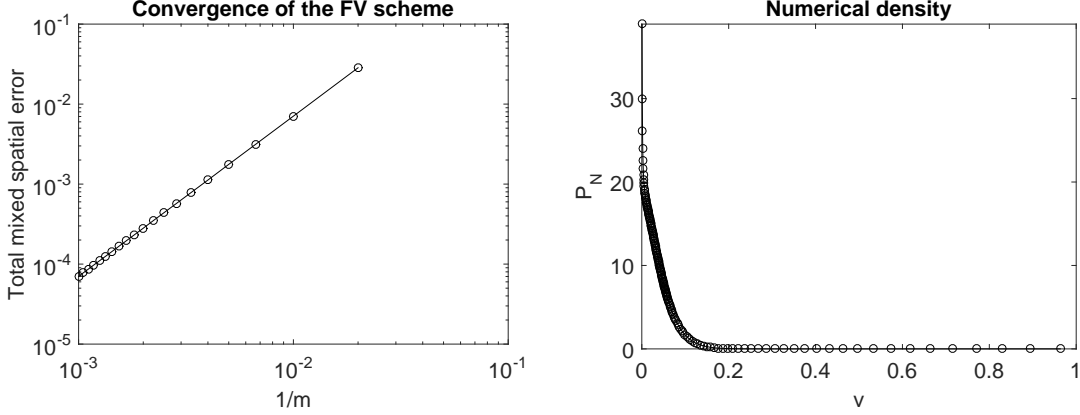


Figure 5: Convergence results within the CIR model. The parameters are given by Set B.

We can now consider the two-dimensional equivalent of the cell averages, volume averages, $\bar{p}_{i,j}(\tau)$ which are defined by

$$\bar{p}_{i,j}(\tau) = \frac{1}{|\Omega_{i,j}|} \int_{\Omega_{i,j}} p(x, y, \tau) dx dy,$$

with $|\Omega_{i,j}| = (x_{i+0.5} - x_{i-0.5})(y_{j+0.5} - y_{j-0.5})$ the area of the corresponding volume. The volume average $\bar{p}_{i,j}(\tau)$ is again a second-order approximation to $p(x_i, y_j, \tau)$, provided that the underlying meshes are smooth, and is the quantity that is approximated by the FV discretization. It is readily verified that

$$|\Omega_{i,j}| \frac{\partial}{\partial \tau} \bar{p}_{i,j}(\tau) = \int_{y_{j-0.5}}^{y_{j+0.5}} \left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma_1^2 p \right) - \mu_1 p \right] \Big|_{x=x_{i-0.5}}^{x=x_{i+0.5}} dy \quad (2.20a)$$

$$+ \int_{x_{i-0.5}}^{x_{i+0.5}} \left[\frac{\partial}{\partial y} \left(\frac{1}{2} \sigma_2^2 p \right) - \mu_2 p \right] \Big|_{y=y_{j-0.5}}^{y=y_{j+0.5}} dx \quad (2.20b)$$

$$+ \left[[\rho \sigma_1 \sigma_2 p] \Big|_{x=x_{i-0.5}}^{x=x_{i+0.5}} \right] \Big|_{y=y_{j-0.5}}^{y=y_{j+0.5}}, \quad (2.20c)$$

and the two-dimensional discretization is based on the approximation

$$|\Omega_{i,j}| \frac{\partial}{\partial \tau} \bar{p}_{i,j}(\tau) \approx \left[\left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma_1^2 p \right) - \mu_1 p \right] \Big|_{x=x_{i-0.5}}^{x=x_{i+0.5}} \right] \Big|_{y=y_j} \frac{\Delta y_j + \Delta y_{j+1}}{2} \quad (2.21a)$$

$$+ \left[\left[\frac{\partial}{\partial y} \left(\frac{1}{2} \sigma_2^2 p \right) - \mu_2 p \right] \Big|_{y=y_{j-0.5}}^{y=y_{j+0.5}} \right] \Big|_{x=x_i} \frac{\Delta x_i + \Delta x_{i+1}}{2} \quad (2.21b)$$

$$+ \left[[\rho \sigma_1 \sigma_2 p] \Big|_{x=x_{i-0.5}}^{x=x_{i+0.5}} \right] \Big|_{y=y_{j-0.5}}^{y=y_{j+0.5}}. \quad (2.21c)$$

Equation (2.20) includes several flux terms that are similar to the flux terms in Subsections 2.1 and 2.2. It reflects the fact that the total integral of p over a volume changes only as a result of the flux difference over the volume boundary. This is completely analogous with the one-dimensional interpretation of (2.3). If the total flux over the boundary of the spatial domain is zero, i.e. if

$$\begin{aligned} 0 &= \int_{y_{\min}}^{y_{\max}} \left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma_1^2 p \right) - \mu_1 p \right] \Big|_{x=x_{\min}}^{x=x_{\max}} dy \\ &+ \int_{x_{\min}}^{x_{\max}} \left[\frac{\partial}{\partial y} \left(\frac{1}{2} \sigma_2^2 p \right) - \mu_2 p \right] \Big|_{y=y_{\min}}^{y=y_{\max}} dx \\ &+ \left[[\rho \sigma_1 \sigma_2 p] \Big|_{x=x_{\min}}^{x=x_{\max}} \right] \Big|_{y=y_{\min}}^{y=y_{\max}}, \end{aligned}$$

then the total integral of p over the entire domain is constant in time.

Let $\mathbf{P}_{i,j}(\tau)$ denote approximations to the exact value $\bar{p}_{i,j}(\tau)$, denote by $\mathbf{P}(\tau)$ the $m_1 \times m_2$ matrix with entries $\mathbf{P}_{i,j}(\tau)$ and let

$$P(\tau) = \text{vec}[\mathbf{P}(\tau)],$$

where $\text{vec}[\cdot]$ denotes the operator that turns any given matrix into a vector by putting its successive columns below each other. The bold notation is only introduced to indicate the subtle difference between the matrix form and the vectorised form of the approximations. Similarly to the one-dimensional discretization in Subsections 2.1 and 2.2, we discretize (2.21) in the following way by introducing numerical fluxes

$$\mathbf{P}'_{i,j}(\tau) = [f_{i-0.5,j}(P, \tau) - f_{i+0.5,j}(P, \tau)] \frac{2}{\Delta x_i + \Delta x_{i+1}} \quad (2.22a)$$

$$+ [f_{i,j-0.5}(P, \tau) - f_{i,j+0.5}(P, \tau)] \frac{2}{\Delta y_j + \Delta y_{j+1}} \quad (2.22b)$$

$$+ \sum_{i_1, j_1=0}^1 (-1)^{i_1+j_1} f_{m, i-0.5+i_1, j-0.5+j_1} \frac{2}{\Delta x_i + \Delta x_{i+1}} \frac{2}{\Delta y_j + \Delta y_{j+1}}, \quad (2.22c)$$

for $1 \leq i \leq m_1, 1 \leq j \leq m_2$. For the ease of presentation, denote

$$\begin{aligned} \mu_{1,i \pm 0.5,j} &= \mu_1(x_{i \pm 0.5}, y_j, \tau), & \sigma_{1,i,j} &= \sigma_1(x_i, y_j, \tau), \\ \mu_{2,i,j \pm 0.5} &= \mu_2(x_i, y_{j \pm 0.5}, \tau), & \sigma_{2,i,j} &= \sigma_2(x_i, y_j, \tau), \\ \sigma_{1,i \pm 0.5,j \pm 0.5} &= \sigma_1(x_{i \pm 0.5}, y_{j \pm 0.5}, \tau), & \sigma_{2,i \pm 0.5,j \pm 0.5} &= \sigma_2(x_{i \pm 0.5}, y_{j \pm 0.5}, \tau). \end{aligned}$$

