Counterparty credit risk in bilateral hedging within energy markets

Joonatan Linkola

School of Science

Thesis submitted for examination for the degree of Master of Science in Technology.
Espoo 7.7.2025

Supervisor

Prof. Ahti Salo

Advisor

D.Sc. (Tech.) Edoardo Tosoni

Aalto University
School of Science

Aalto University
School of Science

Aalto University, P.O. BOX 11000, 00076 AALTO
www.aalto.fi
Abstract of the master's thesis

Author Joonatan Linkola

Title Counterparty credit risk in bilateral hedging within energy markets

Degree programme Master's Programme in Mathematics and Operations Research

Major Systems and Operations Research **Code of major** SCI3055

Supervisor Prof. Ahti Salo

Advisor D.Sc. (Tech.) Edoardo Tosoni

Collaborative partner Fortum Oyj

Date 7.7.2025 **Number of pages** 77 **Language** English

Abstract

As the global energy consumption grows at an accelerating rate, decarbonization measures and sustainable and reliable energy become increasingly important. Bilateral Power Purchase Agreements (PPAs) have gained noticeable traction as hedging instruments in recent years, as firms want to secure electricity at a pre-agreed price for up to 25 years in advance. However, while bilateral PPAs mitigate market risk that stems from the volatility and uncertainty of future electricity prices, they introduce counterparty credit risk, which refers to the risk that either party may fail to meet their contractual obligations.

Quantifying counterparty credit risk from the perspective of an electricity producer and retailer is essential to ensure that PPA terms are reasonable and that the PPAs are ultimately completed without premature terminations. Two common methods to mitigate this risk are setting a collateral requirement for the counterparty and including a credit value adjustment (CVA) add-on to the contract price. Determining the right balance of these requirements is challenging but also critical: too low requirements expose the electricity producer and retailer to excessive risk, whereas overly high requirements pose unnecessary financial strain on the counterparty and may potentially encourage it to turn to competitors.

This thesis presents a decision tree model that probabilistically estimates the earnings coming from a single bilateral PPA. A case study compares the expected earnings from accepting the contract to those from declining it, as well as to hypothetical but realistic alternative hedging strategies. The collateral requirement and the CVA add-on are then used to adjust the earnings to an acceptable risk level. In the end, a sensitivity analysis of the earnings confidence level is conducted, and both the model and the case study results are critically evaluated.

Keywords Counterparty credit risk, bilateral hedging, energy markets, decision trees

Aalto University
School of Science

Aalto-yliopisto, PL 11000, 00076 AALTO
www.aalto.fi
Diplomityön tiivistelmä

Tekijä Joonatan Linkola	
Työn nimi Counterparty credit risk in bilateral hedging within energy markets	
Koulutusohjelma Master's Programme in Mathematics and Operations Research	
Pääaine Systems and Operations Research	**Pääaineen koodi** SCI3055
Työn valvoja Prof. Ahti Salo	
Työn ohjaaja TkT Edoardo Tosoni	
Yhteistyötaho Fortum Oyj	
Päivämäärä 7.7.2025 **Sivumäärä** 77	**Kieli** Englanti

Tiivistelmä

Kun globaali energiankulutus kasvaa entistä nopeammin, hiilidioksidipäästöjen vähentäminen sekä kestävä ja luotettava energia nousevat entistä tärkeämmiksi puheenaiheiksi. Bilateraaliset eli kahdenväliset sähkönostosopimukset (PPA-sopimukset, engl. Power Purchase Agreement) ovat viime vuosina yleistyneet sähkön hinnanvaihtelun suojauskeinoina, kun yritykset haluavat hankkia sähköä ennalta sovittuun hintaan jopa 25 vuodeksi eteenpäin. Vaikka sähkön hinnanvaihteluihin liittyvät riskit vähenevät PPA-sopimuksien myötä, näihin sopimuksiin kuitenkin sisältyy vastapuoliluottoriskiä. Tällä tarkoitetaan riskiä siitä, että jompikumpi osapuoli ei täytä sopimusvelvoitteitaan.

Sähköntuottajan ja -myyjän näkökulmasta vastapuoliluottoriskin kvantifiointi on keskeisessä roolissa, jotta PPA-sopimusten ehdot olisivat kohtuullisia ja sopimukset onnistuisivat. Kaksi yleistä tapaa pienentää tätä riskiä ovat vastapuolelta vaadittava vakuus sekä sopimushinnan CVA-lisä (engl. credit value adjustment). Oikean tasapainon löytäminen on haastavaa mutta myös olennaista: liian alhaiset vaatimukset altistavat sähköntuottajan ja -myyjän turhan suurelle riskille, kun taas liialliset vaatimukset voivat rasittaa vastapuolta taloudellisesti ja jopa ajaa sen kilpailjoille.

Tämä diplomityö esittelee päätöspuumallin, joka arvioi yksittäisen PPA-sopimuksen tuottoja. Tapaustutkimuksessa hyväksytyn sopimuksen ennustettuja tuottoja vertaillaan tilanteeseen, jossa se hylättäisiin, sekä hypoteettisiin mutta realistisiin vaihtoehtoisiin suojaussopimuksiin. Tämän lisäksi sopimuksen tuotot nostetaan hyväksyttävälle riskitasolle vakuutta ja CVA-lisää käyttämällä. Lopuksi tuottojen luottamustasoa tarkastellaan herkkyysanalyysin avulla, ja päätöspuumallia sekä tapaustutkimuksen tuloksia arvioidaan kriittisesti.

Avainsanat Vastapuoliluottoriski, bilateraalinen suojaaminen, energiamarkkinat, päätöspuut

Acknowledgements

I am sincerely grateful to my supervisor, Professor Ahti Salo, for his invaluable guidance and support, whether from pointing out mathematical imprecisions to catching blatantly obvious typos I managed to overlook. His insights and attention to detail have shaped this thesis in more ways than I can count, and it wouldn't be what it is without him.

I want to thank my advisor, D.Sc. (Tech.) Edoardo Tosoni, for his near-daily mentoring and assistance throughout the project. I could always rely on his professional insights and advice, even during the busiest of times. I would also like to thank the entire team at Fortum Oyj; working with all of you has been a pleasure.

Finally, I want to thank my wonderful partner, Rebecca, for her constant encouragement and support throughout my studies and this thesis project. You are the best.

Espoo, 7 July 2025

Joonatan Linkola

Contents

Abstract 3

Abstract (in Finnish) 4

Acknowledgements 5

Contents 6

List of abbreviations 8

1 Introduction 9

2 Literature review 14
- 2.1 Credit risk models . . . 14
  - 2.1.1 Firm-value models . . . 14
  - 2.1.2 Reduced-form models . . . 17
  - 2.1.3 Credit risk in power markets . . . 19
- 2.2 Modeling counterparty credit risk . . . 21
  - 2.2.1 Scenario-based modeling . . . 22
  - 2.2.2 Forward curve modeling . . . 24
  - 2.2.3 Credit value adjustment . . . 26
  - 2.2.4 Industry standards . . . 27
- 2.3 Credit ratings and credit rating agencies . . . 28
- 2.4 Decision analysis and decision trees in energy sector risk management 31

3 Problem context 34
- 3.1 Parameters . . . 36
- 3.2 Collateral requirement and CVA add-on impacts . . . 37
- 3.3 Cash flow analysis of the underlying process . . . 38
- 3.4 Model formulation . . . 41
  - 3.4.1 Scenario modeling: spot prices and counterparty defaults . . . 42
  - 3.4.2 Monthly cash flows in the model . . . 44
  - 3.4.3 Earnings distributions and risk measures . . . 46

4 Case study 51
- 4.1 Data . . . 51
- 4.2 Evaluating the decision tree . . . 53

4.3 Collateral requirement and CVA add-on optimization 55
4.3.1 An illustrative example 57
4.3.2 Base case: confidence level $\alpha = 10\%$ 59
4.3.3 Sensitivity analysis 61
4.4 Critical reflections and potential future development 66

5 Conclusions **70**

References **72**

List of abbreviations

CVA	Credit Value Adjustment
CVaR	Conditional Value at Risk
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortization
MtM	Mark-to-Market
OTC	Over-the-Counter
PD	Probability of Default
PPA	Power Purchase Agreement
VaR	Value at Risk

1 Introduction

The energy sector is one of the most critical infrastructures in our modern society. Most, if not all, other industries depend on electricity and fuels in some form, which highlights the energy sector's importance and the need to keep it resilient and secure. The electrical power industry forms a key part of the energy sector, and it includes both the electricity wholesale and retail markets. In the wholesale market, electricity producers sell electricity in large quantities to electricity providers and large industrial consumers, such as steel manufacturers and data centers. Conversely, in the retail market, electricity providers sell electricity to smaller businesses and households.

The global electricity consumption has risen continuously over the past half a century. In 2023, global electricity consumption reached 27 000 TWh, and has increased at a relatively linear rate of approximately 3% annually since 2010 (Statista, 2025). Recent estimates point towards a similar increasing trend: electricity consumption is estimated to range between 31 000 and 36 000 TWh by 2030 (increase of 11-33% from 2023), and between 52 000 and 71 000 TWh by 2050 (increase of 93-163% from 2023) (Statista, 2023). It is clear that a constant growth of electricity consumption would not be possible without the support of electricity production and electrical grids; reliable, affordable, and sustainable electricity is needed.

Electricity differs from many other commodities, which introduces a unique challenge for electricity producers and retailers: electricity cannot be stored. As a result, electricity markets have to be extremely adaptive and flexible to constant changes in supply and demand. Whenever supply is greater than demand, prices increase, and vice versa, causing inherent volatility to electricity. Therefore, sudden economic disturbances can have a major impact on the price of electricity and the entire electricity market.

A recent example that highlighted the volatility of electricity markets came during the COVID-19 pandemic, which started in early 2020. As the virus began to spread globally, many countries initiated immediate and strict lockdown restrictions. This slowed down the global economic activity noticeably, which came with decreased electricity demand and thus a drop in electricity prices. This effect was particularly noticeable in the European electricity markets, as price drops were the largest compared to other parts of the world, and the European average electricity price level was the lowest it had been in six years (Zhong et al., 2020).

Once countries were able to relieve the lockdown measures, transportation and industrial activity began to approach their normal levels, and there was again a higher electricity demand. However, due to production cuts and diminished investment earlier in the pandemic, the supply side was unable to keep up with the sudden increase in demand. This supply shortage led to record-high electricity prices in 2021, and these supply-side challenges were further escalated after the Russian invasion of Ukraine, which started in early 2022.

This recent global crisis underscored the high volatility of the electricity markets. Electricity prices can fluctuate immensely as a result of imbalances in supply and demand during social and geopolitical incidents. Furthermore, wholesale electricity markets operate on a day-ahead basis, meaning that electricity prices become available for both buyers and sellers only a day in advance, which adds to the volatility. Accurately estimating future electricity prices or causes of supply and demand imbalances is extremely challenging but also very important to prevent such incidents as effectively as possible.

To minimize financial losses in extreme events where prices could suddenly plummet, electricity producers have traditionally hedged their future electricity sales price. Hedging is a common strategy where the goal is to reduce potential losses by purchasing an investment that offers an opposite position to an existing investment. A perfect hedge would be exactly inversely correlated with the underlying investment—when the value of the original investments rises by some amount, the value of the hedge decreases by the same amount, and vice versa—effectively removing all risk related to price movement. Electricity producers hedge their future sales price by selling electricity volumes in advance using financial instruments such as futures contracts or Power Purchase Agreements (PPAs).

Energy companies have commonly traded and hedged through exchanges such as Nasdaq or the European Energy Exchange (EEX), where a central clearinghouse manages all transactions anonymously between buyers and sellers. However, in recent years, more companies have negotiated hedging agreements as over-the-counter (OTC, i.e., privately negotiated) derivatives transactions. Such bilateral agreements are often more flexible than traditional agreements traded in exchanges, because the contractual terms can be negotiated more freely and be tailored to each unique customer.

In hedging agreements, both parties agree on the electricity volume and price for the duration of the contract, which reduces the market risk related to electricity price volatility. However, while hedging agreements reduce market risk stemming from electricity price volatility, they give rise to another type of risk: counterparty credit risk. Counterparty credit risk can be defined as the risk that either contracting party fails to meet their financial obligations, such as contractual settlements or physical electricity delivery. In bilateral trading, both parties are naturally exposed to the risk that the other party may default. This exposure, known as counterparty credit exposure, refers to the potential financial losses that may occur in the event of the other party's default. It is dynamic because it changes over time, depending on the financial statuses of both parties and the mark-to-market (MtM) valuation of the underlying agreement (Canabarro and Duffie, 2003).

There are many ways market participants can mitigate this risk. A common safeguard is to post a collateral requirement to the contracting counterparty for the duration of the agreement. This collateral (e.g., funds or assets) acts as security for the contract duration, such that if either party were to default, the collateral could then be used to cover the other party's financial losses. Another strategy from the seller's perspective is to include a credit value adjustment (CVA) add-on to the contract price, which is an additional charge that compensates the seller for the counterparty credit risk embedded in the contract.

Counterparty credit risk is a topic of extensive discussion in the electrical power industry, and this trend is especially apparent because bilateral contracts have become increasingly common in recent years. As more and more decarbonization measures are implemented worldwide to boost the global clean energy transition, access to affordable and stable energy is very important. The European Commission underlines this trend in Europe, as it is, together with the European Investment Bank (EIB), launching a €500 million pilot programme to support the uptake of corporate PPAs in small and medium-sized enterprises and energy-intensive industries. As part of this programme, the EIB will counter-guarantee a portion of the PPAs. The EIB is also launching a €1.5 billion "Grids manufacturing package", which is intended to boost production measures in European supply chain businesses. (European Commission, 2025)

The financial strength of the contracting parties, especially in bilateral agreements, is critical to ensure that the agreements are not terminated prematurely. Similarly, ensuring that collateral requirements, CVA add-ons, or other financial safeguards are appropriately set is important; insufficient requirements increase counterparty credit risk, whereas excessive demands impose unnecessary financial strain and may drive potential customers to competitors.

In finance, probabilistic risk models can be used to manage counterparty credit risk and assess appropriate sizes of collateral requirements and CVA add-ons in bilateral contracts. In these models, risk is quantified by two key characteristics: the likelihood and impact of an event. Probabilistic risk models simulate different market conditions and counterparty default events to produce a distribution of potential outcomes. This allows organizations to make informed decisions about their future operations because these decisions will be backed up with realistic market estimations.

This thesis has been written for Fortum Oyj, a large Finnish energy company that produces electricity and trades both physical electricity and financially settled electricity derivatives in the Nordic electricity markets. Fortum hedges a portion of its electricity production to mitigate market risk coming from potential longer periods of low electricity prices. Some of the hedging agreements Fortum is using are large long-term PPAs. To mitigate counterparty credit risk emerging from these PPAs, Fortum may impose additional financial requirements on the contracting counterparties in the form of collateral requirements and CVA add-ons to secure these contracts at a risk level considered acceptable by Fortum.

The scope of this thesis considers a single large long-term PPA and its estimated earnings. The earnings originating from accepting the contract are compared to those from declining it, and to some alternative hedging strategies. Additionally, the collateral requirement and the CVA add-on are optimized from Fortum's perspective, with the goal to ensure that downside earnings from accepting the contract meet predefined target risk levels. These target levels also reflect competitiveness aspects: Fortum wants to limit its financial risk but also offer contractual terms that are attractive to potential customers. Finally, a sensitivity analysis on the confidence level is performed to examine how changes in risk attitude affect the optimization. In this context, the confidence level refers specifically to how much of the lower (downside) tail of the distribution is taken into account when calculating the downside earnings.

The thesis builds on a large internally used probabilistic financial simulation model that estimates the company's future financial position by simulating various market conditions and potential counterparty default events within its bilateral portfolio. As part of this thesis, the model is extended to facilitate the analysis of earnings from a single PPA over its entire contract horizon.

To support decision-making on whether to accept a presented PPA and under what terms, the problem is structured as a decision tree with three top-level decision branches. Two of these represent possible actions that Fortum may take regarding a contract proposal: either *decline* the contract and instead hedge the same volume using a portfolio of smaller contracts with many counterparties, or *accept* the contract and the associated counterparty credit risk. On the contrary, the third branch represents a set of reference alternatives that serve as benchmarks for the original contract. These can provide insights into how the earnings would change under altered assumptions, for example, if the contracting counterparty had a higher credit rating or if the contract duration was different. The decision tree is then evaluated using a hypothetical but realistic case PPA, after which the collateral requirement and the CVA add-on are optimized, and sensitivity to the confidence level (and thus risk attitude) is analyzed. Finally, the decision tree model and the results of the case study are critically assessed, and some potential developments for future work are discussed.

The structure of the thesis is as follows. Section 2 reviews the literature and research related to the thesis topic. The problem setting of the thesis and the decision tree model formulation are described in Section 3, while a case study applying the model is presented in Section 4. Finally, a summary of the overall thesis and the case study results is provided in Section 5.

In summary, the following research questions will be addressed in this thesis:

Q1: What are the impacts on earnings of different decision alternatives and reference alternatives regarding a proposed bilateral contract?

Q2: At what levels of collateral requirement and CVA add-on is entering the proposed contract an attractive option?

Q3: How do different confidence levels for the earnings estimates influence the collateral requirement and the CVA add-on?

2 Literature review

This literature review is structured as follows. Section 2.1 explores different types of credit risk models and their extensions in broader financial applications, as well as credit risk applications in the power markets. Section 2.2 outlines common strategies in modeling counterparty credit risk, with a focus on electricity markets. Section 2.3 discusses the use of corporate credit ratings in credit risk analysis, credit rating agencies that assign these ratings to organizations, and addresses criticism towards these agencies. Finally, the review is concluded in Section 2.4, which provides an overview of decision analysis and presents some applications of decision analysis and decision trees in risk management within the energy sector.

2.1 Credit risk models

McNeil et al. (2011) categorize financial credit risk models into two broader categories: *firm-value models* (also known as *structural models*), and *reduced-form models*. The credit risk models in these two categories have one main difference, which is how the probabilities of default (PDs) of the analyzed firms are determined.

In firm-value models, the PDs are determined internally in the model. These models assume that default occurs whenever the firm's asset value falls below some threshold, and hence, the reasons for the default events can also be analyzed to some extent. On the other hand, in reduced-form models, the PDs are determined exogenously from the model, meaning that the models themselves do not capture any information about potential processes leading to default events (McNeil et al., 2011). Sections 2.1.1 and 2.1.2 present various firm-value and reduced-form credit risk models and show how they have been applied and extended in the literature, and Section 2.1.3 presents credit risk applications specifically in the power markets.

2.1.1 Firm-value models

The Merton model (Merton, 1974), nowadays considered the predecessor of all firm-value models, is one of the most widely recognized models in credit risk analysis. The Merton model provides a framework for analyzing the relationships between a firm's asset value, debt structure, and default likelihood. Understanding these relationships is crucial when assessing a firm's overall financial status.

In the Merton model, it is assumed that the firm's asset value follows a continuous stochastic process, typically a geometric Brownian motion. Additionally, the model assumes that the firm is financed by equity (money coming from shareholders) and debt (borrowed money). The debt is represented by a single zero-coupon bond with a given face value and maturity. It is also assumed that the firm has no additional liabilities or complicating factors, such as dividend payments or other debt issuers. The firm is said to default exactly at debt maturity if its asset value falls below the face value of the zero-coupon bond (debt). This default event is modeled using the Black-Scholes framework (introduced by Black and Scholes (1973)), where the firm's equity is treated as a European option on its assets, such that the default event corresponds to this option expiring out of the money.

This option-based framework of the Merton model can provide a better understanding of shareholder and debtholder dynamics. An example is the incentive for shareholders to prefer higher asset volatility because this increases the potential upside, but also the risk, of the investment.

The Merton model has strongly impacted the development of other modern credit risk models and is an important tool in both academic research and practical applications even today. Despite its simple assumptions, such as ignoring dividend payments and modeling default only at bond maturity, the Merton model and its many extensions are still widely used for corporate liability pricing and corporate creditworthiness assessments.

One famous extension was presented by Black and Cox (1976), where the limitation of the Merton model's stiff default timing was addressed. The authors included a so-called time-dependent barrier in the model, which helps capture default events before bond maturity. In this model, a default is assumed to occur whenever the firm's asset value falls below a time-dependent threshold. This extension later became known as the Black-Cox model and allows for more flexible credit risk assessment compared to its predecessor because firms may naturally face solvency issues well before debt maturity.

