On the Fourier cosine series expansion (COS) method for stochastic control problems

M.J. Ruijter[1,2], C.W. Oosterlee[1,3], and R.F.T. Aalbers[2]

[1]Centrum Wiskunde & Informatica, Amsterdam, the Netherlands
[2]CPB Netherlands Bureau for Economic Policy Analysis, Den Haag, the Netherlands
[3]Delft University of Technology, Delft, the Netherlands

June 15, 2011

Abstract

We develop a method for solving stochastic control problems under one-dimensional Lévy processes. The method is based on the dynamic programming principle and a Fourier cosine expansion method. Local errors in the vicinity of the domain boundaries may disrupt the algorithm. An extensive error analysis provides new insights based on which we develop an extrapolation method to deal with the propagation of local errors.

Keywords. Stochastic control problems, dynamic programming principle, Fourier cosine expansion method, error analysis, extrapolation, portfolio-selection problem.

AMS subject classifications. 65T40, 42A10, 60E10, 49J21, 91B25.

1 Introduction

Stochastic optimization can be defined as the optimization of a certain objective function, where an underlying state process is subject to random perturbations. The class of optimization problems can be subdivided into several different types, such as optimal stopping problems or impulse control problems ([14]). In this paper, we focus on stochastic control problems, in which the controller may influence the drift and diffusion terms of the underlying stochastic process.

In financial mathematics, the price of an option can often be formulated as a stochastic optimization problem. In the last decades, financial mathematics has contributed significantly to the theory and the improvement of the numerical methods to solve these problems. The techniques employed are closely related to those in the field of real option problems ([4]), encountered in economics, for example. These can also be represented by stochastic optimization problems.

In [6], an option pricing method for European options, based on Fourier cosine series expansions, was developed. This method was called the COS method and it was extended in [7] to pricing Bermudan, barrier, and American options, and to pricing swing options, which are frequently used in energy markets, in [19]. We will generalize this Fourier cosine technique to solving stochastic control problems, in which the state process can be controlled. Our work builds on both [6] and [7], but it differs on essential points for the accurate treatment of stochastic control problems.

In the resulting method, we need to determine the optimal control law for all possible state values. For that purpose, the value function must be accurately represented in the entire computational domain. It is known, however, that Fourier cosine expansions may be inaccurate near spatial boundaries, particularly outside the expansion interval. We give a detailed insight into the source of these local errors and their evolution. Based on this, we propose an extrapolation technique near the domain boundaries as an accurate solution technique in this context.

We test the method by solving two stochastic control problems of practical interest. The first one is the valuation of an option under uncertain volatility, which was solved in [15] and [16] by fully implicit discretization methods for the corresponding nonlinear partial differential equation (PDE). In contrast, the COS method will be based on the dynamic programming principle.

The second problem we discuss is a consumption-portfolio problem from economics. The model used is a simplified version of the well-known portfolio-selection problem, which is originally formulated and studied by Merton [11]. Here, an agent allocates his wealth to investments in risky or risk-free assets, and to consumption. The objective is to maximize the expected lifetime utility, by choosing consumption.

The outline of this paper is as follows. We start with the concepts and notation of stochastic control problems under multi-dimensional state processes. Then, in Section 3, a method based on the dynamic programming principle and the COS formula is derived for solving stochastic control problems with a one-dimensional underlying Lévy process. Section 4 provides an extensive error analysis and a way to solve possible propagating errors in the backward recursion. The two practical examples come up in Section 5. Finally, a conclusion and our main contributions are presented in Section 6.

2 Stochastic control problems

We consider the problem class of finite horizon stochastic control problems, where the objective function is optimized over a given finite domain. We start with the notation of control problems and some definitions, based on [14]. The numerical method that we will develop relies on the dynamic programming principle, which is explained in Section 2.2.

2.1 Problem description

Let $(\Omega, \mathcal{F}, P)$ be a probability space, $T > 0$ a finite terminal time, $\mathbb{F} = (\mathcal{F}_s)_{0\leq s\leq T}$ a filtration satisfying the usual conditions and W a d-dimensional Brownian motion on the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$. For $t \in [0, T]$, we denote by $\mathcal{T}_{t,T}$ the set of stopping times valued in $[t, T]$. The controlled state process X_t is valued in $\mathbb{R}^n$ and satisfies the stochastic differential equation

$$dX_s = b(s, X_s, \boldsymbol{\alpha}_s)ds + \sigma(s, X_s, \boldsymbol{\alpha}_s)dW_s. \tag{2.1}$$

The process here is a controlled diffusion process, which we use in this section for ease of notation and to stay in line with [14]. Later on, we will work with the class of Lévy processes.

The control process $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_s)_{0\leq s\leq T}$ is progressively measurable with respect to $\mathbb{F}$ and is valued in the control set A, a subset of $\mathbb{R}^\ell$. In the class of stochastic optimization problems that we discuss here, the state process is influenced by the control process, $\boldsymbol{\alpha}$, whose value is determined at any time t based on the available information.

The measurable functions $b : [0, T] \times \mathbb{R}^n \times A \to \mathbb{R}^n$ and $\sigma : [0, T] \times \mathbb{R}^n \times A \to \mathbb{R}^{n\times d}$ satisfy

uniform Lipschitz conditions in A. Let $\mathcal{A}$ denote the set of control processes $\boldsymbol{\alpha}$ that satisfy a square integrability condition. With this notation, an element $\boldsymbol{\alpha} \in \mathcal{A}$ is a process over time, with values in set A. With the conditions on b and σ, for all $\boldsymbol{\alpha} \in \mathcal{A}$ and for any starting condition $(t, x) \in [0, T] \times \mathbb{R}^n$ a unique strong solution to equation (2.1) starting from x at $s = t$ exists, which is denoted by $\{X_s^{t,x}, t \leq s \leq T\}$ ([14]). The dependence of $X_s^{t,x}$ on the control process $\boldsymbol{\alpha}$ is omitted for notational convenience.

We assume that $f : [0, T] \times \mathbb{R}^n \times A \to \mathbb{R}$ and $g : \mathbb{R}^n \to \mathbb{R}$ are two measurable functions that satisfy a lower boundedness or a linear growth condition.

The *gain function* on the finite horizon is defined as

$$J(t, x, \boldsymbol{\alpha}) := \mathbb{E}\left[\int_t^T e^{-\rho(s-t)} f(s, X_s^{t,x}, \boldsymbol{\alpha}_s) ds + e^{-\rho(T-t)} g(X_T^{t,x})\right], \tag{2.2}$$

for all $(t, x) \in [0, T] \times \mathbb{R}^n$ and $\boldsymbol{\alpha} \in \mathcal{A}$. The function f is a so-called *running profit function* and g is a *terminal reward function*. $\rho \geq 0$ is a *discount rate*, which is common in economic and financial problems. The objective of the finite horizon problem is to maximize the gain function over all admissible controls in $\mathcal{A}$. We also introduce the so-called *value function*

$$v(t, x) \quad := \quad \sup_{\boldsymbol{\alpha} \in \mathcal{A}} J(t, x, \boldsymbol{\alpha}). \tag{2.3}$$

For an initial state $(t, x) \in [0, T) \times \mathbb{R}^n$, we say that $\boldsymbol{\alpha}^* \in \mathcal{A}$ is an optimal control if $v(t, x) = J(t, x, \boldsymbol{\alpha}^*)$. A control process $\boldsymbol{\alpha}$ is called a Markovian control if it has the form $\boldsymbol{\alpha}_s = a(s, X_s^{t,x})$ for some measurable function $a : [0, T] \times \mathbb{R}^n \to A$.

The notion of stochastic control problems can easily be extended with the concepts of optimal stopping or impulse control ([13]). Then the controller does not (only) have the disposal of a control process $\boldsymbol{\alpha}$ to optimize his objective, but he can determine the terminal time or can add extra impulses to the state process.

2.2 Dynamic programming principle

An important principle in the theory of stochastic control is Bellman's optimality principle, also called the dynamic programming principle ([14]). It means that if one has taken an optimal control path until some arbitrary observation time θ, then, given this information, it remains optimal to use it after that observation time. The dynamic programming principle is stated as follows:

Result 2.1. *(Dynamic programming principle (DPP) (Finite horizon))[14]*
Let $(t, x) \in [0, T] \times \mathbb{R}^n$. Then we have

$$v(t, x) \quad = \quad \sup_{\boldsymbol{\alpha} \in \mathcal{A}} \mathbb{E}\left[\int_t^\theta e^{-\rho(s-t)} f(s, X_s^{t,x}, \boldsymbol{\alpha}_s) ds + e^{-\rho(\theta-t)} v(\theta, X_\theta^{t,x})\right], \tag{2.4}$$

for any stopping time $\theta \in \mathcal{T}_{t,T}$.

By the dynamic programming principle, one can split the optimization problem into two parts. An optimal control may be obtained by first searching for an optimal control from a time θ given the state value $X_\theta^{t,x}$, in other words, compute $v(\theta, X_\theta^{t,x})$. Then, the quantity

$$\mathbb{E}\left[\int_t^\theta e^{-\rho(s-t)} f(s, X_s^{t,x}, \boldsymbol{\alpha}_s) ds + e^{-\rho(\theta-t)} v(\theta, X_\theta^{t,x})\right]$$

is maximized over all controls on $[t, \theta]$. We will use this principle to set up a numerical approach for stochastic control problems in Section 3.

Remark 2.1. *By the dynamic programming principle, one can derive the well-known Hamilton-Jacobi-Bellman (HJB) equation corresponding to problem (2.3), see [14]. This second-order non-linear partial differential equation is the infinitesimal version of the dynamic programming principle and reads*

$$-\frac{\partial v}{\partial t}(t,x) + \rho v(t,x) - \sup_{a \in A}[\mathcal{L}^a v(t,x) + f(t,x,a)] = 0, \qquad \forall (t,x) \in [0,T) \times \mathbb{R}^n, \tag{2.5}$$

with differential operator of second order

$$\mathcal{L}^a v(t,x) \quad = \quad b(x,a) \cdot D_x v(t,x) + \tfrac{1}{2} tr(\sigma\sigma'(x,a) D_x^2 v(t,x)), \tag{2.6}$$

where $\sigma\sigma'(x,a)$ is an $n \times n$ matrix with components $(\sigma\sigma')_{ij}(x,a) = \sum_{k=1}^{d} \sigma_{ik}(x,a)\sigma_{jk}(x,a)$. $\mathcal{L}^a$ is called the infinitesimal generator associated to the diffusion X_t with constant control a. The vector D_x denotes the gradient of a function and matrix D_x^2 consists of its second derivatives. The terminal condition is $v(T,x) = g(x)$, $\forall x \in \mathbb{R}^n$, resulting from the definition of the value function.

Stochastic control problems may be solved by the use of PDE methods to the corresponding HJB equation. We refer to [8] and [16] for numerical discretization methods. Then, issues about convergence to the correct, viscosity solution arise. The viscosity solution concept was introduced by P.L. Lions ([10]). We refer to [2] for a general introduction to viscosity solutions and some general uniqueness and existence results. As we use the dynamic programming approach, we will not go into details about this.

3 COS method for stochastic control problems

In this section, we set up a general method to solve stochastic control problems under a one-dimensional[1] *Lévy process*, X_t, for which the characteristic function is known. This class contains the constant coefficient jump diffusion processes. The method is based on the dynamic programming principle and uses the so-called COS formula, which was developed in [6] for pricing European options. It results in a recursive algorithm based on the Fast Fourier Transform algorithm. We will explain the COS formula in Section 3.1. We start here with the discrete-time framework of the solution method.

Initial time is denoted by t_0 and T is the terminal time. We take a fixed equidistant grid of control times $t_0 < t_1 < \ldots t_m < \ldots < t_M = T$, with $\Delta t := t_{m+1} - t_m$, and a bounded set of possible control values $A \subset \mathbb{R}^\ell$. As a discrete approximation, we assume that the control processes are constant during the time intervals $[t_m, t_{m+1}]$. At each control time t_m, with $m < M$, one can choose a control value from the set A, which influences the stochastic process during the time interval $[t_m, t_{m+1}]$. This value is denoted by α_m, where the subscript refers to the control time. The choice may depend on the current value of the state process. With this notation, bold faced $\boldsymbol{\alpha}$ denotes a control process and α_m denotes a single control value.

