Efficient Valuation of Energy Derivatives Using the COS Method in a Markov-Modulated Framework

by

Andrea Paniccia

to obtain the degree of Master of Science in Applied Mathematics
at Delft University of Technology,

Student number:	5975468		
Project duration:	January 15, 2025 – November 25, 2025		
Thesis committee:	Dr. F. Fang,	TU Delft,	daily supervisor
	Prof.dr.ir. C. Vuik,	TU Delft,	responsible professor
	Prof.dr. A. Papapantoleon	TU Delft	

Abstract

Energy markets are characterized by unique features such as seasonality, mean reversion, and sudden price spikes, which make the valuation of related derivatives considerably more challenging than in traditional financial markets. This thesis investigates the use of the COS method, a Fourier-based numerical technique, for the efficient valuation of energy derivatives within a Markov-modulated framework.

The study begins with energy quanto options, whose payoff depends on two correlated underlyings incorporating jumps and regime-switching dynamics. By deriving the characteristic functions for two benchmark models and implementing the COS method, the results demonstrate high accuracy and significant computational speed-ups compared to FFT- and Monte Carlo-based benchmarks.

The research work then extends to electricity storage contracts, which feature early-exercise opportunities and operational constraints. In line with the methodology presented by [7], this section replicates their integration of the COS method into a dynamic programming framework to evaluate the option values across multiple contract types and volatility regimes. The obtained prices fall within the confidence intervals of the LSMC benchmark, confirming the robustness of the COS approach in handling path-dependent and constrained problems.

Finally, the pricing framework is generalized to electricity storage under a two-state Markov-modulated model, to better capture the behavior of electricity prices. The extended COS-based algorithm accurately reproduces complex probability densities and achieves stable convergence of the prices while maintaining high computational efficiency even under this more complex setting.

In conclusion, this thesis demonstrates that the COS method is a powerful and computationally efficient alternative to stochastic simulation methods for pricing energy derivatives. It is capable of handling jump dynamics, regime switching, early-exercise features and operational constraints, highlighting its potential for broader applications in energy finance.

Keywords: Energy derivatives, COS method, Markov-modulated models, Electricity

storage contracts, Energy quanto options.

Acknowledgements

I would like to begin by thanking my supervisor, Professor Fang Fang, who has been my greatest source of support throughout this journey. Her insightful ideas, valuable feedback, and extensive knowledge have been fundamental in shaping this thesis. Above all, I am profoundly thankful for her understanding and for always finding the words and solutions to reassure me and guide me back on track.

I would also like to thank the committee members, Professors Kees Vuik and Antonis Papapantoleon, for taking the time to be part of this committee. Another thank you goes to Professor Griselda Deelstra for her support in the initial phase of this project, and to Professor Kees Oosterlee for his contribution regarding electricity storage contracts.

My deepest gratitude goes to my parents, to whom I will be eternally thankful for allowing me to pursue my dreams, despite all the sacrifices they had to make to make this possible, and for always giving me the freedom to choose my own path and be myself. I would also like to thank the rest of my family, starting with my grandmothers, Teresa and Giuliana, my aunt Pina, my cousins Kikka and Yaya, and my uncle Adriano, for always supporting me and loving me unconditionally.

And now, finally, it's time to thank my friends. To my best and closest friends, Laura, Lucrezia, Enzo, Ledia, Ludovica, Valeria, Martina, Greta, Giada, Giorgia, Francesca, and Vittoria, thank you for always being my greatest supporters. You all welcomed me into your lives without ever making me feel out of place, walking beside me through every journey, whether near or far. You have been the warmest refuge for my insecurities and have always been in the front row to celebrate every one of my achievements. If I have never felt alone, it is thanks to you, who have given me countless memories to look back on, with the awareness that I will probably never find people as special as you. You are my greatest point of reference, a compass that guides me home whenever I am not.

Lastly, I would also like to thank all the people who have accompanied me throughout my university journey in Delft, our Dutch friends, Cecilia, and especially Ludovica, Eleonora, and Elisa, who have been with me since the very first day of university and have made my studies a truly wonderful adventure. To Ludovica and Eleonora in par-

ticular, thank you for giving me the most peaceful and joyful home I could have ever asked for. You will always have a special place in my heart.

Il ringraziamento più grande va ai miei genitori, a cui sarò eternamente grato per avermi permesso di inseguire i miei sogni, nonostante tutti i sacrifici che hanno dovuto fare per rendere tutto ciò possibile, e per avermi sempre lasciato libero di decidere ed essere. Vorrei ringraziare anche il resto della mia famiglia, a partire dalle mie nonne, Teresa e Giuliana, mia zia Pina, le mie cugine Kikka e Yaya, e mio zio Adriano, per avermi sempre supportato e amato incondizionatamente.

E ora, finalmente, è giunto il momento di ringraziare i miei amici. Ai miei migliori amici e gli amici più stretti, Laura, Lucrezia, Enzo, Ledia, Ludovica, Valeria, Martina, Greta, Giada, Giorgia, Francesca, Vittoria, grazie per essere sempre stati i miei più grandi sostenitori. Mi avete tutti accolto nelle vostre vite senza mai farmi sentire di troppo, accompagnandomi in ogni percorso, da vicino o da lontano. Siete stati il rifugio più accogliente per le mie insicurezze e siete sempre stati in prima fila per festeggiare ogni mio traguardo. Se non mi sono mai sentito solo è grazie a voi, che mi avete donato infiniti ricordi da riguardare con la consapevolezza che probabilmente non troverò mai delle persone speciali come voi. Siete il mio più grande punto di riferimento, una bussola che mi guida a casa ogni volta che non ci sono.

Infine, vorrei ringraziare tutte le persone che mi hanno accompagnato nel mio percorso universitario a Delft, i nostri amici olandesi, Cecilia e soprattutto Ludovica, Eleonora ed Elisa, che sono con me dal primo giorno di università e che hanno reso il mio percorso di studi una bellissima avventura. A Ludovica ed Eleonora, nello specifico, voglio dire grazie per avermi regalato la convivenza più serena e piacevole che potessi chiedere, vi porto sempre nel mio cuore.

Delft, 25 November 2025
Andrea Paniccia

Contents

List of Figures

List of Tables

List of Abbreviations

1st FTAP 1st fundamental theorem of asset pricing. 13

2nd FTAP 2nd fundamental theorem of asset pricing. 13, 51

ATM at-the-money. 19

BM Brownian Motion. 51, 68, 69, 71

CDF cumulative distribution function. 7, 8

CPP compound Poisson process. 18, 32, 33

CTMC continuous-time Markov chain. 14

EMM equivalent martingale measure. 12, 13, 105

FFT Fast Fourier Transform. i, 20, 97

GBM Geometric Brownian Motion. 4, 68, 69, 72

ITM in-the-money. 19

LSMC Least Squares Monte Carlo. i, vi–x, 58, 59, 62–64, 66, 83, 86–95, 97, 98, 105–108

MM martingale measure. 12

MMAP Markov-modulated additive process. 14

MRJD mean-reversion jump-diffusion. 69

MRS Markov regime-switching. 4, 5, 69, 70

OTM out-of-the-money. 19

OU Ornstein-Uhlenbeck. 51, 52, 59, 69, 71, 72, 83–85

PDF probability density function. 7, 8, 20–22, 37, 40, 43, 83–85

SDE stochastic differential equation. 68, 69

w.r.t with respect to. viii, 12, 94, 105

1 Introduction

The pricing of financial derivatives has been a pillar of modern quantitative finance for a long time, with a rich body of theory and methods developed for traditional markets. On the other hand, energy markets are characterized by many unique features that distinguish them from traditional markets. They have undergone significant changes over the years, including market liberalization and the introduction of renewable energy sources. The latter, for instance, have caused a sharp increase in volatility. Electricity and other energy commodities cannot be stored physically on a large scale, and exhibit some recurrent features, like seasonality, mean reversion and sudden price spikes, all of which significantly make their dynamics more complex. In addition to this, energy derivatives are often also conceived as practical tools for risk management, tailored to the operational and contractual realities of the energy sector. These distinctive characteristics require adjustments to pricing methodologies and motivate the exploration of whether techniques successful in conventional markets can be effectively applied in energy markets.

This thesis focuses on two types of contracts that represent the growing complexity of the sector. The first are energy quanto options, whose payoff is similar to the product of two options, each written on an energy-related underlying, such as temperature, electricity or gas. Unlike standard contracts, which only account for price risk of the underlying, these options are designed to protect against both volumetric and price risks arising in the energy markets. These two types of risk are strongly correlated. For instance, weather conditions can simultaneously influence energy demand and spot prices: during an unexpectedly warmer winter, the reduced heating demand leads to a decline in energy consumption, yielding a decline in the energy price. Another example arises from the interdependence between different energy sources: when natural gas prices fall, the cost of producing electricity from gas-fired power plants decreases, leading to an increase in supply on the electricity market. This increase in supply can then lead to a decrease in electricity prices, demonstrating how fluctuations in gas markets indirectly affect electricity prices. Energy quanto options are designed to help option holders to hedge against such risks.

The second more complex class of instruments are electricity storage contracts. As mentioned earlier, electricity prices are very sensitive to changes in supply and demand because, unlike most commodities, electricity cannot be stored on a large scale. Because of this, electricity markets display the unique traits mentioned above. In addition to this, the growth of renewable energy sources increases price volatility due to their intermittent feature, making it more difficult to maintain a stable grid frequency. One solution to these problems is electricity storage, which is constantly evolving. This solution is not only useful for this purpose, but also interesting from a financial perspective: for instance, electricity can be bought at a low price when the supply is high, and sold later when the prices are increased to due higher demand. The valuation of these contracts is complicated by the fact that they can be exercised at multiple predefined dates, they take into account physical constraints of the storage like charging/discharging efficiencies and capacity limits, and they also include contractual penalties.

To ensure a reliable pricing of these derivatives, it is fundamental to develop accurate models that can capture the main features of the energy price dynamics. In fact, standard common models, like the Geometric Brownian Motion (GBM) are unable to capture sudden spikes or the mean-reverting behavior. We decide to work with Markov-modulated models, also known as Markov regime-switching (MRS) models, which provide a more realistic representation of the market by allowing the prices to evolve under different regimes, each having its own parameters or even a distinct underlying model. This choice, however, increases the complexity of the pricing problem.

This thesis extends the COS method, initially introduced in [14], to price the aforementioned two types of energy derivatives under Markov-modulated models. The COS method is a Fourier-based numerical method, which was proven to be extremely efficient and reliable for a wide range of applications, especially in financial markets. Yet, its application in the field of energy markets, especially with the additional complexity given by MRS models, is still an area that leaves room for exploration. The goal of this thesis is to investigate the possibility, versatility, speed, and accuracy of applying the COS method to energy derivatives of escalating model and contract complexity. This growing complexity is reflected in the structure of this thesis.

Chapter 2 gives an overview of the main mathematical tools that are fundamental to this research work, including stochastic calculus, financial mathematics and numerical methods.

Chapter 3 focuses on energy quanto options. This is a natural starting point for our analysis as they share structural similarities with European contracts, yet already have some complexity due to the use of MRS models.

Chapter 4 turns to electricity storage contracts, whose evaluation becomes more

challenging due to early exercise features and the real-world operational constraints. In this chapter, I replicate and analyze the existing literature, specifically [7], where these contracts are priced using the COS method.

Chapter 5 presents the main contribution on electricity storage contracts: we further extend the COS method, on top of the extension made in [7], by adopting a MRS framework to capture the real market behavior. This choice significantly improves the validity of the model assumption, but also adds challenges to the numerical valuation. Chapters 3, 4, 5 share a common structure: each begins with theoretical background on the relevant model or contract, proceeds with the model methodology, and concludes with results and discussion.

Finally, Chapter 6 summarizes the findings of the thesis and outlines the main conclusions.

2 Mathematical Framework

In this chapter, we introduce the mathematical instruments used throughout the thesis.

2.1 Random Variables

In this section, we present fundamental concepts from probability theory and stochastic calculus, which form the mathematical foundation of financial modeling. The material is based on [26].

Definition 2.1.1. Given a random variable X, its **cumulative distribution function (CDF)** is defined as:

$$F_X(x) = \mathbb{P}[X \leq x],$$

and if it is continuous, its **probability density function (PDF)** is defined as:

$$f_X(x) = \frac{dF_X(x)}{dx}.$$

Definition 2.1.2. Given X a continuous real-valued random variable with PDF $f_X(x)$, its **expected value** $\mathbb{E}[X]$ is defined as:

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f_X(x) dx,$$

assuming that the integral $\int_{-\infty}^{\infty} |x| f_X(x) dz$ exists and is finite.
The **variance** of X is defined as:

$$\mathrm{Var}[X] = \int_{-\infty}^{\infty} (x - \mathbb{E}[X])^2 f_X(x) dx,$$

assuming that the integral exists.

In computational finance, the density functions for random variables such as interest rates or stock prices are what one typically works with since these are necessary for

the calculation of expectations and variances. While there are closed-form densities for a few simple processes, most financial models lack them. But for some of these, the characteristic functions are analytically tractable, from which one can obtain useful information about expectations and other parameters. We therefore also define the characteristic function in this section.

Definition 2.1.3. Given a continuous random variable X, its **characteristic function** $\phi_X(u)$, for $u \in \mathbb{R}$, is defined as:

$$\phi_X(u) = \mathbb{E}[e^{\mathrm{i}uX}] = \int_{-\infty}^{\infty} e^{\mathrm{i}ux} f_X(x) dx,$$

where i is the imaginary unit, such that $\mathrm{i}^2 = -1$.

The characteristic function $\phi_X(u)$ uniquely determines the distribution of X. In particular, if X admits a PDF $f_X(x)$, then $\phi_X(u)$ is its Fourier transform.

Given $\phi_X(u)$, the k-th moment of X can be obtained using:

$$m_k(\cdot) = (-\mathrm{i})^k \frac{d^k}{du^k} \phi_X(u)\Big|_{u=0}.$$

Another useful function is the cumulant characteristic function $\zeta_X(u)$, defined as the logarithm of the characteristic function $\phi_X(u)$:

$$\zeta_X(u) = \ln\phi_X(u) = \ln\mathbb{E}[e^{\mathrm{i}uX}].$$

Given $\zeta_X(u)$, the **k-th cumulant** $\zeta_k(\cdot)$ can be obtained using:

$$\zeta_k(\cdot) = (-\mathrm{i})^k \frac{d^k}{du^k} \zeta_X(u)\Big|_{u=0}. \tag{2.1}$$

Definition 2.1.4. The **joint CDF** of two random variables X and Y, is the function $F_{X,Y}(\cdot,\cdot) : \mathbb{R}^2 \to [0,1]$, defined as:

$$F_{X,Y}(x,y) = \mathbb{P}[X \leq x, Y \leq y]$$

If X and Y are continuous variables, the **joint PDF** of X and Y is a function $f_{X,Y}(\cdot,\cdot) : \mathbb{R}^2 \to \mathbb{R}^+ \cup \{0\}$, defined as:

$$f_{X,Y}(x,y) = \frac{\partial F_{X,Y}(x,y)}{\partial x \partial y},$$

which means that for any event A, it follows that:

$$\mathbb{P}[(X,Y) \in A] = \int\int_A f_{X,Y}(x,y) dx dy.$$

As the joint pdf is a probability distribution function, the following properties hold:

- $f_{X,Y}(x,y) \geq 0$ for any $x, y \in \mathbb{R}$;
- $\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{X,Y}(x,y)dxdy = 1$.

Given the joint distribution of (X,Y), the expectation of a function $h(X,Y)$ can be computed as follows:

$$\mathbb{E}[h(X,Y)] = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h(x,y) f_{X,Y}(x,y)dxdy.$$

2.2 Stochastic processes and Martingale property

A **stochastic process** $X(t)$ is a collection of random variables indexed by a time variable t.
Suppose we are given a sequence of dates $T_1, T_2, \ldots, T_m$. Up to the current date, we have observed the values taken by a stochastic process $X(t)$. This means the historical path of the process is known. In contrast, the future evolution of the process is unknown and can only be explored through simulation, based on an assumed dynamics of asset prices.
To mathematically describe the information available about the process up to a certain time T_i, we use the concept of a *sigma-algebra*, which captures all events that can be determined from observations made up to that time. A sequence of sigma-algebras is referred to as a *filtration*. Given observation dates $T_1, T_2, \ldots, T_i$, the filtration generated by the stochastic process X up to time T_i is defined as

$$\mathcal{F}(T_i) := sigma\left(X(T_j) : 1 \leq j \leq i\right).$$

This captures the information accumulated from the process X up to time T_i. Because the information grows over time as more observations become available, we have a nested sequence of sigma-algebras:

$$\mathcal{F}(T_1) \subseteq \mathcal{F}(T_2) \subseteq \cdots \subseteq \mathcal{F}(T_i).$$

When a process is said to be $\mathcal{F}(T)$-*measurable*, it means that its values are known (or observable) at any time $t \leq T$, based on the information available up to time T.
A stochastic process $X(t)$, for $t \geq 0$, is said to be *adapted* to a filtration $\mathcal{F}(t)$ if

$$sigma(X(t)) \subseteq \mathcal{F}(t) \quad \text{for all } t \geq 0.$$

This means that the process does not incorporate knowledge of future events: at any given time t, the value $X(t)$ is determined by the information available up to that point.

A central example of a stochastic process is the Brownian motion, whose formal definition is given below.

Definition 2.2.1. A real-valued process $\{W(t),\ t \geq 0\}$ is called **Brownian motion** if:

- $W(0) = 0$;
- the increments are *normally distributed*: for all $0 \leq s < t,\ W(t) - W(s) \simeq N(0,\ t - s)$;
- the increments are *independent*: for all $0 \leq t_0 < t_1 < \ldots < t_n$, the random variables $Y(t_i) = W(t_i) - W(t_{i-1}),\ i = 1, \ldots, n$ are independent;
- the map $t \to W(t)$ is *continuous*.

Definition 2.2.2. A **d-dimensional Brownian motion** is a process $\{(W_1(t),\ \ldots, W_d(t)) : t \geq 0\}$, where $\{W_i(t) :\ t \geq 0\},\ i = 1,\ \ldots, d$ are mutually independent Brownian motions.

Before introducing martingales, we first recall the concept of *conditional expectation*, which plays a central role in describing how expectations evolve with information.

Definition 2.2.3 (Conditional Expectation)**.** Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a probability space, and $X : \Omega \to \mathbb{R}$ an integrable random variable, i.e., $\mathbb{E}[|X|] < \infty$. Let $\mathscr{G} \subseteq \mathcal{F}$ be a sigma-algebra. Then $\mathbb{E}[X|\mathscr{G}]$ is defined as the unique random variable $Y : \Omega \to \mathbb{R}$ which satisfies:

- Y is $\mathscr{G}$-measurable;
- Y is integrable;
- Y satisfies the "defining property of conditional expectation", i.e., $\forall G \in \mathscr{G}$:

$$\int_G Y d\mathbb{P} = \int_G X d\mathbb{P}.$$

When $\mathscr{G} = sigma\{X_1,\ \ldots, X_n\}$ with X_i random variables, then $\mathbb{E}[X|\mathscr{G}] = \mathbb{E}[X|X_1,\ \ldots, X_n]$.

Theorem 2.2.1. *$(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and $X : \Omega \to \mathbb{R}$ an integrable random variable. Let $\mathscr{G} \subseteq \mathcal{F}$ be a sigma-algebra. The conditional expectation $\mathbb{E}[X|\mathscr{G}]$ has the following properties:*

- $\mathbb{E}[\mathbb{E}[X|\mathscr{G}]] = \mathbb{E}[X]$;

- *Linearity:* $\mathbb{E}[aX+bY|\mathscr{G}] = a\mathbb{E}[X|\mathscr{G}]+b\mathbb{E}[Y|\mathscr{G}]$, *with* $a,\ b \in \mathbb{R}$ *and* $X,\ Y$ *integrable random variables;*
- *If* X *is* $\mathscr{G}$*-measurable, then* $\mathbb{E}[X|\mathscr{G}] = X$*;*
- *For* $c \in \mathbb{R}$, $\mathbb{E}[c|\mathscr{G}] = c$*;*
- *Conditioning on a trivial sigma-algebra:* $\mathbb{E}[X|\{\emptyset, \Omega\}] = \mathbb{E}[X]$*;*
- *Positivity:* $X \geq 0$ *implies* $\mathbb{E}[X|\mathscr{G}] \geq 0$*;*
- *Monotonicity: if* $X \leq Y$*, then* $\mathbb{E}[X|\mathscr{G}] \leq \mathbb{E}[Y|\mathscr{G}]$*;*
- $|\mathbb{E}[X|\mathscr{G}]| \leq \mathbb{E}[|X||\mathscr{G}]$*;*
- *If* X *and* $\mathscr{G}$ *are independent, then* $\mathbb{E}[X|\mathscr{G}] = \mathbb{E}[X]$.

A key property when working with stochastic processes is the *martingale* property. We present its definition next.

Definition 2.2.4. Consider a probability space $(\Omega, \mathcal{F}, \mathbb{Q})$, where Ω is the set of all possible outcomes, $\mathcal{F}(t)$ is the sigma-algebra, and $\mathbb{Q}$ is the probability measure.
A right-continuous $X(t)$ process with left limits (cádlág) for $t \in [0, T]$ is called a **martingale** with respect to the filtration $\mathcal{F}(t)$ under the measure $\mathbb{Q}$, if $\forall t < \infty$ the following holds:

$$\mathbb{E}[|X(t)|] < \infty,$$

and

$$\mathbb{E}[X(t)|\mathcal{F}(s)] = X(s), \quad \text{with } s < t,$$

where $\mathbb{E}[\cdot|\mathcal{F}]$ is the conditional expectation under measure $\mathbb{Q}$.

2.3 Fundamental theorems of asset pricing

In this section, we present the fundamental theorems of asset pricing, which describe the conditions under which a financial market is arbitrage-free and complete. Let us introduce the mathematical setup for this section. We consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, the time axis for a multi-period model $\mathbb{T} = \{0, 1, \ldots, T\}, T \in \mathbb{N}$, and d-assets (for $i \in \{1, \ldots, d\}$), such that $S^i : \Omega \times \mathbb{T} \to \mathbb{R}_{\geq 0}$ and $(\omega, t) \to S^i_t(\omega)$ is the value of the asset i at time point t in case event ω occurs, and S^i_0 is the deterministic value. The assets prices $S = (S^0_t, S^1_t, \ldots, S^d_t)_{t\in\mathbb{T}}$ are adapted to the filtration $(\mathcal{F}_t)_{t\in\mathbb{T}}$.

Definition 2.3.1. Let $\mathbb{P}$ and $\mathbb{Q}$ be two probability measures on the space $(\Omega, \mathcal{F})$. The measure $\mathbb{Q}$ is called:

- **equivalent** to $\mathbb{P}$ ($\mathbb{Q} \sim \mathbb{P}$) if $\mathbb{P}(A) = 0 \iff \mathbb{Q}(A) = 0,\ \forall A \in \mathcal{F}$;
- **absolutely continuous** with respect to (w.r.t) $\mathbb{P}$ ($\mathbb{Q} \ll \mathbb{P}$) if $\mathbb{P}(A) = 0 \implies \mathbb{Q}(A) = 0,\ \forall A \in \mathcal{F}$;
- **singular** w.r.t $\mathbb{P}$ ($\mathbb{Q} \perp \mathbb{P}$) if $\exists A \in \mathcal{F}$: $\mathbb{Q}(A) = 0$ and $\mathbb{P}(A) = 1$.

Definition 2.3.2. A probability measure $\mathbb{Q}$ on $(\Omega, \mathcal{F})$ is a **martingale measure (MM)** for S if:

- $\mathbb{E}_{\mathbb{Q}}[|S_t^i|] < \infty, \quad \forall t \in \mathbb{T}, \forall i = 1, \ldots, d$;
- $S_{t-1}^i = \mathbb{E}^{\mathbb{Q}}\left[\frac{S_t^i}{B_t} \mid \mathcal{F}_{t-1}\right], \quad \forall t \in \mathbb{T}, \forall i = 1, \ldots, d.$

where B_t is the numéraire or discount factor (which is often the risk-free money market account).
If $\mathbb{Q}$ is a MM and it holds that $\mathbb{Q} \sim \mathbb{P}$, then $\mathbb{Q}$ is an **equivalent martingale measure (EMM)**.

Definition 2.3.3. We define the **set of all EMMs** as follows:

$$\mathcal{P} = \{\mathbb{Q} | \mathbb{Q} \sim \mathbb{P},\ \mathbb{Q} \text{ is a MM for } S\}.$$

Definition 2.3.4. An **investment strategy** $\bar{\xi} = (\xi^0, \xi^1, \ldots, \xi^d) \in \mathbb{R}^{d+1}$ in the multi-period model is a predictable stochastic process $\bar{\xi} = (\bar{\xi}_t)_{t \in \mathbb{T}}$. The investment strategy $\bar{\xi}$ is related to the **value process** $V = V^{\xi}$, as follows:

$$V_t = V_t^{\xi} = \bar{\xi}_t \cdot \bar{S}_t = \sum_{i=0}^{d} \xi_t^i \cdot S_t^i, \quad \forall t \in \{0, 1, \ldots, T\},$$

where $\bar{S}_t = (S_t^0, S_t^1, \ldots, S_t^d)$ is the vector of asset prices at time $t \in \mathbb{T}$. V_t is the value of the portfolio at time $t \in \mathbb{T}$.
The investment strategy is self-financing if: $\delta_t = \bar{\xi}_{t+1} \cdot \bar{S}_t - \bar{\xi}_t \cdot \bar{S}_t = 0, \qquad \forall t \in \mathbb{T}$.

Definition 2.3.5. An investment strategy $\bar{\xi} \in \mathbb{R}^{d+1}$ with value process V is called an **arbitrage** if:

- There is no initial cost: $V_0 \leq 0$;
- There are no losses with certainty: $\mathbb{P}(V_T \geq 0) = 1$;
- There is gain with positive probability: $\mathbb{P}(V_T > 0) > 0$.

Theorem 2.3.1. *The **1st fundamental theorem of asset pricing (1st FTAP)** states that a financial market is abritrage-free if and only if there exists an EMM $\mathbb{Q}$:*

$$\textit{No arbitrage} \iff \mathcal{P} \neq \emptyset.$$

Definition 2.3.6. We say that a European claim C is **replicable** if there exists a self-financing strategy $\bar{\xi}$ such that:

$$C = \overline{\xi}_T \cdot \overline{S}_T \qquad \mathbb{P} - \text{a.s.}$$

A financial market is **complete** if every claim $C : \Omega \to \mathbb{R}$ is replicable.

Theorem 2.3.2. *The **2nd fundamental theorem of asset pricing (2nd FTAP)** states that an arbitrage-free financial market is complete if and only if there exists a unique EMM:*

$$\textit{complete market} \iff |\mathcal{P}| = 1.$$

2.4 Markov-modulated additive model

In this section, inspired by [2], we introduce a risk-neutral model for futures contract whose prices are governed by Markov-modulated additive processes. The increments are independent but non-stationary.

We consider a complete probability space $(\Omega, \mathcal{F}, \mathbb{Q})$, where $\mathbb{Q}$ is the risk-neutral probability measure, with respect to which we take the expectations $\mathbb{E}[\cdot]$. We introduce a time interval $[0, \tau]$, and we equip the probability space with a $\mathbb{Q}$-complete, right-continuous filtration $\{\mathcal{F}_t\}_{t\in[0,\tau]}$.

A filtration $\{\mathcal{F}_t\}_{t\geq 0}$ is said to be $\mathbb{Q}$-complete if it contains all subsets of $\mathbb{Q}$-null sets in $\mathcal{F}$, so every event of probability zero and its subsets are included in $\mathcal{F}_t, \forall t \in [0, \tau]$. A filtration is right-continuous if $\mathcal{F}_t = \bigcap_{s>t} \mathcal{F}_s,\ \forall t \in [0, \tau]$.

First, we give the definition of **additive processes** [10]:

Definition 2.4.1. A stochastic process $(X_t)_{t\geq 0}$ on $\mathbb{R}^d$ is called an **additive process** if:

- $X_0 = 0$.
- it is cadlag, meaning that it is right-continuous with left limits.
- its increments are independent: for every increasing sequence of times $t_0,\ ..., t_n$, the random variables $X_{t_0},\ X_{t_1} - X_{t_0},\ ..., X_{t_n} - X_{t_{n-1}}$ are independent.
- it is stochastically continuous: $\forall \epsilon > 0,\ \mathbb{P}[|X_{t+h} - X_t| \geq \epsilon] \underset{h\to 0}{\to} 0$.

Note that a Lévy process is a special case of an additive process where the increments are not only independent but also stationary.

We now introduce the concept of Markov processes, inspired by [22], which form a fundamental basis for the models developed in the following sections.

Definition 2.4.2. Let S be a countable, nonempty set equipped with the discrete metric, so that right-continuity of functions $X : \mathbb{R}_{\geq 0} \to S$ is well-defined. A **continuous-time Markov chain (CTMC)** on S is a stochastic process $X = (X(t))_{t\geq 0}$ specified by:

- an initial distribution $\mathbf{p} = (p_i)_{i\in S}$, where $p_i = \mathbb{P}(X(0) = i)$;
- a rate matrix $Q = (q_{i,j})_{i,j\in S}$, where:
  1. $q_{i,j} \geq 0$ for all $i \neq j \in S$,
  2. $q_{i,i} = -\sum_{j\neq i} q_{i,j}$, so that each row sums to zero: $\sum_{j\in S} q_{i,j} = 0$.

The process satisfies the Markov property with transition rates governed by Q: for all $i, j \in S$, for all $t \geq 0$ such that $\mathbb{P}(X(t) = i) > 0$, and for all small $h > 0$,

$$\mathbb{P}(X(t + h) = j \mid X(t) = i) = \delta_{i,j} + q_{i,j}h + o(h),$$

where $\delta_{i,j}$ is the Kronecker delta (i.e., $\delta_{i,j} = 1$ if $i = j$, and 0 otherwise).

The rate $q_{i,j}$ describes the instantaneous transition rate from state i to state j. For $i \neq j$, it represents how quickly the process jumps from i to j, while $-q_{i,i}$ gives the total rate of leaving state i.

Note: The symbol $\mathbb{1}_A$ denotes the indicator function for a condition A. It is defined as:

$$\mathbb{1}_A = \begin{cases} 1 & \text{if condition } A \text{ is true} \\ 0 & \text{if condition } A \text{ is false.} \end{cases}$$

With these preliminary definitions in place, we now introduce the framework of **Markov-modulated additive processes (MMAPs)**: consider a Markov process $\{M(t)\}_{t\in[0,\tau]}$, continuous in time and with a finite-state space $\mathscr{S} = \{1, ..., J\}$, generator matrix $\mathcal{Q} = \{q_{jl}\}_{J\times J}$ (q_{jl} is the transition rate from state j to state $l \neq j$) and initial probability vector $\mathbf{p} \in \mathbb{R}^J$. A two dimensional process $\mathbf{X}(t, \tau) = (X_1(t, \tau),\ X_2(t, \tau))$, is a MMAP when $\mathbf{X}(0, \tau) = \mathbf{0}$ and when $M(s) = j$, for $s > 0$, $\mathbf{X}$ behaves like a two dimensional additive process $\mathbf{Z}_j$, for $j = 1, ..., J$. These additive processes are assumed to be independent and have the following characteristic exponent ([21]):

$$\Phi_j(s,t,\tau,\mathbf{u}) = -\frac{1}{2}\left\langle \mathbf{u}^T, \left(\mathbf{C}_j(s,\tau) - \mathbf{C}_j(t,\tau)\right)\mathbf{u}^T \right\rangle \quad + \\ + \int_s^t \int_{\mathbb{R}^2} \left(e^{\mathrm{i}\langle \mathbf{u},\mathbf{z}\rangle} - 1 - \mathrm{i}\,\langle \mathbf{u},\mathbf{z}\rangle\, \mathbb{1}_{|\mathbf{z}|\leq 1}\right) \ell_j(dz,dy)$$

such that:

$$\mathbb{E}\left[e^{\mathrm{i}<\mathbf{u},\ \mathbf{Z}_j(s,\tau)-\mathbf{Z}_j(t,\tau)>}\right] = e^{\Phi_j(s,t,\tau,\mathbf{u})}\ ,$$

where $\mathbf{u} = (u_1, u_2) \in \mathbb{R}^2$, $0 \leq s \leq t$, $\mathbf{u}^T$ is the transpose of $\mathbf{u}$ and Φ_j is assumed to be differentiable w.r.t. time. The matrix $\mathbf{C}_j(s,\tau)$ represents the cumulative covariance between s and τ of the Gaussian component of the additive process $\mathbf{Z}_j$. So $\mathbf{C}_j(t,\tau) - \mathbf{C}_j(s,\tau)$ behaves like the covariance of the increment $\mathbf{Z}_j(t,\tau) - \mathbf{Z}_j(s,\tau)$. $\mathbf{C}_j(s,\tau)$ has the following properties:

1. $\mathbf{C}_j(s,\tau)$ is a 2×2 positive definite matrix
2. $\mathbf{C}_j(0,\tau) = 0$
3. $\mathbf{C}_j(s,\tau) \to \mathbf{C}_j(t,\tau)$ as $s \to t$
4. $\mathbf{C}_j(t,\tau) - \mathbf{C}_j(s,\tau)$ is positive definite for $s \leq t$.