Since the actual form of the fluxes in (2.21) is completely similar to the form of the fluxes in Subsection 2.4, we define the numerical fluxes by

$$f_{i \pm 0.5,j}(P, \tau) = f_{a,i \pm 0.5,j}(P, \tau) + f_{d,i \pm 0.5,j}(P, \tau),$$

with

$$f_{a,i-0.5,j}(P, \tau) = \mu_{1,i-0.5,j} \frac{\mathbf{P}_{i-1,j}(\tau) + \mathbf{P}_{i,j}(\tau)}{2} \approx \mu_{1,i-0.5,j} p(x_{i-0.5}, y_j, \tau),$$

and

$$f_{d,i-0.5,j}(P, \tau) = \frac{\frac{1}{2}\sigma_{1,i-1,j}^2 \mathbf{P}_{i-1,j}(\tau) - \frac{1}{2}\sigma_{1,i,j}^2 \mathbf{P}_{i,j}(\tau)}{\Delta x_i} \approx -\frac{\partial}{\partial x} \left(\frac{1}{2}\sigma_1^2(x, y_j, \tau) p(x, y_j, \tau) \right) \Big|_{x=x_{i-0.5}},$$

for $2 \leq i \leq m_1, 1 \leq j \leq m_2$. Moreover

$$f_{i,j \pm 0.5}(P, \tau) = f_{a,i,j \pm 0.5}(P, \tau) + f_{d,i,j \pm 0.5}(P, \tau),$$

where

$$f_{a,i,j-0.5}(P, \tau) = \mu_{2,i,j-0.5} \frac{\mathbf{P}_{i,j-1}(\tau) + \mathbf{P}_{i,j}(\tau)}{2} \approx \mu_{2,i,j-0.5} p(x_i, y_{j-0.5}, \tau),$$

and

$$f_{d,i,j-0.5}(P, \tau) = \frac{\frac{1}{2}\sigma_{2,i,j-1}^2 \mathbf{P}_{i,j-1}(\tau) - \frac{1}{2}\sigma_{2,i,j}^2 \mathbf{P}_{i,j}(\tau)}{\Delta y_j} \approx -\frac{\partial}{\partial y} \left(\frac{1}{2}\sigma_2^2(x_i, y, \tau) p(x_i, y, \tau) \right) \Big|_{y=y_{j-0.5}},$$

for $1 \leq i \leq m_1, 2 \leq j \leq m_2$. Finally, for the mixed spatial derivative we define

$$\begin{aligned} f_{m,i-0.5,j-0.5}(P, \tau) &= \rho \sigma_{1,i-0.5,j-0.5} \sigma_{2,i-0.5,j-0.5} \frac{\mathbf{P}_{i-1,j-1}(\tau) + \mathbf{P}_{i-1,j}(\tau) + \mathbf{P}_{i,j-1}(\tau) + \mathbf{P}_{i,j}(\tau)}{4} \\ &\approx \rho \sigma_1(x_{i-0.5}, y_{j-0.5}, \tau) \sigma_2(x_{i-0.5}, y_{j-0.5}, \tau) p(x_{i-0.5}, y_{j-0.5}, \tau), \end{aligned} \quad (2.23)$$

for $1 \leq i \leq m_1 + 1, 1 \leq j \leq m_2 + 1$, where it is assumed that

$$\mathbf{P}_{0,j}(\tau) := \mathbf{P}_{1,j}(\tau), \quad \mathbf{P}_{m_1+1,j}(\tau) := \mathbf{P}_{m_1,j}(\tau), \quad \mathbf{P}_{i,0}(\tau) := \mathbf{P}_{i,1}(\tau), \quad \mathbf{P}_{i,m_2+1}(\tau) := \mathbf{P}_{i,m_2}(\tau), \quad (2.24)$$

such that the general formula is naturally extended at the boundaries of the spatial domain. The numerical flux term (2.23) can be viewed as the result of applying the discretization for the advection part first in the x -direction and then in the y -direction.

The semidiscretization is completed by defining boundary conditions and discretizations at the boundaries of the truncated domain. Since p is again a density function, it follows that

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \left[\frac{\partial}{\partial \tau} p \right] dx dy = 0.$$

Inserting the right hand side of the PDE (2.1), this can be rewritten as

$$\begin{aligned} 0 &= \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} \left[\frac{\partial^2}{\partial x^2} \left(\frac{1}{2} \sigma_1^2 p \right) - \frac{\partial}{\partial x} (\mu_1 p) \right] dx \right) dy \\ &+ \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} \left[\frac{\partial^2}{\partial y^2} \left(\frac{1}{2} \sigma_2^2 p \right) - \frac{\partial}{\partial y} (\mu_2 p) \right] dy \right) dx \\ &+ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\partial^2}{\partial x \partial y} (\rho \sigma_1 \sigma_2 p) dx dy. \end{aligned}$$

Analogously to the one-dimensional case it is assumed that the boundaries are chosen sufficiently far away from the spot value (X_0, Y_0) or that they are defined by a natural truncation of the spatial domain. The condition above is then approximated by

$$\begin{aligned} 0 &= \int_{y_{\min}}^{y_{\max}} \left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma_1^2 p \right) - \mu_1 p \right] \Big|_{x=x_{\min}}^{x=x_{\max}} dy \\ &+ \int_{x_{\min}}^{x_{\max}} \left[\frac{\partial}{\partial y} \left(\frac{1}{2} \sigma_2^2 p \right) - \mu_2 p \right] \Big|_{y=y_{\min}}^{y=y_{\max}} dx \\ &+ \int_{y_{\min}}^{y_{\max}} \int_{x_{\min}}^{x_{\max}} \frac{\partial^2}{\partial x \partial y} (\rho \sigma_1 \sigma_2 p) dx dy. \end{aligned} \tag{2.25}$$

Note that by assuming that

$$\rho \sigma_1 \sigma_2 p \Big|_{x=x_{\min}, y=y_{\min}} = \rho \sigma_1 \sigma_2 p \Big|_{x=x_{\min}, y=y_{\max}} = \rho \sigma_1 \sigma_2 p \Big|_{x=x_{\max}, y=y_{\min}} = \rho \sigma_1 \sigma_2 p \Big|_{x=x_{\max}, y=y_{\max}} = 0, \tag{2.26}$$

the last integral, corresponding with the mixed derivative term, is always equal to zero. Next, we generalise the idea that there are no fluxes at the boundaries, i.e. that no mass is coming in or going out at the boundaries. In light of this it is assumed that the following boundary conditions hold:

$$\begin{aligned} \left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma_1^2 p \right) - (\mu_1 p) \right] \Big|_{x=x_{\min}} &= 0, & \text{for } y_{\min} \leq y \leq y_{\max}, \\ \left[\frac{\partial}{\partial x} \left(\frac{1}{2} \sigma_1^2 p \right) - (\mu_1 p) \right] \Big|_{x=x_{\max}} &= 0, & \text{for } y_{\min} \leq y \leq y_{\max}, \\ \left[\frac{\partial}{\partial y} \left(\frac{1}{2} \sigma_2^2 p \right) - (\mu_2 p) \right] \Big|_{y=y_{\min}} &= 0, & \text{for } x_{\min} \leq x \leq x_{\max}, \\ \left[\frac{\partial}{\partial y} \left(\frac{1}{2} \sigma_2^2 p \right) - (\mu_2 p) \right] \Big|_{y=y_{\max}} &= 0, & \text{for } x_{\min} \leq x \leq x_{\max}. \end{aligned}$$

By combining this boundary conditions with the assumption (2.26) it follows that condition (2.25) is satisfied.

For the discretization of the one-dimensional fluxes in (2.21a) and (2.21b) at the boundaries of the truncated domain, the approach from the one-dimensional case is generalised. By using

a similar discretization of the boundary conditions, one ends up with numerical fluxes which are zero at the boundaries, i.e. with

$$f_{0.5,j}(P, \tau) = f_{m_1+0.5,j}(P, \tau) = f_{i,0.5}(P, \tau) = f_{i,m_2+0.5}(P, \tau) = 0,$$

for $1 \leq i \leq m_1, 1 \leq j \leq m_2$. The fluxes stemming from the mixed derivative term, see (2.21c), are discretized at the boundaries by using (2.23) in combination with (2.24). It is readily verified that if

$$\begin{aligned} 0 &= \rho\sigma_{1,1,1}\sigma_{2,1,1}\mathbf{P}_{1,1}(\tau) = \rho\sigma_{1,1,m_2}\sigma_{2,1,m_2}\mathbf{P}_{1,m_2}(\tau) \\ &= \rho\sigma_{1,m_1,1}\sigma_{2,m_1,1}\mathbf{P}_{m_1,1}(\tau) = \rho\sigma_{1,m_1,m_2}\sigma_{2,m_1,m_2}\mathbf{P}_{m_1,m_2}(\tau), \end{aligned}$$

which is the semidiscrete version of (2.26), then the total numerical flux over the boundary of the spatial domain is equal to zero and the total numerical mass is kept constant in time, i.e.

$$\sum_{i=1}^{m_1} \sum_{j=1}^{m_2} \mathbf{P}_{i,j}(\tau) |\Omega_{i,j}| = \text{constant}, \quad \text{for } \tau \geq 0.$$

As stated above, some processes are naturally bounded, e.g. the general variance process from Section 1 can never become negative. Suppose for example that the process corresponding with the y -variable in the PDE (2.1) is bounded from below. Then, y_{\min} is naturally taken equal to this lower boundary. Moreover, it can happen that this lower boundary is attainable (cf. the variance process from Section 1 with $\alpha < 1/2$) and probability mass can stack up at this boundary. This can cause instabilities in the approximation of the mixed derivative term near this boundary. In order to deal with this, if for example the boundary y_{\min} is attainable, the central FV scheme in the y -direction for the "mixed derivative fluxes" (2.21c) at $y_{\min} + \frac{1}{2}\Delta y_2$ are replaced by a first-order forward scheme. More precisely, the $f_{m,i\pm 0.5,1.5}(P, \tau)$ from above are then replaced by

$$f_{m,i-0.5,1.5}(P, \tau) = \rho\sigma_{1,i-0.5,1.5}\sigma_{2,i-0.5,1.5} \frac{\mathbf{P}_{i-1,2}(\tau) + \mathbf{P}_{i,2}(\tau)}{2},$$

for $1 \leq i \leq m_1 + 1$, where

$$\mathbf{P}_{0,2}(\tau) := \mathbf{P}_{1,2}(\tau), \quad \mathbf{P}_{m_1+1,2}(\tau) := \mathbf{P}_{m_1,2}(\tau).$$