Remark 3.1. *For diffusion processes, as in (2.1), the stochastic process evolves according to the following dynamics:*

$$dX_s \quad = \quad b(s, X_s, \alpha_m)ds + \sigma(s, X_s, \alpha_m)dW_s, \qquad \text{for } s \in [t_m, t_{m+1}]. \tag{3.1}$$

In the examples in Section 5 we will use processes of this form, to be precise, geometric Brownian motions. Their log-transformed dynamics belong to the class of Lévy processes, which we consider from now on.

[1] Extension of the methodology to higher dimensional state processes is part of forthcoming research.

The value function reads

$$v(t,x) := \max_{\boldsymbol{\alpha}\in\hat{\mathcal{A}}} \mathbb{E}\left[\int_t^T e^{-\rho(s-t)} f(s, X_s^{t,x}, \boldsymbol{\alpha}_s) ds + e^{-\rho(T-t)} g(X_T^{t,x})\right]. \tag{3.2}$$

$\hat{\mathcal{A}} \subset \mathcal{A}$ denotes the set of all possible control paths $\{\alpha_m\}_{m=0}^{M-1}$, where α_m is valued in the control set A. The terminal condition is $v(T,y) = g(y)$. We deal with a discrete-time stochastic control problem, with M control times. Convergence of the numerical solution to the solution of the original problem (2.3) is achieved by increasing the number of time steps (value of M).

The dynamic programming principle now gives:

$$\begin{aligned} v(t_{m-1},x) &= \max_{\alpha_{m-1}\in A} \mathbb{E}^{t_{m-1},x}\left[\int_{t_{m-1}}^{t_m} e^{-\rho(s-t)} f(s, X_s, \alpha_{m-1}) ds + e^{-\rho\Delta t} v(t_m, X_{t_m})\right] \\ &= \max_{\alpha_{m-1}\in A}\left[\int_{t_{m-1}}^{t_m} e^{-\rho(s-t)} \mathbb{E}^{t_{m-1},x}[f(s, X_s, \alpha_{m-1})] ds + e^{-\rho\Delta t}\mathbb{E}^{t_{m-1},x}\left[v(t_m, X_{t_m})\right]\right]. \end{aligned} \tag{3.3}$$

For ease of notation, we use the form $\mathbb{E}^{t,x}[X_s]$ instead of $\mathbb{E}[X_s^{t,x}]$. We presume that the first term in the maximization operator, i.e. the time integral, is known analytically, or can be approximated using, for example, a numerical integration rule. We denote this value by $F(t_{m-1}, x, \alpha_{m-1})$ and call it the *profit function*. The expectation in the second term, which we call the *continuation value*, is denoted by $c(t_{m-1}, x, \alpha_{m-1})$. So, we use the notation

$$v(t_{m-1},x) = \max_{\alpha_{m-1}\in A}[F(t_{m-1}, x, \alpha_{m-1}) + c(t_{m-1}, x, \alpha_{m-1})]. \tag{3.4}$$

3.1 Fourier cosine expansion formula (COS formula)

Next, we explain the method of choice to approximate the continuation value, in the backward recursion,

$$\begin{aligned} c(t_{m-1},x,\alpha) &= e^{-\rho\Delta t}\mathbb{E}\left[v(t_m, X_{t_m}) | X_{t_{m-1}} = x, \alpha_{m-1} = \alpha\right] \\ &= e^{-\rho\Delta t}\int_{\mathbb{R}} v(t_m, y) f(y|x,\alpha) dy. \end{aligned} \tag{3.5}$$

The numerical method is based on series expansions of the value function at the next time level and the density function. The resulting equation is called the COS formula, due to the use of Fourier-*cos*ine series expansions. Fourier series expansions and their convergence properties have been discussed in [1].

The conditional density function $f(y|x,\alpha)$ decays to zero rapidly as $y \to \pm\infty$, so that we can, for given x, truncate the infinite integration range of the expectation to some interval $[a,b] \subset \mathbb{R}$ without loosing significant accuracy. This gives the approximation

$$\begin{aligned} c_1(t_{m-1},x,\alpha|[a,b]) &= e^{-\rho\Delta t}\int_a^b v(t_m, y) f(y|x,\alpha) dy \\ &= e^{-\rho\Delta t}\int_a^b v(t_m, y) {\sum_{k=0}^{+\infty}}' G_k(x,\alpha)\cos\left(k\pi\frac{y-a}{b-a}\right) dy. \end{aligned} \tag{3.6}$$

The notation c_i is used for the different approximations of c and keeps track of the numerical errors that set in from each step. In the second equation in (3.6), the conditional density is replaced by its Fourier cosine expansion in y on $[a,b]$, with series coefficients $\{G_k(x,\alpha)\}_{k=0}^{+\infty}$ defined by

$$G_k(x,\alpha) := \frac{2}{b-a}\int_a^b f(y|x,\alpha)\cos\left(k\pi\frac{y-a}{b-a}\right)dy. \tag{3.7}$$

$\sum'$ in (3.6) indicates that the first term in the summation is weighted by one-half. We interchange summation and integration and define:

$$V_k(t_m) := \frac{2}{b-a}\int_a^b v(t_m,y)\cos\left(k\pi\frac{y-a}{b-a}\right)dy, \tag{3.8}$$

which are the Fourier cosine series coefficients of $v(t_m,y)$ on $[a,b]$. This results in

$$c_1(t_{m-1},x,\alpha|[a,b]) = \frac{b-a}{2}e^{-\rho\Delta t}\sum_{k=0}^{+\infty}{}' G_k(x,\alpha)V_k(t_m). \tag{3.9}$$

Truncation of the series summation gives us the approximation

$$c_2(t_{m-1},x,\alpha|[a,b],N) = \frac{b-a}{2}e^{-\rho\Delta t}\sum_{k=0}^{N-1}{}' G_k(x,\alpha)V_k(t_m). \tag{3.10}$$

The coefficients $G_k(x,\alpha)$ can now be approximated as follows

$$\begin{aligned} G_k(x,\alpha) &\approx \frac{2}{b-a}\int_{\mathbb{R}} f(y|x,\alpha)\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &= \frac{2}{b-a}\mathrm{Re}\left(\varphi\left(\frac{k\pi}{b-a}\Big|x,\alpha\right)e^{-ik\pi\frac{a}{b-a}}\right) := F_k(x,\alpha). \end{aligned} \tag{3.11}$$

$\mathrm{Re}\,(.)$ denotes taking the real part of the input argument. $\varphi(.|x,\alpha)$ is the conditional *characteristic function* of X_{t_m}, given $X_{t_{m-1}} = x$ and $\alpha_{m-1} = \alpha$. The density function of a stochastic process is usually not known, but often its characteristic function is known (see [5], [6]). Moreover, for Lévy processes it follows that:

$$\varphi(u|x,\alpha) = \varphi(u|0,\alpha)e^{iux} := \varphi_{levy}(u|\alpha)e^{iux}. \tag{3.12}$$

Inserting the above equations in (3.10) gives us the *COS formula* for approximation of $c(t_{m-1},x,\alpha)$:

$$\begin{aligned} \hat{c}(t_{m-1},x,\alpha|[a,b],N) &:= c_3(t_{m-1},x,\alpha|[a,b],N) \\ &= e^{-\rho\Delta t}\sum_{k=0}^{N-1}{}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{k\pi}{b-a}\Big|\alpha\right)e^{ik\pi\frac{x-a}{b-a}}\right)V_k(t_m). \end{aligned} \tag{3.13}$$

Since the terms $V_k(t_m)$ are independent of x, we can calculate the continuation value for many values of x simultaneously.

The value function is now approximated by

$$\hat{v}(t_{m-1},x) := \max_{\alpha_{m-1}\in A}[F(t_{m-1},x,\alpha_{m-1}) + \hat{c}(t_{m-1},x,\alpha_{m-1})]. \tag{3.14}$$

3.2 Recursion formula for coefficients V_k

The algorithm for solving stochastic control problems is based on the recursive recovery of the coefficients V_k, starting with the coefficients at the terminal time:

$$V_k(t_M) \quad = \quad \frac{2}{b-a}\int_a^b v(T,y)\cos\left(k\pi\frac{y-a}{b-a}\right)dy, \tag{3.15}$$

for which we assume that an analytic solution is available. This is the case for, among others, exponential and polynomial terminal reward functions. These coefficients are used for the approximation of the continuation value at time t_{M-1}.

Next we consider the coefficients that are used to approximate the continuation value at time t_{m-1}, for $m \leq M-1$. The value function, equation (3.4), at time t_m appears in the terms $V_k(t_m)$ and we need to find an optimal control law for all state values $y \in [a,b]$. We propose two techniques for this:

- Firstly, suppose that the set A of possible control values is finite, $A = \{\alpha^1, \ldots, \alpha^q, \ldots, \alpha^K\}$, where K is a finite number. Then it may be possible to determine sub-domains $\mathcal{D}_m^q \subset [a,b]$, so that for each $y \in \mathcal{D}_m^q$ it is optimal to choose control α_m^q at control time t_m. The subscript of α_m^q indicates the time level and the superscript the control value. This approach will be applied to the problem in Section 5.1.
- If this procedure cannot be applied, or in case the control set A is a continuous range, then the interval $[a,b]$ is divided into sub-intervals, which span the interval $[a,b]$. They are denoted by $\mathcal{D}_m^q$, $q = 1,2,\ldots,K$, where K is a finite number. On each sub-interval, $\mathcal{D}_m^q$, we determine the optimal control, α_m^q, for the time interval $[t_m, t_{m+1}]$. We here assume that the control value is constant over the spatial sub-interval. With many sub-intervals this approximation may be sufficiently accurate. This approach will be used in the application in Section 5.2.

In both approaches, we split the integral for the definition of V_k into different parts:

$$\begin{aligned} & V_k(t_m) \\ = \quad & \frac{2}{b-a}\sum_{q=1}^{K}\int_{\mathcal{D}_m^q} F(t_m,y,\alpha_m^q)\cos\left(k\pi\frac{y-a}{b-a}\right)dy + \frac{2}{b-a}\sum_{q=1}^{K}\int_{\mathcal{D}_m^q} c(t_m,y,\alpha_m^q)\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ := \quad & \sum_{q=1}^{K} U_k(t_m,\mathcal{D}_m^q,\alpha_m^q) + \sum_{q=1}^{K} C_k(t_m,\mathcal{D}_m^q,\alpha_m^q), \qquad (m \neq M). \end{aligned} \tag{3.16}$$

Here assume that the terms U_k are known analytically[2],

$$U_k(t_m,z_1,z_2,\alpha) \quad = \quad \frac{2}{b-a}\int_{z_1}^{z_2} F(t_m,y,\alpha)\cos\left(k\pi\frac{y-a}{b-a}\right)dy, \tag{3.17}$$

where z_1 and z_2 denote the boundaries of interval $\mathcal{D}_m^q$. In practical applications these terms may be independent of time. The coefficients C_k at time t_m can be approximated by using the coefficients V_k from the next time level, as we will explain shortly. This results in a backward recursion of the coefficients V_k.

[2]This is the case for, amongst others, exponential and polynomial functions F. If these terms are not known analytically, they can be approximated by numerical integration rules or discrete Fourier cosine transforms.