In addition to this, ℓ_j is a measure on $(0,\infty) \times \mathbb{R}^2$, such that:

1. $\int_0^t \int_{\mathbb{R}^2} (1 \wedge |z|^2) \ell_j(dz,dy) < \infty$
2. $\ell_j(\{t\} \times \mathbb{R}^2) = 0$ for $t \geq 0$.

The two differentials dz and dy refer to integration over two distinct dimensions: dz denotes integration over the jump sizes $z \in \mathbb{R}^2$, while dy denotes integration over time.

A fundamental assumption is made, following [13], for the finiteness of the expectation $\mathbb{E}[e^{<\mathbf{a},\mathbf{X}(t,\tau)>}]$, where $\mathbf{a} \in \mathbb{C}^2$:

Assumption 2.4.1. *There exists a constant $C > 2$ such that $\forall \mathbf{a} \in \mathbb{C}^2$ with $\mathrm{Re}(\mathbf{a}) \in [-C, C]^2$*

$$\int_0^\tau \int_{|z|>1} e^{<\mathrm{Re}(\mathbf{a}),z>} \ell_j(dz,dy) < \infty, \quad j = 1, ..., J,$$

where $\mathrm{Re}(\mathbf{a})$ denotes the real part of the components of the vector $\mathbf{a}$.

This implies that:

1. $\mathbb{E}[|e^{<\mathbf{a},\mathbf{Z}_j(t,\tau)>}|] < \infty \quad \forall t \in [0,\tau]$
2. $\mathbb{E}[e^{<\mathbf{a},\mathbf{Z}_j(t,\tau)>}] = e^{\phi_j(0,t,\tau,\mathbf{ia})} \quad \forall j = 1,...,J.$

The following result will be useful to determine the joint characteristic function of Markov-modulated additive processes.

Lemma 2.4.1. *Consider a Markov-modulated drift process*

$$\Gamma_1(t,\tau) = \int_0^t \sum_{j=1}^{J} \mu_j(s,\tau)\mathbb{1}_{M(s)=j} ds,$$

where $\mu_j(t,\tau)$ are real-valued deterministic functions which satisfy the integrability condition $\int_0^\tau |\mu_j(s,\tau)|ds < \infty$. Consider a Markov-modulated process

$$\Gamma_2(t) = \int_0^t \sum_{j=1}^{J} \gamma_j \mathbb{1}_{M(s)=j} ds,$$

where $\mu_j,\ j = 1,...,J$ are constants. Then, $\forall \mathbf{a} \in \mathbb{C}^2$ that satisfies Assumption 2.4.1 and $\forall\ 0 \le t \le \tau$:

$$\mathbb{E}[e^{\Gamma_1(t,\tau)+\Gamma_2(t)+<\mathbf{a},\mathbf{X}(t,\tau)>}] = \mathbf{p}e^{(Q+D)t+\int_0^t A(s,\tau)ds}\mathbf{1},$$

where $\mathbf{1}$ is a column vector of ones, D and $A(s,\tau)$ are diagonal matrices with entries $D_{jj} = \gamma_j$ and $A_{jj}(t,\tau) = \mu_j(t,\tau) + \frac{\partial \phi_j}{\partial y}(t,y,\tau,\mathbf{ia})\Big|_{y=t}$.

Proof. A proof can be found in [2].

Although Markov-modulated additive models can be defined for an arbitrary number d of processes $\{X_l(t,\tau)\},\ l = 1,...,d$, we focus our analysis on the case where $d = 2$. Under the measure $\mathbb{Q}$, we introduce a two-dimensional price process that models futures prices of contracts $\mathbf{F}(t,\tau) = (F_1(t,\tau), F_2(t,\tau))$:

$$F_l(t,\tau) = F_l(0,\tau)e^{\Lambda_l(t,\tau)+X_l(t,\tau)}, \quad l = 1,2,$$

where τ is the delivery time, $t \in [0,\tau]$ is the current time, X_l is the Markov-modulated additive process and Λ_l denotes the drift process, defined as:

$$\Lambda_l(t,\tau) = \int_0^t \sum_{j=1}^{J} \mu_{j,l}(s,\tau)\mathbb{1}_{M(s)=j} ds, \quad l = 1,2,$$

where $\mu_{j,l}(t,\tau)$ is a real-valued and deterministic function which satisfies the integrability condition $\int_0^t |\mu_{j,l}(s,\tau)|ds < \infty$. We define the following column vector:

$$\mu_l(t,\tau) = (\mu_{1,l}(t,\tau),...,\mu_{J,l}(t,\tau))^T, \quad t \le \tau, \quad l = 1,2.$$

The following proposition states a martingale condition for futures prices.

Proposition 2.4.1. *Let $\mathbf{e}_l$ be the l-th standard basis vector of $\mathbb{R}^2$. The futures price process $\mathbf{F}_l(t,\tau)$ is a martingale under $\mathbb{Q}$ if the vectors $\mu_l(t,\tau)$ are chosen as:*

$$\mu_{j,l}(t,\tau) = -\frac{\partial \phi_j}{\partial y}(t,y,\tau,\mathrm{i}\mathbf{e}_l)\bigg|_{y=t}, \quad l=1,2.$$

Proof. A proof can be found in [2].
Lastly, we introduce a Markov-modulated interest rate process:

$$r(t) = \sum_{j=1}^{J} r_j \mathbb{1}_{M(t)=j},$$

where the coefficients r_j are constants. From this, we introduce the integrated interest rate process

$$U(t) = \int_0^t r(s)ds.$$

2.5 Poisson Process

In this section, we introduce the Poisson process which can be used to extend the Black-Scholes framework by governing discontinuous jumps, allowing for more realistic modeling of sudden movements in asset prices. To layout the Poisson process, we first introduce the Poisson random variable.

Definition 2.5.1. A **Poisson random variable** X_P counts the number of discrete events that occur independently within a fixed interval of time. The probability of observing $k \geq 0$ occurrences in the time interval is given by

$$\mathbb{P}[X_P = k] = \frac{\xi_p^k e^{-\xi_p}}{k!}.$$

For a Poisson random variable, it holds that:

$$\mathbb{E}[X_P] = \mathrm{Var}[X_P] = \xi_p.$$

Building on this, the Poisson process extends the concept of a Poisson random variable to a continuous-time setting, tracking the number of events that occur over time.

Definition 2.5.2. A **Poisson process** $X_P(t), t \geq 0$, with the rate $\xi_p > 0$ is an integer-valued stochastic process with the following properties:

- $X_P(0) = 0$;
- $\forall 0 < t_1 < ... < t_n$, the increments $X_P(t_1) - X_P(0), ..., X_P(t_n) - X_P(t_{n-1})$ are independent random variables;
- for $s \geq 0, t > 0$ and integers $k \geq 0$, the increments have the Poisson distribution:

$$\mathbb{P}[X_P(s+t) - X_P(t) = k] = \frac{(\xi_p s)^k e^{-\xi_p s}}{k!}.$$

The parameter ξ_p is the rate of the Poisson process and indicates the number of jumps in a time interval.

The Poisson process models the random occurrence of discrete events over continuous time, where each event is identical and contributes a unit increment to the process. While this is suitable for counting the number of events, many real-world phenomena, such as asset price jumps, require not just counting events but also accounting for the magnitude of their impact. This motivates the compound Poisson process, which extends the Poisson process by associating each jump with a random size.

Definition 2.5.3. A **compound Poisson process (CPP)** is a continuous-time stochastic process characterized by random jumps occurring at random times. The arrival times of the jumps follow a Poisson process, while the jump sizes are themselves random, drawn from a specified distribution. Formally, a compound Poisson process with intensity $\lambda > 0$ and jump size distribution G is defined as:

$$Y(t) = \sum_{i=1}^{\mathcal{N}(t)} D_i,$$

where $\{\mathcal{N}(t) : t \geq 0\}$ is a Poisson process with rate λ, and $\{D_i : i \geq 1\}$ is a sequence of independent and identically distributed random variables with distribution function G, independent of $\mathcal{N}(t)$. The process $Y(t)$ thus represents the cumulative effect of random jumps occurring at Poisson-distributed times.

2.6 Options

Financial derivatives are products that are based on the performance of some underlying asset, like a stock or an interest rate. An option is a financial contract that gives the holder the right to buy or sell an underlying asset in the future at a predetermined price. Something that is crucial is that options do not involve any obligation for the holder, but rather the optionality to trade in the asset. This changes when it comes to

the seller of the option, the counterparty of the contract, who is obliged to trade the asset when the holder makes use of the exercise right in the option contract. Options are financial derivatives in the sense that their values can be derived from the performance of another underlying asset.

The first distinction to make on options is between call options and put options. A **call option** gives the holder the right to buy an asset, whereas a **put option** gives the holder the right to sell an asset. The trade can happen at some time $t = T$ in the future and the asset is bought or sold at a predefined amount called *strike price*, denoted by K.

The second distinction that can be made is based on the number of time points in the future when the holder of the option may decide to trade in the asset for the strike price. In the case of **European options**, there is only one time point at which the holder can exercise its right, it is called *maturity* or expiry date, and it is denoted by $t = T$.

Definition 2.6.1. Given a certain asset value $S = S(t)$ and the payoff function $H(T, S)$, the **value of a call option** at time $t = T$ is given by:

$$V_c(T, S) = H(T, S) = \max(S(T) - K, 0)$$

Definition 2.6.2. Given a certain asset value $S = S(t)$ and the payoff function $H(T, S)$, the **value of a put option** at time $t = T$ is given by:

$$V_c(T, S) = H(T, S) = \max(K - S(T), 0)$$

When at time $t < T$, the payoff of the option is equal to zero (for a call option that happens when $S(t) - K < 0$), we say that the option is "out-of-the-money (OTM)", because its intrinsic value is equal to zero. If the intrinsic value is greater than zero, then the option is said to be "in-the-money (ITM)". When the intrinsic value is close to zero ($S(t) - K \approx 0$, for a call option), then the option is said to be "at-the-money (ATM)".

Other options instead include early-exercise features, meaning that there is more than one time point where they can be exercised. **Bermudan options** can be exercised at a set of predetermined dates $T_1, ..., T_n$, with $T_n = T$, whereas **American options** can be exercised at any time during the life of the option, $t_0 \leq t \leq T$. American options can be seen as the continuous time equivalent of Bermudan options, assuming an increasing number of exercise dates with increasingly smaller time intervals $\Delta t = \frac{T}{n} \to 0$ as $n \to +\infty$.

Another important class of derivatives is known as **exotic options**, which are distinguished by non-standard features in their payoff structures. A common example involves *path-dependence*, where the payoff depends not only on the stock price at a specific time, such as $S(t)$ or $S(T)$, but also on the evolution of the stock price over a period of time.

2.7 COS method

Numerical methods are required to be rapid and efficient in pricing options or calibrating financial models, reason why one of the main focuses of research in computational finance is to improve the performance of these methods. To give an overview, numerical methods can be classified in three different categories: partial (integro) differential equation (PIDE) methods, Monte Carlo simulation and numerical integration methods. Numerical integration methods rely on a transformation to the Fourier domain, one of the most popular examples is the Carr-Madan method ([8]). Pricing formulas usually involve the PDF, which for many relevant processes is not known, as opposed to the characteristic function which is instead available for many processes. Recalling that the PDF and the characteristic function make a Fourier pair, the pricing problems can be solved in the Fourier domain as long as the characteristic function is available. The integration required by the pricing formulae can be performed by means of the Fast Fourier Transform (FFT), with a computational complexity of $O(N\log_2 N)$, where N is the number of integration points. Quadrature-based methods, such as those relying on the FFT, tend to be inefficient for Fourier-transformed integrals due to the highly oscillatory nature of the integrands, requiring a fine grid to achieve acceptable accuracy.

The COS method relies on the Fourier-cosine (COS) expansion as a more efficient alternative for numerical integration. It has been proven that the COS method significantly improves the speed of pricing both plain vanilla and certain exotic options ([14]), and even barrier and American-style products ([15]). The COS method is more flexible in handling general underlying dynamics. It also allows for simultaneous pricing across a range of strikes and provides an efficient way to recover densities from characteristic functions, which is valuable in applications like calibration, forward-start option pricing, and static hedging.

In this section, we will introduce the COS method to price European options in a one-dimensional setting. Later, we will see how this method can be extended to a two-dimensional setting (3), and also how it can be used to price financial products that resemble Bermudan-style options (4).

To introduce the COS method, we recall Definition 2.1.3, and using the inverse

Fourier transform we find the following expression for the PDF:

$$f(x) = \frac{1}{2\pi} \int_{\mathbb{R}} e^{-\mathrm{i}\omega x} \phi(\omega) d\omega. \tag{2.2}$$

The main idea behind this method is to reconstruct the whole integral in the above expression from its Fourier-cosine series expansion, extracting the series coefficients from the integrand.

Considering a function with support on $[0, \pi]$, its cosine expansion reads:

$$f(\theta) = {\sum_{n=0}^{\infty}}' A_n \cos(n\theta) \quad \text{with} \quad A_n = \frac{2}{\pi} \int_0^{\pi} f(\theta) \cos(n\theta) d\theta, \tag{2.3}$$

where $\sum'$ means that the first term of the summation is weighted by $\frac{1}{2}$. If the function has support on a finite interval $[a, b] \in \mathbb{R}$, the cosine series expansion is obtained through a change of variables:

$$\theta = \frac{x-a}{b-a}\pi; \quad x = \frac{b-a}{\pi}\theta + a.$$

Thanks to this, we can write:

$$f(x) = {\sum_{n=0}^{\infty}}' A_n \cos\left(n\pi \frac{x-a}{b-a}\right) \quad \text{with} \quad A_n = \frac{2}{b-a} \int_a^b f(x) \cos\left(n\pi \frac{x-a}{b-a}\right) dx. \tag{2.4}$$

In order for the Fourier transform to exist, we require the integrands in Equation 2.2 to decay to zero at $\pm\infty$, which also allows us to truncate the integration range in such a way that we do not lose accuracy. Choose $[a, b] \in \mathbb{R}$ such that the truncated integral approximates the infinite integral well:

$$\phi_1(\omega) = \int_a^b e^{\mathrm{i}\omega x} f(x) dx \approx \int_{\mathbb{R}} e^{\mathrm{i}\omega x} f(x) dx = \phi(\omega). \tag{2.5}$$

The subscript 1 in ϕ_1 means that we are dealing with a first approximation, ϕ_i is the i-th numerical approximation. Looking at Equations 2.4 and 2.5, we find:

$$A_n = \frac{2}{b-a} \mathrm{Re}\left\{ \phi_1\left(\frac{n\pi}{b-a}\right) \cdot \exp\left(-\mathrm{i}\frac{n\pi a}{b-a}\right) \right\} \tag{2.6}$$

$$\approx \frac{2}{b-a} \mathrm{Re}\left\{ \phi\left(\frac{n\pi}{b-a}\right) \cdot \exp\left(-\mathrm{i}\frac{n\pi a}{b-a}\right) \right\} = F_n, \tag{2.7}$$

where $\mathrm{Re}\{\cdot\}$ denotes the real part of the argument. Now, we replace the coefficient A_n by F_n in the cosine series expansion of $f(x)$ on $[a, b]$:

$$f_1(x) = {\sum_{n=0}^{\infty}}' F_n \cos\left(n\pi \frac{x-a}{b-a}\right) \tag{2.8}$$

and truncate the series summation:

$$f_2(x) = \sum_{n=0}^{N-1}{}' F_n \cos\left(n\pi\frac{x-a}{b-a}\right). \tag{2.9}$$

The starting point for European option pricing is the risk-neutral valuation formula:

$$v(x,t_0) = e^{-r\Delta t}\,\mathbb{E}^{\mathbb{Q}}[v(t,T)\mid x] = e^{-r\Delta t}\int_{\mathbb{R}} v(y,T)f(y\mid x)dy,$$

where v is the option value, $\Delta t = T - t_0$, $\mathbb{E}^{\mathbb{Q}}[\cdot]$ is the expectation taken under the risk-neutral measure $\mathbb{Q}$, x and y are the state variables at time t_0 and T, respectively; $f(y \mid x)$ is the PDF of y given x, and r is the risk-free interest rate. Since the PDF decays to zero as $y \to \pm\infty$, we can truncate the infinite truncation range to $[a,b] \subset \mathbb{R}$ without losing accuracy.

$$v_1(x,t_0) = e^{-r\Delta t}\int_a^b v(y,T)f(y\mid x)dy. \tag{2.10}$$

Recalling what we mentioned in the beginning of this section, since the characteristic function is usually known more than the PDF, we replace $f(y \mid x)$ by its Fourier-cosine expansion:

$$f(y\mid x) = \sum_{n=0}^{+\infty}{}' A_n(x)\cos\left(n\pi\frac{y-a}{b-a}\right) \tag{2.11}$$

with

$$A_n(x) = \frac{2}{b-a}\int_a^b f(y\mid x)\cos\left(n\pi\frac{y-a}{b-a}\right) \tag{2.12}$$

so that

$$v_1(x,t_0) = e^{-r\Delta t}\int_a^b v(t,T)\sum_{n=0}^{+\infty}{}' A_n(x)\cos\left(n\pi\frac{y-a}{b-a}\right)dy. \tag{2.13}$$

By interchanging summation and integration, we define:

$$V_n = \frac{2}{b-a}\int_a^b v(y,T)\cos\left(n\pi\frac{y-a}{b-a}\right)dy, \tag{2.14}$$

resulting in

$$v_1(x,t_0) = \frac{1}{2}(b-a)e^{-r\Delta t}\cdot\sum_{n=0}^{+\infty}{}' A_n(x)V_n. \tag{2.15}$$

Furthermore, due to the rapid decay of the coefficients V_n, we truncate the summation:

$$v_2(x,t_0) = \frac{1}{2}(b-a)e^{-r\Delta t} \cdot {\sum_{n=0}^{N-1}}' A_n(x)V_n. \tag{2.16}$$

Similarly to Equations 2.6 and 2.7, we replace the coefficients A_n by F_n and obtain the COS formula for general underlying processes:

$$v(x,t_0) \approx v_3(x,t_0) = e^{-r\Delta t} {\sum_{n=0}^{N-1}}' \text{Re}\left\{\phi\left(\frac{n\pi}{b-a};x\right) e^{-\text{i}n\pi\frac{a}{b-a}}\right\} V_n. \tag{2.17}$$

2.7.1 Coefficients V_n for European options

The last step for European option pricing with COS method is to determine the coefficients V_n. We consider the log-asset prices $x = \ln(S_0/K)$ and $y = \ln(S_T/K)$, with S_t being the underlying asset price at time t and K the strike price. The payoff of a European option in log-asset price is

$$v(y,T) = [\alpha \cdot K(e^y - 1)]^+, \text{ with } \alpha = \begin{cases} 1 & \text{for a call} \\ -1 & \text{for a put.} \end{cases}$$

Following the derivations in [14], we obtain

$$V_n = \begin{cases} \frac{2}{b-a}K(\chi_n(0,b) - \psi_n(0,b)) & \text{for a call} \\ \frac{2}{b-a}K(-\chi_n(a,0) + \psi_n(a,0)) & \text{for a put,} \end{cases}$$

where

$$\chi_n(c,d) = \frac{1}{1+\left(\frac{n\pi}{b-a}\right)^2}\Big[\cos\left(n\pi\frac{d-a}{b-a}\right)e^d - \cos\left(n\pi\frac{c-a}{b-a}\right)e^c + \frac{n\pi}{b-a}\sin\left(n\pi\frac{d-a}{b-a}\right)e^d - \frac{n\pi}{b-a}\sin\left(n\pi\frac{c-a}{b-a}\right)e^c\Big] \tag{2.18}$$

$$\psi_n(c,d) = \begin{cases} \left[\sin\left(n\pi\frac{d-a}{b-a}\right) - \sin\left(n\pi\frac{c-a}{b-a}\right)\right]\frac{b-a}{n\pi} & n \neq 0 \\ d-c & n = 0. \end{cases} \tag{2.19}$$

2.7.2 Choice of the integration interval $[a,b]$

As for the choice of the integration interval $[a,b]$, [14] proposes the following:

$$[a,b] = \left[c_1 - L\sqrt{c_2 + \sqrt{c_4}}\ ;\ c_1 + L\sqrt{c_2 + \sqrt{c_4}}\right], \tag{2.20}$$

with c_n the n-th cumulant of $\ln(S_T/K)$. The formula in Equation 2.20 is accurate within $T = 0.1$ and $T = 10$. When pricing call options, the accuracy of the method is notably sensitive to the choice of the parameter L in Equation 2.20. This sensitivity stems from the fact that the call option payoff increases exponentially with the log-stock price, which can lead to significant cancellation errors for large values of L. Put options, on the other hand, do not exhibit this issue, as their payoff is bounded by the strike price K. To avoid numerical instability when pricing call options, one may either select $L \in [7.5, 10]$, or alternatively use the more robust put-call parity relation:

$$v^{\text{call}}(\mathbf{x}, t_0) = v^{\text{put}}(\mathbf{x}, t_0) + S_0 e^{-qT} - K e^{-rT}.$$

2.8 Numerical Approximation of Cumulants via Finite Differences

In order to compute the cumulants (Equation 2.1) of the underlying stochastic process, it is necessary to evaluate various derivatives of the cumulant-generating function. In many practical settings, these derivatives cannot be obtained in closed form, especially when dealing with more intricate stochastic models.

To address this challenge, we approximate the derivatives numerically using **central finite difference** schemes ([32]). This approach is both accurate and easy to implement, making it well-suited for our purposes. Central differences are preferred over forward or backward differences because they offer a higher order of accuracy for the same stencil width. These formulas are derived by expanding the function f in a Taylor series around the point x_i and combining the resulting expressions to eliminate lower-order terms. The result is an efficient and robust way to compute derivatives needed for cumulant evaluation.

Given a sufficiently smooth function f, and a uniform step size h, the central difference formulas used in this thesis are the following:

First-order derivative:

$$f'(x_i) \approx \frac{f(x_{i+1}) - f(x_{i-1})}{2h}, \quad \text{with error } \mathcal{O}(h^2).$$

Proof. To derive the central difference approximation of the first derivative, we use Taylor expansions around the point x_i:

$$f(x_{i+1}) = f(x_i) + hf'(x_i) + \frac{h^2}{2}f''(x_i) + \frac{h^3}{6}f^{(3)}(x_i) + \mathcal{O}(h^4),$$
$$f(x_{i-1}) = f(x_i) - hf'(x_i) + \frac{h^2}{2}f''(x_i) - \frac{h^3}{6}f^{(3)}(x_i) + \mathcal{O}(h^4).$$

Subtracting the two equations:

$$f(x_{i+1}) - f(x_{i-1}) = 2hf'(x_i) + \frac{2h^3}{6}f^{(3)}(x_i) + \mathcal{O}(h^5).$$

Dividing by $2h$, we obtain:

$$f'(x_i) \approx \frac{f(x_{i+1}) - f(x_{i-1})}{2h}, \quad \text{with error } \mathcal{O}(h^2).$$

Second-order derivative:

$$f''(x_i) \approx \frac{f(x_{i+1}) - 2f(x_i) + f(x_{i-1})}{h^2}, \quad \text{with error } \mathcal{O}(h^2).$$

Proof. Using the same Taylor expansions:

$$f(x_{i+1}) = f(x_i) + hf'(x_i) + \frac{h^2}{2}f''(x_i) + \frac{h^3}{6}f^{(3)}(x_i) + \frac{h^4}{24}f^{(4)}(x_i) + \mathcal{O}(h^5),$$
$$f(x_{i-1}) = f(x_i) - hf'(x_i) + \frac{h^2}{2}f''(x_i) - \frac{h^3}{6}f^{(3)}(x_i) + \frac{h^4}{24}f^{(4)}(x_i) + \mathcal{O}(h^5).$$

Adding the two equations:

$$f(x_{i+1}) + f(x_{i-1}) = 2f(x_i) + h^2 f''(x_i) + \frac{h^4}{12}f^{(4)}(x_i) + \mathcal{O}(h^5).$$

Rearranging gives:

$$f''(x_i) \approx \frac{f(x_{i+1}) - 2f(x_i) + f(x_{i-1})}{h^2}, \quad \text{with error } \mathcal{O}(h^2).$$

Fourth-order derivative:

$$f^{(4)}(x_i) \approx \frac{f(x_{i+2}) - 4f(x_{i+1}) + 6f(x_i) - 4f(x_{i-1}) + f(x_{i-2})}{h^4}, \quad \text{with error } \mathcal{O}(h^2).$$

Proof. To derive a second-order accurate formula for the fourth derivative, we use Taylor expansions at points $x_{i+2}, x_{i+1}, x_i, x_{i-1}, x_{i-2}$. The relevant expansions are:

$$f(x_{i+2}) = f(x_i) + 2hf'(x_i) + \frac{(2h)^2}{2!}f''(x_i) + \frac{(2h)^3}{3!}f^{(3)}(x_i) + \frac{(2h)^4}{4!}f^{(4)}(x_i) + \mathcal{O}(h^5),$$
$$f(x_{i+1}) = f(x_i) + hf'(x_i) + \frac{h^2}{2}f''(x_i) + \frac{h^3}{6}f^{(3)}(x_i) + \frac{h^4}{24}f^{(4)}(x_i) + \mathcal{O}(h^5),$$

$$f(x_{i-1}) = f(x_i) - hf'(x_i) + \frac{h^2}{2}f''(x_i) - \frac{h^3}{6}f^{(3)}(x_i) + \frac{h^4}{24}f^{(4)}(x_i) + \mathcal{O}(h^5),$$
$$f(x_{i-2}) = f(x_i) - 2hf'(x_i) + \frac{(2h)^2}{2!}f''(x_i) - \frac{(2h)^3}{3!}f^{(3)}(x_i) + \frac{(2h)^4}{4!}f^{(4)}(x_i) + \mathcal{O}(h^5).$$

Taking a linear combination:

$$f^{(4)}(x_i) \approx \frac{f(x_{i+2}) - 4f(x_{i+1}) + 6f(x_i) - 4f(x_{i-1}) + f(x_{i-2})}{h^4}, \quad \text{with error } \mathcal{O}(h^2).$$

This formula eliminates the contributions of the lower-order derivatives and isolates the fourth derivative, with second-order accuracy.

3 Energy Quanto Options: Applying the COS method in a Markov-modulated Framework

In this chapter, we study Energy Quanto options under two different models and demonstrate how to apply the COS method in this set-up, which has not been done in existing literature. Results show improved performance compared to the methods proposed in [2].

3.1 Payoff and discounted characteristic function

Energy quanto options have a payoff which resembles the product of the payoff of two Asian options if the energy spot or weather indexes are considered to be underlying assets. However, it was shown in [5] that if energy futures contracts are used as the underlying assets, then the payoff becomes the product of the payoffs of two European vanilla options. The payoff of an Energy Quanto option follows:

$$V(F_1, F_2, t) = \mathbb{E}\left[e^{-U(t)}\max((F_1(t,\tau) - K_1,\ 0)\max(F_2(t,\tau) - K_2,\ 0))\right],$$

where $t \leq \tau$ is the maturity of the option, F_1 and F_2 are the current futures prices, and K_1 and K_2 are the strike prices.

To model futures prices, we introduce the same model as described in Section 2.4. Since the interest rate process $U(t)$ is not constant, we consider the discounted characteristic function of $\left(\ln\left(\frac{F_{\ell_1}}{K_1}\right), \ln\left(\frac{F_{\ell_2}}{K_2}\right)\right)$ to reduce the dimensionality of the pricing problem by one:

$$\begin{aligned}\hat{\psi}(\omega_1, \omega_2) &= \mathbb{E}\left[e^{-U(T)+\mathrm{i}<(\omega_1,\omega_2),(\ln(F_{\ell_1}/K_1),\ln((F_{\ell_2}/K_2))>}\right] \\ &= \mathbb{E}\left[e^{-U(T)+\mathrm{i}(\omega_1\ln(F_{\ell_1}(0)/K_1)+\omega_2\ln(F_{\ell_2}(0)/K_2))+\mathrm{i}(\omega_1\Lambda^{\ell_1}+\omega_2\Lambda^{\ell_2})+\mathrm{i}(\omega_1 X^{\ell_1}+\omega_2 X^{\ell_2})}\right]\end{aligned}$$

Defining: $s_1 = \Lambda^{\ell_1} + X^{\ell_1}$ and $s_2 = \Lambda^{\ell_2} + X^{\ell_2}$

$$= e^{\mathrm{i}(\omega_1 \ln(F_{\ell_1}(0)/K_1) + \omega_2 \ln(F_{\ell_2}(0)/K_2))} \int_{\mathbb{R}} \int_{\mathbb{R}} e^{\mathrm{i}\omega_1 s_1 + \mathrm{i}\omega_2 s_2} \underbrace{\int_{\mathbb{R}} e^{-u} f(s_1, s_2) du}_{\hat{f}(s_1, s_2)} ds_1 ds_2,$$

where $\hat{f}(\cdot,\cdot)$ is the discounted density. Using Lemma 2.4.1 and the fact that $\Lambda_l(t,\tau) = \int_0^t \sum_{j=1}^J \mu_{j,l}(s,\tau) \mathbb{1}_{M(s)=j} ds$, we recognize:

$$\Gamma_2(t) = -U(t) = -\int_0^t \sum_{j=1}^J r_j \mathbb{1}_{M(s)=j} ds$$

$$\Gamma_1(t,\tau) = \mathrm{i}(\omega_1 \Lambda^{\ell_1} + \omega_2 \Lambda^{\ell_2}) = \int_0^t \sum_{j=1}^J \mathrm{i}(\omega_1 \mu_{j,1}(s,\tau) + \omega_2 \mu_{j,2}(s,\tau)) \mathbb{1}_{M(s)=j} ds.$$

This implies that:

$$\hat{\psi}(\omega_1, \omega_2) = e^{\mathrm{i}(\omega_1 \ln(F_{\ell_1}(0)/K_1) + \omega_2 \ln(F_{\ell_2}(0)/K_2))} \cdot \mathbf{p} e^{(Q+D)T + \int_0^T A(s,\tau) ds} \mathbf{1} \tag{3.1}$$

with:

- $D = -\mathrm{diag}(r)$;
- Q is the generator matrix of the Markov process, which in this case is a 2×2 matrix because we are considering two possible states;
- $\mathbf{p} = (p_0, p_1, \ldots, p_J) \in \mathbb{R}^J$ is the initial probability vector;
- $A_{jj}(t,\tau) = \mathrm{i}\omega_1 \mu_{j,1}(t,\tau) + \mathrm{i}\omega_2 \mu_{j,2}(t,\tau) + \frac{\partial \phi_j}{\partial y}(t, y, \tau, \mathrm{i}\mathbf{a})\Big|_{y=t}$, with $\mathbf{a} = (\mathrm{i}\omega_1, \mathrm{i}\omega_2)$;
- $\mathbf{1}$ is a vector whose elements are all equal to 1.

This discounted characteristic function will be used throughout our work together with the COS method to determine the price of Energy Quanto options.

3.2 Case I: A Gaussian Markov-modulated model for temperature and electricity futures prices

In this section, we introduce a Markov-modulated additive model for temperature and electricity futures prices, following the analysis conducted in [2].

Note: To describe the temperature futures price we use the superscript c (from Celsius, as in [2]), because t, T are already used to describe time and C was used previously to describe a claim.

The **temperature futures price process** $F^c(t,\tau)$ is defined as:

$$F^c(t,\tau) = F^c(0,\tau)e^{\Lambda^c(t,\tau)+X^c(t,\tau)},$$

where Λ^c is the drift process described in Proposition 2.4.1 and

$$X^c(t,\tau) = \int_0^t \sigma^c(u)e^{-\alpha^c(\tau-u)}dW^c(u).$$

In particular, W^c is a standard Brownian motion under the measure $\mathbb{Q}$, $\alpha^c > 0$ is the mean-reversion speed, and $\sigma^c(t)$ is defined as follows:

$$\sigma^c(t) = \ln\left(\left(c_1 + c_2\cos\left(\frac{2\pi t}{365}\right) + c_3\sin\left(\frac{2\pi t}{365}\right)\right)^{1/2}\right),$$

with $c_1,\ c_2,\ c_3$ constants. This volatility function is introduced in [2] based on empirical findings from [3], with the aim of capturing the seasonal component observed in the data.

As for the **electricity futures price process**, we define a Markov-modulated model:

$$F^E(t,\tau) = F^E(0,\tau)e^{\Lambda^E(t,\tau)+X^E(t,\tau)},$$

where

$$X^E(t,\tau) = \int_0^t \sigma^E(M(u))e^{-\int_0^\tau \alpha^E(M(u))du}dW^E(u).$$

In this case, W^E is a standard Brownian motion under $\mathbb{Q}$, α^E is the mean-reversion speed modulated by the Markov process $M(t)$:

$$\alpha^E(M(t)) = \sum_{j=1}^{J} \alpha_j^E \mathbb{1}_{M(t)=j},$$

with α_j^E some positive constants for $j = 1, ..., J$. The volatility σ^E is as well Markov-modulated:

$$\sigma^E(M(t)) = \sum_{j=1}^{J} \sigma_j^E \mathbb{1}_{M(t)=j},$$

with σ_j^E some positive constants for $j = 1, ..., J$. The process Λ^E is the drift process and follows the description of Proposition 2.4.1. The two Brownian motions W^C and

W^E are correlated by a Markov-modulated function $\rho(M(t)) = \sum_{j=1}^{J} \rho_j \mathbb{1}_{M(t)=j}$, with $\rho_j \in (-1, 1)$ constants.
From the definition of the process X^E, we understand that when $M = j$, X^E behaves like the additive process Z_j^E whose dynamics are:

$$dZ_j^E(t,\tau) = \sigma_j^E e^{-\alpha_j^E(t-\tau)} dW^E(t).$$

The following proposition, given in [5], describes the characteristic exponent ϕ_j of the process $\mathbf{Z}_j^{c,E}(t,\tau) = (X^c(t,\tau), Z_j^E(t,\tau))$.

