Potential Future Exposure Modelling for the Carbon Market

February 19, 2024

Author:

R.P.J. van Gestel

Supervisors:

dr. B. Roorda

dr. R.A.M.G. Joosten

A thesis submitted to the Faculty of Behavioural, Management and Social Sciences (BMS) and in partial fulfilment of the requirements for the degree of

MSc Industrial Engineering and Management

Faculty of Behavioural, Management and Social Sciences (BMS)

University of Twente

UNIVERSITY OF TWENTE.

DISCLAIMER

This thesis was written by R.P.J. van Gestel as part of his MSc thesis. The opinions expressed and conclusions drawn in this work are the author's own and do not reflect the view of Shell plc, or any of its subsidiaries.

ACKNOWLEDGEMENTS

This report marks the end of my of Master Industrial Engineering and Management at the University of Twente and my time as a student. Months of hard work have resulted in this thesis, which would not have been possible without the support of people around me.

Foremost, I extend my gratitude to Shell for providing me with the opportunity to conduct my thesis research during my internship. Additionally, I would like to convey my thanks to my company supervisors, Lam Lee Li and Jonathan Moles. Lam Lee's continuous assistance, both in terms of content and navigating within the organization, is deeply appreciated. I express my thanks to Jonathan for his steadfast support and meticulous attention to detail, which played a crucial role in the success of my thesis project. Furthermore, my appreciation extends to all other individuals from Shell who assisted me throughout the project, offering their time.

Support did not solely come internally from Shell. Primarily, I extend my gratitude to Berend Roorda, my supervisor at the University of Twente, for his consistent feedback and guidance during the entirety of the project. His precise reading, stimulating discussions, and comprehensive feedback played a pivotal role in enhancing the quality of my research, enabling necessary adjustments, and pushing the outcomes to a higher level. The value of second opinions cannot be overstated, and I wish to express my appreciation to Reinoud Joosten for providing invaluable input by dedicating time to read and review my thesis.

Lastly, I would like to thank my family and friends for their unconditional support and for helping me to achieve this academic and professional milestone.

Remko van Gestel

Rotterdam, February 2024

ABSTRACT

Within this thesis, we developed a model to facilitate the computation of Potential Future Exposure (PFE) at an individual deal level within the carbon market domain. In addition to other commodities such as oil, gas, and electricity, Shell engages in trading carbon allowances and credits, granting the entitlement to emit greenhouse gases or offset them. The imperative for a specialized PFE model tailored to the carbon markets is underscored by their expanding scale. We start this research by analysing the risk factors influencing carbon market prices and the methodologies available for their modelling. We favour directly modelling market prices rather than relying on alternative risk factors such as temperature, correlations with other commodities, or regulatory influences. The primary rationale behind this preference is the pursuit of a comprehensive PFE model that can be universally applied across all carbon markets. The chosen methodology for calculating PFE is the Monte Carlo simulation, as it accommodates the inclusion of deal optionalities, such as vanilla options. Subsequently, we examine carbon market data. Within one segment of the carbon markets, forward prices are available while in another segment, only spot prices are recorded within Shell. This dichotomy implies that for certain markets, it is feasible to directly model the forward curves, whereas for others, the forward prices must be derived from simulated spot prices. Despite the consideration and testing of various forward curve models, we decide to uniformly model forward curves in all carbon markets by simulating the spot price and extrapolating this value with a fixed percentage to establish a linear forward curve. This determination stems from our principal component analysis, revealing that spot prices contribute to more than 96% of the price variance across all carbon markets. Following the assessment of stochastic spot price models on historical data, the Geometric Brownian Motion, Merton Jump Diffusion Model, and the GBM-hypsec model (where a normal random variable is replaced by a hyperbolic secant random variable in the GBM) emerged as the most suitable options. We developed a PFE model in Python, providing users with the capability to employ these three stochastic models within the price simulation engine. This model is specifically designed for a forward sale deal, involving the receipt of a predetermined quantity of carbon allowances or credits in the future in exchange for an upfront lump sum payment, potentially in conjunction with vanilla (put/call) options. Additionally, the model accommodates swap deals, wherein one category of allowance/credit is exchanged for another.

CONTENTS

1. INTRODUCTION

1.1 BACKGROUND

During the 2007-2008 financial crisis, one of the weaknesses of the global financial system turned out to be inappropriate recognition and pricing of counterparty risk (Hull, 2018b). Most especially, the over-the-counter (OTC) derivatives market created a domino effect prior to, and even more so after the bailout of Bear Stearns and the bankruptcy of Lehman Brothers (Gregory, 2012). Risk management practices at the time turned out to be immature to deal with the high level of complexity of financial interconnections. As a result of this, after the financial crisis, counterparty risk became subjected to much more scrutiny.

Although Shell does not classify as a large financial institution, its presence on the global derivative market is noteworthy, being one of the largest global commodity trading companies. While Shell has a long-standing history in the oil and gas market, in light of global climate change, the company is undergoing a transformation to provide more low-carbon energy such as renewable electricity, biofuels, and hydrogen (Shell plc, 2022).

Part of Shell's trading portfolio are carbon allowances, which can be traded both over the counter (OTC) as well as via market exchanges (Shell plc, 2024). A carbon allowance gives companies the right to emit a certain amount of carbon dioxide. Different emission trading systems (ETSs) have been put in place in several countries/regions (World Bank Group, 2023). These ETSs are so-called regulatory carbon markets (RCM), although there are examples of voluntary ETSs as well (e.g., in Japan). In addition, on the voluntary market (VCM) carbon credits (or offsets) can be traded that are generated through voluntary emissions reduction activities (Refinitiv, 2023).

Restricted by the specific market, the carbon allowances and credits can be traded in the form of spot, future, forward, and option contracts. Those contracts that are traded OTC give rise to, among others, counterparty risk. The exchange-traded contracts are associated with less, practically no, counterparty risk. However, the calculations for margin requirements for these contracts exhibit similarities with the quantification methodology for counterparty risk.

1.2 PROBLEM STATEMENT

Counterparty risk can be quantified by assessing the Potential Future Exposure (PFE), which denotes the potential worst-case positive credit exposure in the future under a specified confidence level (Gregory, 2012). In a broader context, modelling PFE serves diverse objectives, which can vary based on the trading portfolio. For instance, it can be employed to ascertain the amount of credit reserves, impacting the Profit and Loss (P&L) statement, and to analyse the risks associated with potential new contracts (S&P Global, 2022). In certain scenarios, the PFE may also undergo regular comparison with pre-established exposure limits. Notably, for the carbon markets, a suitable PFE model is currently lacking, making it the focus of this research.

1.3 RESEARCH SETUP

Based on the problem description and approach described above, the following research goal has been formulated:

Develop a Potential Future Exposure model for contracts traded on the carbon markets.

To attain the research objective, we have formulated four research questions, as outlined below. Each subsequent chapter in this thesis will address one of these research questions. The second

chapter serves as a literature review focusing on carbon markets, encompassing a high-level understanding of their operations. Furthermore, we provide an overview of various approaches to modelling PFE, aiding in the selection of the most appropriate approach for this research.

We dedicate the third chapter to addressing the question of identifying and modelling the risk factors that encapsulate the market dynamics of carbon assets. Essentially, in this chapter we delve into the simulation methodologies for capturing the price behaviour of these assets. Subsequently, we test these methodologies on historical price data in chapter four.

Prior to advancing from the selection of a modelling approach for carbon risk factors to the construction of the PFE model, understanding of the contracts facilitating the trade of these assets is necessary. The development of the PFE model should occur at the contract level, where diverse obligations to buy or sell assets at different times and potentially under specific conditions may exist. Therefore, the identification of a common framework within carbon contracts and its manifestation in the PFE model's structure is needed. We dedicate chapter five to exploring this aspect. Finally, chapter six will concentrate on the validation procedures applicable to the PFE model.

1. What modelling approach is most suitable for Shell's carbon allowances/credits trading portfolio?
    a. What are the main elements of the carbon markets (that are relevant for Shell)?
    b. Which modelling approaches exist for modelling PFE?
2. What are appropriate risk factors for modelling the market behavior of the carbon products (which influence the MtM value of the contracts traded by Shell)?
    a. How can the time evolution of these risk factors be modelled?
    b. What market behaviours would the model replicate (and which behaviours cannot be replicated in the PFE model)?
3. What is the common backbone in the different carbon contracts and how can this be incorporated into the structure of the PFE model?
4. Which tests can be performed to test and validate the model?

2. THEORETICAL CONTEXT

In this chapter we provide a theoretical introduction with a dual focus on the operational dynamics of carbon markets and the categorization of overarching PFE modelling methodologies. In the initial section we explore the key strategies employed in carbon pricing, designed to discourage greenhouse gas (GHG) emissions, and promote the transition towards sustainable energy. In the latter part we delve into three PFE modelling techniques, leading to the selection of the most appropriate approach for this research in the next chapter.

2.1 CARBON MARKETS

GHGs, encompassing gases like carbon dioxide, contribute to the greenhouse effect and play a role in global climate change. In financial contexts, the term "carbon" is commonly used, extending to other GHGs through the concept of carbon dioxide equivalent—a scaling of a gas's greenhouse effect to that of carbon dioxide.

At a high level, the categorization of carbon pricing mechanisms can be distinguished into indirect and direct approaches. Indirect carbon pricing encompasses "instruments that change the price of products associated with carbon emissions in ways that are not directly proportional to the relative emissions associated with those products" (World Bank Group, 2023). While these instruments typically serve broader socio-economic objectives, such as reducing air pollution and augmenting revenues, they often convey a price signal on carbon emissions. Illustrative instances of indirect carbon pricing include commodity and fuel taxes.

Conversely, direct carbon pricing primarily seeks to establish price signals for reducing GHG emissions (World Bank Group, 2023). A carbon tax imposing levies on GHG emissions to incentivize reductions, is a notable example. Another approach is the implementation of an ETS by governments, where a cap is set on GHG emissions for covered entities. These entities must surrender allowances or emission units for their emissions at the end of a compliance period. Typically, these allowances correspond to the emission of one tonne of carbon dioxide equivalent and can be traded among entities and often with other institutions.

2.1.1 Emission Trading Schemes

Several ETS types are in place worldwide, of which the cap-and-trade ETSs currently cover the largest emissions volume (World Bank Group, 2023). In this scheme, the total net volume of GHG emissions in predefined sectors is limited by the government by a cap. Allowances are sold, for example via auctions, to entities and/or are distributed for free. The total volume of allowances sold or distributed on this so-called primary market is the cap. On the secondary market, allowances can be traded. Examples of this type of emission scheme are the EU ETS and the California Cap-and-Trade Program.

A different approach is to allocate a performance benchmark to entities, often as units of emissions per unit of output (emission intensity), which is called a rate-based ETS. This does not limit total emissions but rather the emission intensity. Entities can earn emission units when they achieve a lower emission intensity than the benchmark. Canada’s Output-Based Pricing System and China’s National ETS are examples of rate-based systems.

A number of ETSs have so-called stabilization mechanisms, which allow governments to intervene when supply and demand of allowances create undesirable price levels (World Bank Group, 2023). For example, this can be an auction reserve level preventing the sale of allowances below a floor, or a safety valve releasing extra allowances when a ceiling threshold is reached.

A feared side effect of ETSs is carbon leakage; the shift of emission from an ETS-covered region to another region with no or less strict regulations. A border carbon adjustment (BCA) mechanism is a policy instrument that allows a government to demand a carbon price for imported goods from other jurisdictions that reflects the emissions associated with the good's production. This should create a level playing field for producers within the jurisdiction where BCA is put in place.

Within certain ETSs banking and/or borrowing is allowed. Banking allows entities to hold a surplus of allowances from a previous trading period and surrender them in future periods. On the other hand, if an entity believes future mitigation efforts will be cheaper in the future, borrowing allows the entity to surrender allowances in the current compliance period, which will be deducted from their future surrendered allowances.

2.1.2 Carbon Credit Mechanisms

Carbon credit mechanisms differ from ETSs because it concerns systems where voluntary emission reduction activities generate tradable credits (or offsets) (World Bank Group, 2023). Such activities can for example include direct air carbon capture or capturing methane from landfills. In addition, projects preventing deforestation can also yield carbon credits. Figure 1 shows the difference between the RCM (or compliance market) and the VCM. Both markets operate separately, but in certain ETSs it is possible to replace a limited amount of carbon allowances with carbon credits. An ETS is often a compliance market, but this does not have to be the case per se.

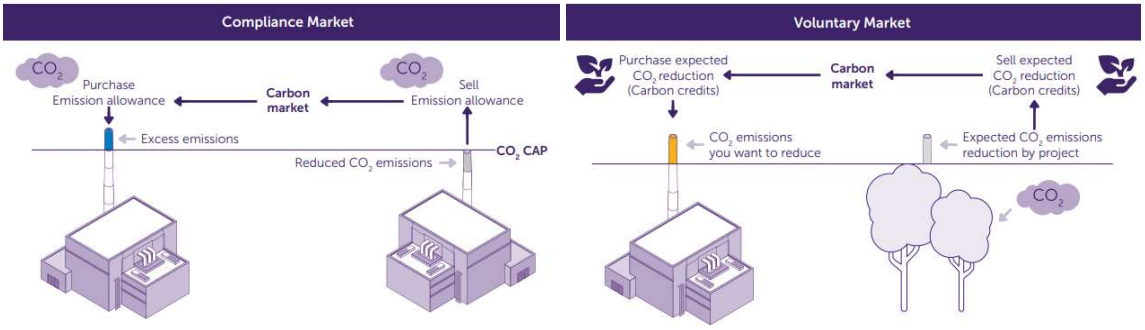

FIGURE 1 *Comparison between regulated (compliance) and voluntary market. Adapted from (Autoriteit Financiële Markten, 2023).*

Carbon credits are traded on exchanges, via brokers or intermediates, and OTC (Autoriteit Financiële Markten, 2023). Due to the limited degree of standardization, the OTC trade makes up the bulk of the VCM (Refinitiv, 2023). In the voluntary OTC carbon market, relatively straightforward spot and forward contracts are traded. Additionally, other OTC derivatives are traded as well, of which carbon swaps are an example. Offset-allowance swaps allow companies to exchange allowances and offsets. Brokers are often investment banks, although investment funds or speculators may also act as brokers (Autoriteit Financiële Markten, 2023).

Prices within the VCM vary widely influenced by vintage (age), quality, type of certification, type of carbon offsetting project, and the country where the project is executed , ranging from a few cents to USD 20 per megatonne of CO_2 equivalent (MtCO2e) (Autoriteit Financiële Markten, 2023). Nevertheless, the prices in the VCM are far below those in the European RCM, mainly caused by a surplus of credits build up in recent years (Ecosystem Marketplace, 2022).

In conclusion, there is no single carbon market. This term is not strictly defined and can refer to the worldwide carbon industry as a whole or to a specific geographical region (Zhu & Chevallier, 2017). Depending on the specific carbon market in which they operate, governments, non-governmental

organizations, and other public and private stakeholders trade allowances (RCM) and carbon credits (VCM). This may be driven by regulations, a voluntary desire to offset an organization's carbon footprint, or even speculative incentives.

2.2 SHELL'S CARBON PORTFOLIO

The trading portfolio of Shell is spread out over both the regulatory and voluntary markets, as shown in Table 1. For most of them, the market prices can be accessed publicly. However, for some markets, prices are not publicly available, and for confidentiality reasons these (prices of the) markets have been redacted. A summary of the regulatory characteristics of each of the markets with publicly reported prices is presented in Appendix A.2. For the prices of the New Zealand ETS we used a Python script published on GitHub to web scrape weekly prices reported on the website of CarbonNews (CarbonNews, 2024; Theecanmole, 2024).

Next to these regulatory differences, a few additional aspects can be identified. First each market may contain a range of market components. For example, within the European ETS the tradable components are split in the four regulatory compliance periods. Even more market components exist in voluntary markets where each carbon offset project gives rise to a new market component. Since the goal of this research is to develop an overarching PFE model, a discussion of the individual market components is out of scope. The second aspect relates to the liquidity and type of price information that is available. Table 1 shows that for five of the fifteen markets, only spot price information is available. Moreover, for four of the components the price information is extremely illiquid.

TABLE 1 *Availability of forward and spot prices for carbon markets.*

Market	Full name	Regulatory/voluntary	Liquid forward/spot
CMV	Carbon Market V	-	Illiquid spot
CMW	Carbon Market W	-	Illiquid spot
CMX	Carbon Market X	-	Spot
CMY	Carbon Market Y	-	Illiquid spot
CMZ	Carbon Market Z	-	Forward (& spot)
NZETS	New Zealand ETS	Regulatory	Spot (Theecanmole, 2024)
CCA	California Carbon allowances	Regulatory	Forward (& spot)
EUETS	European Union ETS	Regulatory	Forward (& spot)
GEO	Global Emissions Offsets	Voluntary	Forward (& spot)
ICEVER	Exchange traded VER	Voluntary	Forward (& spot)
N-GEO	Nature-Based Global Emissions Offsets	Voluntary	Forward (& spot)
RGGI	Regional Greenhouse Gas Initiative	Regulatory	Forward (& spot)
UKETS	United Kingdom ETS	Regulatory	Forward (& spot)
WCA	Washington Carbon Allowance	Regulatory	Forward (& spot)

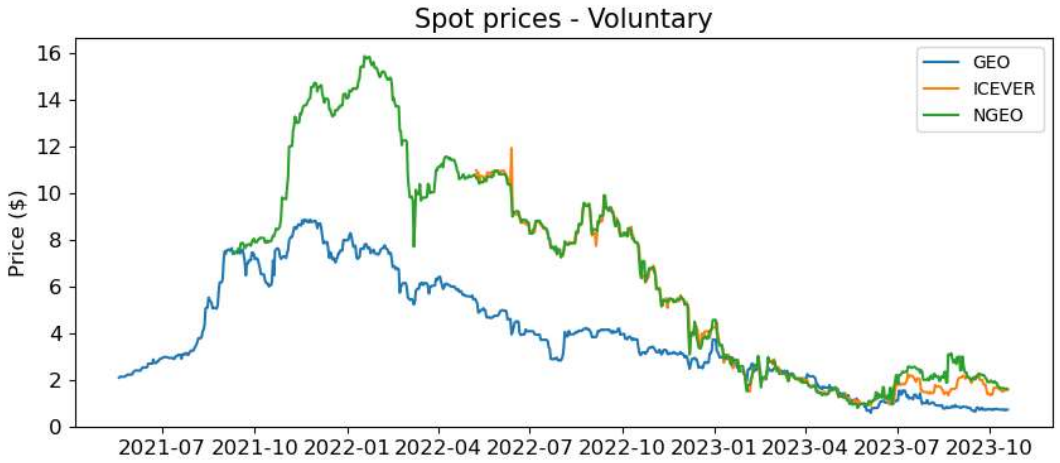

FIGURE 2 *Historical spot prices of liquid public regulatory carbon market components.*

Historical public spot prices are given for one market component in the RCM and VCM in Figures 2 and 3 respectively. In these figures, the illiquid markets are excluded and only one component per market is shown as there are many components which have (often) similar or equal prices. To prevent ambiguity, the specific component names are excluded. The figures clearly show the different price levels for the various markets. The pronounced correlation observed between ICEVER and NGEO prices is a result of their shared characteristics as carbon offset certificates, despite being traded on separate exchanges.

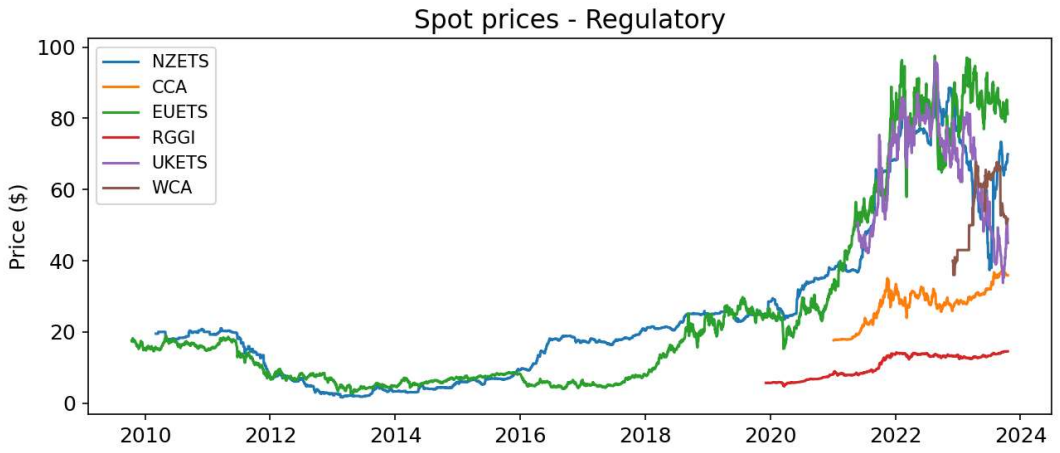

FIGURE 3 *Historical spot prices of liquid voluntary carbon market components.*

2.3 FINANCIAL RISK CLASSIFICATION

On a high level, financial risk can be divided into four main types: market risk, liquidity risk, operational risk, and credit risk (Gregory, 2012). Movement of market prices gives rise to market risk, which may result in changes in the value of a financial product. Quantification of market risk is often performed with the value-at risk (VAR) approach, for which multiple methodologies have been developed. The X% VAR indicates that one is X% certain that the loss will not be higher than this value (given the underlying modelling assumptions).

Liquidity risk relates to a company's inability to make cash payments when they become due and is split up into liquidity trading and liquidity funding risk (Hull, 2018b). The former concerns the ease at

which a company can unwind trading book positions, whereas the latter concerns the ease at which a company can fund contractual payments. The VAR metric can be adjusted to reflect the trading liquidity risk (Hull, 2018b).

Operational risk, on the other hand, is less strictly defined and harder to quantify. However, within the financial sector, numerous instances have shown that operational wrongdoing can lead to millions or even billions of losses (Hull, 2018b). The Basel Committee on Banking Supervision defines operation risk as "the risk of loss resulting from inadequate or failed internal processes, people, and system or from external events" (Basel Committee on Banking Supervision, 2021).

Lastly, credit risk arises due to the possible inability or reluctance of a counterparty to make payments or fulfil contractual obligations (Hull, 2018b). This is often referred to as a default (the exact definition depends on the jurisdiction), but also counterparty credit quality deterioration is considered to be credit risk.

Fundamentally, counterparty risk is classified as a subtype of credit risk with two main distinctions (Gregory, 2012). In relation to counterparty risk, the future value of a contract is often uncertain, whereas for other types of credit risk the amount at stake is generally known with a high degree of certainty. Furthermore, the value of a contract can both be positive or negative, meaning that counterparty is bilateral. Credit risk covers those instances where only one of two parties faces the risk of default. In other words, counterparty risk is a combination of credit risk and market risk (Gregory, 2012).

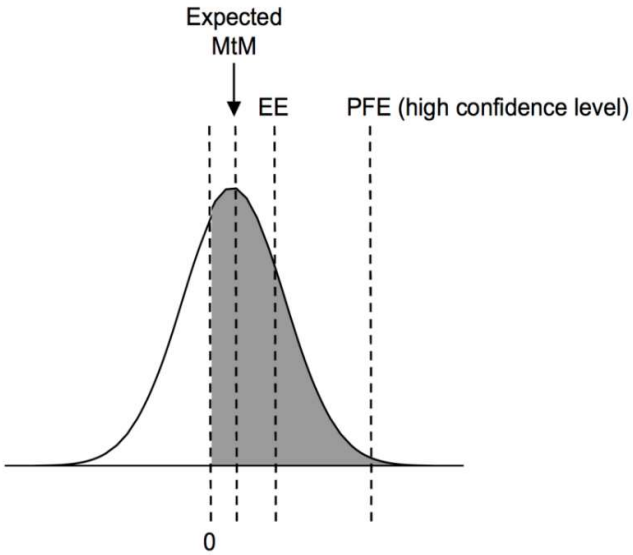

FIGURE 4 *Illustration of the exposure metrics EE and PFE (Gregory, 2012).*

For a long time, financial institutions' counterparty risk management was focused on trading with only the soundest parties (and sometimes even ignoring those risks for high-quality rated counterparties) (Gregory, 2012). However, the global financial crisis showed that even 'too big to fail' institutions can actually default.

A metric to quantify counterparty risk is Potential Future Expose (PFE), which indicates a potential worst-case positive credit exposure in the future (often years) under a certain confidence level (Gregory, 2012). This is illustrated in Figure 4, which shows a probability distribution of a contract's value. The Expected Mark-to-Market (MtM) indicates the expected value of this contract. Exposure at time t may be defined as $E(t) = Max(MtM(t),0)$. The expected exposure (EE) is therefore the expectation of positive MtM values (the shaded area). The PFE is subsequently the tail exposure of

the positive MtM values. Formalized, the α-percentile of the PFE at time t is the value $PFE(t)_\alpha$ such that:

$$P(E(t) \geq PFE(t)_\alpha) \leq 1 - \alpha. \tag{2.1}$$

A different representation is:

$$PFE = \int_0^\infty MtM(t)P(MtM)dt, \tag{2.2}$$

where $P(MtM)$ denotes the probability of MtM.

Various ways of mitigating/limiting counterparty risk exist. The most common ones are discussed below (Gregory, 2012):

- Netting: This allows positive contract values to be offset by negative contract values on a portfolio/counterparty level. Payment netting allows a company to net cash flows that occur on the same day. In the context of counterparty risk, closeout netting is more relevant, which allows the termination of all contracts between the solvent and defaulted counterparties and offsetting the transaction values.
- Collateralization: By demanding the posting of collateral, a collateral agreement limits the (netted) exposure. Only when the collateral giver defaults, becomes the collateral receiver economic owner of the collateral. This may lead to liquidity and operational risk since a collateral management function must be in place. Simultaneously, higher margin requirements increase the working capital of a company. To limit operational costs, collateral posting often occurs in blocks and with predefined minimum transfer amounts.
- Hedging: An institution can take on a position to offset the credit exposure to a certain counterparty. For example, Credit Default Swaps (CDSs) provide insurance for credit exposure in exchange for periodic payments (usually based on a reference rate plus an additional spread). Due to the volatility of the hedging instruments, this may give rise to additional market risk.

2.4 SETTLEMENT RISK

When a collateral call or final settlement is not followed by the associated payment, often there is a grace period before the counterparty is deemed to be officially in default (Gregory, 2012). Theoretically, counterparty risk usually relates to pre-settlement risk, i.e. the default risk before a contract expires. Yet, there is also a counterparty default risk during the settlement process. This settlement risk can arise due to timing differences in each party's contractual obligation during at final settlement. Although default prior to expiration is generally much more likely than at the settlement date, settlement risk can bring about a larger exposure (Gregory, 2012). In principle, all derivative contracts have both pre-settlement and settlement risk. However, the balance between them varies per contract. Long maturity contracts, for example, typically have mainly pre-settlement risk.

2.5 MODELLING POTENTIAL FUTURE EXPOSURE

This section covers an overview of the three main methods for quantifying exposure. At the core, this exercise involves balancing two effects (Gregory, 2012). The first is the increasing uncertainty about market variables, thereby increasing risk since one looks further into the future. Secondly, when financial instruments have cashflows spread over a certain period, this reduces the risk profile. In any case, calculating exposure requires a balance between sophistication and operational considerations such as computation time and flexibility of a model.

2.5.1 Add-on Method

The most straightforward methodology starts with the current exposure of a contract or portfolio and adds a component that should reflect the uncertainty of the PFE in the future (Gregory, 2012). The add-on must take into account the relevant time horizon and volatility of the underlying. This methodology formed the basis of the Basel I accords and is still used in the standardized approach for counterparty credit risk in the Basel III regulations. The first versions of the add-on method struggled to properly reflect the effects of netting and collateral. Although more mature add-on approaches do take into account these two effects (Basel Committee on Banking Supervision, 2014), it is still a general approach that cannot incorporate very specific aspects of contracts such as the timing of cash flows.

2.5.2 Parametric Method

The parametric method, also known as the semi-analytical method or model-building approach, involves adopting a (joint) probability distribution of relevant risk factor(s) (Hull, 2018b). This approach is suitable when the exposure is linearly dependent on the risk factors. The PFE can then be determined as illustrated in Figure 4. However, often this approach is based on simplified assumptions regarding the risk factors and more complicated distributions of the market variables cannot be included. Similar to the add-on method, it is not possible to consider contract particularities.

2.5.3 Monte Carlo Method

When the probability distribution of the risk factors is non-normal, contains correlations, and/or the financial product has a complicated structure, Monte Carlo simulation may be used (Gregory, 2012). Computationally, this method is more demanding, but it can cope with many complexities as netting, collateralisation, and transaction specifics. The process of modelling PFE with Monte Carlo simulation will vary depending on the complexity of the financial instruments, but below a general outline is given (Gregory, 2012):

1. *Understanding the financial contracts:* One should preferably start by thoroughly understanding the structure and characteristics of the portfolio. This involves identifying the types of derivatives, their underlying assets, relevant contractual terms, and the nature of their cash flows.
2. *Parameter estimation:* The next step is to gather relevant data for the underlying assets and relevant market variables (e.g., interest rates, volatilities, correlations). A model must be selected to simulate the risk factors, and its parameters must be estimated from the collected data.
3. *Scenario generation:* A large number of market scenarios should be created over the time horizon that is of interest by using the stochastic models.

4. *Valuation:* For each scenario, value the derivatives portfolio using the models and market inputs. This yields a distribution of potential portfolio values for a range of future time points.
5. *Collateral and netting agreements:* Consider the impact of netting agreements and collateral arrangements, which can reduce the potential future exposure.
6. *Estimating Potential Future Exposure:* Once you have the distribution of portfolio values, calculate the potential future exposure under a given confidence level for each time point by taking the maximum positive value for each scenario. This represents the peak potential loss under each market condition.
7. *Model Validation:* Validate the PFE model, including the mathematical models used, data inputs, and the assumptions made.

2.5.4 Parameter Estimation: Historical data versus Market Data

Pricing of financial instruments is generally done under the so-called risk-neutral measure, which is justified by hedging arguments (Hull, 2018b). Hence, the expected return required by individual investors is often not relevant for pricing. A direct result of this is that market-implied parameters such as drifts and volatilities need not correspond to real distributions. For risk management purposes, one should in general focus on the real-world measure (Gregory, 2012). Hence, for modelling PFE real-world probabilities and parameters should be used.

Modelling PFE encompasses the simulation of future market scenarios with historical parameters and the revaluation of trades at each date using risk-neutral valuation. In practice, however, sometimes risk-neutral parameters are used for determining the drift of certain underlying risk factors (e.g. when there is a lack of historical data) (Gregory, 2012). In contradiction to a VaR analysis where the risk horizon is relatively short (10 days for example), the drift is important while assessing credit exposure, due to the much longer time horizon (Gregory, 2012). The influence of volatility follows a square root of time scaling, whereas the scaling is more linear for the drift. Over a long time horizon, the effect of the drift will be dominant.

2.6 SPOT, FUTURES & FORWARD PRICES

In the world of commodities and financial markets, spot, futures, and forward prices are fundamental concepts. The spot price refers to the market price of an asset that is bought or sold with immediate delivery, whereas the futures and forward prices are agreed-upon prices for delivery of that asset at a future date. Theoretically, the futures and forward prices are different. Futures, standardized regarding delivery specifications, amounts and quality (grade), are traded on exchanges and often deemed to be liquid. In addition, futures contracts are generally settled daily, meaning potential gains and losses on a futures position are settled against the margin account that the company must hold at an exchange or broker. Forward contracts are traded OTC and are often tailored to suit a company's interests.

Classically, futures (or forward) prices can be seen as the intersection of expected future supply and demand. Yet it is not uncommon for forward and spot rates to be far apart in the market (Gregory, 2012). In the case of commodity prices, seasonal effects and storage costs can cause deviations between spot and futures prices. One of the many reasons for future prices to be higher than spot prices (contango) is when demand is low and inventory levels are high (Hull, 2018b). Likewise, future prices can be lower than spot prices (backwardation) when the demand is high, and inventories are low.

The difference between a futures price and the spot price of a commodity or asset is known as the basis. Typically, as the expiration date of the futures contract approaches, the futures price will converge with the spot price, as illustrated in Figure 5. This phenomenon is due to the impending physical delivery or settlement. If a large discrepancy existed between the spot and futures prices at expiry, it would create opportunities for arbitrage – a trader could simultaneously buy low and sell high to lock in risk-free profits. In practice, the spot price of a commodity may not always be directly observable, for example when it is determined the less liquid OTC market.

