Layout optimization for offshore windfarm selling on the day-ahead market

GAUTHIER Matthis

RWE, Delft University of Technology

TU Delft

Layout optimization for offshore windfarm selling on the day-ahead market

by

GAUTHIER Matthis

to obtain the degree of Master of Science in Sustainable Energy Technology
at the Delft University of Technology,
in collaboration with RWE Offshore Wind,
to be defended publicly on Wednesday October 14, 2025.

Student number: 6084877
Project duration: December 30, 2024 – September 30, 2025
Supervisors: B. Nguyen, RWE supervisor
J. Iori, TU Delft supervisor
M. Baricchio, TU Delft supervisor

Wind Energy Group, Faculty of Aerospace Engineering, Delft University of Technology.

TU Delft

Preface

This thesis is submitted in fulfillment of the requirements for obtaining the Master of Science degree in Sustainable Energy Technology at Delft University of Technology. The work has been conducted in collaboration with RWE, inside the resource assessment team and focuses on the economic optimization of offshore wind farm layouts operating in the day-ahead electricity market.

The motivation behind this research lies in the growing importance of integrating renewable energy sources into electricity markets. As the penetration of wind energy increases, market participants face new challenges related to price volatility, cannibalization effects, and the need for revenue-oriented project design. The societal relevance of this research is also substantial. Indeed, improving the alignment between wind energy production and electricity demand can help address one of the main limitations of wind power today, its intermittency and lack of availability when demand is highest.

The study combines methods from wind resource assessment, optimization, and energy market modeling. It aims to provide insights not only for academic research but also for wind farm developers, policymakers, and investors seeking to enhance decision-making under market uncertainty.

I would like to express my sincere gratitude to everyone who supported me during the course of this thesis project. I am especially thankful to **Bruno Nguyen** for his near-daily supervision and guidance, and for making this project both challenging and deeply engaging. I am also grateful to my TU Delft supervisors, **Jenna Iori** for her dedication and continuous support in making this thesis an enjoyable and rewarding experience, and **Matteo Baricchio** for his enthusiasm and insightful discussions throughout the development of this work.

My sincere thanks go to all members of the **RWE French Offshore Wind Office** for welcoming me warmly from day one and making me feel immediately part of the team. I would also like to thank the **Resource Assessment Team** for their role in enabling this research and for the opportunity to present my work at the Wind Energy Science Conference in Nantes. Special thanks to **Abhinav Kapila** for the many calls and support during the optimization runs, and to **Sam Williams** for his constant involvement in ensuring that our experience at RWE was as rewarding as possible.

I would further like to acknowledge **Mihir Mehta** for his valuable insights on his market price scenario model, and **Julian Quick** for his help in the parameterization of the Stochastic Gradient Descent algorithm.

Finally, I extend my deepest gratitude to my loved ones, who have continuously believed in me and allowed me to pursue this work under the best possible conditions. Special thanks to my family, my close friends, and Luna. Their unwavering support and encouragement have made this achievement possible.

Paris, September 2025
Matthis Gauthier

Summary

This thesis investigates whether and how optimizing offshore wind farm layouts for expected day-ahead market revenue rather than the conventional Annual Energy Production (AEP) objective affects technical and economic performance. Using historical hourly wind and day-ahead price data, wake and power models, and a layout optimization framework, the study compares layouts optimized for AEP with layouts optimized for expected revenue across several locations and three case studies.

Key findings show that revenue-optimized layouts deliver small but consistent improvements in projected revenues while keeping reasonable AEP performance. Site-dependent examples illustrate this trade-off: at one site, the revenue-optimized layout increases revenue by 0.22% while reducing AEP by 0.44%, and at another site, revenue increases by 0.20% with a negligible change in AEP (0.03%). Layouts optimized on an earlier period (2015–2020) retain performance when evaluated on later data (2022–2024), with objective-function losses remaining below 0.12%. Scenario experiments further show that long-term market structure predictions matter: changing the assumed correlation between wind speed and electricity prices in future scenarios can impact revenue substantially.

Overall, the study concludes that incorporating price information into layout optimization yields economic benefits and enhances robustness to realistic market variability. The results recommend that developers consider revenue-based objectives alongside traditional energy metrics when designing offshore wind farms operating in merchant electricity markets.

Contents

List of Figures

List of Tables

1 Introduction

The global shift toward renewable energy sources has intensified in recent years due to climate change concerns and the depletion of fossil fuels. Offshore wind farms play a critical role in transitioning towards sustainable energy systems due to their high capacity and reliability. Compared to onshore wind sources, offshore installations generally experience stronger and more consistent wind speeds, allowing for greater energy output. Additionally, the strategic placement of these farms in maritime zones reduces land use conflicts and minimizes the visual and noise impacts often associated with terrestrial wind turbine generators (WTGs). As a result, offshore wind energy is developing rapidly, with an additional 11 GW of capacity installed in 2023, taking the total installed capacity worldwide to more than 75 GW [1].

1.1. Motivation and scope

The energy transition intensifies the need to integrate large shares of renewable generation into liberalized electricity markets. Offshore wind in particular plays an increasingly central role in decarbonizing electricity systems. This trend raises new technical and economic questions: how should wind farms be designed to maximize value in electricity markets with hourly variability, and exposed to growing wind energy penetration?

A key market mechanism motivating this work is *cannibalization*: as renewable generation increases, periods with high wind output can depress spot prices, reducing the market value of additional wind generation during those periods. This effect creates a mismatch between maximum energy output and maximum revenue, motivating a revenue-aware design approach [2, 3]. A concise introduction to the cannibalization concept is given here; a fuller description appears later in the market trends section.

The design aspect addressed by this work is the optimization of the layout of the wind farm. Optimization is the mathematical process of selecting the best configuration of design variables to minimize or maximize an objective function subject to constraints. In the broadest sense, an optimization problem can be stated as:

$$\min_{x \in \mathcal{X}} \; f(x) \quad \text{s.t.} \quad g_i(x) \leq 0, \; h_j(x) = 0,$$

where x denotes the decision variables, f the objective function, and g_i, h_j the constraints. In the wind-farm context, the decision vector x typically collects turbine coordinates (and possibly control settings), f might be energy production or revenue, and the constraints capture spacing and boundary limits.

The schematic in Figure 1.1 visually represents the fundamental goal of optimization: finding the point (or set of points) that minimizes an objective function. In convex problems, there is a single global minimum that is easy to find. However, in non-convex problems like wind farm layout optimization, multiple local minima may exist. The figure shows how an optimizer starting from different initial positions may become trapped in a local minimum rather than reaching the global one. This underlines the importance of algorithm choice and initialization strategy, particularly when dealing with complex, nonlinear design problems such as offshore wind farm layout optimization.

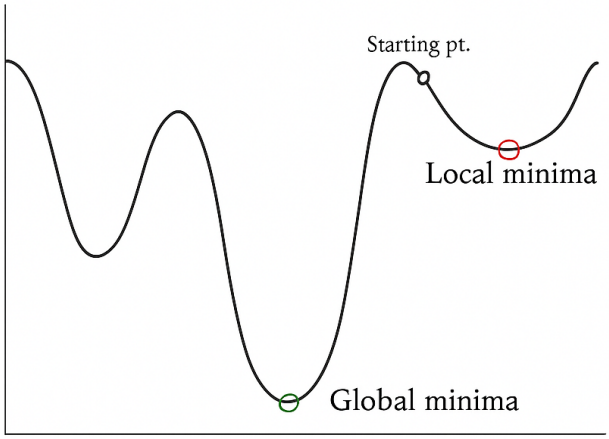

Figure 1.1: Schematic illustration of a generic optimization problem

1.2. The wind farm layout optimization problem

Wake losses are due to the extraction of energy from the flow to activate the blades, decreasing downstream velocity, and increasing turbulence intensity, reducing the efficiency of downstream WTG. Optimizing offshore wind farm layouts is paramount to maximizing the energy output by strategically positioning WTG to minimize these wake losses [4].

1.2.1. Standard objective functions

A common practice in the industry is to optimize the layout of offshore wind farms by maximizing the annual energy production (AEP), particularly for farms selling electricity at fixed prices, as this equates to optimizing revenues.
For instance, Pérez et al. [5] introduce such a methodology. They model wind distribution using a wind rose and discretize it into directional bins to define the frequency of each wind speed and direction. They set the coordinates of each WTG as optimization variables and define constraints of minimum spacing between WTGs and boundaries of the wind farm area. Their optimization is applied to a test-case study and increases the annual energy production by 3.5 % compared to the base layout. Similarly, Gebraad et al. [6] optimize the AEP of a wind plant by coupling a yaw control system optimization and a layout optimization. In their model, WTG coordinates remained optimization variables, but they also incorporated WTG yaw angles as additional variables. When considering the case study of the redesign of an offshore wind plant, this resulted in a significant increase in AEP of 5. 3 % compared to the baseline layout.

However, optimizing layouts using AEP as the objective function is limited because it does not account for the costs associated with WTG, their installation, and their operation. Consequently, it is not the most relevant metric for maximizing revenue from a developer's perspective. A more suitable metric is the Levelized Cost Of Energy (LCOE). This is especially the case with more complex site conditions, with variations in bathymetry or ground types, and with costs of WTG implementations varying in the design space.

A commonly used calculation of the LCOE is:

$$\text{LCOE} = \frac{\sum_{t=0}^{T} \frac{I_t + O_t + F_t}{(1+r)^t}}{\sum_{t=0}^{T} \frac{E_t}{(1+r)^t}}, \tag{1.1}$$

where I_t are investment expenditures, O_t operating expenditures, F_t fuel/variable costs (if any), E_t the energy produced in year t, r the discount rate, and T the project lifetime.

For example, Stanley et al. [7] analyze the impact of different objective functions and optimization algorithms on wind farm performance. Their research shows that when the number of WTGs is not fixed, optimizing for AEP can lead to significant profit losses. This is due to wake effects, where the addition of extra WTGs reduces the efficiency of others by lowering the electricity output per WTG.
Furthermore, Gualtieri [8] compares LCOE, capacity factor, and AEP as objective functions for wind farm layout optimization. By analyzing a wide range of commercially available WTG and optimizing layouts for different configurations, the study concludes that AEP optimization alone does not lead to a lower LCOE, which is often found to be the primary goal. However, it was discovered that, in most wind scenarios, optimizing for capacity factor produces results comparable to LCOE-based optimization.
Another scenario involving a non-homogeneous wind farm (different WTG models in the farm) was studied by Ziyaei et al. [9], where two different types of WTG and a variable number of WTGs were available. The study applied both a single-objective algorithm (minimizing the cost of energy) and a multi-objective algorithm (balancing cost minimization and power output maximization). The wind farm layout was discretely modeled, using a grid of possible WTG positions. The results showed that layouts optimized by both algorithms were similar, indicating no clear advantage of using a multi-objective approach for this problem.

1.2.2. Optimizing for revenues

Although LCOE, AEP, and capacity factor are the predominant objective functions in the industry as of today, the rate of wind energy development and the integration perspective of numerous wind farms into the day-ahead market raise some concerns regarding their efficiency in determining the most profitable wind farm layout for developers, should they sell on the day-ahead market.

As offshore wind technology matures, improvements in turbine efficiency, farm design, and operational strategies continue to enhance the overall performance of projects. According to Cox [10], LCoE for offshore wind halved between 2016 and 2020, and this downward trend is expected to continue, with BloombergNEF forecasting a further 22% reduction by 2035 [11]. This increasing cost-efficiency has prompted governments to reconsider subsidy schemes, leading offshore wind farms to enter competitive electricity markets such as the day-ahead market [10].

Several offshore wind farm tenders have already been issued without any subsidy mechanisms. The first such tender in the Netherlands was launched in 2017 [12]. Although uncertainties led some companies to refrain from bidding in this round [13], the trend has persisted. For example, Germany awarded nearly 1 GW of zero-subsidy licenses for projects scheduled to be operational in 2026 [14].

As the industry expands and the conditions of the electricity sale change, it becomes crucial to explore how incorporating revenues in electricity markets as an objective function can influence the optimization of wind farm layouts. Indeed, Loth et al. [15] and Dykes [16] present the limit of the LCOE as a driving objective function for renewable energies and introduce the Cost of Valued Energy (COVE) as an alternative for renewable energy developers. It is explained that LCOE does not take into account the variability of hourly electricity prices when COVE weights electricity production with the spot market price. Here, the spot market refers to the actual selling price of electricity, taking into account day-ahead markets and intraday trading markets. COVE is calculated as:

$$COVE = \frac{\text{Annual Costs}}{\int pG\,dt}, \tag{1.2}$$

where $p = \frac{P}{P_{avg}}$ is the actual hourly spot selling price divided by the annual averaged electricity price, G is the hourly energy generation, and Annual Costs are all project-related expenditures distributed

evenly per year over the plant lifetime.

This is, in particular, essential in a high renewable energy penetration market (above 5% according to Loth et al. [15]), where high generation of energy by renewable sources can lead to low and even negative prices, considerably reducing the profit for developers. Test cases have demonstrated that minimizing LCOE can significantly increase COVE compared to COVE-optimized designs.

In their layout optimization for revenue, Stanley et al. [7] consider a fixed electricity price instead of a variable hourly spot price data, to model the situation of a wind farm under a feed-in-tariff (FiT), while taking into account a variable number of WTG. One of this study's results is that layouts optimized for annual profit as an objective function tend to have slightly higher LCOE than LCOE-optimized ones but manage to yield significantly more energy, resulting in considerable improvements in revenues. This FiT has been one of the main electricity selling mechanisms for offshore wind farms, along with Contracts for Difference (CfDs), both allowing them to sell their electricity at constant rates instead of on the spot market. This explains the standard consideration of fixed electricity prices for layout optimization in the literature. But the increasing share of renewable energies makes these contracts harder to secure with more recent projects, as is explained in the next Section.

Optimizing WTG design for revenue in the day-ahead market has been researched by Mehta et al. [17]. A simple market model is used to easily generate several market scenarios and several metrics including LCOE, Net Present Value (NPV), Profitability Index (PI), Modified Internal Rate of Return (MIRR), and COVE are compared. This study concludes that the benefits raised by designing WTGs for profits remain marginal compared to the LCOE-optimized designs for most market scenarios. However, they underline that the difference is more significant for very low prices market scenarios.

To the author's best knowledge, wind farm layout optimization considering revenues in the day-ahead market has not been extensively studied in the literature. However, Zhang & Jiang [18] explored this aspect by proposing a joint optimization of the number, type, and layout of WTG for an offshore wind farm participating in the day-ahead market. A bi-level optimization model is proposed, where the upper-level model focuses on investment optimization, maximizing profit by selecting the optimal number, type and location of WTG, while the lower-level model represents the day-ahead electricity market clearing. This electricity market model includes conventional generating units with fixed generation capacity and offering prices. Regarding the selection of the number and type of WTG, results show their crucial role in maximizing profit, finding a difference of profit of 38% between the best and worst combination of type and number of WTG. However, the analysis does not include a comparison between the optimal layout found and a reference layout, optimized for AEP, for example, or profit in a fixed-price market.

Recently, Nguyen et al. [4] propose a layout optimization method to account for wind farm participation in the day-ahead electricity market and in the secondary reserve markets. They use the market data from the year 2023 and build an inner optimization algorithm to determine the power sold on the day-ahead market and on the secondary market. It is found that, compared to the base layout of the farm, the optimized layout for profit in day-ahead and reserve markets features wider spacing between rows of WTGs, while maintaining a similar outer boundary layout (i.e., WTGs placed along the farm perimeter). However, WTGs within each row are positioned more closely together, resulting in denser rows but increased inter-row spacing. Performance-wise, the layouts optimized for combined day-ahead and reserve market profit outperformed AEP-optimized layouts, both in profit and in AEP. The conclusions are the same when the different layouts optimized for 2023 are faced with 2024 wind and market data. While this research shows the potential advantages of optimizing layouts for revenue with wind farms participating in electricity markets, it does not directly compare day-ahead-only optimized layouts with AEP-optimized ones. Comparisons between AEP and profit are made on the joint day-ahead and reserve market. However, profits are displayed on the day-ahead market for both layout types (day-ahead-only and joint market optimization), showing that layouts optimized for both markets also perform strongly in the day-ahead market alone. This supports the idea that similar improvements from AEP optimizations could be found when considering the day-ahead market optimization only.

1.2.3. Optimization algorithms

Wind farm layout optimization is a particularly difficult problem for several reasons. First, the problems are non-convex, meaning that multiple layouts can yield local optima, making it hard to guarantee a global optimum. This is displayed in Figure 1.1, where a simple algorithm following the gradient decrease starting in the initial given point wouldn't find the global minima. Second, the increased size of today's offshore wind farms also increases the number of design variables (each WTG's position, and sometimes also their control settings). Third, evaluating candidate layouts is computationally expensive because wake-loss calculations require modeling the interactions among all the WTGs. Researchers have developed and utilized various optimization algorithms to address this challenge, each with its own strengths and limitations. Early approaches relied on mathematical programming techniques like gradient-based methods, which use derivatives of the objective function to guide the search for an optimum. While these often perform well, they require objective functions to be continuous, which can not be the case depending on the wake models. More recent methods incorporate data-driven algorithms and metaheuristic algorithms. Metaheuristics, such as genetic algorithms or particle swarm optimization, are derivative-free and flexible, making them well-suited for the non-convex nature of wind farm layout problems. They are efficient in avoiding local optima and handling discontinuous objective functions, but they can also come with slower convergence.