The total spatial discretization (2.22) yields a large system of differential equations. By making use of the *Kronecker product*, this system of differential equations can be written as a system of ODEs

$$P'(\tau) = A(\tau)P(\tau), \tag{2.27}$$

for $\tau > 0$ and with given matrix $A(\tau)$. Analogously to the one-dimensional case, since the values $\mathbf{P}_{i,j}(\tau)$ can be seen as approximations of the cell averages $\bar{p}(x_i, y_j, \tau)$, it is natural to define the initial vector as $P(0) = \text{vec}[\mathbf{P}(0)]$ where

$$\mathbf{P}_{i,j}(0) = \begin{cases} \frac{1}{|\Omega_{i,j}|} & \text{if } (X_0, V_0) \in \Omega_{i,j}, \\ 0 & \text{else.} \end{cases}$$

Note that the matrix A can be split as

$$A = A_0 + A_1 + A_2,$$

where A_1 , respectively A_2 , represents the discretization of the spatial derivatives in the first, respectively second, spatial dimension. The matrix A_0 represents the discretization of the mixed spatial derivative term in (2.1). It is readily verified that A_1 is tridiagonal, A_2 is essentially tridiagonal and A_0 has in general nine non-zero elements per row and column. In Section 3 this structure is used to perform an effective time discretization of the general system of ODEs (2.27).




Figure 6: Illustration of the non-uniform grid around $(S_{1,0}, S_{2,0})$ for the Black–Scholes example and the actual values $S_{1,0} = S_{2,0} = 100$, $m_1 = m_2 = 50$.

2.5 Numerical experiments for 2D forward Kolmogorov equations

In this section the performance of the 2D FV discretization is tested for two practical examples. For the first experiment we consider the two-dimensional Black–Scholes model which can be described by the system of SDEs

$$\begin{cases} dS_{1,\tau} = rS_{1,\tau}d\tau + \sigma_{1,BS}S_{1,\tau}dW_\tau^{(1)}, \\ dS_{2,\tau} = rS_{2,\tau}d\tau + \sigma_{2,BS}S_{2,\tau}dW_\tau^{(2)}, \end{cases}$$

with $dW_\tau^{(1)} \cdot dW_\tau^{(2)} = \rho d\tau$, $-1 \leq \rho \leq 1$ and $r, \sigma_{1,BS}, \sigma_{2,BS}$ strictly positive constants. The corresponding forward Kolmogorov equation is given by

$$\frac{\partial}{\partial \tau} p = \frac{\partial^2}{\partial s_1^2} \left(\frac{1}{2} \sigma_{1,BS}^2 s_1^2 p \right) + \frac{\partial^2}{\partial s_1 \partial s_2} \left(\rho \sigma_{1,BS} \sigma_{2,BS} s_1 s_2 p \right) + \frac{\partial^2}{\partial s_2^2} \left(\frac{1}{2} \sigma_{2,BS}^2 s_2^2 p \right) - \frac{\partial}{\partial s_1} (r s_1 p) - \frac{\partial}{\partial s_2} (r s_2 p),$$

for $s_1, s_2, \tau > 0$ and with $p(s_1, s_2, 0) = \delta(s_1 - S_{1,0})\delta(s_2 - S_{2,0})$ for given values $S_{1,0}, S_{2,0}$. The exact solution is known analytically and can be written as

$$p(s_1, s_2, \tau) = n_2(\log(s_1/S_{1,0}), \log(s_2/S_{2,0}), \tau) \frac{1}{s_1} \frac{1}{s_2}, \quad \text{for } s_1 > 0, s_2 > 0, \tau > 0,$$

where this time $n_2(x, y, \tau)$ is the density function of a two-dimensional normally distributed random variable with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ given by

$$\boldsymbol{\mu} = \begin{bmatrix} (r - \frac{1}{2}\sigma_{1,BS}^2)\tau \\ (r - \frac{1}{2}\sigma_{2,BS}^2)\tau \end{bmatrix}, \quad \boldsymbol{\Sigma} = \begin{bmatrix} \sigma_{1,BS}^2\tau & \rho\sigma_{1,BS}\sigma_{2,BS}\tau \\ \rho\sigma_{1,BS}\sigma_{2,BS}\tau & \sigma_{2,BS}^2\tau \end{bmatrix}.$$

Similarly to the domain truncation in the one-dimensional numerical experiment from Subsection 2.3, the spatial domain is truncated to $[S_{1,\min}, S_{1,\max}] \times [S_{2,\min}, S_{2,\max}] = [0, 30S_{1,0}] \times [0, 30S_{2,0}]$. The Cartesian spatial grid,

$$(s_{1,i}, s_{2,j}) \quad \text{for } 1 \leq i \leq m_1, 1 \leq j \leq m_2,$$

is constructed by considering the spatial grid from Subsection 2.3 in both spatial dimension. In Figure 6 this spatial grid is illustrated within the region $[0, 5S_{1,0}] \times [0, 5S_{2,0}]$ for $S_{1,0} = S_{2,0} = 100$ and $m_1 = m_2 = 50$. The FV discretization from Subsection 2.4 then defines approximations

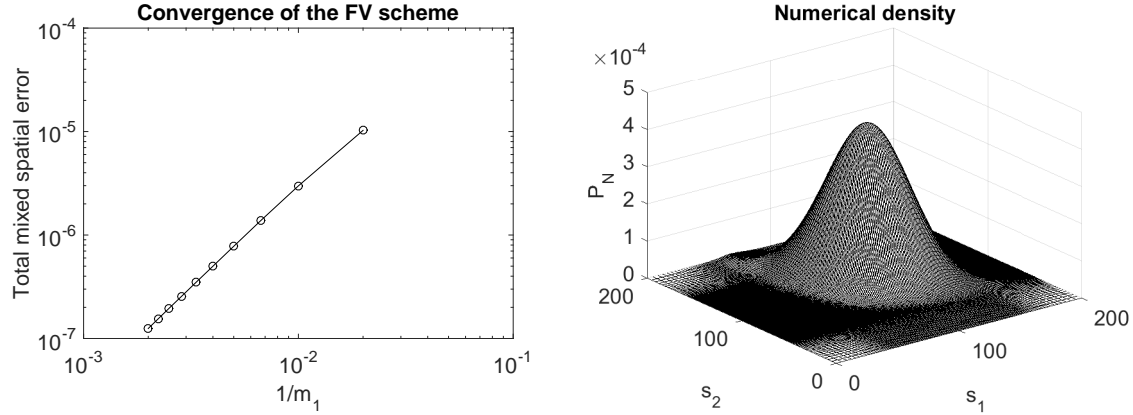


Figure 7: Convergence results within the 2D Black-Scholes model. The parameter values are $r = 0.03, \sigma_{1,BS} = 0.2, \sigma_{2,BS} = 0.25, \rho = -0.7, T = 1$.

$\mathbf{P}_{i,j}(\tau)$ to $p(s_{1,i}, s_{2,j}, \tau)$ and we compute the total mixed spatial error

$$\max_{1 \leq i \leq m_1, 1 \leq j \leq m_2} \epsilon_{i,j}(m),$$

where

$$\epsilon_{i,j}(m) = \begin{cases} \left| \frac{p(s_{1,i}, s_{2,j}, T) - \mathbf{P}_{i,j}(T)}{p(s_{1,i}, s_{2,j}, T)} \right| & \text{if } p(s_{1,i}, s_{2,j}, T) > 1, \\ |p(s_{1,i}, s_{2,j}, T) - \mathbf{P}_{i,j}(T)| & \text{else.} \end{cases}$$

The values $\mathbf{P}_{i,j}(T)$ are calculated by applying the Hundsdorfer-Verwer time stepping method (see Section 3) with a large number of time steps such that the temporal discretization error is negligible. In Figure 7 the total mixed spatial error is shown for the realistic parameter values $r = 0.03, \sigma_{1,BS} = 0.2, \sigma_{2,BS} = 0.25, \rho = -0.7, T = 1$ and for the number of spatial grid points $m_1 = m_2 = \{50, 100, \dots, 500\}$. The convergence plot indicates that the FV discretization is second-order convergent with respect to the current initial-boundary value problem.