For approximation of the value function at time t_{M-2} we need the coefficients $V_k(t_{M-1})$. We will use the approximated values, $\hat{c}(t_{M-1}, y, \alpha)$, to approximate the terms $C_k(t_{M-1}, z_1, z_2, \alpha)$. This approximation is then denoted by $\hat{C}_k(t_{M-1}, z_1, z_2, \alpha)$ and results in the coefficients:

$$\hat{V}_k(t_{M-1}) \quad := \quad \sum_{q=1}^{K} U_k(t_{M-1}, \mathcal{D}^q_{M-1}, \alpha^q_{M-1}) + \sum_{q=1}^{K} \hat{C}_k(t_{M-1}, \mathcal{D}^q_{M-1}, \alpha^q_{M-1}). \tag{3.18}$$

On the integrands of terms $\hat{C}_k$ we can apply again the Fourier cosine series expansion by inserting equation (3.13):

$$\begin{aligned}
&\hat{C}_k(t_{M-1}, z_1, z_2, \alpha) \\
&= \frac{2}{b-a} \int_{z_1}^{z_2} \hat{c}(t_{M-1}, y, \alpha) \cos\left(k\pi \frac{y-a}{b-a}\right) dy \\
&= e^{-\rho \Delta t} \frac{2}{b-a} \int_{z_1}^{z_2} \left({\sum_{j=0}^{N-1}}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{j\pi}{b-a} \middle| \alpha \right) e^{ij\pi \frac{y-a}{b-a}} \right) V_j(t_M) \right) \cos\left(k\pi \frac{y-a}{b-a}\right) dy \\
&= e^{-\rho \Delta t} \mathrm{Re}\left({\sum_{j=0}^{N-1}}' \varphi_{levy}\left(\frac{j\pi}{b-a} \middle| \alpha \right) V_j(t_M) \cdot M_{k,j}(z_1, z_2) \right),
\end{aligned} \tag{3.19}$$

where the elements of matrix $M(z_1, z_2)$ are given by:

$$M_{k,j}(z_1, z_2) := \frac{2}{b-a} \int_{z_1}^{z_2} e^{ij\pi \frac{y-a}{b-a}} \cos\left(k\pi \frac{y-a}{b-a}\right) dy. \tag{3.20}$$

Finally, we end up with the vector form

$$\hat{V}(t_{M-1}) \quad = \quad \sum_{q=1}^{K} U(t_{M-1}, \mathcal{D}^q_{M-1}, \alpha^q_{M-1}) + \sum_{q=1}^{K} e^{-\rho \Delta t} \mathrm{Re}\left(M(\mathcal{D}^q_{M-1}) \mathbf{w}^q \right), \tag{3.21}$$

where

$$\mathbf{w}^q = \{w^q_j\}_{j=0}^{N-1} \quad \text{with} \quad w^q_j = \varphi\left(\frac{j\pi}{b-a} \middle| \alpha^q_{M-1} \right) V_j(t_M), \quad w^q_0 = \tfrac{1}{2} \varphi(0 | \alpha^q_{M-1}) V_0(t_M). \tag{3.22}$$

The parameters of the matrices M are the boundary values of their respective integration ranges.

For the other coefficients, $V_k(t_m)$, $1 \le m \le M-2$, the approximations $\hat{c}(t_m, y, \alpha)$ *and* $\hat{V}_k(t_{m+1})$ will be used to approximate the terms $C_k(t_m, z_1, z_2, \alpha)$. The same arguments give the following numerical approximation of the Fourier cosine coefficients at time t_m:

$$\hat{V}(t_m) \quad = \quad \sum_{q=1}^{K} U(t_m, \mathcal{D}^q_m, \alpha^q_m) + \sum_{q=1}^{K} e^{-\rho \Delta t} \mathrm{Re}\left(M(\mathcal{D}^q_m) \hat{\mathbf{w}}^q \right), \quad (m = 1, \ldots, M-2), \tag{3.23}$$

where

$$\hat{\mathbf{w}}^q = \{\hat{w}^q_j\}_{j=0}^{N-1} \quad \text{with} \quad \hat{w}^q_j = \varphi\left(\frac{j\pi}{b-a} \middle| \alpha^q_m \right) \hat{V}_j(t_{m+1}), \quad \hat{w}^q_0 = \tfrac{1}{2} \varphi(0 | \alpha^q_m) \hat{V}_0(t_{m+1}). \tag{3.24}$$

An additional error is introduced because the coefficients are approximated using the approximated elements $\hat{V}_j(t_{m+1})$. We will examine this evolving error in Sections 4.2 and 4.3 and propose a more accurate approximation for the Fourier-coefficients $V_k(t_m)$.

3.3 Algorithm

The matrix-vector products $M\mathbf{w}$ in the terms $\hat{C}$ can be computed by a Fourier-based algorithm, as stated in the next result:

Result 3.1. *(Efficient computation of $\hat{C}(t_m, z_1, z_2, \alpha)$)[7]*
The matrix-vector product $M(z_1, z_2)\mathbf{w}$ can be computed in $O(N \log_2 N)$ operations, with the help of the Fast Fourier Transform (FFT) algorithm.

The key insight of this efficient computation is the equality

$$M_{k,j}(z_1, z_2) = -\frac{i}{\pi}\left(M^c_{k,j}(z_2, z_2) + M^s_{k,j}(z_1, z_2)\right), \tag{3.25}$$

where matrix M^c is a Hankel matrix and M^s a Toeplitz matrix ($M^c_{i,j} = M^c_{i-1,j+1}$ and $M^s_{i,j} = M^s_{i+1,j+1}$). The matrices $M^c_{k,j}$ and $M^s_{k,j}$ can be found in [7]. The special matrix structure enables us to use the FFT algorithm for the matrix-vector products ([7]). If a process does not possess the property in equation (3.12), the FFT algorithm cannot be employed in a straightforward way (see [19]).

We can recover the terms $\hat{V}_k(t_m)$ recursively, starting with $V_k(t_M)$. The algorithm to solve the discrete-time stochastic control problem (3.2) backwards in time then reads:

Algorithm 1. *(COS method for stochastic control problems)*
Initialization:
Calculate coefficients $V_k(t_M)$ for $k = 0, 1, \ldots, N-1$.
***Main loop to recover** $\hat{V}(t_m)$**:** For $m = M-1$ to 1:*

- *Determine the sub-domains $\mathcal{D}^q_m$ for which the optimal control value is α^q_m, or determine the optimal control values α^q_m for given sub-domains $\mathcal{D}^q_m$.*
- *Compute $\hat{V}(t_m)$, equation (3.23), with the help of the FFT algorithm.*

Final step:
Compute $\hat{v}(t_0, x_0)$ by inserting $\hat{V}_k(t_1)$ into equation (3.13).

The computational complexity of the algorithm is $O\Big(MKN \log_2 N\Big)$, as we need to compute M time steps, and K sub-intervals.

Remark 3.2. *We elaborate on the differences between using the COS method for pricing Bermudan and barrier options and for solving stochastic control problems. In Algorithm 1, we search for an optimal control law for all state values y in the computational domain $[a, b]$. For this, the numerical continuation values need to be accurate over the entire interval. Significant errors may arise, however, in the vicinity of the boundaries, as we will show in Sections 4.1 and 5.2. In Section 4.2, a remedy for this problem will be proposed. When pricing barrier and Bermudan call options, as in [7], one searches for the early-exercise points, where the continuation value equals the payoff. For this task, interval $[a, b]$ can be chosen large so that the early-exercise points are not close to the interval boundaries. When pricing financial options, this type of inaccuracy does thus not occur in practice, e.g., for usual values of maturity T.*

Another difference in applying the COS method to stochastic control problems is the dependency of the characteristic function on the control values. As this function is evaluated often during the optimization procedure, it may be time-consuming.

4 Error analysis and extrapolation technique

In this section, we analyze the error of the COS method for stochastic control problems and base our analysis on [6], [7], and [17]. Errors are introduced by the COS formula and by evolution through time via the coefficients $\hat{V}_k$ and a possibly incorrect control $\boldsymbol{\alpha}$. We start with the local error where backward recursion of the approximated terms, $\hat{V}_k$, and control are not taken into account, in Section 4.1. In the financial context, this corresponds to a European option with an uncontrolled asset price process. We give an example in which the COS formula is inaccurate in the vicinity of a domain boundary (Section 4.1.1). This may give difficulties with the recursive recovery of the Fourier cosine coefficients $V_k(t_m)$. In Section 4.2, we propose an improved approximation for $V_k(t_m)$, which is more accurate than $\hat{V}_k(t_m)$ from (3.23). Finally, the propagating error in the backward recursion is studied and bounded.

4.1 Local error COS formula

We define the local error of the COS formula for the continuation value by

$$\epsilon_{COS}(t_{m-1}, x, \alpha|[a,b], N) := c(t_{m-1}, x, \alpha) - \hat{c}(t_{m-1}, x, \alpha|[a,b], N). \tag{4.1}$$

The above notation includes the parameters used for the approximations, namely $[a,b]$ and N. The error

$$\max_{\alpha_{m-1} \in A} |\epsilon_{COS}(t_{m-1}, x, \alpha_{m-1}|[a,b], N)| \tag{4.2}$$

bounds the absolute error of the approximated value function $\hat{v}(t_{m-1}, x)$, assuming that the correct optimal control law has been chosen and that the function $F(t_{m-1}, x, \alpha)$ is known analytically, or can be approximated sufficiently accurate.

We first assume that the terms $V_k(t_m)$ are exact. An upper bound for the error of the European option pricing COS formula with respect to the truncation range and the convergence rate, in dependence of N, has been derived in [6]. Errors are introduced in three steps, we discuss them one after the other:

1. The integration range truncation error:

$$\epsilon_1(t_{m-1}, x, \alpha|[a,b]) := c(t_{m-1}, x, \alpha) - c_1(t_{m-1}, x, \alpha|[a,b]) = e^{-\rho \Delta t} \int_{\mathbb{R}\setminus[a,b]} v(t_m, y) f(y|x, \alpha) dy. \tag{4.3}$$

If $v(t_m, y) f(y|x, \alpha)$ is sufficiently small outside the interval $[a,b]$, then the error ϵ_1 can be ignored.

2. The series truncation error on $[a,b]$:

$$\begin{aligned} \epsilon_2(t_{m-1}, x, \alpha|[a,b], N) &:= c_1(t_{m-1}, x, \alpha|[a,b]) - c_2(t_{m-1}, x, \alpha|[a,b], N) \\ &= \frac{b-a}{2} e^{-\rho \Delta t} \sum_{k=N}^{+\infty} G_k(x, \alpha) V_k(t_m). \end{aligned} \tag{4.4}$$

The convergence rate of Fourier cosine series depends on the properties of the approximated functions in the expansion interval. Information about different convergence types can be found in [1].

For N large, so that $|V_k(t_m)| \leq 1$ for $k \geq N$, the error can be bounded by

$$\begin{aligned} |\epsilon_2(t_{m-1}, x, \alpha|[a,b], N)| &= \left| \frac{b-a}{2} e^{-\rho\Delta t} \sum_{k=N}^{+\infty} G_k(x,\alpha) V_k(t_m) \right| \\ &\leq \frac{b-a}{2} e^{-\rho\Delta t} \sum_{k=N}^{+\infty} |G_k(x,\alpha)|. \end{aligned} \tag{4.5}$$

The error is dominated by the series truncation error of the coefficients G_k.

With the theory in [1], we find that the error converges exponentially for density functions in the class $C^\infty([a,b])$:

$$|\epsilon_2(t_{m-1}, x, \alpha|[a,b], N)| \leq [\,] \exp(-(N-1)\nu), \qquad \forall\, x, \alpha, \tag{4.6}$$

where $\nu > 0$ is a constant and the empty brackets $[\,]$ denote factors that vary slower than exponentially in N. A density function with discontinuity in one of its derivatives results in an algebraic convergence:

$$|\epsilon_2(t_{m-1}, x, \alpha|[a,b], N)| \leq \frac{P}{(N-1)^{n-1}}, \qquad \forall\, x, \alpha, \tag{4.7}$$

where $P > 0$ and $n \geq 1$ are constants.