The Black-Cox model has been widely applied in corporate credit risk modeling, and it has seen various extensions in the literature that build on its foundation. Katz and Shokhirev (2010) generalized the Black-Cox model by introducing diffusion in a linear potential with a radiation boundary condition. This extension allows the model to capture the possibility of a firm avoiding default, even when its asset value momentarily drops below its liabilities. The authors mention that this improves the model's ability to estimate credit risk, especially for shorter time horizons.

The Black-Cox model has also been extended to incorporate recovery risk, as demonstrated by Cohen and Costanzino (2017), who developed the Stochastic Recovery Black-Cox model. This two-factor firm-value model introduced a second recovery risk driver. In their work, the authors were able to separate the default and recovery risk components, which allows for a more detailed modeling of defaultable bonds and credit default swaps.

Feldhütter and Schaefer (2018) introduced a novel calibration method for the Black-Cox model by including broader historical credit spread data than was used previously. This enhanced the Black-Cox model's ability to match real observed credit spreads. Their approach minimized the distance between historical default rates and those coming from the model across several maturities and credit ratings. This use of broader historical data allowed the authors to narrow the variation between model predictions and historical credit spreads, which improved the Black-Cox model's overall predictive power.

Another famous and widely applied extension of the Merton model is the KMV model (Crosbie and Bohn, 1993). This model was developed by the KMV Corporation (the abbreviation comes from the surnames of the founders: Stephen Kealhofer, John McQuown, and Oldřich Vašíček), which was a firm specializing in credit risk analysis and quantitative modeling in the 1990s, before being acquired by Moody's in 2002.

The Merton and KMV models have one notable difference: their approach to determining the PDs of firms. As was discussed previously, the Merton model determines the PD based on the relationship between a firm's asset value and its debt obligations; however, the KMV model instead estimates a firm's expected default frequency (EDF), which represents how likely the firm is to default during a one-year period. In the KMV model, distance to default (DD) measures the number of standard deviations the firm's asset value is away from its debt obligations (i.e., the default threshold). EDF is then derived using DD and historical default data and included in place of the Merton model's default formula. By doing this, the model better accounts for the firm's total debt obligations and matches observed default behavior more accurately. (Crosbie and Bohn, 1993)

The KMV model has seen a lot of use in various applications throughout the years. With the rise of computational advancements, machine learning methods have also been applied to the KMV model to improve its predictive power. Li et al. (2023) integrated convolutional neural networks, long short-term memory networks, and an attention mechanism into the KMV model. Wu and Wu (2016) studied the KMV model in conjunction with random forests (RFs) to develop credit rating prediction models, and demonstrated that the addition of RFs is a beneficial one. A similar application came from Yeh et al. (2012), who integrated the KMV model's DD and EDF with RFs and rough set theory for credit rating predictions.

One aspect from Yeh et al. (2012) worth noting is the classification accuracy of credit rating predictions. They mention that combining DD and EDF with financial variables in their model is a key factor in why it outperforms many of the models that only use traditional financial variables. This underlines the importance of market indicators when assessing credit risk. The study also used RFs in a novel way in the predictive variable selection step, which greatly simplified the feature selection process and ensured that only relevant variables were ultimately used.

2.1.2 Reduced-form models

In reduced-form models, the credit risk of firms is modeled without any information about their asset or debt valuations. Instead, observable market data is used for these valuations, and the firm-specific PDs are determined exogenously from the model. The Jarrow-Turnbull model (Jarrow and Turnbull, 1995) was one of the first reduced-form models. The default event is modeled as a Poisson process, controlled by a stochastic default intensity (also referred to as the hazard rate). This default intensity is used to determine the firm's PD at any given time, assuming that no prior default has occurred. The benefits of using this default intensity stem from its ability to account for many realistic market factors and to allow random and sudden changes in default risk.

The key innovation of the Jarrow-Turnbull model compared to earlier firm-value models came from the stochastic term structures of default-free interest rates and credit risk spreads, which are given as exogenous input. These capture many aspects of credit risk that were previously not accounted for. They provide a foundation for an arbitrage-free pricing framework, and they are key in pricing complex derivatives, such as options on risky assets. This (at the time) novel feature overcame many of the limitations of earlier firm-value models because it effectively provided a more practical approach for pricing risky derivatives while separating the PD modeling from the firm's implied financials. (Jarrow and Turnbull, 1995)

The Jarrow-Lando-Turnbull (JLT) model (Jarrow et al., 1997) is an extension of the Jarrow-Turnbull model, in which the possibility of credit rating transitions is incorporated with the use of a Markov process. This addition is intuitively very appealing because the original assumption of a constant credit rating is often not particularly realistic. This is especially true when modeling credit risk over long time horizons, which could span decades in many real-life applications. Unlike the default intensities that were part of the original Jarrow-Turnbull model, the JLT model treats defaults as discrete-state Markov chains of the underlying credit ratings. This allows for realistic and endogenous credit rating migrations of firms, which result in changes in credit risk and PDs over time. The authors mention that this refinement in the JLT model produces results that align with observed credit rating data and helps it better capture realistic credit risk dynamics.

Millossovich (2002) presented a further extension to the JLT model, where a stochastic recovery rate is included. This recovery rate represents the amount that one is expected to recover after a default. In the model, the recovery rate depends on the firm's rating at the time of default. The default state is expanded into multiple classes, each corresponding to a different recovery rate based on the firm's possible earlier ratings. These classes reflect various bankruptcy outcomes with their distinct recovery rates. This inclusion allows for a more nuanced treatment of recovery dynamics because, rather than being a fixed parameter, the recovery rate is determined by the firm's credit quality before default. Additionally, the model accommodates the possibility of firms transitioning between these default states, which can capture default-related complexities at a finer granularity than before.

A more recent modification to the Jarrow-Turnbull model was proposed by Krabichler and Teichmann (2024), who analyzed a financial market consisting only of zero-coupon bonds subject to both credit and liquidity risk. The Jarrow-Turnbull model was extended by incorporating the relationship between credit and liquidity risk through a foreign exchange analogy. Their framework included possible delays in the recovery process that resulted from liquidity squeezes, which can lead to uncertainty in recovery rates at the time of default. A key feature of this approach is the introduction of two filtrations: the more idealistic and impractical of the two represents complete information about the entire financial system (i.e., the full state of recovery, defaults, and credit quality), while the more realistic one represents limited, observable market information, such as liquidity squeezes. This extended model can capture such real-world complexities, where liquidity issues and delayed recoveries complicate credit pricing.

The Duffie-Singleton model (Duffie and Singleton, 1999) builds upon earlier reduced-form models. The model has a more generalized structure that models the default hazard rate, the recovery rate, and the risk-free rate by using a set of state variables following a Markov process. In this model, defaultable bond prices are modeled as exponentially affine functions of the state variables. A key feature in the Duffie-Singleton model is that the recovery rate is assumed to be a fixed fraction of the bond's market value. Unlike many earlier models, this model allows credit spreads and default intensities to be negatively correlated with default-free interest rates. This improves its flexibility and computational tractability and makes it more suitable for both econometric estimation and market calibration. (Schlögl and Schlögl, 2010)

Takahashi et al. (2001) utilized the Duffie-Singleton model to value convertible bonds (CBs) with default risk included. They applied the model using data from the Japanese CB market and compared its performance with other models. They demonstrated that the model could price different securities, such as CBs, non-convertible corporate bonds, and equities, issued by the same firm. The authors modeled the default hazard rate as a function of the underlying stock price and recognized that PD is often negatively correlated with the stock price. They also noted that while the model assumed the risk-free interest rates to be deterministic to keep the model simpler, it could easily be extended to accommodate more complex term structures.

2.1.3 Credit risk in power markets

Counterparty credit risk can vary considerably between different fields and their respective markets, depending on the types of market participants and the kinds of agreements that are common. The power markets are unique in that OTC derivatives (such as bilateral PPAs) are widely used for both hedging and trading, which together with the special characteristics of electricity (such as high price volatility, seasonality, non-storability, and major role of physical delivery) bring their unique counterparty credit risk challenges to market participants, for example, when compared to traditional financial markets. These aspects are also visible at the firm level because geopolitical incidents or other unexpected sudden price fluctuations can drastically alter the market environment and pose immense financial stress on individual firms. The increase of renewable energy sources in recent years introduces further complexities (e.g., different trading horizons and forecast uncertainty) to the mix. These are just some of the reasons why effective credit risk management is crucial to maintain the robustness of the energy system.

One does not have to look too far back in history to find an example of extreme energy market stress. In 2022, the natural gas prices reached record-highs, totaling approximately three times higher than what the average price had been for the last few years. The main cause for this incident was geopolitical tensions, which further escalated during the Russian invasion of Ukraine at the beginning of 2022 and led to large-scale disruptions in the energy supply side. This event caused major counterparty credit stress and even defaults throughout the European energy markets, notably also for large energy firms, which faced immense financial pressure during this time. Furtuna et al. (2022) also highlighted that another reason for the increased counterparty credit issues in Europe was the widespread margin calls that were set off by the extreme price volatility.

Furtuna et al. (2022) also brought up an important and interesting observation. While the extreme price volatility during this time increased counterparty credit risk in firms, another major contributor to the financial strain came from the inability to hedge effectively. In their study, the authors showed how large firm-level exposures can emerge as a result of sudden price increases in the market. Especially firms that rely on long-term hedging can experience major liquidity issues in such situations. In energy sector credit risk management, it is important to account for extreme adverse situations, which may be unlikely but certainly possible, as the crisis showed. Many market participants had underprepared for such extreme market conditions, which is also one contributor to the total impact of the crisis. Their study highlighted the importance of real-time risk assessment and scenario-based stress testing in credit risk management processes. Overall, the event demonstrated that risk management practices may require some readjustment.

Denton et al. (2003) examined credit risk management in the energy sector, with a focus on the power markets. They highlighted Value at Risk (VaR) as a traditional risk metric that is commonly used to estimate credit exposure. However, according to them, VaR has its clear limitations and may not be suitable to fully capture all the complexities of the electricity markets. The authors noted: "*90% of the trades in large power producers' books may be physical trades or hedges on them, with only 10% 'derivative' trades that are deemed speculative and have to be marked to market.*" They also emphasized that the main source of credit risk in the energy markets is related to counterparty risk, whether involving physical delivery obligations or making contractual payments for agreed trades.

An approach for estimating counterparty credit exposure in long-term PPAs was presented by Edge (2015). Their method offers a computationally efficient way to estimate exposures in PPAs because it accounts for varying yearly volumes and pricing dependencies that are common in PPAs. The author also mentions that their method takes into account challenges related to long contract horizons and high price volatility, which adds value to the method. Although they note that the current model is tailored to evaluate single contracts only, it could be extended, although with increased computational complexity, to incorporate the accumulation of credit risk across a portfolio of contracts with the same counterparty.

Chang et al. (2015) studied collateral requirements in the electricity markets. They presented a reserve-forecast approach to tackle a common deficiency in collateral requirement methods, which often relied on static historical data. The method was then applied in a case study in which ERCOT data with predictive modeling and changing system conditions were used in credit risk estimation. The case study revealed a strong correlation between system-wide reserves and real-time electricity prices and showed that forecasting reserves up to a week ahead noticeably improved the collateral requirement sizing. Setting these requirements accurately is important because overly high collateral requirements can cause unnecessary financial strain on counterparties and thus increase the likelihood of the contracts terminating prematurely. Chang et al. (2015) also pointed out challenges with this forecast-based method. As an example, the credit calculations using such methodologies need to be transparent in order not to allow for strategic market behavior using available credit estimates.

Mori and Umezawa (2007) incorporated random forests to enhance credit rating prediction in the energy markets. They trained an RF model on financial data and demonstrated that their model could predict the credit ratings of 19 Japanese energy companies. They also showed that the RF model outperformed both of their reference machine learning models, one of which used multilayer perceptrons and the other classification and regression trees, by approximately 10% and 5%, respectively. The good performance of the RF model was largely attributed to its ability to handle high-dimensional data and capture complexities that the other models were not able to capture equally well. The authors also drew attention to the RF model's robustness against overfitting in the case study, which is an important feature of robust classification models.

2.2 Modeling counterparty credit risk

Counterparty credit risk plays a central role in all financial and energy markets, where insufficient risk management practices can lead to significant financial losses and instability among market participants. Accurate and reliable counterparty credit risk modeling is particularly important in the energy markets due to the high use of OTC derivatives for both hedging and trading. Risk models enable market participants to determine exposures and compare them against limits, estimate and hedge against market risk more effectively, and meet economic and regulatory capital requirements. All three of these activities are necessary to prevent and mitigate losses in unexpected adverse situations.

2.2.1 Scenario-based modeling

A general framework for scenario-based counterparty credit exposure modeling is outlined by Zhu and Pykhtin (2007). In this context, scenarios refer specifically to different simulated market and counterparty default outcomes. The authors explain that the scenario-based framework consists of three steps: scenario generation, instrument valuation, and portfolio aggregation.

In the first step, various potential market conditions and counterparty default events are simulated for a range of future dates, which are often referred to as *time buckets*. The goal of these scenarios is to account for various risk factors (e.g., interest rates and commodity prices) across a wide range of potential market outcomes to gain a comprehensive understanding of the general expected turnout while also accounting for more unlikely and adverse outcomes. They explain that these varying market conditions are commonly obtained with the use of historical market data and stochastic processes. Two methods that can be used to generate these scenarios are discussed in the paper: *Path-Dependent Simulation*, which accounts for continuous development of market conditions throughout the simulation horizon, and *Direct Jump to Simulation Date*, where conditions for the simulated date are obtained without using information from previous conditions. (Zhu and Pykhtin, 2007)

In scenario generation, keeping computational constraints in mind is extremely important. The portfolio of contracts for a firm may be large, and the contract time horizons may span many years, even several decades. This is why one cannot simply generate scenarios for all possible combinations of counterparty defaults at any arbitrary time granularity, because this would lead to an astronomical number of scenarios and thus computationally infeasible simulation times. Therefore, using around a few thousand simulated scenarios evaluated at discrete simulation dates often offers a good balance between modeling accuracy and computational tractability. (Zhu and Pykhtin, 2007)

After the scenarios have been generated, the next step in the framework is instrument valuation, where all trades within the firm's portfolio are evaluated in each simulated market scenario. Because the number of scenarios is typically very large, it is important to use simplified valuation models and other suitable approximations to ensure that the framework captures the most critical characteristics and risk profiles of the entire portfolio. (Zhu and Pykhtin, 2007)

In the final step of the framework, portfolio aggregation, the various simulated scenario outcomes are aggregated together with their corresponding probabilities to produce a distribution of possible outcomes. These distributions enable the calculation of key exposure metrics, such as current exposure (CE), expected exposure (EE), and potential future exposure (PFE), all of which are important statistics in counterparty credit risk management. The authors note that this framework offers a powerful and flexible tool to assess counterparty credit risk because it can account for various market conditions and portfolio structures. (Zhu and Pykhtin, 2007)

Mausser and Rosen (2001) studied portfolio credit risk optimization using a similar scenario-based approach to the one outlined above. They presented three different optimization models that focus on restructuring portfolios to improve the tradeoff between risk and return. The first model creates the trade risk profile and then optimizes the hedge position of a single asset. The second model considers all positions simultaneously to minimize the regret factor of the portfolio. The third model uses parametric programming to construct a credit risk-return efficient frontier. The authors demonstrated how scenario-based optimization can help mitigate credit risk in bond portfolios and showed that regret is an attractive risk measure.

In contrast to the framework described by Zhu and Pykhtin (2007), Wang and Ziegel (2021) studied scenario-based risk metrics, with a focus on economic scenarios and risk adjustments for expected shortfall (ES, also known as conditional value at risk, CVaR) at the portfolio level. They presented various risk metrics (such as Max-ES, Max-VaR, and their variants) and emphasized their reliability under different market conditions. They also presented a case study based on empirical data, in which the usefulness of these risk metrics in high-stress economic scenarios was demonstrated. The results were also linked to Basel III regulatory requirements to bring attention to the importance of scenario-based analysis in capital buffer adjustments. Overall, their work contributed to improving resilience against financial shocks.

In their paper, Skoglund and Chen (2016) examined the impacts of rating momentum for macroeconomic stress testing and scenario-based analysis for a credit risk application. According to them, one downside of traditional credit risk models is their assumption of a Markov property, meaning that changes in credit ratings depend only on the previous state and not past rating movements. The authors demonstrated that accounting for past rating movements is crucial to estimate potential credit loss projections more accurately. They compared their rating momentum model to standard Markovian models using nine-quarter (CCAR and DFAST) and twelve-quarter (EBA) macroeconomic stress test scenarios. In the case study, they showed that disregarding rating momentum led to noticeable underestimation of losses, especially when considering portfolios consisting mainly of lower credit ratings.

They highlighted the importance of this effect in regulatory stress testing frameworks so that more realistic credit rating migrations at the portfolio level would be accounted for.

A recent computational improvement to the scenario-based modeling approach was presented by Matsakos and Nield (2024), who incorporated quantum computing to enhance the simulation of different risk factors. Specifically, they implemented various firm-value and reduced-form credit risk models and simulated credit rating migrations and survival probabilities within these models over extended horizons using quantum circuits. Their findings showed how quantum computing can offer increased accuracy and new modeling solutions that are impossible with traditional computers. However, they also highlighted the current limitations of their approach, such as the available qubits and the depth of the circuit, which may limit the practical applicability of their methodology in the near future.

2.2.2 Forward curve modeling

Forward curve modeling is common in the energy markets because these curves are used in a range of applications, such as commodity pricing, risk management practices, and decision-making processes. The forward curve represents the expected price evolution of some commodity over time, and it can for example be used to price a commodity with a future delivery period according to today's market environment. Forward curves also naturally have an important role in counterparty credit risk management. A few examples of their important use cases include valuations of derivative instruments, calculations of portfolio exposures, and assessments of different hedging strategies. In their extensive book, Eydeland and Wolyniec (2002) explored various methods for constructing forward curves and pointed to their importance in risk management practices.

Among the many methods they discussed, Eydeland and Wolyniec (2002) presented a single-factor model that relied on geometric Brownian motion with zero drift. The attractiveness of the model came from its simplicity and applicability for standard derivative instrument pricing. However, due to its simplistic nature, the model did not fully capture the relationships between different types of forward contracts and could not account for the dynamics of the entire forward curve.

With the goal to address these limitations, the authors also explored the Heath-Jarrow-Morton (HJM) multi-factor model (introduced by Heath et al. (1992)), where the entire stochastic evolution of the forward curve is accounted for. The HJM multi-factor model includes several sources of randomness, each represented by their deterministic perturbation function, and therefore offers greater flexibility compared to the single-factor model from earlier. This addition allows for a more well-rounded analysis of price movements across different maturities.

Another example is the Schwartz–Smith two-factor model (Schwartz and Smith, 2000), in which the logarithm of the forward prices is decomposed into two parts: a long-term equilibrium level, which represents the evolution of spot prices during longer time periods, and a short-term deviation, which represents more temporary random fluctuations. This decomposition gives an intuitive interpretation of the forward curve because it retains the lognormality of the forward prices and captures key features (e.g., mean-reversion and term structure dynamics) of common commodity markets. However, despite its many benefits, the approach presents challenges, particularly in the energy markets, because estimating the needed parameters can be difficult with the limited available historical information. (Eydeland and Wolyniec, 2002)

Lucia and Schwartz (2002) studied the unique characteristics of electricity price movements and their implications on derivatives pricing. They focused on forward curve modeling in the Nordics (specifically, Nord Pool) and highlighted the important role of systematic seasonal components in the forward curve creation process. This seasonality is particularly apparent in the Nordics because business and weather cycles have significant impacts on the overall electricity demand throughout the year. The authors analyzed data from Nord Pool and constructed one- and two-factor models with a specially designed deterministic seasonal component. Both models were able to account for the high price volatility and storability challenges of electricity. The authors emphasized the importance of including a sinusoidal component in the forward curves in the Nordics because it is key to allowing for comprehensive and accurate electricity forward price modeling.

Audet et al. (2002) proposed a parameterized model for electricity forward curves, which was specifically tailored for the Nordic electricity markets. Their study addressed challenges unique to Nord Pool, such as the heavy reliance on hydro and the distinct seasonal variability caused by changing weather conditions. The parameterized model accounted for time dependency in spot price volatility, maturity effects, and correlations between forward curves. They also mentioned that their model is suitable for many practical use cases, such as planning hedging strategies, optimizing power plant operations, and pricing derivatives, because it offers deep insight into spot and forward price relationships in the Nordics.