Based on many empirical studies across various markets on the relationship between spot and futures prices, the general belief is that the futures price is a biased estimation of the future spot price, which is a contradiction to the efficient market hypothesis (Gregory, 2012). A paper by the International Monetary Fund (IMF) showed that for ten commodities, futures-based forecasts are hard to beat and, for most horizons, perform at least as well as the geometric Brownian motion (GBM) (Reichsfeld & Roache, 2011). Still, the futures price does not statistically significantly outperform the GBM benchmark. Furthermore, the forecasting ability of futures falls when increasing the forecast horizon, likely a result of lower liquidity in the back end of the futures curve.

Where there is a relatively straightforward formula to determine a futures price from the spot price for futures on stocks, this is different in commodity markets. In these markets, the shape of the forward curve reflects the perception of the market participants regarding current and future developments (Borovkova & Geman, 2009). For the carbon and EAC markets, this includes political and regulatory changes (Palao & Pardo, 2021). From an arbitrage point of view, the futures price can be calculated with the following cost-of-carry relationship (Borovkova & Geman, 2009)

In what follows, the terminology of futures and forwards will be used interchangeably, as Shell's trading portfolio contains both types of contracts. In the context of PFE, both the spot price as well as the forward curve are of importance. Imagine a contract in which Company A promises to deliver one carbon allowance to company B in six months and twelve months. The MtM value of this contract for company B directly after origination can be determined without any uncertainty and is the sum of the 6-month forward and 12-month forward price of the carbon allowance. PFE modelling requires upfront modelling the exposure (and thus the MtM value of the contract) over the duration of the contract. When the MtM value would be determined again after 6 months, this requires information on the 6-monts forward price at that point in time.

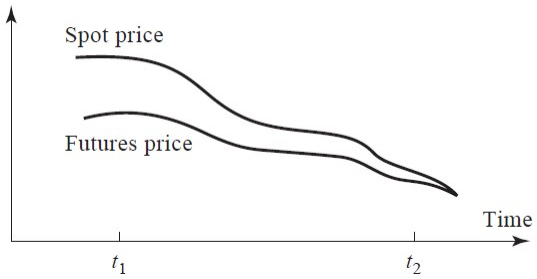

FIGURE 5 *Fluctuation of basis over time (Hull, 2018a).*

3 RISK FACTOR MODELLING

Building upon the overview of the carbon markets and the diverse PFE modelling approaches introduced in the preceding chapter for modeling the requisite risk factors, in this chapter we will progress by first selecting the most appropriate approach. Subsequently, we delve into an examination of the risk factors inherent in carbon markets. Following this analysis, we decide which risk factors will be incorporated into the modeling approach. Finally, we review the theoretical underpinnings of how these risk factors can be effectively modeled.

3.1 MODEL APPROACH SELECTION

The previous chapter described three primary approaches for modelling PFE: the add-on method, the parametric method, and Monte Carlo (MC) simulation. The add-on method is deemed unsuitable due to its over-simplifications and its inability to incorporate considerations related to the timing of cash flows. A similar argument can be posited for the parametric method, as it is unable to account for the intricacies of deal structures and anticipated optionalities. Consequently, the remaining viable alternative is the utilization of Monte Carlo simulation to model the risk factors.

3.2 CARBON MARKET RISK FACTORS

Numerous studies have shown that energy prices, economic activity, temperature conditions, and external factors are the main drivers of carbon prices in the RCM (Alberola et al., 2008; Lovcha et al., 2022; Zhu & Chevallier, 2017). Given the long history of the EU ETS, most of the studies investigating the price drivers focussed on this market. A large part of the worldwide carbon dioxide emissions stem from electricity and heat generation (44% in 2021, with coal plants having a share of 73% of these emissions) (International Energy Agency, 2021). Because power companies can switch between different fossil fuels (coal, natural gas, and oil), an internal price transmission mechanism is induced between the fossil power and carbon markets.

This can be explained by the so-called dark and spark spreads (Alberola et al., 2008). The dark-spread is the profit a coal-powered company makes by selling electricity (taking into account the cost of buying the fuel). The spark spread is the equivalent of a gas-powered plant. When the cost of carbon allowances is included in this calculation the dark and spark spreads must be adjusted and are then called the clean dark and clean spark spread. The equilibrium carbon allowance price refers to the price at which both spreads are equal. Since burning natural gas emits less carbon dioxide than burning coal, above this equilibrium price it is advantageous to switch from coal to natural gas. When the gas price increases relative to the coal price, the equilibrium carbon price will increase too. This leads to a higher demand for coal and thereby for carbon allowances, which drives up the prices of these allowances (Mansanet Bataller et al., 2007).

Despite this energy-carbon price interaction, Batten et al. found that oil, natural gas, coal and electricity prices and the clean/dark spreads accounted for only 11% of the fluctuations in the EUA spot prices (Batten et al., 2021). Hence, there are clearly other factors that influence prices on the carbon markets. One of these additional factors is temperature (Zhu & Chevallier, 2017). Cold winters result in a higher demand for heat and electricity and combined with lower hydroelectricity production in dry periods, this can lead to higher carbon prices. More specifically, not the absolute temperature levels, but rather the unpredicted weather changes appear to influence carbon prices (Batten et al., 2021).

In addition, carbon prices are found to be also subjected to economic activities (Seifert et al., 2008; Zhu & Chevallier, 2017). Higher levels of economic activity generally result in more carbon dioxide

emissions. With a higher demand for carbon allowances and/or carbon credits, prices tend to increase. The reverse effect was observed during the 2007-2008 global financial crisis when the EUA spot price decreased from EUR 20/MtCO2e to EUR 15/MtCO2e.

As a policy-based market, the RCM is greatly influenced by the market mechanism and external factors such as global climate negotiations and regulatory changes (Zhu & Chevallier, 2017). For example, in the last week of April 2006, it became clear that several EU countries had allocated more free allowances compared to what they should have based on actual (or verified) emissions. This led to a drop of 54% in the EUA spot price, which was followed by an additional drop when the European Commission confirmed that the total number of allocated allowances in the EU was 4% too high (Zhu & Chevallier, 2017). Moreover, the stringency and changes in emission targets, changes in allowance allocation methods, or other regulatory adjustments can affect carbon prices (Bredin & Muckley, 2011).

The influence of the above-mentioned price drivers has been studied mainly in the context of the RCM. We expect similar price drivers to affect the VCM to some extent as well. However, a lack of transparency in the VCM, due to a lack of transparency and product uniformity, makes a systematic study difficult (SIX, 2022). Furthermore, public and media scrutiny seem to play a role in price fluctuations as well. Next to macroeconomic and external factors, this has discouraged corporations from purchasing carbon credits over the first half of 2023 (CarbonCredits.com, 2023).

3.3 RISK FACTOR SELECTION

Next, we need to select the risk factors that will eventually be used for modelling PFE. As elucidated in the preceding section, a multitude of risk factors are documented in existing literature, and the potential for discovering additional factors is considerable. This thesis' objective is to formulate a comprehensive PFE model, capable of accommodating the intricacies of most carbon deals. This objective rules out the option to model individual risk factors such as temperature, dark/spark spreads, and correlations with other commodities, as their applicability is contingent upon the unique characteristics of each specific carbon market. These idiosyncrasies arise from the varying geographical locations and regulatory frameworks governing each market. Moreover, most of the research in literature is focussed on (more) developed RCMs such as the EU ETS and most of the other RCMs in Table 1 **Fout! Verwijzingsbron niet gevonden.**are not covered in those studies. Not to mention the VCMs which, due to their opacity, have received even less scientific attention. Consequently, we decide to focus exclusively on modelling the pricing dynamics of carbon allowances and credits as a direct risk factor.

This decision still leaves open the question which price will be modelled; the spot price or the forward price. The answer to this question differs per market component and partially depends on the source of the available forward curve. In certain markets, the forward curve is published by external parties, whereas for others, only the spot price is available. In the latter case, the forward curve must be derived from the spot price.

3.4 SPOT PRICE MODELLING

This section provides a review of four models suitable for the modelling of spot prices that can be used within the PFE model. In the next chapter these models will be fitted on the carbon markets price data. Among the models, the Geometric Brownian Motion stands out as the most conventional and widely employed stochastic model for simulating price data. Furthermore, attention is directed towards the Ornstein-Uhlenbeck and Exponential Ornstein-Uhlenbeck models, as they account for mean reversion—a phenomenon observed in specific commodity markets and potentially applicable

to carbon markets. Finally, the inclusion of the Merton Jump Diffusion Model aims to explore whether the spot prices in carbon markets are best characterized by a discontinuous model.

3.4.1 Geometric Brownian Motion

The most standard and widely used continuous stochastic model is the geometric Brownian motion (GBM). The stochastic differential equation (SDE) for a GBM is given below (Roncoroni et al., 2015):

$$dx_t = \mu x_t dt + \sigma x_t dW_t. \tag{3.1}$$

Here, μ is the drift of x_t, σ is the volatility of x_t, and W_t is a standard wiener process. The distribution of $\ln\left(\frac{X(t)}{X(0)}\right)$ (log return) has the property:

$$\ln\left(\frac{X(t)}{X(0)}\right) \sim N\left(\left(\mu - \frac{\sigma^2}{2}\right)t, \sigma^2 t\right). \tag{3.2}$$

The solution of the above SDE can be found with Ito's lemma (see Appendix B.1) and is given by Equation (8.3) in Appendix B.2. A discrete version of the this solution is given by (Nkemnole & Abass, 2019):

$$X(t) = X(t-\Delta t)\exp\left(\left(\mu - \frac{\sigma^2}{2}\right)\Delta t + \sigma N[0,1]\sqrt{\Delta t}\right). \tag{3.3}$$

One of the common methods to estimate parameters for a given probability distribution based on observed data is the maximum likelihood estimation. This method makes use of the likelihood function, which indicates the support provided by the observed data for the values of the unknown parameters in the probability distribution. This function is computed as the product of the probability density functions for individual observations (Myung, 2003). Maximizing this function may give a solution for the unknown parameters, which in theory make the observed data most probable. These so called maximum likelihood estimators (MLEs) estimated may not be unique, nor will they always exist at all (Myung, 2003).

Differentiating the likelihood function for the GBM (which is derived from the normal distribution probability density function) with respect to μ and σ and setting them equal to zero gives the following expressions (Hull, 2018a):

$$\mu = \frac{1}{n}\sum_{i=1}^{n} \ln\left(\frac{X_i}{X_{i-1}}\right). \tag{3.4}$$

$$\sigma = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(\ln\left(\frac{X_i}{X_{i-1}}\right) - \mu\right)^2}. \tag{3.5}$$

These are the standard maximum likelihood estimators for the mean and the variance of the log price returns.

3.4.2 Ornstein-Uhlenbeck Model

Empirical studies have shown that some economic variables, such as certain commodity prices and inflation rates, have a tendency to be pulled back to a long-run average level over time (Hull, 2018b). This is the phenomenon of mean reversion. The further away from the equilibrium level, the greater the tendency to move towards it. The Ornstein Uhlenbeck (OU) model captures this behaviour. This model is also known as geometric mean reversion (GMR) and is equal to the Vasicek model, other than that in the OU model μ is assumed to be deterministic. In theory, it is also possible to include a time-dependent equilibrium level to incorporate seasonal effects. The SDE is given by (Roncoroni et al., 2015):

$$dx_t = \kappa(\mu - x_t)dt + \sigma dW_t. \tag{3.6}$$

Here, κ is the rate at which x_t converges to μ. The distribution of $X(t)$ can be written as:

$$X(t) \sim N\left(\exp(-\kappa t)X(0) + \mu(1-\exp(-\kappa t)), \frac{\sigma^2}{2\kappa}(1-\exp(-2\kappa t))\right). \tag{3.7}$$

The solution of the OU SDE is given in Appendix B.3, for which the discrete version is given by (Roncoroni et al., 2015):

$$X(t) = \exp(-\kappa t)X(t-\Delta t) + \mu(1-\exp(-\kappa t)) + \sigma N[0,\xi^2], \tag{3.8}$$

with $\xi = \sqrt{\frac{1}{2\kappa}(1-\exp(-2\kappa \Delta t))}$.

The log-likelihood function and the subsequent expressions for the MLEs for the OU process can be found in Equations (8.6)-(8.9) in Appendix B.3

3.4.3 Exponential Ornstein-Uhlenbeck Model

The main limitation of the OU model is that the solutions can be negative. The exponential OU (eOU) model circumvents the possibility of negative process and follows the following SDE:

$$dx_t = \kappa(\mu - \log(x_t))x_t dt + \sigma x_t dW_t. \tag{3.9}$$

The distribution of $X(t)$ can be expressed as (Mejía Vega, 2018):

$$\ln(X(t)) \sim N\left(\exp(-\kappa t)\ln(X(0)) + \left(\mu - \frac{\sigma^2}{2\kappa}\right)(1-\exp(-\kappa t)), \frac{\sigma^2}{2\kappa}(1-\exp(-2\kappa t))\right). \tag{3.10}$$

The OU solution is given in Appendix B.4 and can be approximated with the following equation (Mejía Vega, 2018):

$$\ln(X(t)) = \exp(-\kappa t)\ln(X(t-\Delta t)) + \left(\mu - \frac{\sigma^2}{2\kappa}\right)(1-\exp(-\kappa t)) + \sigma N[0,\xi^2], \tag{3.11}$$

with $\xi = \sqrt{\frac{1}{2\kappa}(1-\exp(2\kappa \Delta t))}$.

In Appendix B.4 the log-likelihood function is given along with the MLEs for the associated parameters.

3.4.4 Merton Jump Diffusion Process

For commodities that exhibit price jumps (e.g., electricity and natural gas) a jump diffusion process is sometimes used to model the behaviour. A well-known example is the Merton Jump Diffusion (MJD) process for which the SDE is given in Equation (3.12) (Gugole, 2016). The jumps in the MJD have two main characteristics: the moment of arrival and the size. When time at which the jump occurs is modelled according to a Poisson distribution and the size is assumed to be log normal. The $-\lambda k$ term does not indicate mean reversion, but this is a term to offset the drift introduced by the jump term.

$$dx_t = (\mu - \lambda k)x_t dt + \sigma x_t dW_t + x_t dQ_t. \tag{3.12}$$

Here, Q_t is a compound Poisson process in the form of:

$$Q_t = \sum_{i=1}^{N_t}(Y_i - 1). \tag{3.13}$$

$N(t)$ is a Poisson process with probability of k jumps occurring equal to $\mathbb{P}(N(t)=k) = \frac{(\lambda t)^k \exp(-\lambda t)}{k!}$, λ is the expected number of jumps per year (intensity). Y_i is a log-normally distributed random variable:

$$\ln(Y_i) \sim N(\mu_y, \sigma_y^2). \tag{3.14}$$

which means that:

$$\mathbb{E}[Q_t] = \lambda k t. \tag{3.15}$$

$$k = \mathbb{E}[Y_i - 1] = \exp\left(\mu_y + \frac{\sigma_y^2}{2}\right) - 1. \tag{3.16}$$

Therefore, Equation (3.12) can also be written in the form of (Gugole, 2016):

$$dx_t = (\mu - \lambda k)x_t dt + \sigma x_t dW_t + x_t\left(\prod_{j=1}^{N_t} Y_i - 1\right). \tag{3.17}$$

The distribution of $\log\left(\frac{X(T)}{X(0)}\right)$ is:

$$\log\left(\frac{X(t)}{X(0)}\right) \sim N\left(\left(\mu - \frac{\sigma^2}{2} - \lambda k + \lambda\mu_y\right)t, \left(\sigma^2 + \lambda\sigma_y^2 + \lambda\mu_y^2\right)t\right). \tag{3.18}$$

The explicit solution, conditional on a given value of N_t, is as follows (Gugole, 2016):

$$X(t) = X(0)\exp\left(\left(\mu - \frac{\sigma^2}{2} - \lambda k\right)t + \sqrt{\sigma^2 + \lambda\sigma_y^2 + \lambda\mu_y^2}N[0,1]\right). \tag{3.19}$$

3.5 FORWARD CURVE MODELLING

On a high level, the different Monte Carlo (MC) methodologies for the modeling of forward curves may be categorized as follows: (1) perturbing the forward curve by utilizing volatilities associated with specific maturities, characterised by an N-distribution; (2) employing principal component analysis; and (3) utilizing stochastic factor models. These modeling techniques will be expounded upon in the forthcoming section.

It is worth noting that within the ambit of the stochastic factor model, an array of alternatives is at one's disposal. To start, one may consider a one-factor, two-factor, or three-factor model, with each factor being a stochastic variable. Empirical findings suggest that, overall, two-factor models exhibit superior performance compared to their one-factor counterparts (Lautier, 2003). However, the same level of clarity is not manifest in the comparison between two-factor and three-factor models (Lautier, 2003). Consequently, we decide to adopt a two-factor model, taking into account that a three-factor model entails a more intricate parameter estimation and simulation process.

Within the realm of two-factor models, various proposed models can be found in the existing literature, designed for the simulation of (commodity) term structures. The Schwartz-Smith model stands out as one of the most frequently utilized models for the characterization of commodity forward curves (Lautier, 2003). Therefore, we elucidated this model in this section and subject it to a comparative analysis with other modeling techniques within the purview of this thesis.

3.5.1 Futures Curve Shocking with (N-)Distribution

A direct approach to shock the futures curve involves calculating the volatility of the curve at each maturity using historical futures curve data. The conventional choice would be to perturb the forward prices with a normal distribution, corresponding to the Geometric Brownian Motion (GBM).

One of the downsides of this approach is that the futures curve will not be kept intact when each of the maturities is shocked separately. Furthermore, the implicit normality assumption that underlies this approach may prove to have a limited validity in case historical market returns exhibit non-

normal behaviour. In principle, a different distribution (e.g. Student's T or Laplace) can be used for shocking the curve to replicate this behaviour.

3.5.2 Principle Component Analysis

An alternative to shocking the futures curve with an SDE is the technique called principal component analysis (PCA). With this technique a random vector (e.g. the daily returns the futures price with a X-month maturity), which is dependent on various correlated components, can be transformed into a random vector with uncorrelated components (Kreinin et al., 1998). An instance where PCA is often applied is for calculating potential changes of a portfolio's value over time as each individual position/contract may be dependent on correlated risk factors. Modelling all those risk factors and their dependencies often proves to be a cumbersome task on a portfolio level. Applying PCA yields a new set of uncorrelated variables that are a linear combination of the returns of the risk factors. Those uncorrelated variables are the principal components.

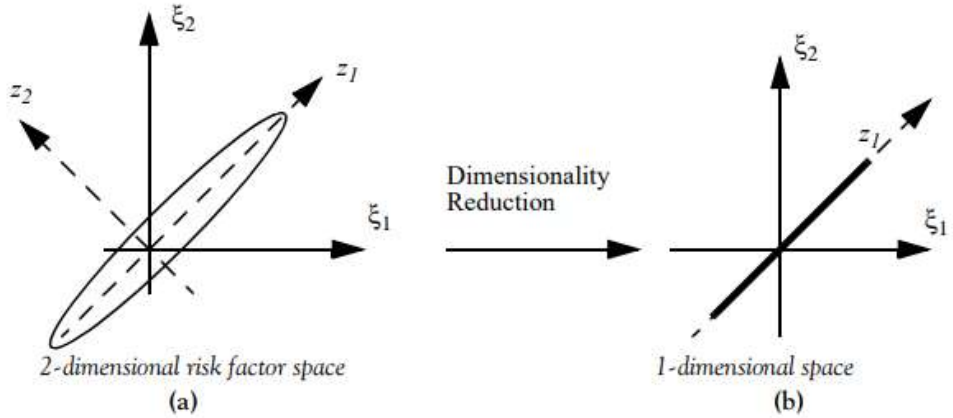

FIGURE 6 *PCA-based dimensionality reduction (Kreinin et al., 1998).*

In Figure 6 PCA is illustrated for a two-dimensional case. The ellipse in Figure 6a corresponds to the joint distribution of two risk factor returns: ζ_1 and ζ_2. These returns are positively correlated, whereas the principal components z_1 and z_2, which can be thought of as two random variables with perpendicular axes, are uncorrelated. In the figure, z_1 contributes more to the overall variance of the joint distribution. A reduced version of the joint distribution is therefore captured only by principal component z_1, as shown in Figure 6b.

To illustrate how PCA can be used to model the futures curve stochastically, suppose the vector $\boldsymbol{R}$ is composed of n historic price returns of a X-month maturity futures price:

$$\boldsymbol{R} = \begin{pmatrix} r_1 \\ r_2 \\ \vdots \\ r_n \end{pmatrix}. \tag{3.20}$$

As a measure of the joint variability of these asset returns, the covariance matrix is given by:

$$\mathrm{Cov}(R) = \begin{pmatrix} Var(r_1) & Cov(r_1, r_2) & \cdots & Cov(r_1, r_n) \\ Cov(r_2, r_1) & Var(r) & \cdots & Cov(r_2, r_n) \\ \vdots & \vdots & \ddots & \vdots \\ Cov(r_n, r_1) & Cov(r_n, r_2) & \cdots & Var(r_n) \end{pmatrix}. \tag{3.21}$$

Dividing each entry of the covariance matrix by the product of the standard deviation of the two corresponding returns yields the correlation matrix:

$$\mathrm{Corr}(\boldsymbol{R}) = \begin{pmatrix} 1 & \rho_{r_1,r_2} & \cdots & \rho_{r_1,r_n} \\ \rho_{r_2,r_1} & 1 & \cdots & \rho_{r_2,r_n} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{r_n,r_1} & C\rho_{r_n,r_2} & \cdots & 1 \end{pmatrix}. \tag{3.22}$$

Here, $\rho_{X,Y} = \frac{cov(X,Y)}{\sigma_X \sigma_Y} = \frac{\mathbb{E}[(X-\mu_X)(Y-\mu_Y)]}{\sigma_X \sigma_Y}$ is the corelation coefficient between variables X and Y. The next step is to determine the orthonormal eigenvectors U_j, which follow the following condition (Kreinin et al., 1998):

$$\mathrm{Cov}(\boldsymbol{R})\boldsymbol{U_j} = \lambda_j \boldsymbol{U_j}. \tag{3.23}$$

for $\lambda_j > 0$, $j = 1,2, \dots, n$, being the eigenvalues.

The eigenvectors can be interpreted as the modes of fluctuation of the random vector v. The principal components are then defined such that v can be written as (Kreinin et al., 1998):

$$v = \eta_1 \sqrt{\lambda_1} U_1 + \eta_2 \sqrt{\lambda_2} U_2 + \cdots + \eta_n \sqrt{\lambda_n} U_n, \tag{3.24}$$

where $z_j = \sqrt{\lambda_j}\eta_j$, $j = 1, \dots, n$ are the principal components.

Equation (3.24) can be used in combination with Monte Carlo simulations, as a single scenario can be generated by drawing independent standard normal variables in the vector $\eta = (\eta_1, \eta_2, \dots, \eta_n)$. The first few principal components generally make up a large part of the total variance (of v in this example) and hence a selection often suffices for Monte Carlo simulations (Kreinin et al., 1998).

In addition to one futures price with a certain maturity, PCA can be applied to the complete futures curve (Blanco, 2002). For each maturity, the principal components can be determined, which may be shocked in accordance with their standard deviation. Consider the data set $F(t_i, T_j)$ with $i = 1 \dots n$ and $j = 1 \dots m$. Define the returns of futures prices at time t with maturity T as follows:

$$r_{ij} = \frac{F(t_i, T_j) - F(t_{i-1}, T_j)}{F(t_{i-1}, T_j)}. \tag{3.25}$$

PCA must then be applied on the return matrix $\boldsymbol{R}$:

$$\boldsymbol{R} = \begin{pmatrix} r_{11} & \cdots & r_{1m} \\ \vdots & \ddots & \vdots \\ r_{n1} & \cdots & r_{nm} \end{pmatrix}. \tag{3.26}$$

Using the eigenvalues and eigenvectors given by PCA on the return matrix, the change in a futures price with a certain maturity follows Equation (3.27) (Roncoroni et al., 2015). MC simulation can subsequently be performed to generate shocked forward curves.

$$\frac{dF(t,T)}{F(t,T)} = \sum_{k=1}^{K} \sigma_k dW = \sum_{k=1}^{K} U_k \sqrt{\lambda_j} dW. \tag{3.27}$$

Here, σ_k represents the contribution of the kth principal component to the total variance. For simulation, the discrete version of Equation (3.27) is given by:

$$F(t + \Delta t, T) = F(t,T) * \exp\left(\sum_{k=1}^{K} \sigma_k \sqrt{\Delta t}\, N[0,1] - \frac{1}{2}\sigma_k^2 \Delta t \right). \tag{3.28}$$

To determine which K principal components are relevant to consider, one can calculate the contribution of the first K factors to the total variance with Equation (3.29).

$$\text{contribution} = \frac{\sum_{k=1}^{K} \lambda_k}{\sum_{j=1}^{m} \lambda_j}. \tag{3.29}$$

Theoretically, any forward curve shape can be constructed by changing three aspects: the level, slope and curvature. Research has shown that for non-seasonal commodities, the first three principal components of the historical futures price returns exactly correspond to these three main aspects (Borovkova & Geman, 2009). Hence the first principal component can be referred to as a parallel shift of the futures curve, the second to the change in the slope, and the third to a change in curvature. In practice, the first three components often make up for most of the observed variance (Borovkova & Geman, 2009). Utilizing the results of PCA for stochastic simulation of a forward curve may involve shocking the principal components according to a normal distribution.

3.5.3 Schwartz-Smith model

In the Schwartz-Smith model, the spot price of the asset is modelled with two stochastic components as given in Equations (3.30)-(3.32) (Schwartz & Smith, 2000). Here, χ_t is the short-term deviation in prices, which is assumed to follow an OU process, which reverts towards zero. The long-term equilibrium price, ζ_t, is assumed to follow a GBM. The Wiener processes are assumed to be correlated according to $dW_\chi dW_\zeta = \rho_{\chi\zeta} dt$

$$\ln(S_t) = \chi_t + \zeta_t, \tag{3.30}$$

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dW_\chi, \tag{3.31}$$

$$d_{\zeta_t} = \mu_\zeta dt + \sigma_\zeta dW_\zeta. \tag{3.32}$$

The model allows for variations in the spot price via χ_t, which would not be expected to remain in the long run. In addition, it stipulates how these short-term deviations from the equilibrium price are expected to disappear. The presence of the long-term equilibrium price, which allows for long-term changes in spot prices, distinguish the S&S model from a pure mean-reversion model. In other words, the model gives room for mean-reversion behaviour in the short term with changes of, and uncertainty in, the equilibrium price.

The processes for χ_t and ζ_t are jointly normally distributed with the expected values and covariances as given in Equation (3.33) and (3.34) respectively (Schwartz & Smith, 2000).

$$\mathrm{E}[(\chi_t, \zeta_t)] = \left[\exp(-\kappa t)\,\chi_0, \zeta_0 + \mu_\zeta t\right], \tag{3.33}$$

$$\mathrm{Cov}[(\chi_t, \zeta_t)] = \begin{bmatrix} (1-\exp(-2\kappa t))\dfrac{\sigma_\chi^2}{2\kappa} & (1-\exp(-\kappa t))\dfrac{\rho_{\chi\zeta}\sigma_\chi\sigma_\zeta}{\kappa} \\ (1-\exp(-\kappa t))\dfrac{\rho_{\chi\zeta}\sigma_\chi\sigma_\zeta}{\kappa} & \sigma_\zeta^2 t \end{bmatrix}. \tag{3.34}$$

From Equation (3.30) the expected value and variance for the log of the spot price can be determined (Schwartz & Smith, 2000):

$$\mathrm{E}[\ln(S_t)] = \exp(-\kappa t)\,\chi_0 + \zeta_0 + \mu_\zeta t, \tag{3.35}$$

$$\mathrm{Var}[\ln(S_t)] = (1-\exp(-2\kappa t))\frac{\sigma_\chi^2}{2\kappa} + \sigma_\zeta^2 t + 2(1-\exp(-\kappa t))\frac{\rho_{\chi\zeta}\sigma_\chi\sigma_\zeta}{\kappa}. \tag{3.36}$$

Before this spot price model can be used to value futures contracts, risk-neutral versions of Equation (3.31) and (3.32) must be given. In Equation (3.37) and (3.38) below the additional parameters $\lambda_\chi$and λ_ζ, are the short-term and long-term risk premiums. Now, the short-term deviations revert towards $-\frac{\lambda_\chi}{\kappa}$ (instead of zero) and the equilibrium price has a drift rate of $\mu_\zeta^* = \mu_\zeta - \lambda_\zeta$. In this risk-neutral process, the joint normal distribution can be characterized with Equation (3.39) and (3.40) (Schwartz & Smith, 2000).

$$d\chi_t = (-\kappa - \lambda_\chi)\chi_t dt + \sigma_\chi dW_\chi^*, \tag{3.37}$$

$$d_{\zeta_t} = (\mu_\zeta - \lambda_\zeta)dt + \sigma_\zeta dW_\zeta^*, \tag{3.38}$$

$$\mathrm{E}^*[(\chi_t, \zeta_t)] = \left[\exp(-\kappa t)\,\chi_0 - (1 - \exp(-\kappa t))\frac{\lambda_\chi}{\kappa}, \zeta_0 + \mu_\zeta^* t\right], \tag{3.39}$$

$$\mathrm{Cov}^*[(\chi_t, \zeta_t)] = \mathrm{Cov}[(\chi_t, \zeta_t)]. \tag{3.40}$$

The logarithm of the spot price in the risk-neutral version is normally distributed with (Schwartz & Smith, 2000):

$$\mathrm{E}^*[\ln(S_t)] = \exp(-\kappa t)\,\chi_0 - (1 - \exp(-\kappa t))\frac{\lambda_\chi}{\kappa} + \zeta_0 + \mu_\zeta^* t, \tag{3.41}$$

$$\mathrm{Var}^*[\ln(S_t)] = \mathrm{Var}[\ln(S_t)]. \tag{3.42}$$

Theoretically, from a non-arbitrage point of view, the futures price is equal to the risk-neutral expected spot price. From the properties of a lognormal distribution, the relationship between the spot price and futures price is determined by (Schwartz & Smith, 2000):

$$\begin{aligned}\ln\big(F(0,T)\big) &= \ln(\mathrm{E}^*[S_t]) = \mathrm{E}^*[\ln(S_t)] + \frac{1}{2}\mathrm{Var}^*[\ln(S_t)] \\ &= \exp(-\kappa T)\,\chi_0 + \zeta_0 + A(T),\end{aligned} \tag{3.43}$$

$$\begin{aligned}A(T) = \mu_\zeta^* T - (1 - \exp(-\kappa T))\frac{\lambda_\chi}{\kappa} \\ + \frac{1}{2}\left((1 - \exp(-2\kappa T))\frac{\sigma_\chi^2}{2\kappa} + \sigma_\zeta^2 T + 2(1 - \exp(-\kappa T))\frac{\rho_{\zeta\chi}\sigma_\chi\sigma_\zeta}{\kappa}\right).\end{aligned} \tag{3.44}$$

The parameter for this model can also be estimated from historic spot and futures prices by using a technique called Kalman filtering. How this Kalman filter works is explained in APPENDIX C. In any case, this approach requires the model equations to be written in the state-space form. Considering χ_t and ζ_t as state variables, the transition equation and measurement equation for the S&S model can be written as Equation (3.45) and (3.46) respectively (Schwartz & Smith, 2000).

$$\boldsymbol{X}_t = \boldsymbol{c} + \boldsymbol{G}_t\boldsymbol{X}_{t-1} + \boldsymbol{\omega}_t, \tag{3.45}$$

$$\boldsymbol{Y}_t = \boldsymbol{d}_t + \boldsymbol{F}_t\boldsymbol{X}_{t-1} + \boldsymbol{v}_t. \tag{3.46}$$

Here:

$\boldsymbol{X}_t = [\chi_t, \zeta_t]'$, a 2 × 1 vector with state variables

$\boldsymbol{Y}_t = \left[\ln(F_{T_1}) \cdots \ln(F_{T_n})\right]'$, a n × 1 vector of observed log futures prices with maturities $T_1 \ldots T_n$

$\boldsymbol{c} = [0, \mu_\zeta \Delta t]'$, a 2 × 1 vector

$\boldsymbol{d}_t = [A(T_1) \cdots A(T_n)]'$, a n × 1 vector

$\boldsymbol{G} = \begin{bmatrix}\exp(-\kappa\Delta t) & 0 \\ 0 & 1\end{bmatrix}$, a 2 × 2 matrix

$$\boldsymbol{F}_t = \begin{bmatrix}\exp(-\kappa T_1) & 1 \\ \vdots & \vdots \\ \exp(-\kappa T_n) & 1\end{bmatrix}$$

$\boldsymbol{\omega}_t$ is a 2 × 1 vector with disturbances with $\mathrm{E}[\boldsymbol{\omega}_t] = [0,0]'$ and $\mathrm{Var}[\boldsymbol{\omega}_t] = \boldsymbol{W} = \mathrm{Cov}[(\chi_t, \zeta_t)]$

$\boldsymbol{v}_t$ is a n × 1 vector of disturbances with $\mathrm{E}[\boldsymbol{v}_t] = [0,0]'$ and $\mathrm{Var}[\boldsymbol{v}_t] = \mathbf{V}$

n_T is the number of time periods in the data

n_{fut} is the number of futures contracts

Like in the paper of Schwartz and Smith, $\boldsymbol{V}$ is assumed to be diagonal (Schwartz & Smith, 2000). This means that the measurement errors are assumed to be uncorrelated.