Stanley et al. [7] analyze different optimization algorithms for wind farm layout optimization, considering a variable number of WTGs. They compare the commonly used greedy algorithm, genetic algorithm, and gradient-based methods to a novel repeated-sweep algorithm. Greedy algorithms work by selecting the best local evolution option at each step, without requiring any additional calculations. They work fast but can lack precision because of their non-consideration for history or further trends. Genetic algorithms work by evolving a population of potential solutions over multiple generations. They select at each step the best solutions, and combine or modify them to recreate new sets of potential solutions. They are commonly used for non-linear and complex problems. Findings show that greedy algorithms are computationally efficient but struggle when increasing the size of wind farms, while genetic algorithms perform well but require high computational resources. The repeated-sweep algorithm, which works with a gridded potential positions layout where WTGs are randomly assigned, removed, or swapped to work, is found to be an effective alternative. It indeed balances well computational cost and solution quality, particularly for large wind farms. The gradient-based algorithm was only studied for a low-complexity wind farm and was found to be efficient, although computationally expensive. The study also highlights that algorithm choice depends on the problem's complexity and the objective functions considered (with AEP, LCOE, and profit being studied in the article).

Similarly, Kunakote et al. [19] conduct a comparative study of twelve metaheuristic algorithms for wind farm layout optimization, evaluating their efficiency in minimizing the LCOE. The study evaluates these algorithms across four test cases, with a combination of variable and fixed numbers of WTG and different wake models. The results highlight that, among all the studied algorithms, moth-flame optimization performs best with fixed WTG numbers, while ant colony optimization for continuous domains excels when the number of WTGs is variable. Other algorithms like particle swarm optimization and differential evolution are also found to be efficient, but can present some issues of often being trapped at a local optimum. It is worth noting, however, that this benchmark focuses exclusively on metaheuristic solvers. Gradient-based methods, which can sometimes offer faster convergence when the objective function is continuous and differentiable, are not included in the comparison.

Other articles in the literature study specific algorithms and their performance for particular test cases. Feng & Shen [20] develop a random search algorithm for wind farm layout, and show that it can help improve layouts, even if already optimized, namely by helping them escape local minima. They also study its robustness, showing constant improvement from base layouts, even with unexpected wind conditions. Long et al. [21] build a data-driven evolutionary algorithm that uses an adaptive differential evolution algorithm to easily evaluate the performance of generated layouts, to avoid computing their wake losses and objective function, and a general regression neural network to filter the layouts to evaluate, ruling out the least efficient ones and saving computational time. This methodology has proved very efficient compared to other benchmark algorithms, especially when considering high-complexity layouts, where the data-driven algorithm's computational time outperforms the other algorithms. Quick et al. [22] also develop an efficient algorithm for high-complexity wind farms. They present a stochastic gradient descent algorithm that uses Monte Carlo simulation to maximize the AEP, which proves to be

faster and to produce better results than the deterministic approaches to which it was compared. Recent work by Baricchio et al. [23] applies a genetic algorithm custom-made for layout optimization, that co-optimizes layout and yaw control. They find that the co-design yields AEP gains between 0.3–0.6%, indicating that improved control integration can enhance layout performance.

Finally, Thomas et al. [24] present a comprehensive evaluation of various optimization techniques applied to a complex wind farm layout problem with the goal of AEP maximization. The methodology involves comparing eight different optimization methods, including gradient-based, gradient-free, and hybrid approaches. Each method was managed by researchers experienced with the respective algorithms to ensure fair comparisons. The results indicate that all methods performed similarly, achieving optimized wake loss values between 15.48% and 15.70%, compared to 17.28% for the unoptimized layout. The optimized layouts consistently showed tightly packed WTGs along the outer boundaries and more spaced WTGs in internal regions. This is shown in Figure 1.2, where an unoptimized layout and an optimized one are displayed for comparison. The discrete exploration-based optimization, a novel gradient-free method, found the layout with the highest AEP among the eight methods. This highlights the effectiveness of both gradient-based and gradient-free methods in improving wind farm performance, and it is found that the main differences lie in computational efficiency.

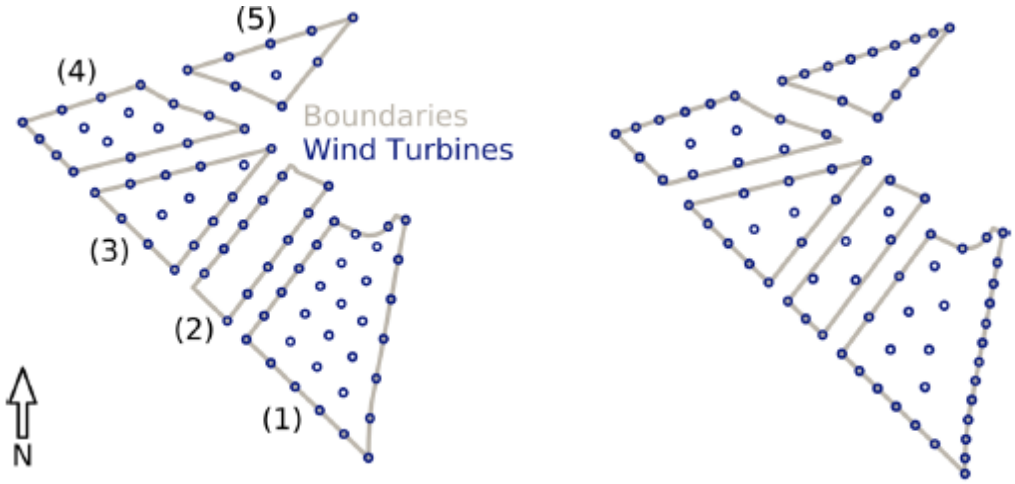

(a) Base layout provided

(b) Layout optimized by discrete exploration-based optimization

Figure 1.2: Comparison of an unoptimized layout and an optimized layout achieving the highest AEP for this case study. Reproduced from Thomas et al. [24]: ”A comparison of eight optimization methods applied to a wind farm layout optimization problem”.

In summary, no single solver is found to be universally superior. The study from Thomas et al. [24] demonstrates the wide variety of optimization strategies used across the literature and shows that, when carefully applied, these diverse methods often produce similar improvements in AEP and comparable turbine-placement patterns on the same case study. The broader literature indicates that both gradient-based and gradient-free approaches can be effective, and algorithm choice is usually driven by problem specifics. For example, gradient-based methods tend to be faster when the objective function is smooth and differentiable, and when derivatives can be obtained efficiently. In contrast, gradient-free algorithms are more suited for discontinuous objectives [24]. Kunakote et al. [19] details the performance comparison of twelve gradient-free algorithms, showing that their effectiveness varies across different case studies. This underscores that case dependence is not only a matter of choosing between gradient-based and gradient-free algorithms, but also applies to the selection of specific algorithms within each family. However, several algorithms like the random selection, the stochastic gradient descent, and the genetic algorithm are consistently found in recent works as effective strategies. These are among the most commonly used in the literature across various scenarios.

1.3. Electricity market forecasts

Beyond technical optimization, the financial viability of offshore wind farms depends heavily on electricity market conditions. To avoid exposing renewable energy developers to variable prices, subsidy mechanisms were implemented; most notably, CfDs and FiTs have been widely used. FiTs guarantee that producers sell their electricity at a fixed, determined price, while double-sided CfDs set a strike price and allow electricity to be sold on the day-ahead market. If the market price falls below the strike price, the producer is compensated for the difference, but if the market price exceeds the strike price, the producer must reimburse the surplus. However, as the share of renewable energy in the mix continues to grow, maintaining such subsidy schemes is more and more challenged by state authorities, and so exposure to fluctuating market prices poses greater financial uncertainty for wind farm developers. With this new paradigm, electricity price forecasting becomes primordial to ensure the profitability of offshore wind farm projects that would be exposed to the electricity markets, as being able to match high electricity price periods with higher production becomes more crucial.

1.3.1. Market trends prediction

The transition to a renewable-based electricity system presents significant pricing challenges. The increasing penetration of renewable energy in the electricity mix has significantly changed spot price dynamics, with wind energy playing a central role in this transformation.

Paraschiv et al. [25] analyze the impact of renewable energy penetration and renewable-enhancing policies on the German EEX day-ahead electricity market, showing that wind and solar energy lower market prices due to the merit order effect. These low prices hurt the traditional electricity producers, because they reduce their margins, and they don't anymore profit from being at the end of the merit order curve.

Similarly, Djørup et al. [26] analyze the effects of a 100% renewable mix with wind energy dominance on electricity markets, particularly in Denmark, where it is the 2050 target, demonstrating that current market structures may not be financially sustainable for high shares of wind power. Indeed, as the share of renewable energy in the electricity mix grows, their typically low bidding prices in day-ahead markets drive spot prices down and reduce profit margins for developers. It is proposed to either redesign the electricity market or to keep subsidies like feed-in tariffs in place. Blickwedel et al. [27] assess the general economic viability of non-subsidized wind farms in Germany, highlighting how fluctuating market prices and uncertain revenue streams may discourage investment in new capacity. Their forecasting model also suggests that without revenue mechanisms such as power purchase agreements or long-term market contracts, profitability remains uncertain. Overall, the role of policy and market-based instruments is critical in shaping future revenue models for renewable energy. Busch et al.[28] emphasize the need for these revenue stabilization mechanisms. These measures are necessary to enhance renewable energy development by limiting the risks for the constructors.

On the other hand, Jansen et al. [29] focus on offshore wind energy and its competitiveness without subsidies, demonstrating that recent cost reductions have made offshore wind increasingly viable in

mature markets. Their study highlights that, in the future, wind farm developers may be able to develop fully merchant offshore wind farms without government subsidies. This would depend on future energy prices and integration costs. For example, in France, a country where subsidies are still the core of the profit for developers, the M0 mechanism [30] was introduced for some tenders to reward wind farms selling electricity in high electricity price hours. This permitted to secure the CfD implementation by encouraging the developers to design their wind farm considering price variations of the day-ahead electricity market. Other technological options have been explored. For example, Mauch et al. [31] explore the possibility of coupling a wind farm with a Compressed Air Energy Storage system (CAES). This would help a wind farm selling in the day-ahead market, as energy would be stored when high-wind speeds but low electricity prices are expected, and distributed back to the grid when it would be more economically profitable. Overall, the study finds that this kind of power plant would not be profitable given the actual costs of the technology, limiting its potential to support the profitability of renewables operating on the electricity markets.

However, a higher penetration of wind energy in the electricity mix can also damage its competitiveness by reducing its market value; this is the cannibalization phenomenon, introduced in Section 1.1. This occurs when increased wind generation capacity leads to lower electricity prices during periods of high wind output, thereby decreasing revenues for wind energy producers.

Indeed, Hirth [2] introduces the concept of the value factor to measure the relative price of variable renewable energies compared to the base price, finding that their market value falls significantly with higher penetration, with wind and solar value factors dropping to around 50-80% of the average power price at high penetration levels. Reichenberg et al. [3] develop an analytical framework to model the correlation between renewable energy production and electricity prices, finding that the cannibalization effect, where high wind energy output depresses market prices and reduces revenues for these wind farms, significantly reduces projected profits relative to investment costs.

According to these studies, the future of the electricity market remains uncertain, especially in the context of growing renewable penetration. But the intermittence of wind energy is predicted to cause and require some changes in policy and market design. In a context where the revenue mechanisms of wind farms are changing, it is necessary to rethink their design to ensure their profitability in a market influenced by variable prices.

1.3.2. Electricity prices forecasting

Long-term electricity price forecasting is critical for wind farm development in volatile, renewable-dominated energy markets. Yet, it presents significant challenges in balancing high temporal resolution with the need to capture evolving market trends. Researchers have developed various methods that reflect these challenges, each targeting different forecast horizons and employing diverse techniques.

As of today, most of the research conducted on Electricity Price Forecasting (EPF) is done on a short-term basis (from one day to the other), to help power plant operators bid in the day-ahead market. For instance, Sætherø's work on constructing hourly price forward curves (HPFCs) forms a foundational approach by merging historical spot prices with futures market data. His method, elaborated in his doctoral thesis [32] and further refined in the collaborative paper with Kiesel and Paraschiv [33], focuses on calibrating seasonal patterns and adjustment functions to produce detailed hourly forecasts. However, this approach is primarily geared toward short (days) to medium-term (several years) applications rather than extending explicitly to multi-decade horizons.

In contrast, Ziel and Steinert [34] introduce a probabilistic framework that simulates supply and demand curves to generate hourly forecasts over medium-term horizons, typically up to three years. Their method delivers probabilistic forecasts and quantifies uncertainty through prediction intervals, effectively capturing risks such as price spikes and periods of negative pricing. Although this represents a significant advancement beyond day-ahead forecasting, its reach is limited when considering the extended timeframes necessary for investment and wind farm development and operation.

De Marcos et al. [35] apply co-integration and vector error correction models to the Spanish electricity market for long-term forecasting of prices. The study addresses the limitations of traditional models, such as the Black-Scholes model, in capturing the complex dynamics of electricity prices over extended periods. Their method uses the relationships among electricity prices, fuel spot prices, and futures data

to produce forecasts typically covering horizons of one year or more. While theoretically robust, such econometric models may struggle to fully account for multi-decade structural changes in the market.

Broader reviews, such as Weron's survey of electricity price forecasting methods [36], provide valuable context by comparing traditional statistical and econometric approaches with modern computational intelligence techniques. Although not focused exclusively on long-term horizons, Weron's work underscores the persistent difficulties in achieving both high granularity and robust uncertainty quantification across different forecast periods.

Expanding the focus to longer-term horizons, Gabrielli et al. [37] propose a data-driven model based on Fourier analysis. By decomposing electricity price time series into their principal periodic components and residual fluctuations, the model predicts the main frequencies of the electricity price profile using regression techniques and incorporates annual values of energy and market quantities as inputs. This approach is designed to support investment planning over horizons of 10 years or more. It demonstrates the ability to generalize across different electricity markets and shows comparable performance to market-based models. The integration of data-driven and market-based approaches is highlighted as a promising direction for improving forecast accuracy and capturing market dynamics.

Additionally, Oh et al. [38] propose a hybrid scenario generation method that targets a 30-year forecast horizon. Their approach decomposes the system marginal price into annual, monthly, and daily components using oil price scenarios to drive the long-term annual trend, neural networks for capturing monthly variations, and statistical methods for daily fluctuations. This multi-scale strategy allows for the synthesis of detailed hourly scenarios over an extended period, although the complexity of integrating multiple modeling techniques can challenge reproducibility.

Finally, Gea-Bermúdez et al. [39] present a mathematical model for electricity price forecast up to 2050, considering high renewable energy penetration, using the market model Balmorel and focused on Denmark. The study emphasizes the importance of including generator unit commitment modeling in such analyses to reflect the realistic operational behavior of energy units, despite the increased computational complexity. The findings suggest that the impact of unit commitment depends on the energy system's configuration and that careful consideration of the modeling approach is crucial for operational planning in those systems.

Together, these studies illustrate that while most of the research on electricity price forecasting is done on short to medium-term forecasts, only a subset is explicitly designed to address the complexities of forecasting over horizons extending to 10 or even 30 years. However, some of these long-term forecasting techniques can be difficult to adapt in this study, due to their complexity.

1.4. Research question

Although several works consider market effects on wind farm performances, to the author's best knowledge, no study explicitly optimizes wind farm layout for revenue with a fixed WTG number under day-ahead only market forecasts. While Mehta et al. [17] consider market forecasts for wind turbine design, Stanley et al. [7] consider layout optimization with a fixed hourly electricity price. Zhang et al. [18] and Nguyen et al. [4] both perform layout optimization with hourly variable electricity prices, but they respectively consider a variable number of WTG and optimize for a joint day-ahead and reserve electricity market.

This gap motivates the work of this report. This study aims to assess how revenue-optimized wind farm layouts perform in comparison to those optimized for Annual Energy Production, and in the context of an uncertain electricity market. Specifically, this research aims to answer the following questions:

- **How does a revenue-optimized wind farm layout compare to AEP-optimized layouts in terms of financial returns and energy generation when participating in the day-ahead market?**
- **How do market conditions influence revenue-based layout optimization?**
- **How resilient are revenue-optimized layouts when subjected to market conditions and wind patterns different from those used in their initial optimization?**

To address these research questions, this study employs a methodology combining wind farm layout optimization, electricity market modeling, and scenario-based analysis. The optimization process uses the TopFarm framework, integrating wake modeling (TurbOPark) and stochastic gradient descent (SGD) algorithms to maximize either AEP or revenue, the latter incorporating variable day-ahead electricity prices. Three case studies are conducted:

- **Case study 1 — Baseline comparison.** Compare layouts optimized for AEP and for expected day-ahead revenue using historical wind and price data for each site; quantify differences in revenue and AEP.
- **Case study 2 — Out-of-sample robustness.** Optimize layouts on an earlier historical period and evaluate performance on later yer-to-be seen wind and price data to measure robustness and objective-function loss.
- **Case study 3 — Market-scenario sensitivity.** Generate parametric future price scenarios (using the parametric approach proposed by Mehta et al. [17], varying the Pearson correlation between wind speed and price) and assess how layouts optimized for a reference scenario perform across alternative market futures.

By addressing these questions, this study seeks to provide valuable insights into the future wind farm layout optimization. Indeed, results could help developers understand how exposure to hourly fluctuating market prices can impact the performance of the farms. It will also test the resilience of revenue-based layouts under unseen price scenarios, to assess how a switch in the objective function could improve future wind farms.

2

Methodology

This chapter presents the methodology used to optimize offshore wind farm layouts under various wind and market conditions. The approach includes data processing, modeling of wind power and revenue, and the application of optimization algorithms under defined constraints. The study is split into three case studies, each providing further information to evaluate the impact of optimizing layouts for revenues.

2.1. Overview

The numerical simulations are conducted within the TopFarm framework [40] that enables wind farm layout optimization. It allows for the use of a vast number of sites, wake models, objective functions, and optimization algorithms, offering flexibility and adaptability for advanced simulation needs. TopFarm integrates external libraries such as PyWake, for modeling wake effects, and OpenMDAO, for optimization routines and multidisciplinary analysis.