Proposition 3.2.1. *For $\boldsymbol{\theta} = (\theta^c, \theta^E) \in \mathbb{R}^2$ and $s \leq t \leq \tau$, the joint characteristic exponent of the process $\mathbf{Z}_j^{c,E}(t,\tau)$ is:*

$$\begin{aligned}\phi_j(s,t,\tau,(\theta^c,\theta^E)) = -\frac{1}{2}\Bigg(&(\theta^c)^2 \int_s^t (\sigma^c(u))^2 e^{-2\alpha^c(\tau-u)} du \\ &+ (\theta^E)^2 \int_s^t (\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-u)} du \\ &+ 2\theta^c\theta^E\rho_j \int_s^t \sigma^c(u)\sigma_j^E e^{-(\tau-u)(\alpha^c+\alpha_j^E)} du.\end{aligned}$$

To apply the COS Method, we determine an explicit expression for the components of Equation 3.1. We start by computing:

$$\begin{aligned}\frac{\partial\phi_j}{\partial t}(s,t,\tau,\boldsymbol{\theta}) &= -\frac{1}{2}\frac{\partial}{\partial t}\Bigg((\theta^c)^2 \int_s^t (\sigma^c(u))^2 e^{-2\alpha^c(\tau-u)}\, du \\ &\qquad + (\theta^E)^2 \int_s^t (\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-u)}\, du \\ &\qquad + 2\theta^c\theta^E\rho_j \int_s^t \sigma^c(u)\sigma_j^E e^{-(\tau-u)(\alpha^c+\alpha_j^E)}\, du\Bigg) \\ &= -\frac{1}{2}\Big((\theta^c)^2(\sigma^c(t))^2 e^{-2\alpha^c(\tau-t)} \\ &\qquad + (\theta^E)^2(\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-t)} \\ &\qquad + 2\theta^c\theta^E\rho_j\sigma^c(t)\sigma_j^E e^{-(\tau-t)(\alpha^c+\alpha_j^E)}\Big).\end{aligned}$$

This means that:

$$\left.\frac{\partial\phi_j}{\partial y}(t,y,\tau,\boldsymbol{\theta})\right|_{y=t} = -\frac{1}{2}\Big((\theta^c)^2(\sigma^c(t))^2 e^{-2\alpha^c(\tau-t)} + (\theta^E)^2(\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-t)}$$

$$+ 2\theta^c\theta_j^E\rho_j\sigma^c(t)\sigma_j^E e^{-(\tau-t)(\alpha^c+\alpha_j^E)}\Big).$$

Recalling the definition of a standard basis vector $\mathbf{e}_i$, whose elements are all equal to 0, expect for the i-th element which is equal to 1, we obtain the following quantities $\forall j = 1, ..., J$:

$$\mu_{j,1}(t,\tau) = -\frac{\partial\phi_j}{\partial y}(t,y,\tau,\mathrm{i}\mathbf{e}_1)\Big|_{y=t} = -\frac{1}{2}(\sigma^c(t))^2 e^{-2\alpha^c(\tau-t)} \tag{3.2}$$

$$\mu_{j,2}(t,\tau) = -\frac{\partial\phi_j}{\partial y}(t,y,\tau,\mathrm{i}\mathbf{e}_2)\Big|_{y=t} = -\frac{1}{2}(\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-t)} \tag{3.3}$$

Since in this case the vector $\mathbf{a} = (i\omega_1, i\omega_2)$, we have:

$$\begin{aligned}\frac{\partial\phi_j}{\partial y}(t,y,\tau,i\mathbf{a}) =& \frac{1}{2}\Big((\mathrm{i}\omega_1)^2(\sigma^c(t))^2 e^{-2\alpha^c(\tau-t)} \\ &+ (\mathrm{i}\omega_2)^2(\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-t)} \\ &+ 2(\mathrm{i}\omega_1)(\mathrm{i}\omega_2)\rho_j\sigma^c(t)\sigma_j^E e^{-(\tau-t)(\alpha^c+\alpha_j^E)}\Big).\end{aligned}$$

Finally, we have an explicit expression for the matrix A in Equation 3.1:

$$\begin{aligned}A_{j,j}(t,\tau) =& \frac{1}{2}\Big(\mathrm{i}\omega_1(\mathrm{i}\omega_1 - 1)(\sigma^c(t))^2 e^{-2\alpha^c(\tau-t)} \\ &+ \ \mathrm{i}\omega_2(\mathrm{i}\omega_2 - 1)(\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-t)} \\ &+ \ 2(\mathrm{i}\omega_1)(\mathrm{i}\omega_2)\rho_j\sigma^c(t)\sigma_j^E e^{-(\tau-t)(\alpha^c+\alpha_j^E)}\Big).\end{aligned}$$

3.3 Case II: A Markov-modulated model for gas and electricity futures prices

In this section, we introduce a Markov-modulated additive model for gas and electricity futures prices that incorporates jump components, following the framework of [2]. The inclusion of jumps is motivated by empirical evidence from UK gas spot price data, which reveals the presence of extreme price spikes, as noted in [3]. The regime-switching structure is based on the approach in [17], where it was demonstrated that such models effectively capture the statistical features of electricity prices.

We define an **electricity price process**:

$$F^E(t,\tau) = F^E(0,\tau)e^{\Lambda^E(t,\tau)+X^E(t,\tau)},$$

where $X^E(t,\tau)$ is an additive process which evolves as:

$$\int_0^t \sigma^E(M(u))e^{-\int_u^\tau \alpha^E(M(u))du}dW^E(u) + \int_0^t e^{-\int_u^\tau \beta^E(M(u))du}d\mathcal{J}^E(u)$$

and Λ^E follows Proposition 2.4.1. In this model, W^E is a standard Brownian motion under $\mathbb{Q}$, α^E represents the Markov-modulated mean-reversion by:

$$\alpha^E(M(t)) = \sum_{j=1}^{J} \alpha_j^E \mathbb{1}_{M(t)=j}$$

and σ^E is the Markov-modulated volatility:

$$\sigma^E(M(t)) = \sum_{j=1}^{J} \sigma_j^E \mathbb{1}_{M(t)=j},$$

with $\alpha_j^E,\ \sigma_j^E$ strictly positive constants for $j = 1, ..., J$.
As for the jump part of the process, β^E is the mean-reversion speed:

$$\beta^E(M(t)) = \sum_{j=1}^{J} \beta_j^E \mathbb{1}_{M(t)=j},$$

with β_j^E some strictly positive constants for $j = 1, ..., J$. J^E is a Markov-modulated compound Poisson process CPP (Definition 2.5.3), which means that when $M = j$, $\mathcal{J}^E$ behaves like the CPP $\mathcal{J}_j^E$:

$$\mathcal{J}_j^E(t) = \sum_{m=1}^{P_j(t)} Y_{j,m}^E \qquad j = 1, ..., J,$$

where every $P_j(t)$ is a non regime-switching Poisson process and $Y_{j,m}^E$ are independent and identically distributed variables for $m = 1, ..., P_j(t)$. The seasonal intensity $\lambda_j(t)$ of the processes $P_j(t)$ is defined as:

$$\lambda_j(t) = c_{1,j} + c_{2,j}\cos\Big(\frac{2\pi t}{365}\Big) \geq 0, \tag{3.4}$$

with $c_{1,j}$ and $c_{2,j}$ some constants. The jump sizes $Y_{j,m}^E$ follow a two-sided exponential distribution with the following density function:

$$g_j^E(y) = p_j^E \eta_{1,j}^E e^{-\eta_{1,j}^E}\mathbb{1}_{y\geq 0} + (1-p_j^E)\eta_{2,j}^E e^{\eta_{2,j}^E y}\mathbb{1}_{y<0},$$

with $p_j^E \in [0,1],\ \eta_{1,j}^E > 0,\ \eta_{2,j}^E > 0,\ j = 1, ..., J$. All the random processes, i.e. the Brownian motion $W^E(t)$, the Markov process $M(t)$, the Poisson process $P_j(t)$ and the jump sizes $(Y_{j,m}^E)_{j\in\mathscr{S}}$ for $m = 1, ..., P_j(t)$, are assumed to be mutually independent.

Analogously, we introduce the dynamics of the **gas futures price**:

$$F^G(t,\tau) = F^G(0,\tau)e^{\Lambda^G(t,\tau)+X^G(t,\tau)},$$

where $X^G(t,\tau)$ is an additive process that evolves as:

$$\int_0^t \sigma^G(M(u))e^{-\int_u^\tau \alpha^G(M(u))du}dW^G(u) + \int_0^t e^{-\int_u^\tau \beta^G(M(u))du}d\mathcal{J}^G(u)$$

and Λ^G is defined as in Proposition 2.4.1. In this model, W^G is a standard Brownian motion under $\mathbb{Q}$, that is related to the above-defined process W^E by: $d[W^E, W^G](t) = \rho(M(t))dt$, where

$$\rho(M(t)) = \sum_{j=1}^{J} \rho_j \mathbb{1}_{M(t)=j}$$

and $\rho_j \in (-1,1)$ are the constant coefficients of correlation.

Additionally, α^G represents the Markov-modulated mean-reversion by:

$$\alpha^G(M(t)) = \sum_{j=1}^{J} \alpha_j^G \mathbb{1}_{M(t)=j}$$

and σ^G is the Markov-modulated volatility:

$$\sigma^G(M(t)) = \sum_{j=1}^{J} \sigma_j^G \mathbb{1}_{M(t)=j},$$

with $\alpha_j^G,\ \sigma_j^G$ strictly positive constants for $j = 1, ..., J$.

Regarding the jump part of the process, the CPP $P_j(t)$ that governs the gas dynamics is taken to be the same as that for electricity. This reflects the close relationship between the two markets, where a sudden jump in gas prices typically triggers a similar jump in electricity prices, and vice versa.

In the same manner as above, we define β^G as the mean-reversion speed:

$$\beta^G(M(t)) = \sum_{j=1}^{J} \beta_j^G \mathbb{1}_{M(t)=j},$$

with β_j^G some strictly positive constants for $j = 1, ..., J$. $\mathcal{J}^G$ is a Markov-modulated compound Poisson process (CPP), which means that when $M = j$, $\mathcal{J}^G$ behaves like the CPP $\mathcal{J}_j^G$:

$$\mathcal{J}_j^G(t) = \sum_{m=1}^{P_j(t)} Y_{j,m}^G \qquad j = 1, ..., J,$$

where $Y_{j,m}^G$ are independent and identically distributed variables for $m = 1, ..., P_j(t)$. The jump sizes $Y_{j,m}^G$ follow a two-sided exponential distribution with the following density function:

$$g_j^G(y) = p_j^G \eta_{1,j}^G e^{-\eta_{1,j}^G} \mathbb{1}_{y\geq 0} + (1 - p_j^G)\eta_{2,j}^G e^{\eta_{2,j}^G y} \mathbb{1}_{y<0},$$

with $p_j^G \in [0, 1]$, $\eta_{1,j}^G > 0$, $\eta_{2,j}^G > 0$, $j = 1, ..., J$.
Furthermore, it is again assumed here that W^G, $M(t)$, $P_j(t)$, $(Y_{j,m}^G)_{j\in\mathscr{S}}$ for $m = 1, ..., P_j(t)$, are all mutually independent.
When $M = j$, X^E and X^G evolve as the additive processes Z_j^E and Z_j^G:

$$\begin{aligned} dZ_j^E(t,\tau) &= \sigma_j^E e^{-\alpha_j^E(\tau-t)} dW^E(t) + e^{-\beta_j^E(\tau-t)} d\mathcal{J}_j^E(t), \\ dZ_j^G(t,\tau) &= \sigma_j^G e^{-\alpha_j^G(\tau-t)} dW^G(t) \ + e^{-\beta_j^G(\tau-t)} d\mathcal{J}_j^G(t). \end{aligned}$$

We further assume that the jump sizes of gas and electricity prices are independent. This allows us to define the characteristic exponent of $\mathbf{Z}_j^{E,G}(t,\tau) = (Z_j^E(t,\tau), Z_j^G(t,\tau))$:

Proposition 3.3.1. *For $\boldsymbol{\theta} = (\theta^E, \theta^G) \in \mathbb{R}^2$ and $s \leq t \leq \tau$, the joint characteristic exponent of the process $\mathbf{Z}_j^{E,G}(t,\tau)$ is:*

$$\begin{aligned} \phi_j(s,t,\tau,\boldsymbol{\theta}) = &-\frac{1}{2}\Bigg((\theta^E)^2 \int_s^t (\sigma_j^E)^2 e^{-2\alpha_j^E(\tau-u)} du + (\theta^G)^2 \int_s^t (\sigma_j^G)^2 e^{-2\alpha_j^G(\tau-u)} du \\ &+ 2\theta^E\theta^G\rho_j \int_s^t \sigma_j^E\sigma_j^G e^{-(\tau-u)(\alpha_j^E+\alpha_j^G)} du\Bigg) \\ &+ \int_s^t \lambda_j(u)\Big[\boldsymbol{\Psi}_j(-\theta^E e^{-\beta_j^E(\tau-u)}, -\theta^G e^{-\beta_j^G(\tau-u)}) - 1\Big] du, \end{aligned}$$

where $\boldsymbol{\Psi}_j(a,b)$ is the joint characteristic function of the joint jump sizes $(Y_{j,m}^G, Y_{j,m}^E)$ for $(a,b) \in \mathbb{R}^2$.

To ensure that Assumption 2.4.1 is satisfied, we further assume that $\eta_{i,j}^E > 2$ and $\eta_{i,j}^G > 2$. Under this condition, if we define $C = \min(\eta_{1,j}^E, \eta_{2,j}^G)$, , then clearly $C > 2$. It follows that for all $\theta^E \in [-C, C]$ and all $u \in [0, \tau]$, we have: $-\eta_{1,j}^E e^{\beta_j^E(\tau-u)} < \theta^E < \eta_{2,j}^E e^{\beta_j^E(\tau-u)}$, which ensures that Assumption 2.4.1 holds.

As in Section 3.2, we derive an explicit expression for the components of Equation 3.1. Firstly, we define the characteristic function $\mathbf{\Psi}_j(a,b)$ of the joint jump sizes $(Y^G_{j,m}, Y^E_{j,m})$ for $(a,b) \in \mathbb{R}^2$:

$$\Psi_j(u_E, u_G) = \mathbb{E}(e^{\mathrm{i}u_E Y^E_{j,m} + \mathrm{i}u_G Y^G_{j,m}})$$

Exploiting the mutual independence between $Y^E_{j,m}$ and $Y^G_{j,m}$:

$$\begin{aligned} &= \mathbb{E}(e^{\mathrm{i}u_E Y^E_{j,m}}) \cdot \mathbb{E}(e^{\mathrm{i}u_G Y^G_{j,m}}) \\ &= \psi^E_j(u_E) \cdot \psi^G_j(u_G) \end{aligned}$$

where $\psi^E_j(u_E)$ and $\psi^G_j(u_G)$ are the characteristic functions respectively of $Y^E_{j,m}$ and $Y^G_{j,m}$.

The two processes $Y^E_{j,m}$ and $Y^G_{j,m}$ follow the same distribution, so we determine one characteristic function that works for both of them:

$$\psi_j(u) = \mathbb{E}(e^{\mathrm{i}uY_{j,m}}) = \int_{-\infty}^{\infty} e^{\mathrm{i}uy} g(y) dy$$

with $g(y)$ the density function of $Y_{j,m}$,

$$\begin{aligned} &= \int_{-\infty}^{\infty} e^{\mathrm{i}uy} (p_j \eta_{1,j} e^{-\eta_{1,j}} \mathbb{1}_{y \geq 0} + (1-p_j)\eta_{2,j} e^{\eta_{2,j} y} \mathbb{1}_{y<0}) dy \\ &= \int_0^{\infty} e^{\mathrm{i}uy} \cdot p_j \eta_{1,j} e^{-\eta_{1,j}} dy + \int_{-\infty}^{0} e^{\mathrm{i}uy} \cdot (1-p_j)\eta_{2,j} e^{\eta_{2,j} y} dy \\ &= p_j \eta_{1,j} \frac{e^{(\mathrm{i}u - \eta_{1,j})y}}{\mathrm{i}u - \eta_{1,j}} \Big|_0^{\infty} + (1-p_j)\eta_{2,j} \frac{e^{(\mathrm{i}u - \eta_{2,j})y}}{\mathrm{i}u + \eta_{2,j}} \Big|_{-\infty}^{0} \\ &= \frac{p_j \eta_{1,j}}{\eta_{1,j} - \mathrm{i}u} + \frac{(1-p_j)\eta_{2,j}}{\eta_{2,j} + \mathrm{i}u} \end{aligned}$$

$$\begin{aligned} \Rightarrow \psi^E_j(u_E) &= \frac{p^E_j \eta^E_{1,j}}{\eta^E_{1,j} - \mathrm{i}u_E} + \frac{(1-p^E_j)\eta^E_{2,j}}{\eta^E_{2,j} + \mathrm{i}u_E} \\ \Rightarrow \psi^G_j(u_G) &= \frac{p^G_j \eta^G_{1,j}}{\eta^G_{1,j} - \mathrm{i}u_G} + \frac{(1-p^G_j)\eta^G_{2,j}}{\eta^G_{2,j} + \mathrm{i}u_G}. \end{aligned}$$

Finally, we have:

$$\Psi_j(u_E, u_G) = \left(\frac{p^E_j \eta^E_{1,j}}{\eta^E_{1,j} - \mathrm{i}u_E} + \frac{(1-p^E_j)\eta^E_{2,j}}{\eta^E_{2,j} + \mathrm{i}u_E} \right) \cdot \left(\frac{p^G_j \eta^G_{1,j}}{\eta^G_{1,j} - \mathrm{i}u_G} + \frac{(1-p^G_j)\eta^G_{2,j}}{\eta^G_{2,j} + \mathrm{i}u_G} \right). \tag{3.5}$$

To apply the COS method, we determine an explicit expression for the components of 3.1. We start by computing:

$$\begin{aligned}\frac{\partial \phi_j}{\partial t}(s,t,\tau,\boldsymbol{\theta}) = \frac{\partial}{\partial t}\Bigg(-\frac{1}{2}\Bigg((\theta^E)^2\int_s^t(\sigma_j^E)^2e^{-2\alpha_j^E(\tau-u)}du + (\theta^G)^2\int_s^t(\sigma_j^G)^2e^{-2\alpha_j^G(\tau-u)}du \\ +2\theta^E\theta^G\rho_j\int_s^t\sigma_j^E\sigma_j^Ge^{-(\tau-u)(\alpha_j^E+\alpha_j^G)}du\Bigg) \\ +\int_s^t\lambda_j(u)\left[\boldsymbol{\Psi}_j(-\theta^Ee^{-\beta_j^E(\tau-u)},-\theta^Ge^{-\beta_j^G(\tau-u)})-1\right]du\Bigg) \\ = -\frac{1}{2}\Big((\theta^E)^2(\sigma_j^E)^2e^{-2\alpha_j^E(\tau-t)} + (\theta^G)^2(\sigma_j^G)^2e^{-2\alpha_j^G(\tau-t)} \\ + 2\theta^E\theta^G\rho_j\sigma_j^E\sigma_j^Ge^{-(\tau-t)(\alpha_j^E+\alpha_j^G)}\Big) \\ +\lambda_j(t)\left[\boldsymbol{\Psi}_j(-\theta^Ee^{-\beta_j^E(\tau-t)},-\theta^Ge^{-\beta_j^G(\tau-t)})-1\right]\end{aligned}$$

where $\lambda_j(\cdot)$ is defined in Equation 3.4 and $\Psi_j(\cdot,\cdot)$ in Equation 3.5.

From the above expression, we obtain:

$$\begin{aligned}\frac{\partial \phi_j}{\partial y}(t,y,\tau,\boldsymbol{\theta})\Big|_{y=t} = -\frac{1}{2}\Big((\theta^E)^2(\sigma_j^E)^2e^{-2\alpha_j^E(\tau-t)} + (\theta^G)^2(\sigma_j^G)^2e^{-2\alpha_j^G(\tau-t)} \\ +2\theta^E\theta^G\rho_j\sigma_j^E\sigma_j^Ge^{-(\tau-t)(\alpha_j^E+\alpha_j^G)}\Big) \\ +\lambda_j(t)\Big[\boldsymbol{\Psi}_j(-\theta^Ee^{-\beta_j^E(\tau-t)},-\theta^Ge^{-\beta_j^G(\tau-t)})-1\Big].\end{aligned}$$

In a similar way to Section 3.2, recalling the definition of the standard basis vectors $\mathbf{e}_1 = (1,0)$ and $\mathbf{e}_2 = (0,1)$, we find the following quantities $\forall j = 1,...,J$:

$$\begin{aligned}\mu_{j,1}(t,\tau) &= -\frac{\partial \phi_j}{\partial y}(t,y,\tau,\mathrm{i}\mathbf{e}_1)\Big|_{y=t} \\ &= -\frac{1}{2}(\sigma_j^E)^2e^{-2\alpha_j^E(\tau-t)} + \Big(c_{1,j}+c_{2,j}\cos\Big(\frac{2\pi t}{365}\Big)\Big)\left[\left(\frac{p_j^E\eta_{1,j}^E}{\eta_{1,j}^E-e^{-\beta_j^E(\tau-t)}}+\frac{(1-p_j^E)\eta_{2,j}^E}{\eta_{2,j}^E+e^{-\beta_j^E(\tau-t)}}\right)-1\right] \\ \mu_{j,2}(t,\tau) &= -\frac{\partial \phi_j}{\partial y}(t,y,\tau,\mathrm{i}\mathbf{e}_2)\Big|_{y=t} \\ &= -\frac{1}{2}(\sigma_j^G)^2e^{-2\alpha_j^G(\tau-t)} + \Big(c_{1,j}+c_{2,j}\cos\Big(\frac{2\pi t}{365}\Big)\Big)\left[\left(\frac{p_j^G\eta_{1,j}^G}{\eta_{1,j}^G-e^{-\beta_j^G(\tau-t)}}+\frac{(1-p_j^G)\eta_{2,j}^G}{\eta_{2,j}^G+e^{-\beta_j^G(\tau-t)}}\right)-1\right]\end{aligned}$$

Lastly, we recall that the vector $\mathbf{a} = (\mathrm{i}\omega_1, \mathrm{i}\omega_2)$, so we have:

$$\frac{\partial \phi_j}{\partial y}(t,y,\tau,\mathrm{i}\mathbf{a})\Big|_{y=t} = \frac{1}{2}\Big((\mathrm{i}\omega_1)^2(\sigma_j^E)^2e^{-2\alpha_j^E(\tau-t)} + (\mathrm{i}\omega_2)^2(\sigma_j^G)^2e^{-2\alpha_j^G(\tau-t)}$$

$$+ 2(\mathrm{i}\omega_1)(\mathrm{i}\omega_2)\rho_j\sigma_j^E\sigma_j^G e^{-(\tau-t)(\alpha_j^E+\alpha_j^G)}\Big)$$
$$+ \left(c_{1,j} + c_{2,j}\cos\left(\frac{2\pi t}{365}\right)\right)\cdot\left[\left(\frac{p_j^E\eta_{1,j}^E}{\eta_{1,j}^E - (\mathrm{i}\omega_1)e^{-\beta_j^E(\tau-t)}} + \frac{(1-p_j^E)\eta_{2,j}^E}{\eta_{2,j}^E + (\mathrm{i}\omega_1)e^{-\beta_j^E(\tau-t)}}\right)\right.$$
$$\left.\cdot\left(\frac{p_j^G\eta_{1,j}^G}{\eta_{1,j}^G - (\mathrm{i}\omega_2)e^{-\beta_j^G(\tau-t)}} + \frac{(1-p_j^G)\eta_{2,j}^G}{\eta_{2,j}^G + (\mathrm{i}\omega_2)e^{-\beta_j^G(\tau-t)}}\right) - 1\right].$$

Now that we have all these elements, we can derive an explicit expression for the matrix A in Equation 3.1.

3.4 COS Method for two-dimensional processes

In the models presented in Sections 3.2 and 3.3, three distinct stochastic processes can be identified: the discounting process $U(t)$, and the two futures price processes, temperature and electricity in Section 3.2, and gas and electricity in Section 3.3, respectively. As a result, the pricing problem is initially three-dimensional. However, by working with the discounted characteristic function (see Equation 3.1), the dimensionality of the problem is effectively reduced to two, as the influence of the discounting process $U(t)$ is analytically integrated out.

The COS method has already been introduced in Section 2.7 as a highly effective numerical technique for pricing one-dimensional European options. In [29], its applicability was further extended to the two-dimensional setting. In this work, we follow their methodology to present the two-dimensional COS method. For the sake of clarity, while introducing the method, we will assume a constant interest rate r, resulting in a formulation analogous to the one-dimensional case presented in Section 2.7. Later, when we apply this method to the Energy Quanto option pricing problem, the discount factor under constant interest rates will no longer appear, as we will make use of the discounted characteristic function instead.

Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space, $T > 0$ a finite terminal time and $\mathbb{F} = (\mathcal{F}_s)_{0\leq s\leq T}$ a filtration. The log-asset prices are described by the two-dimensional process $\mathbf{X}_t = (X_t^1, X_t^2)$ on the filtered probability space. If $g(\cdot)$ is the payoff function, then the value of a European-type option in two dimensions is given by the risk neutral valuation formula:

$$v(t_0, \mathbf{x}) = e^{-r\Delta t}\,\mathbb{E}^{\mathbb{Q}}[g(\mathbf{X}_T) \mid \mathbf{x}] = e^{-r\Delta t}\int_{\mathbb{R}}\int_{\mathbb{R}} g(\mathbf{y})f(\mathbf{y} \mid \mathbf{x})d\mathbf{y},$$

where $\mathbf{x} = (x_1, x_2)$ is the current state, $f(y_1, y_2 \mid x_1, x_2)$ is the conditional PDF, r is the risk-free interest rate, and $\Delta t = T - t_0$. Assuming that the integrand is integrable,

the first approximation is to truncate the integration range to $[a_1, b_1] \times [a_2, b_2] \subset \mathbb{R}^2$ without losing significant accuracy.

$$v_1(\mathbf{x}, t_0) = e^{-r\Delta t} \int_{a_1}^{b_1} \int_{a_2}^{b_2} g(\mathbf{y}) f(\mathbf{y} \mid \mathbf{x}) dy_1 dy_2$$

Replace the conditional density by its Fourier cosine expansion in $\mathbf{y}$ on $[a_1, b_1] \times [a_2, b_2]$

$$= e^{-r\Delta t} \int_{a_1}^{b_1} \int_{a_2}^{b_2} g(\mathbf{y}) \sum_{n_1=0}^{+\infty}{}' \sum_{n_2=0}^{+\infty}{}' A_{n_1,n_2}(\mathbf{x}) \cos\left(n_1\pi \frac{y_1 - a_1}{b_1 - a_1}\right) \cos\left(n_2\pi \frac{y_2 - a_2}{b_2 - a_2}\right) dy_1 dy_2$$

with the coefficients A_{n_1,n_2} defined as:

$$A_{n_1,n_2}(\mathbf{x}) = \frac{2}{b_1 - a_1} \frac{2}{b_2 - a_2} \int_{a_1}^{b_1} \int_{a_2}^{b_2} f(\mathbf{y} \mid \mathbf{x}) \cos\left(n_1\pi \frac{y_1 - a_1}{b_1 - a_1}\right) \cos\left(n_2\pi \frac{y_2 - a_2}{b_2 - a_2}\right) dy_1 dy_2.$$

We exchange summation and integration and define the coefficients V_{n_1,n_2}:

$$V_{n_1,n_2}(T) = \frac{2}{b_1 - a_1} \frac{2}{b_2 - a_2} \int_{a_1}^{b_1} \int_{a_2}^{b_2} g(\mathbf{y}) \cos\left(n_1\pi \frac{y_1 - a_1}{b_1 - a_1}\right) \cos\left(n_2\pi \frac{y_2 - a_2}{b_2 - a_2}\right) dy_1 dy_2.$$

The second approximation is to truncate the series summations:

$$v_2(\mathbf{x}, t_0) = \frac{b_1 - a_1}{2} \frac{b_2 - a_2}{2} e^{-r\Delta t} \sum_{n_1=0}^{N_1-1}{}' \sum_{n_2=0}^{N_2-1}{}' A_{n_1,n_2}(\mathbf{x}) V_{n_1,n_2}(T).$$

Lastly, the coefficients $A_{n_1,n_2}(\mathbf{x})$ are approximated by:

$$F_{n_1,n_2}(\mathbf{x}) = \frac{2}{b_1 - a_1} \frac{2}{b_2 - a_2} \int_{\mathbb{R}} \int_{\mathbb{R}} f(\mathbf{y} \mid \mathbf{x}) \cos\left(n_1\pi \frac{y_1 - a_1}{b_1 - a_1}\right) \cos\left(n_2\pi \frac{y_2 - a_2}{b_2 - a_2}\right) dy_1 dy_2.$$

which can be rewritten using the trigonometric relation $2\cos(\alpha)\cos(\beta) = \cos(\alpha + \beta) + \cos(\alpha - \beta)$ as:

$$F_{n_1,n_2}(\mathbf{x}) = \frac{1}{2}(F^{+}_{n_1,n_2}(\mathbf{x}) + F^{-}_{n_1,n_2}(\mathbf{x})),$$

where

$$\begin{aligned} F^{\pm}_{n_1,n_2}(\mathbf{x}) &= \frac{2}{b_1 - a_1} \frac{2}{b_2 - a_2} \int_{\mathbb{R}} \int_{\mathbb{R}} f(\mathbf{y} \mid \mathbf{x}) \cos\left(n_1\pi \frac{y_1 - a_1}{b_1 - a_1} \pm n_2\pi \frac{y_2 - a_2}{b_2 - a_2}\right) dy_1 dy_2 \\ &= \dots \\ &= \frac{2}{b_1 - a_1} \frac{2}{b_2 - a_2} \mathrm{Re}\left(\psi_{levy}\left(\frac{n_1\pi}{b_1 - a_1}, \pm\frac{n_2\pi}{b_2 - a_2}\right) \exp\left(\mathrm{i}n_1\pi \frac{y_1 - a_1}{b_1 - a_1} \pm \mathrm{i}n_2\pi \frac{y_2 - a_2}{b_2 - a_2}\right)\right). \end{aligned}$$

3.5 Numerical results

In this section, we analyze the performance of the COS method applied to the problem of pricing Energy Quanto options. We aim to compare our results with the benchmark given in the paper [2], where the FFT method was proven to be much more efficient than the Monte Carlo simulation, which is used there as a benchmark. The numerical experiments are performed on a laptop with a processor Intel Core i5, 2.70 GHz and 8 GB RAM.
We define the following generator matrix for the Markov process $M(t)$:

$$Q = \begin{pmatrix} -2 & 2 \\ 2 & -2 \end{pmatrix}.$$

State 1 corresponds to periods of high wind, while State 2 represents low-wind conditions. This modeling choice is motivated by observations from the German energy market, where high wind typically results in increased wind power generation, leading to lower energy prices. From the generator matrix of a continuous-time Markov process, we can find the holding time in each state, which follows an exponential distribution with rate $-Q_{ii}$, that is:

$$T_i \sim \text{exponential}(-Q_{ii}).$$

In this case, the expected time spent in both states is $\mathbb{E}[T_i] = \frac{1}{T_i} = \frac{1}{2}, \text{for } i = 1, 2,$ corresponding to an average duration of half a year in each regime. We assume that the process starts in state 1, setting the initial probability vector as $\mathbf{p} = [1, 0]$.

Case I. Following the analysis in [2], the dynamics of the temperature and electricity futures is governed by the parameter set in Table 3.1. These parameters are based on empirical findings. In particular, [3] fitted an autoregressive AR(1) model with seasonally varying residuals to Stockholm temperature data from January 1, 1961, to May 25, 2006. Regarding the parameters of electric futures, α^E represents the speed of mean reversion and is interpreted through the concept of half-life: the time it takes for the process to decay to half of its long-term mean. It is computed as $\ln(2)/\alpha^E$, which means that in our experiments the half-life time is 14 days for State 1 (where $\alpha^E = \ln(2)/14$) and 10 days for State 2 (where $\alpha^E = \ln(2)/10$). These values of α^E are consistent with the estimates provided in [4]. The volatility parameters σ^E are chosen to reflect the high volatility typically observed in electricity prices.

Table 3.1: Parameter set for temperature and electricity futures

No regime switching				
Temperature	c_1	c_2	c_3	α^c
	4	1.3	0.7	$-\ln(0.8)$
Regime switching				
Electricity	**State 1**		**State 2**	
	α^E	σ^E	α^E	σ^E
	$\ln(2)/14$	0.5	$\ln(2)/10$	1

$r = [0.05, 0.05],\ \tau = 1,\ \rho = -0.5,\ K_1 = K_2 = 50,\ F^c(0,\tau) = F^E(0,\tau) = 50$

Before presenting the option pricing results, we illustrate how the COS method is able to reconstruct the discounted PDF from the discounted characteristic function. This intermediate step provides insight into how the method captures the shape and features of the underlying distribution.