3. The error related to approximating $G_k(x,\alpha)$ by $F_k(x,\alpha)$ (equation (3.11)):

$$\begin{aligned} \epsilon_3(t_{m-1}, x, \alpha|[a,b], N) &:= c_2(t_{m-1}, x, \alpha|[a,b], N) - c_3(t_{m-1}, x, \alpha|[a,b], N) \\ &= \frac{b-a}{2} e^{-\rho\Delta t} {\sum_{k=0}^{N-1}}' (G_k(x,\alpha) - F_k(x,\alpha)) V_k(t_m) \\ &= -e^{-\rho\Delta t} \int_{\mathbb{R}\setminus[a,b]} \left[{\sum_{k=0}^{N-1}}' \cos\left(k\pi \frac{y-a}{b-a}\right) V_k(t_m) \right] f(y|x,\alpha) dy \\ &= -e^{-\rho\Delta t} \int_{\mathbb{R}\setminus[a,b]} \hat{v}(t_m, y) f(y|x,\alpha) dy. \end{aligned} \tag{4.8}$$

Remark 4.1. *Note that the Fourier cosine series expansions in Section 3.1 are defined for $y \in [a,b]$, whereas here the function $\hat{v}(t_m, y)$ is evaluated on $\mathbb{R}\setminus[a,b]$. Here we denote by function $\hat{v}(t_m, y)$, for $y \in \mathbb{R}\setminus[a,b]$, the symmetric extension of the Fourier cosine series expansion outside the expansion interval. This value will usually be different from $v(t_m, y)$, even if N tends to infinity.*

The integration range truncation error, ϵ_1, enters by truncation of the infinite domain to the finite domain $[a,b]$. Conversely, error ϵ_3 is due to replacing the finite domain by an infinite domain in equation (3.11). The third error 'compensates', completely or partly, for the first error. Addition of both errors gives

$$\epsilon_1(t_{m-1}, x, \alpha|[a,b]) + \epsilon_3(t_{m-1}, x, \alpha|[a,b], N) = e^{-\rho\Delta t} \int_{\mathbb{R}\setminus[a,b]} [v(t_m, y) - \hat{v}(t_m, y)]\, f(y|x,\alpha) dy. \tag{4.9}$$

So, error $\epsilon_1 + \epsilon_3$ results from using $e^{-\rho\Delta t}\hat{v}(t_m, y)$ instead of the true discounted value function $e^{-\rho\Delta t} v(t_m, y)$. We can write the local error of the COS formula as

$$\epsilon_{COS}(t_{m-1}, x, \alpha|[a,b], N) = \epsilon_1(t_{m-1}, x, \alpha|[a,b]) + \epsilon_2(t_{m-1}, x, \alpha|[a,b], N) + \epsilon_3(t_{m-1}, x, \alpha|[a,b], N).$$

If, for given x, the integration interval $[a,b]$ is chosen sufficiently wide then the series truncation error ϵ_2 dominates the overall local error. This implies that for smooth density functions the local error converges exponentially to zero, otherwise it goes algebraically.

For a given interval $[a, b]$, the local error may however be large if x is in the vicinity of the domain boundaries, resulting from error $\epsilon_1 + \epsilon_3$. We will show this by an example in the next section. A local error may propagate via the backward recursion for a stochastic control problem.

4.1.1 Significant error in the vicinity of the boundaries

Here, an example of a large error close to the spatial boundaries is presented, when we use the COS formula to price a European *call* option.

In the financial setting, the asset price at time t is denoted by S_t and the strike price by K. We model the asset price by a geometric Brownian motion. The payoff of a call option at terminal time, T, with log-asset price $y = \log(S_T)$, is given by the function

$$g(y) = (e^y - K)^+, \tag{4.10}$$

where $(z)^+ := \max(z, 0)$. Fourier cosine series expansion gives the approximation:

$$\hat{g}(y|[a, b], N) = \sum_{k=0}^{N-1}{}' V_k \cos\left(k\pi \frac{y-a}{b-a}\right), \tag{4.11}$$

in which

$$\begin{aligned} V_k &= \frac{2}{b-a}\int_a^b g(y)\cos\left(k\pi\frac{y-a}{b-a}\right)dy = \frac{2}{b-a}\Big(\chi_k(\log K, b, a, b) - K\psi_k(\log K, b, a, b)\Big), \\ &\quad (a \le \log K \le b). \end{aligned} \tag{4.12}$$

The analytic solution of the functions χ_k and ψ_k can be found in Appendix A. The risk-neutral option pricing formula ([18]) reads

$$v(t_0, x) = e^{-r\Delta t}\mathbb{E}^{t_0,x}\left[v(T, X_T)\right] = e^{-r\Delta t}\int_{\mathbb{R}} g(y)f(y|x)dy, \tag{4.13}$$

where $X_s = \log S_s$ is the log-asset price process, r is the risk-neutral interest rate and $\Delta t := T - t_0$ [3]. The problem considered is a simplified, reduced version of stochastic control problem (2.3). The COS formula yields

$$\begin{aligned} \hat{v}(t_0, x|[a, b], N) &:= e^{-r\Delta t}\sum_{k=0}^{N-1}{}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{k\pi}{b-a}\right)e^{ik\pi\frac{x-a}{b-a}}\right)V_k \\ &= e^{-r\Delta t}\int_{[a,b]} g(y)f(y|x)dy + e^{-r\Delta t}\int_{\mathbb{R}\setminus[a,b]} \hat{g}(y)f(y|x)dy - \epsilon_2(t_0, x|[a, b], N). \end{aligned} \tag{4.14}$$

We take N sufficiently large, so that error ϵ_2 can be neglected. The equation above shows that a significant error is introduced if $\hat{g}(y)f(y|x)$ is not close to $g(y)f(y|x)$ outside the expansion interval $[a, b]$, which is the case in the example to follow.

For the call option under geometric Brownian motion an analytic solution is available, i.e. the Black-Scholes price, so that the numerical option value can be compared with the exact solution. The following parameters are used for the tests in this section:

$$K = 100,\ S_0 = 100\,(x_0 \approx 4.6),\ r = 0.1,\ q = 0,\ T = 0.1,\ \sigma = 0.25,\ [a, b] = [3.82,\ 5.40],\ N = 2^{10}. \tag{4.15}$$

[3] Note that there is no control process $\boldsymbol{\alpha}$ and no running profit function in this test.

The parameters and the interval $[a, b]$ are adopted from [6]. Log-asset price, x, is varied and the results are shown in Figure 4.1. In the left-side plot, function $g(y)$, the series expansion, and its extension outside $[a, b]$, which is symmetric in a and b, are presented. Function $\hat{g}(y)$ resembles the true payoff function well at the left-hand side of a, as the function is constant there. Hence, no error is introduced in the vicinity of that boundary. However, at the right-side of b a difference between the two functions is observed, which gives a significant error $\epsilon_1+\epsilon_3$ at that boundary. This is shown in the right-side plot, where the exact Black-Scholes price and the COS approximation are presented.

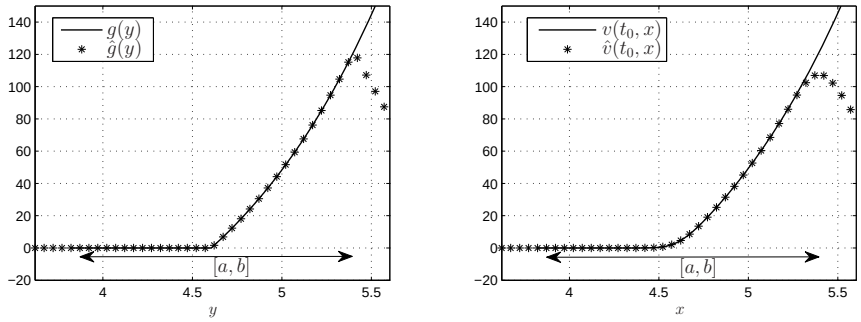

Figure 4.1: Example of a significant error, $\epsilon_1 + \epsilon_3$, close to one of the boundaries.

The numerical option value for initial log-asset price, x_0, is highly accurate, with an error less than 10^{-14}. However, inaccuracies at some place in the domain $[a, b]$ may seriously affect the backward recursion of the Fourier-coefficients, $V_k(t_m)$, when a stochastic control problem is solved recursively.

4.2 Extrapolation method

The coefficients $V_k(t_m)$ are recovered recursively, backwards in time. In that case, the local error, ϵ_{COS}, described in the previous section, may propagate through time. Here, we propose a technique to deal with this issue. In Section 4.3, we prove that the error of the approximated Fourier coefficients converges exponentially in N, for probability density functions in the class $C^\infty([a, b])$.

Recall that the approximated Fourier cosine coefficients at time t_{m-1} are given by

$$\hat{V}_k(t_{m-1}) \quad = \quad \frac{2}{b-a}\int_a^b \Big[\max_{\alpha_{m-1}\in A}[F(t_{m-1}, x, \alpha_{m-1}) + \hat{c}(t_{m-1}, x, \alpha_{m-1})]\Big] \cos\left(k\pi\frac{x-a}{b-a}\right) dx.$$

From this definition it follows that inaccurate numerical continuation values at time t_{m-1} may affect the choice of the optimal control value and the coefficients, which in turn affect the continuation value at time t_{m-2}, and so on.

The idea to deal with the propagating error is to determine the area in which inaccurate approximate values from the COS method occur. In this area, we employ an *extrapolation technique* to compute a value with improved accuracy, using the accurate numerical continuation values from the neighboring region.

In practical applications, it may be possible to determine the area in which $\hat{c}(t_{m-1}, x, \alpha)$ is inaccurate, assuming that coefficients $V_k(t_m)$ are exact. The density function, together with the value function, give the desired information. For instance, suppose we can calculate a value x^*, so

that the continuation value is well approximated for $x \in [x^*, b]$ and is inaccurate for $x \in [a, x^*]$ [4]. The continuation function $c(t_{m-1}, x, \alpha)$, on $[a, x^*]$, can then be approximated by an extrapolation technique. For this, we employ a second-order Taylor expansion in x^*:

$$c^{ex}(t_{m-1}, x, \alpha) := \hat{c}(t_{m-1}, x^*, \alpha) + \hat{c}_x(t_{m-1}, x^*, \alpha)(x - x^*) + \tfrac{1}{2}\hat{c}_{xx}(t_{m-1}, x^*, \alpha)(x - x^*)^2. \quad (4.16)$$

The derivatives can easily be computed in this setting, as:

$$\hat{c}_x(t_{m-1}, x^*, \alpha) = e^{-\rho\Delta t} {\sum_{k=0}^{N-1}}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{k\pi}{b-a}\Big|\alpha\right) e^{ik\pi\frac{x^*-a}{b-a}} \frac{ik\pi}{b-a}\right) V_k(t_m), \quad (4.17)$$

$$\hat{c}_{xx}(t_{m-1}, x^*, \alpha) = e^{-\rho\Delta t} {\sum_{k=0}^{N-1}}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{k\pi}{b-a}\Big|\alpha\right) e^{ik\pi\frac{x^*-a}{b-a}} \left(\frac{ik\pi}{b-a}\right)^2\right) V_k(t_m). \quad (4.18)$$

We denote the extrapolated continuation value by

$$\tilde{c}(t_{m-1}, x, \alpha) := \begin{cases} c^{ex}(t_{m-1}, x, \alpha), & \text{for } x \in [a, x^*], \\ \hat{c}(t_{m-1}, x, \alpha), & \text{for } x \in [x^*, b]. \end{cases} \quad (4.19)$$

The local error of the COS formula *with* extrapolation technique is denoted by

$$\tilde{\epsilon}_{COS}(t_{m-1}, x, \alpha|[a, b], N) := c(t_{m-1}, x, \alpha) - \tilde{c}(t_{m-1}, x, \alpha|[a, b], N), \quad (4.20)$$

and we have:

$$\tilde{\epsilon}_{COS}(t_{m-1}, x, \alpha|[a, b], N) = O((x - x^*)^3), \qquad \text{for } x \in [a, x^*]. \quad (4.21)$$

We use continuation value $\tilde{c}$ to determine the optimal control law and to approximate the terms C_k by

$$\tilde{C}_k(t_{m-1}, z_1, z_2, \alpha) := \frac{2}{b-a}\int_{z_1}^{z_2} \tilde{c}(t_{m-1}, y, \alpha)\cos\left(k\pi\frac{y-a}{b-a}\right) dy. \quad (4.22)$$

The corresponding Fourier coefficients are denoted by $\tilde{V}_k(t_{m-1})$.