2.2. Models and Data

2.2.1. Site Characterization and Modeling Assumptions

Four offshore wind-farm projects are selected as case studies for the layout optimizations. These are not hypothetical designs, but existing or in-development projects chosen to represent a diversity of geographical and technical conditions. Figure 2.1 shows their locations.

Figure 2.1: Map of the selected wind-farm sites: Méditerranée II (France), Borssele III&IV (Netherlands), Nordsee Ost (Germany), and Thor (Denmark).

Table 2.1: Characteristics of the original wind-farm sites

Site	Bidding Zone / Country	Location	Number of WTGs	WTG Type
Thor	DK1 / Denmark	22 km off the Danish west coast	72	15 MW
Nordsee Ost	DE / Germany	57 km north of Heligoland	48	6.15 MW
Borssele III & IV	NL / Netherlands	22 km offshore Zeeland	77	9.5 MW
Méditerranée II	FR / France	22 km offshore Gulf of Fos (Western Mediterranean), floating project,	Unknown	Unknown

Each site with its own specificity, as summarized in Table 2.1. The Borssele III & IV project, originally comprising 77 WTGs and multiple inclusion zones, is used as the principal reference case because of its spatial complexity and frequent use in previous layout-optimization studies [24].

To help with computational time and problem description, several simplifications are introduced. Internal exclusion zones are removed, allowing each site to be treated as a single, continuous buildable area. This avoids WTG subdomain allocation issues and allows the algorithm to focus purely on optimizing the WTGs' locations. Then the original site is scaled down from originally 77 WTGs to 30, which speeds up the computation time and allows optimization runs to finish in 8 to 10 hours. To be representative of existing wind farms and to avoid the spread of WTGs, the furthest away from one another without accounting for wake effects, the boundaries of the site retrieved from Thomas et al. [24] are also scaled down. The aim is to achieve a power density of about 10 MW/km^2, which is reasonable for modern offshore farms. The resulting scaled-down site is shown in Figure 2.2.

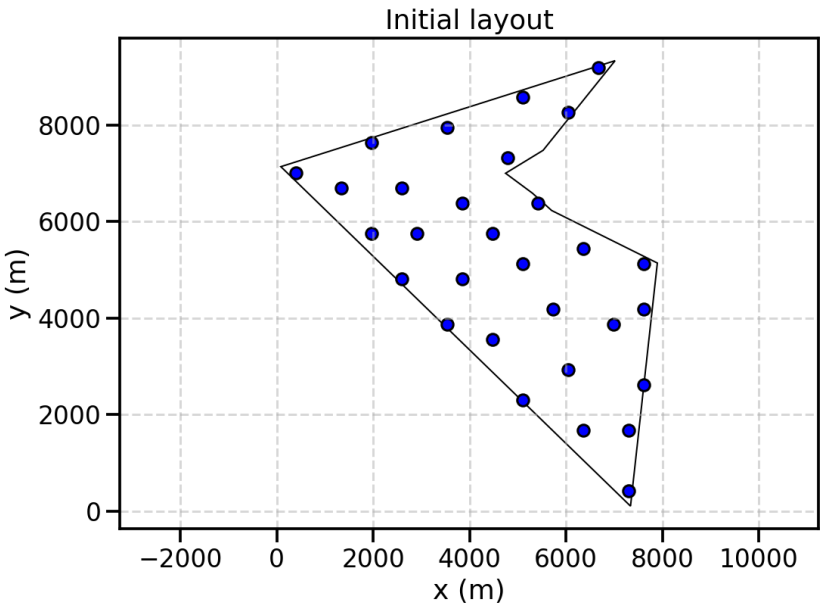

Figure 2.2: Scaled-down “mini” Borssele III & IV site used for layout optimization.

In all simulations, 10 MW IEA reference WTGs [41] originating from IEA Task 37 are used, with a hub height of 119 m, and a rotor diameter of 198 m. This reference WTG is selected due to its frequent use in the literature, and its power and thrust curve can be seen in Figure 2.3.

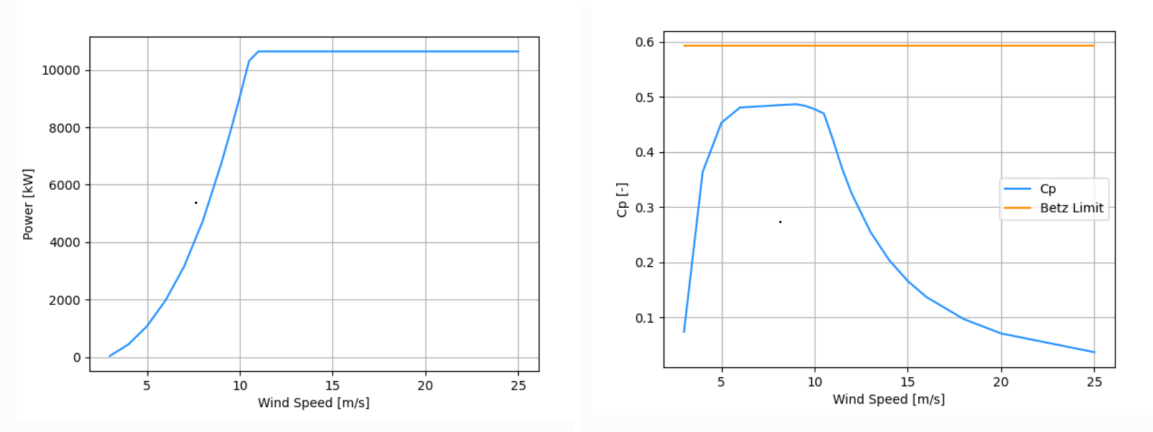

Figure 2.3: Power curve (left) and power coefficient (Cp) curve (right) of IEA 10 MW reference WTG taken from [42]

Wind resources are assumed to be spatially uniform across the site. This means that wind speed distributions, directional statistics, and turbulence levels are homogeneous, with wake losses being the only source of spatial variability.

2.2.2. Wind and Electricity Price Data

Historical wind and electricity price data are collected for the longest available periods in order to capture a broad range of values and ensure robust modeling of both wind resource and market conditions.

Wind data are obtained from the EMD-WRF database [43], which provides long-term mesoscale data of wind conditions across Europe. EMD-WRF is generated using the Weather Research and Forecasting (WRF) model, which assimilates historical meteorological observations into a numerical weather model. This process produces high-resolution, hourly wind fields that are suitable for long-term wind energy assessments. The dataset provides, in particular, wind speeds at a reference height of 100 m above ground level for each site, covering the period 1999 to 2025. Wind speeds are later extrapolated to the

WTG hub height of 119 m using the empirical power law:

$$U(z) = U(z_{\text{ref}}) \cdot \left(\frac{z}{z_{\text{ref}}} \right)^{\alpha}, \tag{2.1}$$

where z is the hub height, $z_{\text{ref}} = 100$ m, and α is the shear exponent. For offshore conditions, vertical shear is generally weak due to the low surface roughness of the sea, and a representative value of $\alpha = 0.10$ is adopted here. This standard approach ensures that the wind data are consistent with the physical characteristics of the wind WTGs used in the simulations.

Electricity price data are retrieved from Ember Energy which sources data from ENTSO-E [44]. ENTSO-E (the European Network of Transmission System Operators for Electricity) coordinates electricity transmission across Europe and provides data on wholesale day-ahead electricity prices, collected from multiple transmission system operators across European bidding zones. Ember Energy aggregates these ENTSO-E data and other market sources to produce validated historical price time series, spanning the years 2015 to 2025.

Both wind and price datasets are retrieved as hourly time series. To ensure comparability across years, the electricity price data are inflation-adjusted so that all values are expressed in constant 2025 euros. Annual inflation rates are taken from Eurostat [45], and the adjustment prevents the optimization algorithm from being biased toward more recent years with higher nominal prices.

Before use, the raw time series are checked and cleaned. Occasional irregularities occur in the form of missing timestamps or duplicate entries. Gaps are corrected by applying linear interpolation between neighboring values, ensuring temporal continuity. Duplicates are removed directly, since in all cases the repeated entries contain identical values. This cleaning process guarantees a one-to-one match between wind and price time series and avoids inconsistencies during optimization.

Given the size of the datasets and computational limits, the cleaned data are then transformed from time series into a frequency distribution. Wind direction is discretized into 12 bins of 30° each, while wind speed is discretized into 26 bins (1 m/s increments from 0 to 24 m/s, plus a final bin for all speeds above 24 m/s). For each wind condition bin, the corresponding electricity prices are gathered and averaged, and a probability of occurrence is computed from the frequency of the wind data. This results in a concise, probabilistic representation of the wind-price distribution at each site, which enables rapid evaluation of many layouts.

To illustrate this transformation, Table 2.2 presents an example of the resulting wind-price frequency matrix for the French location dataset. Each row corresponds to a wind direction sector (12 bins of 30°), and each column corresponds to a wind speed bin (1 m/s increments up to 24 m/s). The entries indicate the frequency of occurrence in percentage for each wind condition, with a total sum across the table of 100.

Table 2.2: Example of the wind–price distribution matrix (French location). Values represent percentages of occurrence per wind condition bin. Dots indicate omitted bins for readability.

Direction Bin	-0.5–0.5	0.5–1.5	1.5–2.5	…	21.5–22.5	22.5–23.5	23.5–24.5
-15–15	0.011	0.190	0.628	…	0.107	0.011	0.000
15–45	0.019	0.190	0.342	…	0.000	0.000	0.000
45–75	0.004	0.213	0.415	…	0.000	0.000	0.000
⋮	⋮	⋮	⋮	⋱	⋮	⋮	⋮
315–345	0.019	0.346	0.514	…	0.167	0.084	0.034

2.2.3. Wake and Power Modeling

The wake losses in this study are computed using the TurbOPark model (Turbulence Optimized Park model) [46], an engineering model designed for efficient and robust simulation of large wind farms. Like traditional wake models, TurbOPark uses a Gaussian wind speed deficit profile. The specificity of this

model is the use of a non-linear wake expansion formulation, driven by the turbulence intensity in the wake.

The normalized velocity deficit at a downstream position x and lateral distance r from the wake centerline when considering the wake of one WTG can be written as

$$\frac{\Delta U(r,x)}{U_\infty} = C(x)\exp\left(-\frac{r^2}{2\sigma^2(x)}\right), \tag{2.2}$$

where U_∞ is the free mean wind speed at hub height, $\sigma(x)$ is the wake width, which grows with the distance from the rotor, and $C(x)$ is the peak deficit at the center of the wake and is computed as follows:

$$C(x) = 1 - \sqrt{1 - \frac{C_T}{8\left(\sigma(x)/D\right)^2}}, \tag{2.3}$$

with D being the rotor diameter, and C_T the thrust coefficient. Wake deficits from multiple upstream WTGs are then combined using linear superposition, where the wind speed deficits are summed to compute the wind speed at each location of the design space.

The TurbOPark model has been shown to improve accuracy for modern large-scale offshore wind farms while maintaining low computational cost, making it suitable for layout optimization studies [47]. To illustrate the wake structure predicted by TurbOPark, Figure 2.4 shows the normalized wind speed deficit $\Delta U/U_\infty$ induced by the TurbOPark wake model, behind a single 180 m diameter wind turbine. The plot corresponds to a horizontal slice at hub height and highlights how the wake expands laterally and recovers with distance downstream. Note that the deficit near $x/D = 0$ is artificially low compared to the true velocity drop because TurbOPark is used to model far wake (after 2 to 3 rotor Diameters downstream).

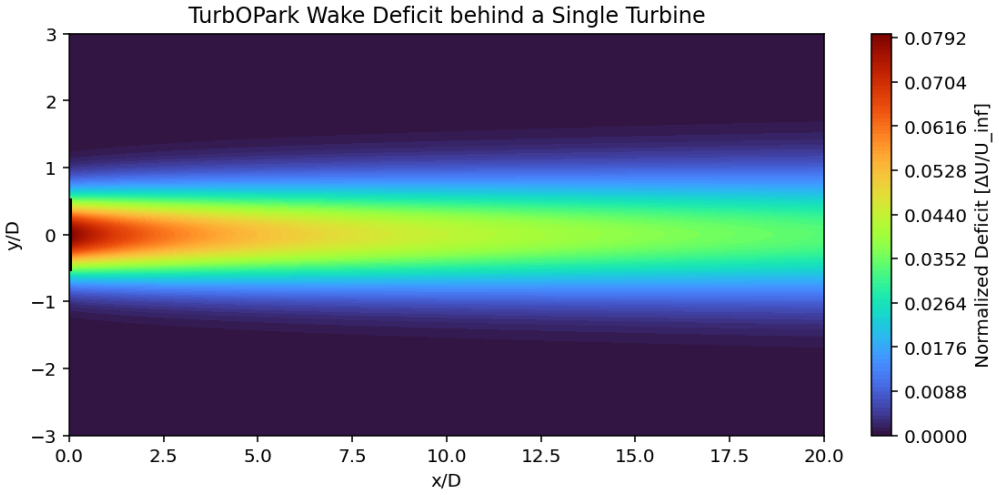

Figure 2.4: Normalized wake deficit $\Delta U/U_\infty$ behind a single 180 m diameter wind turbine, computed using the TurbOPark model.

Power production per WTG is calculated for each bin of wind speed and direction using the power curve of the 10 MW IEA reference WTG. The effective WTG power outputs are aggregated over all bins using the wind data probabilities to compute both energy and revenue outputs.

2.3. Optimization Problem

The wind farm layout optimization problem aims to determine the spatial arrangement of wind WTGs that maximizes a chosen performance metric while satisfying operational and physical constraints. The key components of this problem include the selection of objective functions, the definition of design variables (WTG positions), the consideration of constraints (e.g., minimum spacing, site boundaries),

and the choice of an optimization algorithm. In this study, we consider two distinct objective functions: AEP and revenue, which are optimized independently to assess trade-offs between energy output and economic return.

2.3.1. Objective Functions

Two objective functions are considered:

- **Annual Energy Production (AEP)**: Represents the total energy generated by the wind farm in one year, normalized per MW of installed capacity (MWh/MW). Mathematically:

$$\text{AEP} = \sum_{i=1}^{N} \sum_{j=1}^{N_{\text{bins}}} P_i(v_j, \theta_j) f(v_j, \theta_j) T_{\text{year}}, \tag{2.4}$$

  where $P_i(v_j, \theta_j)$ is the power output of WTG i for wind speed v_j and direction θ_j, $f(v_j, \theta_j)$ is the occurrence frequency of that bin, and T_{year} is the total hours in one year (8760 h).

  For each wind speed and wind direction, the wind farm production is estimated by proceeding from the most upstream WTG to the most downstream WTG and estimating wind speed deficits at the downstream WTGs.

  The computation is performed in four steps:

  - Binning the wind climate into discrete wind speed and direction bins,
  - Calculating WTG power per bin, including wake effects,
  - Weighting the power outputs by the frequency of each bin,
  - Summing over all WTGs and bins to obtain total AEP.

- **Revenue**: Represents the expected annual income from selling electricity on the day-ahead market, in k€/MW/year. It is calculated as

$$\text{Revenue} = \sum_{a=1}^{N} \sum_{n=1}^{N_{\text{bins}}} P_a(v_n, \theta_n) f(n) T_{\text{year}}\, \pi(v_n, \theta_n), \tag{2.5}$$

  where $P_a(v_n, \theta_n)$ is the power output of WTG a, $f(n)$ the occurrence frequency of bin n, T_{year} the total hours in one year (8760 h), and $\pi(n)$ the average electricity price for that bin. The computation follows the same sequence as for AEP, with the additional step of multiplying the power in each bin by the corresponding electricity price.

These objective functions are used independently to generate optimized layouts. Layouts optimized for one objective are then evaluated against the other metric to assess trade-offs.

2.3.2. Design Variables and Constraints

The design variables are the coordinates of the $N = 30$ WTGs:

$$\mathbf{s} = (x_1, y_1, \ldots, x_N, y_N),$$

representing their Easting and Northing positions within the farm.

To initialize the layout, an algorithm is implemented to ensure an efficient starting point:

1. A grid with 200 m spacing is built within the site boundary.
2. One grid point is chosen randomly.
3. Remaining WTGs are placed iteratively at the grid point furthest from all previously placed WTGs, ensuring an even initial spread.

This method allows the creation of an evenly spread grid that represents a balanced start for the optimization algorithm. A randomizing step is also added to the initial positions setter. Before each run of the optimization, a random number between 0 and 500 meters is added to each value of the set of WTGs coordinates. The goal is to help the algorithm explore the entirety of the design space and

consider a variety of initial positions while maintaining an even spread of the WTGs. Figure 2.2 is an example of a starting position for the optimization.

A series of constraints is then applied to these design variables to ensure the feasibility of the final layout.

- **Boundary Constraints:** For a WTG located at position (x_i, y_i) within a site defined by its boundaries $[x_{\text{min}}, x_{\text{max}}]$ and $[y_{\text{min}}, y_{\text{max}}]$, the following equation is enforced:

$$x_{\text{min}} \leq x_i \leq x_{\text{max}}, \quad y_{\text{min}} \leq y_i \leq y_{\text{max}}$$

- **Minimum Spacing Constraints:** For any two WTGs i and j with positions (x_i, y_i) and (x_j, y_j), the distance between them must satisfy:

$$\sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \geq 4D$$

  where D is the rotor diameter of the WTGs, 198 m here.

Constraints are enforced via a penalty mechanism: if a constraint is violated, a penalty term is added to the objective function:

$$\text{F}_{\text{penalized}} = \text{F}_{\text{original}} - \lambda \sum_{\text{violations}} \delta_i, \tag{2.6}$$

where δ_i measures the extent of violation, λ is a scaling factor and F is the objective function. This allows the solver to explore the design space while naturally discouraging infeasible solutions.