To summarize, the model parameters can be estimated with a Kalman filter using historic futures prices with a range of maturities, which will give all model parameters $(\kappa, \sigma_\chi, \lambda_\chi, \mu_\zeta, \sigma_\zeta, \mu_\zeta^*, \rho_{\zeta\chi})$. Subsequently, future spot prices can be simulated with Equations (3.30) to (3.32), by using the parameters not adjusted for risk-neutrality. From these future spot prices, the future futures prices can be determined with Equation (3.47), for which the risk-neutral parameters must be used. The Kalman algorithm is explained in APPENDIX C.

$$\begin{aligned}\ln\big(F(t,T)\big) = \exp(-\kappa T)\,\chi_t + \zeta_t + \mu_\zeta T - (1-\exp(-\kappa T))\,\frac{\lambda_\chi}{\kappa} \\ + \frac{1}{2}\Bigg((1-\exp(-2\kappa T))\,\frac{\sigma_\chi^2}{2\kappa} + \sigma_\zeta^2 T \\ + 2(1-\exp(-\kappa T))\,\frac{\rho_{\zeta\chi}\sigma_\chi\sigma_\zeta}{\kappa} \Bigg).\end{aligned} \tag{3.47}$$

4 MARKET DATA ANALYSIS

In this chapter we address the two subquestions arising from the second research question: 'How can the time evolution of the selected risk factors be modelled?' and 'What market behaviours would the model replicate (and which behaviour cannot be replicated in the PFE model)?'. In essence, various models for simulating spot and forward prices, as discussed in the preceding chapter, will undergo testing using historical carbon market data. Before delving into these assessments, in the first three sections of this chapter we perform a statistical analysis of the market data.

4.1 STATISTICAL ANALYSIS FUNDAMENTALS

Instead of the straight-forward return, the logarithmic return, as given in Equation (4.1), is often used in financial analysis. One of the advantages of the log returns is that they are additive over periods; the overall log return is the sum of individual log returns over smaller periods. In this thesis, returns will refer to the log returns, given by:

$$r = \log\left(\frac{X_i}{X_{i-1}}\right) \tag{4.1}$$

In addition to the mean and variance, skewness and kurtosis are two descriptive statistics that may be used to for time series analysis of returns. Skewness, see Equation (4.2), indicates the level of symmetry in a distribution with a value of zero generally indicating perfect symmetry (Derindere Köseoğlu et al., 2022). A positive skewness commonly indicates that the tail of the distribution is on the right side of the distribution and vice versa. However, in the case of a long tail on one side and a fat tail on the other side, skewness will not give a good measure for symmetry, as both sides of the distribution may balance out and the skewness will be close to zero.

$$S(r) = \frac{\sum_{i=1}^{n}(r_i - \mu_r)^3}{(n-1)\sigma_r^3}. \tag{4.2}$$

Kurtosis is a measure of the 'tailedness' of a distribution. The formula to calculate kurtosis is given in Equation (4.3). The standard normal distribution has kurtosis of 3 and referred to as mesokurtic. The excess kurtosis is defined as the kurtosis minus three. A negative excess kurtosis (kurtosis below 3) corresponds with a platykurtic distribution, meaning that this distribution has fewer and/or less extreme outliers compared to a normal distribution (Derindere Köseoğlu et al., 2022). It does not necessarily mean that the distribution is flat-topped as sometimes claimed. On the other hand, a distribution with positive excess kurtosis is leptokurtic and has more outliers than the normal distribution.

$$K(r) = \frac{\sum_{i=1}^{n}(r_i - \mu_r)^4}{(n-1)\sigma_r^4}. \tag{4.3}$$

4.1.1 Quantile-Quantile plots

A Quantile-Quantile plot (or QQ plot) can also aid in comparing two distributions as it plots their corresponding quantiles against each other (Derindere Köseoğlu et al., 2022). The more identical the two distributions, the closer will the points in the QQ plot be to the identity line $y = x$. This plot shows the same information as a histogram, but in a different fashion. If, in a certain area, the trend of the points is flatter than the identity line, the distribution on the horizontal axis is more dispersed than the distribution on the vertical axis. Vice versa, if at some point the trend is steeper than the identity line, the distribution on the vertical axis is more dispersed than the distribution on the horizontal axis.

4.1.2 Normality tests

In an additional attempt to assess whether data comes from a normal distribution, statistical normality tests can be performed. Two well-known types are the Jarque-Bera and Shapiro-Wilk tests (Derindere Köseoğlu et al., 2022). The first one tests whether the sample data have kurtosis and skewness matching with a normal distribution. If the test statistic, given in Equation (4.4), has a value far from zero, this indicates that the data are not from a normal distribution. Under the null hypothesis is $JB \sim \chi^2(2)$. Hence, if the p-value (calculated from the chi squared distribution) is above the significance level, the null hypothesis that the sample data have a normal distribution cannot be rejected. The Shapiro-Wilk test also has as null hypothesis that the sample data come from a normal distribution. Its test statistic is shown in Equation (4.5), which is assumed to be equal to one under the null hypothesis. SW's distribution does not have a name and p-values (based on the samples test statistic) should be read from a Shapiro-Wilk table (Derindere Köseoğlu et al., 2022)

$$JB = \frac{n}{6}\left(S^2 + \frac{1}{4}(K-3)^2\right) \tag{4.4}$$

$$SW = \frac{(\sum_{i=1}^{n} a_i x_i)^2}{\sum_{i=1}^{n}(x_i - \mu_x)^2} \tag{4.5}$$

Here S is skewness and K is kurtosis, x_i are the ordered sample data (from small to large), and a_i are constants calculated from the mean, variance and covariance of the sample (with size n) from a normally distributed sample.

4.1.3 Goodness of Fit Test

The Cramér–von Mises criterion is a statistical test that assesses the goodness of fit between an observed sample distribution and a theoretical distribution function. It is particularly useful when dealing with continuous probability distributions. The test is based on the empirical cumulative distribution function of the observed sample and compares it to the cumulative distribution function of the theoretical distribution. The null hypothesis is that the samples come from a given theoretical distribution function.

Multiple goodness-of-fit tests, including the Kolmogorov-Smirnov (KS) and Anderson-Darling (AD) tests, are available. The choice of the Cramér-von Mises (CvM) test was primarily driven by practical considerations. When assessing the fit of a model's distribution (either spot price or forward price model) against the observed return distribution, the *Scipy* library in Python restricts the KS and AD tests to comparing the observed return distribution with a specific distribution type. These tests do not allow the specification of the (cumulative) density function but only the type of distribution. In contrast, the CvM test in Python enables the comparison of the observed distribution with a manually inputted cumulative density function, making it the preferred choice.

Kernel density estimation (KDE) is a non-parametric statistical method used for estimating the probability density function of a random variable based on a finite sample set. It is particularly useful for visualizing the underlying distribution of a dataset and can be employed as a complement to histograms. The probability density function given by KDE can be regarded as a smoothed variant of a histogram, but there is no guarantee it correctly represents the true (unknown) distribution.

4.2 SPOT PRICES

A histogram for the daily log-returns is given in Figure 7 for the three carbon markets where the forward curve is only driven by the spot price (CMX, CMZ, and NZETS). These markets are not very liquid on a daily basis; hence the returns are shown on a weekly or monthly basis. Still, for a relatively large number of weeks/months a zero return is observed in these markets. The CMZ

market appears to be the most illiquid which is why the histogram for this market is given for a monthly timeframe. Figure 7 also gives an empirical distribution function of the observed returns derived with KDE.

In Figure 7 a normal distribution is given in addition to an equal weighted mean and variance. This distribution corresponds with a GBM that has a drift equal to the mean returns. We can see that the normal distribution does not fit well to the observed distribution, most especially for the NZETS market. Since, the observations seem to follow a leptokurtic behaviour, a fitted hyperbolic secant distribution is given in Figure 7 as well. This distribution, which shares many properties with the normal distribution but has a leptokurtic shape, appeared to be the best fit from all the distributions present in the Python library *Scipy*.

Several studies have demonstrated that market return distributions, particularly daily returns, deviate from normality. A few of these studies have explored adjustments to the Wiener process in the GBM by substituting it with alternative stochastic distributions, such as the Student's t-distribution (Dhesi et al., 2021; Konlack Socgnia & Wilcox, 2014; Nkemnole & Abass, 2019). Notably, however, no literature was identified that specifically replaced the Wiener process with a hyperbolic secant distribution.

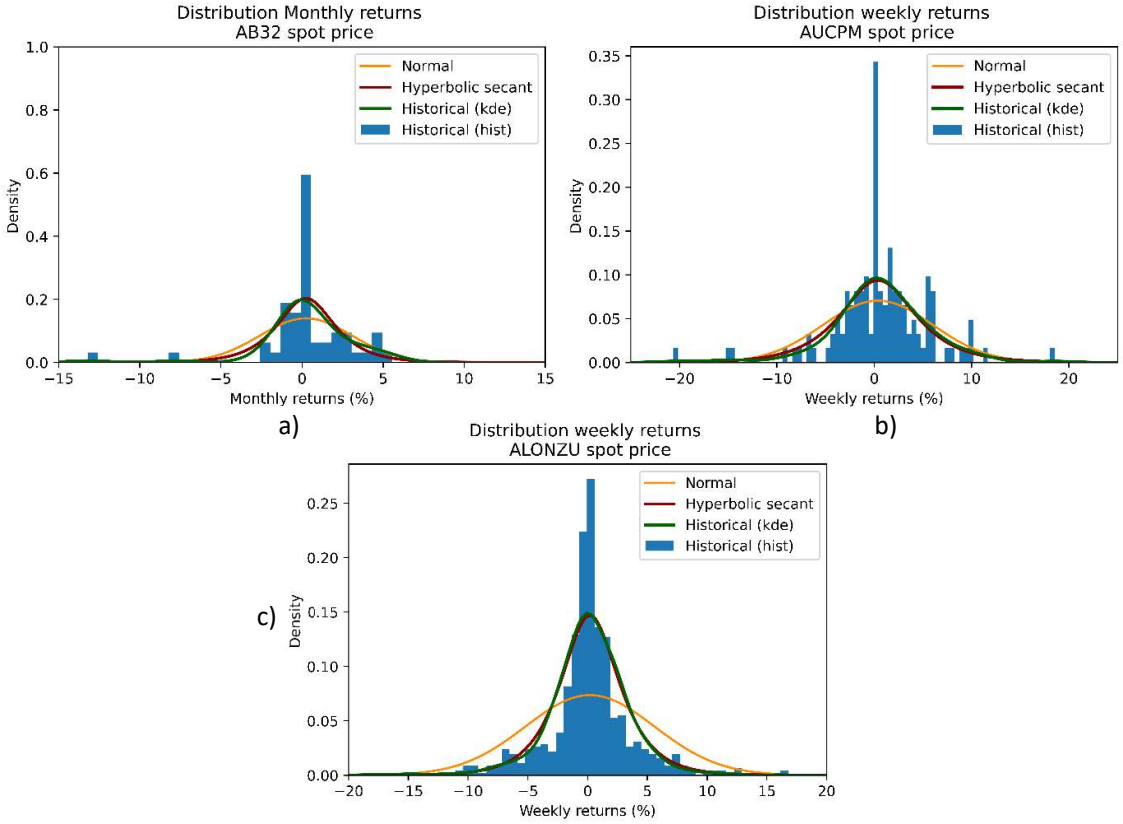

FIGURE 7 *Histograms and fitted distributions for weekly/monthly returns for a) CMZ, b) CMX, c) NZETS (as of halfway November 2023). Normal is based on sample mean and variance.*

TABLE 2 *Statistical summary for all carbon market spot prices (as of halfway November 2023). P-values below 1e-10 are denoted as 0.*

Market	Mean return (%)	Variance return (%)	Skewness	Excess kurtosis	Jarque-Bera p-value	Shapiro-Wilk p-value

CMZ (monthly)	0.24	2.83	-2.04			8.80	0	1.2e-7
NZETS (weekly)	0.19	5.41	1.02			26.7	0	0
CMX (weekly)	0.40	5.57	-1.60			8.30	0	0
CCA (weekly)	0.52	3.95	-0.60			4.80	0	0
EUETS (daily)	0.11	3.58	-0.98			14.0	0	0
GEO (weekly)	-1.46	15.5	-0.35			3.1	0	8.2e-5
ICEVER (weekly)	-5.61	25.6	-3.7			25.6	0	0
N-GEO (weekly)	-4.22	26.4	-4.49			36.4	0	0
RGGI (weekly)	0.47	2.73	-0.09			5.5	0	0
UKETS (daily)	-0.07	3.26	0.07			4.1	0	0
WCA (weekly)		0.72	5.00	-0.53	4.0	1.3e-8	1.3e-4	

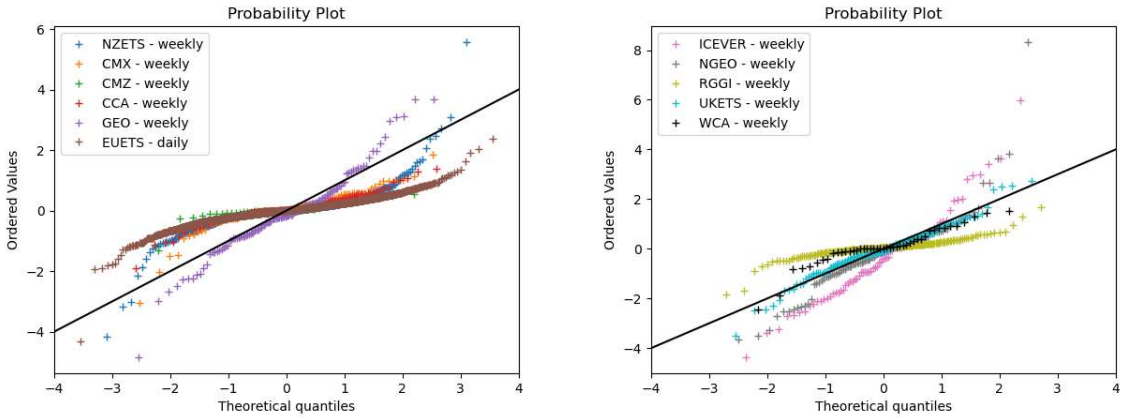

FIGURE 8 *QQ-plots of returns for all carbon market (as of halfway November 2023).*

The non-normal behaviour of the carbon market returns is confirmed by the p-values of the Jarque-Bera and Shapiro-Wilk tests in Table 2. In addition to these p-values, this table also gives the mean and variance of the returns, as well as the skewness and kurtosis of the observed distributions for all the carbon markets. The only exception is the UKETS market, which has a Shapiro-Wilk p-value of 0.87, and may be regarded as normally distributed. The p-value of the Jarque-Bera test is 0.046, meaning the normal distribution null-hypothesis would be rejected at a 5% significance level. Nevertheless, the observed distribution, given in Appendix D.1, shows a close resemblance to the normal distribution. As will be explained soon, this appendix shows the histograms for all the carbon markets in addition to those in Figure 7. Here we find that, although not directly supported by the data in Table 2, the ICEVER returns appear to be relatively close to the normal (GBM) distribution.

Table 2 shows that in most carbon markets the returns are either slightly positively or negatively skewed, meaning the distributions are not perfectly symmetrical. Higher deviations from the normal distributions are observed in terms of kurtosis. Most of the carbon markets have positive excess kurtosis (a leptokurtic distribution), while the ICEVER and UKETS have negative kurtosis (platykurtic).

The QQ-plots for the carbon market returns in Figure 8 present the same data as shown in the histograms, with the theoretical quantiles referring to those of a normal distribution. Recall that a steeper trend in the market returns compared to the black diagonal means that the observations are more dispersed than the normal distribution, and vice versa. For most of the markets in Figure 8 the trend near the mean is flatter (than the diagonal) while towards the tail, this converts into a steeper trend. However, this is not the case for the returns of the GEO and ICEVER markets.

There is one additional important remark that should be made, being that the shape of the distribution may change when the returns are calculated over a different time step. This can be illustrated well with a longer-existing and more liquid market as the EUETS. The histogram of the daily returns for this market (Figure 9a) exhibits leptokurtic behaviour, whereas this becomes less prevailing on a weekly basis (Figure 9b).

To illustrate this effect for more carbon markets as well, Table 3 shows the same statistical results as given in Table 2 but for different sampling intervals. Only those markets are included where there is enough liquidity and number of datapoints to sample on different intervals. In general, the skewness and kurtosis of the return distributions move towards those of a normal distribution when the sampling interval is increased.

Furthermore for the EUETS market, the p-value of the Jarque-Bera and Shapiro-Wilk tests are 0.17 and 0.11 respectively on a monthly basis, whereas they are practically zero on a daily basis. The returns in the UKETS market show the same behaviour when moving from a daily sampling timestep to weekly.

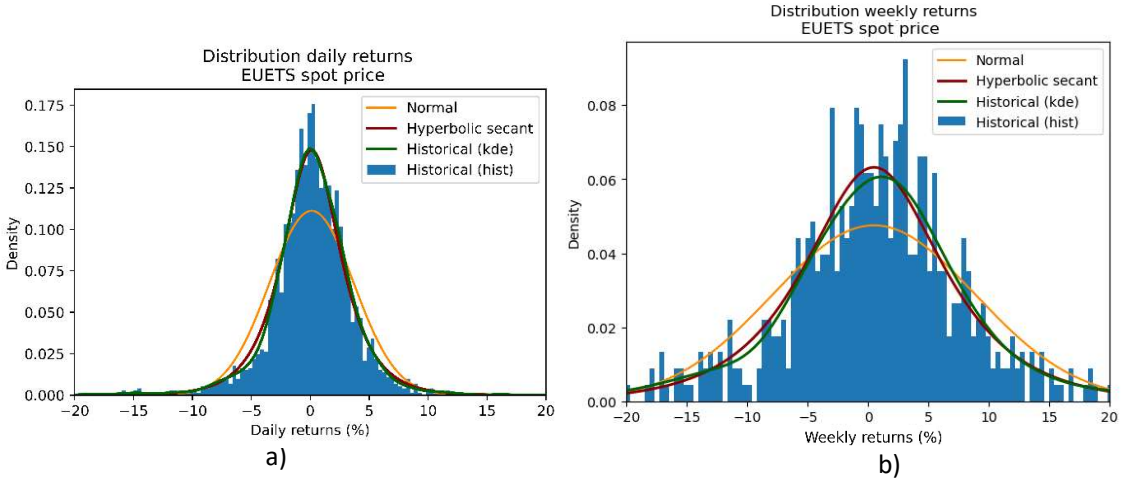

FIGURE 9 *Effect of changing changing return timestep for EUETS.*

TABLE 3 *Statistical summary for selected carbon market spot prices to illustrate effect of sampling interval (as of halfway November 2023). P-values below 1e-10 are denoted as 0.*

Market	Mean return (%)	Variance return (%)	Skewness	Excess kurtosis	Jarque-Bera p-value	Shapiro-Wilk p-value
NZETS (weekly)	0.19	5.41	1.02	26.7	0	0
NZETS (monthly)	0.89	12.0	0.52	5.2	0	5.2e-10
CMX (weekly)	0.40	5.57	-1.60	8.30	0	0
CMX (monthly)	1.69	14.2	-1.1	2.3	2.1e-3	3.7e-2
EUETS (daily)	0.11	3.58	-0.98	14.0	0	0
EUETS (weekly)	0.42	7.11	-0.85	6.1	0	0
EUETS (monthly)	2.77	17.6	-0.37	0.06	1.7e-1	1.1e-1
RGGI (monthly)	2.06	4.96	0.72	3.1	1.2e-5	1.8e-2
UKETS (daily)	-0.07	3.26	0.07	4.1	0	0
UKETS (weekly)	-0.27	7.14	0.09	0.4	6.1e-1	8.6e-1

In any case, the argument is clear that during the parameter estimation process of the SDEs (as discussed in Section 3.3) the timestep used to calculate the market returns, will be an important factor to consider. If one wants to use the PFE model (that will be built in this thesis) with a daily timestep for the EUETS or UKETS (for example to determine expected margin requirements), a different SDE might be required compared to modelling PFE on a weekly or monthly basis.

The observation that, in a liquid market, monthly returns tend to approximate a normal distribution can be elucidated through the Central Limit Theorem (CLT). This theorem states that as the sample size increases, the sum of independent random variables will converge towards a normal distribution, irrespective of the initial distribution of these variables (Mohammad Rafiqul, 2018). This principle is applicable to monthly (and yearly) returns, which constitute the sum of daily log returns. Nevertheless, it is crucial to note that, due to the inherent illiquidity of most carbon markets,

monthly returns frequently deviate from normal distribution characteristics, and there are often insufficient data to construct a distribution of yearly returns. Furthermore, the CLT assumes the daily log returns are independent and identically distributed, which may not necessarily be true.

4.3 FUTURES/FORWARD PRICES

The forward prices pertaining to carbon markets are derived from the energy trade and risk management (ETRM) database and subsequently configured into a matrix structure. Each row within this matrix corresponds to a historical date during which prices were observed, while each column aligns with a distinct month-to-maturity. It is noteworthy that the stored data is not exhaustive. This can be attributed, in part, to the absence of quotes for carbon allowances/credits across all months in a given year. For instance, with respect to the EUETS, only four quotes are available for maturity dates in the year 2024—specifically for March 1st, June 1st, September 1st, and December 1st . Consequently, on December 1st, 2023, these quotes manifest as 3-, 6-, 9-, and 12-month forward prices. Subsequently, one month later (January 1st), the quotes for these contracts translate into 2-, 5-, 8-, and 11-month forward prices. As such, the stored forward prices lack a fixed month-to-maturity.

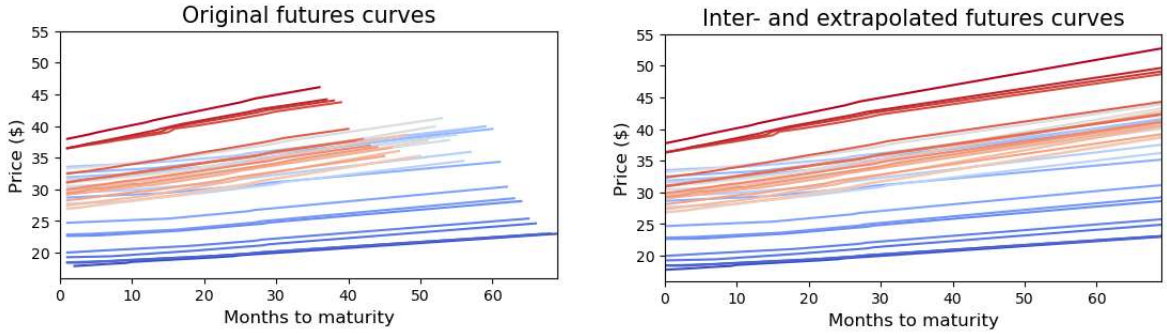

FIGURE 10 *Historical CCA forward curve with monthly sampling frequency (oldest=blue, most recent=red).*

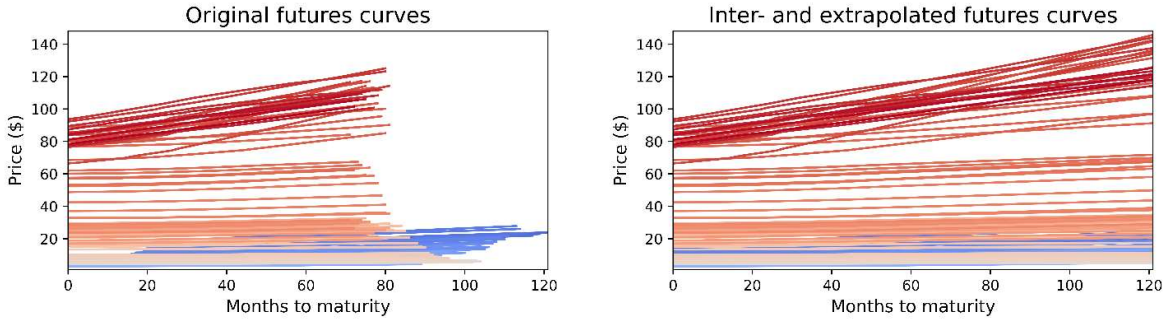

FIGURE 11 *Historical EUETS forward curve with monthly sampling frequency (oldest=blue, most recent=red).*

Moving further along the forward curve, the frequency of price quotes per annum diminishes across all carbon markets. Additionally, the historical forward curves stored exhibit considerable variability in length. To illustrate this for the CCA market, the earliest historical forward curve encompasses a 66-month forward price, whereas the most recent forward curve, as of the present writing, only captures forward prices up to 33 months.

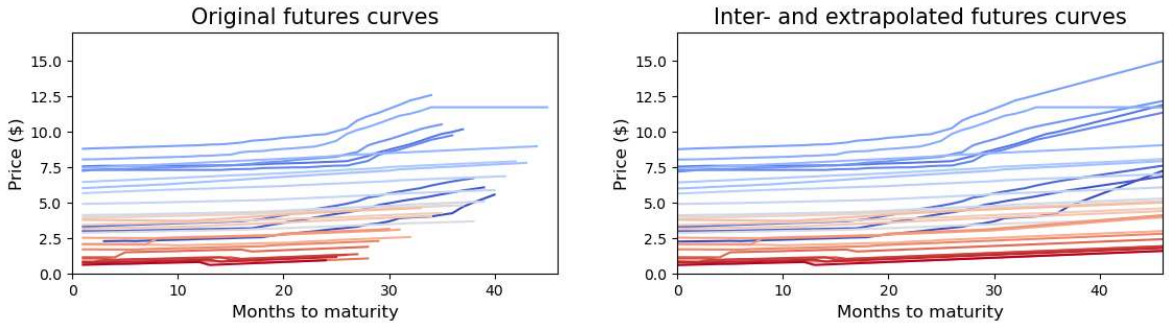

FIGURE 12 *Historical GEO forward curve with monthly sampling frequency (oldest=blue, most recent=red).*

The application of Principal Component Analysis (PCA) on historical forward curves necessitates a matrix without gaps. This prerequisite is similarly applicable to the parameter estimation process of the Schwartz-Smith model. As illustrated in Figure 10 to Figure 17, the forward curves exhibit a notable linearity, rendering interpolation a suitable method for populating certain gaps within these curves. Given this characteristic, the decision was also made to extrapolate the forward curves to address all data gaps. While this extrapolation appears justifiable based on the curve shapes, it is imperative to maintain awareness of this procedure's influence when scrutinizing the subsequent statistical analyses in this chapter, particularly in the context of longer month maturity forward prices.

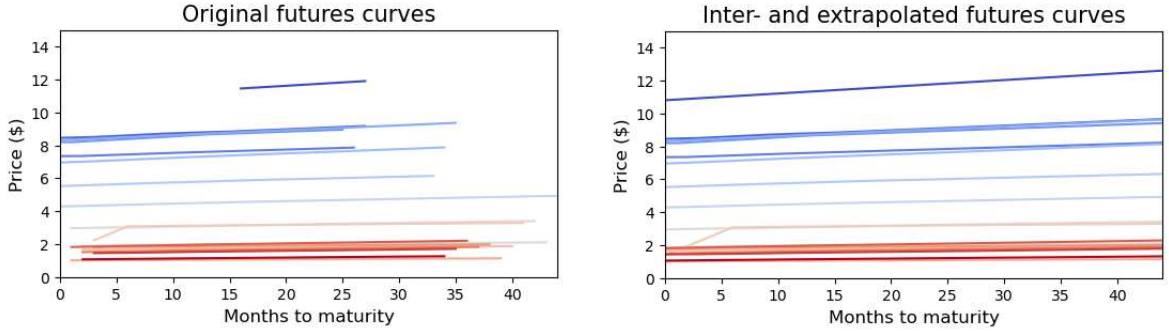

FIGURE 13 *Historical ICEVER forward curve with monthly sampling frequency (oldest=blue, most recent=red).*

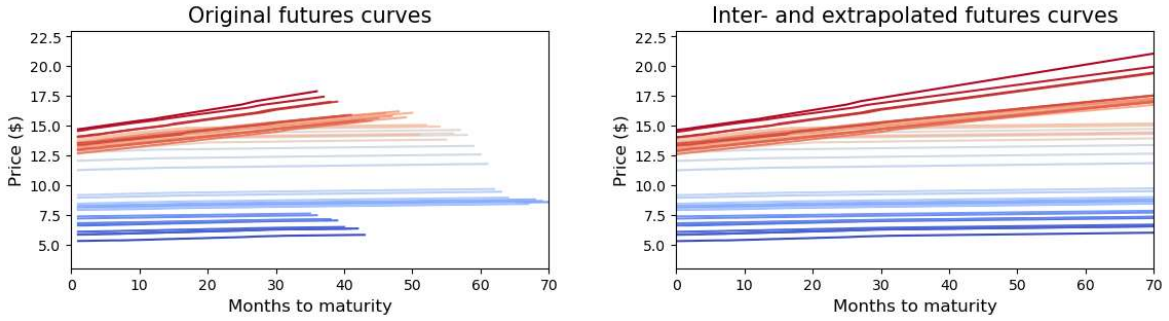

FIGURE 14 *Historical RGGI forward curve with monthly sampling frequency (oldest=blue, most recent=red).*

The integrity of N-GEO curves, as depicted in Figure 17, appears compromised. Subsequent verification through examination of data on the CME website, where prices for the identical market are quoted (CME group, 2023), substantiates the existence of this corruption. Consequently, this particular market will not undergo further analysis.

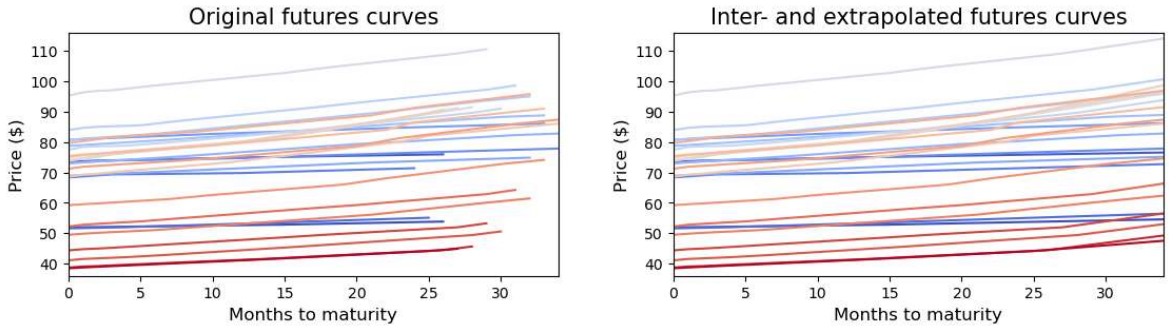

FIGURE 15 *Historical UKETS forward curve with monthly sampling frequency (oldest=blue, most recent=red).*

As previously articulated, most historical carbon forward curves exhibit linearity. Notably, the GEO curves (Figure 12) represent an exception, wherein a discernible curve is evident, particularly in the end portions of the older curves. Nevertheless, the most recent curves, denoted in red, demonstrate a notably heightened linearity. Consequently, the extrapolation of the older historical futures curves is deemed to be less reliable.