Another study focusing on the Nordic power markets came from Koekebakker and Ollmar (2005), who explored forward curve dynamics in Nord Pool. In their case study, they used historical futures and forward contract data with maturities up to two years, traded during the period 1995–2001. The authors analyzed the volatility structure of forward prices using a multi-factor term structure model within the HJM framework. In their work, the authors constructed a continuous forward curve using discrete forward and futures prices. They then applied principal component analysis to further study the factor structure of the curve. The analysis showed that the two-factor model was able to explain approximately 75% of the price variation, which is considerably lower than in other commodity markets, where the value is often approximately 95%. Furthermore, the correlation between short and long horizon forward prices is smaller compared to other markets, which is likely a result of the non-storability of electricity. The authors concluded that modeling the entire forward curve in the electricity markets may not be as beneficial as in other markets, and highlighted related challenges, such as hedging long-term commitments with short-term contracts.

2.2.3 Credit value adjustment

Credit value adjustment (CVA) is a risk measure that accounts for the potential losses resulting from counterparty defaults when estimating counterparty credit risk. Zhu and Pykhtin (2007) define CVA as the difference between the portfolio's risk-free valuation and its true value (also referred to as the "risk-adjusted" value), which takes into account the possibility of a counterparty defaulting. In other words, CVA can be thought of as the market value of counterparty credit risk. The authors also highlighted the importance of incorporating counterparty risk into portfolio valuations through CVA. In practice, CVA can be estimated using the risk-neutral expected value of discounted losses when counterparty risk has been accounted for. (Zhu and Pykhtin, 2007; Gregory, 2007)

Although the CVA framework is a well-founded and widely used counterparty credit risk measure, it has certain flaws. Cherubini (2013) introduced an improvement to the CVA framework, which accounts for the dependencies between credit risk and underlying asset price. The main focus of the paper was the use of copulas, which were utilized to study these dependencies. The inclusion of these dependency structures resulted in a model that accounted for extreme wrong-way risks, such as perfect positive or negative correlation between the counterparty's PD and market movements. This extended approach provided a more well-rounded view of CVA and emphasized how dependency structures can have a significant impact on risk profiles, particularly in long-term contracts. The author also presented a case study in which the framework was applied to a 30-year swap contract with a BBB-rated counterparty. The case study showed that counterparty credit risk could be measured more accurately with the improved framework, as it captured correlations between credit and market risk.

Another comprehensive framework was introduced by Ballotta et al. (2019), who studied the CVA of equity and commodity products. They utilized reduced-form models together with Lévy processes to incorporate risk mitigation strategies such as collateral and margin requirements. Their model included Monte Carlo simulation and methods based on Fourier transform for CVA calculations, and captured key dependencies stemming from counterparty default event timings and underlying position valuations.

Many other authors have also addressed flaws in the CVA framework in recent years. An example was presented by Xiao (2015), who employed a Least Square Monte Carlo method for CVA calculations. They explained that the strength of the model lay in its default time calculations. Previous models had heavily relied on a static default time, whereas their model addressed this by accounting for the probability distribution of various default times. This allowed the model to better capture CVA-related wrong-way and right-way risks. A similar approach was adopted by Trinh and Hanzon (2022), who introduced the Monte Carlo-Tree method for option pricing and CVA calculations. Their approach accurately priced CVAs for American options and accounted for unilateral and wrong-way risk, using a combination of Monte Carlo methods and binomial tree methods. Both of the aforementioned methods (Xiao, 2015; Trinh and Hanzon, 2022) provided effective solutions for pricing defaultable derivatives.

2.2.4 Industry standards

While theoretical models provide numerous ways to analyze and assess counterparty credit risk, it is also important to examine models that have been widely applied in practice. This section provides a short summary of industry standards related to counterparty credit risk in the energy sector.

The energy sector has many standard practices for different processes, such as setting collateral requirements and performing CVA calculations, both of which were discussed in previous sections. However, the exact methodologies vary between companies, as they often do not disclose their counterparty credit risk management practices. Nevertheless, some general conclusions can be inferred from the literature and market guidelines.

A common theme in such risk management processes is the replacement cost. This refers to the total cost of replacing a contract with an equally attractive one at current market prices in a situation where the contracting counterparty defaults or otherwise fails to fulfill their obligations (Zhu and Pykhtin, 2007). The replacement cost is strongly linked to counterparty credit risk: if a contracting counterparty defaults, the other party would naturally seek an equally attractive replacement contracting opportunity at that time, so that the impact of the default on the anticipated default-free valuation would be minimal and the realized outcome would resemble the initial expectation as closely as possible. Additionally, related to commonly used risk metrics, Capponi (2012) mentions that a 95% confidence level is widely considered standard market practice in PFE calculations within the energy markets.

Collateralization agreements rely on periodic MtM valuations of the underlying contracts because these valuations reflect the current market prices of the underlying assets. In practice, the way in which the MtM valuation is calculated is initially agreed upon by the involved parties, and then the agreement's MtM valuation is evaluated in a consistent manner throughout the contract's delivery period. If the total MtM valuation of a contract exceeds some threshold, one party may be requested to post collateral (e.g., cash or other assets) as a security for the contract (Hull, 2006). In the event of default, the replacement cost for the remaining contract can also be estimated using the contract's MtM valuation.

It is clear that the MtM valuation and credit exposure are very closely linked. In their comprehensive book, Eydeland and Wolyniec (2002) also stated the following: *"if the market moves in the right direction (i.e., the contract ends up being in-the-money), the contract holder becomes exposed to counterparty credit risk"*, which further highlights the connection of the two quantities. Moreover, many regulatory frameworks, such as the Basel III framework (Bank for International Settlements (BIS), 2010) and the European Market Infrastructure Regulation (EMIR) (European Parliament and Council, 2012), are important in shaping the market environment for individual participants. Transparency is important, and these entities aim to keep the markets as stable and transparent as possible.

2.3 Credit ratings and credit rating agencies

Credit ratings are evaluations of the creditworthiness of an entity (e.g., a firm or a government). These ratings are designed to provide an understandable creditworthiness label that can guide investment decisions and aid risk management, so that investors or institutions do not need to perform independent and extensive bottomup credit analysis every time they assess a counterparty or consider an investment opportunity. Publicly available credit ratings are often provided by external agencies that specialize in credit assessments, but institutions also often have their private credit assessment processes to support publicly available data.

Credit ratings are often represented with a letter-based scale. In many rating systems, the highest possible rating is represented by AAA. It is then followed by AA and A, and then by similar categories for both B and C, and the rating scales often end at D, which represents the default of the entity. Many rating systems may also specify further granularity by assigning pluses or minuses after the letters (e.g., BB+ or CCC−). A high rating (e.g., AA) signifies that the entity is highly creditworthy with a strong ability to meet its financial obligations, whereas a lower rating (e.g., C) signifies that the entity is less reliable and has a higher likelihood of default. (Ng and Mohamed, 2021)

Private agencies that specialize in assigning public credit ratings to various entities are referred to as credit rating agencies (CRAs). They assign these ratings after extensively analyzing an entity's financial stability, its debt structure, and reflecting these on current market conditions. The three most well-known CRAs globally (also referred to as the "Big Three") are Standard & Poor's (S&P Global Ratings), Moody's (Moody's Investors Service, 2025), and Fitch Ratings (Fitch Ratings, 2023), which dominate the rating agency industry with a collective market share of approximately 95% (Ng and Mohamed, 2021). These CRAs share many similarities, but their methodologies and rating scales also have notable differences, which is why their credit rating assessments can result in considerably different credit ratings given for the same entity.

CRAs have played an important role in financial markets because their credit ratings have been widely relied on by many market participants throughout history. However, despite this, CRAs have been criticized for numerous reasons. For one, as the three major CRAs dominate the market by a considerable margin, it can be challenging for competitors to gain recognition. The limited competitiveness aspect can also hinder new innovations within the credit rating industry. Another common point of criticism towards CRAs is the lack of transparency in their specific methodologies, which can make it difficult to cross-check or validate the ratings. (Haspolat, 2015)

There is also an inherent concern with publicly assigned credit ratings that needs to be addressed. CRAs are private companies with the main service of assessing the creditworthiness of other entities and assigning credit ratings to them. This structure can create inherent conflicts of interest because financial incentives may disturb their objectivity. Furthermore, CRAs have historically failed to predict financial crises and caused market instability with overly optimistic ratings and delays in credit rating downgrades. (Haspolat, 2015)

The 2007–2008 global financial crisis highlighted even more flaws in the credit rating system. The "Big Three" significantly underestimated the credit risk of mortgage-backed securities and collateralized debt obligations before mid-2007. Once this was realized, CRAs quickly downgraded many of these securities, which triggered concerns among market participants and ultimately led to rapid and widespread sales in the debt market. This event was later identified as the key driver of the crisis. As an example, structured finance securities, which accounted for 35% of the outstanding U.S. bond market debt in December 2008, had predominantly very high credit ratings, most being rated AAA. However, as market conditions started to deteriorate, a total of over 36 000 tranches were downgraded, including many of these previously AAA-rated securities. These failures were attributed to errors in credit analysis, overreliance on quantitative models, and conflicts of interest within CRAs. (Benmelech and Dlugosz, 2010; Haspolat, 2015; DeHaan, 2017; Ng and Mohamed, 2021)

Another example of the problematic presence of CRAs in financial markets came from the 2010 eurozone sovereign debt crisis. Their credit assessments were a major contributor to shaping the market environment and borrowing costs. Sovereign spreads (measured relative to the German Bund) were influenced by regional risk, country-specific credit risk, and spillover effects from Greece. The generally increased risk-aversion among market participants increased the demand for Bunds, which widened credit spreads even for fiscally stable countries such as Austria, Finland, and the Netherlands (De Santis, 2014). For fiscally weaker GIIPS countries (Greece, Ireland, Italy, Portugal, and Spain), this effect was even stronger. In particular, Greece's credit rating deterioration caused widespread financial distress in many countries. This entire event was further escalated by inconsistencies between credit ratings assigned by the three major agencies; for example, Fitch issued more favorable ratings than Moody's and Standard & Poor's during this time. (Altdörfer et al., 2019)

2.4 Decision analysis and decision trees in energy sector risk management

Decision analysis spans a wide range of quantitative approaches that can be utilized to assess the risks, benefits, and implications of decision-making under uncertainty using mathematical tools. Keeney (1982) summarized the intuition behind decision analysis nicely in the following way: *"[Decision analysis is] a formalization of common sense for decision problems which are too complex for informal use of common sense."* Decision analysis can be very complicated, especially in large organizations, because the implications of decisions must be considered with care. Furthermore, there is rarely a single objectively superior decision alternative in complex decision problems. For instance, decision problems may have many sources of uncertainty (e.g., market conditions or regulatory changes), decision-makers and shareholders may have vastly different risk preferences, and there may be other constraints (e.g., financial, operational, reputational, or legal) at play. (Keeney, 1982)

Decision trees (Magee, 1964) are a widely used tool within decision analysis. A decision tree is a graphical representation of the underlying decision problem, where different (often sequential) decisions, sources of uncertainty, and potential decision outcomes can be examined. A decision tree consists of three types of nodes: decision nodes (squares) represent decisions with certain alternatives to choose between, chance nodes (circles) represent uncertain events with different outcomes and probabilities, and terminal nodes (triangles) represent the final outcomes of the tree, after which no decisions or random events occur. Terminal nodes often have a monetary value or another type of utility attached to them, which corresponds to the benefit of ending up at that node. (Rokach and Maimon, 2005).

To construct a decision tree to support decision-making, it is important to first fully understand all possible decisions, sources of uncertainty, and other nuanced intricacies related to the decision problem. The tree typically starts from an initial decision node (e.g., representing whether to proceed with an investment) and then branches out to subsequent decision and chance nodes. All branches in the tree eventually terminate at the terminal nodes, which have some known payoffs. By using these payoffs and the probabilities within the tree, the tree is evaluated backwards from the terminal nodes: expected value calculations are performed at each chance node, and at the decision nodes, the decision alternative with the highest expected monetary value (EMV) or utility is chosen. This procedure is repeated until optimal strategies have been identified for all decision nodes. Figure 1 presents a simple decision tree for an investment opportunity and its optimal solution with respect to

EMV. It is also worth mentioning that while decision trees are versatile quantitative tools in decision analysis, they are also popular machine learning tools and are used in a wide range of machine learning applications (Navada et al., 2011; Somvanshi et al., 2016).

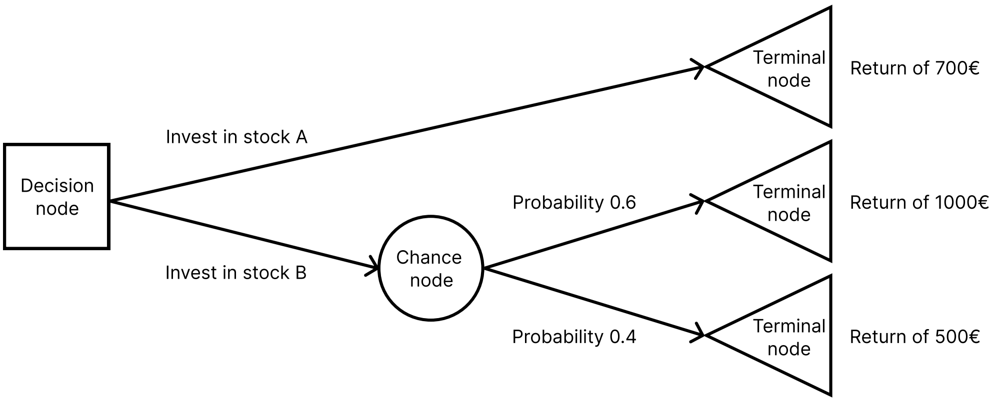

Figure 1: A simple decision tree representing an investment opportunity with one decision node (square), one chance node (circle), and three terminal nodes (triangles), which represent the possible outcomes of the investment. The EMV of investing in stock A is 700€, whereas the EMV of investing in stock B is $0.6 \cdot 1000€ + 0.4 \cdot 500€ = 800€$. Note that stock B is expected to perform better (in terms of EMV) but is also riskier due to a lower worst-case return.

An interesting application using decision trees in the energy sector risk analysis was presented by Mosquera et al. (2008), who studied the medium-term risks that electricity producers face in their operations. They first identified five risk factors: natural gas prices, coal prices, CO_2 emission prices, demand, and hydro conditions. They then used these risk factors to generate a large number of market scenarios using Monte Carlo simulations. The simulation results were then analyzed with the help of decision trees to gain information about the relationships between the output variables (such as profits or electricity prices) and the risk factors. The authors also applied the framework in a case study using data from the Spanish electrical system in 2006, but they emphasized that the case study is just a numerical example and should not be taken as the main takeaway from the paper. Instead, their key contribution was the decision tree framework itself, which can support risk management and hedging decisions from the perspective of an electricity producer.

Another risk management study in the electricity markets came from Kettunen et al. (2009), who presented a multistage optimization framework through which electricity retailers could optimize their hedging portfolio. They constructed a scenario tree using the following risk factors: forward prices and forward premiums, expected future loads, conditional standard deviations and mean reversion parameters of spot price and load, and correlation parameters. To assess the framework, they presented a case study where mean-CCFAR (conditional-cash-flow-at-risk) efficient frontiers were constructed using data from Nord Pool. They demonstrated that different hedging strategies resulted in different balances of expected procurement costs and risk exposure, which could help to choose the most suitable strategy based on a decision-maker's risk attitude. The authors also highlighted that the inclusion of spot price and load correlations in their methodology was important because ignoring these correlations may overlook many extreme scenarios and lead to underestimated risks.

Another approach that applies scenario trees was presented by Lorca and Prina (2014). They concentrated on locational electricity prices and their stochastic evolution over time in the context of portfolio optimization for a power producer. They constructed a scenario tree based on a time series model of historical electricity prices. This tree was then used in a stochastic optimization framework to optimize the power producer's portfolio structure with respect to a balance of expected profits and risk. The authors highlighted that it is very important to construct such portfolios with as many low-correlated elements as possible (that is, diversify the portfolio) because this considerably reduces its overall risk.

Rocha and Kuhn (2012) presented a similar application, in which they studied portfolio optimization of an electricity retailer to minimize market risk. They formulated the problem as a multistage mean-variance optimization model to better account for the unpredictable nature of electricity demand and prices. They also approximated this optimization model using linear decision rules (LDRs) to reduce computational complexity. The authors presented a case study in which they constructed mean-variance efficient frontiers and subsequently analyzed the sensitivity of different risk factors. Based on the results, the LDR approach was a viable alternative to previous scenario tree methods, and the model's adaptivity was especially beneficial in settings with high spot price volatility.

3 Problem context

This thesis has been written for Fortum Oyj, a large Finnish energy company that produces electricity and trades both physical electricity and financially settled electricity derivatives in the Nordic electricity markets. The focus of this thesis is on assessing the counterparty credit risk that comes with signing a large bilateral PPA, which serves as a hedge for Fortum. The goal is to evaluate the expected earnings coming from the bilateral PPA under market and counterparty credit risk, and analyze whether the earnings are sufficient from a risk management perspective, specifically from Fortum's point of view, when compared to alternative hedging strategies.

The setting of the thesis is as follows. Fortum is planning to generate and sell a volume V_g of electricity during a given time period. The entire volume V_g could be sold in the spot market, but as discussed in the literature review in Section 2, this would expose Fortum to high market risk due to the volatile nature of electricity. As a risk mitigation procedure, Fortum hedges a portion $V_h < V_g$ of the total production through bilateral contracts with a large number of different counterparties.

This study considers a situation where a counterparty has reached out to Fortum with a contract proposal c, which would hedge a portion $V_c < V_h$ of the total hedging target (the specific contractual terms are described in detail in the following sections). Especially if V_c is a significant portion of the total hedging target, Fortum is faced with a non-trivial decision regarding the proposed contract c: whether to *decline* it and instead choose to hedge V_c through a diversified portfolio of smaller contracts with various counterparties (later referred to as the *portfolio*), or *accept* it and the counterparty credit risk that comes with it. This decision problem is modeled as a decision tree.

Before diving into the exact methodology of the decision tree model, we first discuss the types of risks related to the proposed contract c. The first obvious source of risk in such bilateral agreements is counterparty credit risk, which arises when either party is unable to fulfill their financial obligations, for example, in the event of a default. The market risk related to the uncertain spot price of electricity also plays a crucial role. Some of the contracts in the portfolio are scheduled to be signed later than the proposed contract c, which means that their exact contract prices are uncertain at that point. Furthermore, if the contracting counterparty defaults, some of the volume V_c will be sold at spot price in the remaining months of the contract period. Any other types of risks, such as ones related to the liquid reserves of either party, are not accounted for in this thesis.

This thesis also considers two types of counterparty credit risk mitigation strategies that Fortum may employ during the negotiation phase of the PPA. The first method is a collateral requirement. The collateral requirement is a fixed amount of money that either party may require from the other for the duration of the contract as a security for their perceived counterparty credit risk. If either party defaults, this collateral can be used to cover the other party's losses. The creditworthiness of the counterparty plays a key role in sizing the collateral requirement: a counterparty with high creditworthiness may have higher thresholds before collateral is required, whereas counterparties with lower creditworthiness may have stricter limits. The second method is a CVA add-on, which acts as a price adjustment to the proposed contract price. Fortum could effectively charge a higher electricity price if it estimates its exposure to counterparty credit risk as too high.

Fortum may also be interested in analyzing alternative situations before deciding whether to accept the proposed contract c and with what terms. This additional analysis can help Fortum understand the true value of its product (electricity) better, and allow it to price the contract accordingly. Numerous alternative situations could prove insightful if analyzed. For example, how would the expected earnings differ if the contracting counterparty had a higher credit rating? Or what if Fortum would hedge only a portion of the volume V_c with this counterparty and the remaining portion with another counterparty or through the portfolio? Additionally, if Fortum deems the counterparty credit risk associated with signing the proposed contract c as too high, how can this risk be mitigated or managed to reach an acceptable risk level?

To address these questions and study different hedging strategies, the contractual earnings from each strategy are estimated probabilistically. Based on these earnings estimates, a collateral requirement and a CVA add-on may be included in the contract and optimized (from Fortum's perspective) so that the downside earnings from accepting the contract c increase to an acceptable risk level, making the contract an attractive decision alternative. Finally, a sensitivity analysis on the earnings confidence level is performed to see how changes in risk attitude affect the optimization. Here, the confidence level refers specifically to the width of the downside tail of the distributions when calculating downside earnings. In later sections, the decision-maker (Fortum) is referred to as the *power producer*, and the probabilistic decision tree model is referred to as the *model*.