For the second example we consider the popular Heston model [13], i.e. the SV model (1.3) with $\psi(v) = \sqrt{v}$ and $\alpha = 1/2$. The underlying density function satisfies the forward Kolmogorov equation

$$\frac{\partial}{\partial \tau} p = \frac{\partial^2}{\partial x^2} \left(\frac{1}{2} v p \right) + \frac{\partial^2}{\partial x \partial v} (\rho \xi v p) + \frac{\partial^2}{\partial v^2} \left(\frac{1}{2} \xi^2 v p \right) - \frac{\partial}{\partial x} \left((r_d - r_f - \frac{1}{2} v) p \right) - \frac{\partial}{\partial v} (\kappa (\eta - v) p), \quad (2.28)$$

for $x \in \mathbb{R}, v > 0, \tau > 0$ and with initial condition $p(x, v, 0) = \delta(x) \delta(v - V_0)$ where $X_0 = 0$ and V_0 is given. To the best of our knowledge, no analytical solution is available in the literature for the density function p that satisfies (2.28). In order to test the performance of our FV discretization, we compute a reference solution with an alternative discretization method described in [7]. The latter method is based on rewriting

$$p(x, v, \tau) = p_1(x, \tau | V_{SV, \tau} = v) p_2(v, \tau), \quad (2.29)$$

where $p_2(v, \tau)$ denotes the one-dimensional density of the volatility process, and $p_1(x, \tau | V_{SV, \tau} = v)$ denotes the conditional density of X_τ given the variance value $V_{SV, \tau} = v$. In [7] the characteristic function

$$\psi(\omega | V_{SV, \tau} = v) = \mathbb{E}[\exp(\mathbf{i} \omega X_{SV, \tau}) | V_{SV, \tau} = v]$$

corresponding with $p_1(x, \tau | V_{SV, \tau} = v)$ is given in semi-analytical form and it is stated that $p_2(v, \tau)$ is given by (2.18). By using the COS-method [6] we approximate the conditional density function $p_1(x, \tau | V_{SV, \tau} = v)$ and our reference solution p_{ref} is then defined via (2.29).

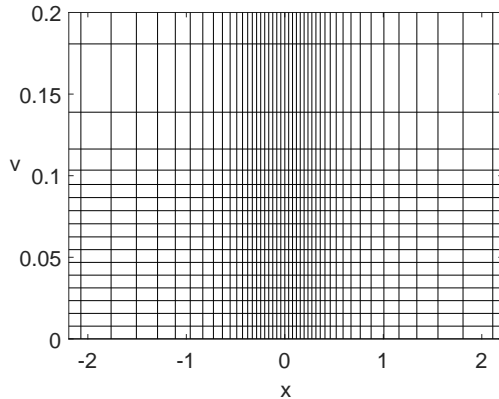


Figure 8: Illustration of the non-uniform grid around (X_0, V_0) for the Heston example and the actual values $X_0 = 0, V_0 = 0.0625, m_1 = 2m_2 = 50$.

In the numerical experiment, we opt to truncate the domain in the x -direction to $[X_0 - \log(30), X_0 + \log(30)]$. The spatial domain in the v -direction is truncated to $[0, 15]$, analogously to the CIR example. We consider spatial grids

$$-\log(30) = x_1 < x_2 < \dots < x_{m_1} = \log(30),$$

$$0 = v_1 < v_2 < \dots < v_{m_2} = 15,$$

with $m_1 = 2m_2$, which are similar to the ones described in [10] and such that there exist indices i_0, j_0 such that $(x_{i_0}, v_{j_0}) = (X_0, V_0)$. In Figure 8 the spatial grid is shown for $X_0 = 0, V_0 = 0.0625$ and the sample value $m_1 = 2m_2 = 50$. Applying the FV discretization then yields approximations $\mathbf{P}_{i,j}(\tau)$ to $p(x_i, v_j, \tau)$. From equation (2.29) and the CIR example it follows that if $q = \frac{2\kappa\eta}{\xi^2} - 1 < 0$, then the density function can tend to infinity as v tends to zero and at $v = 0$ the exact density is not defined. Analogously to the remark in Subsection 2.3, it is readily seen that adequately comparing the difference between $p(x_i, v_2, T)$ and $\mathbf{P}_{i,2}(T)$ then becomes difficult for increasing values of m_2 . The error is therefore computed on similar spatial domains. Let v_{low} again be the smallest non-zero grid point in the v -direction if the total number of grid points in that direction is $m_2 = 50$. For given m_2 , let j_1 be the lowest index such that $v_{j_1} \geq v_{low}$ and define the total mixed spatial error by

$$\max_{1 \leq i \leq m_1, j_1 \leq j \leq m_2} \epsilon_{i,j}(m),$$

where

$$\epsilon_{i,j}(m) = \begin{cases} \left| \frac{p_{ref}(x_i, v_j, T) - \mathbf{P}_{i,j}(T)}{p_{ref}(x_i, v_j, T)} \right| & \text{if } p_{ref}(x_i, v_j, T) > 1, \\ |p_{ref}(x_i, v_j, T) - \mathbf{P}_{i,j}(T)| & \text{else.} \end{cases}$$

The values $\mathbf{P}_{i,j}(T)$ are once more approximated by applying the Hundsdorfer–Verwer time stepping method (see Section 3) with a small time step such that the temporal discretization error is negligible.

For the actual numerical experiments we consider an extension of the two sets of parameters used in Subsection 2.3. The extensions are also taken from [7], used in [26], and given by

	κ	η	ξ	ρ	r_d	r_f	X_0	V_0	T
Set C	5	0.16	0.9	0.1	0.1	0	0	0.0625	0.25
Set D	1.15	0.0348	0.39	-0.64	0.04	0	0	0.0348	0.25

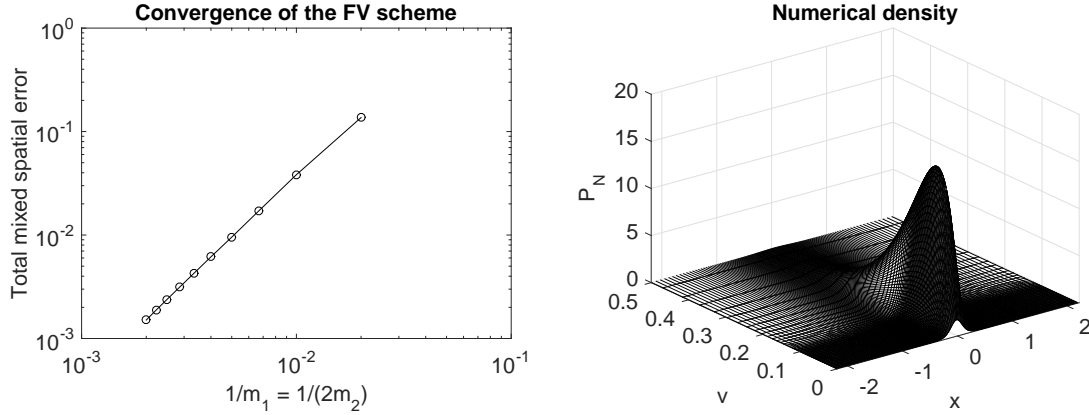


Figure 9: Convergence results within the Heston model. The parameters are given by Set C.

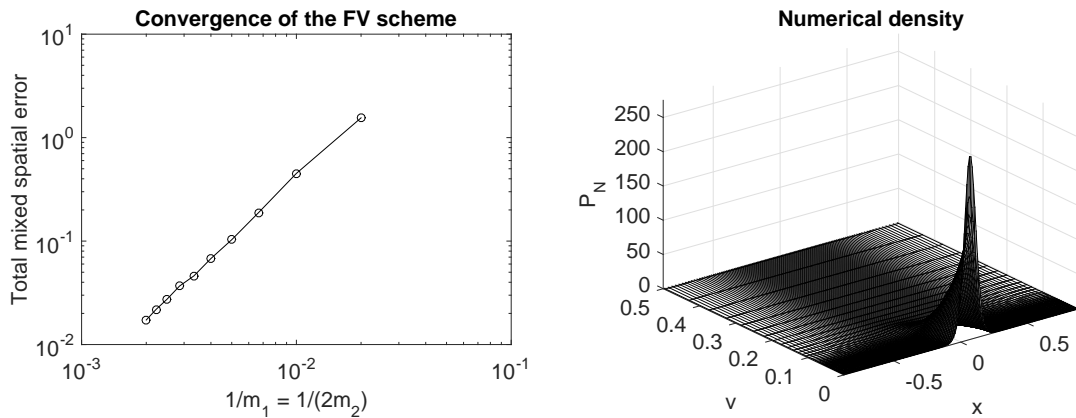


Figure 10: Convergence results within the Heston model. The parameters are given by Set D.

Recall that for Set C we have that $q = 0.98$ and for Set D it holds that $q = -0.47$. In the left plot of Figure 9, respectively Figure 10, the total mixed spatial error is shown for the parameters of Set C, respectively Set D, and for the number of spatial grid points $m_1 = 2m_2 = \{50, 100, \dots, 500\}$. In the right plots, the corresponding numerical solutions are shown for $m_1 = 2m_2 = 200$. The convergence plots indicate that the FV discretization is convergent with respect to the current initial-boundary value problems. Additional experiments again suggest that the FV discretization is second-order convergent if the Feller condition is satisfied. If $q < 0$ the order of convergence can drop to one. Please note that the conclusions of the numerical experiments are essentially unchanged for different values of v_{low} as long as it is defined via one of the coarsest grids considered in the experiment. The two-dimensional tests also confirm that the total numerical mass is, indeed, kept constantly equal to one, even if the Feller condition is strongly violated.