2.3.3. Optimization Algorithm

Optimization is performed using a stochastic gradient descent (SGD) algorithm [22]. The principle of SGD is to replace the exact gradient of the objective function with an unbiased stochastic approximation. Instead of evaluating the full wind climate (all wind speed and direction bins) at every iteration, which would be computationally heavy, the algorithm draws a small random subset of conditions and estimates the gradient from these samples. This significantly reduces the computational burden while still providing a search direction that, on average, points toward improved layouts. Over many iterations, the randomness cancels out and the algorithm converges to a near-optimal solution.

Formally, the gradient of the objective function (AEP or revenue) is approximated via Monte Carlo sampling of K wind speed and direction pairs:

$$\frac{dF}{ds} \approx \frac{8760}{K} \sum_{k=1}^{K} \frac{\partial F}{\partial s}\left(s, u_\infty^{(k)}, \theta^{(k)}\right). \tag{2.7}$$

where K is the number of sampled wind conditions, fixed at 50 in this study, s contains all WTG coordinates, and $(u_\infty^{(k)}, \theta^{(k)})$ are the free-stream velocity and direction for sample k.

At each iteration, the WTG positions are updated according to:

$$s_{i+1} = s_i - \eta_i \cdot \frac{\hat{m}_i}{\sqrt{\hat{v}_i}}, \tag{2.8}$$

where $\hat{m}_i$ and $\hat{v}_i$ are exponential moving averages of the gradient and its square, and η_i is the learning rate.

In summary, SGD optimization sacrifices the precision of a full gradient evaluation in favor of a cheaper stochastic approximation, which is sufficient to guide the optimizer toward high-performing wind farm layouts.

To accelerate the convergence, there is a speed-up option within the SGD algorithm, allowing the reuse of previous gradient samples for a few iterations and reducing the number of Monte Carlo samples as the algorithm progresses. This reduces computational cost without sacrificing too much accuracy. This also smooths the convergence of the algorithm.

Proper parameter tuning is also crucial: setting changes favoring convergence speed may lead to unstable behavior and non-convergence, while overly conservative values increase runtime without significant benefit. To tune the algorithm in the most efficient way possible, several test simulations were run to find the best set of parameters for the specific site and the available computational power:

- **Initial learning rate (η_0)**: Although studies suggest the use of an initial learning rate of $D/5$, best performing layouts were found with higher values of η_0, and the value is set at 300. This is found to allow the algorithm to explore more of the design space before converging towards a solution.
- **Maximum iterations (`max_iter`)**: Defines the maximum number of iterations before the stopping of the algorithm, and is set in this study to 2500. Even if the simulations often stopped before 2000 iterations, this number influences the decay of the learning rate, and having a longer max_iter value allowed for a slower decrease of this learning rate and better results.
- **Threshold (`sgd_thresh`)**: This threshold parameter determines how far the learning rate can decrease before stopping the optimization. In this study, it was set to 0.025, meaning that the optimization runs while $\eta/\eta_0 > 0.025$. Then, after reaching this threshold, an additional number of iterations is computed, with only constraint component gradients, to ensure that the final solution is feasible. The value of 0.025 allowed the algorithm to stop when no significant progress was anymore possible.
- **Constant learning-rate phase**: Starting with a fixed learning rate for an initial number of iterations encourages broader exploration and helps escape local minima. But testing different values of constant learning rate iterations showed very little improvement in final performance while increasing the computational time. This value was then kept at 0 for the optimizations.

This configuration ensures efficient exploration of the design space and convergence to high-performance layouts. SGD was chosen for its computational efficiency and better observed performance compared to other TopFarm drivers, such as Random Choice, SLSQP, and COBYLA. These performances are, for example, increased by 3% to 6%, depending on the sites and objective functions, when compared with the COBYLA algorithm.

A typical evolution of the objective function during optimization is shown in Figure 2.5.

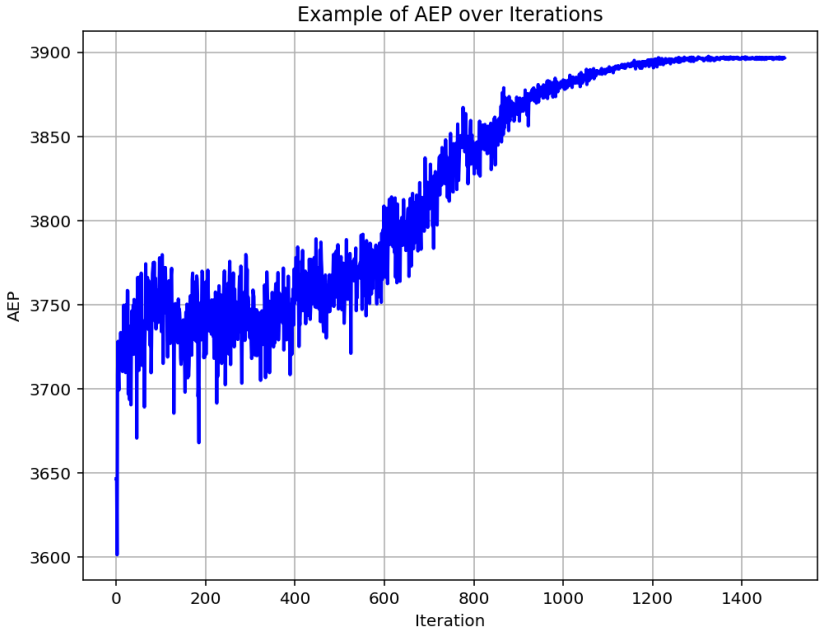

Figure 2.5: Typical evolution of the objective function during SGD optimization.

2.4. Case Studies

This section describes the numerical experiments used to assess the optimization methodology. A single numerical experiment is characterised by:

1. **Objective function:** either AEP or expected annual revenue (k€/MW/year).
2. **Optimization period and data:** the wind (historical or synthetic / scenario) climate and price data used during the optimization phase.
3. **Evaluation period and data:** the (possibly different) wind climate and price data used to compute final performance metrics of the obtained layouts.

. Table 2.3 summarises the three case studies considered in this work.

Table 2.3: Overview of simulation phases

Case study	Optimization period	Evaluation period	Objective function(s)	Purpose
1) Same-period comparison	2022–2024	2022–2024	AEP, Revenue	Compare layouts optimised for AEP and for revenue under identical wind and market conditions
2) Cross-period comparison	2015–2020	2022–2024	AEP, Revenue	Test robustness of layouts optimised on past data when evaluated on future conditions
3) Forecast-based scenario test	2035–2040	2035–2040	AEP, Revenue	Assess sensitivity of revenue-optimised layouts to alternative market scenarios

For the first case study, the objective is to directly compare layouts optimized for revenue with those optimized for AEP. The aim is to quantify the performance differences between the two, assess the degree of compatibility or conflict between the objectives, and observe how each criterion influences the resulting layouts.

The second case study focuses on robustness. By evaluating the layouts on a future time interval, we can assess how well designs optimized for AEP or revenue preserve their performance when subjected to different wind and market conditions.

Finally, the third case study is designed to reflect the perspective of wind farm developers. Here, the goal is to investigate the impact of electricity market forecasts, particularly assumptions about wind energy penetration, on layout optimization. Layouts are optimized for specific 2035–2040 market scenarios, but evaluated against alternative electricity price trajectories for the same period, to highlight the sensitivity of revenue performance to forecast uncertainties.

2.4.1. Numerical experiment protocol

To account for the stochastic nature of the algorithm and the sensitivity to initial layouts and local minima, each optimization run (defined by a choice of objective, location, and period) is solved 20 times with different randomized starting positions. For each case study, we follow the same reproducible protocol:

1. *Define inputs:* all the necessary numerical experiments are listed in an Excel file (20 per configuration), with the optimization period, the evaluation period, the location and the objective function.
2. *Solve optimization problems:* each entry of this Excel file is fed to the optimization algorithm, a full optimization (SGD driver) is run, and the final layouts are retained.
3. *Evaluate solutions:* each final layout is evaluated using the evaluation period data to compute both AEP and revenue (even when a layout was produced using only one of those objectives). This provides cross-metric comparison and robustness assessment.

4. *Aggregate results:* statistics across the 20 runs (mean, standard deviation, best/worst) for each metric (AEP and revenue) are then computed.

2.4.2. Iterations count

The number of optimization runs is shown in table 2.4.

Table 2.4: Run-count summary of the simulations

Case study	Objectives	Runs per configuration	Locations	Total runs
1) Same-period comparison	AEP, Revenue	20	**4**	$2 \times 20 \times \mathbf{4} = 160$
2) Cross-period comparison	AEP, Revenue	20	**4**	$2 \times 20 \times \mathbf{4} = 160$
3) Forecast-based scenarios	AEP, Revenue	20	**1**	$2 \times 20 \times \mathbf{1} = 40$
Total (all case studies)				**360** runs

Each optimization run takes approximately 8- 10 hours on one core; therefore, the total time for the optimization campaign is between 2740 and 3420 hours. This computational time is considerably reduced with parallelization and the use of a multi-core solving cluster.

2.5. Market Scenario Generation

To evaluate how wind farm layouts perform under different market conditions, synthetic hourly electricity price scenarios are constructed. A key driver for these scenarios is the expected cannibalization effect, where increased wind penetration lowers wholesale electricity prices during high wind periods [3]. This effect can alter the economic attractiveness of different wind speed and direction sectors: high-wind conditions, typically favorable for energy yield, may coincide with depressed prices, while low-wind periods may see higher prices. Accounting for this dynamic is essential when comparing AEP- and revenue-optimized layouts.

Forecasting future hourly prices with full system models, such as Balmorel [39], is possible but requires significant data and modeling effort. Instead, this study adopts the simpler but flexible stochastic approach of Mehta et al. [17], which generates plausible hourly price series by controlling a small set of key statistical parameters: the long-term mean price μ, the price volatility (expressed as a coefficient of variation, CV), and the target Pearson correlation ρ between hourly prices and local wind speeds.

2.5.1. Scenario construction procedure

Each synthetic price series is obtained in four steps:

1. **Long-term baseline:** A deterministic baseline $p_{\text{base}}(t)$ is built to reflect the long-term expected mean price μ for each hour of the projection horizon $p_{\text{base}}(t) = \mu$.
2. **Seasonal and diurnal patterns:** Hourly scaling factors are applied to replicate observed seasonality and intraday variability:

$$p_{\text{trend}}(t) = p_{\text{base}}(t) \cdot f_{\text{season}}(m(t)) \cdot f_{\text{diurnal}}(h(t)),$$

   where $m(t)$ is the calendar month and $h(t)$ is the hour of the day. f_{season} and f_{diurnal} are derived from historical ENTSO-E price data.
3. **Stochastic perturbation with controlled correlation:** A zero-mean random perturbation $y(t)$ with prescribed volatility σ_p and correlation ρ with the standardized wind speed time series $W(t)$ is constructed. Let z be a standard normal random vector independent of W. Then we decompose z into a component parallel to W and an orthogonal component:

$$w = \frac{z \cdot W}{\|W\|^2} W, \quad w_\perp = z - w.$$

   A correlated perturbation is formed as

$$y = \rho \frac{w}{\|w\|} + \sqrt{1-\rho^2} \frac{w_\perp}{\|w_\perp\|}.$$

Finally, y is scaled to the desired standard deviation σ_p and shifted to zero mean.

4. **Price series assembly:** The final hourly price scenario is

$$P_\rho(t) = p_{\text{trend}}(t) + y(t).$$

Multiple realizations are created by varying the random seed to test sensitivity to unpredictable short-term price noise.

This procedure directly enforces the target long-run mean and volatility while prescribing the linear correlation ρ between price and wind speed. It allows constructing scenarios such as:

- $\rho = 0$ — no cannibalization (wind has no impact on prices);
- $\rho = -0.3$ — baseline current market conditions;
- $\rho = -0.6$ — stronger future cannibalization.

2.5.2. Advantages of the method

Compared to full system models, this approach:

- isolates the price–wind correlation as a tunable parameter, directly linking market scenarios to the physical drivers of cannibalization;
- preserves realistic intra-annual and daily price patterns based on historical data;
- enables generation of large scenario ensembles at low computational cost, suitable for layout optimization studies.

An illustration of a resulting scenario is provided in Figure 4.17, showing the hourly price series after seasonal and diurnal modulation and the stochastic correlated perturbation. This scenario generation is used throughout Case Study 3 to evaluate the robustness of revenue-optimized layouts to possible future market structures.

3

Analysis of the wind and price distribution of the sites

Before running the first case study, it is essential to examine the site characteristics of each location to better understand the expected impact of changing the objective function. In wind farm design, the wind rose is typically used to characterize the wind resource at a site, providing the frequency distribution of wind speeds and directions. For each study area, the wind rose is plotted in Figure 3.1, highlighting the prevailing wind directions and speed classes.

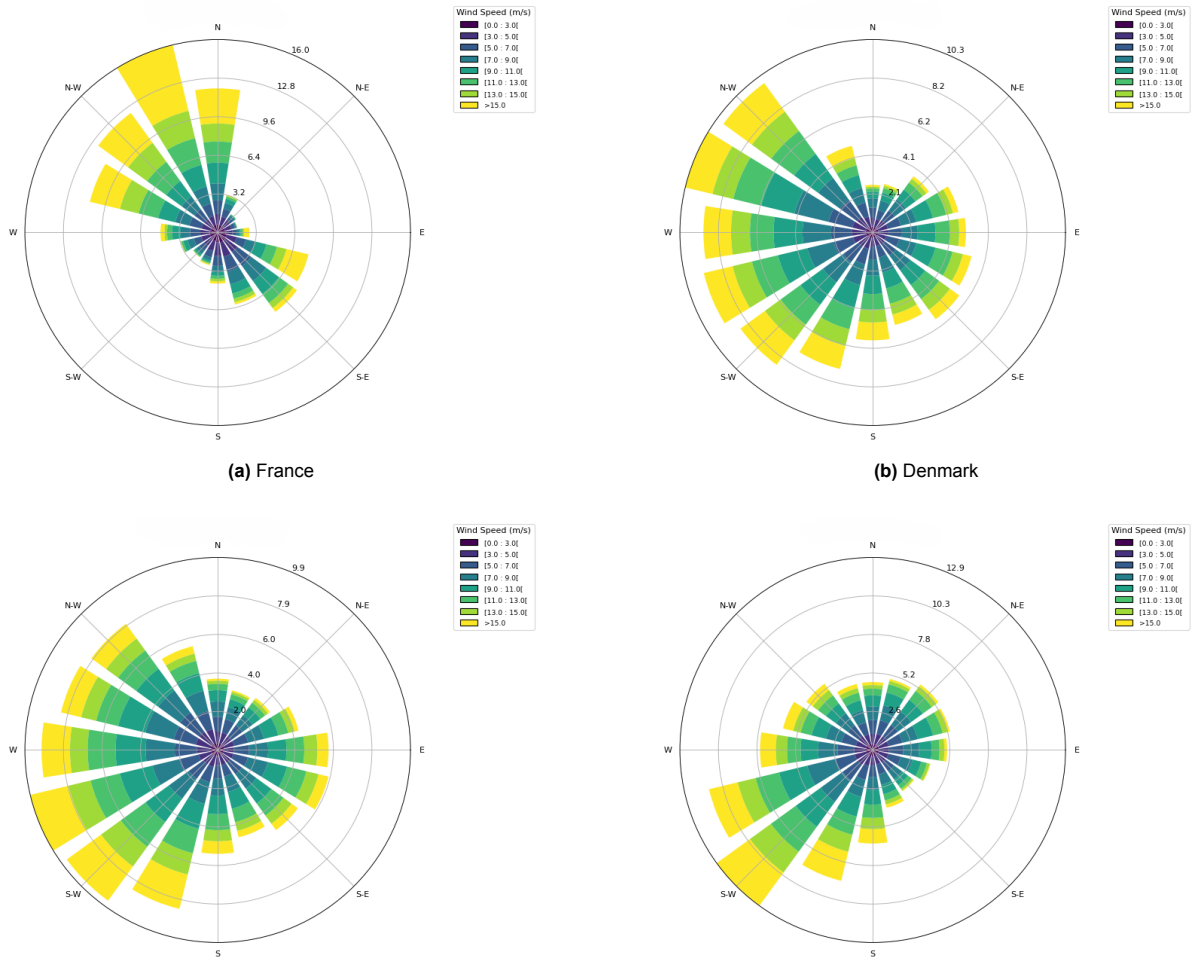

(a) France

(b) Denmark

(c) Germany

(d) Netherlands

Figure 3.1: Wind roses for each studied location for the time period 01/01/2022 to 31/12/2024.

However, the focus of this study is on the difference between energy yield and revenue. To capture this, we compute two analogous "roses". Specifically, for each wind speed bin v_i and direction bin θ_j (with annual frequency $f(v_i, \theta_j)$), the expected gross energy output per turbine is

$$E(v_i, \theta_j) = f(v_i, \theta_j)\, P(v_i, \theta_j),$$

where $P(v_i, \theta_i)$ is the turbine power curve output at speed v_i and direction θ_j. The corresponding annual revenue is

$$R(v_i, \theta_j) = E(v_i, \theta_j)\, c(v_i, \theta_j),$$

where $c(v_i, \theta_j)$ is the average electricity price in that speed–direction bin. Summing these over all speed bins for a given direction θ_j yields the directional energy rose $E_j = \sum_i E(v_i, \theta_j)$, and similarly the revenue rose $R_j = \sum_i R(v_i, \theta_j)$. In Figure 3.2, the normalized energy and revenue roses (values normalized to their respective maxima) are plotted for each site. This visualization illustrates the comparison between the directional distribution of potential power production (left panels) and the distribution of potential revenue (right panels) at each location.