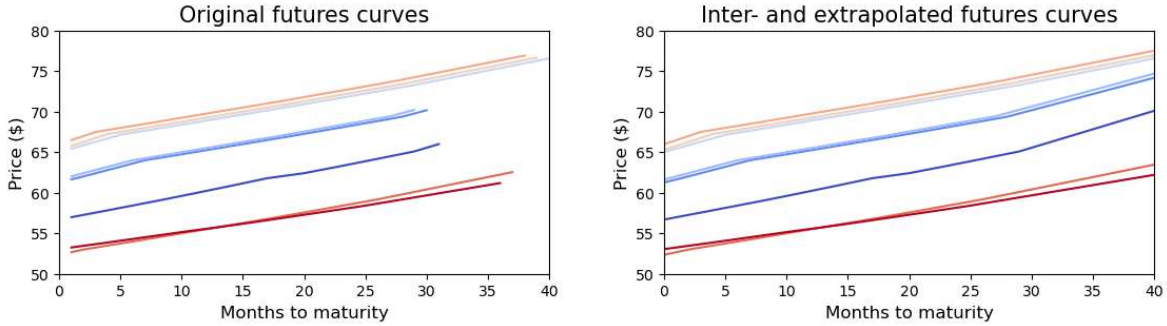

FIGURE 16 *Historical WCA forward curve with monthly sampling frequency (oldest=blue, most recent=red).*

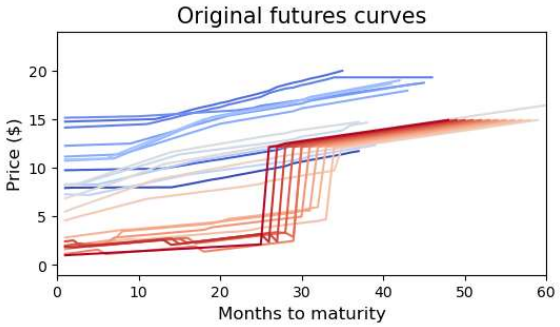

FIGURE 17 *Historical N-GEO forward curve with monthly sampling frequency (oldest=blue, most recent=red). Recent data is corrupted.*

PCA was executed on the historical forward curves, and the outcomes, in terms of the cumulative explained variance of the principal components, are presented in Figure 18. This figure shows the cumulative explained variance for historical forward curves, sampled on both weekly and monthly intervals. Evidently, the findings substantiate the assertion of a prevailing linear trend, as the first principal component contributes to over 96% of the total observed variance in forward prices across each carbon market. In many instances, this percentage is even higher. As previously mentioned, the

first principal component corresponds to a parallel shift in the curves, denoting a change in the spot price within this context.

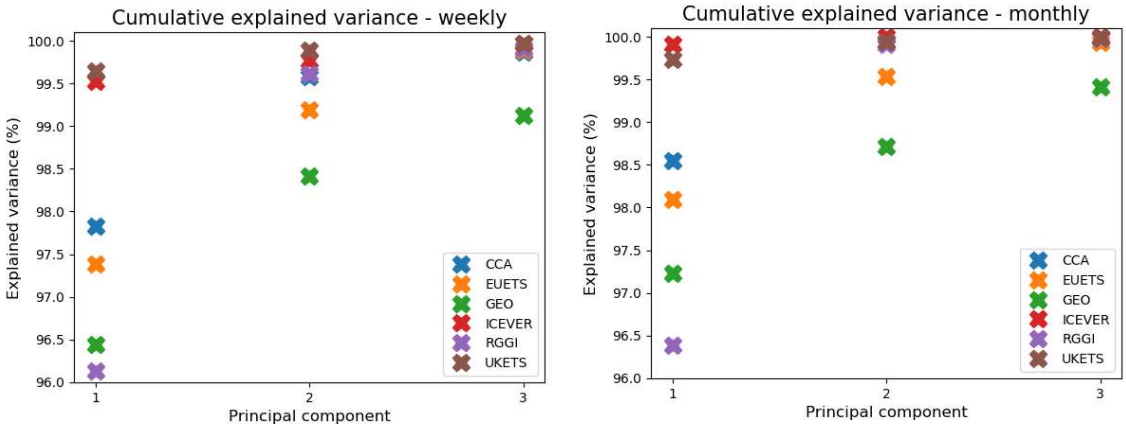

FIGURE 18 *Cumulative explained variance of first three principal components for historical forward curves sampled on a weekly and monthly basis.*

Notably, the carbon markets exhibiting the lowest contributions for the first principal component are GEO and RGGI. In the case of GEO, this observation is likely attributed to a change in curvature towards the end of older forward curves.

For RGGI, Figure 14 illustrates a change in the slope of forward curves as one moves ahead in time. This is substantiated by the notable cumulative explained variance at the second principal component in Figure 18. The difference between the cumulative explained variance of the first and second principal components represents the explained variance of the second principal component, serving as an indicator for the alteration in slope evident in the historical curves.

Table 12 to Table 18 in Appendix D.2 provide a comprehensive summary of the statistical attributes characterising the distribution of returns from the historical forward curves. The returns are computed on a per-month-to-maturity basis, involving a comparison of prices for consecutive periods, such as the prices of the consecutive 5-month forward contracts. The table presents these statistical measures for three distinct month-to-maturities: one at the start of the forward curve, another in the middle, and a third near the end. Furthermore, a differentiation has been established between weekly and monthly forward price returns. For this analysis, the inter- and extrapolated forward curves have been used.

The null hypotheses for the Jarque-Bera and Shapiro-Wilk tests are consistently rejected at the 95% confidence level when conducted at a weekly sampling interval. However, for the majority of markets, this null hypothesis is not rejected when the tests are applied the monthly sampled data. An exception to this trend is observed in the case of the RGGI market, where the null hypothesis is also rejected at the 95% confidence level for the monthly returns. Notably, for the WCA market, the Jarque-Bera test yields a rejection of the null hypothesis. In the case of the Shapiro-Wilk test, the null hypothesis is only rejected for the 5-month-to-maturity category. However, it is important to note that for all three maturities, the p-values for this test in this market is relatively close to 5% for all three maturities.

At the weekly sampling interval, the p-value of the Cramér–von Mises test, which assesses the empirical distribution against a theoretical hyperbolic secant distribution, is notably high. Additionally, it becomes evident, both visually and from the results, that this distribution more accurately characterizes daily returns in comparison to the normal distribution. Similarly, at the monthly sampling interval, the p-value remains relatively high, leading to the non-rejection of the null hypothesis positing a hyperbolic secant distribution. As said, the normality null hypothesis is also

not rejected for most markets at the monthly sampling interval, but with much lower p-values. These findings imply that even at a monthly sampling frequency, the hyperbolic secant distribution might offer a more suitable fit than the normal distribution; however, this conclusion lacks unambiguous statistical substantiation.

The empirical density functions, derived from the histogram of returns, are depicted in Figures 27 to 33 (Appendix D.2). Consistent with observations from the spot price analysis, distinctions in the shapes of daily and monthly forward price returns are discernible, primarily manifested through alterations in skewness and kurtosis. Notably, minimal variations in the mean and variance of the annualized returns are evident, attributable to inherent randomness.

The above results indicate that the spot price is by far the most important factor driving the changes of the forward curves for all carbon markets, with the RGGI having the largest change in shape. However, this change is still relatively limited. Another way of coming to the same conclusion is by looking at the continuous empirical density functions in Figure 42 to Figure 48. We see that although there is a change between the weekly and monthly distributions, there is more overlap between the distributions of the three months-to-maturities that are shown. This means that the different months-to-maturity forward prices move in a relatively similar fashion.

4.4 FORWARD CURVE – MODEL PARAMETER ESTIMATION

While we have concluded in the preceding section that the driving force behind the evolution of forward curves in all carbon markets is predominantly attributed to changes in the spot price, we want to assess the appropriateness of the Schwartz-Smith model and PCA as models for simulating forward curves. As shocking the forward curve with the normal distribution closely resembles the PCA approach, it is not examined separately. Consequently, this section is dedicated to examining the simulation of forward curves through both the Schwartz-Smith model and PCA. These modelling techniques undergo testing in the simulation of weekly and monthly forward price returns within the RGGI and UKETS markets.

4.4.1 Schwartz-Smith Model Parameter Estimation

We developed a Python script to undertake the parameter estimation of the Schwartz-Smith model using the Kalman filter approach which is detailed in APPENDIX C. The estimated parameters derived from this process were subsequently employed to simulate future spot prices utilizing Equations (3.30) to (3.32) without incorporating risk adjustments. Subsequently, we applied Equation (3.47) to ascertain forward prices based on the simulated future spot prices, necessitating the utilization of risk-neutral parameters.

Prior to executing simulations on carbon market data, we validated the parameter estimation section of the Python script. Parameter estimation was conducted on data that closely approximated the dataset utilized in the original work by Schwartz and Smith, as presented in the attached MSc thesis (Goodwin, 2013). The resultant parameters closely mirrored those reported in the paper by Schwartz and Smith (Schwartz & Smith, 2000).

TABLE 4 *Estimated Schwartz-Smith model parameters on the historical RGGI forward curve returns.*

	κ	σ_χ	λ_χ	μ_ζ	σ_ζ	μ_ζ^*	$\rho_{\chi\zeta}$
Weekly	2.672	0.6289	0.3798	0.3199	0.1111	-4.125e-3	0.2838
Monthly	0.9369	0.2546	2.892e-2	0.3142	3.471e-2	1.222e-3	0.2092

The historical RGGI forward curves encompass prices extending up to a maturity of 70 months, resulting in an inter- and extrapolated forward curve matrix with 70 columns. Notably, estimation of the Schwartz-Smith model parameters using this relatively extensive matrix exceeded a time frame of 10 minutes and yielded unstable results, characterized by a high dependence on the initial estimates for the parameters. Consequently, we decided to employ a subset of maturity months columns, specifically selecting the 1st, 10th, 20th, 30th, 40th, 50th, 60th, and 70th months. The parameters resulting from this approach are detailed in Table 4.

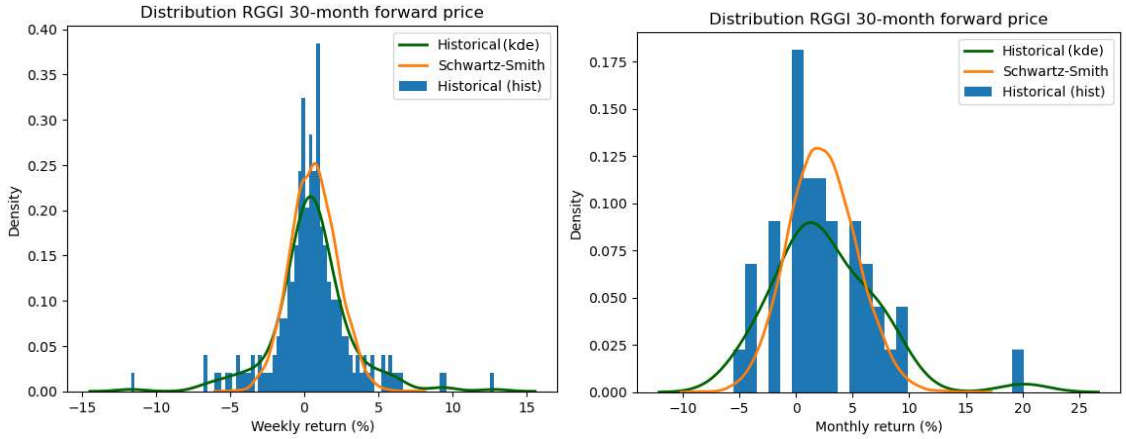

FIGURE 19 *Weekly and monthly RGGI 30-month returns predicted by the Schwartz-Smith model compared to historical returns.*

In Figure 19, the distributions of weekly and monthly returns, derived from the Schwartz-Smith model for the 30-month RGGI forward price, are presented alongside the observed historical distribution. It is evident from the analysis that the predictions generated by the Schwartz-Smith model inadequately capture the historical return pattern, both on a weekly and monthly sampling interval. Specifically, the tails of the Schwartz-Smith distribution appear to be too thin, and in the vicinity of the mean, the distribution exhibits an elevated profile. Similar shortcomings are observed in the results of the 70-month forward price returns, as illustrated in Figure 49 in Appendix E.1.

Additionally, it is noteworthy that the parameter estimation conducted on weekly and monthly sampled data yields disparate parameter values. This does not imply the model works incorrectly, but it limits the explanatory power of the estimated parameters in a real-world context.

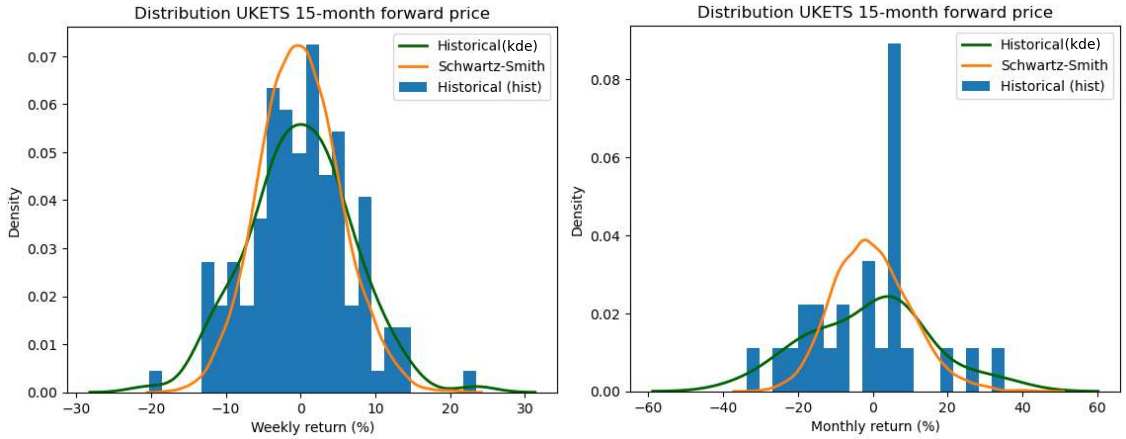

FIGURE 20 *Weekly and monthly UKETS 15-month returns predicted by the Schwartz-Smith model compared to historical returns.*

TABLE 5 *Estimated Schwartz-Smith model parameters on the historical UKETS forward curve returns.*

	κ	σ_χ	λ_χ	μ_ζ	σ_ζ	μ_ζ^*	$\rho_{\chi\zeta}$
Weekly	9.272e-2	0.3538	0	3.453e-3	0.1332	4.443e-3	0.3922
Monthly	6.218e-2	0.270715	0	-2.757e-2	0.1777	-2.324e-3	0.3669

We took similar steps for the UKETS market, involving a subset of the 1st, 5th, 15th, 20th, 25th, 30th, and 35th-month forward prices. The resulting parameters are detailed in Table 5, and the distributions for the 15-month forward prices are depicted in Figure 20. While the estimated parameters exhibit less disparity compared to the RGGI market, the distributions generated by the Schwartz-Smith model still do not align well with observed market data returns. This observation holds true when examining the weekly and monthly returns distributions for the 35-month forward price, as illustrated in Figure 50 in Appendix E.1.

As mentioned at the beginning of this chapter, the p-value of the CvM test serves as a means to evaluate the fit of the distribution predicted by a model with the observed data. Table 6 provides the CvM p-values for the test comparing the Schwartz-Smith return distribution with the observed data for both markets discussed in this section, with additional details in Appendix E.1. The results indicate that the Schwartz-Smith model performs inadequately for RGGI, as evidenced by the ability to reject the null hypothesis that the observed data comes from the Schwartz-Smith return distribution at a 5% significance level. For the UKETS, the null hypothesis cannot always be rejected, but the relatively small p-values suggest a suboptimal fit. Similar p-values in other carbon markets for Schwartz-Smith simulations further support this conclusion. Considering these findings, it was determined that the Schwartz-Smith model is not a suitable candidate for modelling the forward curves, based on the following arguments:

- The parameter estimation process is notably time-consuming, exceeding 20 minutes for certain carbon markets. These durations are observed even when employing a limited selection of forward prices, focusing solely on specific maturities, whereas an ideal scenario would involve utilizing the complete set of prices.
- The estimated parameters exhibit limited explanatory power, due to the change in value when varying the sampling interval.
- The return distribution demonstrates poor alignment with the observed data, as indicated by the CvM p-values.

TABLE 6 *P-value CvM test for weekly and monthly simulated RGGI and UKETS forward prices with Schwartz-Smith model. Middle month refers to 30 and 15 and last month refers to 70 and 35 for RGGI and UKETS respectively.*

	RGGI		UKETS	
	Weekly	Monthly	Weekly	Monthly
Middle month forward	2.4e-2	1.8e-2	0.08	0.06
Last month forward	1.7e-2	1.9e-2	0.05	0.02

4.4.2 Principal Component Analysis

We developed an additional Python script to introduce shocks to the forward curve of the RGGI and UKETS markets, following the formulation in Equation (3.28). The resulting distributions from this approach for these markets, observed on a weekly and monthly interval, are illustrated in Figure 21 and Figure 22, respectively. Note that the histograms in these figures are constructed with a different number of bins compared to Figure 19 and Figure 20, resulting in a slightly different shape. Nevertheless, the underlying data and KDE estimate of the distribution remain the same. In this methodology, the principal components are subjected to shocks based on a normal distribution, and

consequently, the resulting distribution fails to capture the leptokurtic characteristics observed in the weekly RGGI returns. The monthly-sampled RGGI returns exhibit a distribution closer to that of a normal distribution, leading to a PCA distribution that bears a stronger resemblance to the observed distribution. However, similar to the case of weekly returns, the PCA return distribution inadequately captures the mean of the distribution.

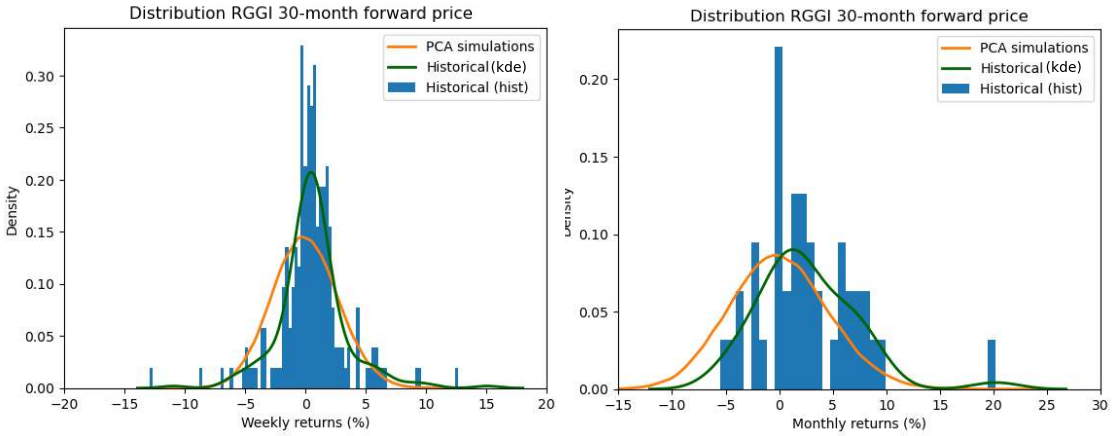

FIGURE 21 *Weekly and monthly RGGI 30-month returns predicted by PCA simulations compared to historical returns.*

Consistent with observations made for spot prices, the observed return distributions for UKETS exhibit a close resemblance to a normal distribution and lack a pronounced mean return. Consequently, the PCA forward curve simulation model yields a relatively good fit with the observed distribution, as we in Figure 22. This pattern persists in the results for forward prices with a longer maturity (Figure 51 and Figure 52 in Appendix E.2).

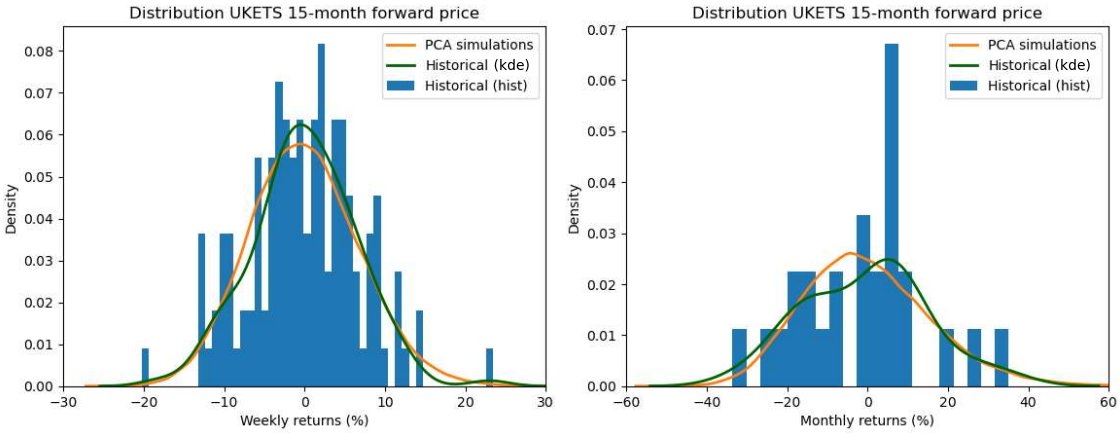

FIGURE 22 *Weekly and monthly UKETS 15-month returns predicted by PCA simulations compared to historical returns.*

The CvM test p-values, which assess the fit of the PCA return distribution to the observed data, are provided in Table 7. These findings validate the visual observation that the PCA return distribution aligns more closely with the observed data compared to the Schwartz-Smith model for the UKETS market. Notably, the p-values are remarkably high, indicating a relatively strong fit. In contrast, for the RGGI market, the p-values are low. However, a significant portion of this discrepancy stems from the PCA model's inability in capturing the mean return. Adjusting the observed distribution by subtracting the mean return from each data point leads to a substantial increase in the CvM p-values. While the weekly simulated returns exhibit p-values below the 5% significance level,

attributed to the leptokurtic behaviour of observed market returns, the monthly returns present p-values surpassing this threshold.

TABLE 7 *P-value CvM test for weekly and monthly simulated RGGI and UKETS forward prices with PCA.*

	RGGI		UKETS	
	Weekly	Monthly	Weekly	Monthly
Middle month forward	1.5e-5 (Zero mean=2.2e-3)	8.1e-4 (Zero mean=0.72)	0.97	0.55
Last month forward	1.4e-5 (Zero mean =1.1e-3)	5.9e-4 (Zero mean=0.32)	0.89	0.53

4.5 SPOT – MODEL PARAMETERS ESTIMATION

In this section, we delve into the parameter estimation (or model fitting) results of the spot price models. However, it is required to acknowledge a constraint stemming from the relatively brief duration over which historical prices have been observed. Given the limited time series data, the estimation of the drift term becomes impractical. For instance, when considering the annualized average weekly returns of 47% for WCA, 26% for CCA, and -63% for GEO, it is unreasonable to assume that these observed drifts accurately reflect the drifts for the ensuing 5 or 10 years—the prospective time horizon for which the PFE model could potentially be executed. Consequently, in this section, the parameter estimation process is executed on returns where the mean return has been subtracted. This adjustment ensures that the drift term in the models equates zero.

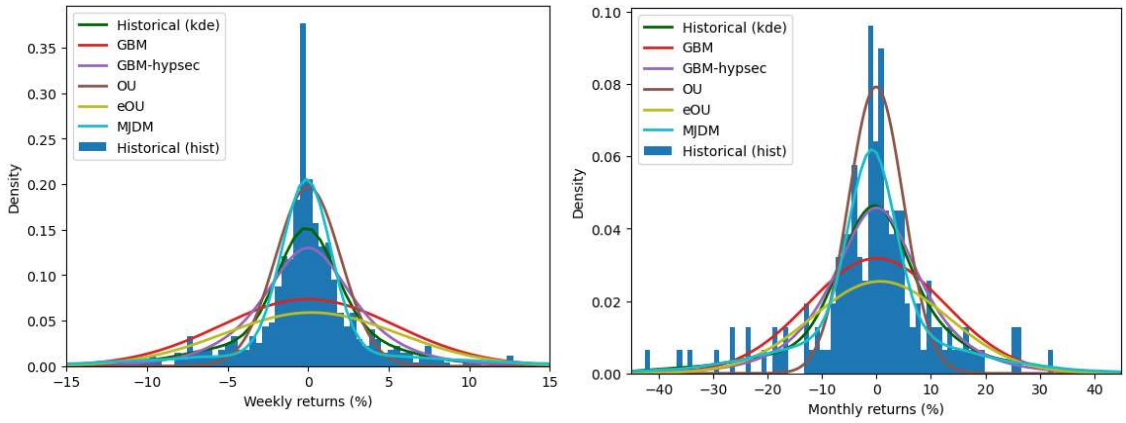

FIGURE 23 *Fitted spot price model distributions for NZETS.*

Figure 23 shows the histogram of weekly and monthly returns for the NZETS spot price, accompanied by KDE density function. Additionally, the density functions of the fitted price simulation models are given. The GBM-hypsec model denotes a GBM where the Wiener process is substituted with a stochastic process characterized by a hyperbolic secant distribution featuring a zero mean and a predefined variance. This variance is determined by fitting a hyperbolic secant function to the observed return distribution.

The first conclusion that can be drawn is that the OU and eOU models exhibit limited capability in accurately capturing market return behaviour. This observation is substantiated by Table 8, which provides the log-likelihood of the observed prices predicted by various models for a weekly and monthly sampling interval, along with the p-value of the Cramér–von Mises (CvM) test. Notably, both the OU and eOU models yield the lowest log-likelihoods and p-values. This trend is not specific to NZETS spot prices; rather, across all carbon spot prices, these models consistently demonstrate the least favourable performance.

The lack of mean reversion becomes even more apparent when looking at the Hurst components (H) in Table 19 in APPENDIX F. In short, the Hurst component is a measure indicating whether a time series is purely random ($H = 0.5$), trending ($H > 0.5$), or mean reverting ($H < 0.5$). Since none of the values in Table 19 are below 0.5, there is no evidence of mean reversion. Consequently, the mean reverting models are excluded from subsequent analyses in this chapter.

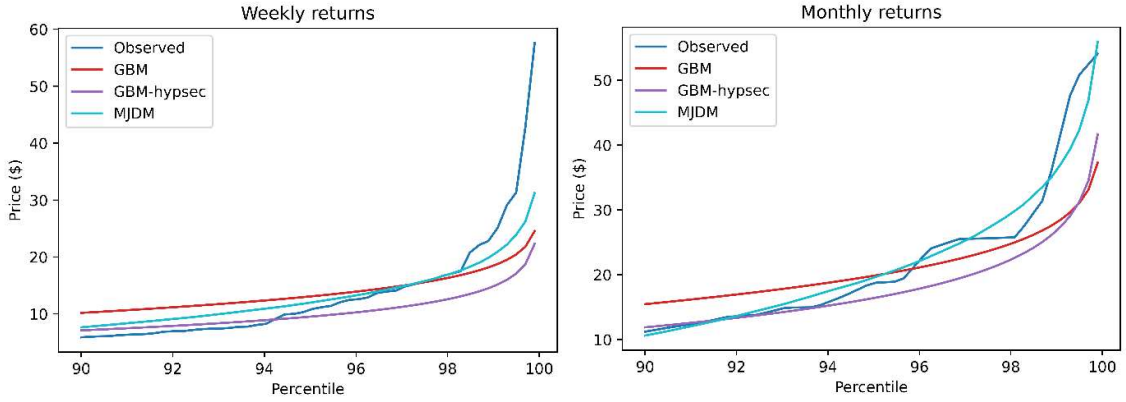

FIGURE 24 *Right tail percentiles of fitted model distributions for NZETS.*

TABLE 8 *NZETS spot price model fitting results.*

		GBM	GBM-hypsec	OU	eOU	MJDM
Log-likelihood	Weekly	-734	X	-1243	-908	-598
	Monthly	-297	X	-406	-333	-295
p-value CvM	Weekly	0	0	0	0	1.9e-1
	Monthly	1.1e-3	5.5e-2	1.6e-2	6.4e-4	4.0e-1

In terms of both log-likelihood and the CvM p-value, the MJDM appears to offer the most accurate representation of the observed distribution for the weekly and monthly sampled returns in the NZETS market. Notably, the log-likelihood for the GBM-hypsec model is excluded from Table 8 due to its non-standard nature, resulting in the absence of a well-defined density function. Although one might anticipate a higher p-value for the GBM-hypsec model compared to the MJDM model, given its superior alignment with the historical KDE density function (Figure 23), it is crucial to emphasise that the CvM test evaluates the theoretical distribution against the observed data rather than against the KDE density function.

The log-likelihood and CvM p-value are comprehensive metrics that consider the entire distribution, making them insensitive to specific characteristics of the distribution's tails. Consequently, Figure 24 is introduced to examine the 90th to 99.9th percentiles of the observed and fitted distributions. The percentiles of the observed distribution are represented by an irregular line, as linear interpolation is employed to estimate these percentiles due to the limited number of observations in the tail region. The observed percentiles should be regarded as proxies rather than a precisely defined line due to the interpolation and limited data points. Modelling PFE implicitly places a focus on the tail of the simulated return distribution and hence it is desirable to have a proper match between the model used for simulation and the observed distribution.

For both the weekly and monthly returns, the MJDM appears to provide the most accurate representation of the tail distribution in comparison to the observed distribution. It is noteworthy that the 99.9th percentile of observed weekly returns significantly exceeds the corresponding percentile of the MJDM distribution. This elevated observed percentile is attributed to a singular exceptionally high weekly return of 58%. However, discerning whether this discrepancy arises from the MJDM model's inability to accurately capture the extreme tail of the distribution (for the weekly

returns) or if this isolated data point is unrepresentative of the entire distribution remains indeterminate.

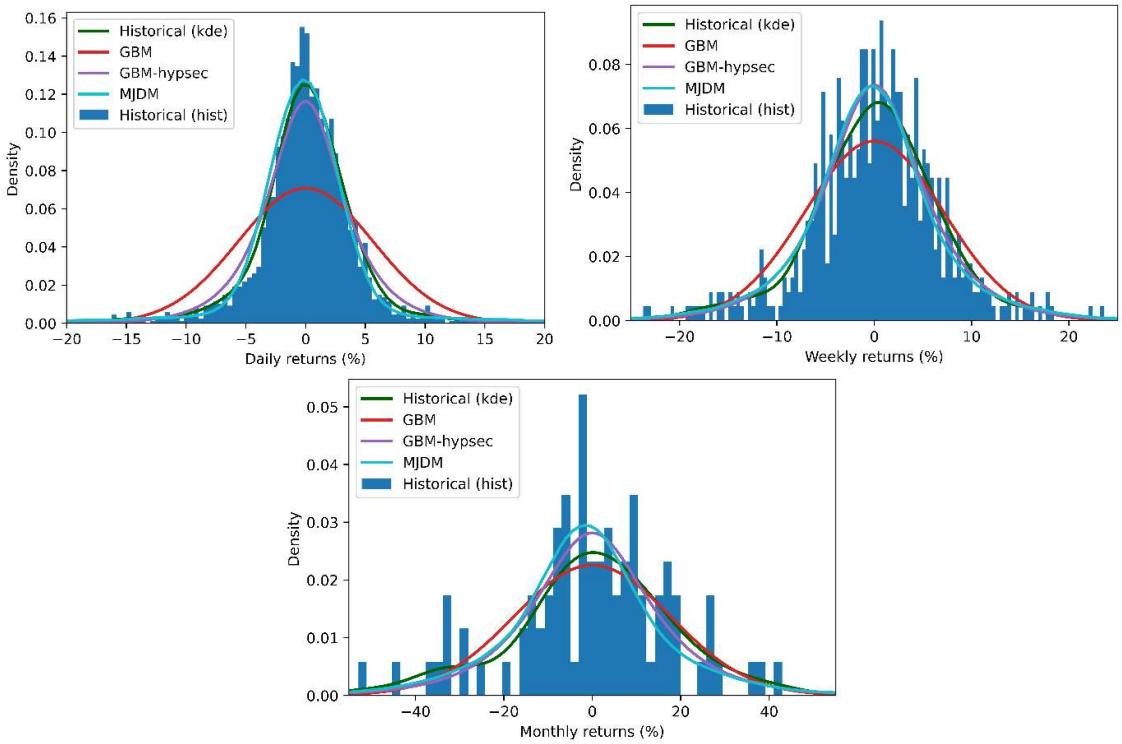

FIGURE 25 *Fitted spot price model distributions for EUETS.*

TABLE 9 *EUETS spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Daily	-2251	X	-2085
	Weekly	-730	X	-830
	Monthly	-202	X	-222
p-value CvM	Daily	0	7.4e-3	7.5e-2
	Weekly	6.1e-3	3.7e-1	1.1e-1
	Monthly	4.9e-1	3.1e-1	2.0e-1

The model fitting outcomes for the EUETS, at a daily interval, as given in Figure 25 and Table 9, distinctly reveal superior performance of the GBM-hypsec and MJDM models compared to the standard GBM. In the context of weekly sampled returns, the GBM-hypsec model appears to have a slight edge, as indicated by both the CvM p-value and percentile profile in Figure 26. Notably, considering our earlier conclusion that the EUETS likely follows a normal distribution for monthly returns, the GBM model emerges as the more effective choice for this sampling period.

Similar results for the other carbon markets are given in APPENDIX F and will be discussed below shortly. The idea of the final PFE model is that it can be used to select the most appropriate price simulation model before the actual PFE calculations are performed. Hence, it is not required to fit all

the components in the various carbon markets. However, for the components in Table 11 (Appendix D.1) a preselection of the price simulation model will be given.