The density is obtained through the same Fourier-cosine expansion technique used for pricing, with the integration interval $[a, b]$ determined based on the cumulants of the log-price. These cumulants, as explained in Section 2.8, are computed numerically using finite differences applied to the logarithm of the characteristic function.

We present both a 3D surface plot (Figure 3.1a) and a contour plot (Figure 3.1b) of the discounted density considering only one maturity. The impact of different maturities is further explored in the option pricing analysis that follows.

Overall, the recovered density appears smooth and well-behaved, confirming that the COS method is effective in recovering the density function of the underlying process.

We proceed to apply the COS method for pricing. We determine the price $V(F_1, F_2, T)$ at two different maturities: $T = 1/2$ and $T = 1/12$, where time is measured in years. To benchmark, we refer to the Monte Carlo confidence intervals reported in [2] as well as the FFT prices. For $T = 1/2$, the FFT reference price is **18.553**, while the Monte Carlo estimate is **18.670**, with a 95% confidence interval of [18.503, 18.836] (standard error = 0.085). We applied the COS method with various values of N, in the range $N \in \{32, 64, 128, 256\}$. The results, along with the computational times, are shown in Table 3.2. We observe that the COS price matches the FFT price up to three digits (the maximum precision available in the reference prices) already at $N = 64$.

Although the computational environments differ (Intel Core i5 2.70 GHz with 8 GB RAM in our experiments vs. Intel Core i7 2.80 GHz with 12 GB RAM in [2]),

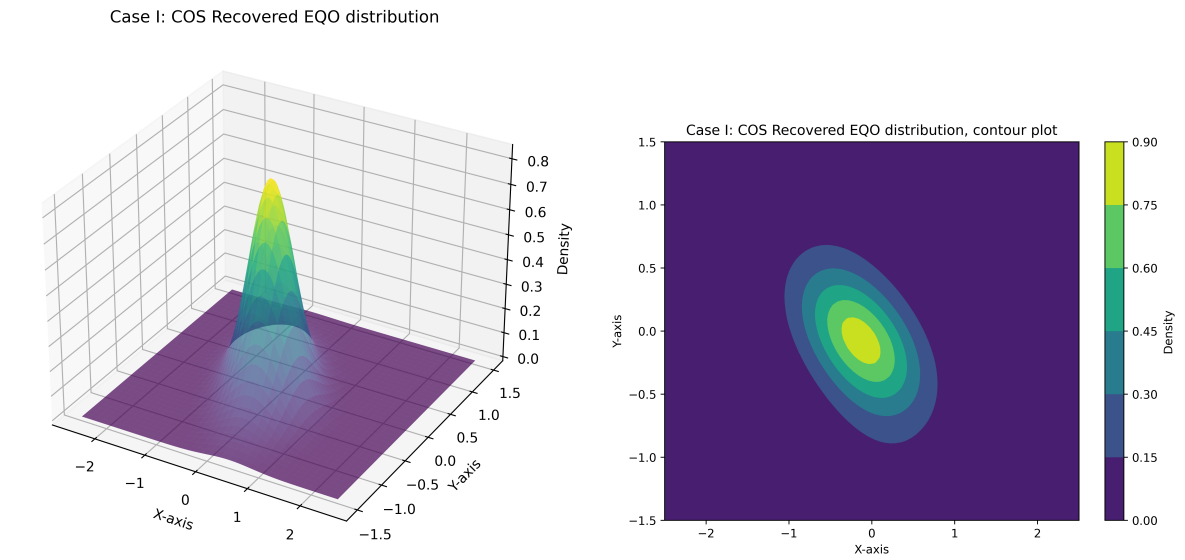

(a) 3D surface plot.

(b) Contour plot.

Figure 3.1: Case I, maturity $T = 1/2$: Discounted density recovery using COS method.

the COS method shows significantly faster computation times in our implementation compared to the benchmarks reported using Monte Carlo and FFT. This suggests that the COS method is highly efficient, though a direct comparison is not fully conclusive due to potential differences in hardware and implementation details. Specifically, our implementation completes with $N = 64$ in approximately 10.85 seconds, whereas the method in [2] requires around 9 minutes (540 seconds), indicating a speed-up of nearly 50 times. While a one-to-one comparison is not entirely possible, it is worth noting that the hardware used in [2] is more powerful, which suggests that the COS method could be even more efficient than this speed-up indicates.

With greater values of N, we observe changes in the digits beyond the third decimal place compared to the case $N = 64$. However, since our reference values are only reported with three-digit accuracy, a full numerical comparison is not possible. Nevertheless, due to the high accuracy and fast convergence of the COS method, we can regard its output, especially at higher N, as a reliable benchmark for option prices. The COS method not only offers computational efficiency but also achieves greater numerical precision than FFT-based approaches, making it particularly suitable for producing reference values in option pricing studies.

Notably, the COS price at $N = 32$ falls outside the Monte Carlo confidence interval, while for $N \geq 64$, the COS prices consistently lie within it.
We include two plots: Figure 3.2a showing how the COS price approaches the Monte Carlo estimate as N increases (including the confidence interval), and Figure 3.2b show-

ing convergence to the FFT reference.

Method	**Price**	**Confidence Interval / Notes**	N	**Time (s)**
Monte Carlo	18.670	[18.503, 18.836]	–	370 min
FFT	18.553	–	–	9 min
COS	24.9491152377	–	32	0 min 2.75 sec
COS	18.5535916338	–	64	0 min 10.85 sec
COS	18.5534245356	–	128	0 min 40.26 sec
COS	18.5534245356	–	256	2 min 45.84 sec

Table 3.2: Energy Quanto option, case I, $T = 1/2$. Comparison of option pricing methods: Monte Carlo, FFT, and COS with varying N.

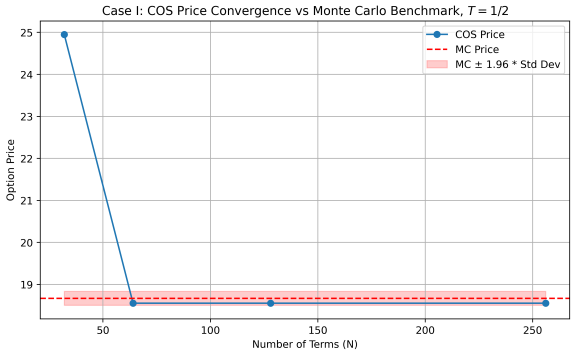

(a) COS price convergence toward the Monte Carlo estimate (with confidence interval).

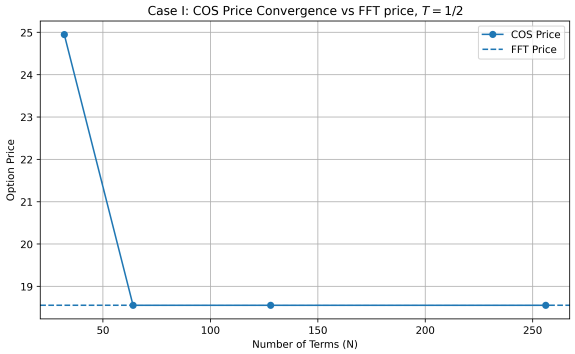

(b) COS price convergence toward the FFT reference.

Figure 3.2: Case I, maturity $T = 1/2$: COS method convergence behavior compared to benchmark prices.

When $T = 1/12$, the FFT reference price is **3.351**, whereas the Monte Carlo estimate is **3.353** with 95% confidence interval [3.332, 3.374] (standard error 0.011). We applied the COS method with values of $N \in \{32, 64, 128, 256\}$, and the results together with the computational times are shown in Table 3.3. Just like before, the COS prices match the FFT reference price already at $N = 64$. Our implementation with $N = 64$ takes 10.23 seconds to perform, whereas the FFT method requires 8 min (480 seconds), meaning that our method is nearly 47 times faster, while always keeping in mind that a one-to-one comparison cannot be done due to the different hardware used. Again, for higher values of N, we observe a change in the digits beyond the third decimal place. As for the Monte Carlo confidence interval, the value we get with the COS method for $N = 32$ does not fall in the interval, but for values of $N \geq 64$, it does.

The COS price convergence to the Monte Carlo estimate as N increases (including the confidence interval) is displayed in Figure 3.3a, while the convergence to the FFT reference is shown in Figure 3.3b.

Method	**Price**	**Confidence Interval / Notes**	N	**Time (s)**
Monte Carlo	3.353	[3.332, 3.374]	–	146 min
FFT	3.351	–	–	8 min
COS	3.4016299454	–	32	0 min 3.79 sec
COS	3.3508289044	–	64	0 min 10.23 sec
COS	3.3508253075	–	128	0 min 45.65 sec
COS	3.3508253075	–	256	2 min 54.61 sec

Table 3.3: Energy Quanto option, case I, $T = 1/12$. Comparison of option pricing methods: Monte Carlo, FFT, and COS with varying N.

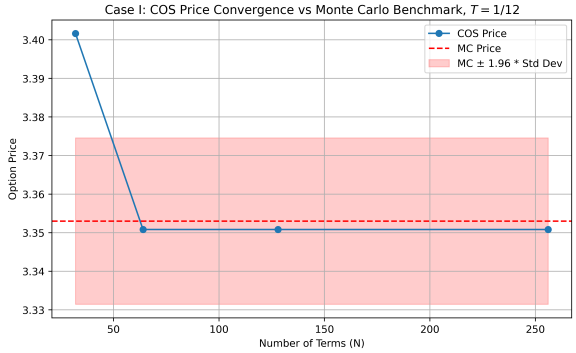

(a) COS price convergence toward the Monte Carlo estimate (with confidence interval).

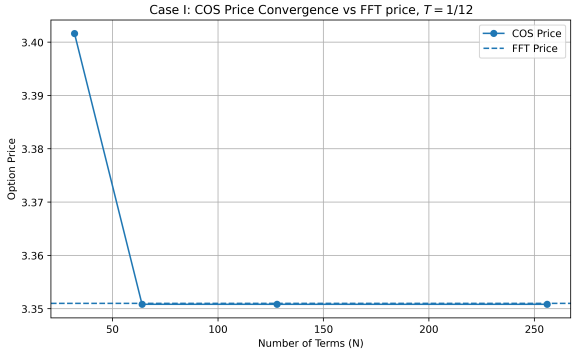

(b) COS price convergence toward the FFT reference.

Figure 3.3: Case I, maturity $T = 1/12$: COS method convergence behavior compared to benchmark prices.

Case II. Now, we carry out an analysis analogous to the one in the previous section, this time focusing on the process described in Section 3.3. The set of parameters describing this model is provided in Table 3.4. In this setting, both futures price processes F_1 and F_2 are Markov-modulated and incorporate jumps.

As done for Case I, before presenting the option pricing results for the second model, we examine how the COS method reconstructs the discounted PDF from the corresponding discounted characteristic function. As with the first case, this step offers a clear view of how well the method captures the characteristics of the underlying process.

The density is computed using the same Fourier-cosine expansion technique applied in pricing, with the integration range $[a, b]$ selected based on the cumulants of the log-price. These cumulants are again obtained numerically via finite differences applied to the logarithm of the characteristic function, as explained in Section 2.8.

We provide both a 3D surface plot (Figure 3.4a) and a contour plot (Figure 3.4b) of the discounted density for a representative maturity. The influence of maturity on the shape of the density will be further examined in the pricing results that follow.

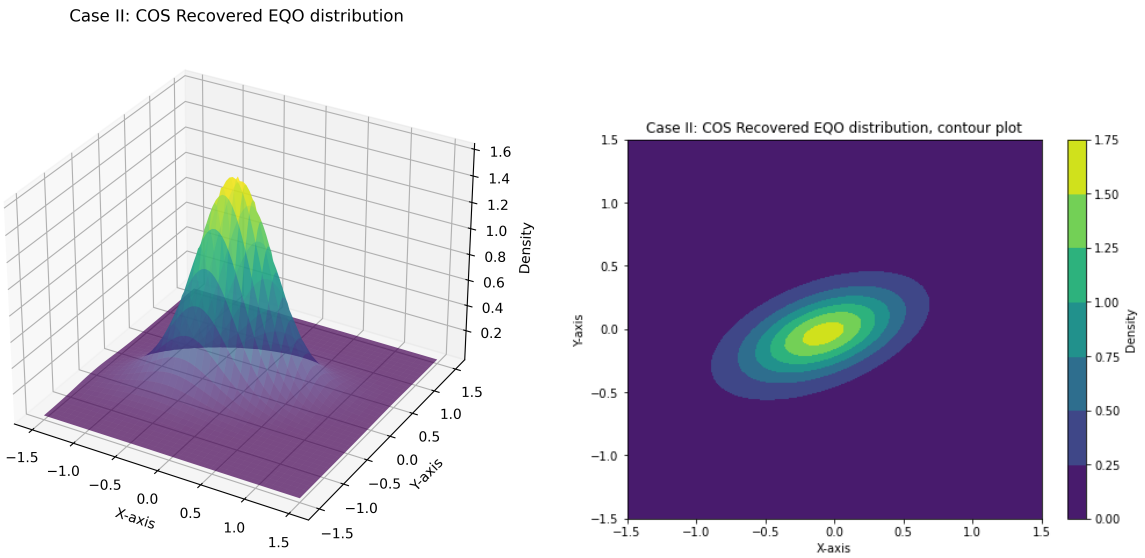

(a) 3D surface plot.

(b) Contour plot.

Figure 3.4: Case II, maturity $T = 1/2$: Discounted density recovery using COS method.

Once again, the recovered density looks smooth and well-structured, indicating that the COS method is even able to effectively capture the more complex dynamics introduced by jumps and the Markov-modulated framework.

Moving on to the pricing stage, we determine the price $V(F_1, F_2, T)$ at two different maturities: $T = 1/2$ and $T = 1/12$, where time is measured in years.

Table 3.4: Parameter set for gas and electricity futures

Electricity	c_1	c_2	η_1^E	η_2^E	α^E	σ^E	β^E
State 1	0.1	0.1	ln(30)	ln(40)	ln(2)/14	0.5	ln(2)/7
State 2	0.1	0.1	ln(100)	ln(150)	ln(2)/10	1	ln(2)/3
Gas	c_1	c_2	η_1^G	η_2^G	α^G	σ^G	β^G
State 1	0.1	0.1	ln(20)	ln(30)	ln(2)/14	0.3	ln(2)/7
State 2	0.1	0.1	ln(80)	ln(120)	ln(2)/10	0.7	ln(2)/3

$r = [0.02, 0.02],\ \tau = 1,\ \rho = 0.5,\ K_1 = K_2 = 50,\ p^E = p^G = 0.5,\ F^G(0,\tau) = F^E(0,\tau) = 50$

For maturity $T = \frac{1}{2}$, the FFT reference price is **138.877**, while the Monte Carlo estimate is **139.100** with a 95% confidence interval [138.330, 139.870] (standard error = 0.393). As in case I, we apply the COS method using different values of $N \in \{32, 64, 128, 256\}$. The corresponding prices and computational times are reported in Table 3.5.

Method	**Price**	**Confidence Interval / Notes**	N	**Time (s)**
Monte Carlo	139.100	[138.330, 139.870]	–	3748 min
FFT	138.877	–	–	13 min
COS	160.4004477460	–	32	0 min 4.77 sec
COS	139.3738556706	–	64	0 min 15.99 sec
COS	138.8770187308	–	128	1 min 4.43 sec
COS	138.8770190583	–	256	3 min 51.83 sec

Table 3.5: Energy Quanto option, case II, $T = 1/2$. Comparison of option pricing methods: Monte Carlo, FFT, and COS with varying N.

In this setting, which involves a slightly more complex process, the COS prices match the FFT reference value up to three digits when $N = 128$, with a computational time of 1 minute and 4 seconds. This compares to 13 minutes for the FFT implementation, yielding a speed-up factor of approximately 13. However, we can observe how the COS price is in the Monte Carlo interval already when $N = 64$.
Compared to the previous model, where the COS method matched the FFT reference price already at $N = 64$, here convergence to the FFT value is slightly slower, requiring $N = 128$. This is likely due to the increased complexity of the underlying dynamics, as both futures price processes in this case are Markov-modulated and include jumps. These features lead to a more complex structure in the characteristic function, which may require a higher number of expansion terms (N) for the COS method to accurately

capture the price dynamics. Despite this, the COS price already falls within the Monte Carlo confidence interval at $N = 64$, indicating that acceptable accuracy is achieved at a relatively low computational cost. In practice, this means that even if full numerical convergence (in terms of matching FFT) is slower, the method still delivers reliable pricing estimates efficiently.
The COS price convergence to the Monte Carlo estimate as N increases (including the confidence interval) is displayed in Figure 3.5a, while the convergence to the FFT reference is shown in Figure 3.5b.

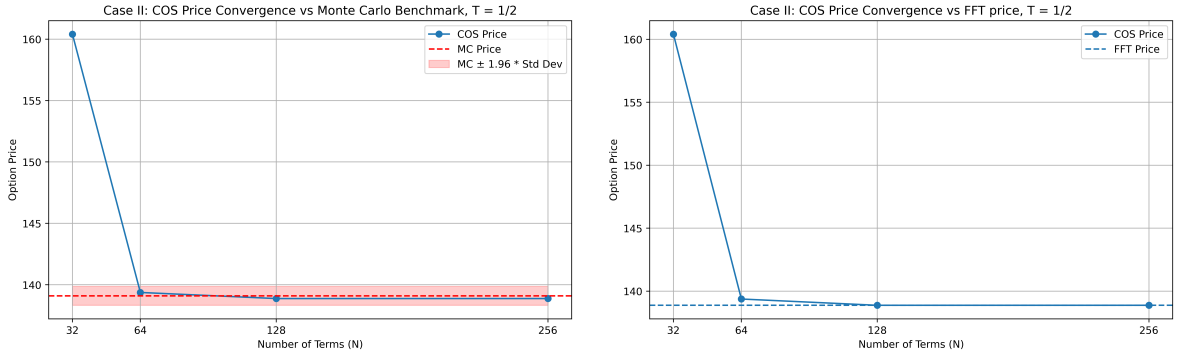

(a) COS price convergence toward the Monte Carlo estimate (with confidence interval).

(b) COS price convergence toward the FFT reference.

Figure 3.5: Case II, maturity $T = 1/2$: COS method convergence behavior compared to benchmark prices.

For maturity $T = \frac{1}{12}$, the FFT reference price is **13.078**, whereas the Monte Carlo estimate is **13.096** with 95% confidence interval $[13.007,\ 13.186]$ (standard error = 0.046). We applied the COS method with values of $N \in \{32, 64, 128, 256\}$, and the results together with the computational times are displayed in Table 3.6.

Method	**Price**	**Confidence Interval / Notes**	N	**Time (s)**
Monte Carlo	13.096	[13.007, 13.186]	–	1552 min
FFT	13.078	–	–	6 min
COS	12.6216808811	–	32	0 min 4.29 sec
COS	13.0688976880	–	64	0 min 15.25 sec
COS	13.0703934311	–	128	1 min 3.87 sec
COS	13.0703934155	–	256	4 min 19.51 sec

Table 3.6: Energy Quanto option, case II, $T = 1/12$. Comparison of option pricing methods: Monte Carlo, FFT, and COS with varying N.

In this case, the COS method approximates the FFT price up to the second digit already at $N = 128$, but even at $N = 256$, the third digit does not match. This behavior differs from the previous cases, where agreement up to three digits was already reached at $N = 64$. The slower convergence here is likely related to the added complexity of the model, including both jumps and regime-switching dynamics, which can lead to slower decay in the Fourier coefficients used in the COS expansion.

Interestingly, the way the COS price evolves as N increases suggests that it does not fully converge to the FFT value. Given the known accuracy and fast convergence of the COS method, this deviation might indicate that the COS result is in fact more reliable. In this sense, we can reasonably treat the COS output at higher N as a trustworthy reference price, even when it slightly differs from the FFT result.

As for the Monte Carlo comparison, the COS price at $N = 32$ falls outside the confidence interval, but already at $N = 64$ it lies well within it. This confirms that even with a low number of expansion points, the COS method delivers a reliable approximation, underlining its efficiency compared to simulation-based approaches.

Figure 3.6a illustrates how the COS price converges to the Monte Carlo estimate as N increases, including the corresponding confidence interval. The convergence towards the FFT reference price is shown in Figure 3.6b.

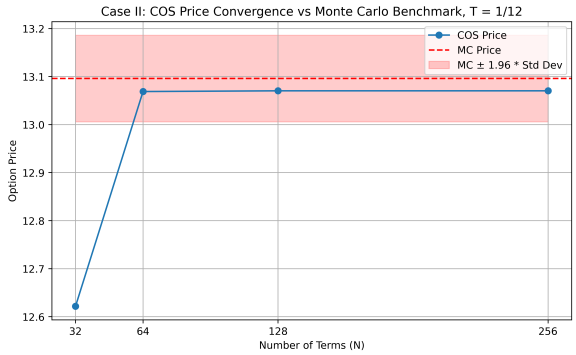

(a) COS price convergence toward the Monte Carlo estimate (with confidence interval).

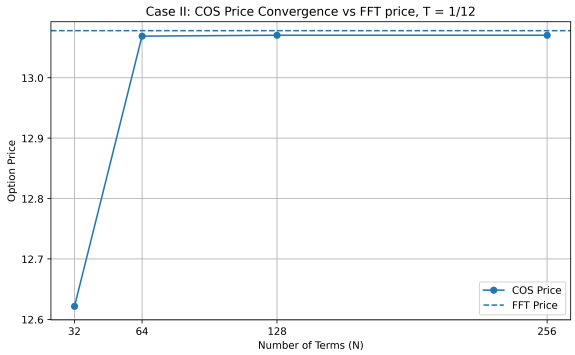

(b) COS price convergence toward the FFT reference.

Figure 3.6: Case II, maturity $T = 1/12$: COS method convergence behavior compared to benchmark prices.

4 Electricity Storage Contracts: Replication of the COS Method in [7]

In this section, we study Electricity Storage contracts and demonstrate how the COS method can efficiently approximate the prices of these options. This research work is a direct continuation of our work in Section 3, where we treated Energy Quanto options as European-style options. The COS method has been applied to price electricity storage contracts in [7]. In this section, we focus on the replication of their method. In the next section, we extend it to regime-switching models.

4.1 Option details

As mentioned earlier, these contracts resemble Bermudan-type options for certain features, namely, they are contracts where electricity can be bought or sold in the electricity market at predefined dates. In particular, the holder of the option has to decide at every time step whether to sell electricity, buy electricity, or do nothing. However, these contracts have additional features that distinguish them from other financial derivatives:

- physical limitations of the electricity storage, like capacity and endurance;
- efficiency of the electricity storage;
- limitations on the quantity of electricity that can be stored or released;
- the payoff can be negative when electricity is stored;
- penalty functions that are activated if certain contract conditions are not fulfilled by the holder.

With this information, we define a storage contract. Consider B exercise dates, where t_0 is the initial time and $\{t_1, \ldots, t_B\}$ is the set of exercise dates such that $0 = t_0 < t_1 < \ldots < t_B = T$ and the time difference between two steps is constant:

$\Delta t = \Delta t_m = t_{m+1} - t_m$. In any of these time steps, the holder of the contract can take an action, but these contracts include an additional date, the **settlement date** t_{B+1}, at which it is not possible to exercise.
On the settlement date, the holder of the options has to pay a penalty if the agreed amount of energy is not present in the storage. This penalty is represented by $q_s(S_{t_{B+1}},\ e(t_{B+1}))$, where $e(t_m)$ is the amount of energy in the storage at time t_m and S_{t_m} is the spot price at time t_m.

As mentioned earlier, the holder of the electricity storage contract can take three actions in each time step, and each of them corresponds to a change in energy level in storage. If the holder stores electricity, then $\Delta e(t_m) = e(t_{m+1}) - e(t_m) > 0$, if electricity is released, then $\Delta e(t_m) < 0$, and if nothing is done, then $\Delta e(t_m) = 0$. At time t_0, the holder cannot take any action, which means that $\Delta e(t_0) = 0$. Based on the action taken, there are different payoffs:

$$g(t_m, S_{t_m}, \Delta e(t_m)) = \begin{cases} -\overline{c}(S_{t_m})\Delta e(t_m), & \text{if } \Delta e(t_m) > 0, \\ 0, & \text{if } \Delta e(t_m) = 0, \\ -\overline{p}(S_{t_m})\Delta e(t_m), & \text{if } \Delta e(t_m) < 0, \end{cases} \tag{4.1}$$

where $\overline{c}(S_{t_m})$ is the cost of storing electricity and $\overline{p}(S_{t_m})$ is the profit of releasing electricity. The contract also includes the **efficiency** of the storage η, so the cost and profit functions are defined as follows: $\overline{c}(S_{t_m}) = \frac{S_{t_m}}{\eta}$, $\overline{p}(S_{t_m}) = S_{t_m}$. This means that in order to sell 1 unit of electrical energy, we need to buy $\frac{1}{\eta} > 1$ units.

As for the **physical limitations** to the capacity of the storage, we define a minimum and maximum level of energy for the storage: $e^{\max}$ and $e^{\min}$, such that $\forall t_m,\ m \in \{0, \ldots, B+1\}$, it holds that $e^{\min} \leq e(t_m) \leq e^{\max}$.

We also take into account the **operational restrictions** on the minimum and maximum energy level changes at an exercise moment, respectively, $i_{\text{op}}^{\min}$ and $i_{\text{op}}^{\max}$. Moreover, we also consider the required minimum energy that can be released, $i_{\text{market}}^{\min}$. This results in the following boundaries for the energy level changes at an exercise moment:

$$\Delta e(t_m) \in [i_{\text{op}}^{\min};\ i_{\text{market}}^{\min}] \cup [0; i_{\text{op}}^{\max}], \quad \forall m \in \{1, \ldots, B\},\ \text{with } i_{\text{op}}^{\min} \leq i_{\text{market}}^{\min} \leq 0 \leq i_{\text{op}}^{\max}.$$

The **endurance** of the battery is also considered, so an interval for energy changes is set in order to preserve the life of the battery, that is $[i_b^{\min}, i_b^{\max}]$. In case the holder decides to take an action that involves a change in the energy level that lies outside this interval, then there is another penalty function denoted by $q_b(\Delta e(t_m))$, which only depends on the action $\Delta e(t_m)$. To consider all these restrictions together, we define two sets of actions, the allowed actions $\mathcal{A}(t_m, e(t_m))$ and the allowed actions that don't involve a penalty $\mathcal{D}(t_m, e(t_m))$. In particular, the set of allowed actions at time $t_m,\ \forall m \in \{1, \ldots, B\}$ is defined as:

$$\mathcal{A}(t_m, e(t_m)) = \{\Delta e \mid e^{\min} \leq e(t_m) + \Delta e \leq e^{\max},\ \text{and } \Delta e \in [i_{\text{op}}^{\min}, i_{\text{market}}^{\min}] \cup [0, i_{\text{op}}^{\max}]\},$$

and the set of allowed actions without penalty at time t_m, $\forall m \in \{1, \dots, B\}$ is defined as:

$$\mathcal{D}(t_m, e(t_m)) = \{\Delta e \mid e^{\min} \leq e(t_m) + \Delta e \leq e^{\max}, \text{ and } \Delta e \in [i_{\mathrm{b}}^{\min}, i_{\mathrm{market}}^{\min}] \cup [0, i_{\mathrm{b}}^{\max}]\}.$$

The value of the electricity storage contract at time t_0 is equal to the discounted future payoffs and penalties, considering that the holder chooses the optimal action at every exercise moment. The pricing formula reads:

$$v(t_0, S_{t_0}) = \max_{\Delta e^*} \mathbb{E}^{\mathbb{Q}}\Big[\sum_{m=1}^{B} e^{-r(t_m - t_0)}(g(t_m, S_{t_m}, \Delta e(t_m)) + q(\Delta e(t_m))) + e^{-r(t_{B+1} - t_0)} q_s(t_{B+1}, S_{t_{B+1}}, e(t_{B+1}))\Big], \tag{4.2}$$

where $\mathbb{Q}$ is the risk-neutral pricing measure, r is the risk free interest rate, and $\Delta e^* = [\Delta e^*(t_1), \dots, \Delta e^*(t_B)]$ denotes the set of optimal actions.

4.1.1 Risk-neutral measure in the electricity market

We recall the 2nd FTAP, as stated in Theorem 2.3.2, which asserts that the completeness of a market is characterized by the uniqueness of the risk-free pricing measure. However, the electricity market is **incomplete** due to its unique characteristics. Despite this, several approaches exist to address such incompleteness. Following the framework proposed in [7], we assume that there exists a risk premium to compensate for risk.

This premium is modeled as $\lambda\sigma(t, X_t)$, where λ represents the market price of risk associated with the state variable X_t, and $\sigma(t, X_t)$ denotes the volatility of the process. This risk premium is deducted from the real-world drift of the electricity price process under the risk-neutral measure.

4.1.2 Electricity price dynamics

We model the electricity price process through a polynomial stochastic process and a polynomial map. Further details on polynomial maps can be found in Appendix A of [7]. We now introduce the Ornstein-Uhlenbeck (OU) process, which is our stochastic polynomial process of choice. The OU process is described by the following stochastic differential equation:

$$dX_t = \kappa(\theta - X_t)dt + \sigma dW_t, \tag{4.3}$$

where $t \geq 0$, κ is the mean reversion rate, θ is the long-run mean, σ is the volatility of the process and W_t is a Brownian Motion (BM) under the real world measure $\mathbb{P}$. This

process is normally distributed, $X_t \sim N(\mathbb{E}(X_t), \mathrm{Var}(X_t))$, with the following conditional mean and variance:

$$\mathbb{E}(X_t|\mathcal{F}_0) = X_{t_0}e^{-\kappa(t-t_0)} + \theta(1 - e^{-\kappa(t-t_0)})$$
$$\mathrm{Var}(X_t|\mathcal{F}_0) = \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa(t-t_0)})$$

with X_{t_0} being the initial value. The characteristic function of the OU process is given by:

$$\phi(u|x, \Delta t) = e^{\mathrm{i}uxe^{-\kappa\Delta t} + A(u,\Delta t)}$$

where

$$A(u, \Delta t) = \frac{1}{4\kappa}(e^{-2\kappa\Delta t} - e^{-\kappa\Delta t})(u^2\sigma^2 + ue^{\kappa\Delta t}(u\sigma^2 - 4\mathrm{i}\kappa\theta)).$$

To model the **electricity spot price**, we consider a stochastic polynomial process X_t, specifically a OU process, and we define the spot price at time t as follows:

$$S_t = \Phi(X_t) = H(X_t)^T\mathbf{p},$$

where $H(X_t)$ is a vector of basis functions for the space of polynomials preserved by the polynomial process (e.g. $H(X_t) = (1, X_t, X_t^2, \ldots, X_t^n)^T$, for a one-dimensional polynomial of degree $n \in \mathbb{N}$), and $\mathbf{p}$ is the vector whose elements are the coefficients of the polynomial map.

4.2 The COS method for electricity storage contracts

In this section, we present the dynamic pricing algorithm used to value electricity storage contracts and describe how the COS method is employed to approximate the continuation values within this framework for each possible energy level $e \in E$.

4.2.1 The dynamic pricing algorithm

The starting point for pricing electricity storage contracts is to discretize the capacity of the electricity storage into N_e equally spaced energy levels. Specifically, the step size between two consecutive levels is defined as $\delta = \frac{e^{\max} - e^{\min}}{N_e}$, resulting in the set of feasible energy levels: $E = \{e^{\min}, e^{\min} + \delta, e^{\min} + 2\delta, \ldots, e^{\max} - \delta, e^{\max}\}$. Moreover, it is

assumed that any action taken by the holder of the contract at a given time t_m changes the energy level by a multiple of δ.

The pricing algorithm proceeds *backward* in time, and at each time step, the contract value must be computed for all energy levels $e \in E$, as the energy level at any time is not known in advance. Accordingly, we denote the value of the contract at time t_m by $v(t_m, S_{t_m}, e(t_m))$, where S_{t_m} is the electricity price and $e(t_m)$ is the energy in storage. As described in Section 4.1, at the settlement date t_{B+1}, no further actions can be taken by the holder, but a penalty $q_s(S_{t_{B+1}}, e(t_{B+1}))$ applies if the storage level does not meet the contract terms. Hence, the value of the contract at settlement date is equal to the penalty function:

$$v(t_{B+1}, S_{t_{B+1}}, e) = q_s(t_{B+1}, S_{t_{B+1}}, e), \quad \forall e \in E. \tag{4.4}$$

Proceeding backward in time, the holder can choose an action $\Delta e \in \mathcal{A}$ at each time t_m, selecting the one that maximizes the contract's value. To determine the value resulting from an action, we first compute the *continuation value*, which captures the expected future value of the contract given the post-action energy level $e(t_{m+1})$.