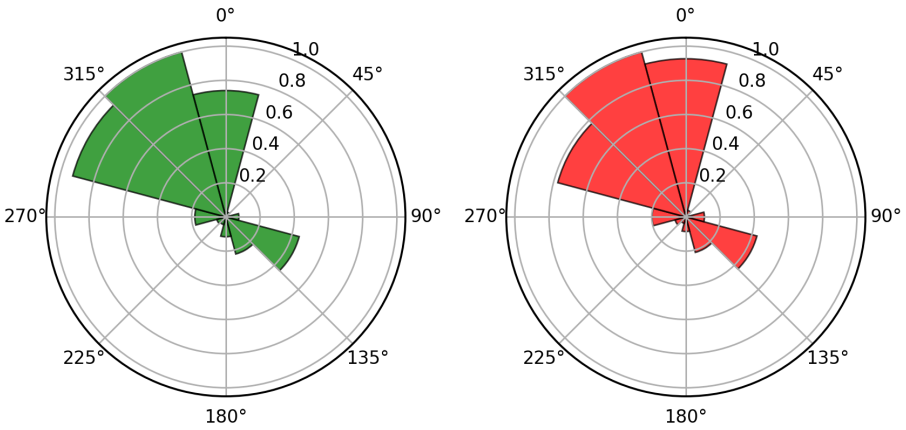

(a) France – Normalized directional potential energy (left) and revenue (right)

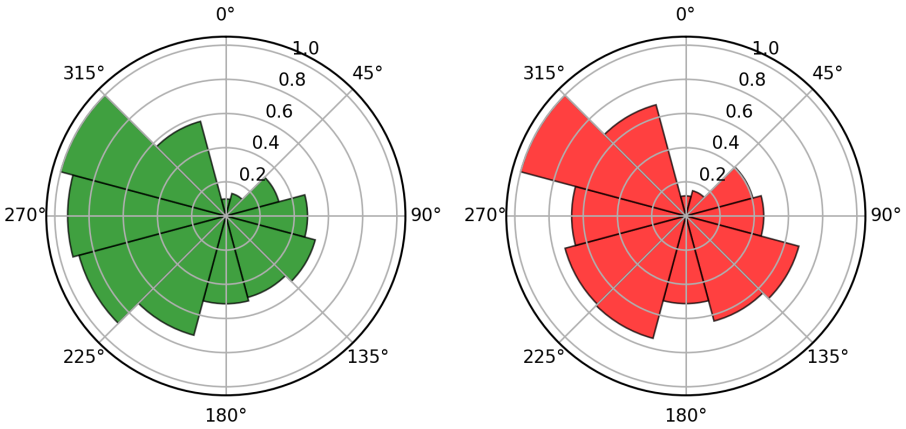

(b) Denmark – Normalized directional potential energy (left) and revenue (right)

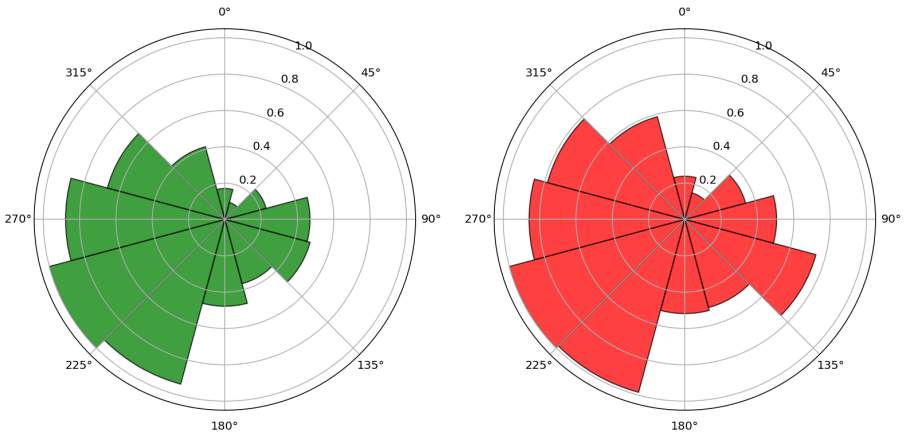

(c) Germany – Normalized directional potential energy (left) and revenue (right)

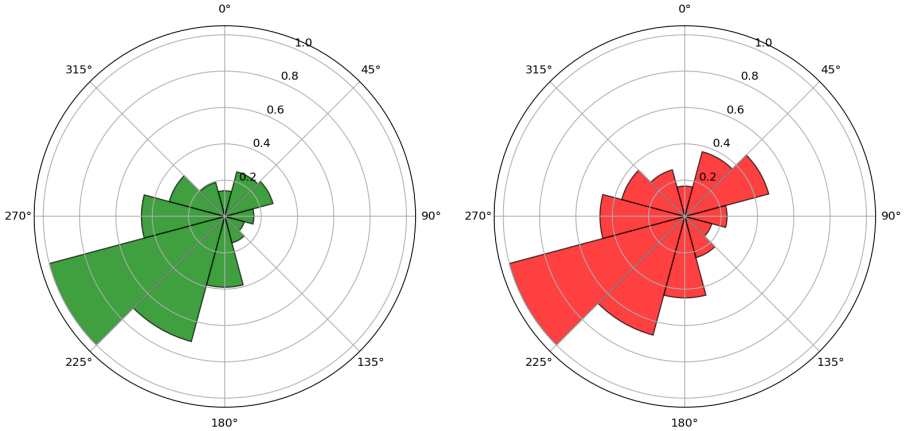

(d) Netherlands – Normalized directional potential energy (left) and revenue (right)

Figure 3.2: Comparison of normalized potential energy and revenue roses for each location.

From Figure 3.2, one observes that some sites have energy and revenue roses of similar shape (e.g. the Netherlands), whereas others differ more (France, Denmark, Germany). Visual comparison alone can be inconclusive, so we use a quantitative metric. We adopt the 1D Wasserstein distance (Earth Mover's Distance, without unit) to measure the dissimilarity between the energy and revenue distributions [48]. Intuitively, this distance represents the "work" needed to transform one distribution into the other by moving probability mass between bins. In one dimension, the 1-Wasserstein distance has the form

$$W_1(Q_1, Q_2) = \int_{-\infty}^{\infty} |F_{Q_1}(x) - F_{Q_2}(x)|\, dx,$$

where F_{Q_1} and F_{Q_2} are the cumulative distribution functions. In our discrete binned case, we compute two distances:

$$W_{\text{dir}} = \sum_{j=1}^{N_{\text{dir}}} |F_{Q_1,\text{dir}}(j) - F_{Q_2,\text{dir}}(j)|\,, \quad W_{\text{speed}} = \sum_{i=1}^{N_{\text{speed}}} |F_{Q_1,\text{speed}}(i) - F_{Q_2,\text{speed}}(i)|\,.$$

Here, $F_{Q_1,\text{dir}}(j)$ is the cumulative distribution of Q_1 per directional bin, and $F_{Q_2,\text{dir}}(j)$ is the analogous cumulative distribution for Q_2; similarly, $F_{Q_1,\text{speed}}(i)$ and $F_{Q_2,\text{speed}}(i)$ are the cumulative distributions over wind speed bins instead of direction bins. A larger W indicates a stronger discrepancy (i.e. a large displacement of distribution "mass" between bins), implying that optimizing for Q_2 may yield a layout significantly different from optimizing for Q_1. Conversely, a small W suggests the two objectives should favor similar layouts. The resulting values of W_{dir} and W_{speed} for each site are given in Table 3.1.

Table 3.1: 1D Wasserstein distances between Q_1 and Q_2 distributions for wind direction and speed

Location	$\boldsymbol{W_{\text{dir}}(Q_1, Q_2)}$	$\boldsymbol{W_{\text{speed}}(Q_1, Q_2)}$
France	0.469	0.225
Denmark	0.279	1.046
Germany	0.257	0.917
Netherlands	0.396	0.693

One can notice from Table 3.1 that, for all locations, the Wasserstein distance with respect to wind speed bins is greater than the one with respect to directional bins ($W_{\text{speed}} > W_{\text{dir}}$). This indicates that differences between Q_1 and Q_2 distributions are stronger for the wind speed dimension than for the wind direction. In other words, optimizing for Q_2 is likely to emphasize certain wind speed ranges more than optimizing for Q_1, for example, favoring high-value but less frequent wind speeds due to their association with higher market prices. By contrast, the preferred wind directions of Q_1 and Q_2 are more aligned, leading to lower values of W_{dir}.

But by treating wind direction and speed independently, these results cannot capture how misalignments in both dimensions interact. As a result, 1D Wasserstein distances may underestimate the overall divergence between Q_1- and Q_2-oriented distributions, potentially leading to misleading conclusions regarding the sensitivity of layout optimization to the choice of objective function.

To overcome the limitations of 1D metrics, we compute the 2D Wasserstein distance between the joint distributions of wind direction and wind speed.

Let $P(i,j)$ and $Q(i,j)$ represent the normalized joint probability distributions of energy and revenue, where i indexes direction bins and j indexes speed bins. Each bin can be associated with a point in the 2D plane, using the center values of the direction and speed bins.

M is then defined as the cost matrix that contains the distance between every pair of bins in P and Q. The distance between bin (i,j) in P and bin (k,l) in Q is computed using the Euclidean formula:

$$M_{i,j;k,l} = \sqrt{(x_i - x_k)^2 + (y_j - y_l)^2}, \tag{3.1}$$

where x_i and x_k are the direction bin centers, and y_j and y_l are the speed bin centers. This represents the "effort" required to move a unit of probability from one bin to another.

The 2D Wasserstein distance W_{2D} is obtained by finding the optimal transport plan γ that moves the probability mass from P to Q with minimal total cost:

$$W_{2D} = \min_{\gamma \in \Pi(P,Q)} \sum_{i,j,k,l} \gamma_{i,j;k,l} \, M_{i,j;k,l}, \tag{3.2}$$

subject to the constraints that all "mass" in P is moved to Q and the row/column sums of γ match the distributions.

Table 3.2: 2D Wasserstein distance between energy and revenue distributions for all sites

Location	W_{2D}
France	14.100
Denmark	8.530
Germany	7.852
Netherlands	11.942

Table 3.2 shows that the largest overall difference between energy- and revenue-oriented distributions occurs at the French site ($W_{2D} = 14.100$), followed by the Netherlands (11.942), Denmark (8.530), and Germany (7.852). While this ranking differs from the W_{speed} computation, W_{dir} seems to play an important role in the characterization of the sites, as the ranking remains the same for W_{2D}. This indicates that the difference between the sites of this study is mainly driven by differences in direction bins rather than speed bins.

For France and the Netherlands, these results indicate that layout optimization under revenue objectives could produce configurations that differ considerably from AEP-optimized layouts. In contrast, Denmark and Germany display smaller 2D distances, suggesting that changing the objective function is likely to have a more limited impact. However, it is important to interpret these metrics carefully. For example, at the Dutch site, the largest directional shift in revenue emphasis occurs in a sector roughly opposite the main energy-producing direction (see Figure 3.2d, where the strongest increase in red is in the 45°–75° sector). Although the optimization problem is not perfectly symmetric, there is still a partial symmetry—for instance, wake effects under northerly wind are mirrored under southerly wind. While it is not straightforward to quantify how this influences performance differences between the two objective functions, such symmetry may limit the actual impact.

Overall, the 2D Wasserstein distance provides a more comprehensive measure of misalignment, capturing interactions between wind direction and speed, although direction misalignment seems to drive the differences in this study. It is still found to be the most informative metric for quantifying differences between energy and revenue in this study. Nevertheless, it should always be interpreted in context, alongside visual inspection of the roses, to fully understand the implications for layout optimization.

4

Results

In this chapter, the case studies are presented and evaluated with the aim of answering the research questions. Each case study is first briefly reintroduced, and then results are presented as performance evaluations, or layout comparisons.

4.1. Case Study 1

For the first case study, the layouts are optimized for the time period 01/01/2022– 31/12/2024. The selection of this time frame was constrained by the availability of data. Day-ahead electricity prices are available only from 01/01/2015 up to the retrieval date (31/03/2025). The time period 01/01/2020 to 31/12/2022 was excluded from the optimization window, as it was characterized by abnormally high electricity prices driven by geopolitical factors. Similarly, the data from 2025 were not considered. This choice is justified by the fact that the transformation of time-series data into probabilistic input files requires balanced seasonal representation. Including only the first months of 2025 would artificially overweight winter months (e.g., January and February would be represented three times, while summer months only twice), biasing the probabilistic distributions.

4.1.1. Performance Comparison

The first research question is addressed here: **To what extent does switching the objective function from AEP to revenue affect the resulting AEP and revenue of the optimized layouts?** Figures 4.1, 4.2, 4.3, and 4.4 illustrate this relationship, showing scatter plots of AEP versus revenue for each optimization run alongside box plots that summarize the results for both objective functions.

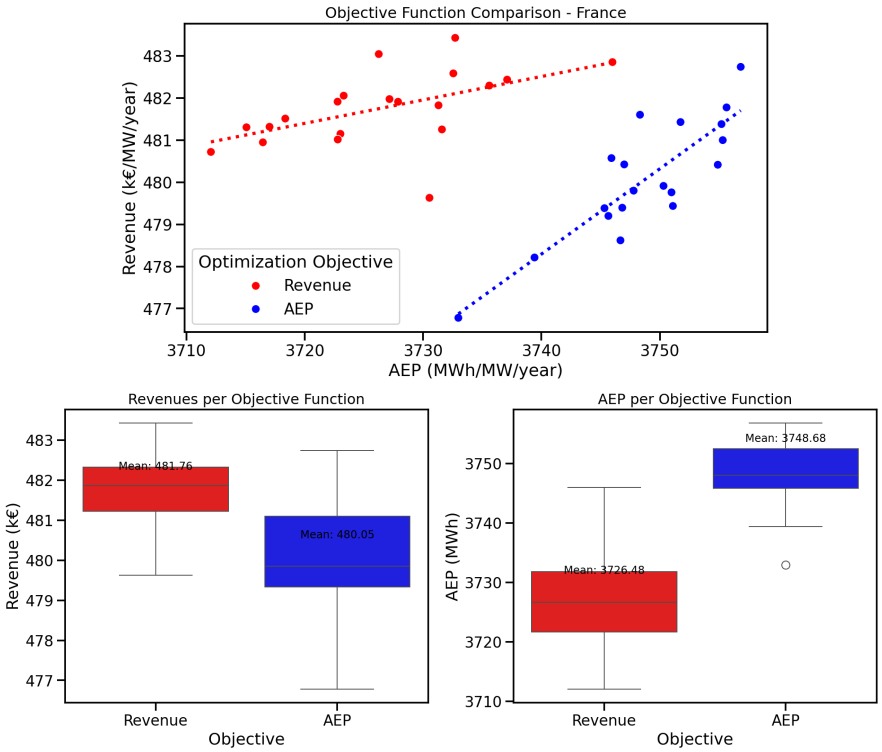

Figure 4.1: Scatter plots and boxplots for the French project.

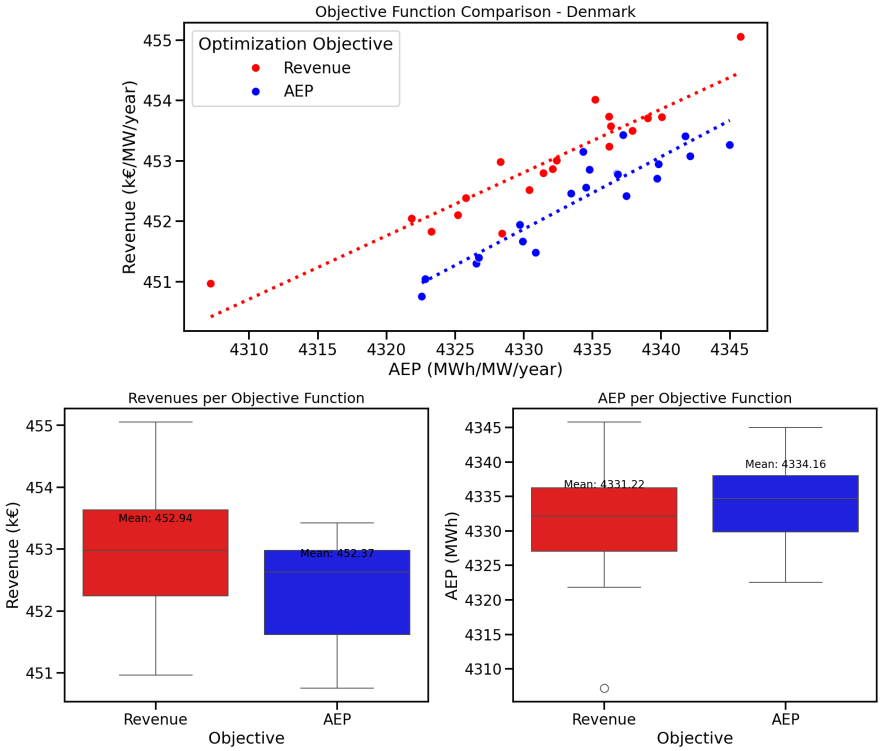

Figure 4.2: Scatter plots and boxplots for the Danish project.