3.1 Parameters

The following *contractual parameters* describe the proposed contract c:

V_c	The total volume of electricity (MWh) in the proposed contract c.
$m_{\text{start}}, m_{\text{end}}$	The start and end months of the proposed contract c. The contractual payments occur once a month (see Section 3.3), meaning that the contract horizon can be represented at a monthly granularity.
P_c	The sales price of electricity (€/MWh) in the proposed contract c. The price is assumed to be constant throughout the contract horizon.
R_c	The credit rating of the contracting counterparty (determined internally within Fortum), which is represented as an integer from 0 (lowest credit quality) to 5 (highest credit quality). The credit rating only reflects the estimated creditworthiness of the counterparty, and not, for instance, its size or equity. It is assumed that the credit rating of the counterparty does not change, and Fortum's methodology for determining these ratings is not discussed in this thesis.

Unlike the fixed contractual parameters described above, the collateral requirement and the CVA add-on are decision variables in the problem because they are not part of the initial proposed contract c. Rather, they are additional requirements that the power producer may choose to include in the contractual terms during the negotiation phase, and will be part of the legal contract once both parties have agreed on the terms. We define:

$\mathcal{C}_c$	The collateral requirement (€) issued to the contracting counterparty to secure the proposed contract c.
CVA_c	The CVA add-on (€/MWh) that will be added to the sales price P_c, such that the final contractual price will be $P_c + \text{CVA}_c$. The CVA add-on acts as extra compensation for the power producer's perceived counterparty credit risk.

In addition to the contractual parameters and decision variables described above, we define $m_{\text{sign}} \leq m_{\text{start}}$ as the month in which the definitive contract is signed by both parties (i.e., the time of decision) and the legal contractual obligations go into effect. The month m_{sign} is also the first month in the *contract lifetime*, which we define as the set of months $M_{\text{lifetime}} = \{m_{\text{sign}}, \ldots, m_{\text{end}}\}$. The contract lifetime covers the entire time horizon that is relevant from the legal agreement's perspective. Additionally, we define the *delivery period* as the set of months $M_{\text{delivery}} = \{m_{\text{start}}, \ldots, m_{\text{end}}\}$. This period includes all months during which electricity is sold and delivered.

3.2 Collateral requirement and CVA add-on impacts

The collateral requirement and the CVA add-on are decision variables from the power producer's perspective. Both of these are additional financial requirements the power producer may add to the contractual terms if it estimates its counterparty credit risk as too high without them. By including either one or both of these in the final contractual terms, the power producer can increase its expected earnings to a level where the counterparty credit risk is acceptable (subject to some reference level). At first glance, optimizing these decision variables may seem trivial because one can add arbitrarily large additional requirements to the contractual terms and effectively increase the expected earnings without any limit. However, the power producer must naturally offer competitive contractual terms, or otherwise any reasonable customer would instead negotiate a similar contract with a competitor. Therefore, a good starting point for the power producer is to have some reference downside risk levels in mind and size the collateral requirement and the CVA add-on according to these levels.

The collateral requirement is a fixed amount of money (in reality, it may also be some other assets) that the counterparty places in a secure bank account for the duration of the contract when the contract is signed (i.e., in m_{sign}). If the contract is completed without a premature termination, the counterparty will receive the full amount after the delivery period has ended. However, if the counterparty defaults during the delivery period, some or all of the collateral is transferred to the power producer as compensation for the prematurely terminated agreement. Conversely, the CVA add-on is simply a fixed price increase that is added to the proposed contract price P_c. The precise ways in which these decision variables affect the earnings are described in Section 3.3.

In reality, it is possible that the counterparty would also post a collateral requirement to the power producer because they are naturally also affected by counterparty credit risk. However, the financial implications resulting from such collateral requirements for the power producer are not modeled in this thesis. Furthermore, the model does not account for the default risk of the power producer in any way, and hence it is assumed not to default.

In the final signed contract, the collateral requirement will be visible and function in the way described above. However, the same does not hold for the CVA add-on: in reality, it corresponds to the difference between the final agreed price and the proposed price P_c, and is not viewed as its own pricing component. In this thesis, the proposed contract and the negotiation reactions to it are viewed from the power producer's perspective, which is why the proposed price P_c and the CVA add-on CVA_c are kept as separate pricing components for the sake of clarity.

In practice, even without initially requesting collateral, the power producer may receive some of the owed money from the counterparty after a default if the counterparty later becomes able (and willing) to pay some of its remaining debt. This phenomenon is referred to as the recovery rate. However, in this thesis, the recovery rate is assumed to be zero, and therefore, the power producer will not receive any delayed debt payments after a counterparty default.

Additionally, it is assumed that the collateral retains its full value and does not gain interest when held in the bank account. Also, unlike in reality, it is assumed that a collateral requirement can be imposed only once and is secured simultaneously, and that this collateral requirement occurs in m_{sign}.

3.3 Cash flow analysis of the underlying process

Before formulating the decision tree model, we describe all cash flow sources that can occur during the delivery period. First of all, any electricity production costs or other recurring outflows are not considered in the contractual scope because the focus is on earnings coming from the sales of the contractual volume V_c. Furthermore, the earnings are measured as EBITDA (earnings before interest, taxes, depreciation, and amortization), and all cash flows occur at a monthly granularity. We introduce the following additional notation:

H_m — The number of hours in month m, which accounts for the impacts of possible leap years and daylight saving hours.

v_m — The sales volume (MWh) allocated to month $m \in M_{\text{delivery}}$. The monthly volumes are calculated as $v_m = V_c \cdot \frac{H_m}{\sum_{m' \in M_{\text{delivery}}} H_{m'}}$, meaning that the monthly volumes are directly proportional to the number of hours the corresponding month has.

$P^{\text{spot}}_{h,m}$ — The spot price of electricity (€/MWh) during hour h in month $m \in M_{\text{lifetime}}$.

m_{default} The default month of the counterparty, if a default occurs.

$\text{MtM}_{m_{\text{default}}}$ The mark-to-market (MtM) valuation (€) of the remaining part of the proposed contract c in the month when the counterparty defaults, assuming that the proposed contract c has been accepted and signed.

$\mathcal{K}$ The set of contracts $k \in \mathcal{K}$ in the portfolio. It is assumed that there is no default risk within the portfolio. In reality, the structure of the portfolio evolves as often as every month and is not tied in any way to the proposed contract c. In this thesis, the portfolio refers to such a subset of the real-world portfolio whose total volume equals V_c, so that it can be used as a comparable hedging strategy when declining the proposed contract c.

$P_{k,m}$ The sales price of electricity (€/MWh) in a contract $k \in \mathcal{K}$ in month m. Because some contracts in the portfolio are scheduled to be signed in a later month than m_{sign}, their contract prices are uncertain at m_{sign} as they depend on the future market prices in their respective signing months. Also, just like with the proposed contract c, these prices are assumed to be constant for each contract k.

$v_{k,m}$ The volume of electricity in a contract $k \in \mathcal{K}$ in month m. As mentioned, the total volume of all contracts in the portfolio during M_{delivery} equals V_c.

In any given month during the delivery period M_{delivery}, the cash flow can originate from the following sources:

(i) From the portfolio, when the proposed contract c is declined. In the portfolio, some of the contracts are signed in m_{sign} and others are scheduled to be signed in a later month. In this case, the cash flow in month m is given by

$$CF_m^{\text{portfolio}} = \sum_{k \in \mathcal{K}} P_{k,m} \cdot v_{k,m}. \tag{1}$$

(ii) From the contractual terms, when the proposed contract c is accepted and the contracting counterparty has not defaulted. In this case, the cash flow in month m is given by

$$CF_m^c = (P_c + \text{CVA}_c) \cdot v_m. \tag{2}$$

(iii) From a payment to the contracting counterparty (outflow) or from a collateral settlement (inflow) in the (possible) month when the contracting counterparty defaults m_{default} and the contract is therefore terminated prematurely. Whether this cash flow is an outflow or an inflow depends on the sign of the MtM valuation of the contract at m_{default}. If the MtM valuation is negative, the power producer is required to make a payment to the defaulting counterparty, which is equal to the MtM valuation. However, if the MtM valuation is positive, the power producer receives a collateral settlement instead, and its size is equal to whichever is smaller: the MtM valuation or the posted collateral. Therefore, if the power producer did not include a collateral requirement in the contractual terms, no money will be received even if the MtM valuation is positive.

In reality, the power producer and the counterparty may not agree on the exact MtM valuation of the remaining contract at the time of a default event because there may be a lack of available market data far enough in the future. Any uncertainty related to this is not accounted for in this thesis, and the MtM valuations used in the model are estimates by the power producer. Additionally, both parties may agree upfront not to make any MtM valuation-related payments to the other party if either were to default. However, this thesis assumes that these payments are made according to the estimated MtM valuation.

In this case, the cash flow in the default month m_{default} is given by

$$CF^{\text{term}}_{m_{\text{default}}} = \begin{cases} \text{MtM}_{m_{\text{default}}}, & \text{if } \text{MtM}_{m_{\text{default}}} < 0 \\ \min\{\text{MtM}_{m_{\text{default}}}, \mathcal{C}_c\}, & \text{if } \text{MtM}_{m_{\text{default}}} \geq 0 \end{cases} \tag{3}$$

(iv) From the spot market, in the default month m_{default} and the remaining months of the contract. The spot price changes hourly in the spot market, and the monthly volume v_m is assumed to be evenly distributed across the hours in that month. In this case, the cash flow in month m is given by

$$CF^{\text{spot}}_m = \sum_{h=1}^{H_m} \frac{v_m}{H_m} \cdot P^{\text{spot}}_{h,m}. \tag{4}$$

Given the monthly cash flows in Equations 1–4, we can describe what the contractual earnings over the entire contract lifetime would be in practice. Related to this, we define *total earnings* as the sum of monthly earnings over the delivery period. Because the total earnings differ depending on whether the contract is declined or accepted and whether the counterparty defaults or not, we can categorize possible total earnings over the contract lifetime into three different categories: i) the contract is declined, ii) the contract is accepted, the counterparty does not default and hence the contract is fully completed, and iii) the contract is accepted, but the counterparty defaults during the contract lifetime. The total earnings in each of the three cases can be expressed using Equations 1–4 in the following way:

$$e^{\text{decline}} = \sum_{m \in M_{\text{delivery}}} CF_m^{\text{portfolio}}, \tag{5}$$

$$e^{\text{accept, no default}} = \sum_{m \in M_{\text{delivery}}} CF_m^c, \tag{6}$$

$$e^{\text{accept, default}} = \sum_{m=m_{\text{start}}}^{m_{\text{default}-1}} CF_m^c + CF_{m_{\text{default}}}^{\text{term}} + \sum_{m=m_{\text{default}}}^{m_{\text{end}}} CF_m^{\text{spot}} \tag{7}$$

3.4 Model formulation

In this section, we present the decision tree model in its entirety. The graphical representation of the decision tree can be seen in Figure 2; however, note that the specific notation used in Figure 2 has not been defined yet and is instead introduced in Sections 3.4.1–3.4.3.

The decision tree model has three top-level decision branches, all of which stem from the initial decision node (see Figure 2). The two *main branches*, which we refer to as the *decline* branch and the *accept* branch, correspond to decision alternatives regarding the proposed contract c that the power producer can choose between. On the other hand, the third top-level branch, the *references* branch, represents a collection of n reference alternative branches, denoted $\textit{reference}_1, \ldots, \textit{reference}_n$. Each of these has an associated *reference alternative contract*, denoted $c'_1, \ldots, c'_n$, that can be used as a benchmark to compare against the *accept* branch total earnings. The reference alternative contracts c'_n may differ in terms of contractual parameters from the proposed contract c, so we define the total volume of electricity, the fixed price of electricity, and the credit rating of the contracting counterparty for reference alternative contracts c'_n as $V_{c'_n}$, $P_{c'_n}$, and $R_{c'_n}$, respectively.

In the model, the total earnings in each reference alternative branch are evaluated similarly to the total earnings in the *decline* and *accept* branches. However, because these branches represent hypothetical references for the *accept* branch total earnings, they are not actual decision alternatives that the power producer can choose between. Instead, they provide additional information on how the expected earnings may change under different assumptions, for example, regarding a counterparty with a higher credit rating or a different delivery period than the one with the proposed contract c. In this way, the power producer can analyze the impacts of accepting the contract c from different angles before making any definitive decisions.

Because the *references* branch is utilized for additional comparisons, the number of *reference*$_n$ branches may vary between practical PPA case studies, depending on what the power producer considers valuable to compare against at any given time. The way in which the total earnings distributions in each branch are obtained is presented in Section 3.4.3. However, before that, the precise methodology by which monthly earnings are simulated in the model needs to be introduced, which is done in Sections 3.4.1 and 3.4.2.

3.4.1 Scenario modeling: spot prices and counterparty defaults

The model considers a set of 135 independent *spot price scenarios*, denoted as

$$S = \{s_j \mid j = 1, \ldots, 135\}, \tag{8}$$

where each scenario s_j consists of a sequence of monthly spot price estimates. These scenarios are defined exogenously to the model and are used as common market price estimates throughout the decision tree model, and are thus not tied to any particular contract or decision tree branch. Unlike the real spot market, which operates at an hourly granularity as discussed in Section 3.3, these spot price scenarios are instead defined at a monthly granularity and represent average monthly spot price estimates. It is assumed that all spot price scenarios are equally likely to be realized and hence their probabilities are given by $p_{s_j} = \frac{1}{135}$ for all j. The methodology for determining these spot price scenarios is beyond the scope of this thesis.

The model also accounts for the possibility that a contracting counterparty (either in the *accept* or *reference*$_n$ branches) may default during the contract lifetime. In principle, the time bucketing is not tied to the granularity of the cash flows, but these default events are modeled at the same monthly granularity for maximum precision. However, despite the aligned granularity, these are treated differently in the model, and hence we define a different index set for the default events.

Because any considered counterparty may not default during M_{lifetime}, we represent this special "no default" default event with a special index δ_0. We can now define the set of possible default events as

$$\Delta = \{\delta_0, \delta_{\text{sign}}, \ldots, \delta_{\text{end}}\}, \tag{9}$$

where δ_0 represents no default of the counterparty within M_{lifetime}, and all other non-zero indices δ_i represent a default in month $m_i \in M_{\text{lifetime}}$.

To properly represent the default events based on this underlying structure, we define the set of *default scenarios* as

$$D = \{d_\delta \mid \delta \in \Delta\}, \tag{10}$$

where d_δ represents either no default within the contract lifetime (when $\delta = \delta_0$), or a default event within the contract lifetime (when $\delta \in \{\delta_{\text{sign}}, \ldots, \delta_{\text{end}}\}$). Therefore, the set D contains exactly one default scenario for each month in the contract lifetime and an additional "no default" default scenario.

These default scenarios are relevant in the *accept* and *reference$_n$* branches as a contract (c or c'_n) is signed in these, and hence the default risk of the counterparty needs to be accounted for. The model uniquely determines the probabilities of each default scenario $d_\delta \in D$ realizing from the credit rating of the contracting counterparty (R_c or $R_{c'_n}$). Therefore, we define

$$\begin{aligned} &p_{d_\delta}(R_c), && p_{d_\delta}(R_{c'_n}), \quad \forall d_\delta \in D, \text{ such that} \\ &\sum_{\delta\in\Delta} p_{d_\delta}(R_c) = 1, && \sum_{\delta\in\Delta} p_{d_\delta}(R_{c'_n}) = 1, \quad \forall n. \end{aligned} \tag{11}$$

The methodology for determining these default scenario probabilities p_{d_δ} is outside the scope of this thesis.

The default scenarios $d_\delta \in D$ are not particularly useful by themselves, because in practice, the financial implications of defaults are largely determined by the underlying market conditions at that time. The model assumes that the realizations of spot price and counterparty default scenarios are independent random events. To integrate the default scenarios with underlying market estimations and capture a wider range of potential realized market outcomes when a contract is signed, the spot price and counterparty default scenarios are combined to form a larger set of *combined scenarios*:

$$\Omega = S \times D = \{\omega_{j,\delta} = (s_j, d_\delta) \mid s_j \in S, d_\delta \in D\}. \tag{12}$$

Because of the independence assumption of the spot price and counterparty default scenarios, the probability of a combined scenario $\omega_{j,\delta}$ is the product of the probabilities p_{s_j} and p_{d_δ}:

$$\begin{aligned} &p_{\omega_{j,\delta}}(R_c) = p_{s_j} \cdot p_{d_\delta}(R_c), \quad p_{\omega_{j,\delta}}(R_{c'_n}) = p_{s_j} \cdot p_{d_\delta}(R_{c'_n}), \quad \forall s_j \in S,\ d_\delta \in D, \text{ such that} \\ &\sum_{\omega_{j,\delta} \in \Omega} p_{\omega_{j,\delta}}(R_c) = 1, \qquad\quad \sum_{\omega_{j,\delta} \in \Omega} p_{\omega_{j,\delta}}(R_{c'_n}) = 1, \quad \forall n. \end{aligned} \tag{13}$$

This set of combined scenarios can be used to estimate the total earnings in branches where a contract is signed statistically more robustly due to a larger total number of simulated earnings outcomes.

In later earnings calculations in Section 3.4.2, certain subsets of the default month set Δ, the default scenario set D, and the combined scenario set Ω will be used. Therefore, we define the following:

- The subsets that correspond to no counterparty defaults within the contract lifetime:

$$\Delta_{\text{no default}} := \{\delta_0\}, \quad D_{\text{no default}} := \{d_{\delta_0}\}, \quad \Omega_{\text{no default}} := \{(s_j, d_{\delta_0}) \mid s_j \in S\}. \tag{14}$$

- The subsets that correspond to a contracting counterparty defaulting before the delivery period:

$$\begin{aligned} &\Delta_{\text{before}} := \{\delta \in \Delta \mid \delta \in \delta_{\text{sign}}, \ldots, \delta_{\text{start-1}}\}, \\ &D_{\text{before}} := \{d_\delta \mid \delta \in \Delta_{\text{before}}\}, \quad \Omega_{\text{before}} := \{(s_j, d_\delta) \mid s_j \in S, d_\delta \in D_{\text{before}}\}. \end{aligned} \tag{15}$$

- The subsets that correspond to a contracting counterparty defaulting during the delivery period:

$$\begin{aligned} &\Delta_{\text{during}} := \{\delta \in \Delta \mid \delta \in \delta_{\text{start}}, \ldots, \delta_{\text{end}}\}, \\ &D_{\text{during}} := \{d_\delta \mid \delta \in \Delta_{\text{during}}\}, \quad \Omega_{\text{during}} := \{(s_j, d_\delta) \mid s_j \in S, d_\delta \in D_{\text{during}}\}. \end{aligned} \tag{16}$$

3.4.2 Monthly cash flows in the model

The spot price scenarios s_j and the combined scenarios $\omega_{j,\delta}$ specify the underlying simulated uncertain market environment, based on which the total earnings for the entire delivery period for different market outcomes can be estimated. The model calculates the cash flows in different situations that the power producer may face in a very similar way to how they are determined in reality (see Section 3.3).

The monthly *earnings* in any given branch, referred to as $e^{branch}_{s_j,m}$ or $e^{branch}_{\omega_{j,\delta},m}$ in month $m \in M_{\text{delivery}}$, are determined from the spot price scenarios s_j or the combined scenarios $\omega_{j,\delta}$, the fixed contractual parameters V_c, P_c, and R_c ($V_{c'_n}$, $P_{c'_n}$, and $R_{c'_n}$), as well as the collateral requirement and the CVA add-on, which act as decision variables. As in the real-world case described in Section 3.3, the monthly earnings in the model can originate from four different sources: from the portfolio when the proposed contract c is declined, from contractual earnings when the contract is accepted and no default has occurred, from a payment or collateral settlement in the month when a default occurs, or from the spot market when the contracting counterparty has defaulted. The main difference from the real-world case is that the spot market earnings in the model are not determined by hourly prices. Rather, they are calculated based on monthly average spot price estimates.