Suppose the interval $[z_1, z_2] \subset [a, b]$ can be divided into $[z_1, x^*]$ and $[x^*, z_2]$. Then the corresponding Fourier cosine coefficients read

$$\begin{aligned} \tilde{C}_k(t_{m-1}, z_1, z_2, \alpha) &= \frac{2}{b-a}\int_{z_1}^{x^*} c^{ex}(t_{m-1}, x^*, \alpha)\cos\left(k\pi\frac{y-a}{b-a}\right) dy \\ &+ e^{-\rho\Delta t}\mathrm{Re}\left({\sum_{j=0}^{N-1}}' \varphi_{levy}\left(\frac{j\pi}{b-a}\Big|\alpha\right) V_j(t_m) \cdot M_{k,j}(x^*, z_2)\right), \end{aligned} \quad (4.23)$$

where the analytic solution to the first part can be found in Appendix B. The extrapolation technique can be improved by using a higher-order Taylor expansion. In that case, a similar approach can be applied.

Remark 4.2. *If we know that the continuation value should be of exponential form, which is sometimes the case for the type of problems we are interested in, it may be more accurate to use an exponential extrapolation to approximate function $c(t_{m-1}, x, \alpha)$ on $[a, x^*]$, as follows:*

$$c^{ex}(t_{m-1}, x, \alpha) := \hat{c}(t_{m-1}, x^*, \alpha)\exp\Big(w(y - x^*)\Big), \qquad \textit{for } x \in [a, x^*], \quad (4.24)$$

with

$$w = \frac{\hat{c}_x(t_{m-1}, x^*, \alpha)}{\hat{c}(t_{m-1}, x^*, \alpha)}. \quad (4.25)$$

We will also use this form in an example in Section 5.2.

[4] The methodology above can also be applied if the approximated continuation value is inaccurate in a certain area $[x^{**}, b]$.

4.3 Error propagation in the backward recursion

In this section, we study the error of the Fourier cosine coefficients in the backward recursion. We start with the algorithm *without the extrapolation technique* and define

$$\varepsilon_k(t_m, z_1, z_2, \alpha) := C_k(t_m, z_1, z_2, \alpha) - \hat{C}_k(t_m, z_1, z_2, \alpha). \tag{4.26}$$

The terms U_k are assumed to be exact, so that the error in the Fourier coefficients is given by

$$\varepsilon_k(t_m) := V_k(t_m) - \hat{V}_k(t_m) = \sum_{q=1}^{K} \varepsilon_k(t_m, \mathcal{D}_m^q, \alpha_m^q). \tag{4.27}$$

We assume that an accurate, true control law [5] is found, which is denoted by α_m^*. Here the dependency on the state value is omitted. The error may be larger in case of incorrect control values.

At time lattice $M-1$ we have:

$$\begin{aligned} \varepsilon_k(t_{M-1}, z_1, z_2, \alpha) &= \frac{2}{b-a}\int_{z_1}^{z_2} (c(t_{M-1}, y, \alpha) - \hat{c}(t_{M-1}, y, \alpha)) \cos\left(k\pi \frac{y-a}{b-a}\right) dy \\ &= \frac{2}{b-a}\int_{z_1}^{z_2} \epsilon_{COS}(t_{M-1}, y, \alpha) \cos\left(k\pi \frac{y-a}{b-a}\right) dy. \end{aligned} \tag{4.28}$$

Here we omit the dependency of the error on the interval $[a, b]$ and N. Coefficients $V_k(t_M)$ are assumed to be known analytically, so that the only error introduced by the COS formula is the local error, ϵ_{COS}. This error may be significant close to the domain boundaries a and b. We end up with

$$\varepsilon_k(t_{M-1}) = \frac{2}{b-a}\int_a^b \epsilon_{COS}(t_{M-1}, y, \alpha_{M-1}^*) \cos\left(k\pi \frac{y-a}{b-a}\right) dy. \tag{4.29}$$

Note that the terms $\varepsilon_k(t_{M-1})$ are the Fourier cosine coefficients of the local error at time t_{M-1} on the expansion interval $[a, b]$.

For the approximation of $c(t_{M-2}, x, \alpha)$ the COS formula with coefficients $\hat{V}_k(t_{M-1})$ is used. The approximated value is denoted by $\overline{c}(t_{M-2}, x, \alpha)$. The use of the approximations $\hat{V}_k(t_{M-1})$ in equation (3.23) gives rise to an additional error in $C_k(t_{M-2}, z_1, z_2, \alpha)$:

$$\varepsilon_k(t_{M-2}, z_1, z_2, \alpha) = \frac{2}{b-a}\int_{z_1}^{z_2} (c(t_{M-2}, y, \alpha) - \overline{c}(t_{M-2}, y, \alpha)) \cos\left(k\pi \frac{y-a}{b-a}\right) dy, \tag{4.30}$$

with $\overline{c}$ obtained by inserting $\hat{V}_k(t_{M-1})$ in the COS formula:

$$\begin{aligned} \overline{c}(t_{M-2}, y, \alpha) &= e^{-\rho\Delta t} {\sum_{j=0}^{N-1}}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{j\pi}{b-a}\middle|\alpha\right) e^{ij\pi\frac{y-a}{b-a}}\right) \hat{V}_j(t_{M-1}) \\ &= e^{-\rho\Delta t} {\sum_{j=0}^{N-1}}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{j\pi}{b-a}\middle|\alpha\right) e^{ij\pi\frac{y-a}{b-a}}\right) (V_j(t_{M-1}) - \varepsilon_j(t_{M-1})) \\ &= \hat{c}(t_{M-2}, y, \alpha) - e^{-\rho\Delta t} {\sum_{j=0}^{N-1}}' \mathrm{Re}\left(\varphi_{levy}\left(\frac{j\pi}{b-a}\middle|\alpha\right) e^{ij\pi\frac{y-a}{b-a}}\right) \varepsilon_j(t_{M-1}). \end{aligned} \tag{4.31}$$

[5] which is an optimistic assumption if the continuation value is inaccurate.

The error in the coefficients C_k can now be separated into two parts:

$$\begin{aligned}
&\varepsilon_k(t_{M-2}, z_1, z_2, \alpha) \\
&= \frac{2}{b-a}\int_{z_1}^{z_2} (c(t_{M-2}, y, \alpha) - \hat{c}(t_{M-2}, y, \alpha) + \hat{c}(t_{M-2}, y, \alpha) - \overline{c}(t_{M-2}, y, \alpha)) \cos\left(k\pi\frac{y-a}{b-a}\right) dy \\
&= \frac{2}{b-a}\int_{z_1}^{z_2} (\epsilon_{COS}(t_{M-2}, y, \alpha) + \overline{\epsilon}(t_{M-2}, y, \alpha)) \cos\left(k\pi\frac{y-a}{b-a}\right) dy, \qquad (4.32)
\end{aligned}$$

where

$$\overline{\epsilon}(t_{M-2}, y, \alpha) = e^{-\rho\Delta t}\sum_{j=0}^{N-1}{}' \operatorname{Re}\left(\varphi_{levy}\left(\frac{j\pi}{b-a}\Big|\alpha\right) e^{ij\pi\frac{y-a}{b-a}}\right)\varepsilon_j(t_{M-1}). \qquad (4.33)$$

$\overline{\epsilon}$ resembles the COS formula (3.13), now with Fourier coefficients $\varepsilon_j(t_{M-1})$. For sufficiently large value of N we find

$$\begin{aligned}
&\overline{\epsilon}(t_{M-2}, y, \alpha) \\
&= e^{-\rho\Delta t}\int_{[a,b]} \epsilon_{COS}(t_{M-1}, z, \alpha^*_{M-1}) f(z|y, \alpha) dz + e^{-\rho\Delta t}\int_{\mathbb{R}\setminus[a,b]} \hat{\epsilon}_{COS}(t_{M-1}, z, \alpha^*_{M-1}) f(z|y, \alpha) dz \\
&\approx e^{-\rho\Delta t}\mathbb{E}^{t_{M-2}, y}[\epsilon_{COS}(t_{M-1}, X_{t_{M-1}}, \alpha^*_{M-1})], \qquad (4.34)
\end{aligned}$$

where $\hat{\epsilon}_{COS}$ is the Fourier cosine expansion of the local error. So,

$$c(t_{M-2}, y, \alpha) - \overline{c}(t_{m-1}, y, \alpha) \approx \epsilon_{COS}(t_{M-2}, y, \alpha|[a,b], N) + e^{-\rho\Delta t}\mathbb{E}^{t_{M-2}, y}[\epsilon_{COS}(t_{M-1}, X_{t_{M-1}}, \alpha^*_{M-1})]. \qquad (4.35)$$

The recursive algorithm gives

$$\varepsilon_k(t_m) \approx \frac{2}{b-a}\int_a^b (\epsilon_{COS}(t_m, y, \alpha^*_m) + e^{-\rho\Delta t}\mathbb{E}^{t_m, y}[\epsilon_{COS}(t_{m+1}, X_{t_{m+1}}, \alpha^*_{m+1})]) \cos\left(k\pi\frac{y-a}{b-a}\right) dy. \qquad (4.36)$$

The first part of the error is due to the use of the COS formula at time t_m. The possible propagation of errors from time level t_{m+1} causes the second part to appear. The value of the expectation depends on the drift and diffusion of the stochastic process[6]. A clarifying example is provided in Section 5.2.1. It is not possible to bound the error $\varepsilon_k(t_m)$ by increasing the number of terms in the Fourier series expansions, N, as the error $\epsilon_1 + \epsilon_3$ remains. The propagation of the error may give rise to incorrect results of the algorithm.

Next we discuss the error convergence if we employ the extrapolation methodology from Section 4.2. The error in the terms C_k is redefined by

$$\varepsilon_k(t_m, z_1, z_2, \alpha) := C_k(t_m, z_1, z_2, \alpha) - \tilde{C}_k(t_m, z_1, z_2, \alpha) \qquad (4.37)$$

and

$$\varepsilon_k(t_m) := V_k(t_m) - \tilde{V}_k(t_m) = \sum_{q=1}^{K} \varepsilon_k(t_m, \mathcal{D}^q_m, \alpha^q_m). \qquad (4.38)$$

Result 4.1. *With a sufficiently accurate extrapolation technique, with $[a,b] \subset \mathbb{R}$ chosen sufficiently wide and a probability density function f in $C^\infty([a,b])$, error $\varepsilon_k(t_m)$ converges exponentially in N for $1 \le m \le M-1$.*

The proof of this result is similar to that for pricing Bermudan options, which can be found in [7]. Note that error $\epsilon_1 + \epsilon_3$ has now been reduced. It can also be proved that if the local error converges algebraically, then so does ε.

[6] A large error ϵ_{COS} at the left-hand side of the computational domain may 'disappear' by a large positive drift. It may travel through the domain for other drift terms.

5 Examples

In this section, we apply Algorithm 1 to two different stochastic control problems. In the first example, we calculate the price of a butterfly option under uncertain volatility. Because we are dealing with zero option values at both boundaries, interval $[a, b]$ can be chosen sufficiently large and extrapolation of the continuation value is not required for accuracy.

The second example deals with an optimal consumption-portfolio problem, with which we can demonstrate the impact of a significant error near the spatial boundaries, and its propagation. The continuous-time variant of this stochastic control problem admits an analytic solution. This problem is thus instructive as we can show the propagation of a local error and the improvement by extrapolation of the continuation value.

5.1 Butterfly option under uncertain volatility

The model we use for pricing a butterfly options under uncertain volatility is based on the problem described in [16]. The setting of this problem is the financial option market. The dynamics of the asset price is assumed to evolve according to the following dynamics:

$$dS_s = rS_s ds + \boldsymbol{\alpha}_s S_s dW_s, \qquad S_0 \text{ given.} \tag{5.1}$$

$(\boldsymbol{\alpha}_s)_{0\le s\le T}$ is an uncertain volatility process, which is valued in the interval $[\alpha^-, \alpha^+]$ and r is the risk-neutral interest rate. We consider the worst case for an investor with a long position in a European-style option. Then, the value function reads:

$$v(t,S) \quad = \quad \inf_{\boldsymbol{\alpha}\in\mathcal{A}} J(t,S,\boldsymbol{\alpha}) = \inf_{\boldsymbol{\alpha}\in\mathcal{A}} \mathbb{E}[e^{-r(T-t)} g(S_T)], \tag{5.2}$$

where $g(.)$ is the payoff function of a butterfly option at terminal time T, which is given by:

$$g(S) = (S-K_1)^+ - 2\left(S - \frac{K_1+K_2}{2}\right)^+ + (S-K_2)^+, \tag{5.3}$$

for certain strike prices K_1 and K_2. The pricing problem is now formulated as a stochastic control problem, whereby the process $\boldsymbol{\alpha}_s$ is the control process, with values in $A = [\alpha^-, \alpha^+]$. Note that there is no running profit function included in this problem formulation, but only a terminal reward function.