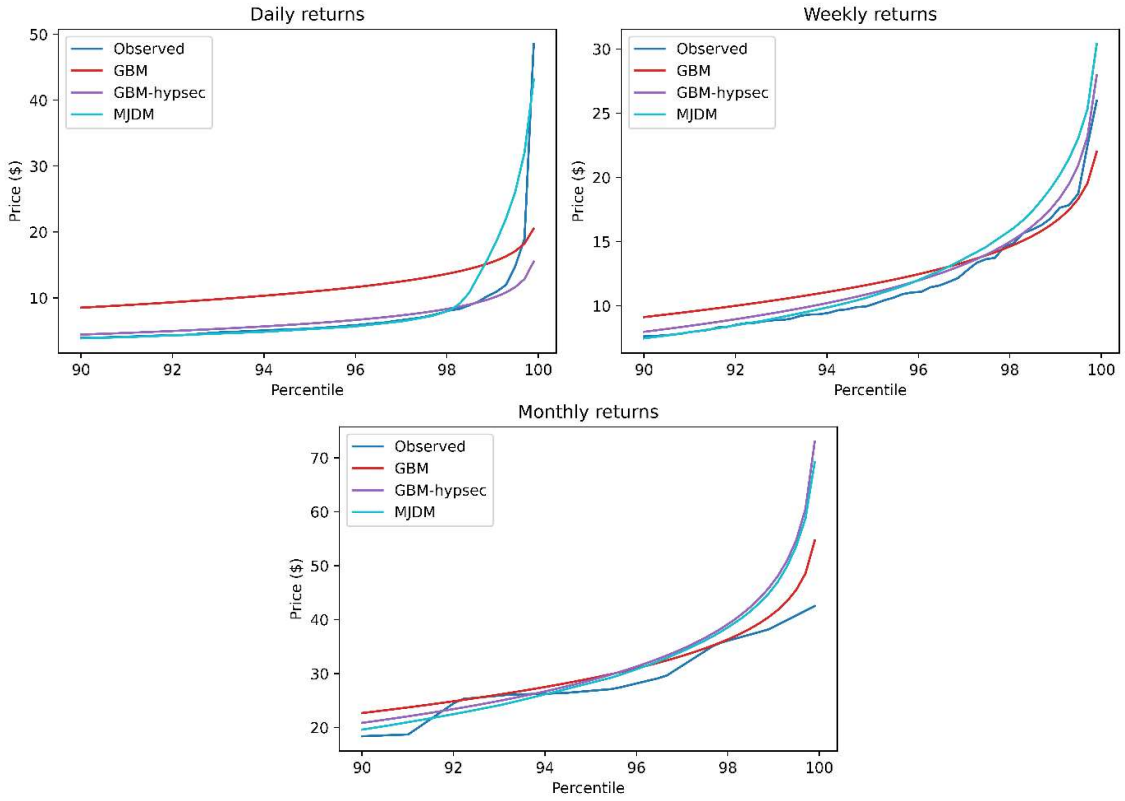

FIGURE 26 *Right tail percentiles of fitted model distributions for EUETS.*

The initial market to discuss is CMX, for which both the observed and fitted distributions, along with the right-tail percentiles, are illustrated in Figure 53. Analysing the log-likelihood and CvM p-value (refer to Table 20) suggests that the MJDM stands out as the most fitting. It is crucial to exercise caution in drawing conclusions due to the relatively limited number of observations. The scarcity of data points particularly renders the right end of the observed percentiles line in Figure 53 less reliable for comparative analysis. Nonetheless, taking into account the log-likelihood and p-value outcomes and noting that the left segment of the observed percentiles line aligns well with that of the MJDM, leads to the conclusion that the MJDM model is the most suitable choice.

Regarding CMZ, with the findings presented in Figure 54 and Table 21, analogous conclusions can be drawn as with CMX. However, it is pertinent to note that the number of observations in the case of CMZ is even more limited. Nevertheless, despite the reduced dataset, the MJDM continues to emerge as the optimal fit for the observed data. This model adeptly captures the characteristic of having a relatively high frequency of months with minimal price changes.

Concerning CCA, the outcomes presented in Figure 55 and Table 22 exhibit a degree of ambiguity. While the log-likelihood scores marginally favour the GBM, the CvM p-value indicates that the MJDM is the superior fit. It is worth noting that the p-value for the GBM-hypsec model also exceeds the 5% significance level. Nonetheless, considering the MJDM's superior accommodation of the high frequency of small price changes, this model has been selected as the most appropriate fit.

In the case of GEO (as illustrated in Figure 56 and Table 23) and NGEO (depicted in Figure 59 and Table 25), the outcomes for the MJDM and GBM-hypsec are closely aligned, rendering it challenging to make a definitive selection between these models.

The model fitting results for ICEVER, as presented in Figure 57 and Figure 58, and Table 24, distinctly convey that the MJDM is most suitable for capturing the characteristics of daily returns. Conversely, when considering weekly returns, the standard GBM model is found to provide a more accurate representation. A comparable conclusion is derived for RGGI (Figure 60 and Figure 61, and Table 26), with the exception that the MJDM is identified as the more appropriate model for describing weekly returns, while the standard GBM is found to be optimal for representing monthly returns.

The best description for the daily returns of the UKETS is provided by the MJDM. This assertion is supported by a notable alignment of the density function, a higher log-likelihood score, and a robust CvM p-value. Correspondingly, the outcomes for weekly returns, as depicted in Figure 62, Figure 63, and Table 27, indicate that the standard GBM is fitting due to the adherence to a normal distribution.

Lastly, the findings for WCA presented in Figure 64 and Table 28 pose challenges in interpretation due to the notably low number of observations. However, the relatively high frequency of minor weekly price changes observed suggests that the MJDM is a more suitable choice, given its ability to better capture this observed behaviour. The above-described selection of price simulation models for the various carbon market components and sampling intervals are summarized in 0.

TABLE 10 *Selected price simulation models for market components in Table 11.*

Market	Daily	Weekly	Monthly
NZETS	X	MJDM	MJDM
CMX	X	X	MJDM
CMZ	X	X	MJDM
CCA	X	MJDM	X
EUETS	MJDM	GBM-hypsec	GBM
GEO	X	MJDM or GBM-hypsec	X
ICEVER	MJDM	GBM	X
N-GEO	X	MJDM or GBM-hypsec	X
RGGI	X	MJDM	GBM
UKETS	MJDM	GBM	X
WCA	X	MJDM	X

4.6 FORWARD VS SPOT PRICE SIMULATION

Initially, it was contemplated that there should be two distinct approaches for simulating forward prices: one tailored for markets with stored historical forward curves and another for markets where only historical spot prices are available. Chapter 3 explored the Schwartz-Smith model and PCA as forward curve simulation options, and these approaches were subsequently applied to carbon markets in this chapter. However, the findings revealed that, for most carbon markets, the forward curves exhibit complete linearity, with minimal observed changes in curvature over time. Only a few carbon markets displayed more pronounced curvature or changes in slope, albeit still relatively limited.

The most significant alteration in the slope of the forward curve was identified in the RGGI market. Nonetheless, 96% of the observed variance in the forward curves for this market was accounted for by changes in spot prices. To substantiate the decision to model forward prices solely with a spot price model, a comparative analysis was conducted between the results of the PCA forward curve simulation approach (utilizing the first three PCs) and the results obtained from the GBM model.

Specifically, the disparity between the 95th percentile predicted by both simulation approaches was documented for both a 1-month and 70-month RGGI future, considering a time step of 10 years. It is noteworthy that the spot price was inflated with a contango inflation percentage, derived from the most recent forward curve.

The disparities observed between the predicted forward price 95% percentiles from the PCA forward curve simulation approach and the GBM model ranged between 3.5% and 5% (across maturities). A parallel analysis was conducted for other carbon markets, revealing similar or even lower differences. Given that these predictions pertain to prices simulated for a 10-year horizon, it is reasonable to infer that the spot price modelling approach is viable and applicable across all carbon markets.

In addition to the aforementioned observations, there are practical considerations that discourage the use of the PCA simulation approach (or Schwartz-Smith model) for modelling the forward curve. Firstly, both methods necessitate the historical forward curves in a matrix format devoid of empty entries. However, as elucidated in Chapter 4, not all maturities have quoted forward prices on each historical date in practice, requiring interpolation for the PCA simulation approach and the Schwartz-Smith model to function. Furthermore, as expounded in Chapter 4, the length of historical forward curves stored in Shell's database undergoes notable changes, mandating extrapolation of the forward curves.

The current inter- and extrapolation methods are linear, assuming straight forward curves. Yet, if market dynamics shift (and curvature emerges), linear extrapolation becomes invalid, necessitating alternative extrapolation techniques. It's imperative to note that in such cases, the spot price approach may no longer be suitable either, introducing model risk. Therefore, the shapes of the forward curves must be systematically and regularly assessed.

The conventional PCA shocking technique presupposes a normal distribution for each PC, akin to the standard GBM. However, actual returns for specific markets and timeframes can deviate substantially from a normal distribution. As we also concluded for the GBM, a potential alternative is to consider a different distribution for the random term to shock the PCs with. Given the time constraints and scope of this study, it appears to be more promising to explore the spot price simulation with the GBM-hypsec and MJDM, which can also take into account non-normal return distributions.

Shaking the forward curve with PCA, possibly incorporating a non-normally distributed randomness factor, presents an avenue for future research. This becomes particularly relevant if forward curves exhibit increased curvature or slope changes in the future. In such a scenario, reconsideration of the intra- and extrapolation methods for the futures matrix becomes imperative. For instance, polynomial interpolation and extrapolation may be more suitable. In this context, the PCA analysis could be executed with lower weights assigned to the intra- and extrapolated values to accord more significance to the original data.

5 PFE MODEL BUILDING

In this chapter, we try to answer the third research question: ''What is the common backbone in the different carbon contracts and how can this be incorporated into the structure of the PFE model?''. As a start, the chapter provides an explanation regarding the incorporation of the drift term in the PFE model and the methodology for calculating volatility in the GBM model. Subsequently, an exploration is conducted into various carbon deal structures, along with optionalities and their impact on the structural composition of the PFE model.

5.1 DRIFT & VOLATILITY

As briefly mentioned in Chapter 2, the drift term of a SDE can yield a significant impact on the PFE over the long term. This influence is also acknowledged in the literature concerning long-term VaR, contrasting with one-day VaR where the drift is typically excluded by definition (Dowd et al., 2004). Many of the carbon market prices examined in this thesis exhibit a discernible upward or downward trend, leading to substantial positive or negative annualized mean returns. Extrapolating these returns into the future for 5 or 10 years may not be reasonable. Moreover, the elevated drifts observed in carbon markets may stem from market perceptions that are no longer relevant. Consequently, we deliberately decide to omit the drift term in the parameter estimation process for the spot price model by subtracting the mean historical return from each individual historical return. This consideration has already been factored into the results presented in Chapter 4.

In practical scenarios, when calculating long-term VaR, the risk-free rate is occasionally employed as a pragmatic estimate for the drift term. This practice is exemplified by the Undertakings for Collective Investment in Transferable Securities (UCITS) regulation of the Committee of European Securities Regulators (CESR) (Committee of European Securities Regulators, 2009). Consequently, the PFE model developed in this thesis provides the flexibility for users to manually input a drift term.

Moreover, it is needed to delve into the methodology for calculating the variance in the context of the standard GBM. While the volatility for other models is determined during the parameter estimation process, a distinct approach is adopted for the standard GBM. We decide to compute the volatility using the exponential weighted moving average (EWMA) with a weighting factor of 0.94, as generally used by companies in a financial risk context (J.P.Morgan/Reuters, 1996)

In this thesis, the EWMA calculation is applied on the same time frame as used to calculate the historical returns. Therefore, if the spot price model parameter estimation is conducted with daily returns, the EWMA calculation utilizes daily volatilities. Conversely, if the parameter estimation is performed with monthly returns, monthly volatilities are employed in the EWMA calculation. The sole criterion is the availability of at least 80 historical volatilities for the EWMA filtering. In cases where this criterion is not met, an average of the volatilities is computed.

5.2 DEAL STRUCTURES

To develop a flexible PFE model in Python capable of accommodating diverse carbon deals within Shell and, if necessary, being easily adaptable to include other structures or optionalities, an examination of a number of carbon deals is undertaken. Due to practical limitations, as of the current writing, this overview remains incomplete, with documentation only received for a subset of the global carbon portfolio.

The deals that have been reviewed involve forward exposures wherein Shell either receives or holds the obligation to deliver carbon allowances/credits at a future date, potentially adjusted by a predetermined value referred to as the forward strike price. Additionally, it is anticipated that contracts may incorporate swap agreements, facilitating the exchange of one type of carbon allowance/credit for another. Furthermore, embedded put and/or call options can be present in these contracts.

5.2.1 Forward Prices

To start, let us focus on explaining the process of forward price modelling in Python. These modelled prices find their place in what is referred to as the "Simulated Prices" matrix. In this matrix, every row corresponds to a specific date designated for PFE calculation (calculation date). For instance, when for a 1-year deal the PFE is calculated monthly, the *Simulated Prices* matrix contains twelve rows. Each column in the matrix represents the maturity date of the forward price, thus aligning the spot prices on the diagonal due to the calculation date coinciding with the maturity date.

As discussed in Section 3.4**Fout! Verwijzingsbron niet gevonden.**, the spot pot price can be simulated with the selected price simulated model at each calculation date. Following this, the forward prices are computed by applying the contango percentage derived from the latest forward curve to inflate the price. Specifically, only those forward prices are computed for which a forward exposure is present in the contract. These spot and forward prices are subsequently inserted into the Simulated Prices matrix, as illustrated in Figure 27, with the spot prices occupying the diagonal and the forward prices populating the upper triangle of the matrix.

	T_1	T_2	$\cdots$	T_X
t_1	S_{t_1}	F_{t_1,T_2}	$\cdots$	F_{t_1,T_X}
t_2		S_{t_2}	$\cdots$	F_{t_2,T_X}
$\vdots$			$\ddots$	$\vdots$
t_X				S_{t_X}

FIGURE 27 *Spot and forward price matrix in Python PFE model. S_t indicates spot price at time t, and $F_{t,T}$ indicates forward price at time t with maturity T.*

5.2.2 Option Prices

When options are included in a contract, it is required to calculate option prices at each calculation date. The method chosen for computing these option prices hinges on how the spot prices are simulated. If the spot prices are generated using the GBM, the Black model comes into play. This model is a modified version of the Black-Scholes option pricing model, where the spot price is substituted with the forward/futures price (Black, 1976). The equations governing the Black model can be found in APPENDIX G.

When the spot prices are simulated using the MJDM, a different option pricing model can be used, which is specifically designed for this price simulation model. This option pricing model can be found in APPENDIX G as well. In Chapter 6, a dedicated analysis will be conducted to compare the results of the Black and Merton MJDM option pricing models, with an emphasis on their application within the carbon markets. It was decided to use the Black option pricing model when the GBM-hypsec model is used to simulate spot prices. There is no dedicated option pricing model for the GBM-hypsec spot price model and the alternative approach of using a Monte Carlo simulation with the GBM-hypsec model to simulate forward prices poses significant computational challenges.

One of the parameters in the Black model is the risk-free rate. In the PFE model, the interest rate term structure needs to be given as input. This allows to value each option accurately by employing the interest rate corresponding to its maturity. When iteratively calculating the PFE at calculation dates over the deal tenor, leading to a reduction in the time to maturity for the options, the progression involves transitioning from the right to the left in the yield curve. In this manner, the T-year maturity aligns with the T-year interest rate. For future model developments, it is recommended to extract the interest rate term structure from a database such that it does not have to be provided as a manual input.

Moreover, the Black model requires volatility as an input, with a preference for the utilization of implied volatility. However, in numerous carbon markets, the absence of a well-established market for option prices contributes to a scarcity of available implied volatility quotes. Such quotes are only prevalent in select markets, such as the EUETS and UKETS. Even when these quotes are accessible, the range of strike prices and maturities may be constrained. Due to this fragmentation, the decision has been made to employ historical GBM volatility as the input for the option pricing model across all carbon markets. It is worth noting that future research endeavours could explore the incorporation of implied volatility quotes for option pricing in markets where they are available.

5.2.1 Swaps

A distinct model is formulated for swap deals, where instead of simulating one price, this process is undertaken for two assets. The MJDM encompasses two separate sources of randomness: the Wiener process (drift) and the Poisson process (jumps). Establishing relationships between both sources of randomness for two assets is not straightforward. For the GBM, this can be accomplished relatively easily by determining the correlation coefficient from historical returns, as the Wiener process is the sole source of randomness. Given that, as will be elucidated in Chapter 6, the MJDM offers limited value for long-term carbon deals, the PFE carbon model for swaps adopts the GBM as the sole price simulation model.

The correlation coefficient is computed using Pearson's method, which involves dividing the covariance of the price returns of two carbon credits/allowances by the product of the standard deviation of the returns of both assets. The price returns used in the correlation calculation should be sampled at the same time interval as that intended for simulating the prices. For instance, if the PFE is computed on a monthly basis (spot prices simulated with a monthly time step), the correlation coefficient should be determined using monthly returns. Subsequently, the Wiener processes within the GBM for both carbon credits/allowances can be derived from Equation (5.1).

$$dW_1 dW_2 = \rho_{1,2} dt \tag{5.1}$$

5.3 PFE CALCULATION

In each iteration of the MC simulation, a MtM value is calculated for the various remaining future exposures (forwards, options, or swaps) at each calculation date. This MtM is the summation of the individual exposures, where negative exposures can offset positive ones, reflecting the netting aspect. After conducting a specific number of runs in the MC simulation, a distribution of the MtM values for the remaining exposure at each calculation date is generated. By selecting the 95^{th} percentile of the positive MtM values at each calculation date, the Potential Future Exposure (PFE) is determined for that specific date, as depicted in Figure 4 and Equation (2.1). The 95% peak PFE represents the highest 95% PFE observed over the duration of the deal.

6 MODEL VALIDATION

The concluding phase of the PFE model's development involves validation. This chapter begins the validation process by primarily focusing on scrutinizing the price simulation engine, as this will drive the exposures in the PFE calculations. Subsequent steps include a review of the option pricing possibilities. Additionally, the model undergoes testing against a carbon deal, for which a basic PFE model had been previously developed within Shell, facilitating a comparative analysis of this PFEs.

6.1 PRICE SIMULATION

The price simulation engine employed in the PFE model utilizes Equations (3.3) and (3.19) for simulating discrete price movements in the case of the GBM and MJDM, respectively. For the GBM-hypsec model, the Wiener process in Equations (3.3) is replaced by a process which samples random variables from a hyperbolic secant distribution. The returns projected by these equations, over a specified time horizon, are compared with the theoretical return distributions outlined in Equations (3.2) and (3.18) for both models. The results are presented in the right plots of Figure 28, Figure 29, and Figure 30, where arbitrarily chosen parameters and a 10-year time horizon are utilized. A close correspondence between the theoretical returns and the returns derived from simulated prices is evident.

Given the known theoretical return distribution, it is possible to derive the theoretical price distribution for the GBM, GBM-hypsec and MJDM. An additional validation step that was performed included to compare the theoretical 95% confidence interval (CI) with the same confidence interval determined from 10,000 simulated price paths for the three models. The results of this comparison are presented in the left plots in Figure 28, Figure 29, and Figure 30 for three models, demonstrating a satisfactory alignment between theoretical and simulated values.

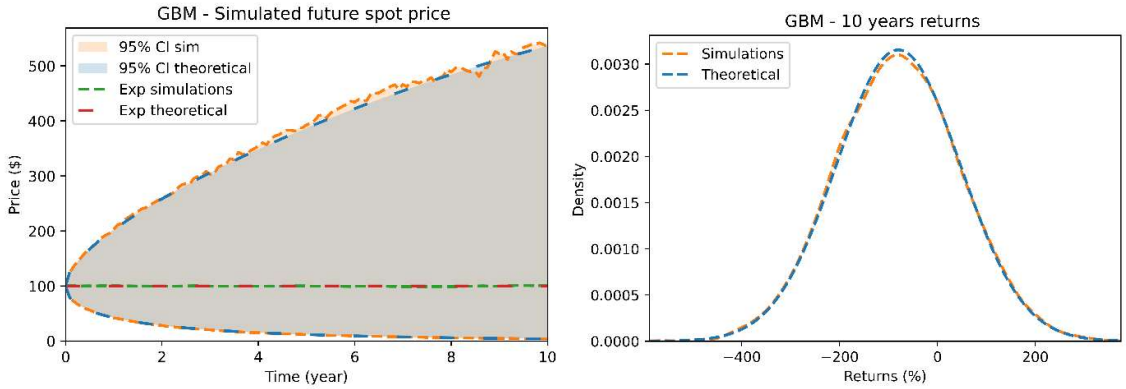

FIGURE 28 *GBM price simulation validation with 10000 simulations. Left) 10-year price return distribution, right) expected value and standard deviation of price paths.*

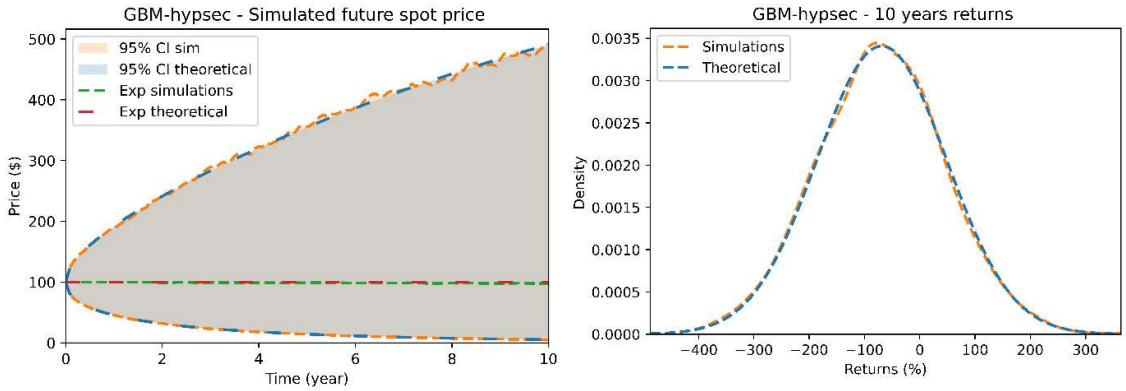

FIGURE 29 *GBM-hypsec price simulation validation with 10000 simulations. Left) 10-year price return distribution, right) expected value and standard deviation of price paths.*

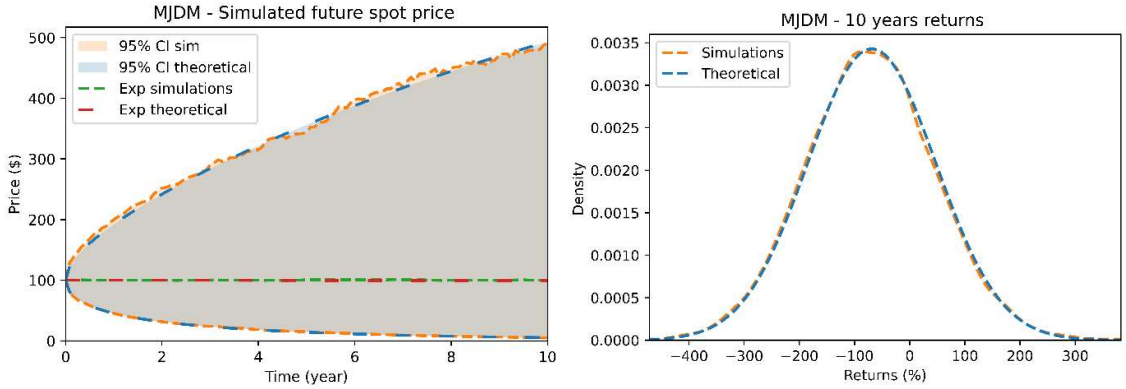

FIGURE 30 *MJDM price simulation validation with 10000 simulations. Left) 10-year price return distribution, right) expected value and standard deviation of price paths.*

6.1.1 Normal Distribution Convergence

The question may arise as to why the return distribution of the GBM-hypsec model in Figure 29 and the MJDM in Figure 30, appear normal, while the GBM-hypsec and MJDM distributions derived from the parameter estimation process in Section 4.5 exhibit non-normal characteristics. The explanation relates to the time horizon over which the returns are simulated. As illustrated in Section 4.5, when simulating returns with these models over a short period (such as one day, week, or month), the resulting returns deviate from normality, displaying a leptokurtic shape. However, with an extended time horizon for price simulation, the log returns gradually converge to the distribution outlined in Equation (3.18).

The required time horizon for the GBM-hypsec and MJDM return distributions of the final price to converge to the theoretical normal distribution is contingent upon the specific parameter values employed in the model. For certain parameter values, this convergence may take a few years, while in other cases, only a few months suffice. Notably, the investigation revealed that, for all carbon markets considered in this thesis, a simulation time horizon of five to six months, at most, was adequate for the return distribution of the final price to converge to the expected normal version. This finding is significant, indicating that the GBM-hypsec model and MJDM do not forecast non-normal return distributions on an annual interval for the carbon markets.

An even more intriguing observation relates to the shape toward which the returns of the MJDM simulations for the carbon markets converge. Remarkably, it was found that, for each carbon market examined in this thesis, the MJDM return distribution converges to the GBM distribution. Since, the

convergence happens already after a very short time horizon, there is little to no difference between the distribution of the prices simulated by the MJDM and the standard GBM for the carbon markets that have been reviewed in this thesis. The only exception is when the prices are simulated for a few days or weeks into the future, as in that case the MJDM produces non-normal return distributions. Figure 31 to Figure 34 illustrate the corresponding 95% confidence intervals for the GBM, GBM-hypsec, and MJDM, and their 10-year final price return distribution. The analysis is conducted using four observed market datasets, highlighting the observable overlap between the GBM and MJDM.

The return distribution and the resulting 95% price path confidence interval of the GBM-hypsec model were found to not consistently overlap with the other two models, as evidenced in Figure 31 to Figure 34. When fitted on the EUETS market data, the three models produced similar outcomes (Figure 32), likely owing to the high level of normality in the observed market data. However, for the other markets presented in the figures below, the return distribution and price path confidence intervals of the GBM-hypsec model diverged from those generated by the other models. Notably, in simulations for the NZETS markets, these differences were most pronounced, while deviations for the RGGI and CCA markets remained relatively modest.

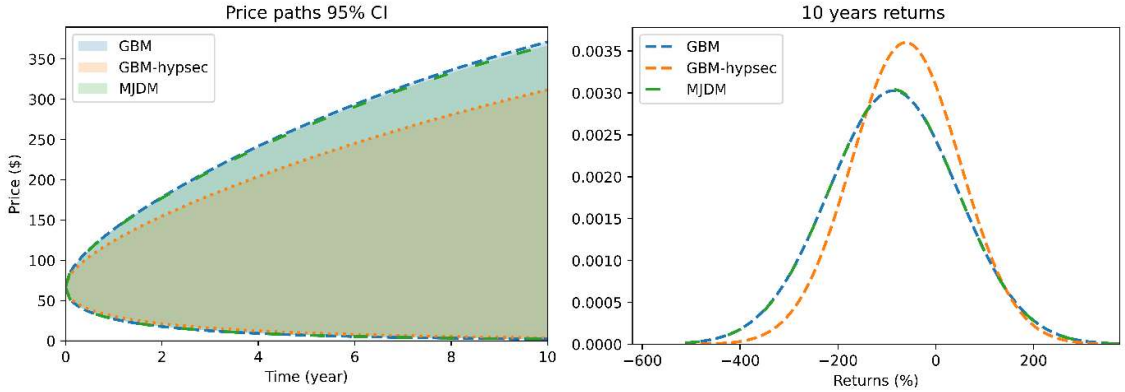

FIGURE 31 *Return distribution and price path 95% confidence interval for NZETS fitted on monthly data.*

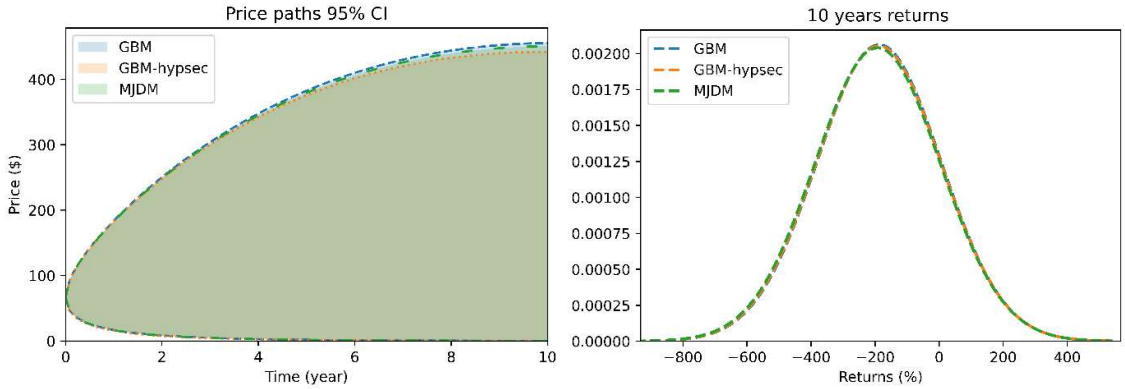

FIGURE 32 *Return distribution and price path 95% confidence interval for EUETS fitted on monthly data.*

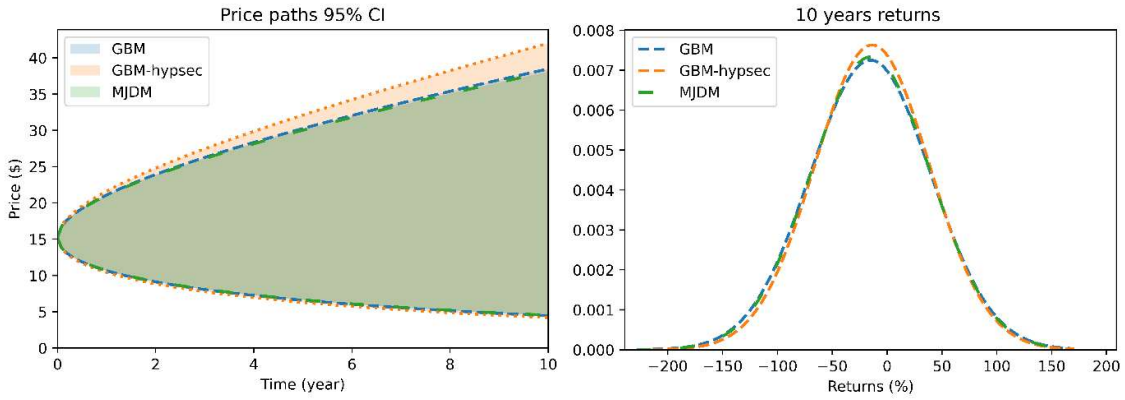

FIGURE 33 *Return distribution and price path 95% confidence interval for RGGI fitted on monthly data.*

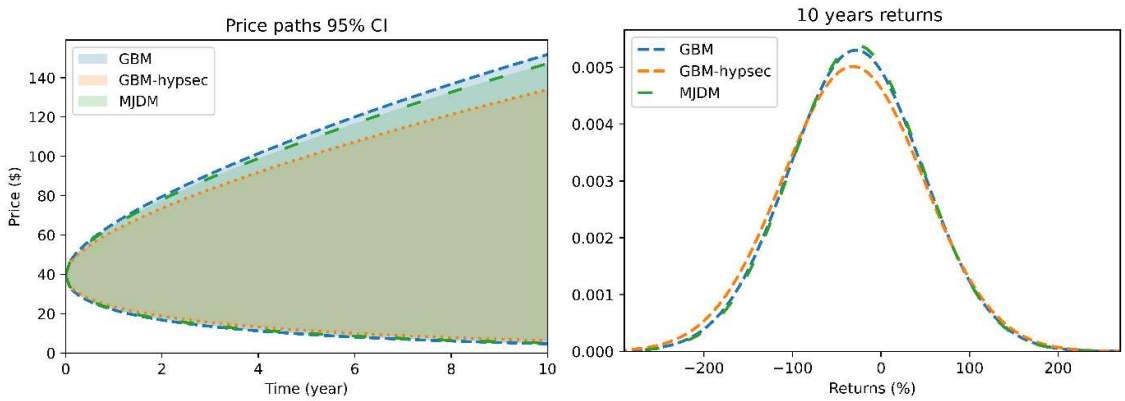

FIGURE 34 *Return distribution and price path 95% confidence interval for CCA fitted on weekly data.*

6.2 OPTION PRICING

For the MJDM an option pricing model has been developed in the literature, which is given in APPENDIX G (Cont & Tankov, 2002). Since the GBM model and MJDM produce similar price paths when fitted on carbon market data, the expectation is that the option prices predicted by the Black model and the MJD option pricing model give similar results too. This was confirmed by calculating option prices under both option pricing models, using the GBM and MJDM parameters estimated from carbon market data. This was done for all carbon markets, and the results for the NZETS market are given in Figure 35. For the other markets, very similar results were obtained, which is not surprising given the observed overlap between the GBM and MJDM in the previous section.