To adapt Equations 1–4 on cash flows to the model's structure, we first need to define $P^{\text{spot}}_{s_j,m}$ as the spot price of electricity in spot price scenario s_j and month m. Now, the model's possible sources of earnings in any month $m \in M_{\text{delivery}}$ can be expressed as:

$$e^{branch,\text{portfolio}}_{s_j,m} = \sum_{k \in \mathcal{K}} P_{k,m} \cdot v_{k,m}, \qquad \forall j, m \in M_{\text{delivery}}, \tag{17}$$

$$e^{branch,c}_{\omega_{j,\delta},m} = (P_c + \text{CVA}_c) \cdot v_m \quad \text{or} \quad e^{branch,c'_n}_{\omega_{j,\delta},m} = P_{c'_n} \cdot v_m, \quad \forall j, m \in M_{\text{delivery}}, \tag{18}$$

$$e^{branch,\text{term}}_{\omega_{j,\delta},m_{\text{default}}} = \begin{cases} \text{MtM}_{m_{\text{default}}}, & \text{if } \text{MtM}_{m_{\text{default}}} < 0 \\ \min\{\text{MtM}_{m_{\text{default}}}, \mathcal{C}_c\}, & \text{if } \text{MtM}_{m_{\text{default}}} \geq 0 \end{cases} \quad \forall j, \tag{19}$$

$$e^{branch,\text{spot}}_{\omega_{j,\delta},m} = P^{\text{spot}}_{s_j,m} \cdot v_m, \qquad \forall j, m \in M_{\text{delivery}}. \tag{20}$$

To formulate the total contractual earnings in these different situations, some assumptions related to the earnings need to be addressed first. The model estimates contractual earnings in nominal terms based on available information in the signing month m_{sign}. This includes the uncertainty factors of the model, which are the future spot price estimates and the credit rating (R_c or $R_{c'_n}$) implied default scenario probabilities. Because this thesis focuses specifically on the counterparty credit risk aspects of the proposed contract c, the model does not discount future earnings. This assumption affects both estimated contractual earnings and optimization of the collateral requirement $\mathcal{C}_c$ and the CVA add-on CVA_c, as there are no time preferences regarding the earnings in the model.

3.4.3 Earnings distributions and risk measures

The decision tree model estimates the total earnings in each decision branch and compares the earnings from accepting the contract to those from declining it, and to other hypothetical but realistic reference hedging strategies. Specifically, because the power producer is risk-averse, the downside earnings at a certain confidence level (i.e., how much of the downside tail is included in the downside earnings calculations) is the focus of this comparison. This reflects the goal of avoiding adverse outcomes over potentially higher but also riskier returns. The model also supports analyzing the effects of the collateral requirement and the CVA add-on to the *accept* branch earnings.

To calculate the total earnings in each branch, we define finite branch-specific sets of total earnings and probabilities as $\mathcal{E}_{scens}^{branch}$, where *branch* refers either to the *decline*, *accept*, or any of the $reference_n$ branches, and *scens* refers to the set of spot price scenarios S or the set of combined scenarios Ω, whichever is the source of uncertainty in the given branch. Each of these sets consists of (total earnings, probability) pairs $(e_{scen}^{branch}, p_{scen}^{branch})$, where *scen* refers to a single spot price scenario s_j or a combined scenario $\omega_{j,\delta}$, depending on the branch.

These pairs $(e_{scen}^{branch}, p_{scen}^{branch}) \in \mathcal{E}_{scens}^{branch}$ define branch-specific discrete probability distributions. Here $e_{scen}^{branch} \in \mathbb{R}$ is the total earnings resulting from the underlying spot price scenario or combined scenario, and p_{scen}^{branch} is its corresponding probability, such that $\sum_{scen} p_{scen}^{branch} = 1$ holds. Therefore, the total number of (total earnings, probability) pairs in any given branch depends on its underlying structure.

To measure and compare the downside earnings of different branches, we first need to introduce two risk measures used in the model: Value at Risk (VaR) and Conditional Value at Risk (CVaR). Because the total earnings in any given branch form a discrete distribution, we define VaR at a confidence level α as

$$\mathrm{VaR}_{\alpha}(\mathcal{E}_{scens}^{branch}) := \inf \left\{ x \in \mathbb{R} \;\middle|\; \sum_{scen: e_{scen}^{branch} \leq x} p_{scen}^{branch} \geq \alpha \right\}. \tag{21}$$

Hence, VaR_{α} corresponds to the smallest total earnings level such that the sum of probabilities of total earnings less than or equal to it is at least α. In other words, VaR_{α} represents a total earnings threshold below which the lowest α-quantile of total earnings lies.

Using the definition for VaR_α, CVaR at a confidence level α can be defined as

$$\mathrm{CVaR}_\alpha(\mathcal{E}_{scens}^{branch}) := \frac{1}{\alpha} \sum_{scen: e_{scen}^{branch} \leq \mathrm{VaR}_\alpha(\mathcal{E}_{scens}^{branch})} p_{scen}^{branch} \cdot e_{scen}^{branch}. \tag{22}$$

The CVaR_α is therefore the weighted average of the lowest α-quantile total earnings.

In this thesis, the risk measure used to compare the total earnings across various branches is CVaR_α at varying confidence levels. We refer to these downside earnings as *$CVaR_\alpha$-earnings*, denoted $\mathrm{CVaR}_\alpha^{decline}$, $\mathrm{CVaR}_\alpha^{accept}$, and $\mathrm{CVaR}_\alpha^{reference_n}$ for the respective branches.

Next, we present the methodology by which the sets $\mathcal{E}_{scens}^{branch}$ are formed and how the corresponding CVaR_α-earnings are calculated based on these sets.

(*Decline*) The first decision branch considers declining the proposed contract c. In this case, the full volume V_c is hedged through the portfolio during the delivery period, with uncertainty only about future spot prices in the form of the spot price scenarios. Therefore, we define the set of (total earnings, probability) pairs in the *decline* branch as

$$\begin{aligned} \mathcal{E}_S^{decline} &= \left\{ \left(e_{s_j}^{decline}, p_{s_j}^{decline} \right) \,\middle|\, s_j \in S \right\}, \quad \text{where} \\ e_{s_j}^{decline} &:= \sum_{m \in M_{\text{delivery}}} e_{s_j,m}^{decline,\text{portfolio}}, \quad p_{s_j}^{decline} := p_{s_j}. \end{aligned} \tag{23}$$

Using this set, $\mathrm{CVaR}_\alpha^{decline}$ can be calculated as

$$\mathrm{CVaR}_\alpha^{decline} = \mathrm{CVaR}_\alpha(\mathcal{E}_S^{decline}). \tag{24}$$

(*Accept*) The second decision branch considers accepting and signing the proposed contract c. In this case, the same underlying uncertainty about future spot prices is present. However, there is also uncertainty about whether the contracting counterparty defaults before the delivery period starts, during the delivery period, or does not default at all during the contract lifetime. The earnings in these different cases are computed as follows:

(i) For combined scenarios in which the default occurs before the delivery period (Ω_{before}), each monthly volume v_m is assumed to be sold at spot price in each month of the delivery period. This gives the following set:

$$\begin{aligned}
\mathcal{E}^{accept}_{\Omega_{\text{before}}} &= \left\{\left(e^{accept}_{\omega_{j,\delta_{\text{before}}}},\ p^{accept}_{\omega_{j,\delta_{\text{before}}}}\right) \,\middle|\, \omega_{j,\delta_{\text{before}}} \in \Omega_{\text{before}}\right\},\ \text{where} \\
e^{accept}_{\omega_{j,\delta_{\text{before}}}} &:= \sum_{m\in M_{\text{delivery}}} e^{accept,\text{spot}}_{\omega_{j,\delta_{\text{before}}},m}, \quad p^{accept}_{\omega_{j,\delta_{\text{before}}}} := p_{\omega_{j,\delta}}(R_c), \quad \forall\delta \in \Delta_{\text{before}}.
\end{aligned} \tag{25}$$

(ii) For combined scenarios in which the default occurs during the delivery period (Ω_{during}), the monthly volumes will be sold at the contracted price in months before default, and at spot price in the default month and months after default. Furthermore, a payment (outflow) or a collateral settlement (inflow) will occur in the default month, the size of which depends on the initial collateral requirement and the MtM valuation of the contract at that time. Let $m^{\delta_{\text{during}}}$ be the default month of any given combined scenario $\omega_{j,\delta_{\text{during}}} \in \Omega_{\text{during}}$. This gives the following set:

$$\begin{aligned}
\mathcal{E}^{accept}_{\Omega_{\text{during}}} &= \left\{\left(e^{accept}_{\omega_{j,\delta_{\text{during}}}}, p^{accept}_{\omega_{j,\delta_{\text{during}}}}\right) \,\middle|\, \omega_{j,\delta_{\text{during}}} \in \Omega_{\text{during}}\right\},\ \text{where} \\
e^{accept}_{\omega_{j,\delta_{\text{during}}}} &:= \sum_{m=m_{\text{start}}}^{m^{\delta_{\text{during}}}-1} e^{accept,c}_{\omega_{j,\delta_{\text{during}}},m} + e^{accept,\text{term}}_{\omega_{j,\delta_{\text{during}}},m^{\delta_{\text{during}}}} + \sum_{m=m^{\delta_{\text{during}}}}^{m_{\text{end}}} e^{accept,\text{spot}}_{\omega_{j,\delta_{\text{during}}},m}, \\
&\qquad \forall m^{\delta_{\text{during}}} \in M_{\text{delivery}}, \\
p^{accept}_{\omega_{j,\delta_{\text{during}}}} &:= p_{\omega_{j,\delta}}(R_c), \quad \forall\delta \in \Delta_{\text{during}}.
\end{aligned} \tag{26}$$

(iii) For combined scenarios in which the contracting counterparty does not default during the contract lifetime ($\Omega_{\text{no default}}$), the earnings are determined from the contractual terms for the entire delivery period. This gives the following set:

$$\begin{aligned}
\mathcal{E}^{accept}_{\Omega_{\text{no default}}} &= \left\{\left(e^{accept}_{\omega_{j,\delta_{\text{no default}}}}, p^{accept}_{\omega_{j,\delta_{\text{no default}}}}\right) \,\middle|\, \omega_{j,\delta_{\text{no default}}} \in \Omega_{\text{no default}}\right\},\ \text{where} \\
e^{accept}_{\omega_{j,\delta_{\text{no default}}}} &:= \sum_{m\in M_{\text{delivery}}} e^{accept,c}_{\omega_{j,\delta_{\text{no default}}},m}, \\
p^{accept}_{\omega_{j,\delta_{\text{no default}}}} &:= p_{\omega_{j,\delta}}(R_c), \quad \forall\delta \in \Delta_{\text{no default}}.
\end{aligned} \tag{27}$$

We can now formulate $\mathcal{E}^{accept}_{\Omega}$ for the whole *accept* branch as the union of the sets given in Equations 25–27 as follows:

$$\mathcal{E}_{\Omega}^{accept} = \mathcal{E}_{\Omega_{\text{before}}}^{accept} \cup \mathcal{E}_{\Omega_{\text{during}}}^{accept} \cup \mathcal{E}_{\Omega_{\text{no default}}}^{accept}. \tag{28}$$

Using this full earnings distribution for the *accept* branch, $\mathrm{CVaR}_{\alpha}^{accept}$ can be calculated as:

$$\mathrm{CVaR}_{\alpha}^{accept} = \mathrm{CVaR}_{\alpha}(\mathcal{E}_{\Omega}^{accept}). \tag{29}$$

(*References*) The third *references* branch consists of a set of hypothetical but realistic reference alternatives that are used as benchmarks to compare $\mathrm{CVaR}_{\alpha}^{accept}$ against. This set is case-specific and is constructed to give a better overall representation of risks related to selling the volume V_c, to allow the power producer to study the earnings under different assumptions. Some examples of realistic reference alternatives include the following:

- How would the earnings change if the same contract (volume, price, delivery period) were proposed by a counterparty with either higher or lower creditworthiness, and this contract were accepted and signed?
- How would the earnings change if the delivery period started at a different time?
- How would the earnings change if the delivery period were either shorter or longer?
- How would the earnings change if only a portion of the volume V_c were contracted with the initial counterparty, while the remainder were hedged through the portfolio?

Because the specific reference alternatives are case-specific by nature, and the total number of these can vary, the exact total earnings of sets $\mathcal{E}_{\Omega}^{reference_n}$ cannot be comprehensively generalized. However, each $reference_n$ branch resembles the structure of the *accept* branch because they represent accepting and signing a contract c'_n with some counterparty and account for the same spot price and counterparty default uncertainties. The only differences lie in the contractual parameters of the reference alternative contracts c'_n and in the potentially different hedging split.

Therefore, mathematically, the sets $\mathcal{E}_{\Omega}^{reference_n}$ are structured similarly to $\mathcal{E}_{\Omega}^{accept}$ given in Equations 25–28, but with modified contractual parameters depending on what kind of reference alternatives are being analyzed. Therefore, $\mathrm{CVaR}_{\alpha}^{reference_n}$

can be calculated separately for all reference alternatives as

$$\mathrm{CVaR}_{\alpha}^{reference_n} = \mathrm{CVaR}_{\alpha}(\mathcal{E}_{\Omega}^{reference_n}). \tag{30}$$

The reference alternatives that are used in the case study are introduced in Section 4.

An overview of the full decision tree is shown in Figure 2.

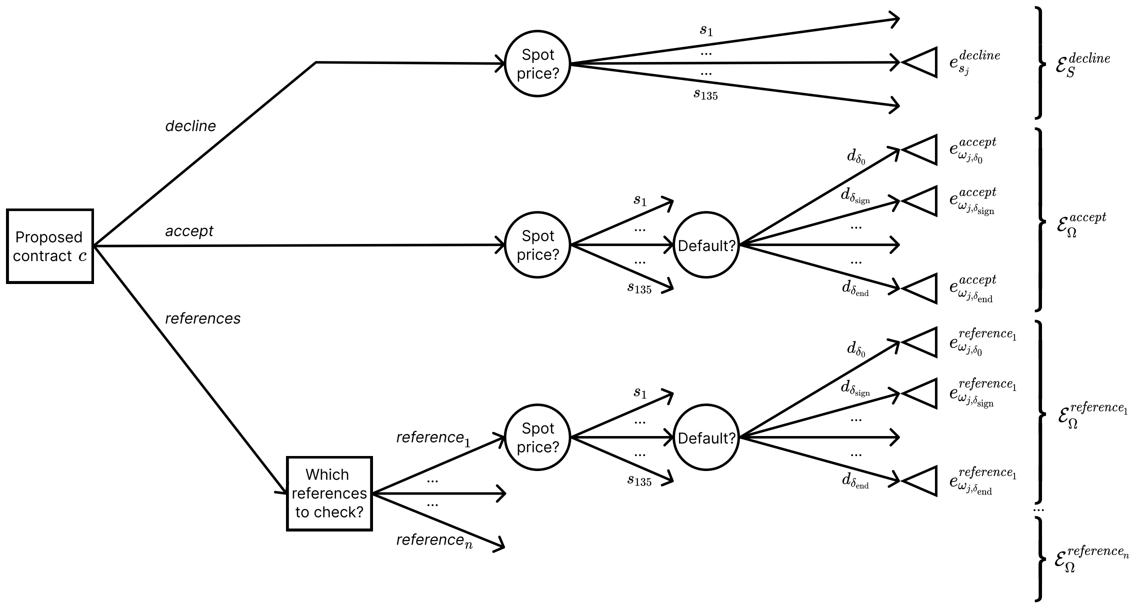

Figure 2: The decision tree for evaluating different decision alternatives regarding a proposed bilateral contract c. The tree consists of three top-level decision branches: *decline*, *accept*, and *references*; the first two represent actual decision alternatives that the power producer may choose between, whereas the third is a collection of n reference alternatives (denoted $\mathit{reference}_1, \ldots, \mathit{reference}_n$) that are used as benchmarks to compare the *accept* branch total earnings against. The total earnings in each branch depend on the contractual parameters V_c, P_c, and R_c ($V_{c'_n}$, $P_{c'_n}$, and $R_{c'_n}$), the delivery period M_{delivery}, and either the spot price scenarios $s_j \in S$ or the combined scenarios $\omega_{j,m} \in \Omega$, depending on the branch.

4 Case study

Next, we present a case study to demonstrate how the decision tree model outlined in Section 3 can support decision-making and negotiations regarding large bilateral hedging contracts, from the perspective of counterparty credit risk management. The case study is structured as follows. First, we present hypothetical but realistic input data for the model in Section 4.1. Using the input data, the decision tree model is evaluated, and the total earnings distributions from each branch are visualized in Section 4.2. Subsequently, in Section 4.3, the collateral requirement and the CVA add-on are optimized with respect to four target levels, and the sensitivity of the used confidence level is studied. Finally, the decision model and the results of the case study are critically assessed in Section 4.4.

4.1 Data

We first provide some context on reported PPA specifications in the Nordics, which will be used as references for the contractual parameters in the case study. Seppälä and Syri (2025) report that in recent years, pay-as-produced PPA prices in Finland and Sweden have ranged around 50 €/MWh, with solar PPAs priced at 48–53 €/MWh and wind PPAs at 50–68 €/MWh. The exact prices depend on the specific region within each country. On the other hand, for baseload PPAs, where the price is tied to current electricity market prices, the prices are estimated to be approximately 10% higher than those of pay-as-produced PPAs. Furthermore, Pexapark states that corporate PPA delivery periods typically range from ten to twenty years.

The contractual parameters for the proposed contract c and the reference alternative contracts c'_n are presented in Table 1. The time of signing the contract is assumed to be m_{sign} = Feb 2025. Therefore, the contract lifetime $M_{\text{lifetime}} = \{\text{Feb 2025}, \ldots, \text{Dec 2039}\}$ consists of 179 months, whereas the delivery period $M_{\text{delivery}} = \{\text{Jan 2030}, \ldots, \text{Dec 2039}\}$ consists of 120 months. The volume V_c corresponds to a continuous delivery of 250 MW throughout the delivery period, which, assuming 8760 hours per year, equals $V_c = 250 \text{ MW} \times 10 \text{ years} \times 8760 \frac{\text{h}}{\text{year}} = 21\,900\,000$ MWh. This volume is then distributed across each month of the delivery period such that the monthly volumes are proportional to the number of hours in each month. The contractual parameters in Table 1 have been selected to resemble realistic PPAs while enabling insightful analyses in later sections.

In this case study, the number of reference alternatives is chosen to be $n = 2$, which results in two reference alternative contracts c'_1 and c'_2. The reference alternative contract c'_1 is otherwise similar to c, with the only difference being a higher credit-rated counterparty. The reference alternative contract c'_2 is also otherwise similar to c, but the total volume is split equally between c'_2 and the portfolio. The volume assigned to the portfolio is referred to as $V_{\text{portfolio}}$.

	$V_c/V_{c'_n}$ (MWh)	$V_{\text{portfolio}}$ (MWh)	m_{start}	m_{end}	$P_c/P_{c'_n}$ (€/MWh)	$R_c/R_{c'_n}$ (0–5)
Proposed contract c	21 900 000	0	Jan 2030	Dec 2039	60	3
Reference alternative contract c'_1	21 900 000	0	Jan 2030	Dec 2039	60	4
Reference alternative contract c'_2	10 950 000	10 950 000	Jan 2030	Dec 2039	60	3

Table 1: Contractual parameters for the proposed contract c and the reference alternative contracts c'_1 and c'_2

As discussed in Section 3.4, the number of spot price scenarios is $|S| = 135$. On the other hand, the number of counterparty default scenarios is $|D| = 180$: one for each month in the contract lifetime, and one for the "no default" case. Therefore, the number of combined scenarios is $|\Omega| = 135 \times 180 = 17\ 550$. We determine the size of each decision branch distribution $\mathcal{E}_{scens}^{branch}$, that is, how many (total earnings, probability) pairs each distribution consists of, as follows:

- *Decline* branch: The only source of uncertainty is the spot price, and hence the size of the distribution $\mathcal{E}_S^{decline}$ is $|S| = 135$.

- *Accept* branch: The sources of uncertainty are the spot price and counterparty default events. The size of the distribution $\mathcal{E}_\Omega^{accept}$ is $|\Omega| = 17\ 550$.

- *Reference*$_1$ branch: Because this branch is otherwise similar to the *accept* branch, with the exception of a higher credit-rated counterparty, the size of the distribution $\mathcal{E}_\Omega^{reference_1}$ is $|\Omega| = 17\ 550$.

- *Reference*$_2$ branch: This branch considers a hypothetical situation, where half of the volume V_c is hedged through the large PPA with the same counterparty as in *accept*, while the remaining half is hedged through the portfolio. Therefore, the earnings from both sources are joined by the spot price scenario, and hence the size of the distribution $\mathcal{E}_\Omega^{reference_2}$ is also $|\Omega| = 17\ 550$.