3 Temporal discretization by ADI methods

In general, spatial discretization by FV methods of initial-boundary value problems for time-dependent convection diffusion equations leads to large systems of stiff ODEs of the type

$$W'(\tau) = F(\tau, W(\tau)) \quad (0 \leq \tau \leq T), \quad W(0) = W_0,$$

with given vector-valued function $F : [0, T] \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ and given vector $W_0 \in \mathbb{R}^m$ where m is the number of volumes. When the function F is stemming from semidiscretization of a multidimensional PDE, classical implicit time stepping schemes such as the *Crank–Nicolson scheme* can often be very time consuming. For the effective time discretization of such semidiscrete systems, operator splitting schemes of the Alternating Direction Implicit (ADI) type are widely considered. In the past decades, several ADI schemes have been studied for the situation where mixed spatial derivatives are present in the convection-diffusion equation. Mixed derivative terms are ubiquitous, notably, in the field of financial mathematics. There they arise due to correlations between the underlying stochastic processes, cf. the general SLV model (1.1).

In this paper we consider the state-of-the-art *Hundsdofer–Verwer (HV) scheme*, [19, 31], which is of the ADI-type. Suppose the semidiscrete operator F is stemming from the semidiscretization of a two-dimensional convection-diffusion equation and suppose F can be decomposed as

$$F(\tau, w) = F_0(\tau, w) + F_1(\tau, w) + F_2(\tau, w) \quad (0 \leq \tau \leq T, w \in \mathbb{R}^m),$$

where F_0 represents the mixed spatial derivative term and F_1 , respectively F_2 , represents all spatial derivative terms in the first, respectively second, spatial direction. Let $\theta > 0$ be a given parameter, $N \geq 1$ the number of time steps and $\tau_n = n\Delta\tau$ with $\Delta\tau = T/N$. Then the HV scheme defines approximation W_n to $W(\tau_n)$ successively for $n = 1, 2, 3, \dots, N$ through

$$\left\{ \begin{array}{l} Y_0 = W_{n-1} + \Delta\tau F(\tau_{n-1}, W_{n-1}), \\ Y_l = Y_{l-1} + \theta\Delta\tau (F_l(\tau_n, Y_l) - F_l(\tau_{n-1}, W_{n-1})), \quad l = 1, 2, \\ \tilde{Y}_0 = Y_0 + \frac{1}{2}\Delta\tau (F(\tau_n, Y_2) - F(\tau_{n-1}, W_{n-1})), \\ \tilde{Y}_l = \tilde{Y}_{l-1} + \theta\Delta\tau (F_l(\tau_n, \tilde{Y}_l) - F_l(\tau_n, Y_2)), \quad l = 1, 2, \\ W_n = \tilde{Y}_2. \end{array} \right. \quad (3.1)$$

The HV scheme (3.1) starts with an explicit Euler predictor stage, followed by two implicit but unidirectional corrector stages. Then a second explicit stage is performed, followed again by two implicit unidirectional corrector stages. The operator F_0 , which contains the mixed spatial derivative term, is always treated in an explicit manner. The implicit stages only handle spatial derivatives in only one spatial dimension. This can lead to a major computational advantage in comparison to classical non-splitted implicit time stepping methods.

Recently, various positive stability results for the HV scheme have been derived relevant to multidimensional convection-diffusion equations with mixed derivative term, see e.g. [15, 16, 17]. Subsequently, it has been proved [18] that, under some natural stability and smoothness assumptions, the HV scheme is second-order convergent with respect to the time step whenever it is applied to semidiscrete two-dimensional convection-diffusion equations with mixed derivative term. The temporal convergence result from [18] has the key property that it holds uniformly in the spatial mesh width.

In this article, the vector $P(0)$ is stemming from an initial function that is nonsmooth. It is well-known that convergence of time discretization methods can then be seriously impaired, cf. [22, 32]. A convergence analysis for the HV scheme relevant to nonsmooth data is still open in the literature. In order to prevent the numerical solution from undesirable behaviour, Rannacher time stepping will be applied, that is, the first few time steps of the HV scheme are replaced by twice as many half time steps of the implicit Euler scheme [23]. Numerical experiments in [10, 14] indicate that the HV scheme then remains second-order convergent. We opt to replace here the first two HV time steps by four half time steps of the implicit Euler scheme. This choice is based on the results in [32] for the Modified Craig–Sneyd scheme which forms another popular ADI scheme for multidimensional convection-diffusion equations with mixed derivative terms. Note that applying Rannacher time stepping to systems of ODEs stemming from multidimensional

convection-diffusion equations can be time consuming if exact solvers are used. In order to limit this drawback, the implicit Euler time steps are performed with a suitable multigrid solver, [12].

4 Calibration of SLV models to the LV model

As stated in Section 1, the goal of the paper is to calibrate state-of-the-art SLV models to its underlying LV model in order to reproduce the known prices for European call and put options. This is done by defining the leverage function σ_{SLV} such that the relationship (1.5) is satisfied. By combining equations (1.5), (1.6) and (1.7), it is readily seen that a highly non-linear PDE needs to be solved. In this section the FV-ADI method is used in combination with an inner iteration to approximate the corresponding leverage function and density function.

In order to use the FV-ADI discretization, one first has to define spatial and temporal grids. Since the initial function of forward Kolmogorov equations is highly non-smooth around the spot value (X_0, V_0) , and the region of interest is also situated around this value, it is natural to consider non-uniform Cartesian grids that are concentrated around the value (X_0, V_0) . If the parameter α from the SLV model is chosen smaller than or equal to $1/2$, then the natural boundary $V_\tau = 0$ can be reached and probability mass stacks up at the boundary $v = 0$. It is then natural to additionally require smaller mesh widths in the v -direction at this boundary. The non-uniform grids define volumes of which the volume average is approximated by the FV scheme. Denote by m_1 , respectively m_2 , the number of spatial grid points in the x -direction, respectively v -direction. We consider spatial grids

$$\begin{aligned} x_{\min} &= x_1 < x_2 < \dots < x_{m_1} = x_{\max}, \\ 0 &= v_1 < v_2 < \dots < v_{m_2} = V_{\max}, \end{aligned}$$

which are similar to the ones described in [10] and such that there exist indices i_0, j_0 such that $(x_{i_0}, v_{j_0}) = (X_0, V_0)$. Denote their corresponding mesh widths by $\Delta x_i, \Delta v_j$ and define volumes

$$\Omega_{i,j} = [x_{i-0.5}, x_{i+0.5}] \times [v_{j-0.5}, v_{j+0.5}],$$

where the values $x_{i-0.5}, v_{j-0.5}$ are defined similar as in Section 2. The spatial grids are the result of a continuous transformation of a uniform underlying grid and it can be shown that the pertinent meshes are smooth. The values $x_{\min}, x_{\max}, V_{\max}$ are chosen sufficiently far away from the spot value such that the boundary conditions from Section 2 can be applied. An example of the spatial grid for the small sample values $m_1 = 2m_2 = 50$ was already shown in Figure 8 for the case where $(X_0, V_0) = (0, 0.0348)$. For the discretization in time we always consider uniform grids $\tau_n = n\Delta\tau$ where the step size is given by $\Delta\tau = T/N$ and N denotes the total number of time steps.

Once the spatial grid and corresponding volumes are defined, the FV discretization from Section 2 is applied. This yields a large system of ordinary differential equations

$$P'(\tau) = A(\tau)P(\tau) = (A_0(\tau) + A_1(\tau) + A_2(\tau))P(\tau) \quad (0 \leq \tau \leq T), \quad (4.1)$$

with given matrices $A_0(\tau), A_1(\tau), A_2(\tau)$ and initial function defined via

$$P_{i,j}(0) = \begin{cases} \frac{2}{\Delta x_{i-1} + \Delta x_i} \frac{2}{\Delta v_{j-1} + \Delta v_j} & \text{if } (i, j) = (i_0, j_0), \\ 0 & \text{else.} \end{cases}$$

The matrices A_0, A_1 contain, however, the unknown function σ_{SLV} at the spatial grid points. The $P_i(\tau)$ can be used in combination with a numerical integration technique to approximate the conditional expectations (1.6) and hence the pertinent leverage function. We opt to perform the numerical integration with the trapezoidal rule and define approximations

$$\mathbb{E}_i(\tau) = \frac{\sum_{j=1}^{m_2} \psi^2(v_j) P_{i,j}(\tau) \frac{\Delta v_j + \Delta v_{j+1}}{2}}{\sum_{j=1}^{m_2} P_{i,j}(\tau) \frac{\Delta v_j + \Delta v_{j+1}}{2}} \approx \mathbb{E}[\psi^2(V_\tau) | X_\tau = x_i], \quad (4.2)$$

where we recall that $\Delta v_1 = \Delta v_{m_2+1} = 0$. Inserting the approximations (4.2) into (4.1) leads to a non-linear system of ODEs.