For ease of computation, the initial stock price was standardized to $100 for both markets, resulting in the intersecting lines for call and put option prices around this value. Notably, the call and put option prices under both models exhibit nearly identical values, even for a short time to maturity for which the MJDM returns differ more from those given by the GBM. The MJD option pricing model involves higher computational demands due to the summation of option prices derived from the Black model with adjusted parameters (see Equation (8.28) and (8.29)). Considering that the option prices from the MJD option pricing model closely align with those from the Black model, it was determined to exclusively utilize the Black model in the final PFE model. Consequently, in instances where carbon prices are simulated using the MJDM, the eventual option prices are computed using the Black model with the usage of GBM volatility.

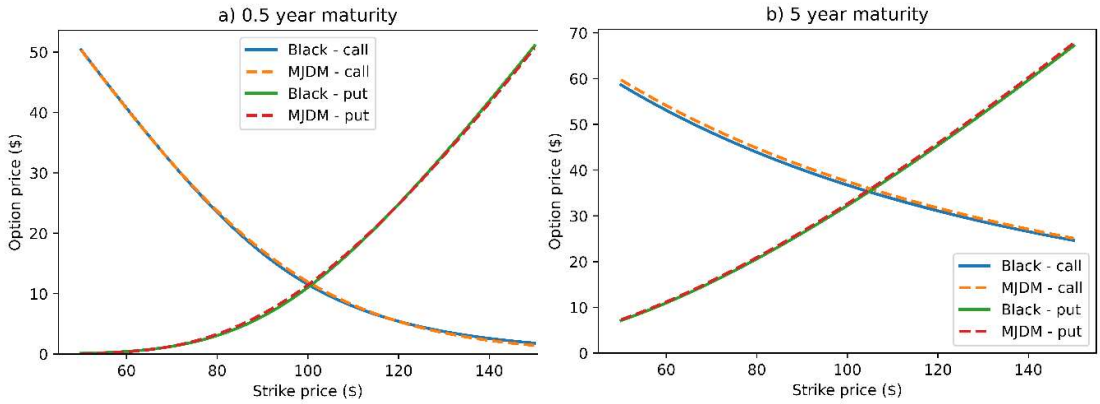

FIGURE 35 *Black and MJDM option prices versus strike price for model parameters estimated on NZETS market data.*

6.3 CARBON DEAL PFE COMPARISON

A team engaged in the trading of carbon credits and allowances within Shell had previously developed a basic PFE model for a carbon deal featuring embedded forwards and options. In this model, subsequently referred to as the Excel PFE model, the forward prices are assumed to follow a lognormal distribution (consistent with the GBM model) with a constant volatility. The construction of the forward curve follows a comparable approach to the one followed in this thesis, albeit with slight variations in the precise calculation. Let us first investigate the differences between the Excel PFE model and the one developed in this thesis, before trying to compare their respective results.

To illustrate this, consider the simulation of the T-year forward price at time t ($F_{t,T}$), assuming a contango assumption of 3%, with the current spot price denoted as S_0. In this thesis, the methodology involves simulating the future spot price with the t-year time step and then inflating this price by $(1 + 0.03 * (T - t))$. Conversely, in the Excel model, the process entails initially inflating the spot price by $(1 + 0.03)^t$ before utilizing the GBM to simulate the spot price after one year. After this simulation, the spot price is then further inflated by $(1 + 0.03)^{T-t}$.

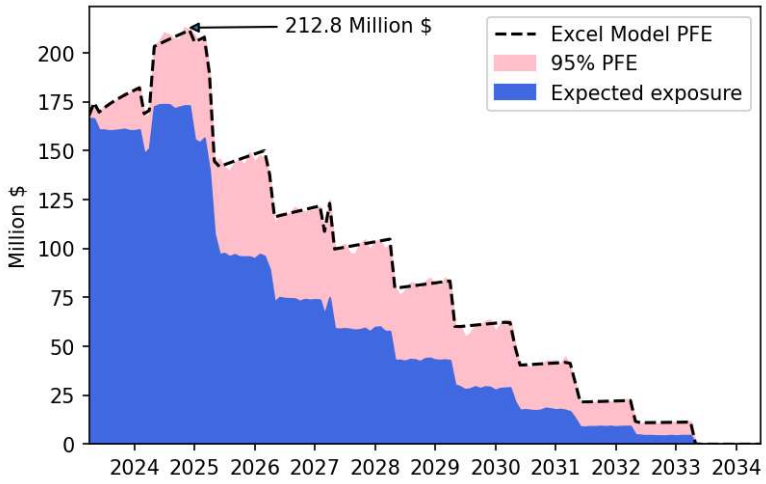

FIGURE 36 *PFE profiles given by model develop in this thesis and Excel model with flat forward curve and without options. Peak PFE in Excel model = 209.8 million.*

This means there are two distinctions in the methods used for constructing the forward curve. To start, the GBM is applied on different price levels: in the model developed in this thesis, it operates

on S_0, while in the Excel model, it operates on $S_0 * \left(1 + i_{contango}\right)^t$. Secondly, the forward curve is presumed to be straight in the model presented in this thesis, whereas it exhibits a slight curvature in the Excel model owing to the compounding effect.

If the Excel model is used for comparison with the PFE model developed in this thesis, we must therefore assume the forward curve to be flat, considering the disparities in the construction approaches of the forward curve. Figure 36 presents the outcomes from both the Excel model and the model developed in this thesis for a specific carbon deal including options. The results from both models exhibit a notable alignment. It was also verified that if the forward curves in the model developed in this thesis are constructed similarly to those in the Excel model, the results are consistent. In such a scenario, the PFE profile is higher on average due to the higher forward prices on average.

7 CONCLUSION & FUTURE RESEARCH

The original research goal we set out is: "*Develop a Potential Future Exposure model for contracts traded on the carbon markets*". We draw our conclusions in Section 7.1. We then make several recommendations for further research in Section 7.2.

7.1 CONCLUSION

First of all, our investigation into the development of a PFE model has unveiled the diverse array of carbon allowances and credits encompassed within the carbon market spectrum. A notable distinction emerges between the regulatory (or compliance) and voluntary carbon markets. Voluntary carbon markets typically exhibit significantly lower prices per metric tonne of carbon dioxide compared to their regulatory counterparts. Furthermore, price liquidity and transparency are generally diminished within these voluntary markets.

We analysed historical carbon market prices and observed a division based on the type of internally stored data at Shell. In the first category, historical forward curves are stored as a whole, while the second category is characterized by the availability of only historical spot prices. Consequently, we examined distinct modelling approaches for each group. However, upon analysing the historical forward curves, in those markets where such data are accessible, it was revealed that the slopes have undergone minimal changes over time and that they exhibit almost negligible curvature. PCA unveiled that, in the reviewed carbon markets, at least 96% of the observed historical variance could be ascribed to fluctuations in the spot price.

The employed models (Schwartz-Smith and PCA simulation) to directly simulate the forward curve, did not offer substantial advantages over the approach wherein the forward curve is linearly extrapolated from a simulated spot price. Additionally, both forward curve modelling approaches suffer from practical limitations due to incompleteness of the historical forward curves that are stored. Consequently, based on these findings, we concluded that the spot price modelling approach is equally applicable for markets with available historical forward curve data as it is for markets with only available spot prices.

After examining a limited number of carbon contracts, we developed a PFE model capable of utilizing the Geometric Brownian Motion (GBM), GBM-hypsec model, and Merton Jump Diffusion Model (MJDM) for price simulation. We propose the GBM-hypsec model as a variation of the normal GBM, where the random normal variable is replaced by a hyperbolic secant random variable to better capture the leptokurtic behaviour of observed return distributions. Similarly, the MJDM replicates the leptokurtic shape of the return distribution by introducing jump effects. This PFE model can be used for forward sale and swap deals, possibly combined with vanilla options, and is designed for potential extensions in future research.

An important observation relates to the central limit theorem (CLT) in combination with the simulated prices using the three aforementioned models. We saw that both GBM-hypsec and MJDM result in a leptokurtic (non-normal) price distribution in the short term (a few months into the future). However, when these models are employed to simulate longer-term price paths, the return distribution rapidly converges to a normal distribution due to the CLT, as the monthly (or daily/weekly) returns are drawn from an identically and independently distributed random variable. As a result of the above, the advantage of GBM-hypsec model and MJDM in capturing the leptokurtic daily/weekly/monthly return distribution observed in certain markets diminishes over time. Additionally, it was noted that MJDM yields nearly the same price distribution as the GBM at a longer time horizon, rendering it redundant compared to GBM. Nevertheless, our PFE model

incorporating the GBM-hypsec model and MJDM does offer Shell an opportunity to more accurately model short-term price movements, albeit beyond the scope of PFE considerations.

7.2 FUTURE RESEARCH

As an initial direction for future research, it may be valuable to incorporate supplementary spot price models into the parameter estimation tool, thereby enabling the evaluation of the performance of a broader spectrum of models. For instance, there are various other jump diffusion models, like the Kou model, where jump sizes adhere to a double-exponential distribution. It would be interesting to investigate whether the price distribution derived from alternative jump-diffusion models converges to the GBM distribution or exhibits distinct characteristics.

Subsequently, given that the PFE model developed in this framework relies on forward prices as risk factors, it becomes impractical to use the model for very illiquid carbon markets. Consequently, for deals in such markets, an alternative model becomes imperative if there is an intention to compute the PFE. This necessitates the identification of specific risk factors for each unique carbon market, with their behaviour being modelled over time in accordance with a probability distribution. Alternatively, scenarios for each risk factor can be formulated, leading to a scenario-based stress test for a carbon deal. However, the latter approach falls short of producing a PFE, as the latter necessitates a probability distribution for the risk factors.

In this thesis, we estimate the model parameters by maximizing the log-likelihood. Another viable parameter estimation technique is the method of moments, wherein the moments of the observed data are aligned with the moments derived from the models by adjusting parameter values. For certain models, least squares fit is also conceivable by reformulating the model's stochastic differential equation (SDE) in a manner that enables a linear fit, facilitating the estimation of parameters. An intriguing avenue for further research would involve exploring the impact of different estimation techniques on parameter values, particularly investigating whether employing an alternative estimation technique could alter the observation that the Merton Jump Diffusion Model's (MJDM) price path distribution converges to that of the GBM over a time horizon of a few months.

If the slope and curvature of the forward curves for the carbon markets undergo changes over time, a different modelling approach may be necessitated. In that case, the PCA forward curve modelling method is recommended over the Schwartz-Smith model. The interpolation and extrapolation technique should be re-evaluated, considering that the linear approach may no longer be suitable. A potential solution could involve developing a PCA variation that assigns less weight to interpolated and extrapolated data points, thereby prioritizing the actual observed forward prices. Additionally, exploring the use of a different distribution to shock the principal components (similar to the GBM-hypsec model) could be considered if the aim is to replicate short-term leptokurtic return behaviour. Alternatively, a more straightforward curve modelling approach could involve shocking each point (or segment) on the curve with the GBM or another model, allowing for the capture of volatility differences across different segments of the curve.

The current PFE model enables the manual input of a drift term, which gives the option to set it equal to the risk-free rate as an estimate for the long-term drift. Moreover, the risk-free rate is utilized as an input for option pricing calculations. Ideally, both parameters would be derived from an interest rate term structure maintained within Shell. This prompts the question of the specific interest rate term structure to be employed and the methodology for selecting the accurate interest rate for options with varying time to maturities. A more comprehensive model would also account for changes in the interest rate term structure over time.

8 BIBLIOGRAPHY

Alberola, E., Chevallier, J., & Chèze, B. t. (2008). Price drivers and structural breaks in European carbon prices 2005–2007. *Energy Policy, 36*(2), 787-797. doi:10.1016/j.enpol.2007.10.029

Autoriteit Financiële Markten. (2023). *Voluntary Carbon Markets - Supervisory issues*. Retrieved from https://www.afm.nl/~/profmedia/files/rapporten/2023/occasional-paper-handel-in-co2.pdf

Basel Committee on Banking Supervision. (2014). *The standardised approach for measuring counterparty credit risk exposures*. Retrieved from https://www.bis.org/publ/bcbs279.pdf

Basel Committee on Banking Supervision. (2021). Revisions to the Principles for the Sound Management of Operational Risk.

Batten, J. A., Maddox, G. E., & Young, M. R. (2021). Does weather, or energy prices, affect carbon prices? *Energy Economics, 96*. doi:10.1016/j.eneco.2020.105016

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics, 3*(1), 167-179. doi:10.1016/0304-405X(76)90024-6

Blanco, C. (2002). Multi-factor models for forward curve analysis: An introduction to principal component analysis.

Borovkova, S., & Geman, H. (2009). Forward Curve Modelling in Commodity Markets. In *Risk Management in Commodity Markets* (pp. 9-32).

Bredin, D., & Muckley, C. (2011). An emerging equilibrium in the EU emissions trading scheme. *Energy Economics, 33*. doi:10.1016/j.eneco.2010.06.009

CarbonCredits.com. (2023). The Collapse of NGEO Carbon Prices: An In-depth Analysis. Retrieved from https://carboncredits.com/the-collapse-of-ngeo-carbon-prices-an-in-depth-analysis/

CarbonNews. (2024). MARKET LATEST: NZUs. Retrieved from https://www.carbonnews.co.nz/story.asp?storyID=29712

CME group. (2023). CBL Nature-Based Global Emissions Offset Futures - Quotes Retrieved from https://www.cmegroup.com/markets/energy/emissions/cbl-nature-based-global-emissions-offset.quotes.html#venue=globex

Committee of European Securities Regulators. (2009). *Annex to CESR's technical advice on the level 2 measures related to the format and content of Key Information Document disclosures for UCITS (Ref. CESR/09-949)*. Retrieved from https://www.esma.europa.eu/sites/default/files/library/2015/11/09_1026_final_kid_srri_methodology_for_publication.pdf

Cont, R., & Tankov, P. (2002). Calibration of Jump-Diffusion Option Pricing Models: A Robust Non-Parametric Approach. *SSRN Electronic Journal*. doi:doi: 10.2139/ssrn.332400

Derindere Köseoğlu, S., Ead, W. M., & Abbassy, M. M. (2022). Basics of Financial Data Analytics. In S. Derindere Köseoğlu (Ed.), *Financial Data Analytics: Theory and Application* (pp. 23-57). Cham: Springer International Publishing.

Dhesi, G., Shakeel, B., & Ausloos, M. (2021). Modelling and forecasting the kurtosis and returns distributions of financial markets: irrational fractional Brownian motion model approach. *Annals of Operations Research, 299*(1). doi:doi: 10.1007/s10479-019-03305-z

Dowd, K., Blake, D., & Cairns, A. (2004). Long-Term Value at Risk. *The Journal of Risk Finance, 5*(2), 52-57. doi:doi: 10.1108/eb022986

Ecosystem Marketplace. (2022). *Ecosystem Marketplace's State of the Voluntary Carbon Markets 2022 Q3*. Retrieved from https://www.ecosystemmarketplace.com/publications/state-of-the-voluntary-carbon-markets-2022/

Goodwin, D. (2013). *Schwartz-Smith Two-Factor Model in the Copper Market: before and after the New Market Dynamics (MSc thesis).*

Gregory, J. (2012). *Counterparty Credit Risk and Credit Value Adjustment: A Continuing Challenge for Global Financial Markets*: Wiley.

Gugole, N. (2016). Merton jump-diffusion model versus the black and scholes approach for the log-returns and volatility smile fitting. *International Journal of Pure and Applied Mathematics*. doi:10.12732/ijpam.v109i3.19

Hull, J. C. (2018a). *Options, Futures, and Other Derivatives* (9th Ed.).

Hull, J. C. (2018b). *Risk Management and Financial Institutions*: Wiley.

International Energy Agency. (2021). Greenhouse Gas Emissions from Energy Highlights. Retrieved from https://www.iea.org/data-and-statistics/data-tools/greenhouse-gas-emissions-from-energy-data-explorer

J.P.Morgan/Reuters. (1996). *RiskMetrics (TM)—Technical Document*. Retrieved from https://www.msci.com/documents/10199/5915b101-4206-4ba0-aee2-3449d5c7e95a

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering, 82*(1), 35-45. doi:10.1115/1.3662552

Konlack Socgnia, V., & Wilcox, D. (2014). A Comparison of Generalized Hyperbolic Distribution Models for Equity Returns. *Journal of Applied Mathematics*. doi:doi: 10.1155/2014/263465

Kreinin, A., Merkoulovitch, L., Rosen, D., & Zerbs, M. (1998). *Principal Component Analysis in Quasi Monte Carlo Simulation*. Retrieved from http://www.financerisks.com/filedati/WP/ALGO_PAPER/arq-principal.pdf

Lautier, D. (2003). *Term Structure Models of Commodity Prices: A Review*.

Lovcha, Y., Perez-Laborda, A., & Sikora, I. (2022). The determinants of CO2 prices in the EU emission trading system. *Applied Energy, 305*, 117903. doi:10.1016/j.apenergy.2021.117903

Mansanet Bataller, M., Pardo Tornero, Á., & Valor, E. (2007). CO2 Prices, Energy and Weather. *Environmental Economics eJournal*.

Mejía Vega, C. A. (2018). Calibration of the exponential Ornstein–Uhlenbeck process when spot prices are visible through the maximum log-likelihood method. Example with gold prices. *Advances in Difference Equations, 2018*(1), 269. doi:10.1186/s13662-018-1718-4

Mohammad Rafiqul, I. (2018). Sample Size and Its Role in Central Limit Theorem (CLT). *International Journal of Physics and Mathematics, 1*(1), 37-47. doi:doi: 10.31295/ijpm.v1n1.42

Myung, I. J. (2003). Tutorial on maximum likelihood estimation. *Journal of Mathematical Psychology, 47*(1), 90-100. doi:https://doi.org/10.1016/S0022-2496(02)00028-7

Nkemnole, B., & Abass, O. (2019). Estimation of geometric Brownian motion model with a t-distribution–based particle filter. *2019, 12*(1). doi:10.4102/jef.v12i1.159

Palao, F., & Pardo, Á. (2021). The inconvenience yield of carbon futures. *Energy Economics, 101*. doi:10.1016/j.eneco.2021.105461

Refinitiv. (2023). *Carbon Market Year In Review 2022* Retrieved from https://www.refinitiv.com/content/dam/marketing/en_us/documents/gated/reports/carbon-market-year-in-review-2022.pdf

Reichsfeld, D. A., & Roache, S. K. (2011). Do Commodity Futures Help Forecast Spot Prices? .

Roncoroni, A., Fusai, G., & Cummins, M. (2015). *Handbook of Multi-Commodity Markets and Products: Structuring, Trading and Risk Management*: Wiley.

S&P. (2021). Voluntary carbon markets: how they work, how they're priced and who's involved. Retrieved from https://www.spglobal.com/commodityinsights/en/market-insights/blogs/energy-transition/061021-voluntary-carbon-markets-pricing-participants-trading-corsia-credits

S&P Global. (2022). Vega risk in potential future exposure. Retrieved from https://www.spglobal.com/marketintelligence/en/mi/research-analysis/vega-risk-in-potential-future-exposure.html

Schwartz, E., & Smith, J. E. (2000). Short-Term Variations and Long-Term Dynamics in Commodity Prices. *Management Science, 46*(7), =. doi:10.1287/mnsc.46.7.893.12034

Seifert, J., Uhrig-Homburg, M., & Wagner, M. (2008). Dynamic behavior of CO2 spot prices. *Journal of Environmental Economics and Management*. doi:10.1016/j.jeem.2008.03.003

Shell plc. (2022). Sustainability Report 2022.

Shell plc. (2024). Emission Trading. Retrieved from https://www.shell.com/business-customers/aviation/aviation-consultancy-services/emission-trading.html

SIX. (2022). *Voluntary Carbon Market*. Retrieved from https://www.six-group.com/dam/download/company/report/whitepapers/six-whitepaper-voluntary-carbon-markets-en.pdf

Staures, M. N. J., Aduda, J. A., & Momeya, R. (2020). Modeling Underlying Assets Log-return in Merton Jump-Diffusion Framework. *Journal of Applied Mathematics & Bioinformatics*.

The Guardian. (2023). Revealed: more than 90% of rainforest carbon offsets by biggest certifier are worthless, analysis shows

Retrieved from https://www.theguardian.com/environment/2023/jan/18/revealed-forest-carbon-offsets-biggest-provider-worthless-verra-aoe

Theecanmole. (2024). New Zealand Emission Unit Prices in the NZ Emissions Trading Scheme 2010 to date. Retrieved from https://github.com/theecanmole/nzu

Valdivieso, L., Schoutens, W., & Tuerlinckx, F. (2009). Maximum likelihood estimation in processes of Ornstein-Uhlenbeck type. *Statistical Inference for Stochastic Processes, 12*(1). doi:10.1007/s11203-008-9021-8

World Bank Group. (2023). *State and Trends of Carbon Pricing 2023*. Retrieved from https://openknowledge.worldbank.org/handle/10986/39796

Zhu, B., & Chevallier, J. (2017). *Pricing and Forecasting Carbon Markets: Models and Empirical Analyses*. Cham: Springer International Publishing.

APPENDIX A. CARBON MARKETS

A.1 CARBON CREDITS: PROJECT PHASES & TYPES

The origin of the VCM lies in the Kyoto Protocol established in 1997. Next to binding emission reduction commitments, the protocol introduced so-called flexible mechanisms. This includes the Clean Development Mechanism (CDM) and the Joint Implementation (JI) Projects, both designed to act at the governmental level. Under the CDM, more developed (Annex I) countries could buy emission reductions, called Certified Emission Reductions (CERs), from developing countries. The JI is similar to the CDM and allowed developed countries to trade Emission Reduction Units (ERUs) with each other. The high administrative burden of the CDM and JI, inability to initiate certain activities, concerns regarding quality, and other shortcomings, lead to the evolution of other carbon mechanisms (Autoriteit Financiële Markten, 2023). These comprise the so-called VCM. In principle, the VCMs are unregulated, and everyone is free to create carbon credits. However, in practice, there are a few large standard-setting players. The largest two are Verra and Gold Standard with an estimated market share of 70% and 15% respectively (Autoriteit Financiële Markten, 2023). The remaining market share is divided between smaller parties such as the Climate Action Reserve and the American Carbon Registry.

A carbon credit project starts with a design document stating the general idea of the project and the methodology of the standard to be followed. The design document must be validated by an accredited auditor to verify the design compared to the relevant methodology, check whether the projected amount of carbon reductions is realistic, and evaluate potential harmful side effects. Annual reports must be submitted which are validated by an accredited auditor (often another auditor than the one involved in the validation stage). If this process is successful, carbon credits are created and issued. This process is also shown in Figure 37.

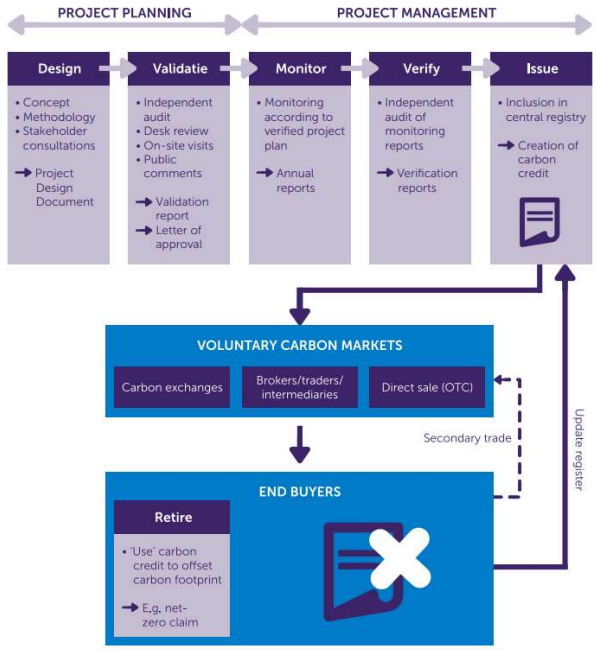

FIGURE 37 *Mechanism of voluntary carbon credits. Taken from (Autoriteit Financiële Markten, 2023).*

Whether the credits are issued frontloaded, backloaded or spread over the project's life cycle differs per project. There can be a delay between achieving the reductions and the creation of credits. One of the main distinctions between the different issue types of carbon credits is whether the project aims at reducing emissions that otherwise would have taken place, or at the removal of carbon

dioxide from the atmosphere. Furthermore, there is a distinction between technology-based and nature-based credits. Technology-based projects utilize technologies such as (direct) carbon capture and storage. On the other hand, nature-based credits are earned through the reduction or removal of emissions by restoring or protecting nature.

The lack of standardisation in the VCM is one of the major issues. The different standard setters (e.g. Verra and Gold Standard) use different methodologies for determining the amount of carbon emission reductions. This causes a range of prices for different carbon credits, which means that there is a lack of granular data in the VCM. Most of the criticism relates to the additionality, which is "the reduction of a certain amount of emissions which would not have occurred without the funding of those credits" (Autoriteit Financiële Markten, 2023). The additionality refers to the deviation of emissions compared to a baseline scenario. Human judgement plays an important role in this non-trivial and complex calculation (Ecosystem Marketplace, 2022). The different methodology of standard setters is subjected to an ongoing debate. There are various examples of projects where reduction are claimed to have occurred without the issue of carbon credits or where the advertised reductions would not have been met (The Guardian, 2023). Furthermore, there is no unified register for carbon credits and hence there are fears that some credits are double counted as they pass hands.

A.2 SHELL CARBON TRADING PORTFOLIO

A.2.1 Regulated Carbon Market

EU Emission Trading Scheme (EUETS)

The oldest ETS currently covers emissions from stationary equipment (e.g., electricity/heat generation, energy-intensive industries, maritime transport, and aviation under the EEA), amounting to around 38% of the EU's total emissions. In certain sectors, only entities above a certain size are included and there are other exceptions for governments to exclude other entities if they put in place other measures for emission reductions. A cap is placed on the number of allowances brought into circulation, which differs between the aviation and stationary installations sectors. Both caps decrease with a linear reduction factor of approximately 4.3% from 2024 onwards (until 2030). In phase IV (2021-2030) auctioning is the primary method for the distribution, accounting for 57% of the cap, which will increase over time.

The EEX is the common auctioning platform. On this and the ICE exchange, spot, futures, and futures options can be traded. Depending on the sector, entities receive a number of allowances by free allocation. This happens based on performance benchmarks, reflecting the 10% most efficient installations of each sector. Banking is allowed unlimitedly whereas borrowing is not. However, implicit borrowing is possible, meaning that allowances allocated in the current trading period can be used for compliance in the previous trading period. From 2021 the use of carbon credits is not allowed anymore. The Market Stability Reserve (MSR) is designed to keep additional allowances in reserve or release more of them to prevent too low or too high prices. When the total number of allowances in circulation (TNAC) is higher than 833 million, 24% (12% from 2024 onwards) of this surplus is taken into the reserve. When the TNAC falls below 400 million, 100 million allowances are taken from the reserve to the market. In 2023 the gradual introduction of the Carbon Border Adjustment Mechanism (CBAM) started. Over time, this will increasingly tax the carbon emissions of goods imported to the EEA.

UK Emission Trading Scheme (UKETS)

After leaving the EU ETS, the UK's ETS came into effect in 2021, covering a large number of entities in the industrial, power, and aviation (domestic and to Switzerland and the EEA) sectors. Auctioning is the primary allowance allocation means. This, and the subsequent trading of UKA futures happens on the ICE Futures Europe exchange. Banking is allowed and borrowing only implicitly, like in the EU ETS. The use of carbon credits is not permitted at this time, but there are plans to change this. The cost containment mechanism (CCM) should avoid spikes in allowance prices by allowing regulators to auction additional allowances. The CCM is triggered when the monthly average carbon price is more than 3 times the 2-year average carbon price for six consecutive months. When the CCM is triggered, the authorities have a large degree of freedom in their response. Among others, this may include bringing forward allowances from future years and drawing allowances from reserves (market stability reserve or new entrants' reserve).

New Zealand Unit (NZU)

the New Zealand ETS has currently the largest emission coverage. Only for emissions-intensive and trade-exposed activities free allocation is granted, the remaining allowances are auctioned on the EEX and New Zealand Exchange (NZX). Banking is allowed but borrowing is not. Exclusively with governmental approval international or domestic carbon credits can be used within the ETS. Within the auctioning process, a price floor is introduced which is currently set at NZD 20. In addition, a cost containment reserve is triggered at a predetermined price (currently NZD 81) which will allow additional allowances to be released.

California Carbon Allowance (CCA)

California's cap-and-trade ETS covers around 75% of the state's GHG emissions. It operates together with Quebec's ETS in the Western Climate Initiative (WCI). Allowances are made available in sealed-bid auctions. On the secondary market, allowances and financial derivatives based thereupon are traded on the ICE, CME Group and Nodal Exchange platforms. The auction price floor was set at USD 22.21 in 2023 and will increase by 5% annually in surplus to the inflation. Borrowing is not possible, but banking is allowed under a holding limit.

Carbon credits from selected offset protocols may be used up to 4% of a company's obligation until 2025, but this limit will rise to 6% for the period 2026-2030. Each year, some allowances are placed in the allowance price containment reserve (APCR). This mechanism consists of two reserve tiers and a price ceiling. Allowances from the APCR are auctioned if the price from the preceding quarter is greater than or equal to 60% of the lowest price tier. The reserve tiers and price ceiling in California's ETS were USD 51.92, USD 66.71, and USD 81.50 respectively in 2023 and will increase annually by 5% plus inflation.

Washington Carbon Allowance (WCA)

Started in January 2023, the Washington cap-and-invest program covers around 70% of the state's emissions. The elements of the program are similar to those of California's program. Auctioning and free allocation (mainly based on benchmarking) are the two methods of distribution. The auctioning is subjected to a price floor of USD 22.20 in 2023, increasing by 5% plus inflation. Furthermore, the program contains an allowance price containment reserve (APCR) from which allowances can be auctioned at pre-set prices when allowance prices increase unexpectedly. The APCR has two price tiers, which are USD 51.90 and USD 66.68 in 2023, which increase by 5% plus inflation. When the

APCR has run out of units, companies with insufficient units can request to buy units at a ceiled price of USD 81.47 in 2023 (which also increases by 5% plus inflation per year).

Regional Greenhouse Gas Initiative (RGGI)

This is an ETS for multiple north-eastern US states, covering the power sector. Covered entities acquire most allowances through actions, although some states have free allocation mechanisms. Banking of allowances is allowed without restrictions. The current regulations reduce the cap in a compliance period to address the number of allowances banked in a previous period. Borrowing is not allowed. Until 2030, 3.3% of a company's obligation can be covered by certain types of carbon credits. A cost containment reserve (CCR) is put in place which releases quantity of allowances to the market in addition to the cap when a certain price trigger is reached. In 2021, the CCR was filled with 10% of the cap. The trigger price is USD 14.88 in 2023 and increases with 7% per year. In 2023 an auction price floor was set at USD 2.50 which will increase by 2.5% annually.

A.2.2 Voluntary Carbon Market

ICE traded Voluntary Emission Reduction (ICEVER)

VERs are similar to the carbon credits mentioned above, except that they are not established in a regulatory context. The VERs are generally certified through a voluntary certification process. The Verified Carbon Standard (VCS, by Verra), the Gold Standard and Climate Action Reserve (CAR) are the three largest certifiers with a combined market share of approximately 90% (Autoriteit Financiële Markten, 2023).

(Nature-Based) Global Emissions Offsets ((N-)GEO)

In an attempt to standardize the the VCM, the CME group launched various Global Emission Offsets (GEOs). There are 5 type of GEOs traded via exchanges and OTC. GEO contracts are based on credits received from three major registries: Verra, the American Carbon Registry, and the Climate Action Reserve (S&P, 2021). The offsets that are delivered under the GEO contracts are CORSIA eligible. This is the Carbon Offsetting and Reduction Scheme for International Aviation set by the UN agency International Civil Aviation Organization (ICAO). GEO contracts are not limited to companies in the aviation sector but adhere to the CORSIA standard.

N-GEOs are comprised of nature-based offset projects from the Verra registry. Nature based solutions offset projects avoid and reduce emissions through nature conservation and restoration. The projects producing N-GEO carbon credits fall under the Agriculture, Forestry, or Other Land Use (AFOLU) category, in contradiction to GEO projects.