Contrary to the problem in equation (2.3), the infimum over gain functions $J(.)$ is taken. However, similar theory and solution algorithms can be developed for minimization problems.

Remark 5.1. *The corresponding Hamilton-Jacobi-Bellman equation reads*

$$-\frac{\partial v}{\partial t}(t,S) + rv(t,S) - \min_{\alpha\in[\alpha^+,\alpha^-]}\left[rS\frac{\partial v}{\partial S}(t,S) + \frac{1}{2}\alpha^2 S^2 \frac{\partial^2 v}{\partial S^2}(t,S)\right] = 0, \qquad \forall (t,S)\in[0,T)\times\mathbb{R}_+, \tag{5.4}$$

which yields

$$\begin{aligned} &\text{if } \tfrac{\partial^2 v}{\partial S^2} \le 0 \quad \Rightarrow \text{take } \alpha = \alpha^+, \\ &\text{if } \tfrac{\partial^2 v}{\partial S^2} > 0 \quad \Rightarrow \text{take } \alpha = \alpha^-. \end{aligned} \tag{5.5}$$

This allows us to restrict the set of possible control values to $A = \{\alpha^-, \alpha^+\}$. We see that the control value, that is the volatility, is a function of the Greek $\Gamma = \frac{\partial^2 v}{\partial S^2}$.

The same partial differential equation is derived for a transaction costs model in [9].

As before, we consider an equidistant time grid, $t_0, t_1, \ldots, t_M = T$, with $\Delta t := t_m - t_{m-1}$. In the numerical approximation a constant volatility $\alpha_m \in \{\alpha^-, \alpha^+\}$ is applied within the time intervals $[t_m, t_{m+1}]$. We switch to the log-asset price process $X_s = \log S_s$, so

$$dX_s = (r - \tfrac{1}{2}\alpha_m^2)ds + \alpha_m dW_s, \qquad \text{for } s \in [t_m, t_{m+1}]. \tag{5.6}$$

We use here the dynamic programming principle to determine

$$\begin{aligned} v(t_{m-1}, x) &= \min_{\alpha_{m-1} \in \{\alpha^-, \alpha^+\}} e^{-\rho \Delta t} \mathbb{E}[v(t_m, X_{t_m}) | X_{t_{m-1}} = x, \alpha_{m-1}] \\ &= \min\left[c(t_{m-1}, x, \alpha^-),\, c(t_{m-1}, x, \alpha^+)\right], \end{aligned} \tag{5.7}$$

with continuation value

$$\begin{aligned} c(t_{m-1}, x, \alpha) &= e^{-\rho \Delta t} \mathbb{E}[v(t_m, X_{t_m}) | X_{t_{m-1}} = x, \alpha_{m-1} = \alpha] \\ &\approx e^{-\rho \Delta t} {\sum_{k=0}^{N-1}}' \operatorname{Re}\left(\varphi_{levy}\left(\frac{k\pi}{b-a}\Big|\alpha\right) e^{ik\pi \frac{x-a}{b-a}}\right) V_k(t_m) := \hat{c}(t_{m-1}, x, \alpha). \end{aligned} \tag{5.8}$$

The characteristic function, $\varphi_{levy}(u|\alpha) = \exp(iu(r - \frac{1}{2}\alpha^2)\Delta t - \frac{1}{2}u^2\alpha^2\Delta t)$, depends on the volatility α.

Algorithm 1 is used to solve the pricing problem. The coefficients at the terminal time are known analytically

$$\begin{aligned} V_k(t_M) &:= \frac{2}{b-a}\int_a^b g(e^y)\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &= \frac{2}{b-a}\Big[\chi_k(\log K_1, b, a, b) - 2\chi_k\Big(\log\frac{K_1+K_2}{2}, b, a, b\Big) + \chi_k(\log K_2, b, a, b) \\ &\quad -K_1\psi_k(\log K_1, b, a, b) + (K_1+K_2)\psi_k\Big(\log\frac{K_1+K_2}{2}, b, a, b\Big) + K_2\psi_k(\log K_2, b, a, b)\Big], \\ &\quad (a \le \log K_1, \log K_2 \le b), \end{aligned} \tag{5.9}$$

see Appendix A for the analytic solution to the functions χ_k and ψ_k. For the other time levels the Fourier coefficients are approximated by

$$\hat{V}_k(t_m) = \frac{2}{b-a}\int_a^b \min\left[\hat{c}(t_m, y, \alpha^-), \hat{c}(t_m, y, \alpha^+)\right]\cos\left(k\pi\frac{y-a}{b-a}\right)dy. \tag{5.10}$$

It is worth mentioning that we do not need to use the extrapolated value $\tilde{V}_k$, with function $\tilde{c}$, from Section 4.2. The reason for this is that the value function converges to zero if the log-asset price goes to plus or minus infinity. So, for sufficiently large intervals $[a, b]$, the value function on time lattice t_{m+1} is almost zero outside the expansion interval. Then, by assuming that N is chosen sufficiently large, the function $\hat{v}(t_{m+1}, x)$ is also accurate outside $[a, b]$ and the local error of the COS formula at time t_m is small for all $y \in [a, b]$. Because of this, extrapolation for the continuation value is not necessary.

We can divide the integration interval $[a, b]$ into sub-domains $\mathcal{D}_m^-$ and $\mathcal{D}_m^+$, for which the optimal control values at control time t_m are α_m^- and α_m^+, respectively:

$$\begin{aligned} \hat{V}_k(t_m) &= \frac{2}{b-a}\int_{\mathcal{D}_m^-}\hat{c}(t_m, y, \alpha_m^-)\cos\left(k\pi\frac{y-a}{b-a}\right)dy + \frac{2}{b-a}\int_{\mathcal{D}_m^+}\hat{c}(t_m, y, \alpha_m^+)\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &:= \hat{C}_k(t_m, \mathcal{D}_m^-, \alpha_m^-) + \hat{C}_k(t_m, \mathcal{D}_m^+, \alpha_m^+). \end{aligned} \tag{5.11}$$

Inserting the COS formula results in the following recursive formula for the coefficients $\hat{C}_k$:

$$\begin{aligned}\hat{C}_k(t_m, z_1, z_2, \alpha) &= \frac{2}{b-a}\int_{z_1}^{z_2} \hat{c}(t_m, y, \alpha)\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &\approx e^{-\rho\Delta t}\mathrm{Re}\left(\sum_{j=0}^{N-1}{}' \varphi_{levy}\left(\frac{j\pi}{b-a}\Big|\alpha\right)\hat{V}_j(t_{m+1})M_{k,j}(z_1, z_2)\right). \end{aligned} \tag{5.12}$$

In the numerical experiment we use the following model parameters:

$$T = 0.25,\ r = 0.1,\ K_1 = 90,\ K_2 = 110,\ S_0 = 100,\ \alpha^- = 0.15 \text{ and } \alpha^+ = 0.25. \tag{5.13}$$

For the integration interval we take $[a, b] = [\log 70, \log 130]$.

M	$\hat{v}(t_0, x_0)$
60	2.3178
120	2.3078
240	2.3027
480	2.3002
960	2.2990

(a) Option values ($N \geq 400$).

		N				
		400	600	800	1000	1200
	60	0.2997	0.3510	0.4365	0.5127	0.5801
	120	0.5062	0.6198	0.7412	0.8670	1.0544
M	240	0.9727	1.1883	1.3898	1.6326	1.8447
	480	1.9219	2.3336	2.7366	3.1954	3.5870
	960	3.8537	4.6722	5.3908	6.2907	7.1339

(b) CPU times (s).

Table 5.1: Results butterfly option pricing model.

In Tables 5.1a and 5.1b, results for different values of M and N are shown. The option values have converged in N up to nine decimal places, for $N \geq 400$. Increasing the number of control times, M, gives convergence to the true option value. The results are highly satisfactory and match the prices in [16] and [15]. The computation time is linear in M and $O(N \log_2 N)$ in the number of terms in the Fourier cosine series expansions. In Result 4.1, we deduced an exponentially converging error in N, assuming that the correct control law was found. In practice, we may however find incorrect control values for small values of N and the convergence result therefore holds for sufficiently large N.

5.2 Optimal consumption path

The second example we discuss is a simplified version of Merton's optimal consumption-portfolio problem [11]. Here, an agent consumes a proportion of his wealth and invests the remaining part in assets, with rate of return, μ, and fixed volatility, σ. The dynamics of the invested capital are given by

$$dK_s = \mu K_s ds + \sigma K_s dW_s. \tag{5.14}$$

Let $\boldsymbol{\alpha}_s Z_s$ denote the amount of wealth consumed at time s, with $\boldsymbol{\alpha}_s$ the control process and wealth Z_s. Taking into account equation (5.14) and consumption gives

$$\begin{aligned} dZ_s &= \frac{Z_s}{K_s}dK_s - \boldsymbol{\alpha}_s Z_s ds \\ &= (\mu - \boldsymbol{\alpha}_s)Z_s ds + \sigma Z_s dW_s. \end{aligned} \tag{5.15}$$

The agent chooses his consumption to maximize his expected discounted utility of consumption over a finite time horizon with terminal time T. The optimal consumption problem can be represented

by the stochastic control problem:

$$v(t,z) = \max_{\boldsymbol{\alpha}\in\mathcal{A}} \mathbb{E}^{t,z}\left[\int_t^T e^{-\rho(s-t)}U(\boldsymbol{\alpha}_s Z_s)ds + e^{-\rho(T-t)}U(Z_T)\right], \tag{5.16}$$

where z is the current wealth level, and $\rho \geq 0$ is the utility discount rate [3].

The *utility function*, $U(.)$, measures the utility gain of consumption $\boldsymbol{\alpha}_s Z_s$. We presume a constant relative risk aversion (CRRA) utility function

$$U(C) = C^\gamma/\gamma, \qquad \gamma \leq 1,\, \gamma \neq 0. \tag{5.17}$$

The exact solution to this continuous-time control problem can be found by the corresponding HJB equation and the verification theorem ([12], [14]):

$$v(t,z) = \frac{b(t)}{\gamma}z^\gamma \quad \text{with} \quad b(t) = \left(\frac{1+(\nu-1)\exp(\nu(t-T))}{\nu}\right)^{1-\gamma}, \tag{5.18}$$

where $\nu = (\rho - \gamma\mu - \frac{1}{2}\gamma(\gamma-1)\sigma^2)/(1-\gamma)$. The optimal consumption law is then given by

$$\boldsymbol{\alpha}^*(t) = b(t)^{1/(\gamma-1)}. \tag{5.19}$$

Note that this optimal control process is independent of current wealth level, z. For the tests in this section we choose the following set of parameters:

$$T = 100,\, \rho = 0.03,\, \gamma = -3,\, Z_0 = 100,\, \mu = 0.04,\, \sigma = 0.1. \tag{5.20}$$

Figure 5.1 shows the optimal control law. At the terminal time, T, the remaining wealth is assumed to be completely consumed. Based on economic arguments the control process $\boldsymbol{\alpha}^*(t)$ goes to one if the time approaches the terminal time. At earlier time levels the control process has reached a steady state.

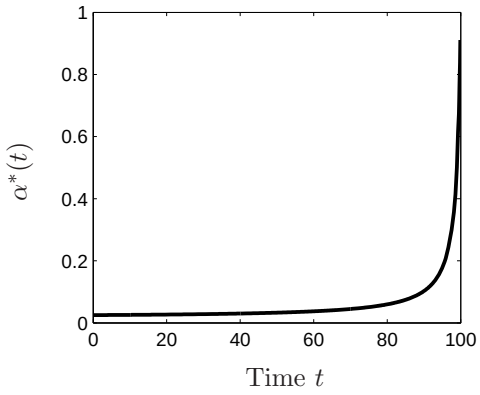

Figure 5.1: Optimal control law $\boldsymbol{\alpha}^*(t)$.