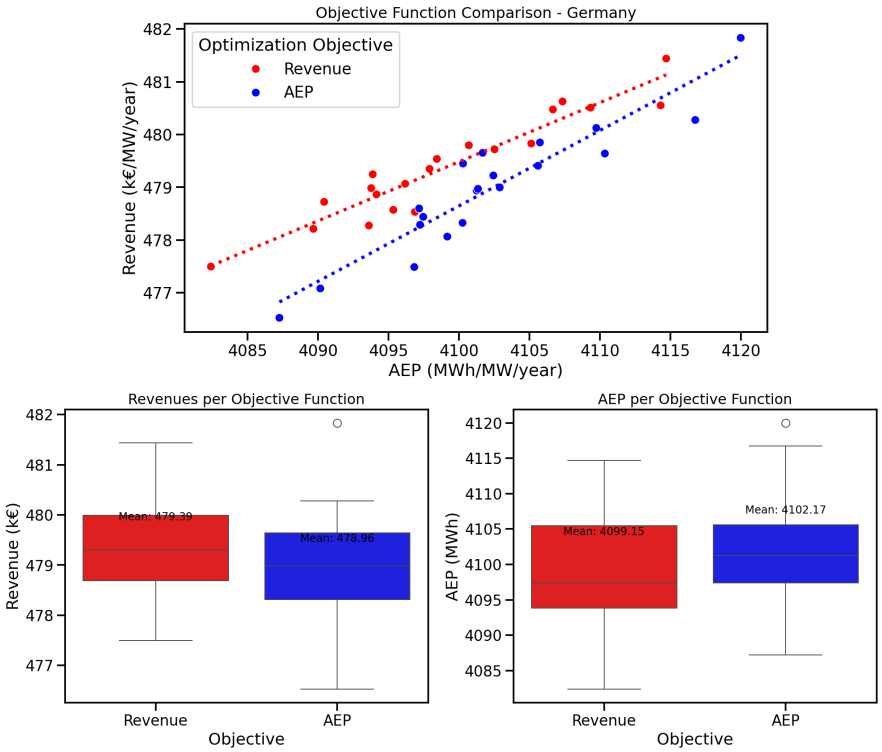

Figure 4.3: Scatter plots and boxplots for the German project.

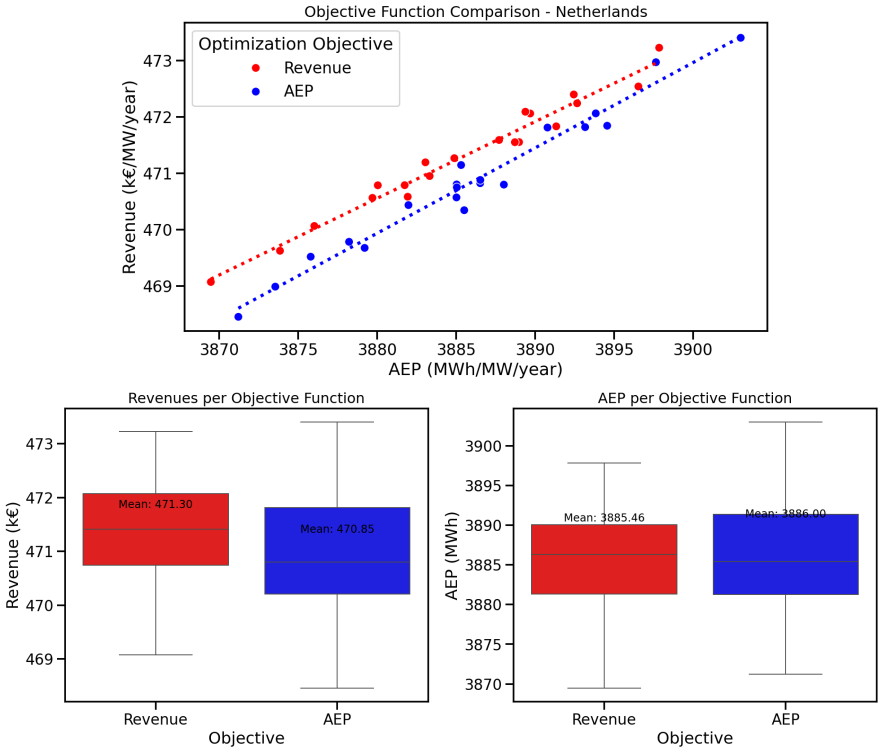

Figure 4.4: Scatter plots and boxplots for the Dutch project.

A first observation from these plots is that high AEP and high revenue tend to be linked. In all cases, the best-performing layouts in terms of AEP are also among the best in terms of revenue for the same

objective function. This is reflected in the regression lines, which all have positive slopes, showing a clear correlation between the two objectives.

The slopes of these regression lines are generally similar across objective functions, with the exception of France. This exception is not considered meaningful in this work, as it disappears in further Case Study 4.2.

Additionally, the expected trade-off is evident: revenue-optimized layouts yield higher revenue on average, while AEP-optimized layouts yield higher AEP. This is clearly illustrated by the red regression line (revenue optimization) lying in the top left section of the plot compared to the blue one (AEP optimization). This can also be observed in the boxplots.

The magnitude of these improvements, however, varies across sites. To quantify this effect, Tables 4.1 and 4.2 show, respectively, the average increase in revenue when using revenue as the objective function, and the corresponding decrease in AEP.

Table 4.1: Comparison of average revenues in k€/MW/year for layouts optimized for AEP vs. revenue, with relative improvement.

Site location	Optimized for AEP	Optimized for revenue	Improvement
France	480.05	481.76	0.36%
Denmark	452.37	452.94	0.13%
Germany	478.96	479.39	0.09%
Netherlands	470.85	471.30	0.10%

Table 4.2: Comparison of average AEP in MWh/MW/year for layouts optimized for AEP vs. revenue, with relative loss.

Site location	Optimized for AEP	Optimized for revenue	Loss
France	3748.7	3726.5	-0.60%
Denmark	4334.2	4331.2	-0.07%
Germany	4102.2	4099.2	-0.07%
Netherlands	3886.0	3885.5	-0.01%

These results confirm the larger trade-off between AEP and revenue in France, while Denmark, Germany, and the Netherlands also exhibit notable but smaller differences.

This ranking is partly consistent with the 2D Wasserstein (Earth Mover's Distance) analysis introduced in Chapter 3. France, which displayed the largest trade-off between AEP- and revenue-optimized layouts, also has the highest 2D Wasserstein distance ($W_{2D} = 14.100$). Similarly, Germany and Denmark, where performance differences are smaller, exhibit comparatively low Wasserstein distances (7.852 and 8.530, respectively).

The Netherlands, however, provides a counterexample. Despite having a relatively large 2D Wasserstein distance ($W_{2D} = 11.942$), the observed trade-off in layout performance is similar to the Germany and Denmark ones, and much smaller than in France. This discrepancy highlights one limitation of using Wasserstein distance as a predictive criterion: while it captures the overall misalignment between energy and revenue roses, it does not always reflect how such differences propagate into layout optimization outcomes. In particular, the Dutch site exhibits a form of directional symmetry: the largest shifts between energy and revenue occur in sectors roughly opposite the main energy-producing directions. Because wake interactions in opposite sectors can be similar, the effective impact on optimized layouts is reduced, even though the statistical distance between roses remains high.

Overall, the 2D Wasserstein distance provides valuable insight into the potential divergence between AEP and revenue optimization, and it successfully identifies France as the most sensitive site. Yet it must be interpreted with caution, as site-specific characteristics such as symmetry in the wind resource can limit the translation of statistical misalignment into actual performance differences.

However, for developers, the most relevant value is not the average of all optimization runs, but the performance of the best layouts, since those would ultimately be selected for implementation. In this optimization setup, it is not possible to ensure the finding of the best possible layout within a set number

of optimizations, and Figures 4.1, 4.2, **??**, and 4.4 show that the best-performing layout depends on individual optimization performance. Notably, the layout yielding the highest revenue is not always obtained from the revenue-optimized runs. For instance, in Germany and the Netherlands, the highest-revenue layout was actually produced by an AEP-optimized run.

This observation highlights an important limitation: while the choice of objective function influences the resulting layouts, the stochastic nature of the optimization and its reliance on local search can mean that the best single design sometimes emerges from the "wrong" objective. To address this, and to provide a fairer comparison, the analysis uses the average of the three best layouts for each objective function, filtering out low-performing solutions (Tables 4.3 and 4.4). It is also worth noting that this effect may diminish with a larger number of optimization runs, as additional runs would increase the likelihood of approaching the true global optimum for each objective.

Table 4.3: Comparison of average revenues in k€/MW/year for the three best layouts optimized for AEP vs. revenue, with relative improvement.

Site location	Optimized for AEP	Optimized for revenue	Improvement
France	482.04	483.11	0.22%
Denmark	453.37	454.26	0.20%
Germany	480.75	480.83	0.02%
Netherlands	472.81	472.72	-0.02%

Table 4.4: Comparison of average AEP in MWh/MW/year for the three best layouts optimized for AEP vs. revenue, with relative loss.

Site location	Optimized for AEP	Optimized for revenue	Loss
France	3755.9	3739.6	-0.44%
Denmark	4343.0	4341.6	-0.03%
Germany	4113.0	4112.8	0.00%
Netherlands	3898.4	3895.7	-0.07%

The trends observed here generally mirror those of the full averages. Denmark stands out and shows a relatively higher revenue gain (comparable to France) and a very small loss in AEP. This is explained by the outstanding performance of a single revenue-optimized layout at 4346 MWh/MW/year, and 55 k€/MW/year. This highlights the importance of running multiple optimization runs to mitigate the risk of outliers dominating the results.

It is also worth noting that the relative magnitude of revenue gains versus AEP losses differs across sites, with no generalizable rule emerging from this study.

In the French case, a 0.22% increase in revenue corresponds to 1.06 k€/MW/year. For a 300 MW site, this translates into approximately 318,000 € per year. While this value is non-negligible, it remains modest compared to major sources of uncertainty, such as wake modeling. Lee and Fields [49] report wake-related loss uncertainties typically ranging between 10% and 20% of overall wkae loss magnitude. Nevertheless, the magnitude of improvements found here is consistent with other algorithmic enhancements reported in the literature. For instance, Nguyen et al. [4] observed a 0.18% increase in revenues when accounting for reserve markets in addition to day-ahead markets for a Belgian offshore wind farm.

Finally, these results can also be contextualized by comparing them to the standard deviation of the layout performances. For France, the standard deviation of revenue is 0.93 k€/MW/year, representing 0.20% of the mean. Improvements of 0.22%–0.36% are therefore on the same order of magnitude, suggesting a significant improvement.

One major conclusion to take from this performance comparison is that the change in revenue and AEP following the change in the objective function ranges from relevant to almost immaterial, depending on the considered location and its wind and market conditions. In particular, while the observed revenue

improvements when switching from AEP to revenue optimization remain modest (below 0.4% in most cases), they are systematic and measurable. These values, although small in absolute terms, can still represent non-negligible additional income when scaled to the size of a commercial offshore wind farm. For instance, the 0.22% increase in revenue observed for France translates to roughly 318,000 € per year for a 300 MW project.

A second important conclusion is that the choice of objective function is not the only driver of layout performance. The stochastic nature of the optimization process means that some of the highest-revenue solutions occasionally emerge from AEP-driven runs. This underlines the need for multiple optimization attempts and robust post-processing of results rather than relying on a single run outcome. Averaging the performance of the best few layouts, as done in this analysis, provides a more reliable basis for comparison and for practical design decisions.

From a developer's perspective, these findings suggest a balanced approach. When time and computational resources are limited, optimizing for AEP alone is unlikely to result in severe economic underperformance, especially in sites where the energy and revenue roses are well aligned. Yet, when market signals indicate a strong potential for wind energy price cannibalization, or when the long-term correlation between wind and prices is expected to increase in absolute value, explicitly optimizing for revenue can bring small but tangible financial benefits.

4.1.2. Impact on Layouts

The objective of this section is to assess how changing the optimization objective (from AEP to revenue) affects the final positions of the WTGs. We focus on the French and Danish sites because these locations exhibited the largest changes in performance metrics in Section 4.1.1. For both sites, the three best layouts obtained for each objective (AEP and revenue) are again analyzed, and their spatial arrangements and directional wake behavior are compared.

Figures 4.5 and 4.6 display the three best layouts found for each objective function at the French and Danish sites, respectively. Each subfigure contains two rows: the top row shows the three best AEP-optimized layouts, and the bottom row shows the three best revenue-optimized layouts. In each panel, a dot marks the turbine hub position and a dashed line indicates the turbine displacement with respect to the initial layout.

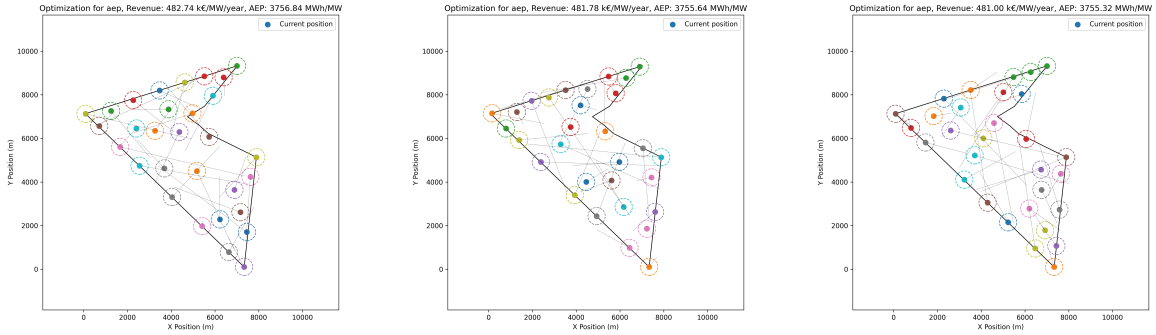

(a) France — top three layouts found when optimizing for AEP (rank 1–3).

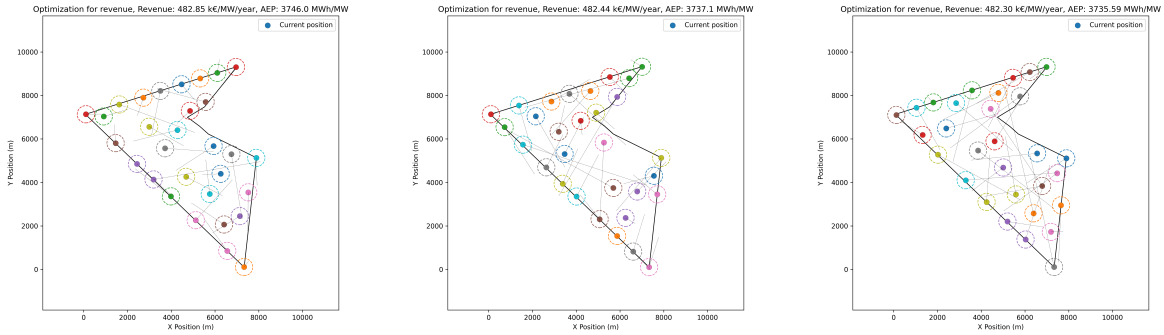

(b) France — top three layouts found when optimizing for revenue (rank 1–3).

Figure 4.5: France — comparison of the best layouts per objective. Top row: best AEP layouts. Bottom row: best revenue layouts. Dots: WTGs; dashed lines: displacement from initial layout; circles: enforced spacing constraint between WTGS.

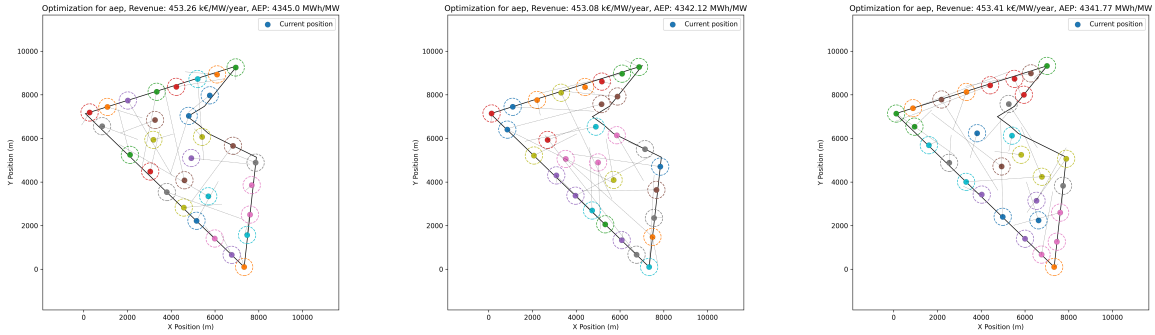

(a) Denmark — top three layouts found when optimizing for AEP (rank 1–3).

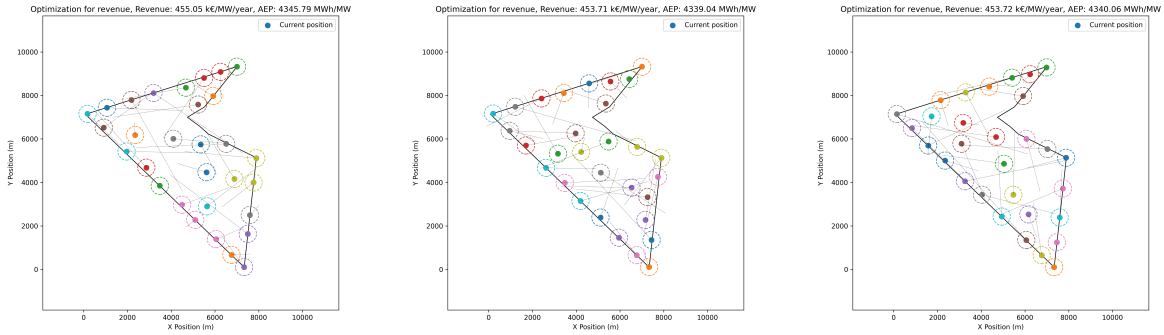

(b) Denmark — top three layouts found when optimizing for revenue (rank 1–3).

Figure 4.6: Denmark — comparison of the best layouts per objective. Top row: best AEP layouts. Bottom row: best revenue layouts. Dots: ; dashed lines: displacement from initial layout; circles: enforced spacing constraint between WTGS.

Visual inspection of the panels in Figures 4.5 and 4.6 reveals only modest positional differences between the two objective sets: WTGs generally occupy similar regions of the farm, and the overall spatial structure is preserved.

First, perimeter WTGs remain broadly stable between AEP- and revenue-optimized layouts, with most WTGs along the farm boundary moving only slightly. This stability is expected since boundary WTGs are less affected by upstream wakes and already benefit from high free-stream exposure. Revenue optimization therefore rarely relocates these machines.