Finally, the confidence levels used in the CVaR$_\alpha$-earnings calculations in Section 4.3 are $\alpha \in \{5\%, 10\%, 15\%, 20\%\}$.

4.2 Evaluating the decision tree

The full decision tree can now be evaluated using the input data presented in Section 4.1. To compare the total earnings in each decision branch, the sets $\mathcal{E}_S^{decline}$, $\mathcal{E}_\Omega^{accept}$, $\mathcal{E}_\Omega^{reference_1}$, and $\mathcal{E}_\Omega^{reference_2}$ are constructed using the methodology described in Section 3.4, and the (total earnings, probability) pairs for each set are plotted as cumulative distributions with cumulative probability on the horizontal axis and total earnings on the vertical axis. These distributions, along with a *default-free hedge* level representing a situation where all earnings uncertainty is eliminated (i.e., a counterparty that cannot default), are shown in Figure 3.

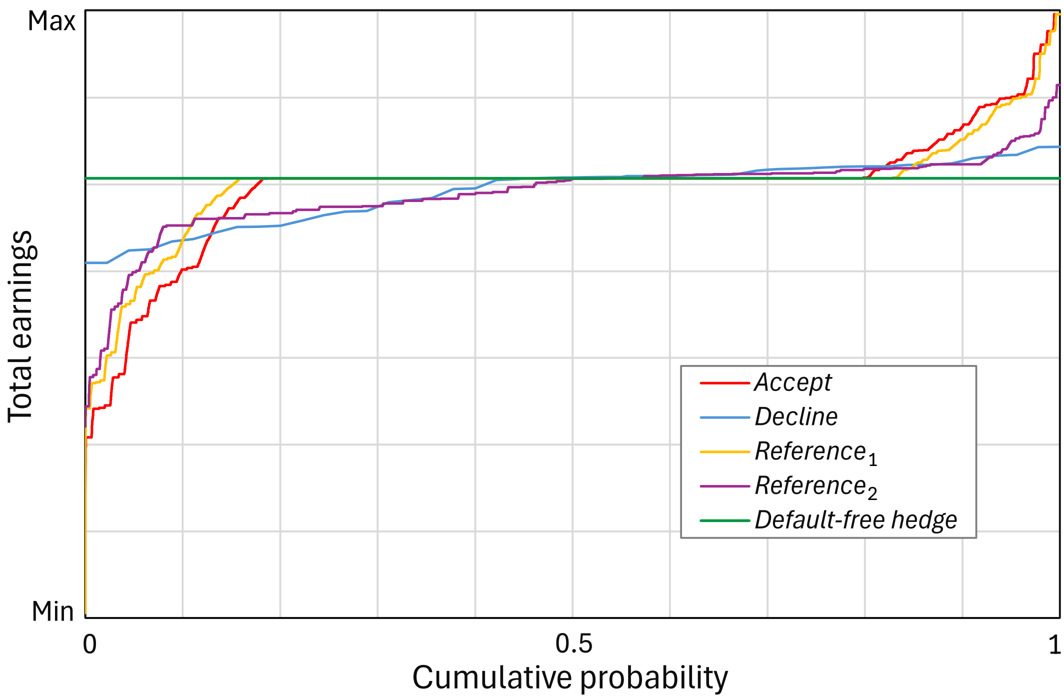

Figure 3: Cumulative distributions of total earnings in different decision branches, including an additional *default-free hedge* level, where all earnings uncertainty is removed. Note that, contrary to convention, cumulative probability is plotted on the horizontal axis and total earnings on the vertical axis to emphasize flat regions corresponding to outcomes with constant total earnings across a range of cumulative probabilities (i.e., the hedging effect).

Note that for confidentiality reasons, the exact total earnings values and cumulative probabilities in Figure 3 have been removed. In addition, the distributions have been subtly modified from the real data to obscure the true relative scales and shapes of the distributions. Total earnings are described only in terms of the overall minimum and maximum values, and for cumulative probabilities, only the minimum, median, and maximum are shown explicitly. As such, Figure 3 is intended purely as a qualitative illustration to highlight general trends and obvious differences between the branches.

Figure 3 provides several interesting and useful insights into the behavior of the different decision branches. The *decline* branch distribution stays relatively close to the *default-free hedge* level throughout the entire distribution. The median value is very similar to that of the *default-free hedge* level, whereas the total earnings start to deviate from this level towards the tails of the distribution. Overall, the distribution follows a relatively linear shape.

In contrast, the *accept* branch distribution has a wide flat region around the median that coincides exactly with the *default-free hedge* level. This region corresponds to all total earnings outcomes where the contracting counterparty does not default, and hence, all earnings come from the completed contract. Outside of this flat region, some default outcomes produce higher total earnings than the *default-free hedge* level; these arise when selling the remaining volume at the spot price after a default is more favorable than the original contract price P_c. Conversely, other default outcomes produce lower total earnings; these correspond either to less favorable spot prices for the remainder of the delivery period, or to situations when significant payments are required to be made to the defaulting counterparty due to a large negative MtM valuation of the contract at the time of default (see Section 3.2).

The total earnings in the $reference_1$ branch, which corresponds to a counterparty with a higher credit rating, are very similar to the *accept* branch total earnings. The effect of the higher creditworthiness is apparent from the distribution shape: the flat region is slightly wider, and the total earnings at the tails of the distribution are closer to the *default-free hedge* level (apart from the absolute extremes, which are equal between the two branches), indicating a lower probability of default.

On the contrary, the $reference_2$ branch can be seen as a middle ground of the *accept* and *decline* branches because the total volume is effectively split equally between such hedging strategies. This effect is also apparent from the distribution shape: the tails of the distribution exhibit similarities to the *accept* branch in both magnitude of total earnings and also the shape of the distribution. However, in the region where the *accept* distribution is flat, the $reference_2$ distribution instead increases approximately linearly from below the *default-free hedge* level to above it, which closely resembles the behavior of the *decline* branch. This part of the distribution likely corresponds mostly to such outcomes where the hedge regarding the large contract is fully complete and uncertainty arises from the volume allocated to the portfolio. The equal split of the volume between the large contract and the portfolio is also visible in the extremes of the distribution; notably, the minimum and maximum total earnings levels of the *reference*$_2$ branch lie approximately in the middle of those of the *decline* and *accept* branches.

Overall, the qualitative performance of the different branches can be summarized in the following way. On the downside (and similarly on the upside, although downside earnings are more critical from a risk management perspective and therefore are the main focus of the study), the *decline* and *reference*$_2$ branches appear to outperform the *accept* branch because their total earnings are closer to the *default-free hedge* level. However, around the median, the *accept* branch performs better because it has a more reliable hedging effect. The *reference*$_1$ branch outperforms the *accept* branch in terms of the hedging effect in all regions where the *accept* distribution deviates from the *default-free hedge* level. However, this is expected because the more creditworthy counterparty in *reference*$_1$ has a lower probability of default, which consequently improves the reliability of the hedge.

4.3 Collateral requirement and CVA add-on optimization

In this section, the evaluated decision tree from Section 4.2 is used to optimize the collateral requirement and the CVA add-on such that $\mathrm{CVaR}_\alpha^{accept}$ matches predefined downside earnings target levels. The following target levels will be used:

i) *Decline* branch CVaR_α: This target level reflects the amount of additional requirements for the counterparty, such that the *accept* branch becomes equally attractive as the *decline* branch in terms of downside risk.

ii) *Reference*$_1$ branch CVaR_α: At this target level, the *accept* branch becomes equally attractive as the *reference*$_1$ branch in terms of downside risk.

iii) *Reference*$_2$ branch CVaR_α: At this target level, the *accept* branch becomes equally attractive as the *reference*$_2$ branch in terms of downside risk.

iv) *Default-free hedge* level: This target level represents a counterparty that cannot default. This target level serves more as a theoretical upper bound, corresponding to a level of additional requirements that would eliminate all counterparty credit risk at the given confidence level. Thus, in practice, it would not be used as a realistic target level for the *accept* branch downside earnings.

As a result of this optimization, indifference curves will be constructed using the optimal pairs of collateral requirement $\mathcal{C}_c$ and CVA add-on CVA_c for each of the four target levels at each of the confidence levels $\alpha \in \{5\%, 10\%, 15\%, 20\%\}$. These curves illustrate all combinations of additional requirements that result in the *accept* branch achieving CVaR_α-earnings which are equal to the corresponding target levels. The indifference curve corresponding to the initial *accept* branch (i.e., a level which equals the total earnings level before any additional requirements are applied) will also be shown, which by construction always passes through the origin.

To visualize the full shapes of the indifference curves, they will be extended with negative CVA add-on values, which would correspond to discounts on the counterparty's proposed sales price P_c. In practice, Fortum would not voluntarily offer such discounts on proposed sales prices, so this extension is included solely for illustrative purposes. The region with non-negative (i.e., realistic) CVA add-ons will later be referred to as the *feasible region*, and the region with negative CVA add-ons as the *infeasible region*. Note that a similar extension cannot be made for collateral requirements, as a negative collateral settlement at the time of default is not well defined.

In the indifference curve figures, the horizontal axis (CVA_c) and the vertical axis ($\mathcal{C}_c$) are bounded to reflect realistic contractual terms. The minimum collateral requirement is set to zero, while the maximum corresponds to the largest positive MtM valuation of the contract across all spot price scenarios and delivery period months in the model; a collateral requirement beyond this maximum would not further improve the *accept* distribution (see Section 3.2). Conversely, the exact minimum and maximum values for the CVA add-on axis remain undisclosed, but the range is limited to values that are considered somewhat reasonable for the contract outlined in Section 4.1. This is done because the CVA add-on could theoretically be increased (or decreased) indefinitely, but analyzing excessively high values would not offer useful insights in contractual negotiations.

The collateral requirement and the CVA add-on optimization and indifference curve creation are first carried out at a confidence level $\alpha = 10\%$. Subsequently, the analysis is repeated for three other confidence levels $\alpha \in \{5\%, 15\%, 20\%\}$ to perform a sensitivity analysis and assess how varying degrees of risk aversion (in the sense of the width of the downside α-quantile used in the CVaR_α-earnings calculations) influence the requirements to meet each target level.

Similar to Figure 3, the exact values of collateral requirements and CVA add-ons are not shown, apart from the zero values, which represent the defined, realistic minimums. However, the indifference curve figures in Sections 4.3.2 and 4.3.3 share the same scaling to allow for visual comparisons between different confidence levels α. Note that the introductory indifference curve in the following Section 4.3.1 uses a different scaling, as it is intended solely for illustrative purposes to present the topic using simplified data.

4.3.1 An illustrative example

We begin with a simple example to demonstrate how the use of the collateral requirement and the CVA add-on impacts the *accept* branch distribution in later optimizations. Figure 4 shows three side-by-side plots, where the *decline* and *accept* distributions as well as the *default-free hedge* level are visible for a part of the tail distribution from Figure 3, and the downside earnings are examined at a confidence level α. The leftmost plot is identical to the corresponding portion of Figure 3. In the middle plot, a collateral requirement has been applied, shifting the *accept* downside earnings to match the *decline* target level. In the rightmost plot, a similar result has been achieved using a CVA add-on.

Instead of displaying total earnings on the vertical axis as in Figure 3, the vertical axes in Figure 4 use *percentual total earnings*, where 0% and 100% correspond to the absolute minimum total earnings across all branches (i.e., the 0%-quantile in Figure 3) and the *default-free hedge* level, respectively. Additionally, the different total earnings and downside earnings target levels in the later optimizations will be represented as percentual total earnings in this way.

At the confidence level α in Figure 4, the downside percentual total earnings of the *accept* and *decline* branches are $\mathrm{CVaR}_{\alpha}^{accept} = 82\%$ and $\mathrm{CVaR}_{\alpha}^{decline} = 88\%$, respectively (leftmost plot). The *accept* CVaR_{α}-earnings are then lifted to 88% separately with a collateral requirement and a CVA add-on (middle and rightmost plots, respectively).

The middle and rightmost plots exhibit clear differences in how the collateral requirement and the CVA add-on change the *accept* distribution. As discussed in Section 3.2, the collateral requirement affects only those outcomes where the counterparty defaults and the MtM valuation of the contract is positive at that time. It is apparent that with the collateral requirement in place, the tail of the distribution shifts noticeably towards the *default-free hedge* level. The flat region of the *accept* distribution also becomes marginally wider because some outcomes where the total earnings were previously slightly below the *default-free hedge* level due to a default are now raised to that level by the collateral.

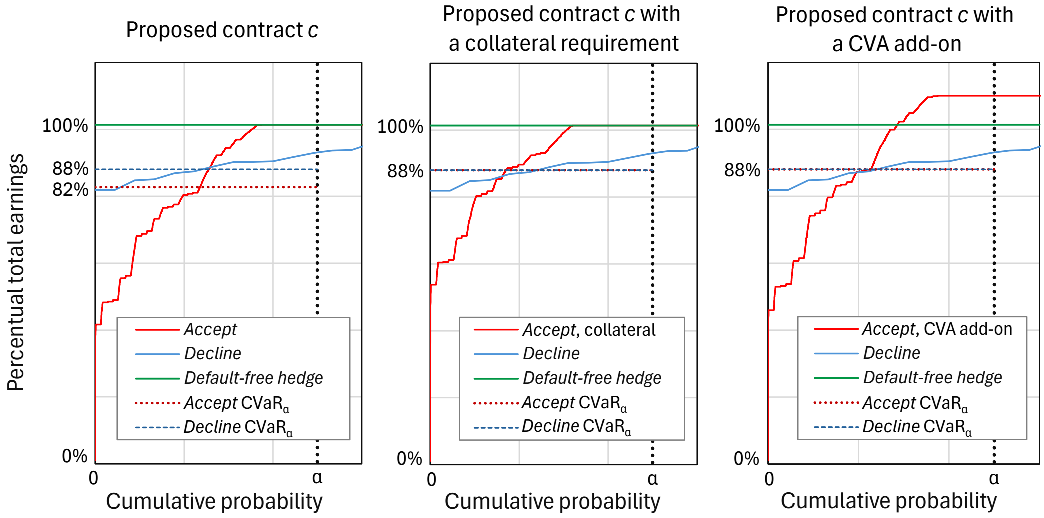

Figure 4: On the left: *Accept* and *decline* distributions as well as the *default-free hedge* level, which are identical to those from Figure 3. In the middle: A collateral requirement is applied to lift $\mathrm{CVaR}_\alpha^{accept}$ to match $\mathrm{CVaR}_\alpha^{decline}$. On the right: The same outcome has been achieved by applying a CVA add-on instead. The vertical axes show percentual total earnings as opposed to total earnings used in Figure 3.

Conversely, while the CVA add-on also shifts the *accept* distribution tail slightly closer to the *default-free hedge* level, the main distinction is that its flat region no longer aligns with the *default-free hedge* level. This is because with a higher sales price given by the CVA add-on, all outcomes where the contract does not terminate prematurely result in higher total earnings than with the proposed contractual price P_c. Furthermore, although the flat region of the *accept* distribution is at a different total earnings level compared to the leftmost plot in Figure 4, its width remains the same because the CVA add-on does not mitigate default-related losses in the same way the collateral requirement does.

The collateral requirement and CVA add-on indifference curve can also be plotted for the *decline* target level, which can be seen in Figure 5. The three points in Figure 5 have been labeled as A, B, and C. Point A (the origin) corresponds to the leftmost plot in Figure 4, where no additional requirements have yet been applied to the contractual terms, and the percentual total earnings level is at 82%. On the other hand, points B and C correspond to the middle and rightmost plots in Figure 4, respectively. Therefore, B corresponds to lifting $\mathrm{CVaR}_\alpha^{accept}$ to 88% by applying a collateral requirement, whereas C achieves the same level by applying a CVA add-on.

As mentioned previously, one can also find combinations of collateral requirements and CVA add-ons, rather than relying on only one or the other, as was done in this

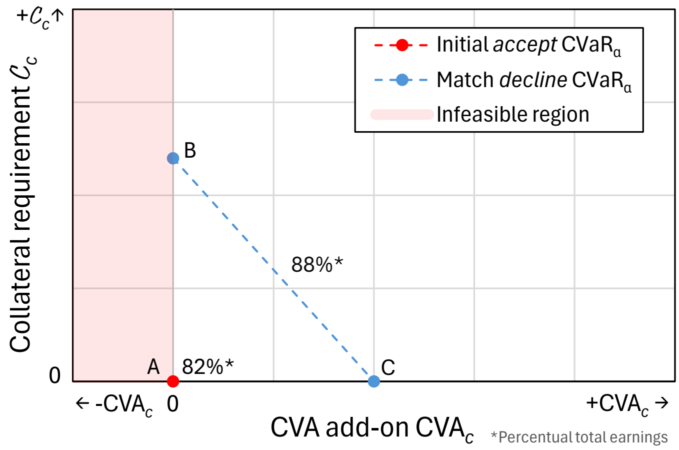

Figure 5: Initial percentual total earnings of the *accept* branch, and the indifference curve for the *decline* target level in this introductory example. Points A, B, and C correspond to the leftmost, middle, and rightmost plots in Figure 4, respectively. Note that the scaling here differs from that used in Sections 4.3.2 and 4.3.3, where the figures share the same scaling for the sake of visual comparisons.

introductory example. This is the approach used in the optimizations presented in Sections 4.3.2 and 4.3.3, where such combinations are identified for the four different target levels outlined earlier.

In the indifference curves in Sections 4.3.2 and 4.3.3, the higher the percentual total earnings of a target level is, the closer that target level is to the *default-free hedge* level, which has percentual total earnings level of 100% by construction. Similarly, the larger the difference in percentual total earnings between the $\mathrm{CVaR}_\alpha^{accept}$ and a target level is, the more additional requirements are needed to lift $\mathrm{CVaR}_\alpha^{accept}$ to that level. Consequently, the corresponding target level indifference curve lies further from the origin.

4.3.2 Base case: confidence level $\alpha = 10\%$

The results of the collateral requirement and CVA add-on optimization with confidence level $\alpha = 10\%$ are shown in Figure 6. The first observation is that the indifference curves for the initial *accept* branch and each of the four target levels are convex throughout the entire region of the plot. This indicates that at lower CVA add-on levels, a small increase in the CVA add-on leads to a larger reduction in the needed collateral requirement to maintain any given downside earnings level. Furthermore, at large CVA add-on values, particularly those needed to match the *decline* and *reference*$_1$ $\mathrm{CVaR}_{10\%}$-earnings, the relationship between the collateral requirement and the CVA add-on appears to become approximately linear with a small negative slope, such that increasing the CVA add-on yields only marginal reductions in the needed collateral requirement.

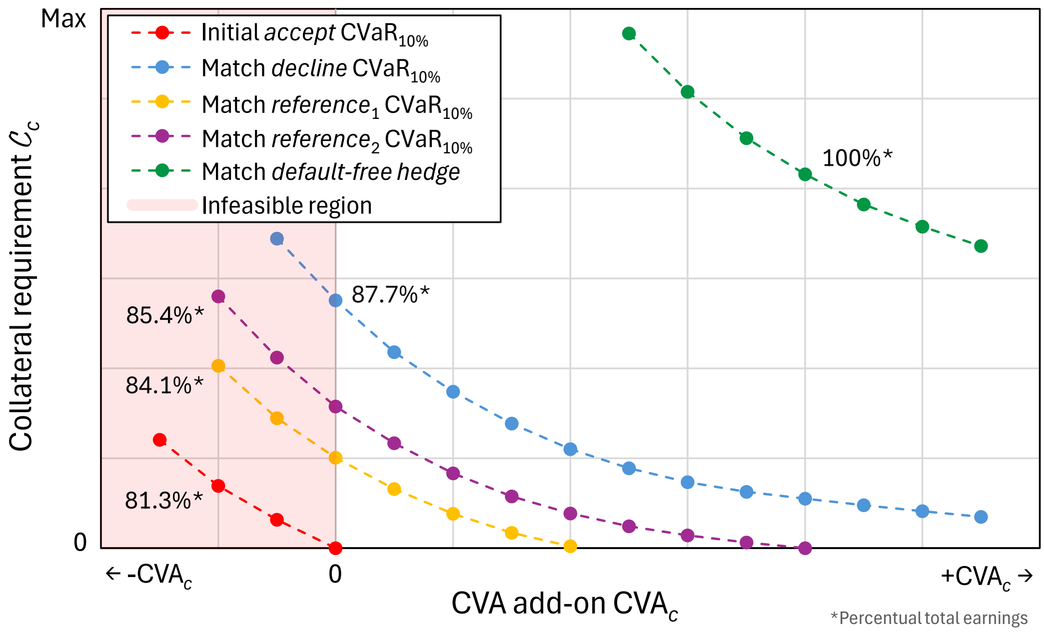

Figure 6: Indifference curves showing combinations of collateral requirement $\mathcal{C}_c$ and CVA add-on CVA_c needed for the initial *accept* $\text{CVaR}_{10\%}$ (red) to match the i) *decline* $\text{CVaR}_{10\%}$ (blue), ii) *reference*$_1$ $\text{CVaR}_{10\%}$ (yellow), iii) *reference*$_2$ $\text{CVaR}_{10\%}$ (purple), or iv) *default-free hedge* level (green). The percentage values on the indifference curves represent percentual total earnings (see Section 4.3.1).