As a final step, the system of ODEs (4.1) is discretized in time with the time stepping scheme from Section 3 and an inner iteration to handle the non-linearity, cf. [28]. By applying the HV scheme, the conditional expectations (4.2) are naturally replaced by their fully discrete versions

$$\mathbb{E}_{n,i} = \frac{\sum_{j=1}^{m_2} \psi^2(v_j) \mathbf{P}_{n,i,j} \frac{\Delta v_j + \Delta v_{j+1}}{2}}{\sum_{j=1}^{m_2} \mathbf{P}_{n,i,j} \frac{\Delta v_j + \Delta v_{j+1}}{2}}, \quad (4.3)$$

and we define the leverage function σ_{SLV} at the spatial and temporal grid by

$$\sigma_{SLV}(x_i, \tau_n) = \frac{\sigma_{LV}(x_i, \tau_n)}{\sqrt{\mathbb{E}_{n,i}}}. \quad (4.4)$$

It is readily seen that at the initial time level, i.e. at $n = 0$, the expression (4.3) is only defined if $i = i_0$. To render the calibration procedure feasible we put

$$\mathbb{E}_{0,i} = \psi^2(V_0) \quad \text{for } 1 \leq i \leq m_1.$$

For strictly positive time levels, i.e. for $n > 0$, the exact density $p(x_i, v_j, \tau_n)$ is always non-negative. By performing the spatial and temporal discretization, however, it is possible that some of the values $\mathbf{P}_{n,i,j}$ become (slightly) negative. In order to prevent the numerical solution from undesirable behaviour, the expression (4.3) is replaced in the calibration procedure by

$$\mathbb{E}_{n,i} = \frac{\sum_{j=1}^{m_2} \psi^2(v_j) |\mathbf{P}_{n,i,j}| \frac{\Delta v_j + \Delta v_{j+1}}{2}}{\sum_{j=1}^{m_2} |\mathbf{P}_{n,i,j}| \frac{\Delta v_j + \Delta v_{j+1}}{2}} \quad \text{for } 1 \leq i \leq m_1, \quad n > 0. \quad (4.5)$$

Theoretically it is possible that the denominator (and hence also the nominator) of (4.5) equals zero and the fully discrete conditional expectation is undefined. In this case we assume that the conditional expectation is locally constant in time and set $\mathbb{E}_{n,i} = \mathbb{E}_{n-1,i}$. Let $Q \geq 1$ be a given integer. For the actual calibration of the SLV model to the LV model, we employ the following numerical procedure.

for n is 1 to N do

 let $P_n = P_{n-1}$ be an initial approximation to $P(\tau_n)$;

 for q is 1 to Q do

 (a) approximate the conditional expectations $\mathbb{E}[\psi^2(V_{\tau_n}) | X_{\tau_n} = x_i]$ by (4.5);

 (b) Define $\sigma_{SLV}(\cdot, \tau_n)$ on the grid in the x -direction by formula (4.4);

 (c) update P_n by performing a numerical time step for (4.1) from τ_{n-1} to τ_n ;

 end

end

Whenever a time step from τ_{n-1} to τ_n with the HV scheme is replaced by two half-time steps of the implicit Euler scheme, the inner iteration above is first performed for the substep from τ_{n-1} to $\tau_{n-1/2} = \tau_{n-1} + \Delta\tau/2$, yielding an approximation of $P(\tau_n - \Delta\tau/2)$ and $\sigma_{SLV}(\cdot, \tau_n - \Delta\tau/2)$. Next, the inner iteration is performed for the substep from $\tau_n - \Delta\tau/2$ to τ_n , yielding an approximation of $P(\tau_n)$ and $\sigma_{SLV}(\cdot, \tau_n)$. Upon completion of the time stepping and iteration procedure above, the original approximation for $\sigma_{SLV}(\cdot, 0)$ is replaced on the grid in the x -direction by $\sigma_{SLV}(\cdot, \tau_1)$. This appears more realistic as the original approximation was actually only valid for the index $i = i_0$.

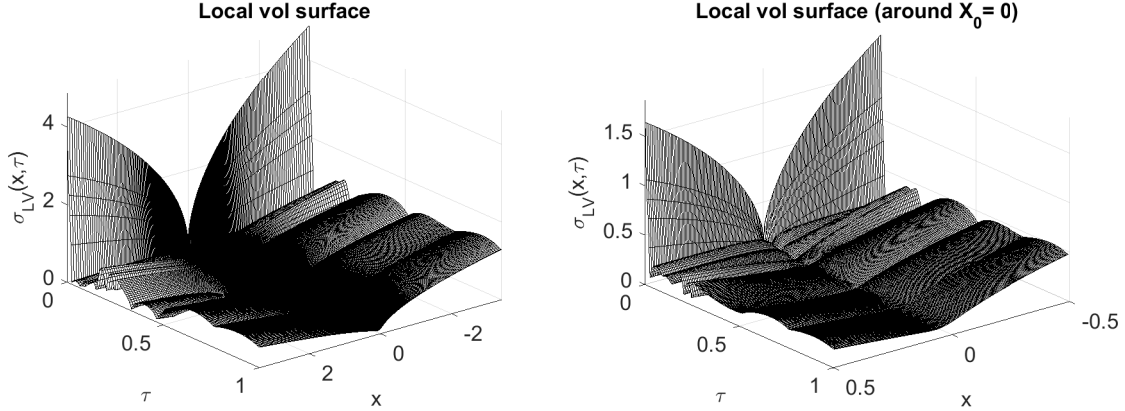


Figure 11: Local volatility function originating from actual EUR/USD vanilla option data (market data as of 2 March 2016). The spot rate $S_0 = 1.08815$.

5 Numerical experiments on the calibration of SLV models

In this section, the effectiveness of the calibration procedure is illustrated by applying it to a practical example. Here, we opt to consider the popular and challenging Heston-based SLV model, i.e. SLV model (1.1) with $\psi(v) = \sqrt{v}$ and $\alpha = 1/2$, to describe the evolution of the EUR/USD exchange rate.

As stated in the introduction, in financial practice it is common to first determine the SV parameters of the underlying SV model and to define the LV function such that the LV model reproduces the known market prices for European call and put options. Afterwards, the calibration procedure aims at matching the SLV model with its underlying LV model, i.e. at obtaining equality (1.4). In this article, we assume that the SV parameters and the LV function are known and we then apply the calibration procedure from Section 4. The performance is illustrated by comparing the densities from the LV model and from the SLV model and by comparing the corresponding option values.

For the actual experiments we consider the following sets of SV parameters:

	κ	η	ξ	ρ	T	V_0
Set E	5	0.16	0.9	0.1	0.25	0.0625
Set F	1.15	0.0348	0.39	-0.64	0.25	0.0348
Set G	1.50	0.0154	0.24	-0.11	1	0.0154

The Sets E and F correspond with the SV parameters from Sets C and D, and are taken from [7]. Set G is taken from [2] and corresponds to the EUR/USD exchange rate for the pertinent maturity (market data as of 16 September 2008). For Set E it holds that $q = \frac{2\kappa\eta}{\xi^2} - 1 = 0.98$ and the process V_τ is strictly positive. For Set F, respectively Set G, it holds that $q = -0.47$, respectively $q = -0.20$, such that $V_\tau = 0$ is attainable. The LV model is completely determined by the LV function, the risk-free interest rates and the spot value S_0 . We assume that the risk-free interest rates are given by

$$r_d = 0.02, \quad r_f = 0.01,$$

and that the LV function is as displayed in Figure 11. The pertinent LV function originates from actual EUR/USD vanilla option data (market data as of 2 March 2016) and is constructed by using an SSVI-type method for implied volatility interpolation, [8]. The corresponding spot rate is given by

$$S_0 = 1.08815.$$

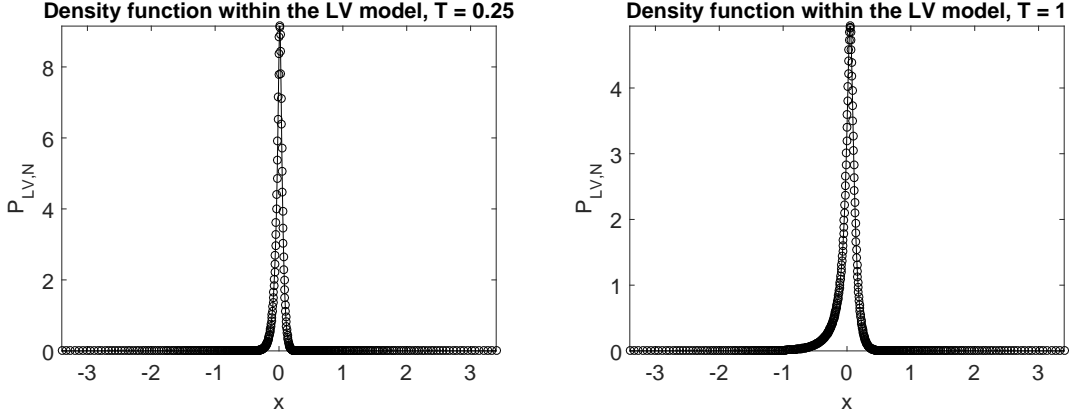


Figure 12: Approximation of the density function $p_{LV}(x, 0.25)$ (left) and $p_{LV}(x, 1)$ (right) by applying the FV discretization from Subsection 2.2 with $m = 400$.