A key insight is that the continuation value depends only on the resulting energy level $e(t_{m+1}) = e(t_m) + \Delta e$, not on the specific combination of the initial state and the action taken. For example, a level of $e(t_{m+1}) = 5$ can be reached from $e(t_m) = 3$ with an action of $\Delta e = 2$, or from $e(t_m) = 6$ with $\Delta e = -1$. This property allows us to compute the continuation value for each potential energy level $e \in E$ just once per time step. The continuation value at time t_m is determined for all energy levels $e \in E$ using the following formula:

$$c(t_m, S_{t_m}, e) = e^{-r\Delta t} \cdot \mathbb{E}^{\mathbb{Q}}[v(t_{m+1}, S_{t_{m+1}}, e)|\mathcal{F}_{t_m}], \quad \forall e \in E, \tag{4.5}$$

where $\mathcal{F}_t = \sigma(S_s : s \leq t)$ is the filtration generated by the electricity price process. With the continuation values in hand, the *contract value* at time $t_m \in \{t_B, \ldots, t_1\}$ is given by:

$$v(t_m, S_{t_m}, e) = \max_{\Delta e \in \mathcal{A}(t_m, e)} \{g(t_m, S_{t_m}, \Delta e) + c(t_m, S_{t_m}, e + \Delta e) + q_b(\Delta e)\}, \quad \forall e \in E. \tag{4.6}$$

Lastly, since at time t_0 the holder cannot perform any action, the contract value is simply equal to the continuation value:

$$v(t_0, S_{t_0}, e(t_0)) = c(t_0, S_{t_0}, e(t_0)) = e^{-r\Delta t} \cdot \mathbb{E}^{\mathbb{Q}}[v(t_1, S_{t_1}, e(t_0))|\mathcal{F}_{t_0}]. \tag{4.7}$$

Summarizing all the steps, the full *dynamic pricing algorithm* is as follows:

1. $v(t_{B+1}, S_{t_{B+1}}, e) = q_s(t_{B+1}, S_{t_{B+1}}, e), \quad \forall e \in E.$

2. $c(t_m, S_{t_m}, e) = e^{-r\Delta t} \cdot \mathbb{E}^{\mathbb{Q}}[v(t_{m+1}, S_{t_{m+1}}, e)|\mathcal{F}_{t_m}], \quad \forall e \in E, m \in \{B, \dots, 0\},$
3. $v(t_m, S_{t_m}, e) = \max_{\Delta e \in \mathcal{A}} \{g(t_m, S_{t_m}, \Delta e) + c(t_m, S_{t_m}, e + \Delta e) + q_b(\Delta e)\}, \quad \forall e \in E, m \in \{B, \dots, 1\},$
4. $v(t_0, S_{t_0}, e(t_0)) = c(t_0, S_{t_0}, e(t_0)), \quad \forall e \in E.$

4.2.2 The COS method approximation

In this section, we present the COS method as a tool to approximate the continuation values in the dynamic pricing algorithm presented above for each allowed energy level $e \in E$. Recalling the notions of the COS method in Section 2.7, it follows that the continuation value can be approximated $\forall e \in E$ and $\forall m \in \{B+1, \dots, 1\}$ as follows:

$$c(t_{m-1}, x, e) = \hat{c}(t_{m-1}, x, e) = e^{-r\Delta t} \sum_{n=0}^{N-1}{}' \text{Re}\left\{\phi\left(\frac{n\pi}{b-a} \mid \Delta t, x\right) e^{\text{i}n\pi\frac{-a}{b-a}}\right\} V_n(t_m, e), \tag{4.8}$$

where the coefficients $V_n(t_m, e)$ are given by:

$$V_n(t_m, e) = \frac{2}{b-a} \int_a^b v(t_m, y, e) \cos\left(n\pi \frac{y-a}{b-a}\right) dy, \quad \forall e \in E. \tag{4.9}$$

At time t_0, the contract value can be approximated by the continuation value at time t_0, which is computed using the coefficients $V_n(t_1, e(t_0))$.

Let us start by deriving the **coefficients** $V_n(t_m, e)$ for $m \in \{B+1, \dots, 1\}$ and $e \in E$. At the settlement date t_{B+1}, the contract value equals the penalty function (Equation 4.4), which yields:

$$V_n(t_{B+1}, e) = \frac{2}{b-a} \int_a^b q_s(y, e) \cos\left(n\pi \frac{y-a}{b-a}\right) dy, \quad \forall e \in E. \tag{4.10}$$

For earlier time steps, the contract value is defined in equation 4.6, leading to the following general form for the coefficients: $V_n(t_m, e)$, $\forall e \in E$:

$$V_n(t_m, e) = \frac{2}{b-a} \int_a^b \max_{\Delta e \in \mathcal{A}(t_m, e)} \{g(t_m, y, \Delta e) + c(t_m, y, e + \Delta e) + q_b(\Delta e)\} \cos\left(n\pi \frac{y-a}{b-a}\right) dy. \tag{4.11}$$

The above integral involves taking the maximum of $Dim(\mathcal{A}(t_m, e))$ functions. The way this is done in our algorithm is by discretizing the integration domain $[a, b]$ into n_y evenly spaced grid points, at each time step t_m and for each energy level e. These points are

placed at the center of small intervals of width $\Delta y = \frac{b-a}{n_y}$, so that each one corresponds to the midpoint of a cell. We construct a valuation matrix of size $(n_y, Dim(\mathcal{A}(t_m, e)))$ in a single, vectorized step. Each entry (i, j) in this matrix contains the contract value at cell i, assuming that action Δe_j is taken.

Once this matrix is built, we find for every grid point the action that maximizes the total value. This gives us a vector of optimal actions, one for each grid cell.
To define the integration intervals, we compress this vector by detecting where the optimal action changes. These breakpoints divide the grid into segments where the same action is optimal throughout. For each segment, we define its integration bounds based on the outermost cells. The lower bound x_i is taken as the center of the first cell in the segment minus $\Delta y/2$, while the upper bound x_{i+1} is the center of the last cell plus $\Delta y/2$. In this way, we ensure that the entire interval $[a, b]$ is covered by a set of non-overlapping sub-intervals, each associated with a single optimal action.
Once the intervals and the corresponding optimal actions are determined, the coefficients $V_n(t_m, e)$ for all $e \in E$ and for the moment $m \in \{B, \dots, 1\}$ can be written as:

$$\begin{aligned} V_n(t_m, e) = \frac{2}{b-a}\Big[&\int_a^{x_1} (g(t_m, y, \Delta e_0^*) + c(t_m, y, e + \Delta e_0^*) + q_b(\Delta e_0^*)) \cos\left(n\pi \frac{y-a}{b-a}\right) dy \\ &+ \int_{x_1}^{x_2} (g(t_m, y, \Delta e_1^*) + c(t_m, y, e + \Delta e_1^*) + q_b(\Delta e_1^*)) \cos\left(n\pi \frac{y-a}{b-a}\right) dy \\ &+ \dots \\ &+ \int_{x_{\mathscr{N}}}^{b} (g(t_m, y, \Delta e_{\mathscr{N}}^*) + c(t_m, y, e + \Delta e_{\mathscr{N}}^*) + q_b(\Delta e_{\mathscr{N}}^*)) \cos\left(n\pi \frac{y-a}{b-a}\right) dy\Big] \\ = \sum_{i=0}^{\mathscr{N}} &\Big[G_n(x_i, x_{i+1}, \Delta e_i^*) + C_n(x_i, x_{i+1}, e + \Delta e_i^*) + Q_n(x_i, x_{i+1}, \Delta e_i^*)\Big], \quad (4.12) \end{aligned}$$

where $x_0 = a$, $x_{\mathscr{N}+1} = b$, $\Delta e_i^* \in \mathcal{A}(t_m, e)$ are the optimal actions in each subinterval.
G_n are the cosine series coefficients of the payoff function, and are defined $\forall e \in E$ and $\Delta e \in \mathcal{A}(t_m, e)$:

$$G_n(x_i, x_{i+1}, \Delta e) = \frac{2}{b-a} \int_{x_i}^{x_{i+1}} g(t_m, y, \Delta e) \cos\left(n\pi \frac{y-a}{b-a}\right) dy. \quad (4.13)$$

C_n are the cosine series coefficients of the continuation value, and are defined $\forall e \in E$ and $\Delta e \in \mathcal{A}(t_m, e)$:

$$C_n(x_i, x_{i+1}, e) = \frac{2}{b-a} \int_{x_i}^{x_{i+1}} c(t_m, y, e) \cos\left(n\pi \frac{y-a}{b-a}\right) dy. \quad (4.14)$$

Q_n are the cosine series coefficients of the penalty functions, and are defined $\forall e \in E$ and $\Delta e \in \mathcal{A}(t_m, e)$:

$$Q_n(x_i, x_{i+1}, \Delta e) = \frac{2}{b-a} \int_{x_i}^{x_{i+1}} q_b(\Delta e) \cos\left(n\pi \frac{y-a}{b-a}\right) dy. \tag{4.15}$$

The coefficients $G_n(x_i, x_{i+1}, \Delta e)$
The coefficients $G_n(x_i, x_{i+1}, \Delta e)$ are defined in Equation 4.13. The payoff function of electricity storage contracts is described in Equation 4.1. However, since the characteristic function of the price process $S_t = \Phi(X_t)$ is not always known in closed form, while the one of the underlying process X_t typically is, we switch to the transformed state variables: $x = \Phi^{-1}(S_{t_{m-1}})$ and $y = \Phi^{-1}(S_{t_m})$.
Using this transformation, the payoff function can be rewritten as:

$$g(t_m, y, \Delta e) = \begin{cases} -\frac{\Phi(y)}{\eta}\Delta e, & \text{if } \Delta e(t_m) > 0, \\ 0, & \text{if } \Delta e(t_m) = 0, \\ -\Phi(y)\Delta e, & \text{if } \Delta e(t_m) < 0, \end{cases} \tag{4.16}$$

Substituting Equation 4.16 into 4.13, we obtain:

$$G_n(x_i, x_{i+1}, \Delta e) = \begin{cases} \frac{2}{b-a}\int_{x_i}^{x_{i+1}} -\frac{\Phi(y)}{\eta}\Delta e \cos\left(n\pi\frac{y-a}{b-a}\right) dy, & \text{if } \Delta e(t_m) > 0, \\ 0, & \text{if } \Delta e(t_m) = 0, \\ \frac{2}{b-a}\int_{x_i}^{x_{i+1}} -\Phi(y)\Delta e \cos\left(n\pi\frac{y-a}{b-a}\right) dy, & \text{if } \Delta e(t_m) < 0. \end{cases} \tag{4.17}$$

In this way, the coefficients $G_n(x_i, x_{i+1}, \Delta e)$ can be computed for any finite-order polynomial map $\Phi(\cdot)$. Assuming a second-order polynomial map $\Phi(\cdot)$ as in [7]:

$$S_t = \Phi(X_t) = \frac{1-\gamma}{2} X_t^2 + \gamma X_t \tag{4.18}$$

we derive a closed-form expression for the coefficients $G_n(x_i, x_{i+1}, \Delta e)$:

- If $n = 0$:

$$G_n(x_i, x_{i+1}, \Delta e) = \begin{cases} \frac{2}{b-a}\frac{\Delta e}{\eta}\Big[-\frac{\gamma}{2}y^2 - \frac{1-\gamma}{6}y^3\Big]_{x_i}^{x_{i+1}}, & \text{if } \Delta e > 0, \\ 0, & \text{if } \Delta e = 0, \\ \frac{2}{b-a}\Delta e\Big[-\frac{\gamma}{2}y^2 - \frac{1-\gamma}{6}y^3\Big]_{x_i}^{x_{i+1}}, & \text{if } \Delta e < 0. \end{cases} \tag{4.19}$$

- If $n > 0$:

$G_n(x_i, x_{i+1}, \Delta e) =$

$$\begin{cases} \dfrac{2}{b-a}\dfrac{\Delta e}{\eta}\left[\begin{array}{c} \frac{1}{2\pi^3 n^3}(a-b)\left(\sin\left(\frac{\pi n(a-y)}{a-b}\right)(2a^2(\gamma-1)\right. \\ -4ab(\gamma-1)+2b^2(\gamma-1)+\pi^2n^2y(\gamma-\gamma y+2\gamma)) \\ \left.+2\pi n(a-b)(\gamma(y-1)-y)\cos\left(\frac{\pi n(y-a)}{a-b}\right)\right) \end{array}\right]_{x_i}^{x_{i+1}} & \text{if } \Delta e > 0, \\ 0, & \text{if } \Delta e = 0, \\ \dfrac{2}{b-a}\Delta e\left[\begin{array}{c} \frac{1}{2\pi^3 n^3}(a-b)\left(\sin\left(\frac{\pi n(a-y)}{a-b}\right)(2a^2(\gamma-1)\right. \\ -4ab(\gamma-1)+2b^2(\gamma-1)+\pi^2n^2y(\gamma-\gamma y+2\gamma)) \\ \left.+2\pi n(a-b)(\gamma(y-1)-y)\cos\left(\frac{\pi n(y-a)}{a-b}\right)\right) \end{array}\right]_{x_i}^{x_{i+1}} & \text{if } \Delta e < 0. \end{cases} \tag{4.20}$$

The coefficients $Q_n(x_i, x_{i+1}, \Delta e)$
The coefficients $Q_n(x_i, x_{i+1}, \Delta e)$, defined in Equation 4.15, incorporate the penalty function $q_b(\Delta e)$, which only depends on the action Δe and not on the electricity price. Moreover, $q_b(\Delta e)$ is non-null only if $\Delta e \in \mathcal{A}\backslash\mathcal{D}$. Therefore, the coefficients $Q_n(x_i, x_{i+1}, \Delta e)$ have the following expression:

- If $n = 0$:

$$Q_n(x_i, x_{i+1}, \Delta e) = \begin{cases} 0, & \text{if } \Delta e \in \mathcal{D} \\ \frac{2}{b-a}\, q_b(\Delta e)\, (x_{i+1} - x_i), & \text{if } \Delta e \in \mathcal{A}\backslash\mathcal{D} \end{cases} \tag{4.21}$$

- If $n > 0$:

$$Q_n(x_i, x_{i+1}, \Delta e) = \begin{cases} 0, & \text{if } \Delta e \in \mathcal{D} \\ \frac{2}{b-a}\, q_b(\Delta e)\, \left(\sin\left(n\pi\frac{a-x_{i+1}}{a-b}\right) - \sin\left(n\pi\frac{a-x_i}{a-b}\right)\right), & \text{if } \Delta e \in \mathcal{A}\backslash\mathcal{D}. \end{cases} \tag{4.22}$$

To define the coefficients $Q_n(x_i, x_{i+1}, \Delta e)$ at the settlement date t_{B+1}, we just need to replace in Equations 4.21 and 4.22 the function $q_b(\Delta e)$ by the function $q_s(\Delta e)$.

The coefficients $C_n(x_1, x_2, t_m, e)$
The coefficients $C_n(x_1, x_2, t_m, e)$, defined in Equation 4.14, are computed by substituting $c(t_m, y, e)$ with its COS approximation from Equation 4.8:

$$C_n(x_1, x_2, t_m, e) = \frac{2}{b-a}\int_{x_i}^{x_{i+1}} c(t_m, y, e)\, \cos\left(n\pi\frac{y-a}{b-a}\right) dy$$

$$= \frac{2}{b-a}\int_{x_i}^{x_{i+1}} e^{-r\Delta t} {\sum_{l=0}^{N-1}}' \mathrm{Re}\left\{\phi\left(\frac{l\pi}{b-a} \mid \Delta t, y\right) e^{\mathrm{i}l\pi\frac{-a}{b-a}}\right\} \hat{V}_l(t_m, e)\, \cos\left(n\pi\frac{y-a}{b-a}\right) dy$$

Recalling that $\phi(u \mid \Delta t, x) = e^{\mathrm{i}u\beta x}\phi(u \mid \Delta t)$

$$= \frac{2}{b-a}\int_{x_i}^{x_{i+1}} e^{-r\Delta t} {\sum_{l=0}^{N-1}}' \mathrm{Re}\left\{\phi\left(\frac{l\pi}{b-a} \mid \Delta t\right) e^{\mathrm{i}l\pi\frac{\beta y-a}{b-a}}\right\} \hat{V}_l(t_m, e)\, \cos\left(n\pi\frac{y-a}{b-a}\right) dy$$

$$= e^{-r\Delta t} {\sum_{l=0}^{N-1}}' \mathrm{Re}\left\{\phi\left(\frac{l\pi}{b-a} \mid \Delta t\right)\right\} \hat{V}_l(t_m, e) \cdot \frac{2}{b-a}\int_{x_i}^{x_{i+1}} e^{\mathrm{i}l\pi\frac{\beta y-a}{b-a}} \cos\left(n\pi\frac{y-a}{b-a}\right) dy$$

We define

$$\mathcal{M}_{n,l}(x_1, x_2) = \frac{2}{b-a}\int_{x_i}^{x_{i+1}} e^{\mathrm{i}l\pi\frac{\beta y-a}{b-a}} \cos\left(n\pi\frac{y-a}{b-a}\right) dy \tag{4.23}$$

and obtain the following closed-form formula for the coefficients $C_n(x_1, x_2, t_m, e)$:

$$C_n(x_1, x_2, t_m, e) = e^{-r\Delta t} {\sum_{l=0}^{N-1}}' \mathrm{Re}\left\{\phi\left(\frac{l\pi}{b-a}\middle|\Delta t\right) \hat{V}_l(t_m, e)\mathcal{M}_{n,l}(x_1, x_2)\right\}, \tag{4.24}$$

where $\hat{V}_l(t_m, e)$ are given in Equation 4.12 for $m \in \{B, \ldots, 1\}$ and in Equation 4.10 at the settlement date.
The coefficients $\mathcal{M}_{n,l}(x_1, x_2)$ can be rewritten as:

$$\mathcal{M}_{n,l}(x_1, x_2) = -\frac{\mathrm{i}}{\pi}(\mathcal{M}^s_{n,l}(x_1, x_2) + \mathcal{M}^c_{n,l}(x_1, x_2)),$$

where:

$$\mathcal{M}^c_{n,l}(x_1, x_2) = \begin{cases} \frac{(x_2-x_1)\pi\mathrm{i}}{b-a}, & \text{if } n = l = 0, \\ \frac{1}{l\beta+n}\left[e^{\frac{((l\beta+n)x_2-(l+n)a)\pi\mathrm{i}}{b-a}} - e^{\frac{((l\beta+n)x_1-(l+n)a)\pi\mathrm{i}}{b-a}}\right], & \text{else.} \end{cases}$$

$$\mathcal{M}^s_{n,l}(x_1, x_2) = \begin{cases} \frac{(x_2-x_1)\pi\mathrm{i}}{b-a}, & \text{if } n = l = 0, \\ \frac{1}{l\beta-n}\left[e^{\frac{((l\beta-n)x_2-(l-n)a)\pi\mathrm{i}}{b-a}} - e^{\frac{((l\beta-n)x_1-(l-n)a)\pi\mathrm{i}}{b-a}}\right], & \text{else.} \end{cases}$$

4.3 Results

In this section, following the analysis in [7], we present various electricity storage contracts and determine their prices using the COS method. As a benchmark, we use the confidence interval in [7] obtained using the Least Squares Monte Carlo (LSMC) method. These confidence intervals have the following form:

$$\text{Confidence Interval} = \left[\overline{V} - z_{\alpha/2}\left(\frac{\overline{\sigma}}{\sqrt{10}}\right); \overline{V} + z_{\alpha/2}\left(\frac{\overline{\sigma}}{\sqrt{10}}\right)\right],$$

where $\overline{V}$ is the sample mean obtained with ten experiments, $\overline{\sigma}$ is the standard deviation and $z_{\alpha/2} = 1.96$ is the critical value for the 95% confidence interval. These confidence intervals are the result of 10 runs of the LSMC method with 25000 trajectories.
For all the different contracts, we use the same polynomial model, obtained by applying the polynomial map defined in 4.18 to the OU underlying process (4.3). The parameters used in this analysis are reported in Table 4.1. As for the different contracts presented here, there are some features that are common to all of them, and these can be found in Table 4.2, while some are specific to each contract as they have different technical characteristics. In our experiments with the COS method, we use

γ	0.5
κ	0.3
θ	10.1
σ	$\{0.3, 0.6, 0.9, 1.2\}$
X_0	10
S_0	$\Phi(X_0)$

Table 4.1: Electricity storage contracts: parameters of the model

Start date	t_0	0	
Maturity date	T	1	
Number of exercise moments	B	50	
Time step	$\Delta t = \frac{T}{B}$	1/50	
Settlement date	t_{B+1}	$T + \Delta t$	
Minimum energy to release in the market	i^{min}_{market}	-0.1	MWh

Table 4.2: Electricity storage contracts: parameters common to all contracts

$N \in \{50, 75, 100, 125, 150, 200, 250\}$ and $L = 10$ for the integration range.

4.3.1 Contract 1: Standard electricity storage

The first contract considers the rechargeable battery, which is the most commonly used storage system for electricity. There are many types of rechargeable batteries, each with its own features, but it is very common for their capacity to be between 0.25 and 50 MWh, their output between 0.1 and 20 MW and their efficiency is at most 95% ([27], [7]). All the characteristics of this contract can be found in table 4.3.
Table 4.4 and Figure 4.1 show the numerical results of contract 1, including the prices obtained using the COS method for the various values of N and also the computational times required. When we consider low volatility, $\sigma = 0.3$, the prices remain close to zero

Starting energy level	$e(t_0)$	7	MWh
Minimum capacity	e^{min}	0	MWh
Maximum capacity	e^{max}	15	MWh
Minimum energy level change	i_{op}^{min}	-6	MWh
Maximum energy level change	i_{op}^{max}	6	MWh
Minimum energy level change without penalty	i_b^{min}	-4	MWh
Maximum energy level change without penalty	i_b^{max}	4	MWh
Efficiency	η	0.95	
Penalty function for (dis)charging too quickly	$q_b(\Delta e)$ if $\Delta e \in \mathcal{A}\backslash\mathcal{D}$	-3	€
Penalty function at t_{B+1}	$q_s(e)$ if $e < e(t_0)$	-350	€

Table 4.3: Electricity storage contracts: Contract 1 parameters

across all values of N, but they only fall within the confidence interval once $N \geq 100$. For $\sigma = 0.6$, the prices are slightly higher compared to the previous case, and again fall inside the interval for $N \geq 100$. When $\sigma = 0.9$, the prices enter the confidence interval already at $N = 75$. Lastly, for high volatility $\sigma = 1.2$, the COS price is within the interval at $N = 50$, falls outside at $N = 75$, and then stabilizes inside the interval for $N \geq 100$. This suggests that convergence is generally stable from $N = 100$ onward, though some irregular behavior can occur at intermediate values, particularly for higher volatilities.
In all four cases, the prices obtained with the COS method show convergence, and for values of $N \geq 200$ they stabilize. For low volatilities the prices are very low, that is because in order to make a profit with the efficiency $\eta = 95\%$ the electricity prices need to increase by $100\% \cdot (\frac{1}{0.95} - 1)$, which is more likely to happen with higher volatilities.

4.3.2 Contract 2: Highly efficient electricity storage

The second type of contract we present is very similar to the first, with the only difference being that the rechargeable battery is assumed to have an efficiency of 100%. These kind of batteries actually exist but are very costly and for this reason they are not employed for energy storage.
All parameters of this contract, except for the efficiency η, are the same as those listed in Table 4.4.

Table 4.5 displays the prices computed using the COS method, along with the corresponding computational times.
In case of low volatility, $\sigma = 0.3$, the price lies outside the confidence interval when $N = 50$, but falls within from $N = 75$ onward. As for the other volatilities, $\sigma = \{0.6, 0.9, 1.2\}$, the COS prices fall within the confidence interval already at $N = 50$.

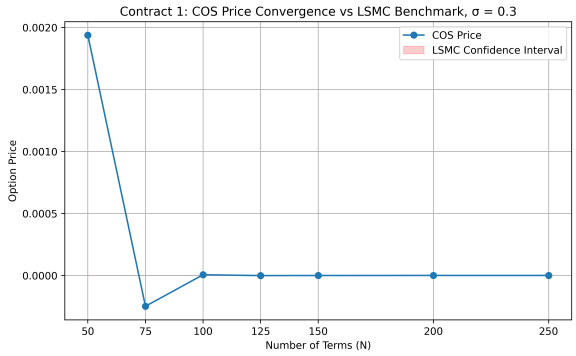

(a) $\sigma = 0.3$

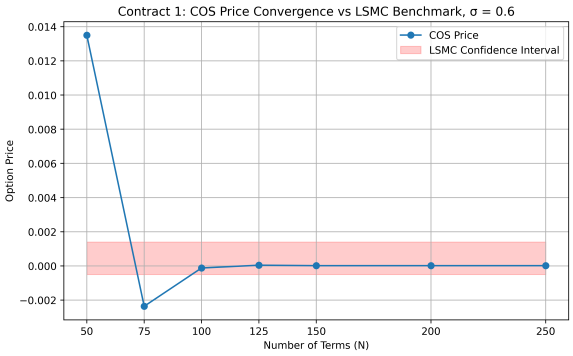

(b) $\sigma = 0.6$

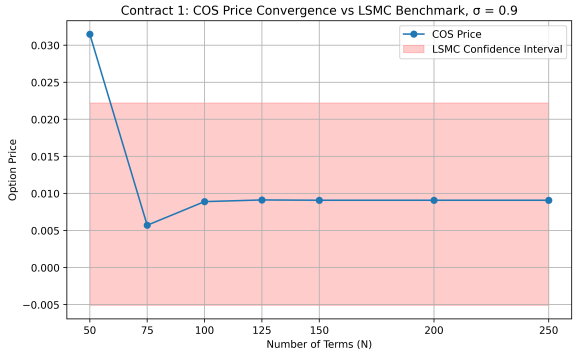

(c) $\sigma = 0.9$

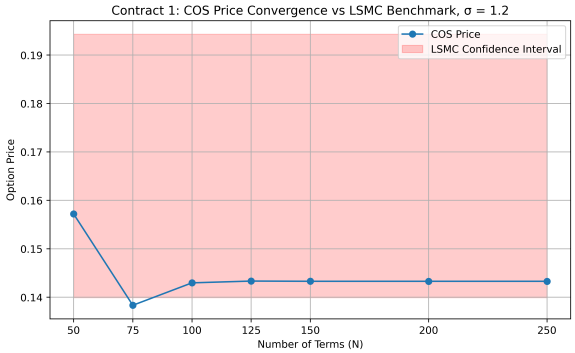

(d) $\sigma = 1.2$

Figure 4.1: Convergence of the COS prices for Contract 1 with different volatility levels

	$\sigma = 0.3$	$\sigma = 0.6$	$\sigma = 0.9$	$\sigma = 1.2$
LSMC 95% C.I.	$[0.0000, 0.0000]$	$[-0.0005, 0.0014]$	$[-0.0051, 0.0222]$	$[0.1399, 0.1943]$
$N = 50$ price	0.00193676	0.01350000	0.03147694	0.15718453
$N = 50$ time	5.15 sec	4.31 sec	5.19 sec	6.54 sec
$N = 75$ price	-0.00024858	-0.00236212	0.00569946	0.13833096
$N = 75$ time	6.27 sec	7.63 sec	8.89 sec	11.53 sec
$N = 100$ price	0.00000556	-0.00012112	0.00888985	0.14296319
$N = 100$ time	8.30 sec	13.75 sec	15.22 sec	14.98 sec
$N = 125$ price	-0.00000114	0.00004170	0.00911301	0.14333134
$N = 125$ time	11.78 sec	18.28 sec	17.47 sec	19.94 sec
$N = 150$ price	-0.00000048	0.00001471	0.00907455	0.14328573
$N = 150$ time	15.58 sec	18.95 sec	23.08 sec	26.65 sec
$N = 200$ price	0.00000000	0.00001487	0.00907439	0.14328401
$N = 200$ time	24.58 sec	30.36 sec	36.62 sec	40.17 sec
$N = 250$ price	0.00000000	0.00001487	0.00907439	0.14328401
$N = 250$ time	50.10 sec	46.02 sec	51.67 sec	58.41 sec

Table 4.4: Electricity storage contracts: COS prices and CPU times for contract 1

In particular, we can observe how for $\sigma = 0.3$ and $\sigma = 1.2$, the prices exhibit an oscillatory pattern at lower values of N: the price is higher at $N = 50$, drops at $N = 75$, and then shows a more stable convergence from $N = 100$. Compared to the first contract, the COS method performs even better here, producing accurate results within the confidence interval with fewer discretization steps, thereby reducing computational time.

4.3.3 Contract 3: Car Park as Power Plant

This third contract proposes car parks as an alternative way to store energy. The analysis in [19] estimates that on average a car is parked 96% of the time, making it suitable for electric vehicles to be used to store electricity during this time. Since a single car may not be always available for this purpose, it is more convenient to consider a greater number of cars together. Indeed, the analysis in [20], estimates that between 92% and 95% of vehicles are available at any instant. The experiments conducted in [7] assumed the car park to contain 150 vehicles with an average capacity of 80 KWh and efficiency of 90%. This results in a total capacity of 12 MWh. Moreover, it is assumed that every vehicle can change up to 25% of its energy level without incurring a penalty, resulting in a total energy change of 3 MWh.

The parameters of contract can be found in Table 4.6, where the penalty function

	$\sigma = 0.3$	$\sigma = 0.6$	$\sigma = 0.9$	$\sigma = 1.2$
LSMC 95% C.I.	$[1.8550, 1.9254]$	$[3.4642, 3.6050]$	$[5.2075, 5.4154]$	$[7.1293, 7.3802]$
$N = 50$ price	1.94087000	3.48988799	5.21513821	7.21464769
$N = 50$ time	4.88 sec	4.49 sec	6.11 sec	8.01 sec
$N = 75$ price	1.89361101	3.49992635	5.22378531	7.13311452
$N = 75$ time	6.73 sec	7.42 sec	11.41 sec	11.81 sec
$N = 100$ price	1.90576522	3.50441877	5.23395406	7.14439669
$N = 100$ time	7.29 sec	11.35 sec	14.91 sec	15.86 sec
$N = 125$ price	1.90645221	3.50349320	5.23399641	7.14486810
$N = 125$ time	10.50 sec	15.28 sec	19.93 sec	21.44 sec
$N = 150$ price	1.90644675	3.50333772	5.23393709	7.14479951
$N = 150$ time	14.48 sec	20.31 sec	24.36 sec	30.40 sec
$N = 200$ price	1.90644748	3.50334804	5.23393509	7.14484789
$N = 200$ time	22.03 sec	31.84 sec	41.39 sec	50.43 sec
$N = 250$ price	1.90644748	3.50334804	5.23393509	7.14484789
$N = 250$ time	33.45 sec	48.00 sec	1 min 2.33 sec	1 min 6.01 sec

Table 4.5: Electricity storage contracts: COS prices and CPU times for contract 2

at settlement date is justified by the fact that the owner of the car must be able to drive at the end of the contract, therefore there has to be a minimum amount of energy in the car battery.

The prices obtained with the COS method and the corresponding computational times are reported in Table 4.7. For the first three volatility levels, $\sigma \in \{0.3, 0.6, 0.9\}$, the COS prices fall outside the LSMC 95% confidence interval when $N = \{50, 75\}$, but they enter the interval from $N = 100$ onward.
In the last case, when $\sigma = 1.2$, the COS price is already within the confidence interval at $N = 75$, and remains stable as N increases.
These results show that the profit from this type of contract remains limited, mainly because the battery efficiency is not high enough to take advantage of the price fluctuations generated by the model. This could change either by assuming a higher efficiency or by adopting a different price dynamics model that better captures features such as cyclical trends or extreme spikes.

Starting energy level	$e(t_0)$	6	MWh
Minimum capacity	e^{min}	0	MWh
Maximum capacity	e^{max}	12	MWh
Minimum energy level change	i_{op}^{min}	-4	MWh
Maximum energy level change	i_{op}^{max}	4	MWh
Minimum energy level change without penalty	i_b^{min}	-3	MWh
Maximum energy level change without penalty	i_b^{max}	3	MWh
Efficiency	η	0.90	
Penalty function for (dis)charging too quickly	$q_b(\Delta e)$ if $\Delta e \in \mathcal{A}\backslash\mathcal{D}$	-10	€
Penalty function at t_{B+1}	$q_s(e)$ if $e < e(t_0)$	-2000	€

Table 4.6: Electricity storage contracts: Contract 3 parameters

	$\sigma = 0.3$	$\sigma = 0.6$	$\sigma = 0.9$	$\sigma = 1.2$
LSMC 95% C.I.	$[0.0000, 0.0000]$	$[-0.0001, 0.0000]$	$[-0.0008, 0.0012]$	$[-0.0044, 0.0020]$
$N = 50$ price	0.00328634	0.00819015	0.00743671	0.00388346
$N = 50$ time	3.38 sec	3.10 sec	5.25 sec	4.97 sec
$N = 75$ price	-0.00044198	-0.00154996	-0.00218187	-0.00317444
$N = 75$ time	4.02 sec	5.59 sec	7.10 sec	8.46 sec
$N = 100$ price	0.00000438	-0.00009245	-0.00016047	0.00026849
$N = 100$ time	6.01 sec	8.00 sec	12.11 sec	11.31 sec
$N = 125$ price	-0.00000245	0.00001853	0.00003542	0.00041088
$N = 125$ time	8.35 sec	15.55 sec	14.92 sec	14.83 sec
$N = 150$ price	-0.00000044	-0.00000001	0.00000084	0.00036958
$N = 150$ time	11.22 sec	13.66 sec	19.45 sec	19.26 sec
$N = 200$ price	0.00000000	0.00000000	0.00000054	0.00036885
$N = 200$ time	14.27 sec	28.78 sec	30.76 sec	29.71 sec
$N = 250$ price	0.00000000	0.00000000	0.00000054	0.00036885
$N = 250$ time	21.86 sec	33.40 sec	43.17 sec	42.85 sec

Table 4.7: Electricity storage contracts: COS prices and CPU times for contract 3

4.3.4 Contract 4: Cost optimization of charging electric vehicles

The fourth type of contract builds on the previous Car Park case. Here, the fleet of electric vehicles is still treated as a form of electricity storage, but the focus is extended to include how the batteries can be charged in a cheap way.
The system's goal is to minimize charging costs by automatically powering the vehicles

during off-peak hours when electricity is cheaper. This approach not only saves money but also relieves pressure on the electricity grid during high demand periods.
The contract requires all vehicles to be fully charged by the settlement date t_{B+1}. A penalty system is in place to enforce this: if a battery is not fully charged on time, an increasing penalty is applied. Furthermore, if the charge falls below a pre-defined critical level (e_{fix}), a much higher, fixed penalty is imposed. Therefore, the penalty function $q_s(e)$ has the following form:

$$q_s(e) = \begin{cases} -1000 \cdot \frac{e^{max}-e}{e^{max}-e_{fix}}, & \text{if } e \geq e_{fix}, \\ -2000, & \text{if } e < e_{fix}. \end{cases} \tag{4.25}$$

As for the other parameters, they remain identical to those of the previous contract and can be found in Table 4.6.