This trade-off between the collateral requirement and the CVA add-on highlights the advantages of potentially reducing the collateral requirement significantly with a relatively small CVA add-on. The indifference curves can therefore serve as useful negotiation tools in commercial settings, where additional requirements may be tailored for specific counterparties. This is particularly valuable, as the credit ratings R_c and $R_{c'_n}$ do not account for the size or equity of the counterparty. For example, a large counterparty with substantial equity may estimate its own default risk as low and hence prefer to post collateral rather than pay a higher total price for the electricity. On the other hand, a smaller counterparty with less equity may not be able to post large collateral amounts and instead prefer a slightly higher final sales price. The indifference curves could also enable the power producer to, for example, offer the counterparty a counter-offer containing multiple collateral requirement and CVA add-on combinations to choose from.

Another observation concerns the percentual total earnings values in Figure 6. The initial *accept* $\mathrm{CVaR}_{10\%}$-earnings (81.3%) is lower than all four target levels of *decline* $\mathrm{CVaR}_{10\%}$ (87.7%), *reference*$_1$ $\mathrm{CVaR}_{10\%}$ (84.1%), *reference*$_2$ $\mathrm{CVaR}_{10\%}$ (85.4%), and *default-free hedge* level (100%), which is also why all target level indifference curves are at least partly inside the feasible region. We also note that the percentual total earnings values increase away from the origin with the indifference curves. This is expected because the closer to the *default-free hedge* level any given target level is, the larger additional requirements are needed to increase *accept* $\mathrm{CVaR}_{10\%}$-earnings to that level.

4.3.3 Sensitivity analysis

The indifference curves for the remaining confidence levels $\alpha \in \{5\%, 15\%, 20\%\}$ can be seen in Figures 7, 8, and 9, respectively. Similar to the case $\alpha = 10\%$, all indifference curves at all remaining confidence levels are convex. However, the slopes of the curves change noticeably when the confidence level is changed. Specifically, as α increases (i.e., a wider tail of the distributions is used to calculate the CVaR_α-earnings), the slopes of the indifference curves decrease; that is, the trade-off between the collateral requirement and the CVA add-on becomes steeper. Conversely, as α decreases, the slopes become flatter. Next, we analyze each of the remaining confidence levels $\alpha \in \{5\%, 15\%, 20\%\}$ separately and make comparisons to the other previously examined levels throughout the analysis.

Looking at Figure 7 reveals the stiffness of the *accept* $\mathrm{CVaR}_{5\%}$-earnings at such a narrow confidence level. Comparing the percentual total earnings of the *accept* $\mathrm{CVaR}_{5\%}$ (66.3%) to the *decline*, *reference*$_1$, and *reference*$_2$ branches (86.6%, 71.1%, and 75.7%, respectively) shows that *accept* $\mathrm{CVaR}_{5\%}$ is initially significantly below the target levels. Moreover, the theoretical upper bound of the *default-free hedge* level is unreachable within the feasible range of collateral requirement and CVA add-on values. Only the other three target levels can be reached, and their corresponding indifference curves are nearly horizontal. This indicates that the CVA add-on has a negligible effect on the total earnings at confidence level $\alpha = 5\%$, and that the collateral requirement dominates the impact on the downside earnings.

The flat shape of the indifference curves suggests that many of the downside earnings outcomes in the worst 5%-quantile are those in which the counterparty defaults early in the contract lifetime when none or only a small portion of the contractual obligations have been realized, and the spot price for the remaining duration is unfavorable. This observation is supported by the fact that the CVA

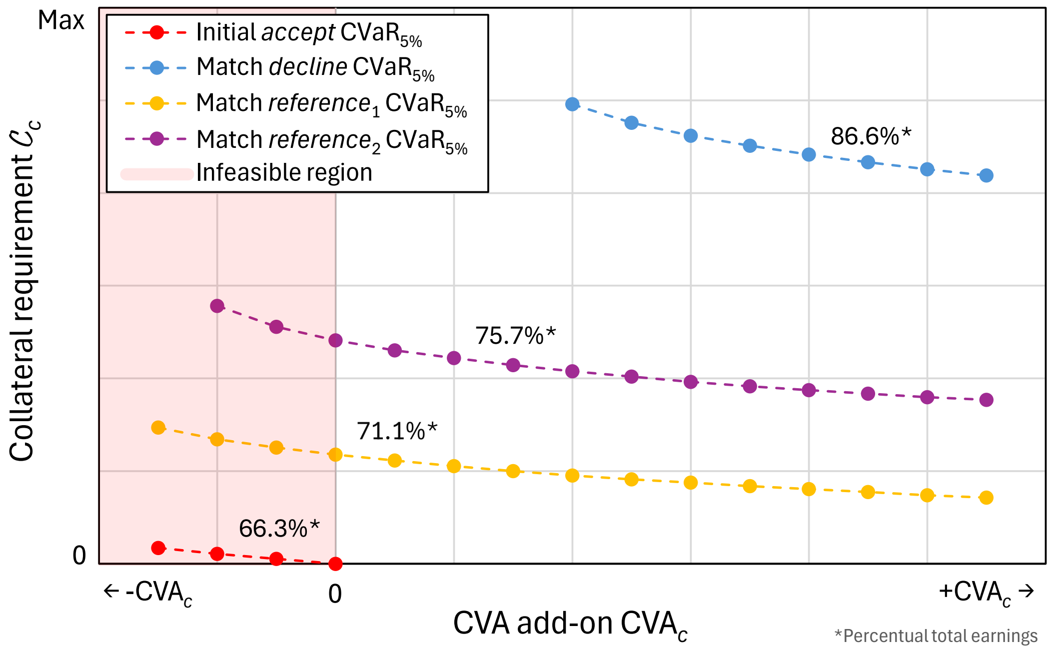

Figure 7: Indifference curves at confidence level $\alpha = 5\%$. Note that the indifference curves are particularly flat at such a narrow confidence level. Furthermore, the curves are generally further away from the origin, with the *default-free hedge* level being outside the feasible region and hence not visible in the plot. As a result, the initial *accept* $\text{CVaR}_{5\%}$ cannot be increased to match the *default-free hedge* level within realistic collateral requirement and CVA add-on bounds.

add-on serves only as a minor component of the total earnings compared to the collateral requirement, meaning that even a large increase in the contract price does not affect the downside earnings considerably (cf. Figure 4). However, because the collateral requirement has an upper bound beyond which it does not affect the *accept* distribution, and only realistic CVA add-on values are used, the potential to reach higher total earnings target levels becomes somewhat limited.

Additionally, the (relatively small) number of (total earnings, probability) pairs in $\mathcal{E}_S^{decline}$ (135 in total) should be kept in mind when comparing *accept* and *decline* branch total earnings at such a narrow confidence level. In the *decline* branch, only seven data points with the lowest earnings fall within the worst 5%-quantile, which introduces some obvious statistical uncertainty—especially as the number of (total earnings, probability) pairs in $\mathcal{E}_\Omega^{accept}$ is significantly larger (17 750 in total). Although the variation in the *decline* branch distribution is not extremely high to begin with (see Figure 3), this data sparsity can affect the robustness of the comparisons and should be considered when interpreting the results.

Overall, the analysis suggests that using such a narrow confidence level may be impractical in reality because it strongly emphasizes the extreme adverse outcomes that are statistically more unstable. As many of the outcomes in the worst 5%-quantile include an early default, the CVA add-on effect becomes negligible and thus difficult to analyze. A higher α (such as $\alpha = 10\%$ from earlier) may provide a more balanced and realistic view of the downside risks.

Figure 8 presents the indifference curves for confidence level $\alpha = 15\%$. Compared to the narrower confidence levels presented earlier, these curves are noticeably steeper and are shifted closer towards the bottom-left corner of the plot. This reflects the higher percentual total earnings of the initial $\mathrm{CVaR}_{15\%}^{accept}$-earnings. Notably, the *accept* $\mathrm{CVaR}_{15\%}$ (87.6%) already exceeds the *reference*$_2$ $\mathrm{CVaR}_{15\%}$ target level (87.3%) without any additional requirements. These observations suggest that the *accept* $\mathrm{CVaR}_{15\%}$ can be more easily increased to the various target levels by applying the collateral requirement and the CVA add-on.

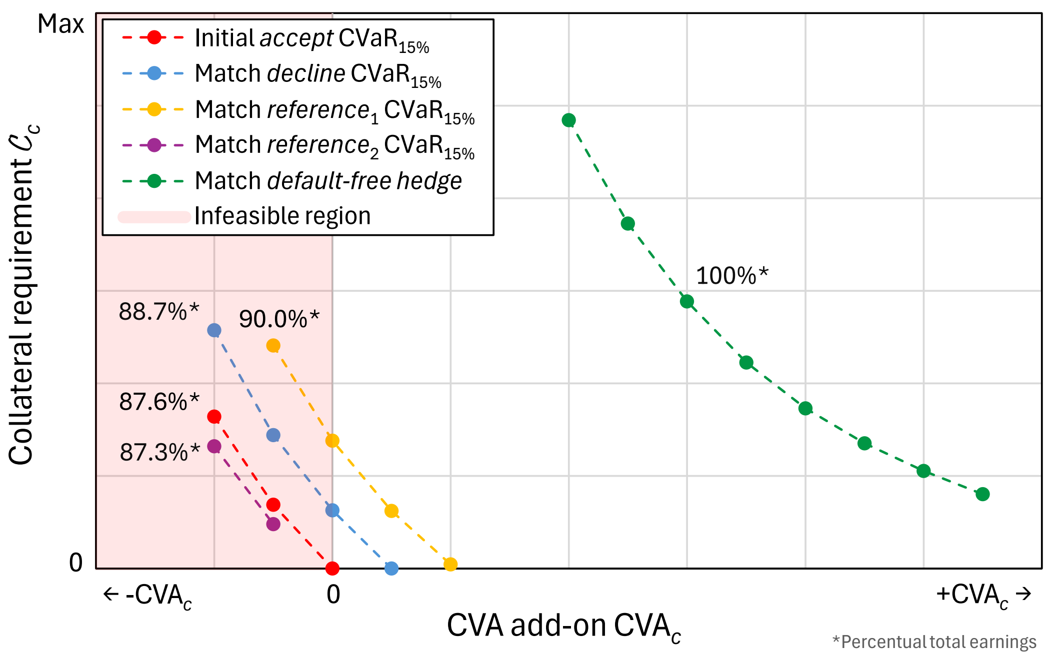

Figure 8: Indifference curves at confidence level $\alpha = 15\%$. The indifference curves are noticeably steeper when compared to those resulting from $\alpha = 5\%$ and $\alpha = 10\%$ (see Figures 7 and 6, respectively). Note that the *accept* $\mathrm{CVaR}_{15\%}$ is already higher than the target level *reference*$_2$ $\mathrm{CVaR}_{15\%}$ without any additional requirements.

As the confidence level becomes wider, more total earnings outcomes with a nearly or fully complete hedge are included in the tail of the distribution. The increased steepness of the curves supports this interpretation: when a larger portion of outcomes within the worst α-quantile result in most of the volume sold at the contractual price, the effectiveness of the CVA add-on on the downside increases relative to that of the collateral requirement. In these outcomes, the collateral requirement has a minimal impact on the total earnings, whereas the CVA add-on significantly improves expected total earnings when the contract is not terminated prematurely (see Figure 4). However, as is illustrated by the *default-free hedge* level indifference curve in Figure 8, increasing the CVA add-on beyond some point to cut down on the collateral requirement becomes less beneficial, which suggests that including a moderate collateral requirement in addition to a CVA add-on may yield more balanced overall additional requirements for the counterparty.

Interestingly, the ordering of the indifference curves (based on the percentual total earnings) is different when compared to those observed at the narrower confidence levels from earlier. At both $\alpha = 5\%$ and $\alpha = 10\%$, the curves follow the ascending order (based on the percentual total earnings of each curve): *accept*, *reference*$_1$, *reference*$_2$, *decline*, and *default-free hedge* level. However, with $\alpha = 15\%$, the ordering is instead *reference*$_2$, *accept*, *decline*, *reference*$_1$, and *default-free hedge* level. This reflects the unique shapes of the distributions (see Figure 3) and that there is no absolute objectively superior hedging strategy (apart from the hypothetical *reference*$_1$ being always favored over *accept*).

This shift gives insights into the differences in the shapes of the total earnings distributions. At this wider confidence level, when looking at the downside tails of the distributions, the *accept* branch performs nearly equally well as the *reference*$_1$ branch. In contrast, the extreme low total earnings downside outcomes in the *decline* and *reference*$_2$ branches lie at noticeably higher total earnings levels (see Figure 3) because both include the portfolio either fully or partially. Furthermore, because a larger portion of the distribution tail is captured at wider confidence levels, the consistently better hedging effect of the *reference*$_1$ makes its target level more difficult to reach with additional requirements when compared to the *decline* and *reference*$_2$ branches.

Finally, Figure 9 presents the indifference curves for confidence level $\alpha = 20\%$. The curves are once again steeper than those from earlier. Notably, the percentual total earnings level of the *accept* $\mathrm{CVaR}_{20\%}$ is the highest among all analyzed confidence levels, reaching 90.7%. This high initial level results in the *accept* $\mathrm{CVaR}_{20\%}$ exceeding both the *decline* and *reference*$_2$ target levels without any additional requirements, meaning that the indifference curves for both of these target levels lie fully within the infeasible region.

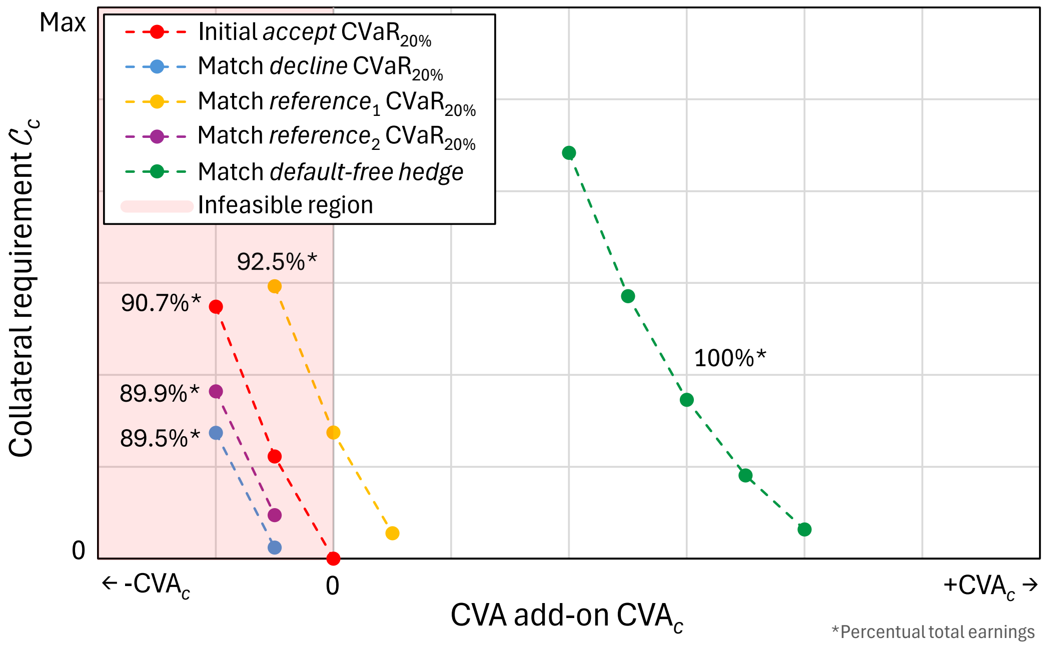

Figure 9: Indifference curves at confidence level $\alpha = 20\%$. Note that the *accept* $\text{CVaR}_{20\%}$ is higher than target levels *decline* $\text{CVaR}_{20\%}$ and *reference*$_2$ $\text{CVaR}_{20\%}$ without additional requirements.

As discussed in the case $\alpha = 15\%$ previously, increasing the value of α (i.e., considering a wider tail of the distribution) includes more total earnings outcomes in which the hedge is nearly or fully complete. The steepness of the curves at $\alpha = 20\%$ suggests that the CVA add-on dominates the effect of the collateral requirement for improving the downside earnings, meaning that for a relatively small CVA add-on, the collateral requirement can be lowered significantly while remaining at the given target level. In contrast, the impact of the collateral requirement is limited, as when widening the confidence level, a smaller relative portion of the total earnings outcomes in the *accept* branch α-quantile are ones considerably affected by the collateral requirement. This is because in many of these outcomes, the contract is completed without a premature termination due to default, and hence, the full collateral amount is returned to the counterparty. Thus, as the confidence level widens, the CVA add-on becomes the main driver for improving the downside earnings.

However, this observation highlights a potential drawback. As the goal of the model is to capture a wide range of total earnings outcomes (including adverse ones) and mitigate the related downside risk, a confidence level as wide as $\alpha = 20\%$ may be too wide for practical use. When the probability mass of nearly and fully complete hedging outcomes begins to dominate the rarer adverse outcomes, the usefulness of the downside risk measure CVaR_α may diminish beyond practical applicability.

Overall, based on the indifference curves at varying confidence levels as depicted in Figures 6–9, the confidence levels $\alpha = 10\%$ and $\alpha = 15\%$ appear to be the most appropriate ones for practical use with this model. They are sufficiently robust to capture the relevant downside outcomes without being too sensitive to extreme low total earnings outcomes, as observed with $\alpha = 5\%$, or diluting the CVaR_α-earnings with too many nearly or fully complete hedging outcomes, as in the case $\alpha = 20\%$. This balance of trade-offs should always be carefully analyzed to ensure that the resulting risk assessment can guide decision-making as realistically and tractably as possible, and that the relationships between the collateral requirement and the CVA add-on are properly understood to support commercial PPA negotiations.

4.4 Critical reflections and potential future development

The decision tree model outlined in Section 3 relies on many assumptions and simplifications that need to be highlighted and critically assessed. Some of these stem from input data availability, others from computational limitations, and some are made up to simplify complex real-world relationships for the sake of computational tractability. This section addresses these assumptions and simplifications and points towards potential future development of the model. Furthermore, other general challenges regarding real-life PPAs (and their modeling) are brought up and discussed.

One obvious potential limitation of the model is the number of spot price scenarios (135 in total). This becomes especially relevant when calculating CVaR_α-earnings under highly narrow confidence levels in the *decline* branch, where counterparty defaults do not play a role in the estimated earnings. In such cases, the number of total earnings outcomes within the α-quantile, from which the CVaR_α-earnings are subsequently computed, is small (potentially too small for statistical reliability), and the results should therefore be interpreted with caution in mind. However, this limitation is less critical with wider confidence levels α, and generally in branches where counterparty defaults are included, as the number of total earnings outcomes is considerably larger (e.g., $135 \times 180 = 17\ 550$ in this case study), resulting in sufficiently accurate CVaR_α-earnings calculations.

The assumption of independence between spot price and counterparty default scenarios is a computational simplification that does not necessarily hold in reality. For instance, if the contracting counterparty is a major electricity consumer, its sudden default may have a significant impact on the total electricity demand on a national or Nordic scale, and hence also on the market price of electricity. However, modeling such dependencies between the spot market and individual market participants realistically and reliably would be extremely challenging, and the modeling of such dependencies was therefore excluded from the scope of this thesis.