Note that, even if the LV function is given, there is often no analytical expression available for the density function p_{LV} or the option values. It is well-known, however, that the density function satisfies the 1D forward Kolmogorov equation

$$\frac{\partial}{\partial \tau} p_{LV} = \frac{\partial}{\partial x^2} \left(\frac{1}{2} \sigma_{LV}^2 p_{LV} \right) - \frac{\partial}{\partial x} \left((r_d - r_f - \frac{1}{2} \sigma_{LV}^2) p_{LV} \right),$$

for $x \in \mathbb{R}, \tau > 0$. By applying the FV discretization described in Subsection 2.2 one defines approximations $P_{LV,i}(\tau)$ of the exact density values $p_{LV}(x_i, \tau)$. Fully discrete approximations $P_{LV,N,i}$ of $p_{LV}(x_i, T)$ are then obtained by applying a suitable time stepping method. In Figure 12 the latter approximations are shown for $\tau = 0.25$, respectively $\tau = 1$, and the practical value $m = 400$.

Once the underlying SV model and LV model are specified, the calibration procedure from Section 4 can be applied. For the actual experiments we consider the discretization parameters

$$m_1 = 400, \quad m_2 = 200, \quad \Delta\tau = 1/200, \quad \theta = \frac{1}{2} + \frac{1}{6}\sqrt{3}, \quad Q = 2.$$

In Figure 13 the resulting discrete leverage function is shown for set G. In order to illustrate the performance of the calibration, we first consider a discrete version of (1.4). More precisely, the discrete numerical densities $P_{LV,N,i}$ from the LV model are compared with

$$P_{SLV,N,i} := \sum_{j=1}^{m_2} P_{N,i,j} \frac{\Delta v_j + \Delta v_{j+1}}{2},$$

which can be seen as the fully discrete approximations of

$$\int_0^\infty p(x_i, v, T) dv,$$

within the SLV model after applying the trapezoidal rule for numerical integration. In the left plots of Figure 14 the approximations $P_{SLV,N}$ are shown for each of the three sets of parameters. In the right plots, the corresponding differences $P_{LV,N} - P_{SLV,N}$ are plotted. Note that the final time $T = 0.25$ for Set E and Set F, and $T = 1$ for Set G. From Figure 14 it is readily seen that the difference between the fully discrete numerical densities is very small and hence that the calibration procedure performs well.

The main goal of the calibration procedure is to define the leverage function in such a way that the LV model and the SLV model define the same fair values for non-path-dependent European

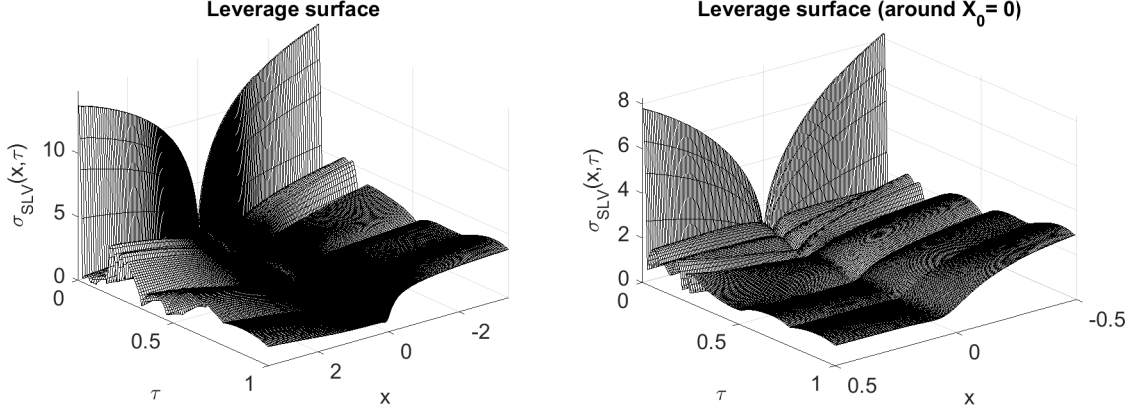


Figure 13: Leverage function stemming from the calibration procedure with local volatility function from Figure 11, SV parameters from Set G and values $m_1 = 400$, $m_2 = 200$, $\Delta\tau = 1/200$, $\theta = \frac{1}{2} + \frac{1}{6}\sqrt{3}$, $Q = 2$.

options. If the leverage function is defined by (1.5), then it follows for the fair value (FaV) of such an option with payoff $u_0(x)$, $x \in \mathbb{R}$, at maturity T that

$$\text{FaV} = e^{-raT} \int_{-\infty}^{\infty} p_{LV}(x, T) u_0(x) dx = e^{-raT} \int_0^{\infty} \int_{-\infty}^{\infty} p(x, v, T) u_0(x) dx dv.$$

Given the approximations $P_{LV, N}$ and \mathbf{P}_N , the pertinent FaV can easily be approximated by applying numerical integration with the trapezoidal rule. In case of the SLV model, it is readily seen that defining the approximated FaV via \mathbf{P}_N and the trapezoidal rule is equivalent to defining the FaV via $P_{SLV, N}$ and the trapezoidal rule. Denote by FaV_{LV} , respectively FaV_{SLV} , the approximated fair values obtained via $P_{LV, N}$, respectively $P_{SLV, N}$. We now compare these approximations for a set of European call options with a range of strikes given by

$$K = 0.75S_0, 0.8S_0, 0.9S_0, S_0, 1.1S_0, 1.2S_0, 1.25S_0.$$

When the strike increases relatively to S_0 , the fair value of European call options tends to zero and it is difficult to adequately compare approximations. In financial practice, European call and put options are often quoted in terms of *implied volatility*. Let $\sigma_{imp, LV}$, respectively $\sigma_{imp, SLV}$, denote the implied volatility (in %) corresponding to FaV_{LV} , respectively FaV_{SLV} . In the following we test the performance of the calibration procedure by calculating the absolute implied volatility errors

$$\epsilon_{imp} = |\sigma_{imp, LV} - \sigma_{imp, SLV}|.$$

In Table 1 these errors are presented for the different SV parameter sets, taking the same values of $m, m_1, m_2, \Delta\tau, \theta, Q$ as above. The somewhat larger values ϵ_{imp} for $T = 0.25$ compared to $T = 1$ can be explained from the fact that the implied volatility is more sensitive to changes in the fair value when the maturity is low. The results in Table 1 confirm that the calibration procedure performs well. They indicate that the fully discrete leverage surface is, indeed, defined such that the SLV model reproduces accurately the known market prices for European call options.

6 Conclusion

Stochastic local volatility (SLV) models constitute state-of-the-art models to describe asset price processes. Their calibration to the underlying local volatility model is, however, highly non-trivial. It incorporates the solution of non-linear forward Kolmogorov equations. In general, no analytical

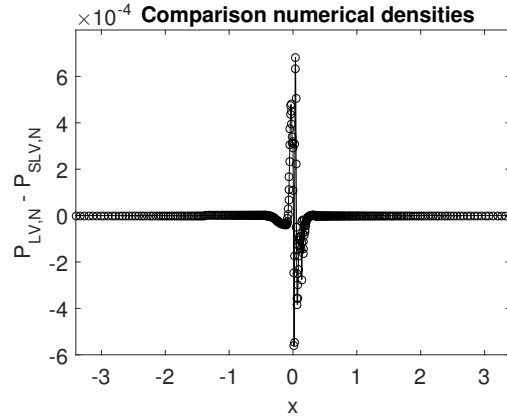
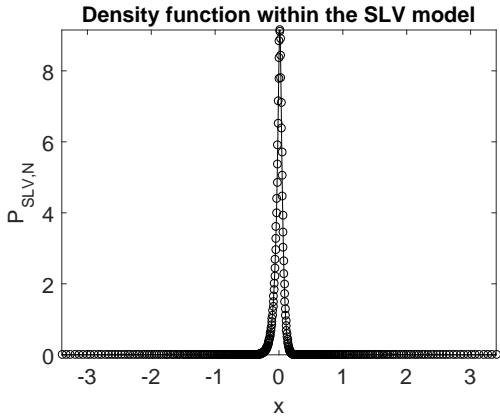


Illustration of the density in case of Set E.





Illustration of the density in case of Set F.