For the interior part of the farm, no clear systematic trend emerges. Some central WTGs are displaced by one or two rotor diameters between the two optimization objectives, but the direction and magnitude of these shifts vary across sites and runs. In some cases, inner WTGs are slightly re-spaced or reoriented, but there is no consistent pattern of movement linked to the change from AEP to revenue optimization. This variability reflects both the complex wake interactions in the dense core of the array and the stochastic nature of the optimization process, making it difficult to generalize how inner WTGs adapt when switching objectives.

To better understand trends and patterns emerging from these layouts, two quantitative diagnostics averaged over the three best layouts for each objective are considered:

- spatial density map obtained via kernel density estimation (KDE)
- wake roses (directional wake loss profiles),

A KDE provides a smooth estimate of the spatial probability density of WTG locations and is useful to detect systematic clustering or shifts between layout ensembles. It is widely used in spatial analyses because it transforms discrete points into a continuous density field that can be compared between ensembles.

KDE maps were computed for the three best layouts under each objective (see Appendix A). The KDEs did not reveal any systematic large-scale displacement between objectives at either site. Small, localized changes are visible (slight increases in density along certain site edges), but no consistent shift of turbines from one sector of the farm to another is observed.

On the other hand, wake roses quantify the directional dependence of wake-induced energy losses and are central to interpreting how layout differences favor some directions. For a given inflow direction θ, $AEP_{\text{layout}}(\theta)$ is the annual energy production associated with the layout. The wake loss for direction θ is then defined as the relative reduction of production with respect to the gross yield without wakes $AEP_{\text{ref}}(\theta)$:

$$\mathrm{WL}(\theta) \;=\; 100\% \times \left(1 - \frac{AEP_{\text{layout}}(\theta)}{AEP_{\text{ref}}(\theta)}\right). \tag{4.1}$$

In this study, to compute the AEP without wake effects, the free-stream production of one turbine at the same inflow conditions is calculated, and multiplied by the number $N = 30$ of wind turbines in the farm:

$$A^{\text{iso}}_{\text{ref}}(\theta) \;=\; N \cdot AEP_{\text{single}}(\theta), \tag{4.2}$$

where $AEP_{\text{single}}(\theta)$ is the annual production of a single isolated turbine under the specific direction wind-speed distribution. The wake loss per direction is therefore obtained by substituting (4.2) into (4.1).

The averaged wake-loss roses for the three best layouts at each site are shown in Figures 4.7 and 4.8. In the left panels of Figures 4.7a and 4.8a, the mean directional wake loss is plotted for the three best AEP-optimized layouts (blue) and the three best revenue-optimized layouts (orange). The right panels show the difference (revenue minus AEP) in wake loss, highlighting directions where one objective yields larger directional losses than the other on average.

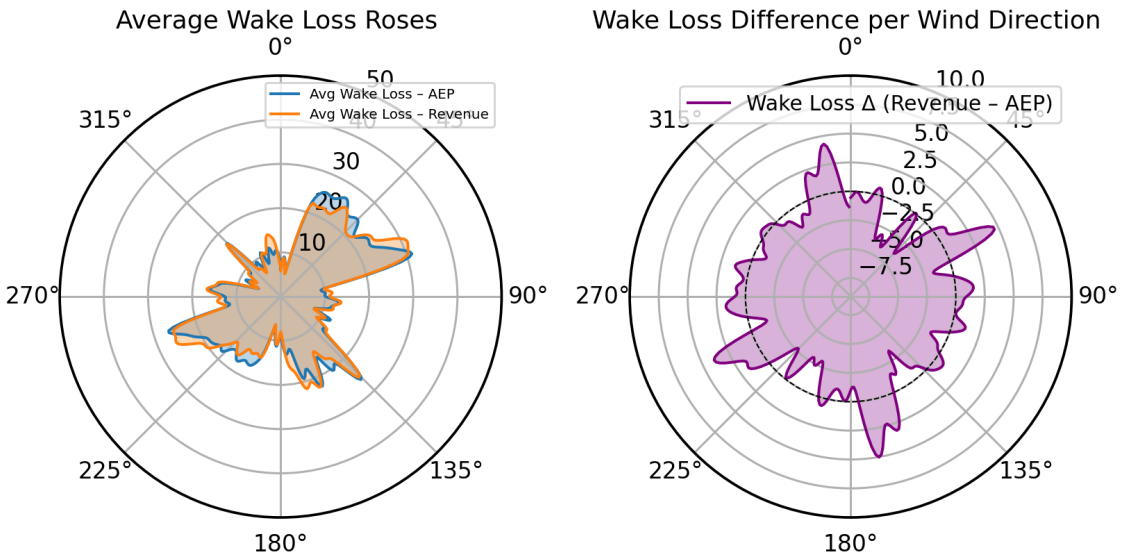

(a) France — left: averaged wake-loss roses (three best AEP layouts in blue; three best revenue layouts in orange); right: difference between wake roses in percents.

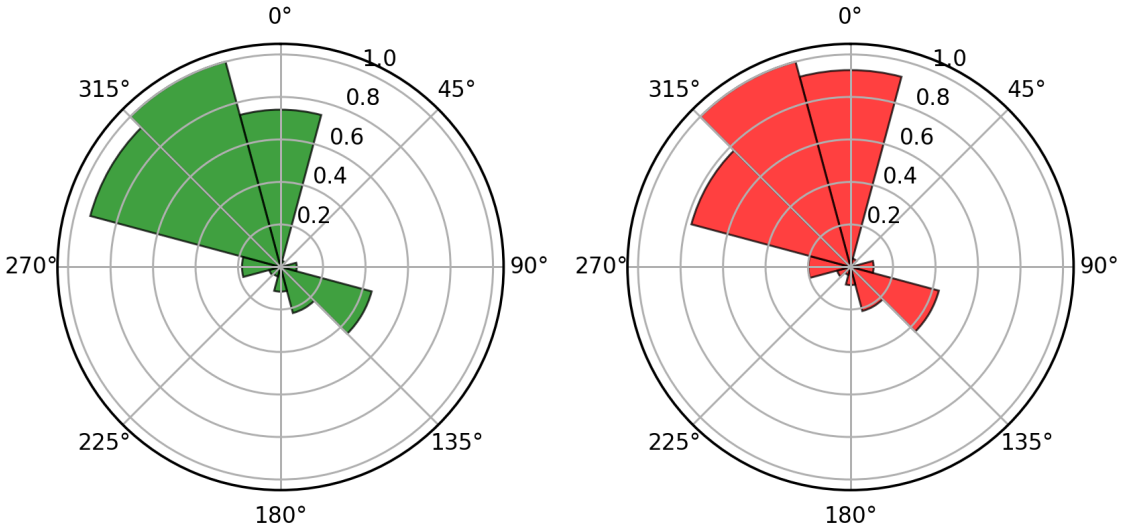

(b) France — normalized directional energy rose (left) and revenue rose (right) used to interpret the wake-rose differences.

Figure 4.7: Wake-rose comparison for France (ensemble average of three best layouts per objective).

The main observable trend is that, for both sites, the largest directional wake losses tend to occur in directions that correspond to little objective function potential (i.e. sectors with low energy or revenue in the objective function roses). This is evident in France for the 70° and 250° sectors (Figure 4.7a) and in Denmark near 0° (Figure 4.8a). Because the baseline production in these directions is small, even moderate absolute wake-induced reductions translate into substantial percentage values in Eq. (4.1); this demands cautious interpretation.

One interesting phenomenon is the very small wake losses exhibited for the French location near 180° (southerly inflow). Two factors explain this: (i) the farm-scale layouts are, for this site, arranged to mitigate wakes from the dominant northerly directions; because wake interactions have some symmetric properties, mitigation strategies that reduce wakes for northing winds also reduce wakes for southing winds (partial directional symmetry), and (ii) the baseline production for 180° is very small in the wind speed distribution, so the ratio in Eq. (4.1) is comparatively smaller than for other directions. The combination of partial symmetry in turbine placement and negligible baseline potential yields the observed near-zero wake-loss values at 180°.

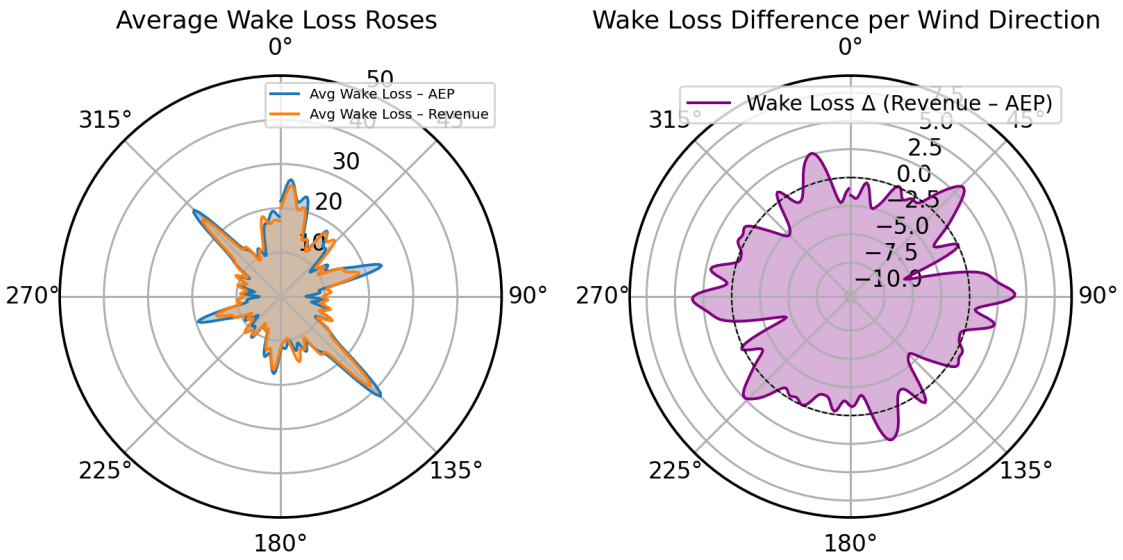

(a) Denmark — left: averaged wake-loss roses (three best AEP layouts in blue; three best revenue layouts in orange); right: difference between wake roses in percents

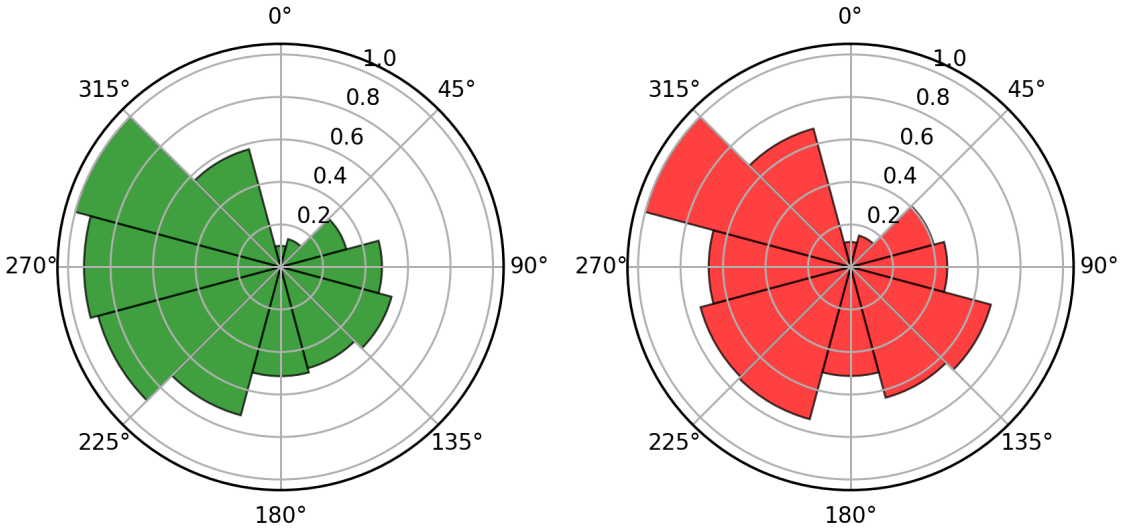

(b) Denmark — normalized directional energy rose (left) and revenue rose (right).

Figure 4.8: Wake-rose comparison for Denmark (ensemble average of three best layouts per objective).

To provide a more localized view, direction-binned histograms of wake loss are presented in Figure 4.9. Each bar corresponds to a direction bin; blue bars show the AEP ensemble, orange bars the revenue ensemble. These histograms clarify where AEP and revenue layouts differ most in their directional exposure to wakes.

In the French site, the histograms and roses reveal alternating intervals where AEP layouts have larger wake loss and intervals where revenue layouts lose more (notably in 15°–75° and 195°–255°). This pattern aligns with the angular shift between the energy and revenue roses (Figures 4.7b and 4.7a): AEP-optimised layouts tend to reduce wakes for directions where energy is concentrated, while revenue-optimized layouts shift toward directions that are more valuable in electricity prices. Because wake losses are generally largest perpendicular to main resource directions, a rotation of the preferred direction produces the observed alternating pattern.

For Denmark, the wake-loss distribution is more even; the histograms show that AEP layouts have larger losses at the highest-loss bins, whereas revenue layouts tend to show slightly larger losses in

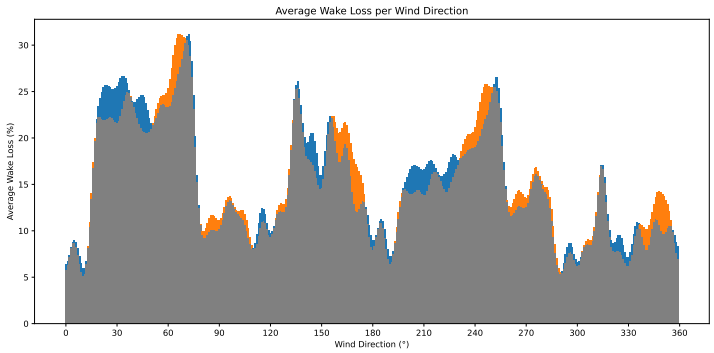

(a) France — histogram of directional wake losses (blue: AEP, orange: revenue, grey: areas of equal wake loss for AEP and revenue).

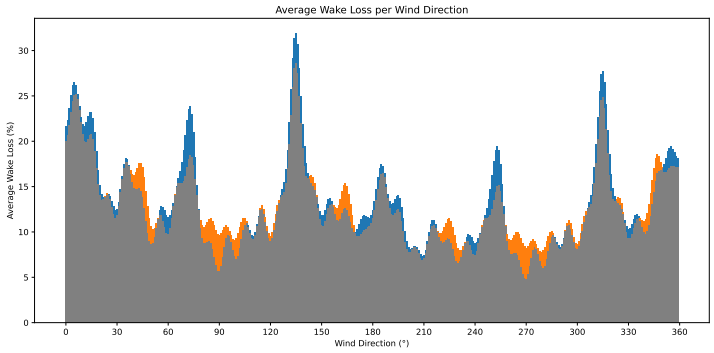

(b) Denmark — histogram of directional wake losses (blue: AEP, orange: revenue, grey: areas of equal wake loss for AEP and revenue).

Figure 4.9: Histograms of wake loss per direction for the averaged three-best layouts.

lower-loss bins. One plausible explanation is that revenue-optimised layouts in this site place relatively more emphasis on lower-speed regimes (where prices are favourable), producing a somewhat flatter wake-loss profile. This can be seen from the one-dimensional Wasserstein distance realtive to speed bins that was presented in Chapter 3 , where Denmark has the highest W_{dir} of 1.046.

4.1.3. Importance of Revenue Consideration

This next section aims to study differences between layouts having similar AEPs but different revenue performances. This is particularly interesting when taking the viewpoint of a wind farm developer. Indeed, even if the developer chooses to use a traditional AEP-based layout optimization algorithm, having an evaluation of expected revenues can help decide between layouts performing very similarly in AEP.

To illustrate why revenue matters even when AEP is almost unchanged, Figure 4.10 plots AEP vs. revenue for all optimization runs at the French site.

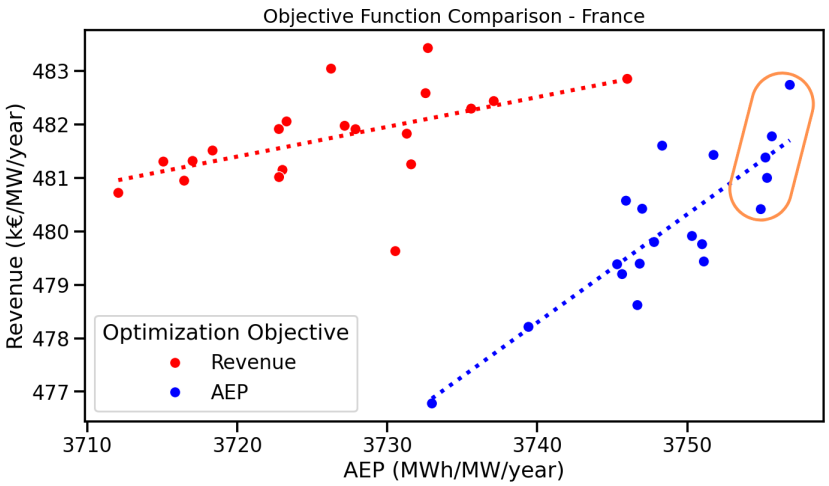

Figure 4.10: Selected layouts with similar AEP for comparison

The circled cluster contains layouts with nearly identical AEP (3754.9 to 3756.8 MWh/MW, only 0.05% variation) but differing revenues (480.415 to 482.74 k€/MW/yr, about 0.48%). In other words, layouts that are similar in energy production can differ by roughly 0.5% in expected revenue. Given the significant uncertainty associated with AEP, such revenue differences can significantly influence the choice among near-equal layouts.