The COS prices and the corresponding computational times are shown in Table 4.8. In the cases where $\sigma \in \{0.3, 0.9\}$, the COS prices fall within the confidence interval already at $N = 50$, and remain inside as N increases, showing a stable convergence.

As for the case of high volatility, $\sigma = 1.2$, $N = 50$ fails to give an approximation that falls inside the interval, but starting from $N = 75$, the prices fall inside and again show stable convergence.

What is different here compared to the other contracts is that when $\sigma = 0.6$, the prices obtained with the COS method remain outside the interval for all the values of N we consider. While the convergence is stable and the approximations get very close to the expected range, the prices never fall inside the interval.

To better understand this issue, we carried out a second analysis using a finer discretization grid for the integration domain $[a, b]$, which is used to determine the optimal actions and their corresponding sub-intervals (Section 4.2.2). In all the experiments conducted above, we used a spacing of $\Delta y = 0.01$, while now we chose $\Delta y = 0.005$. The updated results are reported in Table 4.9, where we can see that the COS prices now fall within the confidence interval already at $N = 50$, and remain inside for larger values of N.

This second analysis was needed because the initial setup, although effective in other volatility scenarios, failed to produce accurate prices when $\sigma = 0.6$, likely due to the discretization being too coarse to capture more subtle features of the value function. By refining the grid, we achieved better accuracy without drastically increasing the computational time. Usually a finer discretization implies a heavier computational load, in practice the increase in time in this second analysis was minor and did not significantly impact the overall performance.

	$\sigma = 0.3$	$\sigma = 0.6$	$\sigma = 0.9$	$\sigma = 1.2$
LSMC 95% C.I.	$[-331.3365, -331.2007]$	$[-330.7876, -330.5472]$	$[-330.3961, -330.0825]$	$[-330.1435, -329.7515]$
$N = 50$ price	-331.24713063	-330.80136362	-330.38054255	-330.15650863
$N = 50$ time	2.19 sec	3.10 sec	3.71 sec	3.91 sec
$N = 75$ price	-331.25488240	-330.81357235	-330.38240915	-330.13710329
$N = 75$ time	3.95 sec	4.96 sec	6.20 sec	7.35 sec
$N = 100$ price	-331.25319110	-330.81399909	-330.38270222	-330.13751554
$N = 100$ time	5.86 sec	6.20 sec	8.01 sec	9.90 sec
$N = 125$ price	-331.25128246	-330.81390311	-330.38261155	-330.13750957
$N = 125$ time	7.57 sec	8.37 sec	11.27 sec	13.53 sec
$N = 150$ price	-331.25127564	-330.81391083	-330.38261126	330.13754087
$N = 150$ time	10.44 sec	12.08 sec	14.21 sec	17.05 sec
$N = 200$ price	-331.25127537	-330.81391164	-330.38261266	-330.13754127
$N = 200$ time	14.05 sec	18.30 sec	21.34 sec	25.82 sec
$N = 250$ price	-331.25127537	-330.81391164	-330.38261266	-330.13754127
$N = 250$ time	24.12 sec	27.51 sec	31.14 sec	36.53 sec

Table 4.8: Electricity storage contracts: COS prices and CPU times for contract 4

	$\sigma = 0.6$
LSMC 95% C.I.	$[-330.7876, -330.5472]$
$N = 50$ price	-330.75709661
$N = 50$ time	3.87 sec
$N = 75$ price	-330.76832012
$N = 75$ time	7.08 sec
$N = 100$ price	-330.76805455
$N = 100$ time	9.97 sec
$N = 125$ price	-330.76793528
$N = 125$ time	12.73 sec
$N = 150$ price	-330.76793902
$N = 150$ time	16.92 sec
$N = 200$ price	-330.76793952
$N = 200$ time	25.16 sec
$N = 250$ price	-330.76793952
$N = 250$ time	36.27 sec

Table 4.9: Electricity storage contracts: COS prices and CPU times for contract 4, $\sigma = 0.6$, finer grid

Overall, the computational times remain low across all the scenarios we analyzed, confirming that the COS method performs efficiently even when pricing options with early-exercise features. In this case, performance is particularly strong given that the method had to account for the specific features of the contract, as well as the discretization of the integration range needed to identify the optimal options and the corresponding sub-intervals (see Section 4.2.2).

5 Electricity Storage Contracts: Applying the COS Method in a Markov-modulated Framework

The drawback of the work done in [7] (Chapter 4) is that the underlying dynamic does not seem to be captured by the assumed stochastic process.

In this final chapter, we select a model that incorporates multiple Markov regimes to describe electricity prices dynamics and adapt the COS pricing algorithm for Electricity Storage contracts in Chapter 4 to account for the Markov-modulated framework.

5.1 Electricity Price model

Electricity spot prices differ fundamentally from those of traditional financial assets and commodities. Their dynamics are shaped not only by economic factors, but also by the physical constraints, such as the fact that electricity cannot be stored on a large scale and the requirement to balance supply and demand in real time. Consequently, electricity prices exhibit several stylized characteristics that must be taken into account in modelling. The following analysis and subsequent model choice are inspired by the work done in [1].

5.1.1 Features

Seasonality is one of the main features shown by electricity prices. It can be observed in different patterns, like intraday, intraweek, or even intrayear. For example, prices are higher in winter rather than in summer because the consumption of electricity is higher in winter for lighting and heating. In general, factors that influence the seasonal behavior of electricity prices are related to climate or business conditions.
This feature can be modeled through different techniques, such as piecewise constant functions ([24], [18]), sinusoidal functions ([24]), or the sum of sinusoidal functions of

different frequencies ([9]).

Spikes are another main feature of electricity prices, and they correspond to rapid changes in electricity spot prices which undergo an intense increase in a very short period of time and then revert to the normal level. The motivation behind the existence of these jumps is that electricity cannot be stored economically. Hence, the fluctuations in demand and supply, which are in turn caused by climate conditions or changes in the prices of fuel, can cause spikes. It was shown in [33] that spikes occur more frequently in electricity prices than for other commodities: for traditional commodities (like crude oil, natural gas and stocks) volatility ranges between 1% and 4%, while for electricity prices volatility goes up to 50%.

Another fundamental feature of these prices is **mean reversion**, meaning that prices will eventually return to their long-term average or mean over time. This characteristic is again more pronounced in this market than in others. In fact, in the interest rate market, the mean reversion rate is weak, in the gas and oil markets, the rate of mean reversion is very slow, so it can take months or years for the prices to revert to their mean level.

5.1.2 Literature Review

In the literature, the spot price S_t or the log-price $\ln S_t$ are typically modelled as the sum of two components:

$$S_t = f(t) + X_t, \quad \ln S_t = f(t) + X_t,$$

where $f(t)$ is a deterministic function of time which represents the seasonal component and $(X_t)_{t\geq 0}$ is a stochastic process. In our experiments we directly deal with the deseasonalized spot price ($X_t = S_t - f(t)$), in order to focus on the stochastic part.
We begin with a short overview of prior models, before turning to the reasoning behind our final choice. At first, electricity prices were modeled accordingly to models for stock prices and interest rates. The most common example is the GBM model, described by the following stochastic differential equation (SDE):

$$dX_t = \mu X_t dt + \sigma X_t dW_t, \tag{5.1}$$

where μ is the drift, σ is the volatility and W_t is a standard BM. The GBM has drawbacks when it comes to modeling electricity prices: it does not properly capture the mean reverting behavior and it does not accurately describe the spikes, neither in size or number.
Some subsequent models focused on taking into account the mean reversion. [30] introduced the Schwartz model, described by the following SDE:

$$dS_t = \theta(\mu - \ln S_t) S_t dt + \sigma S_t dW_t, \tag{5.2}$$

which becomes a OU process (Section 4.1.2) by taking $X_t = \ln S_t$:

$$dX_t = \theta(\alpha - X_t)dt + \sigma dW_t, \tag{5.3}$$

with $\alpha = \mu - \frac{\sigma^2}{2\theta}$. Compared to the GBM model, where any shock has an impact on the price depending on its magnitude, the mean-reverting process returns to its normal level, despite the shocks. Both models are inadequate in describing the spike behavior of the prices. This leads to the introduction of the mean-reversion jump-diffusion (MRJD) models. Jump-diffusion models have the following SDE:

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t + dq(t, X_t), \tag{5.4}$$

where W_t is a standard BM process and $q(t, X_t)$ is a jump process. This form can be adapted to handle the mean-reverting property by considering the drift term in the form: $\mu(t, X_t) = \theta(\mu - X_t)$. MRJD models follow the SDE:

$$dX_t = \theta(\mu - X_t)dt + \sigma X_t dW_t + \mathcal{J}dq_t, \tag{5.5}$$

with θ the speed of mean-reversion, μ the mean-reverting level, σ the volatility, W_t is a standard BM, $\mathcal{J} \sim N(\mu_j, \sigma_j^2)$ is the random jump size and q_t is a Poisson process (Definition 2.5.2) with intensity λ.

This standard model has been widely re-adapted in the literature. [31] distinguished positive and negative jumps with two different Poisson processes with various jump sizes (constant, deterministic and stochastic). A different approach can be found in [17], where the upward or downward jumps are determined using a threshold function: the direction of the spike is positive $(+1)$ if the price X_t is below a certain threshold, and it is negative (-1) if the price X_t is above the threshold. Compared to the processes presented before, MRJD processes perform well in terms of describing the features of electricity prices, such as mean-reversion and spikes. However, while the spikes of electricity prices are sudden, and a jump upward is immediately followed by a jump downward, the prices that follow the MRJD model converge slower back to the mean level. If, on the other hand, we were to consider a higher speed of mean reversion, we would overestimate the parameter in the normal level. To avoid the incorrect specification of parameters, another class of models is considered: the MRS models.

5.1.3 Markov regime-switching models

Let $\mathbb{T}$ be the time index set $[0, T]$, with $T \in (0, \infty)$.
MRS models describe the electricity prices through J different states, and in each of these regimes the electricity price $(X_t)_{t\geq 0}$ is described by different stochastic processes $X_t^i, i \in \{1, \ldots, J\}$. The switch between the regimes is regulated by a Markov chain

$(R_t)_{t\geq 0}$, taking values i in the state space $\mathscr{J} = \{1, 2, \ldots, J\}$, described by the transition matrix $\mathbf{P} = \{p_{ij} : 1 \leq i, j \leq J\} \in \mathbb{R}^{J\times J}$, such that $0 \leq p_{ij} \leq 1$, $\sum_j p_{ij} = 1\ \forall j$, which determines the transition rates from a state to another.
When R_t is in state i, X_t^i is a process with some specific characteristics (like mean and variance), which can vary from state to state.

One of the earliest applications of MRS models to electricity prices is presented in [12], where prices are modeled using a regime-switching mean-reverting process with two types of jumps. A related approach is proposed in [34], where prices follow a regime-switching model incorporating log-normally distributed spikes. In [6], the authors introduce a two-state regime-switching model for deseasonalized average spot prices, with spikes modeled by a Pareto distribution. Further alternatives are explored in [11] and [25], which describe the base regime with a mean-reverting process and the spike regime with a Poisson jump process. Beyond two-regime specifications, [18] proposes a three-regime MRS model, where the base regime captures normal price levels, the spike regime accounts for sudden jumps, and a jump-reversal regime describes how prices revert to their normal level immediately after a spike.

MRS models provide greater flexibility in modeling electricity prices, since each regime is characterized by its own stochastic process, independent of the others. One drawback of this class of models is the high computational cost associated with parameter calibration. For this reason, most applications rely on two-regime specifications, as they reduce computational complexity while still capturing the main features of price dynamics. We also decided to adopt a two-regime specification, following a model proposed in [1]. While it is common practice to model log-prices, the authors argue that it is more methodologically appropriate to model prices directly. One justification is that using log-prices imposes a non-negativity constraint, whereas electricity prices can, in fact, take negative values.

We now introduce our model of choice mathematically. Let $(\Omega, \mathcal{F}_{t\in[0,T]}, \mathbb{P})$ be a filtered probability space with $T < \infty$. We introduce the Markov chain $(R_t)_{t\geq 0}$, which determines the state at time t, with two possible choices: a base regime ($R_t = b$) and a spike regime ($R_t = s$).

The switching mechanism is regulated by a transition probability matrix $\mathbf{P}$, such that:

$$\mathbf{P} = (p_{ij}) = \begin{pmatrix} p_{bb} & 1 - p_{bb} \\ 1 - p_{ss} & p_{ss} \end{pmatrix}. \tag{5.6}$$

We recall the Markov property, according to which the value of the process R_t at time t only depends on the previous value R_{t-1}:

$$\mathbb{P}(R_t = j \mid R_{t-1} = i, R_{t-2} = k, \ldots) = \mathbb{P}(R_t = j \mid R_{t-1} = i). \tag{5.7}$$

We model the stochastic component $(X_t)_{t\geq 0}$ of the electricity price as follows:

$$X_t = \begin{cases} X_{t,b} & \text{if } R_t = b, \\ X_{t,s} & \text{if } R_t = s, \end{cases} \tag{5.8}$$

where $X_{t,b}$ and $X_{t,s}$ are assumed to be mutually independent. The filtration $\mathcal{F}_t$ is generated by both processes $(R_t)_{t\geq 0}$ and $(X_t)_{t\geq 0}$.
The base model $X_{t,b}$ is distributed as a OU process to describe the mean-reverting behavior of the prices.

On the other hand, the spike model $X_{t,s}$ follows a BM with drift:

$$dX_{t,s} = \mu_s dt + \sigma_s dW_{t,s},$$

with $W_{t,s}$ is a standard BM. In this way we have that: $X_{t,s} \sim N(\mu_s t,\ \sigma_s^2 t)$. This choice is convenient for the spike regime because the process admits a closed-form expression for its characteristic function, which will be particularly useful in the numerical experiments employing the COS method.

5.2 COS Method in a Markov-modulated framework

In this section, we start by adapting the COS method to the Markov modulated framework, and then we focus on electricity storage contracts.

5.2.1 Density Recovery

We first focus on the density function of the process X_t, with the aim of reconstructing it using the COS method in this Markov-modulated framework. To determine the marginal density of X_t, we exploit the law of total probability applied to the Markov chain:

$$f(x_t \mid x_0) = \sum_{j=1}^{J} \mathbb{P}(R_t = j) f(x_t \mid R_t = j, x_0),$$

where $f(x_t \mid R_t = j, x_0)$ is the density of the process in a specific state.

To determine $\mathbb{P}(R_t = j)$ we recall that, for Markov chains:

$$\pi_t = \pi_0 \mathbf{P}^t,$$

where π_0 is a row vector of state probabilities at time t_0 and π_t is a row vector of state probabilities at time t. This leads to:

$$f(x_t \mid x_0) = \sum_{j=1}^{J} \pi_t(j) \cdot f(x_t \mid R_t = j, x_0).$$

We recover the individual density functions for each regime from the inverse Fourier integral via cosine expansion. The density function $f_j(x \mid x_0)$ can be obtained as the inverse Fourier transform of the characteristic function of the j-th component of the process X_t:

$$f_j(x \mid x_0) = \frac{1}{2\pi} \int_{\mathbb{R}} e^{-\mathrm{i}ux} \ \varphi^j_{\Delta t}(u \mid x_0) \ du,$$

where i is the imaginary unit such that $\mathrm{i}^2 = -1$.
Following the original procedure to derive the COS approximation (Section 2.7), we find that by truncating the integration range to an interval $[a, b] \in \mathbb{R}$ (such that the density function is negligible in the interval $\mathbb{R}\backslash[a, b]$), the density function can be rewritten as:

$$f_j(x \mid x_0) = \sum_{n=0}^{N-1}{}' F^j_n \cos\left(n\pi \frac{x-a}{b-a}\right), \tag{5.9}$$

with the coefficients

$$F^j_n = \frac{2}{b-a} \mathrm{Re}\left\{ e^{\mathrm{i}n\pi \frac{-a}{b-a}} \varphi^j_{\Delta t}\left(\frac{n\pi}{b-a} \mid x_0\right)\right\}.$$

In our model, we choose the OU process for the base regime and a GBM for the spike regime. For both processes, the conditional characteristic function can be rewritten in the following form:

$$\varphi_{\Delta t}(u \mid x) = e^{\mathrm{i}u\beta x} \varphi_{\Delta t}(u),$$

where $\beta = 1$ for the GBM process and $\beta = e^{-\kappa \Delta t}$ for the OU process.
Hence, we have:

$$F^j_n = \frac{2}{b-a} \mathrm{Re}\left\{ e^{\mathrm{i}n\pi \frac{\beta x_0 - a}{b-a}} \varphi^j_{\Delta t}\left(\frac{n\pi}{b-a}\right)\right\}.$$

5.2.2 Option Pricing

In this section, we adapt the COS pricing formulas to the Markov modulated framework we introduced earlier. Since electricity storage contracts can be exercised at predefined dates, we consider the set of $B+1$ predefined exercise dates $\mathbb{T} = \{t_0, t_1, \ldots, t_B\}$, with $0 = t_0 < t_1 < \ldots < t_B = T$ such that $\Delta t = t_m - t_{m-1}$ is constant.
The risk-neutral formula for the valuation of an option reads:

$$\begin{aligned} v(x, i, t) &= \mathbb{E}^{\mathbb{Q}}\Big[\exp\left(-\int_t^T r(R_s)ds\right) \Phi(X_T) \mid X_t = x, R_t = i\Big] \\ &= \int_{\mathbb{R}} \exp\left(-\int_t^T r(R_s)ds\right) \ \Phi(y) \ f(y \mid x, R_t = i) \ \ dy, \end{aligned}$$

where $\Phi(\cdot)$ is the payoff function of the financial derivative and $f(\cdot \mid x, R_t = i)$ is the transition density function conditioned on the fact that we are in state i at time t.

From this risk-neutral valuation formula, we obtain the following formula for the continuation value of the option at time t_{m-1} in the Markov state i:

$$\begin{aligned}
&c(X_{t_{m-1}} = x, I_{m-1} = i, t_{m-1}) = \\
&= e^{-r\Delta t}\ \mathbb{E}^{\mathbb{Q}}[\hat{v}(X_{t_m}, I_{t_m}, t_m) \mid X_{t_{m-1}} = x, I_{t_{m-1}} = i] \\
&= e^{-r\Delta t}\ \sum_{j=1}^{J} \mathbb{P}(I_{t_m} = j \mid I_{t_{m-1}} = i)\ \mathbb{E}^{\mathbb{Q}}[v(X_{t_m}, t_m) \mid X_{t_{m-1}} = x, I_{t_m} = j, I_{t_{m-1}} = i] \\
&= e^{-r\Delta t}\ \sum_{j=1}^{J} p_{ij}\ \mathbb{E}^{\mathbb{Q}}[v(X_{t_m}, t_m) \mid X_{t_{m-1}} = x, I_{t_m} = j, I_{t_{m-1}} = i] \\
&= e^{-r\Delta t}\ \sum_{j=1}^{J} p_{ij} \int_{\mathbb{R}} v(y, j, t_m)\ f_j(y \mid x) dy
\end{aligned} \tag{5.10}$$

Following the original procedure to derive the COS approximation (Section 2.7), we have that by truncating the integration range to an interval $[a, b] \in \mathbb{R}$ (such that the density function is negligible in the interval $\mathbb{R}\backslash[a, b]$), equation 5.10 can be rewritten as:

$$c_1(x, i, t_{m-1}) = e^{-r\Delta t}\ \sum_{j=1}^{J} p_{ij} \int_a^b v(y, j, t_m)\ f_j(y \mid x) dy. \tag{5.11}$$

Taking the Fourier cosine series expansion, we have:

$$c_1(x, i, t_{m-1}) = e^{-r\Delta t}\ \sum_{j=1}^{J} p_{ij} \int_a^b {\sum_{n=0}^{\infty}}' A_n^j(x)\ \cos\left(n\pi \frac{y-a}{b-a}\right) v(y, j, t_m)\ dy,$$

with

$$A_n^j(x) = \frac{2}{b-a} \int_a^b \cos\left(n\pi \frac{y-a}{b-a}\right) f_j(y \mid x) dy \tag{5.12}$$

being the series coefficient of the Fourier cosine expansion of the density function.

We can truncate the infinite series summation and exchange integration and summation:

$$c_2(x, i, t_{m-1}) = e^{-r\Delta t}\ \sum_{j=1}^{J} p_{ij} \cdot \frac{b-a}{2} {\sum_{n=0}^{N-1}}' A_n^j(x) V_n^j(t_m), \tag{5.13}$$

where

$$V_n^j(t_m) = \frac{2}{b-a} \int_a^b \cos\left(n\pi \frac{y-a}{b-a}\right) v(y, j, t_m) dy.$$

Furthermore, the coefficients $A_n^j(x)$ in equation 5.12 can be approximated by the following:

$$\begin{aligned} D_n^j(x) &= \frac{2}{b-a}\mathrm{Re}\left\{e^{\mathrm{i}n\pi\frac{-a}{b-a}}\int_{\mathbb{R}} e^{\mathrm{i}n\pi\frac{y}{b-a}} f_j(y \mid x)dy\right\} \\ &= \frac{2}{b-a}\mathrm{Re}\left\{e^{\mathrm{i}n\pi\frac{-a}{b-a}}\varphi_{\Delta t}^j\left(\frac{n\pi}{b-a} \mid x\right)\right\}. \end{aligned}$$

Replacing $A_n^j(x)$ by $D_n^j(x)$ in equation 5.13, we get:

$$c_3(x,i,t_{m-1}) = e^{-r\Delta t}\sum_{j=1}^{J} p_{ij} {\sum_{n=0}^{N-1}}' \mathrm{Re}\left\{e^{\mathrm{i}n\pi\frac{-a}{b-a}}\varphi_{\Delta t}^j\left(\frac{n\pi}{b-a} \mid x\right)\right\} V_n^j(t_m). \tag{5.14}$$

Since the characteristic function for affine processes can be rewritten as: $\varphi(u \mid x) = \varphi(u)e^{\mathrm{i}u\beta x}$, we can rewrite equation 5.14:

$$c_4(x,i,t_{m-1}) = e^{-r\Delta t}\sum_{j=1}^{J} p_{ij} {\sum_{n=0}^{N-1}}' \mathrm{Re}\left\{e^{\mathrm{i}n\pi\frac{\beta x-a}{b-a}}\varphi_{\Delta t}^j\left(\frac{n\pi}{b-a}\right)\right\} V_n^j(t_m). \tag{5.15}$$

5.2.3 Dynamic Pricing Algorithm

In this section, we adapt the dynamic pricing algorithm used to value electricity storage contracts (Section 4.2.1) to this Markov-modulated framework. Specifically, we adapt the COS method used to approximate the continuation values for each possible energy level $e \in E$ in each Markov state $i \in \mathscr{J}$.

We recall the dynamic pricing algorithm in a non-Markov framework:

1. $v(t_{B+1}, S_{t_{B+1}}, e) = q_s(t_{B+1}, S_{t_{B+1}}, e), \quad \forall e \in E,$
2. $c(t_m, S_{t_m}, e) = e^{-r\Delta t} \cdot \mathbb{E}^{\mathbb{Q}}[v(t_{m+1}, S_{t_{m+1}}, e) \mid \mathcal{F}_{t_m}], \quad \forall e \in E, m \in \{B, \ldots, 0\},$
3. $v(t_m, S_{t_m}, e) = \max_{\Delta e \in \mathcal{A}} \{g(t_m, S_{t_m}, \Delta e) + c(t_m, S_{t_m}, e + \Delta e) + q_b(\Delta e)\}, \quad \forall e \in E, m \in \{B, \ldots, 1\},$
4. $v(t_0, S_{t_0}, e(t_0)) = c(t_0, S_{t_0}, e(t_0)), \quad \forall e \in E.$

To approximate the second equation of the algorithm, the COS approximation formula for the continuation value in Equation 5.15 can be adapted to consider the energy level $e \in E$, as follows:

$$c_4(x,i,t_{m-1},e) = e^{-r\Delta t}\sum_{j=1}^{J} p_{ij} {\sum_{n=0}^{N-1}}' \mathrm{Re}\left\{e^{\mathrm{i}n\pi\frac{\beta x-a}{b-a}}\varphi_{\Delta t}^j\left(\frac{n\pi}{b-a}\right)\right\} V_n^j(t_m, e). \tag{5.16}$$

So the dependence on the energy level is found in the coefficients $V_n^j(t_m, e)$:

$$V_n^i(t_m, e) = \frac{2}{b-a} \int_a^b v(t_m, y, i, e) \cos\left(n\pi \frac{y-a}{b-a}\right) dy, \quad \forall e \in E, \forall i \in \mathscr{J}. \tag{5.17}$$

Let us start by deriving the **coefficients** $V_n^i(t_m, e)$ for $m \in \{B+1, \ldots, 1\}$, $e \in E$ and $i \in \mathscr{J}$.
At the settlement date t_{B+1}, the contract value is equal to the penalty function (Equation 4.4), which yields:

$$V_n^i(t_{B+1}, e) = \frac{2}{b-a} \int_a^b q_s(y, e) \cos\left(n\pi \frac{y-a}{b-a}\right) dy, \quad \forall e \in E. \tag{5.18}$$

The value at the settlement date is the same in every Markov state, because the penalty function $q_s(\cdot)$ only depends on the energy level and not on the price, which is the state-dependent component. For earlier time steps, the contract value is defined in equation 4.6, leading to the following general form for the coefficients: $V_n^i(t_m, e)$, $\forall e \in E$:

$$V_n^i(t_m, e) = \frac{2}{b-a} \int_a^b \max_{\Delta e \in \mathcal{A}(t_m, e)} \{g^i(t_m, y, \Delta e) + c^i(t_m, y, e+\Delta e) + q_b(\Delta e)\} \cos\left(n\pi \frac{y-a}{b-a}\right) dy. \tag{5.19}$$

The methodology used to solve this integral is explained in section 4.2.2. What changes now compared to the case where the framework is not Markov-modulated is that this has to be done for every Markov state, since the optimal actions and the corresponding sub-intervals could be different based on which Markov state the process is in.

Once the intervals and the corresponding optimal actions are determined, the coefficients $V_n^i(t_m, e)$ for all $e \in E$, for the moment t_m, $m \in \{B, \ldots, 1\}$ in the state $i \in \mathscr{J}$ can be written as:

$$\begin{aligned} V_n^i(t_m, e) = \frac{2}{b-a} \Big[& \int_a^{x_1^i} (g^i(t_m, y, \Delta e_{0^i}^*) + c^i(t_m, y, e + \Delta e_{0^i}^*) + q_b(\Delta e_{0^i}^*)) \cos\left(n\pi \frac{y-a}{b-a}\right) dy \\ & + \int_{x_1^i}^{x_2^i} (g^i(t_m, y, \Delta e_{1^i}^*) + c^i(t_m, y, e + \Delta e_{1^i}^*) + q_b(\Delta e_{1^i}^*)) \cos\left(n\pi \frac{y-a}{b-a}\right) dy \\ & + \ldots \\ & + \int_{x_{\mathscr{N}}^i}^{b} (g^i(t_m, y, \Delta e_{\mathscr{N}^i}^*) + c^i(t_m, y, e + \Delta e_{\mathscr{N}^i}^*) + q_b(\Delta e_{\mathscr{N}^i}^*)) \cos\left(n\pi \frac{y-a}{b-a}\right) dy \Big] \\ = & \sum_{l=0}^{\mathscr{N}^i} \Big[G_n^i(x_{l^i}, x_{l+1^i}, \Delta e_{l^i}^*) + C_n^i(x_{l^i}, x_{l+1^i}, e + \Delta e_{l^i}^*) + Q_n^i(x_{l^i}, x_{l+1^i}, \Delta e_{l^i}^*) \Big], \end{aligned} \tag{5.20}$$

where $x_0 = a$, $x_{\mathcal{M}+1} = b$ and $\Delta e^*_{li} \in \mathcal{A}(t_m, e)$ are the optimal actions in the l-th sub-interval in the Markov state i.

The coefficients $Q_n(x_l, x_{l+1}, \Delta e)$
The coefficients Q_n are the cosine series coefficients of the penalty functions, and are defined in Equations 4.21 and 4.22. There is no difference compared to the non-Markov case because the penalty functions are state-independent.

The coefficients $G_n(x_l, x_{l+1}, \Delta e)$
The coefficients $G_n(x_l, x_{l+1}, \Delta e)$ are defined in Equation 4.13. In Section 4.2.2, the electricity price is modeled using a polynomial map. As a result, the characteristic function of the price process is not always available in closed form. To address this limitation, we employ transformed state variables.

In the current setting, instead, we have a closed-form expression for the characteristic function, so the payoff function is:

$$g(t_m, y, \Delta e) = \begin{cases} -\frac{y}{\eta}\Delta e, & \text{if } \Delta e(t_m) > 0, \\ 0, & \text{if } \Delta e(t_m) = 0, \\ -y\ \Delta e, & \text{if } \Delta e(t_m) < 0, \end{cases} \tag{5.21}$$

Substituting Equation 5.21 into 4.13, we obtain:

$$G_n(x_l, x_{l+1}, \Delta e) = \begin{cases} \frac{2}{b-a}\int_{x_l}^{x_{l+1}} -\frac{y}{\eta}\Delta e\, \cos\left(n\pi\frac{y-a}{b-a}\right) dy, & \text{if } \Delta e(t_m) > 0, \\ 0, & \text{if } \Delta e(t_m) = 0, \\ \frac{2}{b-a}\int_{x_l}^{x_{l+1}} -y\ \Delta e\, \cos\left(n\pi\frac{y-a}{b-a}\right) dy, & \text{if } \Delta e(t_m) < 0. \end{cases} \tag{5.22}$$

We derive a closed-form expression for the coefficients $G_n(x_l, x_{l+1}, \Delta e)$:

- If $n = 0$:

$$G_n(x_l, x_{l+1}, \Delta e) = \begin{cases} \frac{2}{b-a}\frac{\Delta e}{\eta}\Big[-\frac{y^2}{2}\Big]_{x_l}^{x_{l+1}}, & \text{if } \Delta e > 0, \\ 0, & \text{if } \Delta e = 0, \\ \frac{2}{b-a}\Delta e\Big[-\frac{y^2}{2}\Big]_{x_l}^{x_{l+1}}, & \text{if } \Delta e < 0. \end{cases} \tag{5.23}$$

- If $n > 0$:

$$G_n(x_l, x_{l+1}, \Delta e) =$$

$$\begin{cases} -\frac{2}{n\pi}\frac{\Delta e}{\eta}\left[y\,\sin\left(n\pi\frac{y-a}{b-a}\right)+\frac{b-a}{n\pi}\,\cos\left(n\pi\frac{y-a}{b-a}\right)\right]_{x_l}^{x_{l+1}} & \text{if } \Delta e > 0, \\ 0, & \text{if } \Delta e = 0, \\ -\frac{2}{n\pi}\Delta e\left[y\,\sin\left(n\pi\frac{y-a}{b-a}\right)+\frac{b-a}{n\pi}\,\cos\left(n\pi\frac{y-a}{b-a}\right)\right]_{x_l}^{x_{l+1}} & \text{if } \Delta e < 0. \end{cases} \tag{5.24}$$

The coefficients $C_n^i(x_1, x_2, t_m, e)$
The coefficients $C_n^i(x_1, x_2, t_m, e)$, defined in Equation 4.14, are computed by substituting $c(t_m, y, i, e)$ with its COS approximation from Equation 5.16:

$$\begin{aligned} &C_n^i(x_1, x_2, t_m, e) \\ &= \frac{2}{b-a}\int_{x_1}^{x_2} c(t_m, y, i, e)\,\cos\left(n\pi\frac{y-a}{b-a}\right)dy \\ &= \frac{2}{b-a}\int_{x_1}^{x_2} e^{-r\Delta t}\,\sum_{j=1}^{J} p_{ij}\,{\sum_{l=0}^{N-1}}'\mathrm{Re}\left\{e^{\mathrm{i}l\pi\frac{\beta y-a}{b-a}}\varphi_{\Delta t}^j\left(\frac{l\pi}{b-a}\right)\right\}V_l^j(t_{m+1}, e)\,\cos\left(n\pi\frac{y-a}{b-a}\right)dy \\ &= e^{-r\Delta t}\,\sum_{j=1}^{J} p_{ij}\,{\sum_{l=0}^{N-1}}'\mathrm{Re}\left\{\varphi_{\Delta t}^j\left(\frac{l\pi}{b-a}\right)\right\}V_l^j(t_{m+1}, e)\cdot\frac{2}{b-a}\int_{x_1}^{x_2} e^{\mathrm{i}l\pi\frac{\beta y-a}{b-a}}\cos\left(n\pi\frac{y-a}{b-a}\right)dy \end{aligned}$$

We define $\mathcal{M}_{n,l}(x_1, x_2)$ as in equation 4.23 and obtain the following closed-form formula for the coefficients $C_n(x_1, x_2, t_m, e)$:

$$C_n^i(x_1, x_2, t_m, e) = e^{-r\Delta t}\,\sum_{j=1}^{J} p_{ij}\,{\sum_{l=0}^{N-1}}'\mathrm{Re}\left\{\varphi_{\Delta t}^j\left(\frac{l\pi}{b-a}\right)V_l^j(t_{m+1}, e)\mathcal{M}_{n,l}(x_1, x_2)\right\}, \tag{5.25}$$

where $\hat{V}_l(t_m, e)$ are given in Equation 5.20 for $m \in \{B, \ldots, 1\}$ and in Equation 5.18 at the settlement date.