APPENDIX B. SDE SOLUTIONS & LOG-LIKELIHOOD FUNCTIONS

B.1 ITÔ'S LEMMA

The differential of a time-dependent stochastic process can be found by applying Itô's lemma, which is given in Equation (8.2) for a function $f(t,x)$ as shown in Equation (8.1).

$$dx = a(x,t)dt + b(x,t)dz \tag{8.1}$$

$$df = \left(\frac{\partial f}{\partial t}a + \frac{\partial f}{\partial x} + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}b^2\right)dt + \frac{\partial f}{\partial x}b\,dz \tag{8.2}$$

B.2 GEOMETRIC BROWNIAN MOTION

The solution of the GBM SDE is (Roncoroni et al., 2015):

$$X(t) = X(0)\exp\left(\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W(t)\right) \tag{8.3}$$

The discrete log-likelihood function for the GBM is (Nkemnole & Abass, 2019):

$$\mathcal{L}(X,\mu,\sigma) = -\frac{n}{2}\ln(2\pi) - \sum_{i=1}^{n}\ln(X_i\sigma\sqrt{dt}) - \frac{1}{2}\sum_{i=1}^{n}\frac{\left(\ln\left(\frac{X_i}{X_{i-1}}\right) - \left(\mu - \frac{\sigma^2}{2}\right)\Delta t\right)^2}{2\sigma^2\Delta t} \tag{8.4}$$

B.3 ORNSTEIN-UHLENBECK MODEL

The solution of the OU SDE is (Roncoroni et al., 2015):

$$X(t) = \exp(-\kappa t)X(0) + \mu(1-\exp(-\kappa t)) + \sigma\int_0^t \exp\big(-\kappa(t-s)\big)\,dW(s) \tag{8.5}$$

The log-likelihood function for the OU model is (Valdivieso et al., 2009):

$$\begin{aligned}\mathcal{L}(\mu,\kappa,\bar{\sigma}) = &-\frac{n}{2}\ln(2\pi) - n\ln(\bar{\sigma}) \\ &-\frac{1}{2\bar{\sigma}^2}\sum_{i=1}^{n}[X_i - X_{i-1}\exp(-\kappa\Delta t) - \mu(1-\exp(-\kappa\Delta t))]^2\end{aligned} \tag{8.6}$$

With $\bar{\sigma}^2 = \sigma^2\frac{(1-\exp(-2\kappa\Delta t))}{2\kappa}$. Based on this log-likelihood function, the maximum likelihood estimators can be derived (Valdivieso et al., 2009):

$$\mu = \frac{\sum_{i=1}^{n}[X_i - X_{i-1}\exp(-\kappa\Delta t)]}{n(1-\exp(-\kappa\Delta t))} \tag{8.7}$$

$$\kappa = -\frac{1}{\Delta t}\ln\left[\frac{\sum_{i=1}^{n}(X_i-\mu)(X_{i-1}-\mu)}{\sum_{i=1}^{n}(X_{i-1}-\mu)^2}\right] \tag{8.8}$$

$$\bar{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}[(X_i - \mu - \exp(-\kappa\Delta t)\,(X_{i-1}-\mu)]^2 \tag{8.9}$$

B.4 EXPONENTIAL ORNSTEIN-UHLENBECK MODEL

The solution of the eOU model is equal to that of the OU model. The only difference is that $X(t)$ is replaced by $\ln\big(X(t)\big)$ (Mejía Vega, 2018):

$$\ln(X(t)) = \exp(-\kappa t)\ln(X(t)) + \left(\mu - \frac{\sigma^2}{2\kappa}\right)(1 - \exp(-\kappa t)) + \sigma \int_0^t \exp\big(-\kappa(t-s)\big)\, dW(s) \tag{8.10}$$

The log-likelihood function for $\ln(X)$ in the eOU model is (Mejía Vega, 2018):

$$\mathcal{L}(\mu, \kappa, \bar{\sigma}) = -\frac{n}{2}\ln(2\pi) - n\ln(\bar{\sigma}) - \frac{1}{2\bar{\sigma}^2}\sum_{i=1}^{n}[\ln(X_i) - \ln(X_{i-1})\exp(-\kappa\Delta t) - \bar{\mu}(1 - \exp(-\kappa\Delta t))]^2 \tag{8.11}$$

The log-likelihood function for $\ln(X)$ in Equation (8.13) can be converted to the log-likelihood function of X by subtracting $\sum_i^n \ln(X_i)$ (Mejía Vega, 2018). Based on Equation (8.13), the maximum likelihood estimators can be derived (Mejía Vega, 2018):

$$\bar{\mu} = \frac{\sum_{i=1}^{n}[\ln(X_i)] - \exp(-\kappa\Delta t)\sum_{i=1}^{n}[\ln(X_{i-1})]}{n(1 - \exp(-\kappa\Delta t))} \tag{8.12}$$

$$\kappa = -\frac{1}{\Delta t}\ln\left[\frac{\sum_{i=1}^{n}[\ln(X_i)\ln(X_{i-1})] - \bar{\mu}\sum_{i=1}^{n}[\ln(X_i)] - \bar{\mu}\sum_{i=1}^{n}[\ln(X_{i-1})] + n\bar{\mu}^2}{\sum_{i=1}^{n}[\ln(X_{i-1})^2] - 2\overline{\mu\sum_{i=1}^{n}[\ln(X_{i-1})]} + n\bar{\mu}^2}\right] \tag{8.13}$$

$$\bar{\sigma}^2 = \frac{1}{n}\Bigg[\sum_{i=1}^{n}[\ln(X_i)^2] - 2\exp(-\kappa\Delta t)\sum_{i=1}^{n}[\ln(X_i)\ln(X_{i-1})] + \exp(-2\kappa\Delta t)\sum_{i=1}^{n}[\ln(X_{i-1})^2] - 2\bar{\mu}(1 - \exp(-\kappa\Delta t))\left[\sum_{i=1}^{n}[\ln(X_i)] - \exp(-\kappa\Delta t)\sum_{i=1}^{n}[\ln(X_{i-1}]\right] + n\bar{\mu}^2(1 - \exp(-\kappa\Delta t))^2\Bigg] \tag{8.14}$$

With $\bar{\sigma}^2 = \sigma^2\frac{(1-\exp(-2\kappa\Delta t))}{2\kappa}$ and $\bar{\mu} = \mu - \frac{\sigma^2}{2\kappa}$.

B.5 MERTON JUMP DIFFUSION MODEL

The density function and the log-likelihood functions are given below (Staures et al., 2020):

$$f\left(\ln\left(\frac{X_t}{X_{t-\Delta t}}\right)\right) = \sum_{a}^{\infty}\frac{\exp(-\lambda\Delta t)(\lambda\Delta t)^a}{a!}\frac{1}{\sqrt{2\pi\big(\sigma^2\Delta t + \sigma_y^2 a\big)}}\exp\left(-\frac{\left(\ln\left(\frac{X_t}{X_{t-\Delta t}}\right) - \big(\gamma\Delta t + \mu_y a\big)\right)^2}{2(\sigma^2\Delta t + \sigma_y^2 a)}\right) \tag{8.15}$$

$$L(\mu, \lambda, \sigma, \sigma_y, \mu_y) = \prod_{i=1}^{n}\ln\left(f\left(\ln\left(\frac{X_i}{X_{i-1}}\right)\right)\right) \tag{8.16}$$

With a being the number of jumps, $\gamma = \mu - \frac{1}{2}\sigma^2 - \lambda k$ and $k = \mathbb{E}[Y_i - 1]$ (the expected percentage change of the price if a Poisson event occurs).

APPENDIX C. KALMAN FILTER

The Kalman filter is an algorithm which allows to recursively compute an optimal estimator for the states of a linear system in space form (Kalman, 1960). Conditional on the available information up to and including time t, the Kalman filter calculates a posterior (updated) prediction of a state vector's mean and covariance. The available information includes the specified stochastic processes that the state variables follow, a priori prediction for the state vector's mean and covariance, and measurement observations at time t.

The first step in the algorithm is to compute the priori estimator of the state vector's distribution at time t based on the transition equation. This estimator is computed as the expectation of Equation (3.45) (Kalman, 1960):

$$\widehat{\boldsymbol{X}}_{t|t-1} = \boldsymbol{c} + \boldsymbol{G}\widehat{\boldsymbol{X}}_{t-1|t-1} \tag{8.17}$$

Similarly, the priori estimator for the state vector's covariance matrix is computed as the expectation of Equation (3.46) (Kalman, 1960):

$$\boldsymbol{P}_{t|t-1} = \boldsymbol{G}\boldsymbol{C}_{t-1|t-1} + \boldsymbol{G}' + \boldsymbol{W} \tag{8.18}$$

Here, $\boldsymbol{C}_{t-1|t-1}$ is the covariance matrix for the posteriori distribution at the previous stime step. With the measurement equation, the posteriori estimator for the measurement vector ($\widehat{\boldsymbol{Y}}_{t|t-1}$) and the covariance matrix ($\boldsymbol{Q}_{t|t-1}$) can be determined with (Kalman, 1960):

$$\widehat{\boldsymbol{Y}}_{t|t-1} = \boldsymbol{d}_t + \boldsymbol{F}_t'\widehat{\boldsymbol{X}}_{t|t-1} \tag{8.19}$$

$$\boldsymbol{Q}_{t|t-1} = \mathbf{F}'\mathbf{P}_{\mathrm{t|t-1}}\boldsymbol{F}_t + \boldsymbol{V} \tag{8.20}$$

The next step is to compute the posteriori distribution for the state vector (Kalman, 1960):

$$\widehat{\boldsymbol{X}}_{t|t} = \widehat{\boldsymbol{X}}_{t|t-1} + \mathbf{P}_{\mathrm{t|t-1}}\boldsymbol{F}_t\left(\boldsymbol{Q}_{t|t-1}\right)^{-1}\left(\boldsymbol{Y}_t - \widehat{\boldsymbol{Y}}_{t|t-1}\right) \tag{8.21}$$

The covariance matrix of the above updated state vector's estimator is $\boldsymbol{C}_{t|t}$ (Kalman, 1960):

$$\boldsymbol{C}_{t|t} = \boldsymbol{P}_{t|t-1} + \mathbf{P}_{\mathrm{t|t-1}}\boldsymbol{F}_t\left(\boldsymbol{Q}_{t|t-1}\right)^{-1}\boldsymbol{F}_t'\boldsymbol{P}_{t|t-1} \tag{8.22}$$

The unknown parameters in Equations (3.41) and (3.42) and the diagonal elements in the measurement covariance matrix ($\boldsymbol{V}$) can be estimated, under the assumption that the measurement errors are normally distributed, by maximizing the log-likelihood function given below (Schwartz & Smith, 2000):

$$\ln(\mathcal{L}) = \sum_{t=1}^{n_T}\left(-\frac{n_{fut}}{2}\ln(2\pi) - \frac{1}{2}\ln\left(\det\left(\boldsymbol{Q}_{t|t-1}\right)\right) - \frac{1}{2}\boldsymbol{v}_t{}'\left(\boldsymbol{Q}_{t|t-1}\right)^{-1}\boldsymbol{v}_t\right) \tag{8.23}$$

Here, $\boldsymbol{v}_t$ are the errors between the predicted log futures prices, from Equation (3.43), and the observed futures prices.

APPENDIX D. ADDITIONAL MARKET DATA ANALYSIS

D.1 SPOT PRICE

Market	Market Component
NZETS	-
CMX	-
CMZ	-
CCA	-
EUETS	-
GEO	-
ICEVER	-
N-GEO	-
RGGI	-
UKETS	-
WCA	-

TABLE 11 *Names of selected market components per carbon market (redacted).*

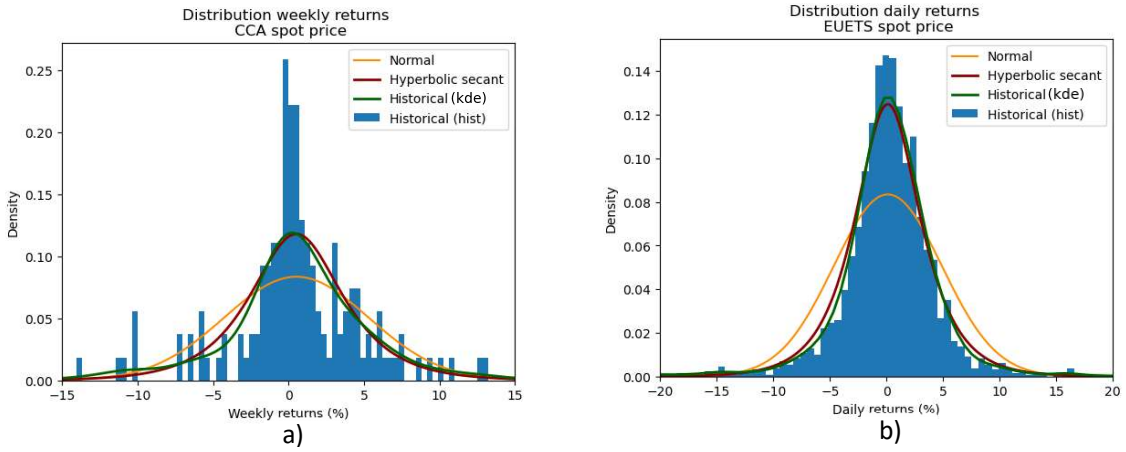

FIGURE 38 *Histograms and fitted distributions for daily/weekly returns for a) CCA and b) EUETS (as of halfway November 2023).*

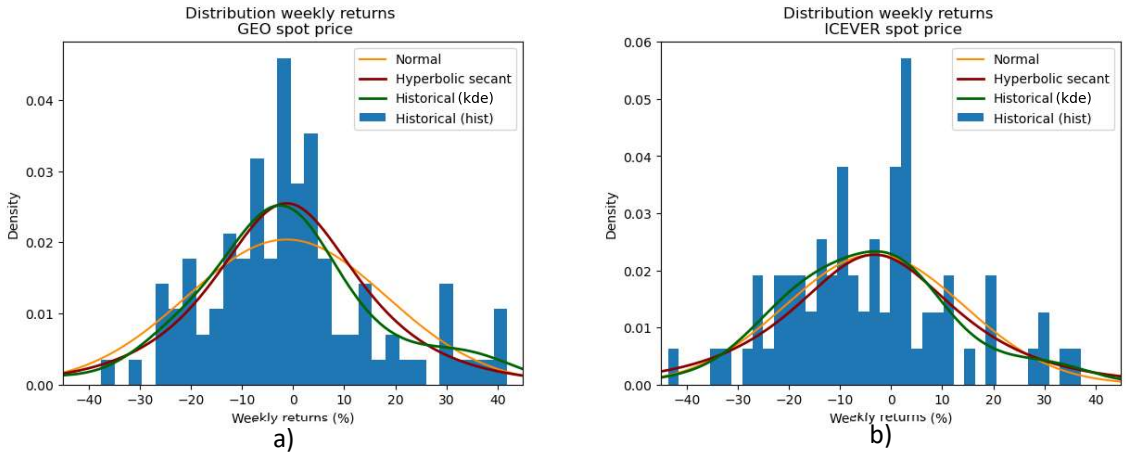

FIGURE 39 *Histograms and fitted distributions for daily/weekly returns for a) GEO and b) ICEVER (as of halfway November 2023).*

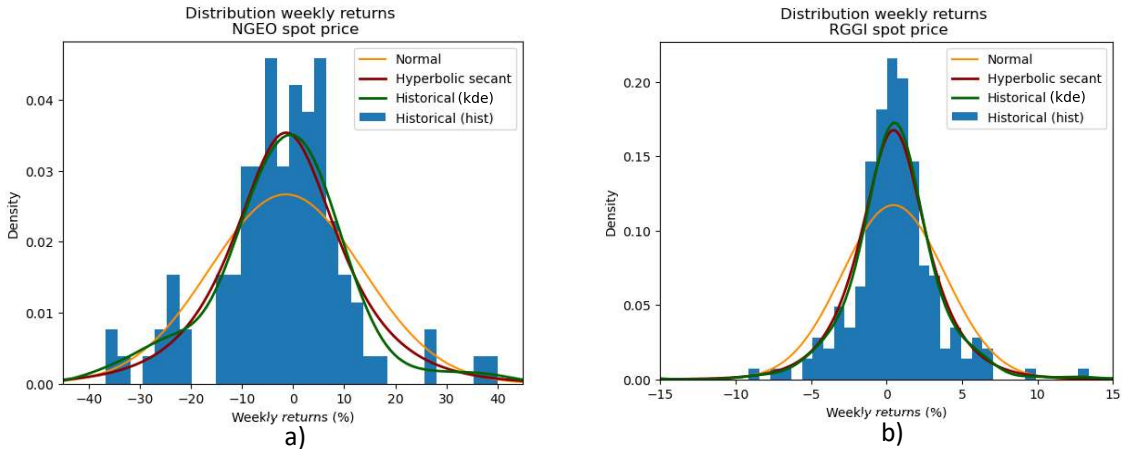

FIGURE 40 *Histograms and fitted distributions for daily/weekly returns for a) NGEI and b) RGGI (as of halfway November 2023).*

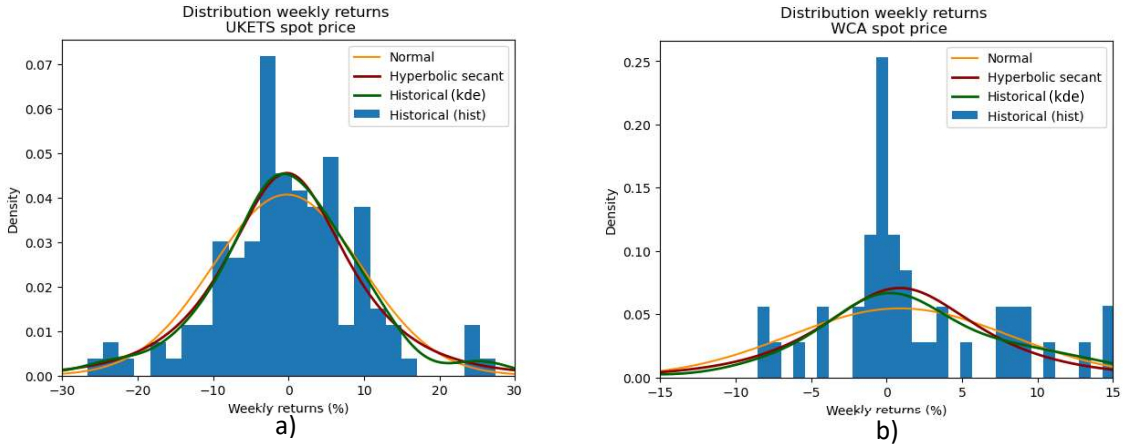

FIGURE 41 *Histograms and fitted distributions for daily/weekly returns for a) UKETS and b) WCA (as of halfway November 2023).*

D.2 FUTURES/FORWARD PRICES

TABLE 12 *Statistics of weekly and monthly CCA forward price returns. Returns are annualized.*

	Weekly			Monthly		
	Mont 5	Month 35	Month 65	Mont 5	Month 35	Month 65
Mean (%)	26.5	27.7	28.0	27.9	29.6	30.1
Variance (%)	28.26	25.75	25.22	23.3	31.1	20.1
Skewness	-0.65	-0.5	-0.6	0.1	0.15	0.2
Kurtosis	4.8	4.5	3.4	-0.5	-0.5	-0.6
Jarque-Bera p-value	1.8e-33	8.1e-30	4.7e-18	8.1e-1	8.2e-1	7.0e-1
Shapiro-Wilk p-value	5.7e-8	5.7e-8	9.9e-7	8.7e-1	7.6e-1	6.5e-1
Cramér-von Mises p-value (hyp. secant)	1.6e-1	7.5e-2	7.8e-2	9.7e-1	9.8e-1	8.3e-1

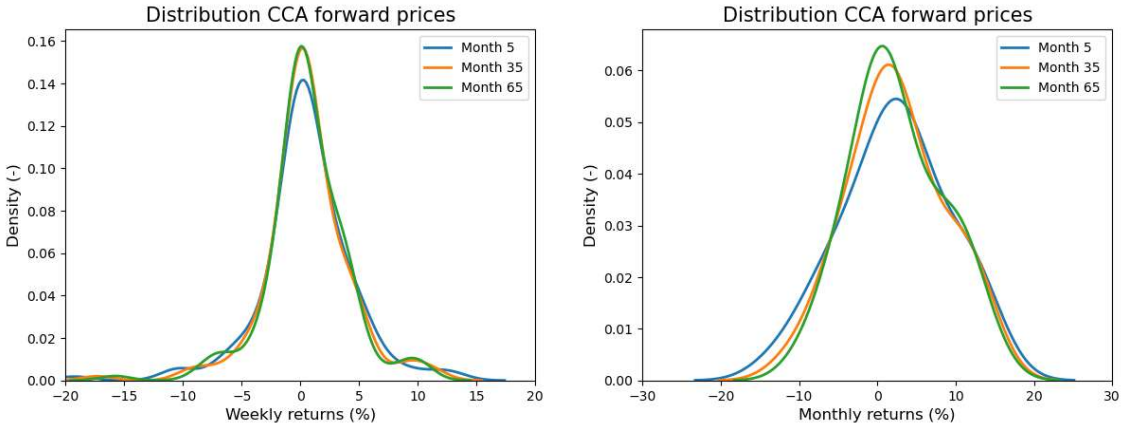

FIGURE 42 *Weekly and monthly CCA forward price returns distribution.*

TABLE 13 *Statistics of weekly and monthly EUETS forward price returns. Returns are annualized.*

	Weekly			Monthly		
	Mont 5	Month 60	Month 110	Mont 5	Month 60	Month 110
Mean (%)	11.2	11.2	11.4	11.7	11.8	11.9
Variance (%)	51.4	45.9	41.2	49.0	47.0	42.8
Skewness	-0.8	-0.8	-0.7	-0.64	-0.7	-0.6
Kurtosis	5.4	5.5	4.8	1.3	1.6	1.6
Jarque-Bera p-value	0	0	0	1.3e-5	1.0e-7	2.0e-6
Shapiro-Wilk p-value	0	0	0	4.0e-4	1.4e-4	7.7e-4
Hyperbolic secant CV p-value	2.9e-1	8.2e-2	7.0e-3	4.2e-1	3.6e-1	3.6e-1

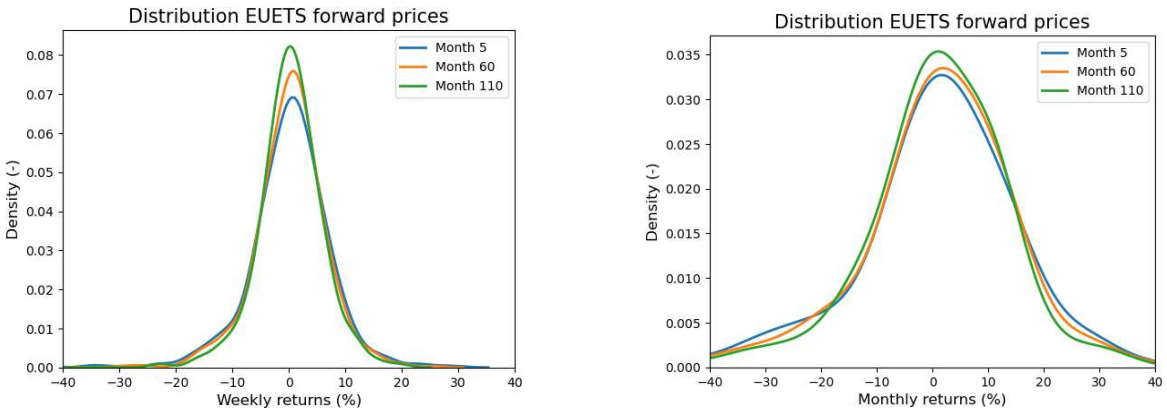

FIGURE 43 *Weekly and monthly EUETS forward price returns distribution.*

TABLE 14 *Statistics of weekly and monthly GEO forward price returns. Returns are annualized.*

	Weekly			Monthly		
	Mont 5	Month 25	Month 40	Mont 5	Month 25	Month 40
Mean (%)	-42.9	-44.3	-50.2	-50.9	-49.5	-57.0
Variance (%)	94.1	78.4	76.4	115	99.5	85.2
Skewness	0.1	0.6	0.8	0.31	-0.2	-0.1
Kurtosis	2.8	1.9	2.9	1.4	1.0	0.1
Jarque-Bera p-value	1.0e-9	1.3e-5	1.8e-5	2.5e-1	5e-1	9.7e-1
Shapiro-Wilk p-value	1.0e-4	6.2e-5	1.8e-5	8.4e-2	4.2e-1	8.4e-1
Hyperbolic secant CV p-value	8.3e-1	1.3e-1	1.3e-1	8.2e-1	9.2e-1	8.8e-1

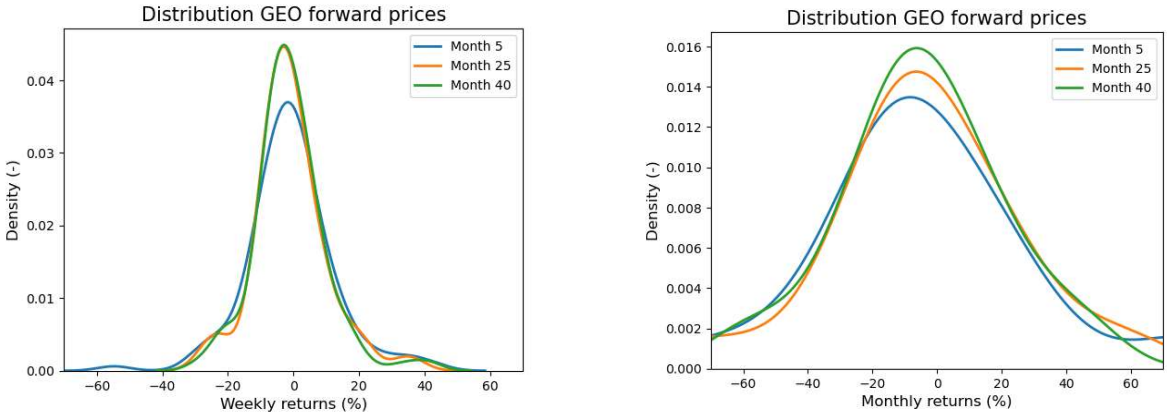

FIGURE 44 *Weekly and monthly GEO forward price returns distribution.*

TABLE 15 *Statistics of weekly and monthly ICEVER forward price returns. Returns are annualized.*

	Weekly			Monthly		
	Mont 5	Month 25	Month 45	Mont 5	Month 25	Month 45
Mean (%)	-170	-167	-167	-163	-160	-160
Variance (%)	117	117	116	117	118	119
Skewness	0.8	0.9	0.9	0.2	0.2	0.2
Kurtosis	1.8	2.1	2.1	0	0	0
Jarque-Bera p-value	3.8e-5	5.1e-6	2.4e-6	9.3e-1	9.5e-1	9.6e-1
Shapiro-Wilk p-value	5.2e-3	2.7e-3	2.2e-3	7.6e-1	7.6e-1	7.6e-1
Hyperbolic secant CV p-value	8.1e-1	7.9e-1	7.9e-1	8.4e-1	9.1e-1	9.3e-1

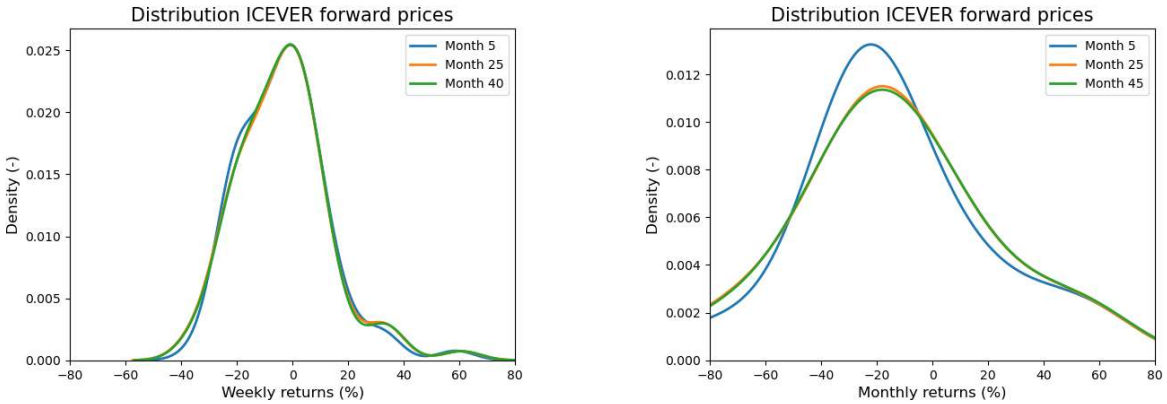

FIGURE 45 *Weekly and monthly ICEVER forward price returns distribution.*

TABLE 16 *Statistics of weekly and monthly RGGI forward price returns. Returns are annualized.*

	Weekly			Monthly		
	Mont 5	Month 35	Month 65	Mont 5	Month 35	Month 65
Mean (%)	24.0	26.5	29.3	28.9	31.5	34.5
Variance (%)	19.5	19.0	20.0	16.4	15.9	17.0
Skewness	0.1	0.2	0.2	1.2	1.2	1.2
Kurtosis	4.9	4.5	4.4	3.3	3.0	2.0
Jarque-Bera p-value	0	0	0	4.0e-7	2.3e-6	3.8e-5
Shapiro-Wilk p-value	0	0	0	5.0e-3	6.7e-3	2.6e-3
Hyperbolic secant CV p-value	2.5e-1	2.0e-1	2.3e-1	8.0e-1	7.4e-1	5.8e-1

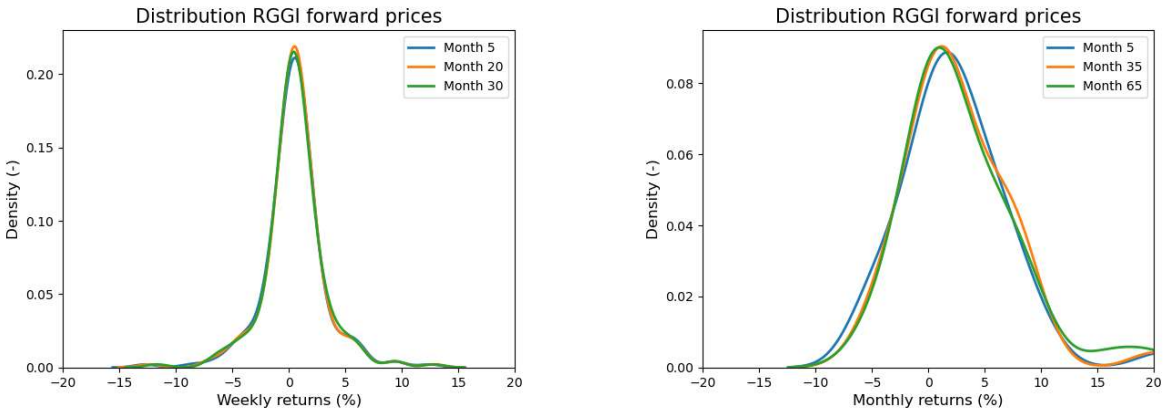

FIGURE 46 *Weekly and monthly RGGI forward price returns distribution.*

TABLE 17 *Statistics of weekly and monthly UKETS forward price returns. Returns are annualized.*

	Weekly			Monthly		
	Mont 5	Month 20	Month 30	Mont 5	Month 20	Month 30
Mean (%)	-7.6	-5.2	-3.8	-8.0	-5.4	-3.3
Variance (%)	50.0	47.5	45.9	54.2	52.0	50.3
Skewness	0.1	0.1	0.1	0.1	0.1	0.1
Kurtosis	0.5	0.6	0.6	-0.1	-0.2	-0.2
Jarque-Bera p-value	4.4e-1	3.7-1	2.9e-1	9.8e-1	9.5e-1	9.3e-1
Shapiro-Wilk p-value	8.2e-1	6.7e-1	5.5e-1	5.9e-1	5.3e-1	5.8e-1
Hyperbolic secant CV p-value	9.1e-1	9.9e-1	9.8e-1	7.8e-1	6.8e-1	7.5e-1

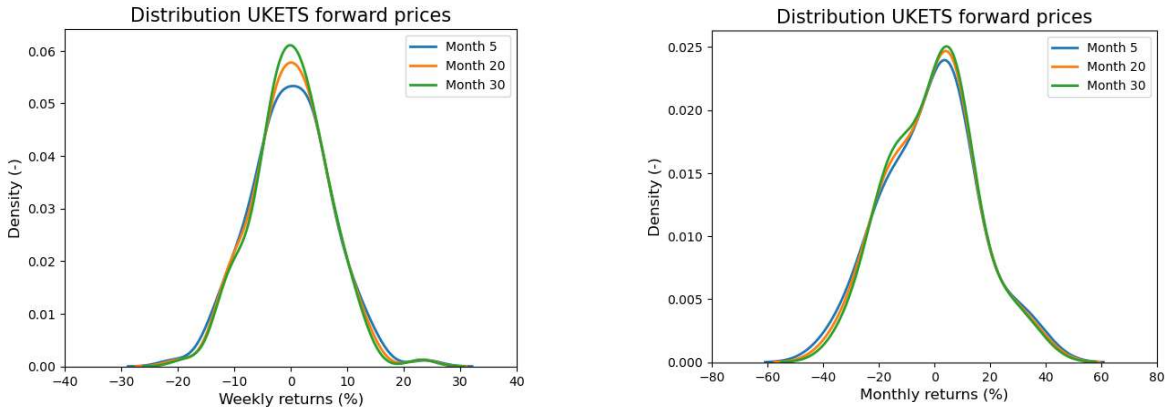

FIGURE 47 *Weekly and monthly UKETS forward price returns distribution.*

TABLE 18 *Statistics of weekly and monthly WCA forward price returns. Returns are annualized.*