For the numerical approach we employ an equidistant grid of control times, $t_0, t_1, \ldots, t_M = T$, with $\Delta t := t_m - t_{m-1}$. At each control time, t_m, one can choose a constant fraction of wealth,

$\alpha_m \in A = [0,1]$, which is consumed during the time interval $[t_m, t_{m+1}]$[7]. We switch to the log-wealth process, $X_s = \log Z_s$, so

$$dX_s = (\mu - \alpha_m - \tfrac{1}{2}\sigma^2)ds + \sigma dW_s, \qquad \text{for } s \in [t_m, t_{m+1}]. \tag{5.21}$$

We rewrite the value function as

$$v(t,x) = \max_{\boldsymbol{\alpha}\in\hat{\mathcal{A}}} \mathbb{E}^{t,x}\left[\int_t^T e^{-\rho(s-t)} U(\boldsymbol{\alpha}_s e^{X_s})ds + e^{-\rho(T-t)}U(e^{X_T})\right], \tag{5.22}$$

where x is the current log-wealth level. $\hat{\mathcal{A}} \subset \mathcal{A}$ denotes the set of all possible control paths $\{\alpha_m\}_{m=0}^{M-1}$, where α_m is valued in the control set A. The dynamic programming principle gives us

$$\begin{aligned} v(t_{m-1},x) &= \max_{\alpha_{m-1}\in A} \mathbb{E}^{t_{m-1},x}\left[\int_{t_{m-1}}^{t_m} e^{-\rho(s-t)}U(\alpha_{m-1}e^{X_s})ds + e^{-\rho\Delta t}v(t_m, X_{t_m})\right] \\ &:= \max_{\alpha_{m-1}\in A}[F(x,\alpha_{m-1}) + c(t_{m-1},x,\alpha_{m-1})]. \end{aligned} \tag{5.23}$$

For the time-independent profit function we find

$$F(x,\alpha) = \frac{\alpha^\gamma}{\gamma\xi(\alpha)}(e^{\Delta t\xi(\alpha)} - 1)e^{\gamma x}, \quad \text{with} \quad \xi(\alpha) = -\rho + \gamma(\mu - \alpha - \tfrac{1}{2}\sigma^2) + \tfrac{1}{2}\gamma^2\sigma^2. \tag{5.24}$$

Applying the COS formula yields

$$\hat{c}(t_{m-1},x,\alpha) = e^{-\rho\Delta t}\sum_{k=0}^{N-1}{}' \operatorname{Re}\left(\varphi_{levy}\left(\frac{k\pi}{b-a}\Big|\alpha\right)e^{ik\pi\frac{x-a}{b-a}}\right)V_k(t_m), \tag{5.25}$$

where the characteristic function is given by

$$\varphi_{levy}(u|\alpha) = \exp\Big(iu(\mu - \alpha - \tfrac{1}{2}\sigma^2)\Delta t - \tfrac{1}{2}\sigma^2u^2\Delta t\Big). \tag{5.26}$$

We use Algorithm 1 to solve the discrete-time stochastic control problem (5.22). The coefficients at time t_M read

$$V_k(t_M) = \frac{2}{b-a}\int_a^b e^{\gamma y}/\gamma \cos\left(k\pi\frac{y-a}{b-a}\right)dy = \frac{2}{b-a}\frac{1}{\gamma^2}\chi_k(a\gamma, b\gamma, a\gamma, b\gamma), \tag{5.27}$$

with the analytic function χ_k in Appendix A.

We divide the interval $[a,b]$ into 200 equally-sized sub-intervals $\mathcal{D}_m^q$ and approximate the optimal control value for the midpoint of each sub-interval, assuming that this value is an accurate approximation for the entire sub-interval.

The terms U_k are time-independent and known analytically:

$$\begin{aligned} U_k(z_1,z_2,\alpha) &= \frac{2}{b-a}\int_{z_1}^{z_2} F(y,\alpha)\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &= \frac{\alpha^\gamma}{\gamma^2\xi(\alpha)}(e^{\Delta t\xi(\alpha)} - 1)\frac{2}{b-a}\chi_k(z_1\gamma, z_2\gamma, a\gamma, b\gamma). \end{aligned} \tag{5.28}$$

[7]For this problem, the true optimal control values are in the set $A = [0,1]$. A wider control set only makes the difficulties that we describe here even more severe.

The coefficients $\hat{C}_k$ are recovered by the coefficients $\hat{V}_k$ from the next time level and the FFT algorithm, as explained in Section 3.2.

We first show in Section 5.2.1 how the local error of the COS formula propagates backwards in time. Then the extrapolation methodology from Section 4.2 is applied to improve the solution (in Section 5.2.2). The following parameters are used for the tests in the next subsections:

$$[a,b] = [-2,\, 8],\, N = 2^{10},\, M = 100. \tag{5.29}$$

5.2.1 Example of a propagating error

Here, we show how local errors propagate when we do not apply extrapolation. Parameter N is chosen sufficiently large, so that error ϵ_2 can be neglected. From the error analysis in Section 4.1 it follows that

$$\hat{c}(t_{M-1},x,\alpha|[a,b],N) = e^{-\rho\Delta t}\int_{[a,b]} v(t_M,y)f(y|x,\alpha)dy + e^{-\rho\Delta t}\int_{\mathbb{R}\setminus[a,b]} \hat{v}(t_M,y)f(y|x,\alpha)dy. \tag{5.30}$$

Function $v(t_M,y)$ resembles $\hat{v}(t_M,y)$ on $[a,b]$ for sufficiently large values of N. However, the approximated continuation value is inaccurate when error $\epsilon_1+\epsilon_3$ in (4.9) is large. The inaccuracy of $\hat{c}(t_{M-1},x,\alpha)$ may give rise to a propagating error. First of all, it may result in an incorrect control value from the maximization operator in (3.14). Secondly, an inaccurate value function and coefficients $\hat{V}_k(t_{M-1})$ give rise to inaccurate numerical continuation values at time t_{M-2}. This is demonstrated by four plots in Figure 5.2.

The upper-left plot presents the terminal reward function, the Fourier cosine expansion, and its extension outside the expansion interval, which is symmetric in $a=-2$ and different from the correct terminal reward function at the left-hand side of a. In the lower-left plot the continuation value is shown for two control values, $\alpha=0.7$ and $\alpha=1.0$. The exact optimal control value for the discrete-time stochastic control problem at time t_{M-1} is $\alpha_{M-1}\approx 0.7$. It is independent of the current log-wealth level. From Figure 5.2 it is clear that significant errors occur in the vicinity of the left-side domain boundary.

The addition of the profit function gives us the graphs in the upper-right plot. As the value function at time t_{M-1} is defined by the maximization operator over all $\alpha_{M-1}\in A$, inaccurate control values will be determined, at least for $x<-1$. Because of this, the approximated value function will be too high. The lower-left plot shows the continuation values at time t_{M-2} for two different control values. The local error at time level $M-1$ has propagated and, in addition, errors from the COS formula occur in the vicinity of the boundary. The correct optimal control value equals $\alpha_{M-2}\approx 0.4$ and, again, incorrect control values will be determined by the maximization operator, if extrapolation is not used here.

The solution for all time steps gives rise to the optimal control values in Figure 5.3a. The correct control law should be independent of x, which is clearly not the case.

5.2.2 Improvement by extrapolation

We use the extrapolation technique from Section 4.2 to deal with the propagating error.

The function $f(y|x,\alpha)$ represents a normal density function of a random variable with distribution

$$\mathcal{N}\Big(x+(\mu-\alpha-\tfrac{1}{2}\sigma^2)\Delta t,\sigma\sqrt{\Delta t}\Big). \tag{5.31}$$

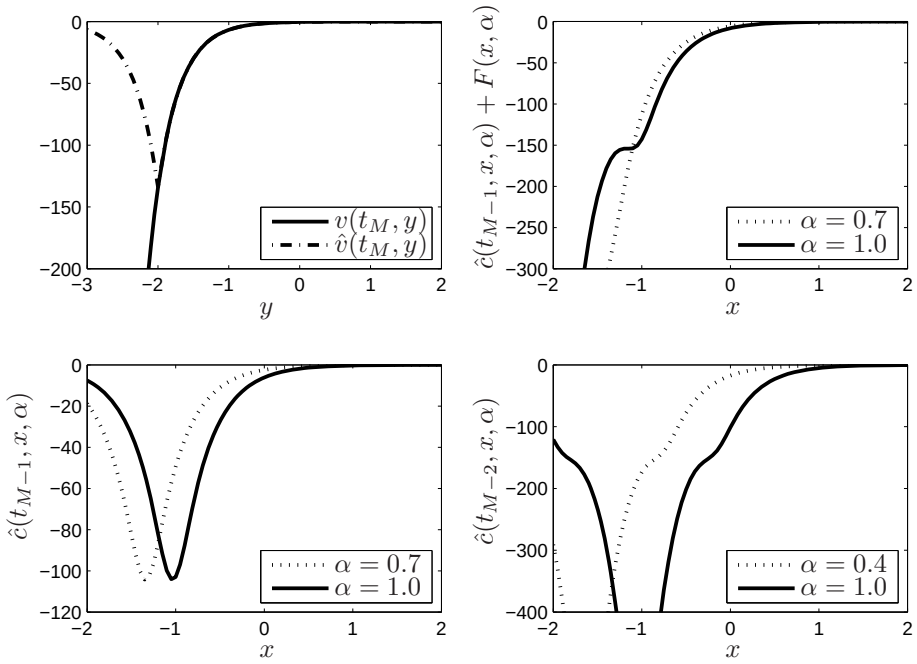

Figure 5.2: The propagation of local errors ($[a, b] = [-2, 8]$).

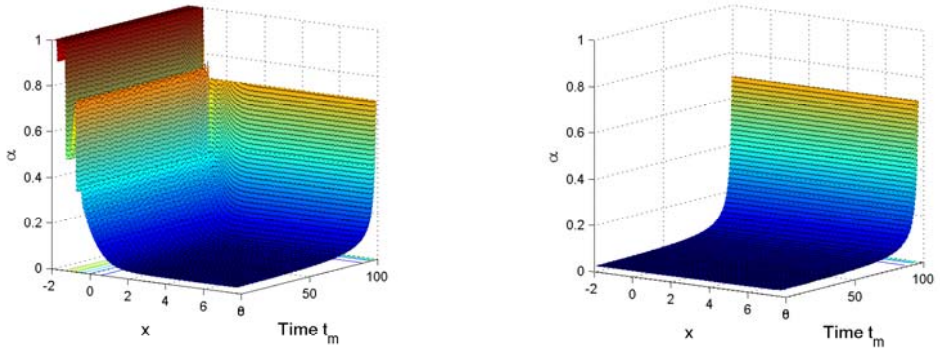

(a) Incorrect values, due to local errors and their propagation.

(b) Correct values, with extrapolation method.

Figure 5.3: Optimal control laws.