5.3 Error Analysis

In this section, we analyze the error of the COS method applied to option pricing. The main results can be found in [14], where this analysis was conducted for the pricing of European options, and in [15], where it was extended to Bermudan options. Our work focuses on the extension of the COS method for pricing electricity storage contracts to

a framework where the underlying prices are Markov-modulated. We start by not considering the additional features that characterize the pricing of these special contracts, as these involve additional features, and hence further approximations. We focus on the case of Bermudan options, assuming that the prices are Markov-modulated, with the purpose of highlighting the effect on the error of the introduction of multiple Markov states.
As stated in the original article ([14]), there are three main steps in the COS method derivation that introduce sources of error:

- The truncation of the integration range:

$$\epsilon_1 = v(x,t_0) - v_1(x,t_0) = \int_{\mathbb{R}\setminus[a,b]} v(y,T) f(y \mid x) dy; \tag{5.26}$$

- The truncation of the series summation on the interval $[a,b]$:

$$\epsilon_2 = v_1(x,t_0) - v_2(x,t_0) = \frac{b-a}{2} e^{-r\Delta t} \sum_{n=N}^{+\infty} A_n(x) \cdot V_n; \tag{5.27}$$

- The approximation of the coefficients $A_n(x)$ by $F_n(x)$:

$$\epsilon_3 = v_2(x,t_0) - v_3(x,t_0) = e^{-r\Delta t} {\sum_{n=0}^{N-1}}' \mathrm{Re}\left\{ \int_{\mathbb{R}\setminus[a,b]} e^{\mathrm{i}n\pi \frac{y-a}{b-a}} f(y \mid x) dy \right\} V_n. \tag{5.28}$$

As for the first source of error, ϵ_1, in Equation 5.26, it is clear that the larger we choose the truncation range, the smaller the error is, because we neglect a smaller portion of the density function.

As for the second source of error, ϵ_2 in Equation 5.27, it is proven in [14] that it converges exponentially for density functions in the class $\mathbb{C}^\infty([a,b])$ and algebraically for density functions that are continuous but differentiable up to a certain order. Exponential convergence means that:

$$|\epsilon_2| < P \cdot \exp(-(N-1)\nu),$$

where $\nu > 0$ is a constant and P varies less than exponentially with N. On the other hand, algebraic convergence reads:

$$|\epsilon_2| < \frac{\overline{P}}{(N-1)^{\gamma-1}},$$

where $\overline{P}$ is a constant and $\gamma \geq \hat{n} \geq 1$, with $\hat{n}$ the algebraic index of convergence of the series coefficients V_n.

Lastly, [14] proves that ϵ_3 in Equation 5.28 can be bounded by: $|\epsilon_3| < |\epsilon_1| + Q|\epsilon_4|$, where Q is a constant independent of N and $\epsilon_4 = \int_{\mathbb{R}\setminus[a,b]} f(y \mid x)dy$, which again depends only on the choice of $[a, b]$.

Combining all these results, the error can be bounded as follows:

$$|\epsilon| < \begin{cases} 2|\epsilon_1| + Q|\epsilon_4| + P \cdot \exp(-(N-1)\nu) & \text{if the density is in the class } \mathbb{C}^\infty(\mathbb{R}) \\ 2|\epsilon_1| + Q|\epsilon_4| + \frac{\overline{P}}{(N-1)^{\gamma-1}} & \text{if the density is continuous but} \\ & \text{differentiable up to a finite order.} \end{cases} \tag{5.29}$$

The framework we introduce, see Section 5.2.1, considers a density function of the form:

$$f(x_t \mid x_0) = \sum_{j=1}^{J} \mathbb{P}(R_t = j) \cdot f(x_t \mid R_t = j, x_0) = \sum_{j=1}^{J} \pi_t(j) \cdot f(x_t \mid R_t = j, x_0),$$

where $f(x_t \mid R_t = j, x_0)$ is the density of the process in the specific state $j \in \mathscr{J}$ and π_t is a row vector of state probabilities at time t.

In our model of choice, see Section 5.1.3, the price follows a OU distribution in the base regime and a normal distribution in the spike regime. Both processes are Gaussian. The Gaussian density function is an infinitely differentiable function ($\mathbb{C}^\infty$) on $\mathbb{R}$. This follows because both the exponential function and all polynomial functions are $\mathbb{C}^\infty(\mathbb{R})$, and compositions and products of $\mathbb{C}^\infty$ functions remain $\mathbb{C}^\infty$. Furthermore, the set of $\mathbb{C}^\infty$ on $\mathbb{R}$ is closed under addition and scalar multiplication. Therefore, any finite linear combination of Gaussian densities (or, more generally, $\mathbb{C}^\infty(\mathbb{R})$ functions) is itself $\mathbb{C}^\infty(\mathbb{R})$ (see [28]). The density function of our price X_t is a linear combination (with the weights given by the probabilities $\pi_t(j)$) of $C^\infty(\mathbb{R})$ density functions in each state $j \in \mathscr{J}$, so we can expect the convergence with this chosen model to be exponential.

Another important factor to analyze is the propagation of the error in the backward induction needed to recursively recover the coefficients $V_n(t_1)$. The following analysis is taken from [15]. We assume that the coefficients $V_n(t_{m+2})$ are exact, meaning that the error in the approximated continuation value $\hat{c}(x, t_{m+1})$, obtained through the COS method, follows Equation 5.29. When the approximated continuation value is used to compute the coefficients $C_n(x_1, x_2, t_{m+1})$, the following error is introduced:

$$\varepsilon(n) = \frac{2}{b-a} \int_{x_1}^{x_2} \epsilon \cos\left(n\pi \frac{x-a}{b-a}\right) dx = \frac{2\epsilon}{b-a} \psi_n(x_1, x_2).$$

The function $\varepsilon(n)$ can be seen as the product of ϵ and the Fourier-cosine series coefficients $\hat{A}_n$ of the function $a(x)$, defined as:

$$a(x) = \begin{cases} 1 & \text{if } x \in [x_1, x_2] \subset [a, b] \\ 0 & \text{if } x \in \mathbb{R}\setminus[x_1, x_2]. \end{cases}$$

The error $\varepsilon(n) = \epsilon \hat{A}_n$ is present in the computation of the coefficients $V_n(t_{m+1})$, and as a consequence there is another error component in the approximated coefficients $\hat{c}(x, t_m)$:

$$\epsilon_5 = \epsilon \; e^{-r\Delta t} \sum_{n=0}^{N-1}{}' \text{Re}\left\{ \varphi\left(\frac{n\pi}{b-a}\right) e^{\text{i}n\pi\frac{x-a}{b-a}} \right\} \hat{A}_n. \tag{5.30}$$

A very interesting approach given in [15], is to look at Equation 5.30 as the application of the COS method to a European option with payoff $a(x)$ and exact value $v_a(x)$. This perspective allows us, using the error analysis for European options, to bound ϵ_5:

$$|\epsilon_5| < |\epsilon| \; |v_a(x) + \epsilon|.$$

Additionally, recalling the risk-neutral valuation formula:

$$e^{r\Delta t} v_a(x) = \int_{\mathbb{R}} a(y) f(y \mid x) dy = \int_{x_1}^{x_2} f(y \mid x) dy \leq \int_{\mathbb{R}} f(y \mid x) dy = 1.$$

As a consequence, we have $v_a(x) \leq e^{-r\Delta t}$, which implies:

$$|\epsilon_5| < |\epsilon| \; |e^{-r\Delta t} + \epsilon| \sim e^{-r\Delta t} |\epsilon|.$$

This means that the local error remains of the same order and that the COS method is stable even when the coefficients are recovered recursively.

5.3.1 Application to Electricity Storage contracts

We want to extend this analysis to the dynamic pricing algorithm for electricity storage contracts. Since the value of these options is determined at every feasible energy level $e \in E$ and at every Markov state $i \in \mathscr{J}$, we need to include this information in the error.
We start with the definition of ϵ_1:

$$\epsilon_1^{e,i} = v(x, e, i, t_0) - v_1(x, e, i, t_0) = e^{-r\Delta t} \; \sum_{j=1}^{J} p_{ij} \int_{\mathbb{R}\setminus[a,b]} v(y, e, j, t_m) \; f_j(y \mid x) dy.$$

Then, we move on to ϵ_3:

$$\begin{aligned} \epsilon_3^{e,i} &= v_2(x, e, i, t_0) - v_3(x, e, i, t_0) \\ &= e^{-r\Delta t} \; \sum_{j=1}^{J} p_{ij} \sum_{n=0}^{N-1}{}' \text{Re}\left\{ \int_{\mathbb{R}\setminus[a,b]} e^{\text{i}n\pi\frac{\beta x - a}{b-a}} f_j(y \mid x) dy \right\} V_n^j(t_m, e) \end{aligned}$$

Assuming $f_j(y \mid x)$ to be real and interchanging summation and integration:

$$
\begin{aligned}
&= e^{-r\Delta t} \sum_{j=1}^{J} p_{ij} \int_{\mathbb{R}\setminus[a,b]} {\sum_{n=0}^{N-1}}' \left[\cos\left(n\pi \frac{\beta x - a}{b-a} \right) V_n^j(t_m, e) \right] f_j(y \mid x) dy \\
&= e^{-r\Delta t} \sum_{j=1}^{J} p_{ij} \int_{\mathbb{R}\setminus[a,b]} \left[{\sum_{n=0}^{+\infty}}' \cos\left(n\pi \frac{\beta x - a}{b-a} \right) V_n^j(t_m, e) \right. \\
&\qquad\qquad \left. - \sum_{n=N}^{+\infty} \cos\left(n\pi \frac{\beta x - a}{b-a} \right) V_n^j(t_m, e) \right] f_j(y \mid x)\, dy
\end{aligned}
$$

Substituting the cosine expansion of $v(y, e, t, t_m)$ in y by $v(y, e, j, t_m)$:

$$
\begin{aligned}
&= e^{-r\Delta t} \sum_{j=1}^{J} p_{ij} \int_{\mathbb{R}\setminus[a,b]} \left[v(y, e, j, t_m) - \sum_{n=N}^{+\infty} \cos\left(n\pi \frac{\beta x - a}{b-a} \right) V_n^j(t_m, e) \right] f_j(y \mid x) dy \\
&= \epsilon_1 - e^{-r\Delta t} \sum_{j=1}^{J} p_{ij} \int_{\mathbb{R}\setminus[a,b]} \left[\sum_{n=N}^{+\infty} \cos\left(n\pi \frac{\beta x - a}{b-a} \right) V_n^j(t_m, e) \right] f_j(y \mid x) dy
\end{aligned}
$$

Since $V_n^j(t_m, e)$ shows at least algebraic convergence (see [14]), we can bound:

$$
\left| \sum_{n=N}^{+\infty} \cos\left(n\pi \frac{\beta x - a}{b-a} \right) V_n^j(t_m, e) \right| \leq \sum_{n=N}^{+\infty} \left| V_n^j(t_m, e) \right| \leq \frac{Q^*}{(N-1)^{\hat{n}-1}} \leq Q^*, \quad \text{for } N >> 1, \hat{n} >> 1.
$$

From this, defining $Q = e^{-r\Delta t} Q^*$, we get:

$$
|\epsilon_3^{e,i}| < |\epsilon_1^{e,i}| + Q \sum_{j=1}^{J} p_{ij} \left| \int_{\mathbb{R}\setminus[a,b]} f_j(y \mid x) dy \right| < |\epsilon_1^{e,i}| + Q|\epsilon_4^i|,
$$

where $\epsilon_4^i = \sum_{j=1}^{J} p_{ij} \int_{\mathbb{R}\setminus[a,b]} f_j(y \mid x) dy$.
As for the second source of error, ϵ_2, by definition we have:

$$
\epsilon_2^{e,i} = v_1(x, e, i, t_0) - v_2(x, e, i, t_0) = \frac{b-a}{2} e^{-r\Delta t} \sum_{j=1}^{J} p_{ij} \sum_{n=N}^{+\infty} A_n^j(x) \cdot V_n^j.
$$

With an analogous reasoning to [14], we assume the density function to be smoother than the payoff function and the coefficients A_n^j (series coefficients of the density) to decay faster than V_n^j (series coefficients of the payoff). Hence, the product of these two sets of coefficients converges faster than either one of them, so:

$$
\left| \sum_{n=N}^{+\infty} A_n^j(x) \cdot V_n^j \right| \leq \sum_{n=N}^{+\infty} |A_n^j(x)|.
$$

Based on the convergence of the series, either geometrical or algebraic, we can introduce the following quantity:

$$|\tilde{\epsilon}_2^{e,j}| < \begin{cases} P_j \cdot \exp(-(N-1)\nu_j) & \text{if the density in state } j \text{ is in the class } \mathbb{C}^\infty(\mathbb{R}) \\ \frac{\overline{P_j}}{(N-1)^{\gamma_j - 1}} & \text{if the density in state } j \text{ is continuous} \\ & \text{but differentiable up to a finite order,} \end{cases}$$

and bound ϵ_2 as follows:

$$|\epsilon_2^{e,i}| < \sum_{j=1}^{J} p_{ij} |\tilde{\epsilon}_2^{e,j}|.$$

Summarizing the effect of the first 3 sources of error, we obtain:

$$\epsilon^{e,i} < 2|\epsilon_1^{e,i}| + Q|\epsilon_4^i| + |\epsilon_2^{e,i}| \tag{5.31}$$

Lastly, we study the influence of the backward induction on the propagation of the error.

We assume, again, that the coefficients $V_n^i(t_{m+2}, e)$ are exact, meaning that the error in the approximated continuation value $\hat{c}(x, e, i, t_{m+1})$, obtained through the COS method, follows Equation 5.31. From the definition of the coefficients $V_n^i(t_{m+2}, e)$, we recall that they also depend on the coefficients G_n^i and Q_n, which are the cosine series coefficients of the payoff function and penalty function, respectively. However, we have a closed form solution for these coefficients, therefore we do not take into account any error with regard to them. When the approximated continuation value is used to compute the coefficients $C_n^i(x_1, x_2, e, t_{m+1})$, the following error is introduced:

$$\varepsilon^{e,i}(n) = \frac{2}{b-a} \int_{x_1}^{x_2} \epsilon^{e,i} \cos\left(n\pi \frac{x-a}{b-a}\right) dx = \frac{2\epsilon^{e,i}}{b-a} \psi_n(x_1, x_2).$$

The function $\varepsilon^{e,i}(n)$ can be seen as the product of $\epsilon^{e,i}$ and the Fourier-cosine series coefficients $\hat{A}_n$ of the function $a(x)$, defined as:

$$a(x) = \begin{cases} 1 & \text{if } x \in [x_1, x_2] \subset [a, b] \\ 0 & \text{if } x \in \mathbb{R} \backslash [x_1, x_2]. \end{cases}$$

The error $\varepsilon^{e,i}(n) = \epsilon^{e,i} \hat{A}_n$ is present in the computation of the coefficients $V_n^i(t_{m+1}, e)$, and as a consequence there is another error component in the approximated coefficients $\hat{c}(x, e, i, t_m)$:

$$\epsilon_5^{e,i} = \epsilon^{e,i}\ e^{-r\Delta t} \sum_{j=1}^{J} p_{ij} {\sum_{n=0}^{N-1}}' \mathrm{Re}\left\{ \varphi^j\left(\frac{n\pi}{b-a}\right) e^{\mathrm{i}n\pi \frac{\beta x - a}{b-a}} \right\} \hat{A}_n. \tag{5.32}$$

As done before, we look at Equation 5.32 as the application of the COS method to a European option with payoff $a(x)$ and exact value $v_a(x)$. This perspective allows us, using the error analysis for European options, to bound ϵ_5:

$$|\epsilon_5^{e,i}| < |\epsilon^{e,i}|\ |v_a(x) + \epsilon^{e,i}| < \ldots \sim e^{-r\Delta t}|\epsilon^{e,i}|.$$

This shows again that the local error remains of the same order and that the COS method is stable even when the coefficients are recovered recursively and we are considering a Markov-modulated framework.

5.4 Results

In this section, we present the results of the implementation of the Markov-modulated framework with regard to the pricing of electricity storage contracts. We start by showing the effects of this new framework on the density recovery and later on the pricing of these options. We analyze the performance of the COS method, compared to the LSMC method, used as a benchmark.

5.4.1 Density recovery

We have introduced the price model in Section 5.1.3 and described the process of density recovery in Section 5.2.1. Our goal here is to showcase how changing the probabilities of the underlying Markov chains affects the marginal density of the electricity price X_t. For these experiments, we consider two different sets of parameters for the normal distribution in the spike regime, the first where the mean is closer to the mean of the OU process, and the second where the mean is more separated from the one of the OU process. These parameters can be found in table 5.1. As for the general setup, the maturity is set to $T = 1$ year, the COS experiments are conducted with $N = 200,\ L = 8$ and the starting state is always set to be the base one ($S_0 = 0$).

We start by analyzing the density with the first set of parameters, and observe how it changes with the Markov chain probabilities.

An interesting first thing to notice in Figure 5.2 is that across all plots, the density recovered via the COS method shows excellent agreement with the analytical solution, validating the accuracy of the method.
Figure 5.1a shows the case where, with a transition probability $p_{bs} = 0$, the process can never transition to the spike state. The condition $p_{sb} = 1$ ensures that if hypothetically the process was in the spike regime, it would immediately transition back to the base state. This means that the process is effectively always confined to the base regime. Therefore, the resulting PDF is identical to the marginal distribution of the pure OU process at T=1, which is a normal distribution with mean $\mu = 20$ and variance $\sigma^2 = 10$.

Table 5.1: Parameters and moments of the OU process and Normal distributions

Process	**Parameters**			**Mean**	**Variance**
OU	κ	θ	σ	20	10
	5	20	10		
	μ_s	σ_s			
	3	5		23	25
Normal	10	5		30	25

The high peak and narrow spread are representative of its relatively low variance.
In figure 5.1b, we can observe the effect of mixing the two regimes. Specifically, the stationary probability of being in the spike state increases, leading to a shift in the mean and an increase in the variance. As for the mean, the central peak of the distribution shifts progressively to the right, moving from 20 towards 23. As for the variance, the peak of the density function becomes lower, and the distribution becomes wider.
In figure 5.1c, the stationary probability of being in the spike state is greater than that of being in the base state, so the previously observed trend is emphasized. The mean of the distribution is now closer to 23 than to 20 and the variance continues to increase, leading to a progressively lower and wider density curve.
Finally, in figure 5.1d, the process is effectively confined to the spike state. The resulting PDF is simply the $N(\mu_s = 23, \sigma_s^2 = 25)$ distribution defined for the spike regime. Again, the density is symmetric and compared to 5.1a, this distribution is significantly wider and has a lower peak.

In the second case, the parameters of the spike regime are adjusted to increase the separation between the means of the two states in order to further investigate the versatility of the Markov-modulated process. The mean of the normal distribution is shifted from $\mu_s = 23$ to $\mu_s = 30$, while all other parameters remain the same.

In figure 5.2a we observe an asymmetric distribution which exhibits a positive skew, given by the long right tail. The PDF has two peaks, one on the left which is more pronounced and centered around the mean of the OU process, and the second, less pronounced, around the mean of the normal distribution. This is because the transition probability of being in the base state $\mathbb{P}(S_t = 0) = \frac{p_{sb}}{p_{bs}+p_{sb}} = \frac{0.7}{0.3+0.7} = 0.7$ is higher than that of being in the spike regime ($\mathbb{P}(S_t = 1) = 0.3$), therefore the overall distribution is dominated by the features of the base regime. The contribution from the spike state is not sufficient to form a distinct second mode, but it nevertheless pulls the right tail of

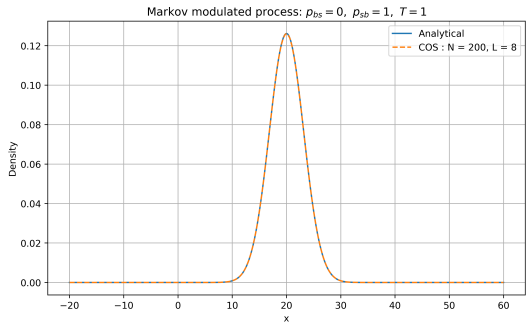

(a) $p_{bs} = 0,\ p_{sb} = 1$

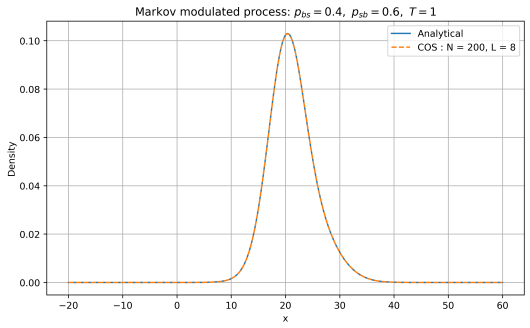

(b) $p_{bs} = 0.4,\ p_{sb} = 0.6$

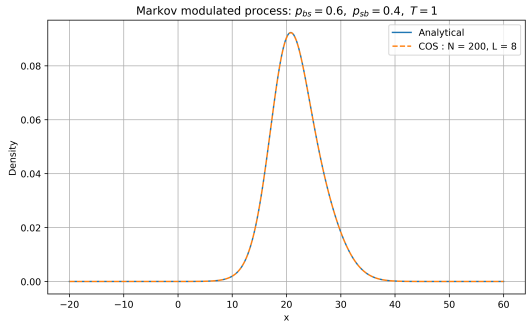

(c) $p_{bs} = 0.6,\ p_{sb} = 0.4$

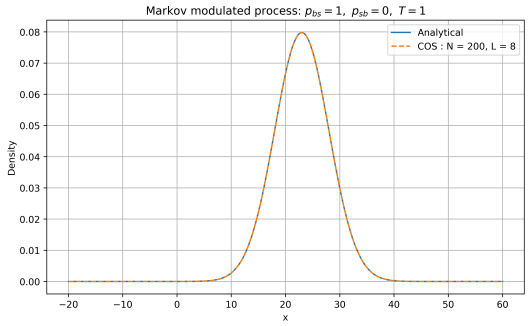

(d) $p_{bs} = 1,\ p_{sb} = 0$

Figure 5.1: Effect of the changing Markov chain probabilities on the density function, with the first set of parameters for the price process

the distribution, creating the observed shoulder and the positive skew.
Figure 5.2b, displays the case where the regimes are perfectly balanced, and there is an evident bimodality caused by the separation between the means which is sufficiently large relative to the standard deviations of the distributions. The difference in the peak heights is due to the difference in variances. On the one hand, the base state concentrates its probability mass into a narrow region because of its small variance, resulting in a high peak. On the other hand, the spike state spreads its probability mass over a wider range because of its larger variance, resulting in a lower and broader peak.
Lastly, figure 5.2c portraits the opposite situation to figure 5.2a, because the stationary probabilities are now inverted. The PDF is unimodal but exhibits strong negative skew. The dominant peak is now located around the mean of the normal distribution and is visibly broader than the dominant peak in Figure 5.2a. A distinct shoulder is present on the left side of the distribution, around the mean of the OU process.

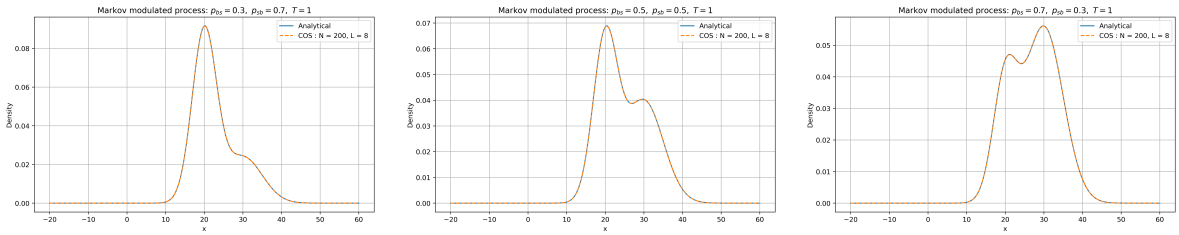

(a) $p_{bs} = 0.3,\ p_{sb} = 0.7$ (b) $p_{bs} = 0.5,\ p_{sb} = 0.5$ (c) $p_{bs} = 0.7,\ p_{sb} = 0.3$

Figure 5.2: Effect of the changing Markov chain probabilities on the density function, with the second set of parameters for the price process

5.4.2 Option pricing

In this section, we present the results on pricing electricity storage contracts in the context where the electricity price is described by a Markov-modulated process, as defined in section 5.1.3. Figure 5.3 displays some simulated paths of the process X_t.

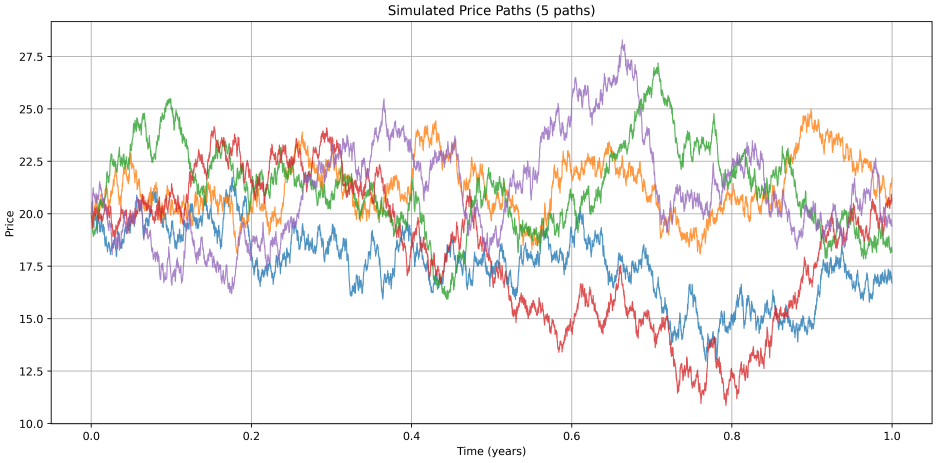

Figure 5.3: Simulation of the Markov-modulated price process X_t with parameters: $\kappa_{base} = 3, \theta_{base} = 20, \sigma_{base} = 10, \mu_{spike} = 3, \sigma_{spike} = 8, p_{bs} = 0.05, p_{sb} = 0.8$, and maturity $T = 1$.

We discuss the results for the same contracts defined in section 4.3. The COS experiments are performed with varying values of the parameters $L \in \{6, 7, 8, 9, 10\}$ and $N \in \{50, 100, 150, 200\}$. As for the benchmark, we use the LSMC confidence interval, defines as:

$$\text{Confidence Interval} = \left[\overline{V} - z_{\alpha/2}\left(\frac{\overline{\sigma}}{\sqrt{10}}\right); \overline{V} + z_{\alpha/2}\left(\frac{\overline{\sigma}}{\sqrt{10}}\right)\right],$$

where $\overline{V}$ is the sample mean obtained with ten experiments, $\overline{\sigma}$ is the standard deviation and $z_{\alpha/2} = 1.96$ is the critical value for the 95% confidence interval. These confidence intervals are the result of 10 runs of the LSMC method with 20000 trajectories.

Contract 1: Standard electricity storage
Contract 1 is described in section 4.3.1. The results can be found in Table 5.2 and in Figure 5.4a.

Table 5.2: Contract 1: LSMC confidence interval and COS prices for various L and N.

LSMC Confidence Interval: $[47.9549, 48.2784]$				
	N values			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	48.03668443	48.0360353	48.03603198	48.03603198
7	48.01308196	48.03603809	48.03603199	48.03603198
8	47.97212243	48.0358958	48.03603215	48.03603198
9	47.9195795	48.03542725	48.03603228	48.03603198
10	47.84752502	48.03502395	48.03602671	48.036032

The COS prices fall within the confidence interval for all values of L when $N \geq 100$. However, when $N = 50$, we see that the greater L, the further the price is from the value to which the COS prices converge. An explanation of this phenomenon is that when L is large but N is small, the series approximation error is dominant. Basically, we are asking the model to stretch a small number of cosine waves over a very wide interval, which results in a poor approximation of the function, leading to inaccurate prices. When L is small, with $N = 50$, the number of terms is more adequate to approximate the function over the smaller range. While the truncation error is technically larger (since we are ignoring a greater portion of the tails), the series approximation error is smaller. In this case, this leads to a price that is much closer to the value to which the COS prices converge.

Contract 2: Highly efficient electricity storage
Contract 1 is described in section 4.3.2. The results can be found in Table 5.3 and in Figure 5.4b.

We can observe here a similar situation to what we described above, with the only difference that when $N = 50$ and the value of L increases, the COS prices are higher, rather than lower, than the value to which they converge. We can also notice how the prices of this contract are higher than the prices of Contract 1, and this is because the 100% efficiency of the battery allows the holder of the contract to better exploit the

Table 5.3: Contract 2: LSMC confidence interval and COS prices for various L and N.

LSMC Confidence Interval: $[71.4886, 71.7919]$				
	N values			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	71.61310678	71.59544324	71.59543802	71.59543802
7	71.68398552	71.59548444	71.59543804	71.59543802
8	71.8213078	71.59576803	71.59543818	71.59543802
9	72.00017226	71.59689983	71.59543961	71.59543802
10	72.21793153	71.5998399	71.59545181	71.59543804

fluctuations in the price process.

Contract 3: Car-park as Power Plant
Contract 3 is described in section 4.3.3. The results can be found in Table 5.4 and in Figure 5.4c.

Table 5.4: Contract 3: LSMC confidence interval and COS prices for various L and N.

LSMC Confidence Interval: $[26.3324, 26.4788]$				
	N values			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	26.44378935	26.44465928	26.44465896	26.44465896
7	26.42628408	26.44467517	26.44465897	26.44465896
8	26.38786368	26.44476296	26.44465891	26.44465896
9	26.32587164	26.44469425	26.44465707	26.44465896
10	26.24915128	26.44366845	26.444657	26.44465895

These results are very similar to what we presented above, and the COS prices have a similar behavior to those of Contract 1. In fact, the COS prices fall within the LSMC interval for all values of L when $N = \{100, 150, 200\}$, and when $N = 50$, they fall within the interval for small values of L, $L = \{6, 7, 8\}$.

Contract 4: Cost optimization of charging electric vehicles
Contract 4 is described in section 4.3.4. The results can be found in Table 5.5 and in Figure 5.4d.

In this case, the COS prices always fall within the LSMC interval, yet still exhibit

Table 5.5: Contract 4: LSMC confidence interval and COS prices for various L and N.

LSMC Confidence Interval: $[-184.9461, -184.6735]$				
	N values			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	-184.79513047	-184.78917776	-184.78917899	-184.78917899
7	-184.81867616	-184.78916646	-184.78917899	-184.78917899
8	-184.85088088	-184.78911584	-184.78917906	-184.78917899
9	-184.88265674	-184.78926493	-184.78918027	-184.78917899
10	-184.91007381	-184.79014249	-184.78917731	-184.78917899

the behavior described previously for $N = 50$.

Tables 5.6, 5.7, 5.8, 5.9 display the computational times required to recover the prices using the LSMC algorithm and the COS method for all four types of contracts. The COS method produces accurate results already when $N = 100$, increasing N leads to only marginal changes in the computed prices, while incurring a higher computational cost. For all contracts, the computational time of the COS method with $N = 100$ and $L = 6$ is approximately 5 times less than that required by the LSMC algorithm. This advantage decreases as L increases, reducing to a speed-up factor of about 4 when $L = 10$.

All these results confirm that the COS method is an extremely efficient option pricing tool, even when the framework is further complicated by the presence of multiple Markov states. The COS method converges rapidly, yielding stable and accurate prices for a modest number of expansion terms N.

A critical advantage of the COS method over the benchmark LSMC algorithm lies in its deterministic nature. As a numerical integration technique, the COS method produces a single and repeatable solution for a given set of parameters. This contrasts with the stochastic nature of the LSMC method. If the experiments with this method are performed multiple times, we might observe a slight difference in the confidence intervals. This is a consequence of the sampling error, meaning that the resulting price and confidence interval are themselves random variables. Therefore, the COS method proves to be a more reliable and practical tool for this application, since it provides consistent and computationally efficient solutions.