	Weekly			Monthly		
	Mont 5	Month 20	Month 35	Mont 5	Month 20	Month 35
Mean (%)	-8.3	-9.5	-10.6	-12.4	-14.7	-18.3
Variance (%)	37.8	36.7	35.6	33.3	30.1	28.0
Skewness	0.1	0.2	0.4	-1.7	-1.6	-1.7
Kurtosis	5.2	5.0	5.0	1.4	1.4	1.5
Jarque-Bera p-value	0	0	0	1.5e-1	1.5e-1	1.5e-1
Shapiro-Wilk p-value	1.9e-5	2.7e-5	5.1e-5	4.3e-3	7.1e-3	5.2e-3
Hyperbolic secant CV p-value	2.2e-1	1.8e-1	2.0e-1	2.3e-1	3.4e-1	3.8e-1

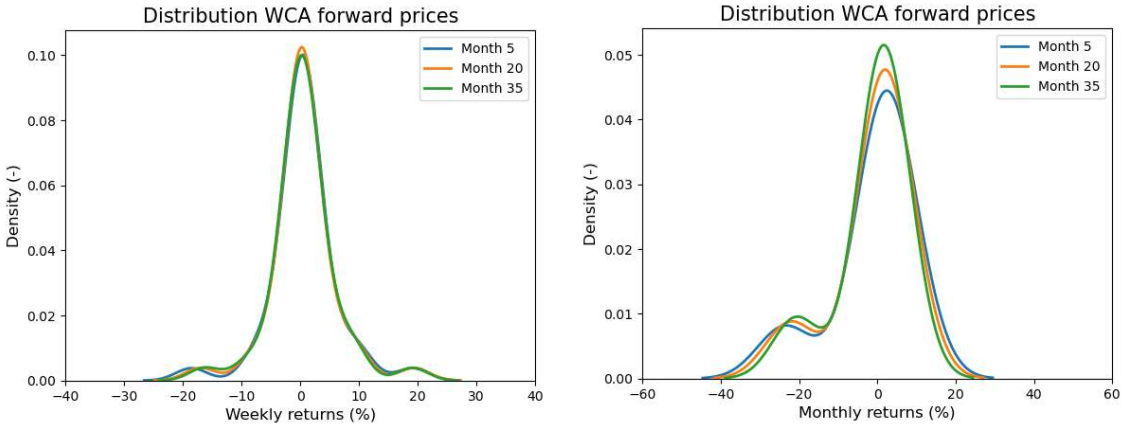

FIGURE 48 ***Weekly and monthly WCA forward price returns distribution**. Right image only based on 8 observations (8 months)*

APPENDIX E. FORWARD CURVE MODEL RESULTS

E.1 SCHWARTZ-SMITH MODEL

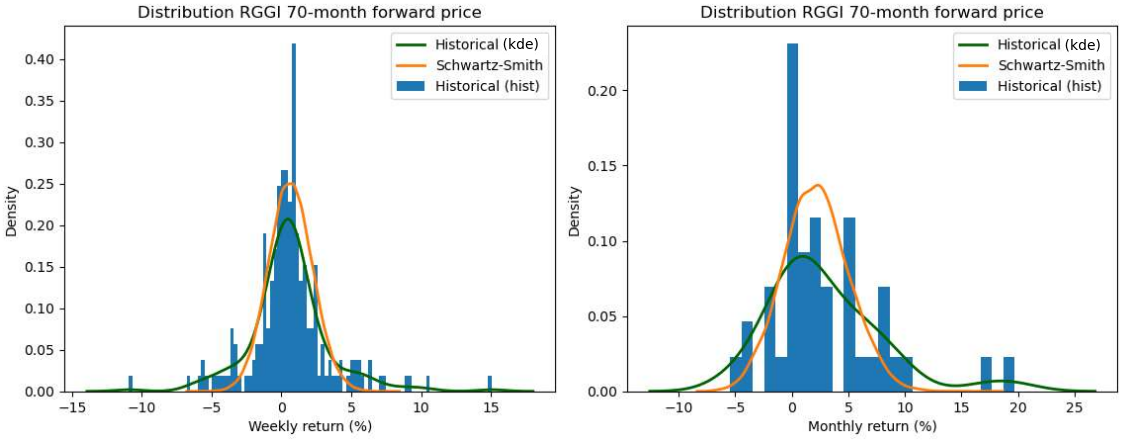

FIGURE 49 *Weekly and monthly RGGI 70-month returns predicted by the Schwartz-smith model compared to historical returns.*

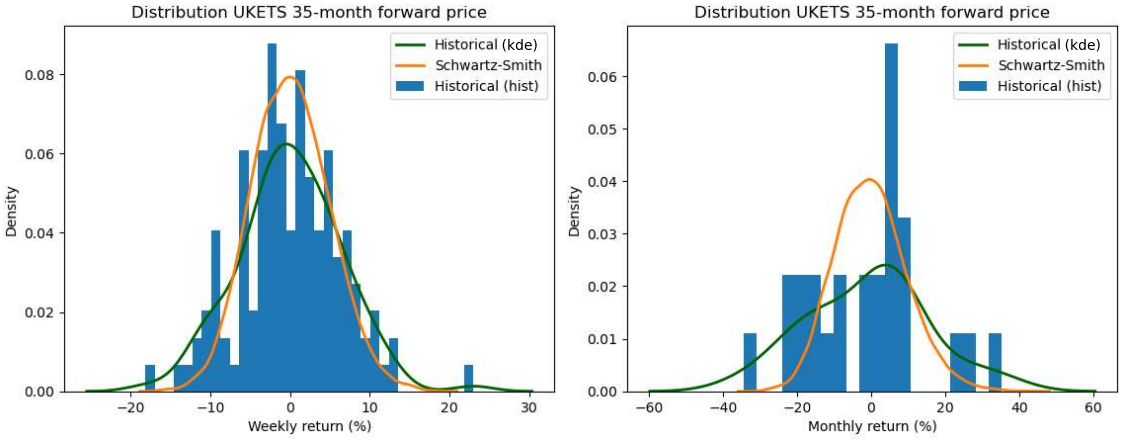

FIGURE 50 *Weekly and monthly UKETS 35-month returns predicted by the Schwartz-smith model compared to historical returns.*

E.2 PRINCIPAL COMPONENT ANALYSIS SIMULATION

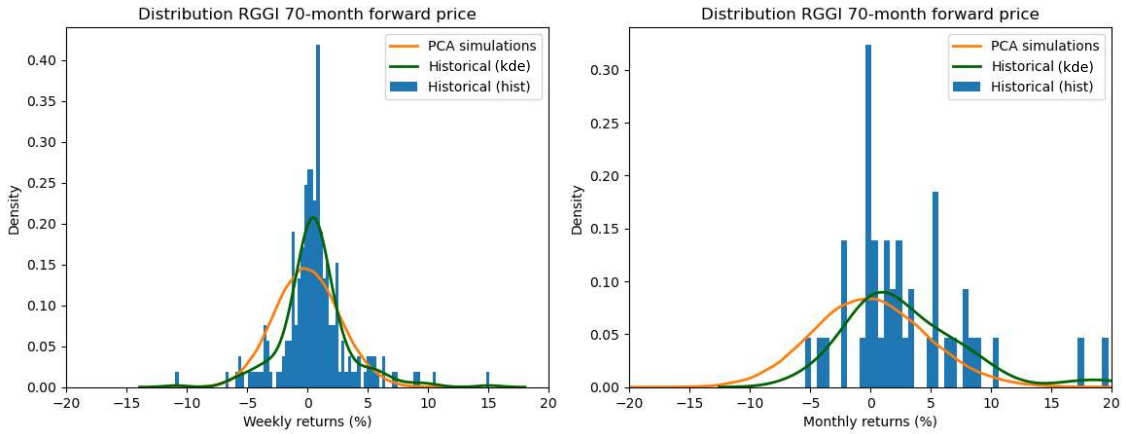

FIGURE 51 *Weekly and monthly RGGI 70-month returns predicted by PCA simulations compared to historical returns.*

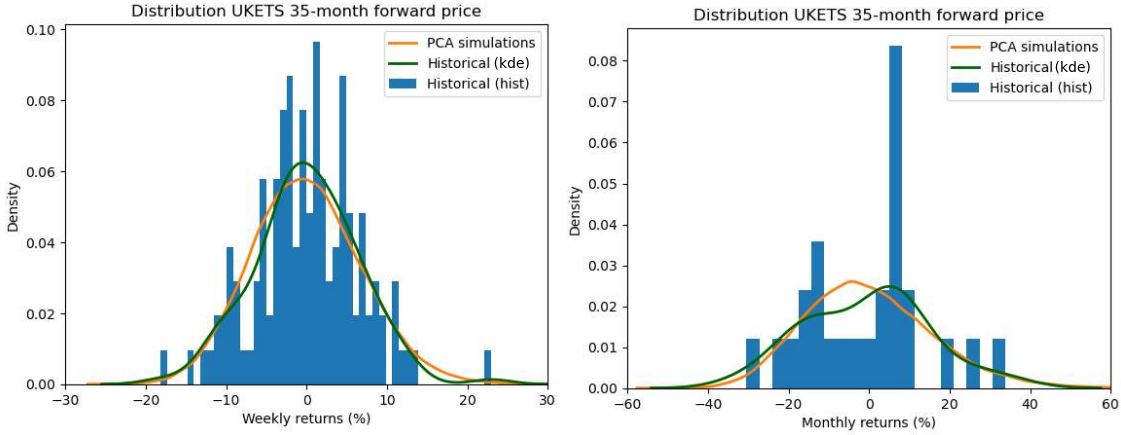

FIGURE 52 *Weekly and monthly UKETS 35-month returns predicted by PCA simulations compared to historical returns.*

APPENDIX F. SPOT PRICE MODEL RESULTS

TABLE 19 *Hurst exponents for historical carbon market spot prices*

Market	Hurst exponent		
	Daily	Weekly	Monthly
NZETS	X	0.85	0.91
CMX	X	X	0.84
CMX	X	x	0.65
CCA	X	0.70	X
EUETS	0.94	0.98	0.99
GEO	X	0.88	X
ICEVER	0.88	0.93	X
N-GEO	X	0.99	X
RGGI	X	0.91	0.93
UKETS	0.60	0.71	X
WCA	x	0.68	X

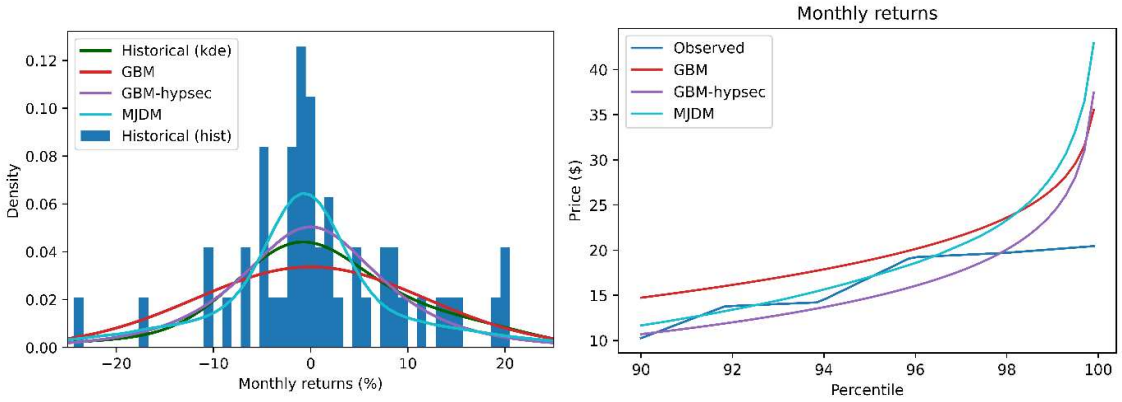

FIGURE 53 *Fitted spot price model distributions and right tail percentiles for CMX.*

TABLE 20 *CMX spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Monthly	-113	X	-108
p-value CvM	Monthly	3.7e-2	2.4e-1	4.5e-1

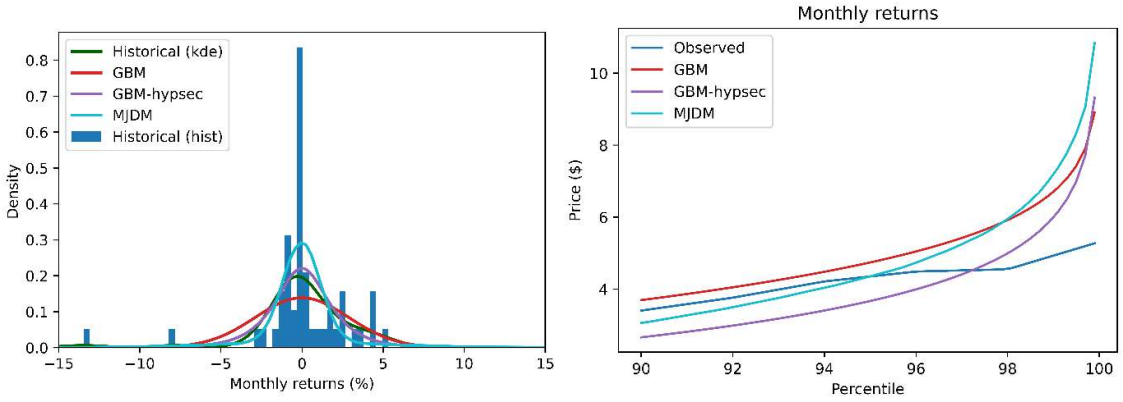

FIGURE 54 *Fitted spot price model distributions and right tail percentiles for CMZ.*

TABLE 21 *CMZ spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood (prices)	Monthly	-19	X	-15
p-value CvM	Monthly	1.9e-2	1.2e-1	1.9e-1

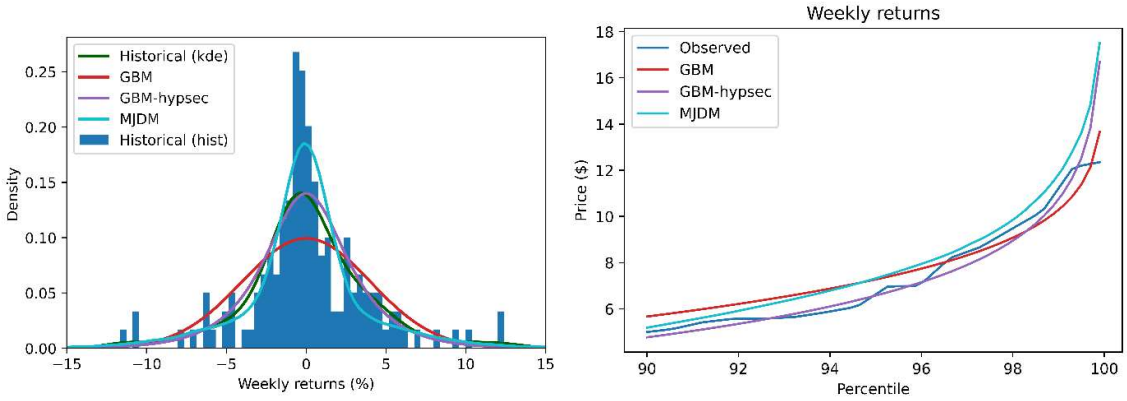

FIGURE 55 *Fitted spot price model distributions and right tail percentiles for CCA.*

TABLE 22 *CCA spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Weekly	-191	X	-205
p-value CvM	Weekly	1.8e-2	1.9e-1	7.3e-1

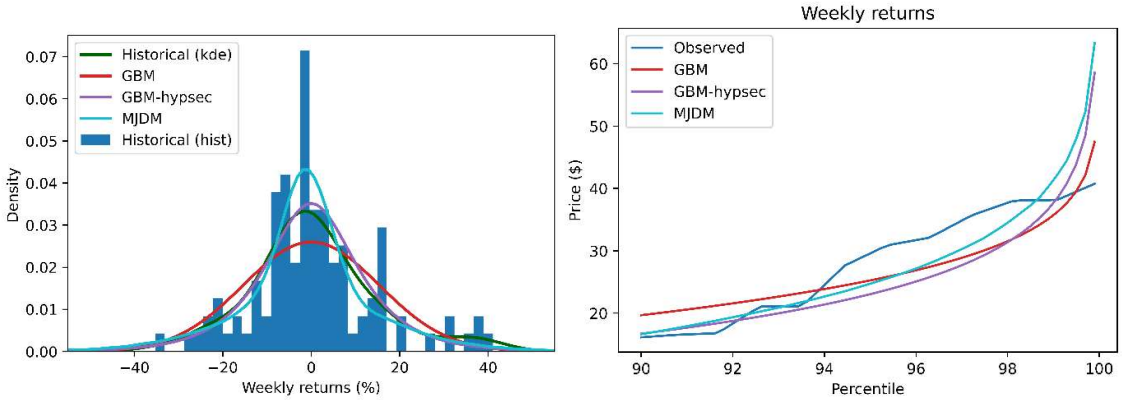

FIGURE 56 *Fitted spot price model distributions and right tail percentiles for GEO.*

TABLE 23 *GEO spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Weekly	-60	X	-78
p-value CvM	Weekly	9.9e-2	4.1e-1	5.6e-1

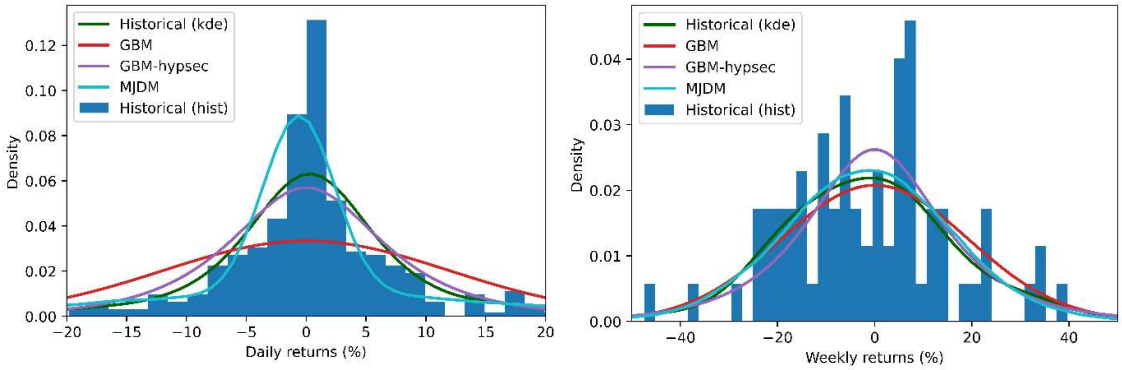

FIGURE 57 *Fitted spot price model distributions for ICEVER.*

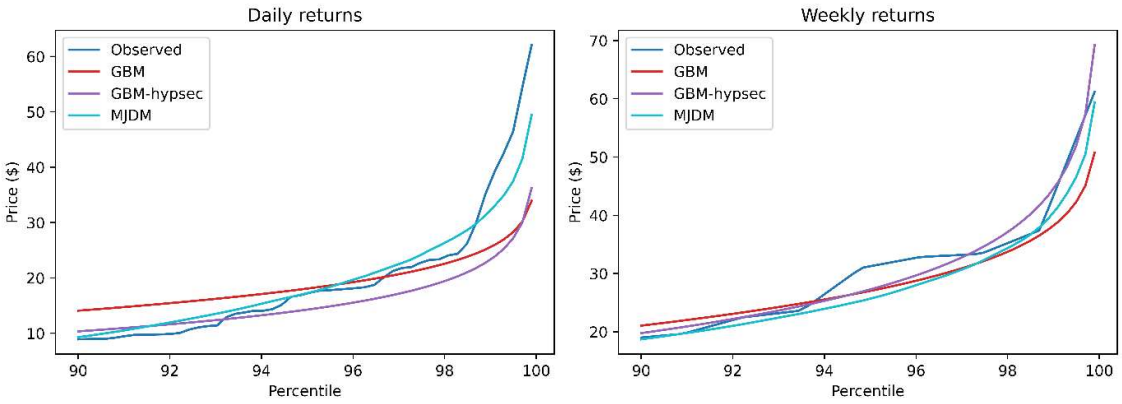

FIGURE 58 *Right tail percentiles of fitted model distributions for ICEVER.*

TABLE 24 *ICEVER spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Daily	-61	X	-39
	Weekly	-35	X	-55
p-value CvM	Daily	0	6.4e-3	1.1e-2
	Weekly	7.0e-1	5.7e-1	6.7e-1

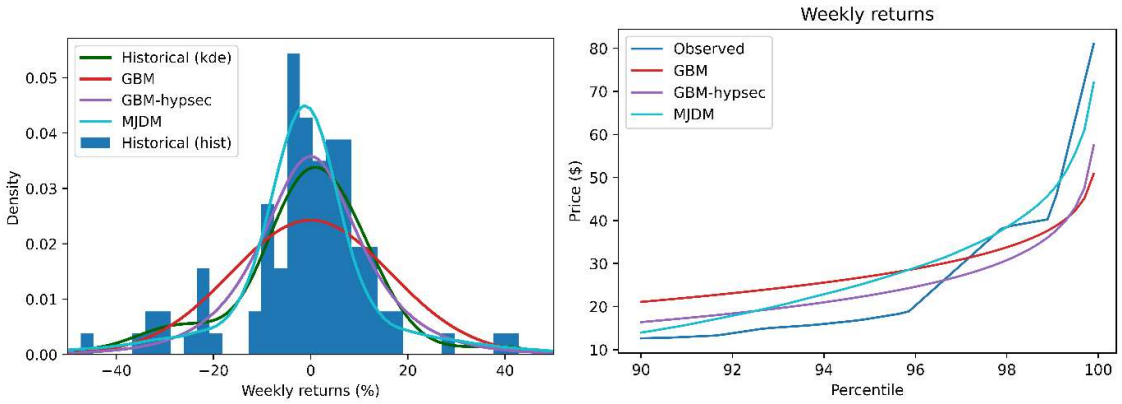

FIGURE 59 *Fitted spot price model distributions and right tail percentiles for NGEO.*

TABLE 25 *NGEO spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Weekly	-89	X	-93
p-value CvM	Weekly	2.0e-2	4.2e-1	1.3e-1

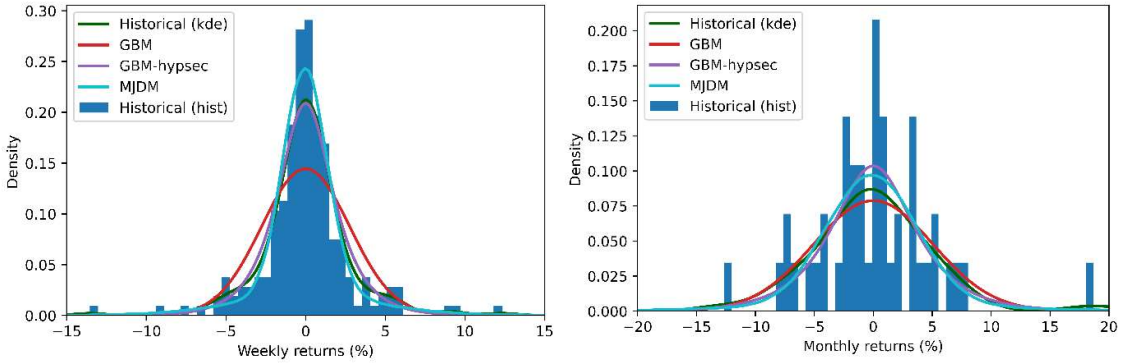

FIGURE 60 *Fitted spot price model distributions for RGGI.*

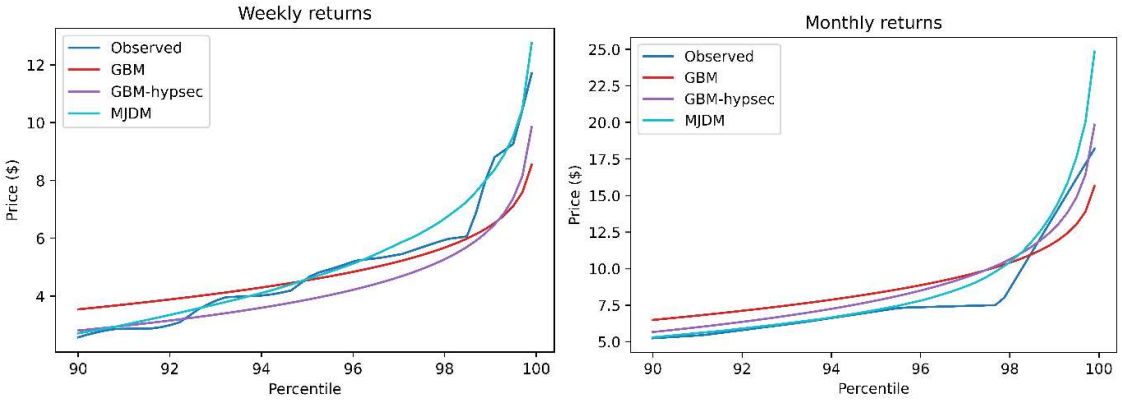

FIGURE 61 *Right tail percentiles of fitted model distributions for RGGI.*

TABLE 26 *RGGI spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Weekly	-114	X	-101
	Monthly	-23	X	-31
p-value CvM	Weekly	1.3e-3	2.1e-1	7.5e-1
	Monthly	6.3e-1	9.7e-1	9.6e-1

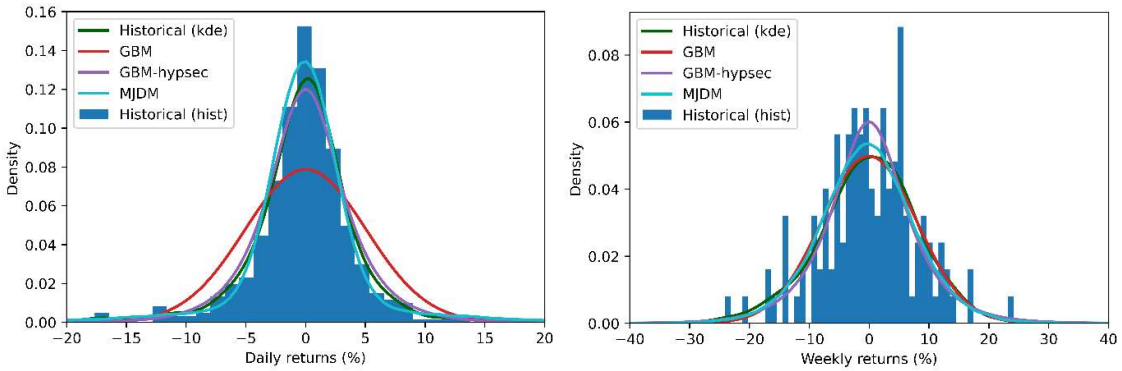

FIGURE 62 *Fitted spot price model distributions for UKETS.*

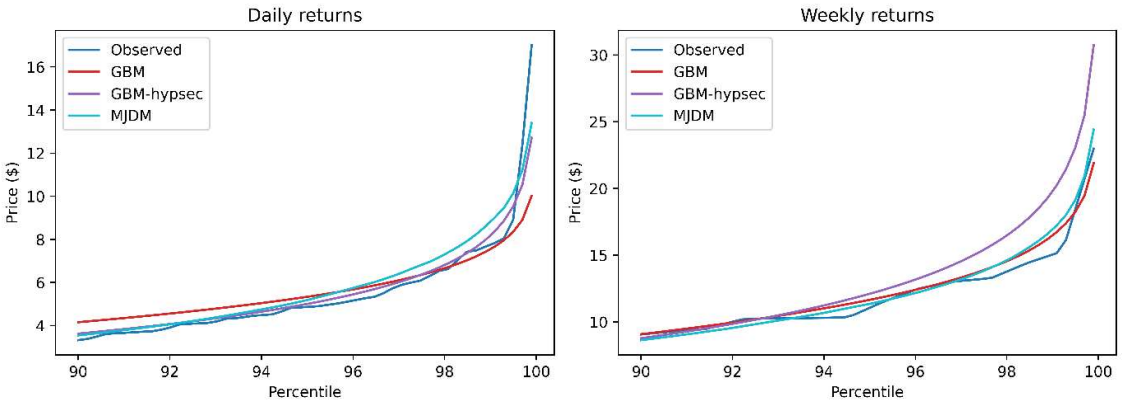

FIGURE 63 *Right tail percentiles of fitted model distributions for UKETS.*

TABLE 27 *UKETS spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Daily	-934	X	-904
	Weekly	-342	X	-374
p-value CvM	Daily	5.8e-3	5.2e-2	3.8e-1
	Weekly	9.9e-1	9.0e-1	9.7e-1

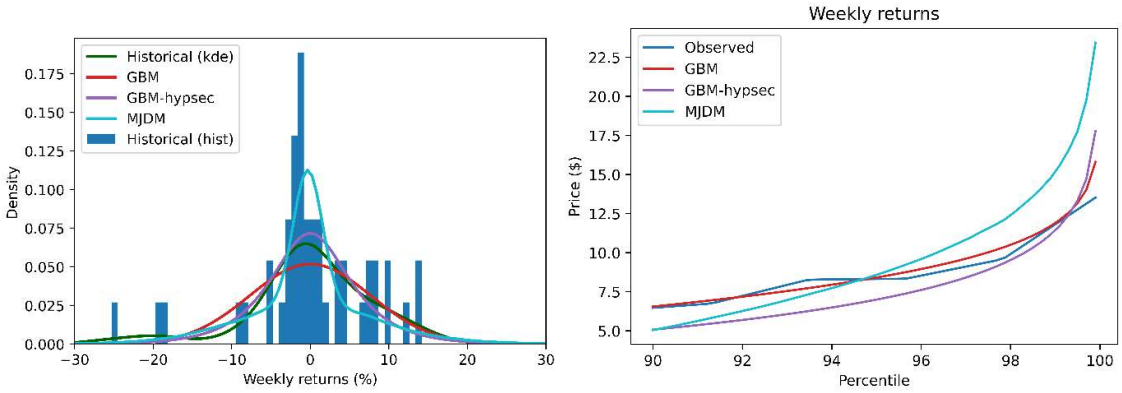

FIGURE 64 *Fitted spot price model distributions and right tail percentiles for WCA.*

TABLE 28 *WCA spot price model fitting results.*

		GBM	GBM-hypsec	MJDM
Log-likelihood	Weekly	-102	X	-102
p-value CvM	Weekly	5.0e-2	1.8e-1	3.6e-1

APPENDIX G. OPTION PRICING MODELS

G.1 BLACK MODEL

If an asset follows the GBM, the option prices can be computed with the following equations (Black, 1976):

$$c_B = \exp(-rT)[FN(d_1) - KN(d_2)] \tag{8.24}$$

$$p_B = \exp(-rT)[KN(-d_2) - FN(-d_1)] \tag{8.25}$$

$$d_1 = \frac{\ln\left(\frac{F}{K}\right) + \left(\frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}} \tag{8.26}$$

$$d_2 = d_1 - \sigma\sqrt{T} \tag{8.27}$$

Here, c_B is the price of a call option, p_B is the price of a put option under the Black model, r is the risk-free rate, T is the maturity, F is the current forward price, K is the strike price, σ the volatility, and $N()$ is the cumulative normal distribution function.

G.2 MERTON JUMP DIFFUSION MODEL OPTION PRICING

If an asset follows the MJDM, the option prices can be computed with the following equations (Cont & Tankov, 2002):

$$c_{MJDM} = \sum_{i=0}^{\infty}\left[\frac{\exp\left(-\bar{\lambda}T\right)\left(\bar{\lambda}T\right)^i}{i!}c_B(\sigma_i, r_i)\right] \tag{8.28}$$

$$p_{MJDM} = \sum_{i=0}^{\infty}\left[\frac{\exp\left(-\bar{\lambda}T\right)\left(\bar{\lambda}T\right)^i}{i!}p_B(\sigma_i, r_i)\right] \tag{8.29}$$

$$\bar{\lambda} = \lambda\exp\left(\mu_y + \frac{1}{2}\sigma_y^2\right) \tag{8.30}$$

$$\sigma_i = \sqrt{\sigma^2 + \frac{i\sigma_y^2}{T}} \tag{8.31}$$

$$r_i = r - \lambda\left(\exp\left(\mu_y + \frac{\sigma_y^2}{2} - 1\right) + \frac{i\mu_y + \frac{i\sigma_y^2}{2}}{T}\right. \tag{8.32}$$

Here, c_{MJDM} is the price of a call option, p_{MJDM} is the price of a put option for the MJDM.