Another assumption that does not hold in real-life PPAs is the unilateral treatment of collateral requirements. In practice, it is perfectly realistic that the contracting counterparty, through its own risk management process, would post a collateral requirement to Fortum as compensation for its perceived counterparty credit risk. As a consequence, this collateral amount would be effectively tied up and unavailable for investments or operational use, and would impact Fortum's total financial flexibility to some extent. Furthermore, related to this, there is also naturally a possibility that Fortum itself may default, which would have major financial implications on a much larger scale than what the model is designed to capture. However, because the model currently only considers the earnings from a single PPA and assumes no internal default, extending the model to incorporate these aspects would be tricky and, therefore, also excluded from the scope of the thesis.

Some of the current model assumptions and simplifications can be relaxed with relative ease. One simple extension would be to allow for fluctuations in the PPA pricing. For instance, if a counterparty proposes a 10-year PPA with one price for the first five years and another price for the remaining five, the current model cannot accommodate this kind of pricing (technically, one could treat such agreements as separate contracts with different prices, simulate their earnings separately, and sum these afterwards, but even this would require many additional manual steps). Similarly, uneven yearly volumes across the delivery period cannot be modeled currently, but their addition would account for more realistic and flexible agreement terms. Both of these features can be incorporated without major changes to the model formulation.

Another relatively straightforward addition to the model would be to incorporate discounting of future cash flows. This would be particularly relevant when the yearly price or volume is uneven, as discussed above. Accounting for the time value of money (meaning that money is more valuable the sooner it is available) would allow for representing the present value of earnings coming from a PPA, if the decision-maker were interested in assessing their discounted value.

The model could also be extended to account for counterparty defaults within the portfolio. While modeling the default of each counterparty in the portfolio $\mathcal{K}$ at a monthly granularity, as is currently done with the contracts c and c'_n in the model, would be computationally infeasible, a simplified approach using broader time and counterparty bucketing could be implemented. For example, yearly default scenarios could be defined for predefined subsets of the portfolio $\mathcal{K}$, which would enable the model to capture some counterparty-related risks in the *decline* branch and hence increase the number of total earnings outcomes in the *decline* distribution from the current 135 to something much more statistically robust.

Another relatively straightforward improvement would be to introduce time-dependency in the credit ratings of contracting counterparties. The contract lifetime in the case study spans nearly 15 years—a duration that is not uncommon in real long-term PPAs. Over such long time horizons, changes in creditworthiness are more than likely. To reflect these aspects, a credit rating transition matrix could be incorporated into the model, which would specify (for instance) the annual probabilities of migrating from one credit rating to another. Moreover, as was also highlighted in the literature review in Section 2, defaults are often preceded by gradual creditworthiness deterioration, which suggests that credit ratings and default events have a clear dependency relationship in reality.

A further refinement of the model relates to the remaining unsold volume after a counterparty default event. In reality, in such a situation, especially before or early in the delivery period, it is somewhat unrealistic to assume that the full remaining volume would be sold solely on the spot market for the remaining delivery period, if the preferred way would instead be to hedge it in some way. In practice, soon after a counterparty default, the remaining volume could possibly either be re-hedged through the portfolio, similar to what is currently assumed in the *decline* branch, or sold as a large PPA to another counterparty, should one become interested in contracting a PPA at the time. This addition would provide more realistic total earnings estimates in all branches that involve a contracting counterparty.

One opportunity for larger future development is to extend the scope of the earnings impact. The model could be extended to consider a portfolio of bilateral contracts with the same counterparty, instead of a single contract as is currently done. The model could also be extended to consider a portfolio of bilateral contracts with many different counterparties. Both of these extensions would bring their own challenges, both conceptually and computationally, even with additional simplifications to the model. Nonetheless, these extensions would add value in terms of the model's scope and are left as opportunities for future work.

The model could also be extended with alternative references where the decision to sign a contract is postponed rather than made in m_{sign}. This would not only account for the earnings structure of the current $\textit{reference}_n$ branches, but also the uncertainty related to finding future opportunities at a later time. For example, the decision-maker may pose the question: *"If we decline the proposed contract and wait for a year or two, what is the probability that similar contract opportunities would arise? Furthermore, if such opportunities were to arise, with what probability would they produce more favorable earnings than the current proposed contract?"* This kind of feature could be incorporated using consecutive decisions (for example) at a yearly granularity, such that this new alternative reference branch would resemble the secretary problem (Freeman, 1983) and could be solved in a similar way, using both historical data and expert judgment to estimate these probabilities. Some further simplifications and assumptions would likely have to be included in the model to support this structure, but this extension could potentially capture more thorough supply and demand dynamics related to large PPAs.

Real PPAs also bring other challenges and risks that should be considered when negotiating contractual terms. One such risk is the cancellation risk of the counterparty. For instance, a start-up may seek to contract a PPA for a project, but if the project ultimately gets canceled, then the power producer may be left without a buyer for the volume of electricity. This risk is different from traditional counterparty credit risk and is not captured in the credit ratings used in the model. As such, case-specific additional qualitative assessment is needed to build confidence in whether to decline or accept a proposed PPA, particularly when dealing with start-up counterparties.

As noted earlier, the credit ratings used in the model do not account for the size or equity of the counterparty. Therefore, some additional processes to support the total earnings distributions coming from the model may be required to fully evaluate the proposed contract and its risks in their entirety. Additionally, the model currently focuses on earnings from a single PPA and does not consider potential existing agreements with the same counterparty. In reality, there may be some counterparty-specific limits related to credit, number of contracts, or total contracted volume, which should be taken into account outside of the model.

Another relevant risk that is not captured in the model is renegotiating risk. In reality, if market prices were to suddenly drop significantly, the contracting counterparty may be reluctant to continue paying the agreed price and may instead try to renegotiate more favorable terms. Additionally, potential reputation risks should be kept in mind; contracting PPAs with shady companies may be more harmful than beneficial in the long run. While these kinds of risks are difficult to quantify, their existence should nonetheless be acknowledged, and they should be mitigated accordingly.

5 Conclusions

This thesis developed and presented a probabilistic decision tree model to estimate earnings associated with a bilateral PPA from the perspective of a power producer and retailer (Fortum Oyj). Once the underlying context had been explained and the model was formulated, it was then applied to a case study, where expected earnings from accepting a hypothetical yet realistic PPA were compared to those from declining it. In addition, the earnings were also compared to two reference hedging strategies. While these references do not represent real decision alternatives for the power producer, they provide insightful comparisons to better understand the counterparty credit risk associated with the PPA.

This procedure addressed the first research question Q1 related to the expected earnings of different decision and reference alternatives. The total earnings distributions from each alternative were plotted together with a default-free hedge level, which revealed their distinct shapes and risk profiles. The hedging effect of the PPA was also apparent from the plot because, in many cases, the simulated earnings remained constant for wide portions of the cumulative distributions.

The case study was designed so that accepting the contract would initially not be the most favorable option in terms of downside earnings among the alternatives. This was done to demonstrate the model's full capabilities, where the contractual collateral requirement and CVA add-on are optimized in response to an initially unattractive contract proposal. Suppose a proposed contract would not sufficiently cover the counterparty credit risk from the power producer's perspective. In such a case, the power producer could, in practice, mitigate this risk by posing a collateral requirement to the counterparty or by including a CVA add-on to the contracted price. These additional requirements allow the power producer to increase the downside earnings to an acceptable level while keeping the requirements reasonable and justified.

This collateral requirement and CVA add-on optimization addressed the second research question Q2 related to finding suitable combinations of these two components to lift the contractual downside earnings to an acceptable risk level. This part of the model resembles the first round of negotiations between the parties, where the power producer reacts to an initial offer by potentially requesting additional compensation. Here, indifference curves for four different earnings target levels were plotted. These represent all pairs of collateral requirement and CVA add-on for which the downside earnings match the corresponding target levels, and provide useful insights into both the relationships between these requirements and the relationships between the different hedging strategies.

The thesis also analyzed the choice of confidence level for expected earnings to find out how changes in the distribution downside tail width (and thereby, the power producer's risk attitude) affect the collateral requirement and CVA add-on optimization. This sensitivity analysis addressed the third research question Q3. Altogether, four different confidence levels were studied, which provided interesting information about the earnings distribution shapes of each decision and reference alternative. Based on the results, confidence levels $\alpha = 10\%$ and $\alpha = 15\%$ appeared to be the most suitable ones for practical use because they were not overly sensitive to extreme adverse outcomes, while not diluting the worst α-quantile of total earnings with excessively many nearly or fully complete hedging outcomes.

Finally, this thesis critically assessed the model's assumptions and computational limitations and outlined a few directions for future development. Some of the limitations stem from inherent challenges of modeling complex energy market and counterparty credit risk dynamics both realistically and accurately, while others arise from the need to keep the model computationally tractable. Similarly, many assumptions were made to simplify the model for tractability reasons, even though they may not fully align with reality. Overall, while the model offers useful quantitative insights, it should be supported with other qualitative risk processes. For instance, the model does not currently account for the counterparty's size or equity, nor any existing trades with the counterparty, all of which are key factors in a comprehensive counterparty credit risk assessment.

In conclusion, this thesis contributes to a growing need for structured quantitative tools to evaluate counterparty credit risk in bilateral PPAs. The presented model is structured as a decision tree to reflect possible decision alternatives that the power producer may choose between regarding a proposed contract. While the model does not fully capture all aspects contributing to contractual risks, it provides a solid foundation for quantitative risk assessment and PPA negotiation support. Future work can build on the model by, for example, introducing more flexible pricing structures, discounting of long-term earnings, and considering portfolio-level impacts of the proposed contract, such as the combined effects of all active agreements with the same counterparty or across all existing counterparties.

References

M. Altdörfer, A. Carlos, A. Guettler, and G. Löffler. The case for a European rating agency: Evidence from the Eurozone sovereign debt crisis. *Journal of International Financial Markets, Institutions and Money*, 58:1–18, 2019.

N. Audet, P. Heiskanen, J. Keppo, and I. Vehviläinen. Modeling of electricity forward curve dynamics. *Ann Arbor, Espoo: University of Michigan, Fortum Power and Heat Oy*, 2002.

L. Ballotta, G. Fusai, and D. Marazzina. Integrated structural approach to credit value adjustment. *European Journal of Operational Research*, 272(3):1143–1157, 2019.

Bank for International Settlements (BIS). Basel III: A global regulatory framework for more resilient banks and banking systems. https://www.bis.org/publ/bcbs189.pdf, 2010. Accessed: 2025-01-30.

E. Benmelech and J Dlugosz. The credit rating crisis. *NBER Macroeconomics Annual*, 24(1):161–208, 2010.

F. Black and J. C Cox. Valuing corporate securities: Some effects of bond indenture provisions. *The Journal of Finance*, 31(2):351–367, 1976.

F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

E. Canabarro and D. Duffie. Measuring and marking counterparty risk. *Asset/Liability Management for Financial Institutions, Institutional Investor Books*, 2003.

A. Capponi. Pricing and mitigation of counterparty credit exposures. In J.-P. Fouque and J. Langsam, editors, *Handbook of Systemic Risk*. Cambridge University Press, 2012.

S. Chang, L. Xie, and J. Dumas. A reserve forecast-based approach to determining credit collateral requirements in electricity markets. In *2015 IEEE Power Energy Society General Meeting*, pages 1–5. IEEE, 2015.

U. Cherubini. Credit valuation adjustment and wrong way risk. *Quantitative Finance Letters*, 1(1):9–15, 2013.

A. Cohen and N Costanzino. Bond and CDS Pricing via the stochastic recovery Black-Cox model. *Risks*, 5(2):26, 2017.

P. J. Crosbie and J. R. Bohn. Modeling default risk. Technical report, KMV LLC, 1993.

R. A. De Santis. The euro area sovereign debt crisis: Identifying flight-to-liquidity and the spillover mechanisms. *Journal of Empirical Finance*, 26:150–170, 2014.

E. DeHaan. The financial crisis and corporate credit ratings. *The Accounting Review*, 92(4):161–189, 2017.

M. Denton, A. Palmer, R. Masiello, and P. Skantze. Managing market risk in energy. *IEEE Transactions on Power Systems*, 18(2):494–502, 2003.

D. Duffie and K. J. Singleton. Modeling term structures of defaultable bonds. *The Review of Financial Studies*, 12(4):687–720, 1999.

P. Edge. An approximation of counterparty credit risk in long term power purchase agreements (PPAs). In *2015 12th International Conference on the European Energy Market (EEM)*, pages 1–4. IEEE, 2015.

European Commission. The Clean Industrial Deal: A Joint Roadmap for Competitiveness and Decarbonisation. https://commission.europa.eu/topics/eu-competitiveness/clean-industrial-deal_en, 2025. Accessed: 2025-03-04.

European Parliament and Council. Regulation (EU) No 648/2012 of the European Parliament and of the Council of 4 July 2012 on OTC derivatives, central counterparties and trade repositories. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32012R0648, 2012. Accessed: 2025-01-30.

A. Eydeland and K. Wolyniec. *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging.*, volume 97. John Wiley & Sons, 2002.

P. Feldhütter and S. M Schaefer. The myth of the credit spread puzzle. *The Review of Financial Studies*, 31(8):2897–2942, 2018.

Fitch Ratings. Rating Definitions. https://www.fitchratings.com/research/fund-asset-managers/rating-definitions-24-04-2023, 2023. Accessed: 2025-01-22.

P. R. Freeman. The secretary problem and its extensions: A review. *International Statistical Review/Revue Internationale de Statistique*, pages 189–206, 1983.

O. Furtuna, A. Grassi, A. Ianiro, K. Kallage, R. Koci, F. Lenoci, A. Sowiński, and F. Vacirca. Financial stability risks from energy derivatives markets. *Financial Stability Review*, 2, 2022.

J. Gregory. *Counterparty credit risk and credit value adjustment: A continuing challenge for global financial markets.* John Wiley Sons, 2007.

F. B. Haspolat. Analysis of Moody's sovereign credit ratings: criticisms towards rating agencies are still valid? *Procedia Economics and Finance*, 30:283–293, 2015.

D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society*, pages 77–105, 1992.

J. C. Hull. *Options, futures, and other derivatives.* 6th ed. Upper Saddle River, NJ: Pearson Prentice Hall, 2006.

R. A. Jarrow and S. M. Turnbull. Pricing derivatives on financial securities subject to credit risk. *The Journal of Finance*, 50(1):53–85, 1995.

R. A. Jarrow, D. Lando, and S. M. Turnbull. A Markov model for the term structure of credit risk spreads. *The Review of Financial Studies*, 10(2):481–523, 1997.

Y. A. Katz and N. V. Shokhirev. Default risk modeling beyond the first-passage approximation: Extended Black-Cox model. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 82(1), 2010.

R. L. Keeney. Decision analysis: an overview. *Operations Research*, 30(5):803–838, 1982.

J. Kettunen, A. Salo, and D. W. Bunn. Optimization of electricity retailer's contract portfolio subject to risk preferences. *IEEE Transactions on Power Systems*, 25(1): 117–128, 2009.

S. Koekebakker and F. Ollmar. Forward curve dynamics in the Nordic electricity market. *Managerial Finance*, 31(6):73–94, 2005.

T. Krabichler and J. Teichmann. The Jarrow and Turnbull setting revisited. *International Journal of Theoretical and Applied Finance*, 27:03n04, 2024.

J. Li, C. Xu, B. Feng, and H. Zhao. Credit risk prediction model for listed companies based on CNN-LSTM and Attention mechanism. *Electronics*, 12(7), 2023.

Á. Lorca and J. Prina. Power portfolio optimization considering locational electricity prices and risk management. *Electric Power Systems Research*, 109:80–89, 2014.

J. J. Lucia and E. S. Schwartz. Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5:5–50, 2002.

J. F. Magee. *Decision Trees for Decision Making*. Brighton, MA, USA: Harvard Business Review, 1964.

T. Matsakos and S. Nield. Quantum Monte Carlo simulations for financial risk analytics: scenario generation for equity, rate, and credit risk factors. *Quantum*, 8: 1306, 2024.

H. Mausser and D. Rosen. Applying scenario optimization to portfolio credit risk. *The Journal of Risk Finance*, 2(2):36–48, 2001.

A. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management*. Princeton University Press, 2011.

R. C. Merton. On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29(2):449–470, 1974.

P. Millossovich. An extension of the Jarrow-Lando-Turnbull model to random recovery rate. *Dipartimento di Matematica Applicata Alle Scienze Economiche (Universita degli Studi di Trieste)*, 2002.

Moody's Investors Service. Rating Symbols and Definitions. https://ratings.moodys.com/rmc-documents/53954, 2025. Accessed: 2025-01-22.

H. Mori and Y. Umezawa. Credit risk evaluation in power market with random forest. In *2007 IEEE International Conference on Systems, Man and Cybernetics*, pages 3737–3742. IEEE, 2007.

N. Mosquera, J. Reneses, and E. F. Sánchez-Úbeda. Medium-term risk analysis in electricity markets: a decision-tree approach. *International Journal of Energy Sector Management*, 2(3):318–339, 2008.

A. Navada, A. N. Ansari, S. Patil, and Sonkamble B. A. Overview of use of decision tree algorithms in machine learning. In *2011 IEEE Control and System Graduate Research Colloquium*, pages 37–42. IEEE, 2011.

A. Ng and M. A. S. Mohamed. Credit Ratings Controversy: A Review. *Journal of Insurance and Financial Management*, 4(3):49–64, 2021.

Pexapark. What is a PPA? The Guide to Power Purchase Agreements. https://pexapark.com/solar-power-purchase-agreement-ppa/. Accessed: 2025-04-30.

P. Rocha and D. Kuhn. Multistage stochastic portfolio optimisation in deregulated electricity markets using linear decision rules. *European Journal of Operational Research*, 216(2):397–408, 2012.

L. Rokach and O. Maimon. Decision Trees. *Data Mining and Knowledge Discovery Handbook*, pages 165–192, 2005.

E. Schlögl and L. Schlögl. Duffie–Singleton Model. *Encyclopedia of Quantitative Finance*, 2010.

E. Schwartz and J. E. Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000.

S. Seppälä and S. Syri. Wind, PV, and Hybrid Power Plant Operation in Competitive Nordic Electricity Market With High Profit Cannibalization. *International Journal of Energy Research*, 2025(1):Article 8850556, 2025.

J. Skoglund and W. Chen. Rating momentum in the macroeconomic stress testing and scenario analysis of credit risk. *Journal of Risk Model Validation*, 11(1):1–26, 2016.

M. Somvanshi, P. Chavan, S. Tambade, and S. V. Shinde. A review of machine learning techniques using decision tree and support vector machine. In *2016 International Conference on Computing Communication Control and Automation (ICCUBEA)*, pages 1–7. IEEE, 2016.

S&P Global Ratings. Understanding Credit Ratings. https://www.spglobal.com/ratings/en/about/understanding-credit-ratings. Accessed: 2025-01-22.

Statista. Electricity consumption worldwide from 2000 to 2022, with a forecast for 2030 and 2050, by scenario. https://www.statista.com/statistics/1426308/electricity-consumption-worldwide-forecast-by-scenario/, 2023. Accessed: 2025-03-14.

Statista. Net electricity consumption worldwide in select years from 1980 to 2023. https://www.statista.com/statistics/280704/world-power-consumption/#statisticContainer, 2025. Accessed: 2025-03-14.

A. Takahashi, T. Kobayashi, and N. Nakagawa. Pricing convertible bonds with default risk: A Duffie-Singleton approach. *Journal of Fixed Income*, 11(3):20–29, 2001.

Y. T. Trinh and B. Hanzon. Option pricing and CVA calculations using the Monte Carlo-Tree (MC-Tree) method. *arXiv preprint arXiv:2202.00785*, 2022.

R. Wang and J. F. Ziegel. Scenario-based risk evaluation. *Finance and Stochastics*, 25(4):725–756, 2021.

H. C. Wu and Y. T. Wu. Evaluating credit rating prediction by using the KMV model and random forest. *Kybernetes*, 45(10):1637–1651, 2016.

T. Xiao. An Accurate Solution for Credit Value Adjustment (CVA) and Wrong Way Risk. *Journal of Fixed Income*, 25(1):84–95, 2015.

C. C. Yeh, F. Lin, and C. Y. Hsu. A hybrid KMV model, random forests and rough set theory approach for credit rating. *Knowledge-Based Systems*, 33:166–172, 2012.

H. Zhong, Z. Tan, Y. He, L. Xie, and C. Kang. Implications of COVID-19 for the electricity industry: A comprehensive review. *CSEE Journal of Power and Energy Systems*, 6(3):489–495, 2020.

S. H. Zhu and M. Pykhtin. A guide to modeling counterparty credit risk. *GARP Risk Review*, 37:16–22, 2007.