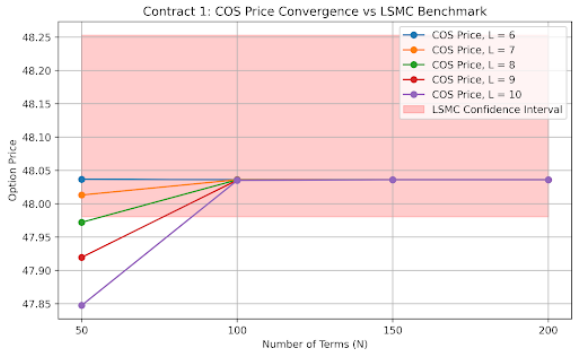

(a) Contract 1

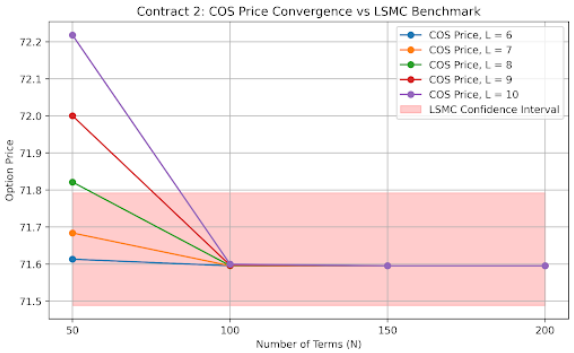

(b) Contract 2

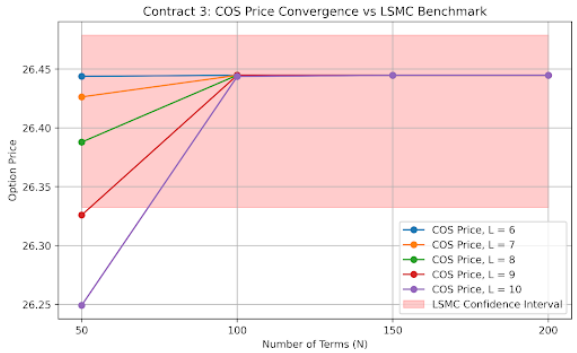

(c) Contract 3

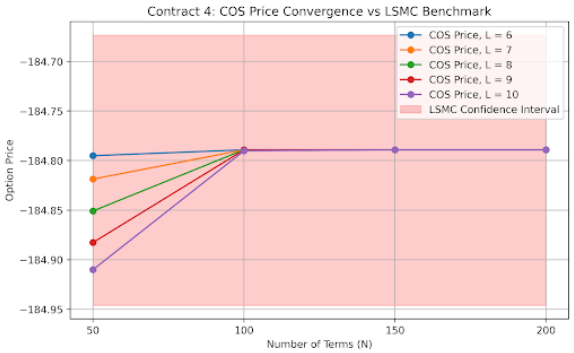

(d) Contract 4

Figure 5.4: Convergence of the COS prices to the LSMC confidence interval with different values of L and N

Table 5.6: Contract 1: Computational times needed to recover the LSMC confidence interval and COS prices for various L and N.

LSMC CPU time: 8 min 44.07 sec				
	COS CPU times			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	0 min 47.82 sec	1 min 37.84 sec	2 min 49.40 sec	4 min 14.69 sec
7	0 min 52.79 sec	1 min 47.26 sec	2 min 57.67 sec	4 min 20.65 sec
8	0 min 57.82 sec	1 min 56.66 sec	3 min 12.21 sec	4 min 37.68 sec
9	1 min 3.19 sec	2 min 5.84 sec	3 min 25.46 sec	4 min 55.65 sec
10	1 min 8.97 sec	2 min 15.34 sec	3 min 39.37 sec	5 min 12.23 sec

Table 5.7: Contract 2: Computational times needed to recover the LSMC confidence interval and COS prices for various L and N.

LSMC CPU time: 8 min 5.90 sec				
	COS CPU times			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	0 min 45.26 sec	1 min 32.59 sec	2 min 32.41 sec	3 min 46.68 sec
7	0 min 50.59 sec	1 min 41.24 sec	2 min 45.74 sec	4 min 2.01 sec
8	0 min 56.17 sec	1 min 52.25 sec	3 min 1.51 sec	4 min 21.17 sec
9	0 min 59.94 sec	2 min 1.40 sec	3 min 14.87 sec	4 min 39.22 sec
10	1 min 6.68 sec	2 min 9.29 sec	3 min 26.91 sec	4 min 57.10 sec

Table 5.8: Contract 3: Computational times needed to recover the LSMC confidence interval and COS prices for various L and N.

LSMC CPU time: 6 min 54.25 sec				
	COS CPU times			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	0 min 32.67 sec	1 min 4.07 sec	1 min 43.17 sec	2 min 32.14 sec
7	0 min 34.81 sec	1 min 12.60 sec	1 min 56.52 sec	2 min 48.80 sec'
8	0 min 39.77 sec	1 min 20.62 sec	2 min 10.67 sec	3 min 7.84 sec'
9	0 min 44.13 sec	1 min 28.68 sec	2 min 21.26 sec	3 min 21.76 sec
10	0 min 47.61 sec	1 min 34.83 sec	2 min 33.41 sec	3 min 37.01 sec

Table 5.9: Contract 4: Computational times needed to recover the LSMC confidence interval and COS prices for various L and N.

LSMC CPU time: 7 min 2.64 sec				
	COS CPU times			
L	**N = 50**	**N = 100**	**N = 150**	**N = 200**
6	0 min 30.99 sec	1 min 3.78 sec	1 min 43.10 sec	2 min 34.30 sec
7	0 min 34.92 sec	1 min 15.89 sec	1 min 57.25 sec	2 min 55.86 sec
8	0 min 39.94 sec	1 min 20.33 sec	2 min 8.92 sec	3 min 4.04 sec
9	0 min 43.92 sec	1 min 26.62 sec	2 min 19.23 sec	3 min 18.68 sec
10	0 min 46.95 sec	1 min 34.34 sec	2 min 28.28 sec	3 min 31.60 sec

5.4.3 Discrepancies in the prices

In the previous section, we have shown results in which the prices for the various contracts obtained with the COS method fall within the LSMC confidence intervals. However, some of our experiments do not show a completely accurate match. We present here an example using a different set of parameters to describe the underlying electricity price process (simulated in Figure 5.5). Figures 5.7 display the convergence

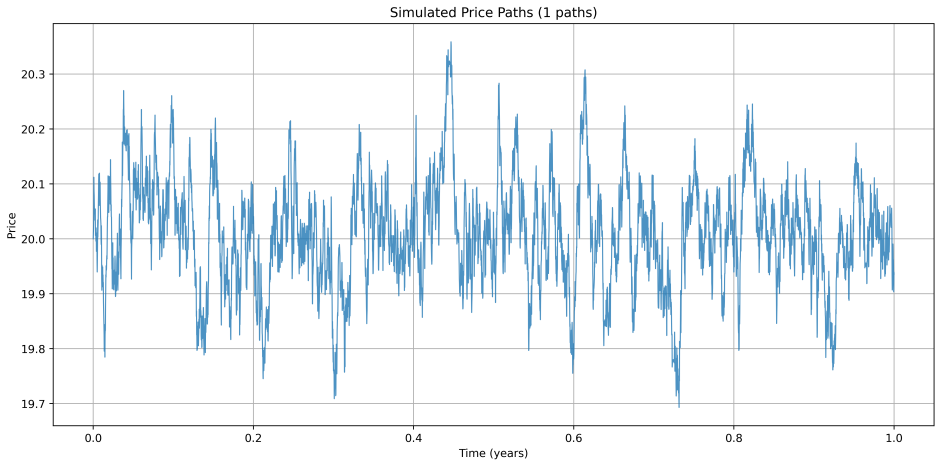

Figure 5.5: Simulation of the Markov-modulated price process X_t with parameters: $\kappa_{base} = 0.5 \cdot 365, \theta_{base} = 20, \sigma_{base} = 0.1\sqrt{365}, \mu_{spike} = 0.001 \cdot 365, \sigma_{spike} = 0.01\sqrt{365}, p_{bs} = 0.01, p_{sb} = 0.3$, and maturity $T = 1$.

of the LSMC prices with an increasing number of paths ($N \in \{10000, 20000, 40000\}$) to the COS value obtained with $N = 200, L = 10$. These experiments are conducted with different maturities $T \in \{0.1, 0.5, 1\}$ and number of exercise dates $B \in \{30, 50, 70\}$.
A consistent pattern that can be observed is that for most of the cases the LSMC price estimates are higher than the COS prices, except for Figure 5.7c where they are lower. We are inclined to attribute the observed discrepancy to the LSMC algorithm rather than the COS method, since the latter accurately reproduces the density functions across all maturities, as shown in Figure 5.6.

This section seeks to explain this mismatch by referencing the biases that are common to the LSMC algorithm, which are explained in more detail in Appendix A. Specifically, there are two main sources of bias in the LSMC estimator:

- Low-bias, which arises from the use of a finite number of basis functions M_{basis}. In short, the regression determines an exercise strategy that is suboptimal, and hence yields an option value which is less than or equal to the true one.

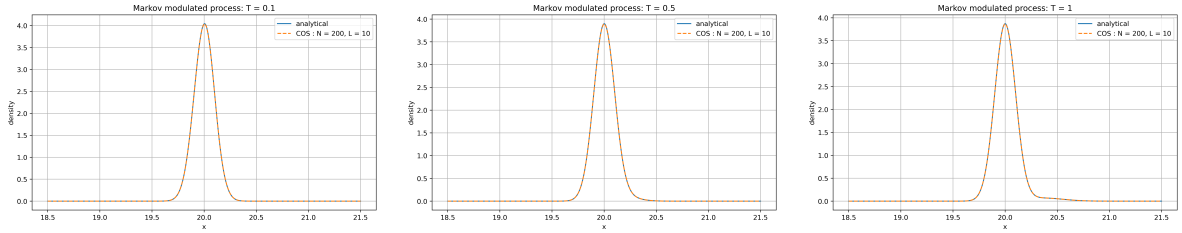

(a) $T = 0.1$ (b) $T = 0.5$ (c) $T = 1$

Figure 5.6: COS recovered density and analytical density for different maturities.

- High-bias, caused by the use of the same set of simulated paths for estimating the regression coefficients and for valuing the option, which leads to overfitting.

In these figures, the high-bias is the dominant effect. We can observe how the confidence intervals become tighter as the number of simulation paths increases, but the mean estimate does not converge to the COS price.

In an attempt to mitigate these discrepancies, we tried two different remedies. First, we increased the number of basis functions from $M_{basis} = 3$ to $M_{basis} = \{6, 9, 12\}$, however, as shown in Figure 5.8, this did not lead to a significant improvement compared to Figure 5.7c. An explanation for this phenomenon might be that the chosen basis functions are inadequate to capture the true shape of the continuation value function for some parameter sets. Another reason is that the number of simulated paths N_{paths} might be insufficient to correctly estimate the regression coefficients for the higher number of basis functions M_{basis}.

Second, we have tried to address the high-bias by using two independent sets of paths: one to estimate the regression coefficients and one to valuate the options. As illustrated in Figure 5.9, this approach also failed to fix the price mismatch, since the price behavior is analogous to the original in Figure 5.7h. This outcome suggests that the number of simulation paths N_{paths} remains insufficient to accurately approximate the optimal exercise strategy and the corresponding option values. Further increasing N_{paths} could potentially improve the accuracy but would also involve a considerable computational cost, given the complexity of the problem.

The bias is likely due to an insufficient number of basis functions M_{basis} and simulation paths N_{paths}. The attempted remedies did not successfully solve the issue, which means that the estimation of the prices through LSMC is particularly challenging for this set of parameters. Achieving higher accuracy would probably require a substantially larger number of simulation paths, a more advanced choice of basis functions, or the use of alternative approximation techniques.

Figure 5.7: Mismatch in the convergence of LSMC prices to the COS price for different maturities (T) and numbers of exercise dates (B).

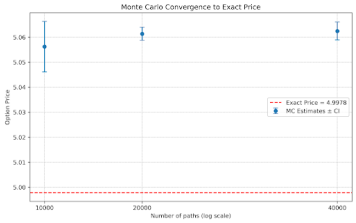

(a) $T = 0.1, B = 30$

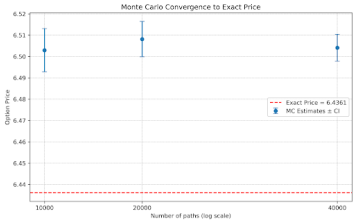

(b) $T = 0.1, B = 50$

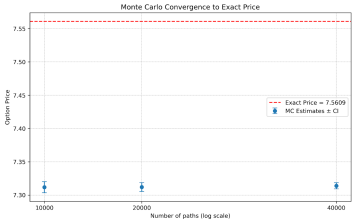

(c) $T = 0.1, B = 70$

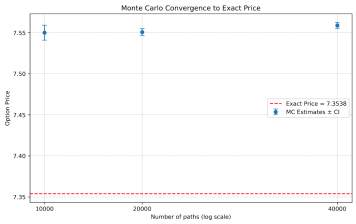

(d) $T = 0.5, B = 30$

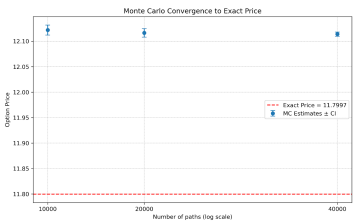

(e) $T = 0.5, B = 50$

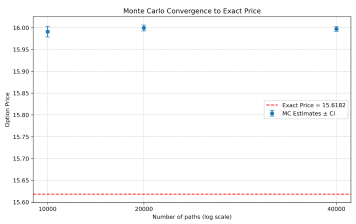

(f) $T = 0.5, B = 70$

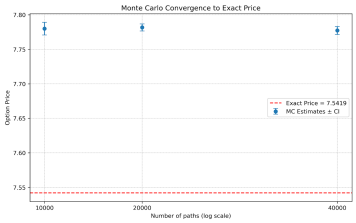

(g) $T = 1, B = 30$

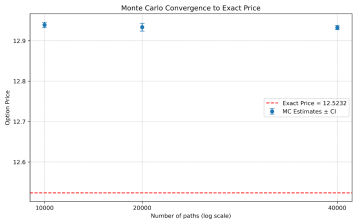

(h) $T = 1, B = 50$

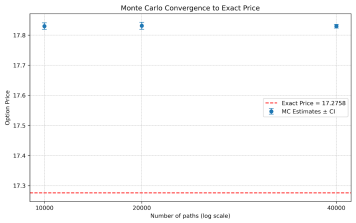

(i) $T = 1, B = 70$

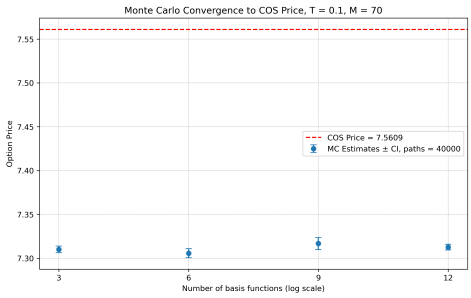

Figure 5.8: Mismatch in the convergence of LSMC prices to the COS price w.r.t the number of basis functions M_{basis}.

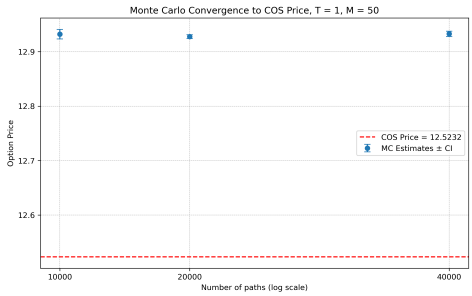

Figure 5.9: Mismatch in the convergence of LSMC prices to the COS price using two different sets of paths for the regression.

6 Conclusions

Energy commodities exhibit some unique features, such as mean-reversion, seasonality and sudden price spikes, that distinguish them from traditional financial assets. This creates challenges for the valuation of energy derivatives.

This thesis investigates and demonstrates the efficiency of the COS Method for pricing instruments within this domain. This goal is achieved gradually by considering cases of escalating difficulty, where each chapter builds upon the last to address increasingly advanced valuation problems.

The starting point of our analysis are Energy Quanto options, European-style contracts whose value depends on two correlated underlying processes which incorporate jumps within a Markov-modulated framework. We analyze two different cases, taken from [2], where these options are priced using a FFT based method. Based on some useful numerical results provided in [2], our analysis focuses on the derivation of the components of the characteristic functions of the two proposed models, which are required to use the COS method. The numerical results prove that the COS method is extremely efficient, as the results not only match the benchmarks, but also provide a significant computational speed-up compared to the FFT and Monte Carlo benchmarks.

The study is then extended to Electricity Storage contracts, which introduce the complexity of early-exercise features and operational constraints. Our analysis of this type of contracts starts from the replication of the main results of [7]. After giving a mathematical definition of the contracts and their physical and operational constraints, the COS method is integrated within the dynamic pricing algorithm and is used to determine the continuation values at each exercise date and at each feasible energy level of the storage. Specifically, the analysis covers four different storage contract types under various volatility scenarios. The produced results converge within the LSMC confidence interval, demonstrating the adaptability of the COS method to handle the recursive nature of early-exercise feature combined with the constraints of energy storage.

Lastly, the final and main contribution of this work is to synthesize the challenges presented in the previous chapters by valuing electricity storage contracts within a two-state Markov-modulated framework. To do so, the COS-based dynamic pricing

algorithm is extended to handle the state-dependent price dynamics. This extension is successful in the sense that the method manages to accurately recover the complex probability densities, and the pricing results for the four different types of contracts show rapid and stable convergence. Moreover, the analysis highlights the interplay between the two COS parameters L and N and proves the COS method's deterministic nature to be more reliable than the stochastic LSMC benchmark. Indeed, while the results for this framework generally converge within the LSMC confidence intervals, the analysis also highlights challenging parameter sets where the LSMC method exhibits a persistent high-bias and fails to converge to the COS price. This observation demonstrates a notable limitation of the benchmark, and further highlights the reliability of the COS method.

This work succeeds in demonstrating the high versatility of the COS method, whose efficiency is not only confined to financial markets. The COS method is proven to be capable of handling the distinct features of energy derivatives, including regime-switching, jump dynamics and also early-exercise features.

Across all experiments, the COS method outperforms the benchmark methods, by orders of magnitude in speed, showcasing computational efficiency. Moreover, a clear advantage over the benchmark LSMC method is that the COS method, as a numerical integration technique, produces a single and repeatable solution. On the other hand, the LSMC method is subject to statistical errors. Our analysis identified challenging scenarios where the LSMC method shows a persistent high-bias that standard remedies failed to correct, highlighting its susceptibility to sampling error and overfitting. In contrast, the COS method's deterministic nature makes it more reliable for producing stable and accurate results.

However, this work has some limitations that can be starting points for future work. First of all, the chosen models for the underlying prices, while sophisticated, remain abstractions of reality. It might be interesting to apply the pricing frameworks developed in this thesis to other models, such as those that incorporate more than two Markov states or alternative dynamics, and see if they can better fit the prices from real markets. Building on this, a significant extension would be to calibrate the Markov-modulated models to real market data for electricity prices, as this study relies on predefined parameters. Another limitation of this work concerns electricity storage contracts. Our analysis relies on a discretization of the storage capacity into a finite number of equally spaced energy levels. This uniform grid may not fully represent the operational realities, therefore future research could explore more complex and realistic discretization schemes and see how that affects the final contract prices and the performance of the algorithm. Directly related to this previous point, this thesis does not include a formal error analysis concerning the energy storage discretization. It would be of great interest to conduct such an analysis and test how the numerical stability is affected by the grid choice. Lastly, our work focuses on energy markets, but we only

consider two classes of options. A natural extension is to adapt the developed Markov-modulated COS framework to other contracts which may involve different constraints or higher-dimensional state spaces, to further investigate its flexibility and scalability.

Bibliography

[1] Rashed Khalid Alghanim. *Modelling Electricity Prices and Pricing Derivatives with Markov Regime Switching Models*. The University of Manchester (United Kingdom), 2019.

[2] Fred Benth, Griselda Deelstra, and Sinem Kozpınar. "Pricing energy quanto options in the framework of Markov-modulated additive processes". In: *IMA Journal of Management Mathematics* 34 (Sept. 2021). DOI: 10.1093/imaman/dpab032.

[3] Fred Benth, Jurate Saltyte Benth, and Steen Koekebakker. *Stochastic Modeling of Electricity and Related Markets*. Nov. 2008. ISBN: 978-981-281-230-8. DOI: 10.1142/6811.

[4] Fred Espen Benth, Álvaro Cartea, and Rüdiger Kiesel. "Pricing forward contracts in power markets by the certainty equivalence principle: Explaining the sign of the market risk premium". In: *Journal of banking & finance* 32.10 (2008), pp. 2006–2021.

[5] Fred Espen Benth, Nina Lange, and Tor Åge Myklebust. "Pricing and hedging quanto options in energy markets". In: *Journal of Energy Markets* 8.1 (2015), pp. 1–35.

[6] Michael Bierbrauer, Stefan Trück, and Rafał Weron. "Modeling electricity prices with regime switching models". In: *International conference on computational science*. Springer. 2004, pp. 859–867.

[7] Boris C Boonstra and Cornelis W Oosterlee. "Valuation of electricity storage contracts using the COS method". In: *Applied Mathematics and Computation* 410 (2021), p. 126416.

[8] Peter Carr and Dilip Madan. "Option valuation using the fast Fourier transform". In: *Journal of computational finance* 2.4 (1999), pp. 61–73.

[9] Alvaro Cartea and Marcelo G Figueroa. "Pricing in electricity markets: a mean reverting jump diffusion model with seasonality". In: *Applied Mathematical Finance* 12.4 (2005), pp. 313–335.

[10] Rama Cont and Peter Tankov. *Financial modelling with jump processes*. Chapman and Hall/CRC, 2003.

[11] Cyriel De Jong. "The nature of power spikes: A regime-switch approach". In: *Studies in Nonlinear Dynamics & Econometrics* 10.3 (2006).

[12] Shijie Deng. *Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes*. University of California Energy Institute Berkeley, 2000.

[13] Ernst Eberlein and Kathrin Glau. *Variational Solutions of the Pricing PIDEs for European Options in Lévy Models*. Jan. 2015. DOI: 10.1080/1350486X.2014.886817.

[14] Fang Fang and Cornelis W Oosterlee. "A novel pricing method for European options based on Fourier-cosine series expansions". In: *SIAM Journal on Scientific Computing* 31.2 (2009), pp. 826–848.

[15] Fang Fang and Cornelis W Oosterlee. "Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions". In: *Numerische Mathematik* 114.1 (2009), pp. 27–62.

[16] Diego Garcıa. "Convergence and biases of Monte Carlo estimates of American option prices using a parametric exercise rule". In: *Journal of Economic Dynamics and Control* 27.10 (2003), pp. 1855–1879.

[17] Hélyette Geman and Andrea Roncoroni. "Understanding the Fine Structure of Electricity Prices". In: *The Journal of Business* 79 (May 2006). DOI: 10.1086/500675.

[18] Ronald Huisman and Ronald Mahieu. "Regime jumps in electricity prices". In: *Energy economics* 25.5 (2003), pp. 425–434.

[19] Willett Kempton and Steven E Letendre. "Electric vehicles as a new power source for electric utilities". In: *Transportation Research Part D: Transport and Environment* 2.3 (1997), pp. 157–175.

[20] Willett Kempton et al. "Vehicle-to-grid power: battery, hybrid, and fuel cell vehicles as resources for distributed electric power in California". In: (2001).

[21] S. Ken-Iti. *Lévy Processes and Infinitely Divisible Distributions*. Cambridge studies in advanced mathematics. Cambridge University Press, 1999. ISBN: 9780521553025. URL: https://books.google.it/books?id=tbZPLquJjSoC.

[22] Gregory F Lawler. *Introduction to stochastic processes*. Chapman and Hall/CRC, 2018.

[23] Francis A Longstaff and Eduardo S Schwartz. “Valuing American options by simulation: A simple least-squares approach”. In: *The review of financial studies* 14.1 (2001), pp. 113–147.

[24] Julio J Lucia and Eduardo S Schwartz. “Electricity prices and power derivatives: Evidence from the Nordic power exchange”. In: *Review of derivatives research* 5.1 (2002), pp. 5–50.

[25] Carlo Mari. “Regime-switching characterization of electricity prices dynamics”. In: *Physica A: statistical mechanics and its applications* 371.2 (2006), pp. 552–564.

[26] Cornelis W Oosterlee and Lech A Grzelak. *Mathematical modeling and computation in finance: with exercises and Python and MATLAB computer codes*. World Scientific, 2019.

[27] Omid Palizban and Kimmo Kauhaniemi. “Energy storage systems in modern grids—Matrix of technologies and applications”. In: *Journal of Energy Storage* 6 (2016), pp. 248–259.

[28] Walter Rudin. *Real and complex analysis*. McGraw-Hill, Inc., 1987.

[29] Marjon J Ruijter and Cornelis W Oosterlee. “Two-dimensional Fourier cosine series expansion method for pricing financial options”. In: *SIAM Journal on Scientific Computing* 34.5 (2012), B642–B671.

[30] Eduardo S Schwartz. “The stochastic behavior of commodity prices: Implications for valuation and hedging”. In: *The Journal of finance* 52.3 (1997), pp. 923–973.

[31] Jan Seifert and Marliese Uhrig-Homburg. “Modelling jumps in electricity prices: theory and empirical evidence”. In: *Review of Derivatives Research* 10.1 (2007), pp. 59–85.

[32] Gordon D Smith. *Numerical solution of partial differential equations: finite difference methods*. Oxford university press, 1985.

[33] Rafal Weron. *Modeling and forecasting electricity loads and prices: A statistical approach*. John Wiley & Sons, 2006.

[34] Rafal Weron, Michael Bierbrauer, and Stefan Trück. “Modeling electricity prices: jump diffusion and regime switching”. In: *Physica A: Statistical Mechanics and its Applications* 336.1-2 (2004), pp. 39–48.

A Appendix: LSMC algorithm

This section focuses on the LSMC algorithm, which is used to create benchmarks for the experiments conducted in this thesis, see Chapters 4 and 5.
A great example of the usefulness of this algorithm can be found in [23], where it was introduced for the first time as a powerful tool to approximate the value of American options. The main idea behind the algorithm is that the conditional expectation, needed to determine the continuation value, can be approximated using least squares applied to simulation. We give here a mathematical framework for this algorithm, inspired by [23]. Consider the usual complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$, the EMM Q, and a finite time horizon $[0, T]$. We define the finite set of K early-exercise time points $0 = t_1 \leq \ldots \leq t_K = T$ and we represent a sample path with ω.
We introduce the notation $C(\omega, s; t, T)$ to describe the path of cash flows generated by the option, conditional on the fact that the option is not exercised until time t and assuming that the holder of the option is always following the optimal stopping strategy. Following the no-arbitrage theory we know that the value of an option, assuming that it cannot be exercised until after time t_k, is equal to the discounted expectation of the remaining cash flows $C(\omega, s; t_k, T)$ w.r.t the measure Q. Mathematically, this reads:

$$F(\omega, t_k) = \mathbb{E}^{\mathbb{Q}} \left[\sum_{j=k+1}^{K} e^{-\int_{t_k}^{t_j} r(\omega,s)ds} C(\omega, t_j; t_k, T) \mid \mathcal{F}_{t_k} \right], \tag{A.1}$$

where $r(\omega, s)$ is the risk free discount rate.
For a time step t_{k-1} we assume that $F(\omega, t_{k-1})$ can be represented as a linear combination of a countable set of $\mathcal{F}_{t_{k-1}}$-measurable basis functions, leading to:

$$F(\omega, t_{k-1}) = \sum_{j=0}^{\infty} a_j L_j(X),$$

where X is the underlying process, $L_j(\cdot)$ is the j-th basis function and a_j are constants. The approximation of $F(\omega, t_{k-1})$ is denoted by $F_{M_{basis}}(\omega, t_{k-1})$, with M_{basis} being the number of the first basis functions used. As stated in [23], it can be proven that the

fitted value of the regression $\hat{F}_{M_{basis}}(\omega, t_{k-1})$ converges in mean square and probability to $F_{M_{basis}}(\omega, t_{k-1})$ as the number of simulation paths N_{paths} increases. In addition to this, it is also stated that $\hat{F}_{M_{basis}}(\omega, t_{k-1})$ is the best linear unbiased estimator of $F_{M_{basis}}(\omega, t_{k-1})$.
However, due to approximation errors, $F_{M_{basis}}(\omega, t_{k-1})$ is not an unbiased estimator of the real continuation value $F(\omega, t_{k-1})$. Specifically, the quality of the approximation depends on two main factors:

- The number of basis functions M_{basis}: a more flexible set of basis functions provides a better representation of the continuation value;
- The number of simulation paths N_{paths}: a larger number of simulated paths improves the estimation of the regression coefficients.

Proposition 1 in [23] states that the true value of the option is greater than or equal to the expected value of the LSMC estimator, implying that the estimator is low-biased. The intuition behind this result is that the true value of the option corresponds to the optimal exercise strategy, while any other (suboptimal) strategy yields a value that is less than or equal to the true one. Since the LSMC algorithm determines an approximate exercise strategy based on a finite number of basis functions, it produces a suboptimal, hence low-biased, estimate.
Although this estimator is biased for a finite number of basis functions M_{basis}, Proposition 2 in [23] shows that it is consistent, meaning that as $M_{basis} \to +\infty$, the LSMC estimate converges in probability to the true value.
A second type of bias, known as high-bias, arises when the same set of paths is used both to estimate the exercise rule and to value the option. This can be explained by the fact that the algorithm overfits the specific realization of paths, resulting in an exercise rule that appears overly favorable for that sample and leads, on average, to an upwardly biased price estimate (Theorem 1 from [16]).

A.1 Application to Electricity Storage contracts

We focus now on the application of this method to price electricity storage contracts, as explained in [7]. The following notation is introduced:

- S_m^i: The asset price of trajectory i at time t_m;
- $CV_m^{i,e}$: The continuation value at energy level $e \in E$ after action $\Delta e \in \mathcal{A}$ is taken at time t_m;
- $CF_m^{i,e}$: The cash flow at time t_m and energy level $e \in E$;

- $ACF_m^{i,e}$: The total value of the future cash flows for times $t_k,\ k \in \{m, \ldots, B+1\}$, at energy level $e \in E$;
- $DACF_m^{i,e} = e^{-r\Delta t} \cdot ACF_{m+1}^{i,e}$;
- $Q_m^{\Delta e} = q_b(\Delta e)$: the penalty function;
- $PO_m^{i,\Delta e} = g(t_m, S_m^i, \Delta e)$: the payoff function.

The LSMC pricing algorithm is described in Algorithm 1, and the LSMC 95% confidence interval is determined with the following formula:

$$\text{CI} : \left[\hat{\mu} - 1.96 \cdot \frac{\hat{\sigma}}{\sqrt{N_{runs}}};\ \hat{\mu} + 1.96 \cdot \frac{\hat{\sigma}}{\sqrt{N_{runs}}}\right],$$

where N_{runs} is the number of independent runs of the algorithm, $\hat{\mu}$ is the sample mean of N_{runs} experiments and $\hat{\sigma}$ is the sample standard deviation.

Algorithm 1 The LSMC algorithm for electricity storage contracts

for $e = 0, \ldots, N_{\text{cap}}$ **do**
  for $i = 1, \ldots, N_{paths}$ **do**
    $CF^{i,e}_{B+1} = q_S(S^i(t_{B+1}), e)$
    $ACF^{i,e}_{B+1} = CF^{i,e}_{B+1}$
  end for
end for
for $m = B, \ldots, 1$ **do**
  $X^i = S^i(t_m) \quad \forall i = 1, \ldots, N_{paths}$
  for $e = 0, \ldots, N_{\text{cap}}$ **do**
    $DACF^{i,e}_m = e^{-r\Delta t} ACF^{i,e}_{m+1} \quad \forall i = 1, \ldots, N_{paths}$
    Estimate regression coefficients β^e by least squares:

$$\beta^e = \arg\min_{\beta} \sum_{i=1}^{N_{paths}} \left(DACF^{i,e}_m - \phi(X^i)^\top \beta \right)^2$$

    Compute continuation value:

$$CV^{i,e}_m = \phi(X^i_m)^\top \beta^e \quad \forall i = 1, \ldots, N_{paths},$$

    where the function $\phi(X) \in \mathbb{R}^{M_{basis}}$ is a vector of basis functions used to approximate the continuation value as a linear combination of nonlinear transformations of the state variable X.
  end for
  for $e = 0, \ldots, N_{\text{cap}}$ **do**
    for $i = 1, \ldots, N_{paths}$ **do**
      $\Delta e^{i,*} = \arg\max_{\Delta e \in \mathcal{A}(t_m, e)} \{ PO^{i,\Delta e}_m + CV^{i,e+\Delta e}_m + Q^{\Delta e} \}$
      $CF^{i,e}_m = PO^{i,\Delta e^{i,*}}_m + Q^{\Delta e^{i,*}}$
      $ACF^{i,e}_m = CF^{i,e}_m + e^{-r\Delta t} ACF^{i,e+\Delta e^{i,*}}_{m+1}$
    end for
  end for
end for
for $e = 0, \ldots, N_{\text{cap}}$ **do**
  $v(t_0, S(t_0), e) = \frac{1}{N_{paths}} \sum_{i=1}^{N_{paths}} DACF^{i,e}_0 = \frac{1}{N_{paths}} \sum_{i=1}^{N_{paths}} e^{-r\Delta t} ACF^{i,e}_1$
end for