We can presume that the continuation value is well approximated on $[x^*, b]$, with

$$x^* := a - (\mu - \alpha - \tfrac{1}{2}\sigma^2)\Delta t + 5\sigma\sqrt{\Delta t}. \tag{5.32}$$

The terminal reward function, the density, and profit functions are all of exponential form. Therefore, we approximate $c(t_{m-1}, x, \alpha)$ on $[a, x^*]$ by employing an *exponential extrapolation technique* [8]:

$$c^{ex}(t_{m-1}, x, \alpha) := \hat{c}(t_{m-1}, x^*, \alpha) \exp\Big(w(y - x^*)\Big), \qquad \text{for } x \in [a, x^*], \tag{5.33}$$

[8] Note that polynomial extrapolation will also work well here.

with

$$w = \frac{\hat{c}_x(t_{m-1}, x^*, \alpha)}{\hat{c}(t_{m-1}, x^*, \alpha)}. \tag{5.34}$$

As proposed in Section 4.2, we will use the improved continuation value, $\tilde{c}(t_{m-1}, x, \alpha)$, to find the optimal control values and to approximate the coefficients $C_k(t_{m-1}, z_1, z_2, \alpha)$. Suppose the interval $[z_1, z_2] \subset [a, b]$ can be divided into $[z_1, x^*]$ and $[x^*, z_2]$. Then the corresponding Fourier cosine coefficients read

$$\begin{aligned} \tilde{C}_k(t_{m-1}, z_1, z_2, \alpha) &= \frac{2}{b-a} \int_{z_1}^{x^*} c^{ex}(t_{m-1}, x^*, \alpha) \cos\left(k\pi \frac{y-a}{b-a}\right) dy \\ &+ e^{-\rho \Delta t} \mathrm{Re}\left(\sum_{j=0}^{N-1}{}' \varphi_{levy}\left(\frac{j\pi}{b-a}\Big|\alpha\right) V_j(t_m) \cdot M_{k,j}(x^*, z_2) \right), \end{aligned} \tag{5.35}$$

where

$$\frac{2}{b-a} \int_{z_1}^{x^*} c^{ex}(t_{m-1}, x^*, \alpha) \cos\left(k\pi \frac{y-a}{b-a}\right) dy = \frac{2}{b-a} \hat{c}(t_{m-1}, x^*, \alpha) \frac{e^{-wx^*}}{w} \chi_k(wz_1, wx^*, wa, wb). \tag{5.36}$$

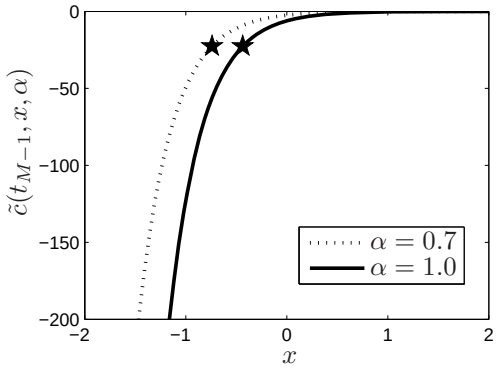

Figure 5.4: Extrapolation of the continuation value ($\bigstar = x^*$).

Figure 5.4 shows the values x^* and the improved continuation values for $\alpha = 0.7$ and $\alpha = 1.0$. The values are accurate, even in the vicinity of boundary $a = -2$, and the correct optimal control values will be determined, see Figure 5.3b for the complete result.

The exact solution to the continuous-time stochastic control problem is $v(t_0, x_0) = -0.8419$. For the discrete-time variant we find $\hat{v}(t_0, x_0) = -0.8419$. Table 5.2 shows the value function for different numbers of control times, M. They converge to the true value. We can conclude that the procedure with exponential extrapolation works highly satisfactorily.

M	10	25	50	100	150	200
$\hat{v}(t_0, x_0)$	-0.8843	-0.8548	-0.8459	-0.8430	-0.8424	-0.8422

Table 5.2: Results for different values of M ($N = 2^{10}$).

Remark 5.2. *The dynamics of the invested capital, equation (5.14), can easily be extended to a jump-diffusion process. Then no analytical solution to the stochastic control problem is available, but we can apply the COS method, as the characteristic function of a jump-diffusion process is known analytically, with the extrapolation technique.*

6 Conclusion

In this paper, we presented a general approach for solving stochastic control problems under a one-dimensional Lévy process. The method relies on the dynamic programming principle and the COS formula ([6]), which is based on Fourier cosine series expansions. A recursive algorithm has been defined based on the recursive recovery of the series coefficients. With the use of a Fast Fourier Transform algorithm we reach a computational complexity of order $O(N \log_2 N)$, per time step, where N denotes the number of terms in the series expansions.

We provided an extensive error analysis, with which we acquired knowledge about the origin and evolution of errors. We demonstrated how significant errors of the COS formula in the vicinity of domain boundaries may arise, and how they may propagate backwards in time. This understanding enabled us to improve the method by introducing an extrapolation method for the area in which the COS formula may give inaccurate continuation values. Extrapolation by Taylor expansion or by exponential extrapolation can easily be applied as the derivatives of approximated continuation values can be computed easily based on the COS formula. An exponentially converging error, in N, is found for a sufficiently accurate extrapolation method, $[a, b] \subset \mathbb{R}$ sufficiently wide and a probability density function in the class $C^\infty([a, b])$.

In [7], the COS method has been employed for pricing Bermudan and barrier options. A difference in using the COS method for stochastic control problems is the dependence of the characteristic function of the stochastic process on the control value, which may be time-consuming. Besides, we need to determine the optimal control law from a finite or continuous control set, for the entire spatial domain. Therefore, the approximated value function needs to be accurate over the complete spatial domain, which is not always the case when using the COS formula. These difficulties have been solved by the extrapolation technique.

We tested our numerical method by two examples, a butterfly option under uncertain volatility and an optimal consumption-portfolio problem. The COS method for stochastic control problems performed highly satisfactory.

Many other problems from finance and the real options context can be represented as a stochastic control problem. This makes our methodology applicable to various practical problems.

Appendix

A Functions χ_k and ψ_k

The functions χ_k and ψ_k are given by:

$$\chi_k(z_1, z_2, a, b) = \int_{z_1}^{z_2} e^y \cos\left(k\pi \frac{y-a}{b-a}\right) dy \qquad \text{and} \qquad \psi_k(z_1, z_2, a, b) = \int_{z_1}^{z_2} \cos\left(k\pi \frac{y-a}{b-a}\right) dy \tag{A.1}$$

and admit the following analytic solutions (for example with Maple 14):

$$\begin{aligned} \chi_k(z_1, z_2, a, b) \quad = \quad & \frac{1}{1+\left(\frac{k\pi}{b-a}\right)^2} \left[\cos\left(k\pi \frac{z_2-a}{b-a}\right) e^{z_2} - \cos\left(k\pi \frac{z_1-a}{b-a}\right) e^{z_1} \right. \\ & \left. + \frac{k\pi}{b-a} \sin\left(k\pi \frac{z_2-a}{b-a}\right) e^{z_2} - \frac{k\pi}{b-a} \sin\left(k\pi \frac{z_1-a}{b-a}\right) e^{z_1}\right], \end{aligned} \tag{A.2}$$

$$\psi_k(z_1,z_2,a,b) = \begin{cases} \left[\sin\left(k\pi\frac{z_2-a}{b-a}\right) - \sin\left(k\pi\frac{z_1-a}{b-a}\right)\right]\frac{b-a}{k\pi}, & \text{for } k\neq 0, \\ z_2 - z_1, & \text{for } k = 0. \end{cases} \tag{A.3}$$

B Estimation of coefficients with extrapolation

Suppose

$$c^{ex}(t_{m-1},x,\alpha) := \hat{c}(t_{m-1},x^*,\alpha) + \hat{c}_x(t_{m-1},x^*,\alpha)(x-x^*) + \tfrac{1}{2}\hat{c}_{xx}(t_{m-1},x^*,\alpha)(x-x^*)^2. \tag{B.1}$$

Then

$$\begin{aligned} &\frac{2}{b-a}\int_{z_1}^{z_2} c^{ex}(t_{m-1},x^*,\alpha)\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &= \frac{2}{b-a}\psi_k(z_1,z_2,a,b)\Big(\hat{c}(t_{m-1},x^*,\alpha) - \hat{c}_x(t_{m-1},x^*,\alpha)x^* + \tfrac{1}{2}\hat{c}_{xx}(t_{m-1},x^*,\alpha)(x^*)^2\Big) \\ &+ \frac{2}{b-a}\xi_k(z_1,z_2,a,b)\Big(\hat{c}_x(t_{m-1},x^*,\alpha) - \hat{c}_{xx}(t_{m-1},x^*,\alpha)x^*\Big) \\ &+ \frac{2}{b-a}\bar{\xi}_k(z_1,z_2,a,b)\tfrac{1}{2}\hat{c}_{xx}(t_{m-1},x^*,\alpha), \end{aligned} \tag{B.2}$$

with

$$\begin{aligned} \xi_k(z_1,z_2,a,b) &= \int_{z_1}^{z_2} y\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &= \frac{b-a}{(k\pi)^2}\left[\cos\left(k\pi\frac{z_2-a}{b-a}\right)(b-a) - \cos\left(k\pi\frac{z_1-a}{b-a}\right)(b-a)\right. \\ &\quad \left. + k\pi\sin\left(k\pi\frac{z_2-a}{b-a}\right)z_2 - k\pi\sin\left(k\pi\frac{z_1-a}{b-a}\right)z_1\right] \end{aligned} \tag{B.3}$$

and

$$\begin{aligned} \bar{\xi}_k(z_1,z_2,a,b) &= \int_{z_1}^{z_2} y^2\cos\left(k\pi\frac{y-a}{b-a}\right)dy \\ &= 2\frac{b-a}{(k\pi)^3}\left[-k\pi z_2(b-a)\cos\left(k\pi\frac{z_2-a}{b-a}\right) + k\pi z_1(b-a)\cos\left(k\pi\frac{z_1-a}{b-a}\right)\right. \\ &\quad \left. + \Big((b-a)^2 - \frac{1}{2}(k\pi z_2)^2\Big)\sin\left(k\pi\frac{z_2-a}{b-a}\right) - \Big((b-a)^2 - \frac{1}{2}(k\pi z_1)^2\Big)\sin\left(k\pi\frac{z_1-a}{b-a}\right)\right]. \end{aligned}$$

References

[1] J. P. Boyd, *Chebyshev and Fourier spectral methods*, Courier Dover Publications, New York, 2001.

[2] M. G. Crandall, H. Ishii, P. L. Lions, A user's guide to viscosity solutions of second order partial differential equations, *Bulletin of the American mathematical society*, Vol. 27, 1-67, 1992.

[3] J. H. Cochrane, *Asset pricing*, Princeton University Press, Princeton, 2005.

[4] A. K. Dixit, R. S. Pindyck, *Investment under uncertainty*, Princeton University Press, Princeton, 1994.

[5] D. Duffie, J. Pan, K. J. Singleton, Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, Vol. 68, 1343-1376, 2000.

[6] F. Fang, C. W. Oosterlee, A novel pricing method for European options based on Fourier-cosine series expansions, *SIAM Journal on Scientific Comp*, Vol. 31(2), 826-848, 2008.

[7] F. Fang, C. W. Oosterlee, Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions, *Numerische Mathematik*, Vol. 114(1), 27-62, 2009.

[8] P. A. Forsyth, G. Labahn, Numerical methods for controlled Hamilton-Jacobi-Bellman PDEs in finance, *Journal of Computational Finance*, Vol 11(2), 1-44, 2007.

[9] H. E. Leland, Option pricing and replication with transaction costs, *The Journal of Finance*, Vol. 40(5), 1985.

[10] P. L. Lions, Optimal control of diffusion processes and Hamilton-Jacobi-Bellman equations part 2: viscosity solutions and uniqueness, *Communications in Partial Differential Equations*, Vol. 8(11), 1229-1276, 1983.

[11] R. C. Merton, Lifetime portfolio selection under uncertainty: The continuous-time case, *Review of Economics & Statistics*, Vol. 51(3), 247-257, 1969.

[12] R. C. Merton, *Continuous-time finance*, Blackwell, 1990.

[13] B. K. Øksendal, A. Sulem, *Applied stochastic calculus and jump diffusion*, Springer, Berlin, 2005.

[14] H. Pham, *Continuous-time stochastic control and optimization with financial applications*, Springer-Verlag, 2009.

[15] S. P. van der Pijl, C. W. Oosterlee, An ENO-based numerical method for the control of dike levels in continuous time, Internal report CWI, 2010.

[16] D. M. Pooley, P. A. Forsyth, K. R. Vetzal, Numerical convergence properties of option pricing PDEs with uncertain volatility, IMA Journal of Numerical Analysis, Vol. 23, 241-267, 2003.

[17] M. J. Ruijter, *Numerical treatment of stochastic control problems by Fourier-cosine series expansions: The dike height problem*, Master thesis, Technische Universiteit Delft, the Netherlands, 2010.

[18] S. E. Shreve, *Stochastic calculus for finance II, Continuous-Time models*, Springer, Berlin, 2008.

[19] B. Zhang, C. W. Oosterlee, An efficient pricing algorithm for swing options based on Fourier cosine expansions, Internal report TU Delft, 2010.