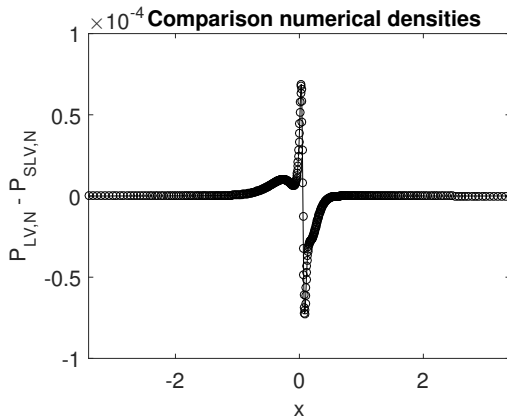
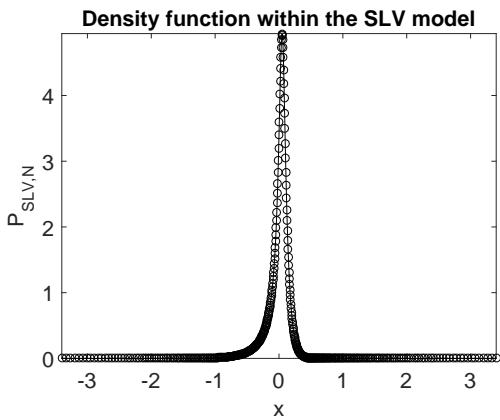


Illustration of the density in case of Set G.

Figure 14: Comparison of the fully discrete density functions $P_{LV,N}$ and $P_{SLV,N}$ for each of the parameters sets and for values $m_1 = 400$, $m_2 = 200$, $\Delta\tau = 1/200$, $\theta = \frac{1}{2} + \frac{1}{6}\sqrt{3}$, $Q = 2$.

	$T = 0.25$	Set E	Set F	$T = 1$	Set G
K/S_0	$\sigma_{imp,LV}$	ϵ_{imp}	ϵ_{imp}	$\sigma_{imp,LV}$	ϵ_{imp}
0.75	19.18	0.1005	0.1208	21.94	0.0021
0.80	18.40	0.0212	0.0454	20.20	0.0015
0.90	15.01	0.0033	0.0154	16.65	0.0008
1.0	11.26	0.0011	0.0030	13.14	0.0004
1.10	11.59	0.0011	0.0153	11.38	0.0003
1.20	13.20	0.0009	0.0937	11.77	0.0003
1.25	14.03	0.0006	0.1888	12.12	0.0003

Table 1: Comparison of the approximated implied volatilities $\sigma_{imp,LV}$ and $\sigma_{imp,SLV}$ for values $m_1 = 400$, $m_2 = 200$, $\Delta\tau = 1/200$, $\theta = \frac{1}{2} + \frac{1}{6}\sqrt{3}$, $Q = 2$.

solution is available and one relies on numerical methods in order to approximate the exact solution. Here, we introduce a finite volume - alternating direction implicit method for the numerical solution of general one-dimensional and two-dimensional forward Kolmogorov equations. The finite volume spatial discretization does not require a transformation of the PDE, which constitutes a main advantage in the calibration of SLV models, and handles the boundary conditions in a natural way. Moreover, the finite volume scheme preserves the crucial property that total mass of a density function is always equal to one. Our numerical experiments for relevant practical applications confirm that the pertinent spatial discretization is convergent. Temporal discretization is performed by using the Hundsdorfer–Verwer ADI scheme. By splitting the semidiscrete system into different parts that represent spatial derivatives in the different spatial dimensions, and by handling spatial derivatives in only one spatial dimension in the implicit substeps, a major computational advantage can be achieved. The non-linearity in the calibration procedure of stochastic local volatility models is handled by introducing an inner iteration. Our numerical experiments reveal that the proposed calibration procedure performs well. The calibrated stochastic local volatility model matches the underlying local volatility model almost exactly, both in terms of the density function and of the implied volatilities of European call options.

Funding

The work by the first author has been supported financially by a PhD Fellowship of the Research Foundation–Flanders.

References

- [1] L. B. G. Andersen and V. V. Piterbarg, *Moment explosions in stochastic volatility models*, Finance Stoch. **11** (2007), pp. 29–50.
- [2] I. J. Clark, *Foreign Exchange Option Pricing: A Practitioner’s Guide*, John Wiley & Sons, 2011.
- [3] J. C. Cox, J. E. Ingersoll and S. A. Ross, *A theory of the term structure of interest rates*, Econometrica **53** (1985), pp. 385–407.
- [4] B. Dupire, *Pricing with a smile*, RISK, January (1994), pp. 18–20.
- [5] B. Engelmann, F. Koster and D. Oeltz, *Calibration of the Heston stochastic local volatility model: a finite volume scheme*, Available at SSRN 1823769 (2012).
- [6] F. Fang and C. W. Oosterlee, *A novel pricing method for European options based on Fourier-cosine series expansions*, SIAM J. Sci. Comput. **31** (2008), pp. 826–848.

- [7] F. Fang and C. W. Oosterlee, *A Fourier-based valuation method for Bermudan and barrier option under Heston's model*, SIAM J. Financial Math. **2** (2011), pp. 439–463.
- [8] J. Gatheral and A. Jacquier, *Arbitrage-free SVI volatility surfaces*, Quant. Financ. **14** (2014), pp. 59–71.
- [9] I. Gyöngy, *Mimicking the one-dimensional marginal distributions of processes having an Ito differential*, Probab. Th. Rel. Fields **71** (1986), pp. 501–516.
- [10] T. Haentjens and K. J. in 't Hout, *Alternating direction implicit finite difference schemes for the Heston–Hull–White partial differential equation*, J. Comp. Finan. **16** (2012), pp. 83–110.
- [11] P. Henry-Labordère, *Calibration of local stochastic volatility models to market smiles*, Risk, September (2009), pp. 112–117.
- [12] V. E. Henson and U. M. Yang, *BoomerAMG: A parallel algebraic multigrid solver and preconditioner*, Appl. Numer. Math. **41** (2002), pp. 155–177.
- [13] S. L. Heston, *A closed-form solution for options with stochastic volatility with applications to bond and currency options*, Rev. Finan. Stud. **6** (1993), pp. 327–343.
- [14] K. J. in 't Hout and S. Foulon, *ADI finite difference schemes for option pricing in the Heston model with correlation*, Int. J. Numer. Anal. Mod. **7** (2010), pp. 303–320.
- [15] K. J. in 't Hout and C. Mishra, *Stability of ADI schemes for multidimensional diffusion equations with mixed derivative terms*, Appl. Numer. Math. **74** (2013), pp. 83–94.
- [16] K. J. in 't Hout and B. D. Welfert, *Stability of ADI schemes applied to convection-diffusion equations with mixed derivative terms*, Appl. Numer. Math. **57** (2007), pp. 19–35.
- [17] K. J. in 't Hout and B. D. Welfert, *Unconditional stability of second-order ADI schemes applied to multi-dimensional diffusion equations with mixed derivative terms*, Appl. Numer. Math. **59** (2009), pp. 677–692.
- [18] K. J. in 't Hout and M. Wyns, *Convergence of the Hundsdorfer–Verwer scheme for two-dimensional convection-diffusion equations with mixed derivative term*, AIP Conf. Proc. **1648** (2015), 850054.
- [19] W. Hundsdorfer, *Accuracy and stability of splitting with Stabilizing Corrections*, Appl. Numer. Math. **42** (2003), pp. 213–233.
- [20] W. Hundsdorfer and J. G. Verwer, *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*, Springer, Berlin, 2003.
- [21] A. Lipton, *The vol smile problem*, RISK, February (2002), pp. 61–65.
- [22] D. M. Pooley, K. R. Vetzal and P. A. Forsyth, *Convergence remedies for non-smooth payoffs in option pricing*, J. Comp. Finan. **6** (2003), pp. 25–40.
- [23] R. Rannacher, *Finite element solution of diffusion problems with irregular data*, Numer. Math. **43** (1984), pp. 309–327.
- [24] Y. Ren, D. Madan and M. Q. Qian, *Calibrating and pricing with embedded local volatility models*, Risk, September (2007), pp. 138–143.
- [25] H. Risken, *The Fokker-Planck Equation: Methods of Solution and Applications*, 2nd edition, Springer, Berlin, 1989.
- [26] M. J. Ruijter and C. W. Oosterlee, *Two-dimensional Fourier cosine series expansion method for pricing financial options*, SIAM J. Sci. Comput. **34** (2012), pp. B642–B671.

- [27] R. Tachet, *Non-Parametric Model Calibration in Finance*, PhD thesis, Ecole Centrale Paris, 2011.
- [28] G. Tataru and T. Fisher, *Stochastic local volatility*, Technical report, Bloomberg (2010).
- [29] D. Tavella and C. Randall, *Pricing Financial Instruments: The Finite Difference Method*, John Wiley & Sons, New York, 2000.
- [30] A. W. van der Stoep, L. A. Grzelak and C. W. Oosterlee, *The Heston stochastic-local volatility model: efficient Monte Carlo simulation*, Int. J. Theor. Appl. Finan. **17** (2014), pp. 1450045-1-1450045-30.
- [31] J. G. Verwer, E. J. Spee, J. G. Blom and W. Hundsdorfer, *A second-order Rosenbrock method applied to photochemical dispersion problems*, SIAM J. Sci. Comput. **20** (1999), pp. 1456-1480.
- [32] M. Wyns, *Convergence analysis of the Modified Craig-Sneyd scheme for two-dimensional convection-diffusion equations with nonsmooth initial data*, To appear in IMA J. Numer. Anal. (2016), <http://dx.doi.org/10.1093/imanum/drw028>.