Figure 4.11 compares the two extreme layouts from this cluster (highest vs. lowest revenue). Both achieve the same direction-averaged wake loss (≈11%), but their directional profiles tilt differently: the higher-revenue layout (left) favors winds near 0° (north), while the lower-revenue layout (right) favors winds near 270° (east–west). This remark is consistent with the pattern in Figure 4.7, and confirms that the difference in revenue performance between layouts is associated by this tilt of the wake rose. This phenomenon has now been observed both when comparing revenue- and AEP-optimized layouts, and when contrasting high-revenue and low-revenue layouts among the AEP-optimized cases.

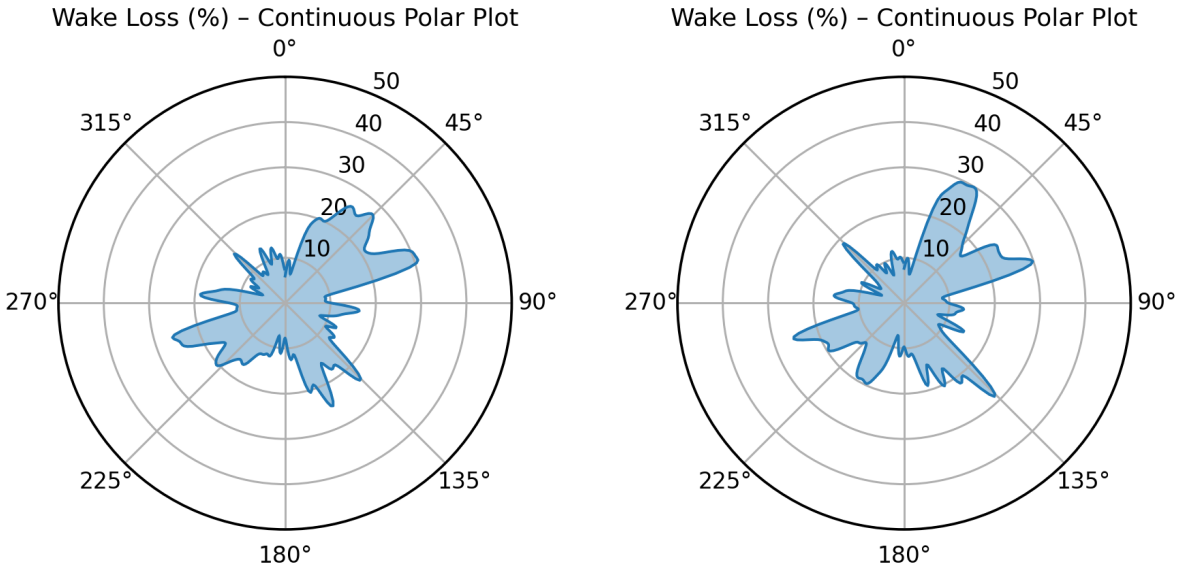

(a) France — Best among selected layouts

(b) France — Worst among selected layouts

Figure 4.11: Selection of wake rose examples for France. Best revenue layout on the left and worst revenue layout on the right

4.1.4. Results Summary

Optimization runs for the 2022–2024 period show that AEP and revenue are strongly correlated: layouts optimized for one objective generally perform well for the other. Revenue-based optimizations consistently yield slightly higher revenues, while AEP-based ones provide marginally higher energy output. The magnitude of these differences remains modest and varies by site. The largest improvement is observed in France, where switching from AEP to revenue optimization increases revenues by about 0.36% on average, with an associated AEP loss of roughly 0.60%. Denmark exhibits smaller but measurable gains, whereas the Netherlands and Germany show minimal differences. Spatially, both optimization strategies produce similar WTG configurations, typically placing turbines near site boundaries to mitigate wake effects. Nevertheless, wake rose analyses reveal directional differences for the French site.

These results broadly correspond to the 2D Wasserstein (EMD) metric: sites with higher W2D values, such as France, show larger differences between energy- and revenue-optimized layouts, while Denmark and Germany, with lower W2D values, exhibit smaller variations. The Netherlands is an exception, where a relatively high W2D does not translate into a significant performance gap, likely due to the symmetry of the shift between energy and revenue roses, which mitigates the impact of switching the optimization objective. This demonstrates both the usefulness and the limitations of the EMD metric. Overall, although average improvements are modest, explicitly incorporating revenue in the optimization can help refine top-performing layouts and support better economic decision-making for offshore wind projects.

4.2. Case Study 2

For this case study, the layouts are optimized for the time period 01/01/2015- 31/12/2020, before being evaluated in 01/01/2022- 31/12/2024 to be compared with the performances of the layouts from Case Study 1. The goal here is to assess how robust each objective function is when presented to real, yet-to-be-seen market and wind data.

4.2.1. Performance Comparison

The conditions for which these layouts have been optimized are relatively close to those of their performance evaluation. Although the study of the difference between conditions between the two time intervals was not pursued in this work, one can still find the corresponding energy and revenue roses in the Appendix B. The performance of these layouts evaluated in the time period 2022 to 2024 is displayed in Figures 4.12, 4.13, 4.14, and 4.15

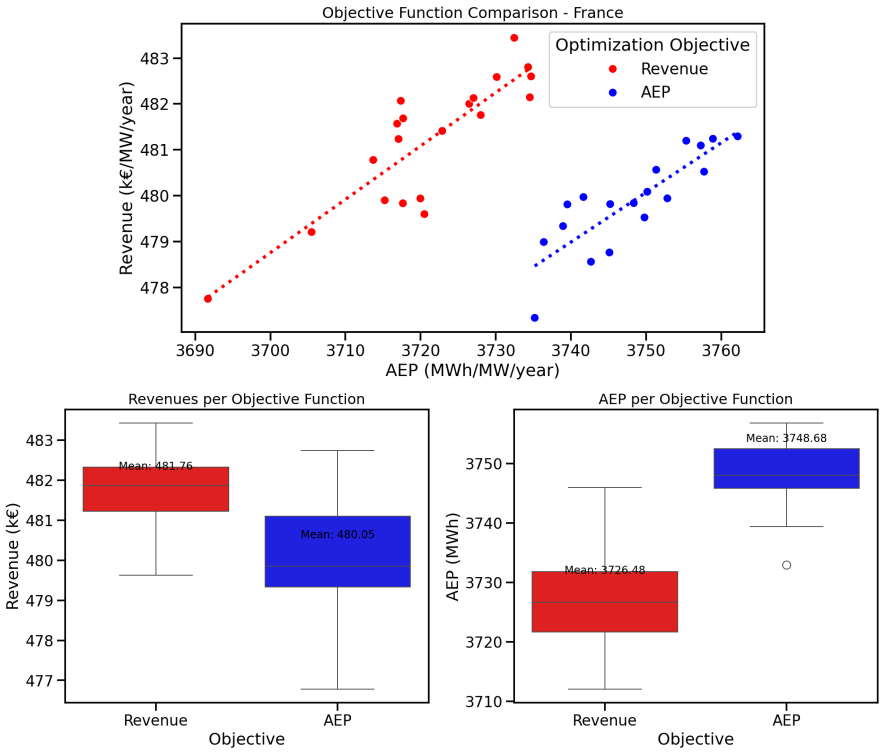

(a) France

Figure 4.12: Scatter plot and boxplot of optimization results for France (2015–2020).

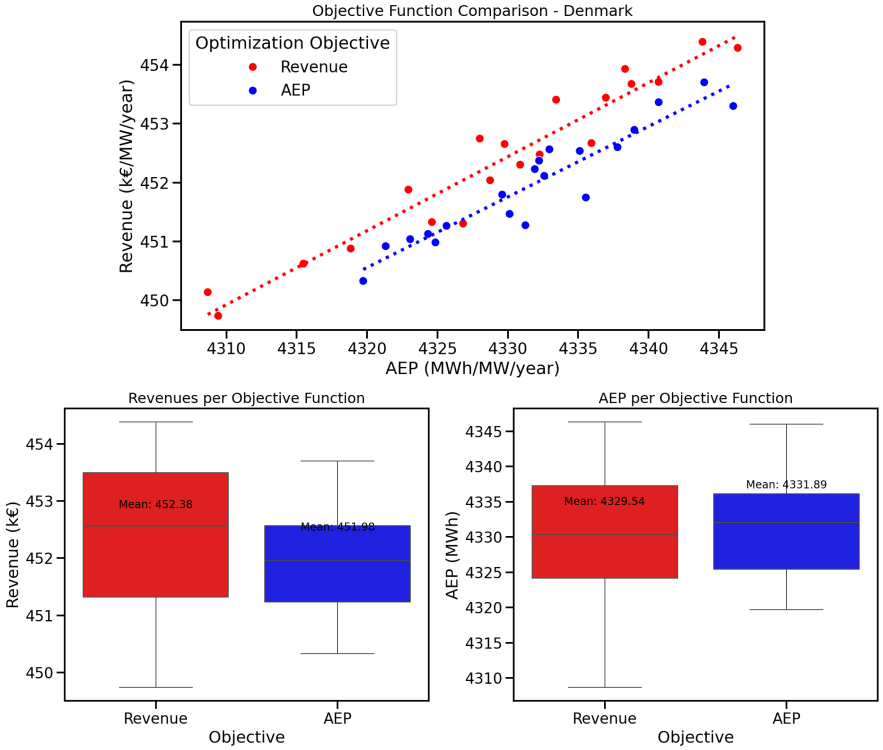

(a) Denmark

Figure 4.13: Scatter plot and boxplot of optimization results for Denmark (2015–2020).

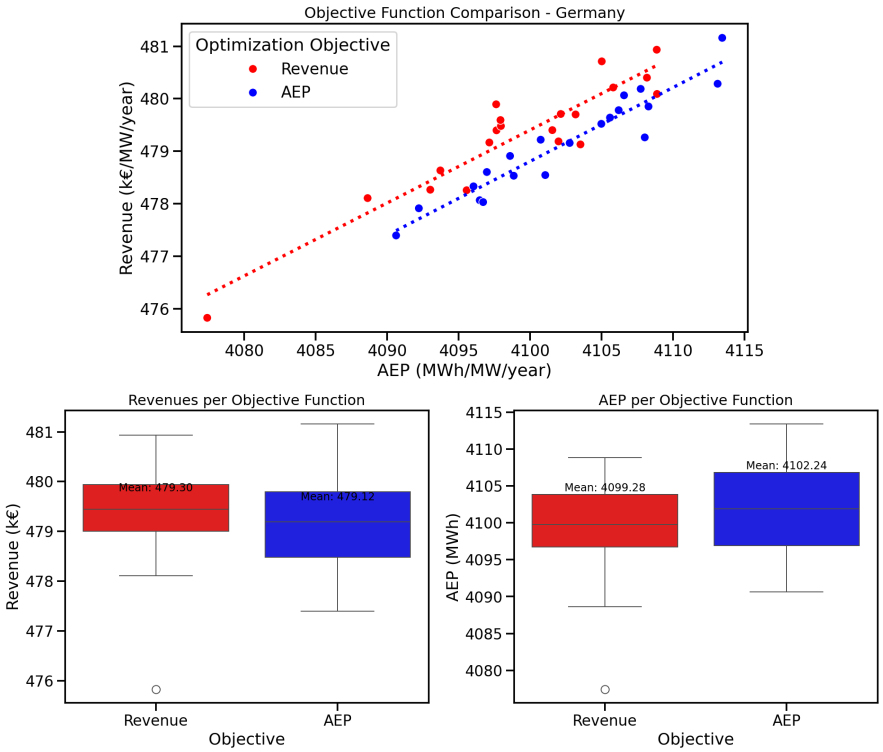

(a) Germany

Figure 4.14: Scatter plot and boxplot of optimization results for Germany (2015–2020).

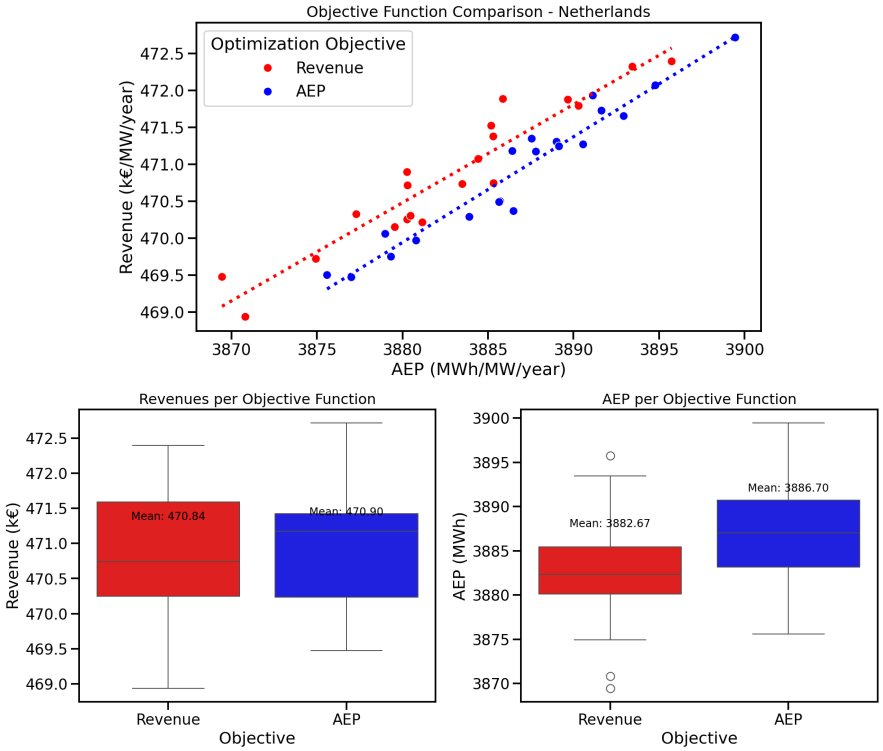

(a) Netherlands

Figure 4.15: Scatter plot and boxplot of optimization results for the Netherlands (2015–2020).

Figures 4.12, 4.13, 4.14, and 4.15 confirm that the positive AEP–revenue correlation persists for all sites. In particular, the regression lines from the French location are now nearly parallel (unlike in Case Study 1), indicating the previous anomaly was due to a few outlier layouts. Otherwise, the qualitative patterns mirror those of Case Study 1.

Tables 4.5 and 4.6 report the average revenue and AEP (with percentage change) for the top-three layouts at each site, evaluated in the time period 2022 to 2024. The results are similar to Case Study 1: revenue-based optimization still manages to yield a revenue gain (a few tenths of a percent) with a reasonable AEP sacrifice. This suggests that, even without knowing future conditions precisely, revenue-optimized layouts perform comparably to those optimized with full foresight when considering these two intervals. An interesting observation is that the magnitude of the trade-offs between AEP and revenue varies across sites. In particular, the French location shows a noticeably larger reduction in AEP when optimizing for revenue compared with the other sites (0.36% increase in revenue, and 0.67% decrease in AEP). However, this effect is not uniform: some sites show only minimal AEP penalties for revenue optimization, and no clear general trend can be established regarding the amplitude of these trade-offs.

Table 4.5: Comparison of average **revenue** (k€/MW/yr) for the three best layouts optimized for AEP vs. revenue (2015–2020 data), with relative improvement.

Site	AEP-optimized	Revenue-optimized	**Improvement (%)**
France	481.24	482.95	0.36%
Denmark	453.45	454.20	0.17%
Germany	480.54	480.68	0.03%
Netherlands	472.24	472.20	-0.01%

Table 4.6: Comparison of average **AEP** (MWh/MW/yr) for the three best layouts optimized for AEP vs. revenue (2015–2020 data), with relative loss.

Site	AEP-optimized	Revenue-optimized	**Loss (%)**
France	3759.61	3734.52	-0.67%
Denmark	4343.56	4343.63	0.02%
Germany	4111.60	4108.63	-0.07%
Netherlands	3895.72	3893.16	-0.07%

Revenue optimization requires both wind resource and electricity market data, whereas AEP optimization relies solely on wind resource information. A plausible concern is that increasing cannibalization effects over time might reduce the performance of revenue-optimized layouts under unforeseen future conditions. However, the time horizons considered in this case study are relatively short, and such structural market shifts may not yet be visible between the optimization and test periods. Within these limited timescales, the transfer functions between the energy potential and revenue potential roses appear to have remained consistent, which explains why revenue-optimized layouts retain performance comparable to those optimized under known conditions.

4.2.2. Performance Comparison between Case Studies 1 and 2

After evaluating the performance of these layouts under the conditions of Case Study 1, it is essential to compare them with those optimized directly for the 2022–2024 period. This comparison helps assess how crucial precise knowledge of future wind and market conditions is for achieving significant revenue improvements.

Figure 4.16 compares the performance of layouts from Case Study 2 (2015–2020 optimization) with those optimized on 2022–2024 data (Case 1). Panel (a) shows the average revenue for each country under both optimization intervals: the 2015–2020 layouts (blue bars) achieve almost the same revenue as the 2022–2024 layouts (orange bars). Revenue losses for the older layouts are tiny (e.g. 0.02% in Germany, up to 0.12% in Denmark). This indicates that the slight shift in input data has a negligible effect on layout revenue. Panel (b) illustrates the worst-case France example: even the layout with the largest drop shows only a modest change in its wake-loss profile.

Figure 4.16 compares the performance of layouts from Case Study 2 (optimized for 2015–2020, blue bars) with those optimized for 2022–2024 (Case Study 1, orange bars). AEP differences remain minimal, with variations below 0.1% across all countries. These small performance variations seem to be caused by individual optimization performances, which could be expected as the potential energy profile is very similar between the two time periods.

Panel (b) shows the comparison of average revenue performance for the same layouts. Revenue losses for layouts optimized with the earlier period remain very small, ranging from 0.02% in Germany to 0.12% in Denmark. This indicates that the moderate shift in underlying wind and market data between these two intervals has only a minor impact on overall revenue performance.

Overall, these results suggest that layouts optimized for an earlier period maintain strong performance under updated conditions. This reinforces the robustness of the optimization framework: close deviations in wind or market inputs do not significantly alter the relative performance of AEP- or revenue-optimized layouts.

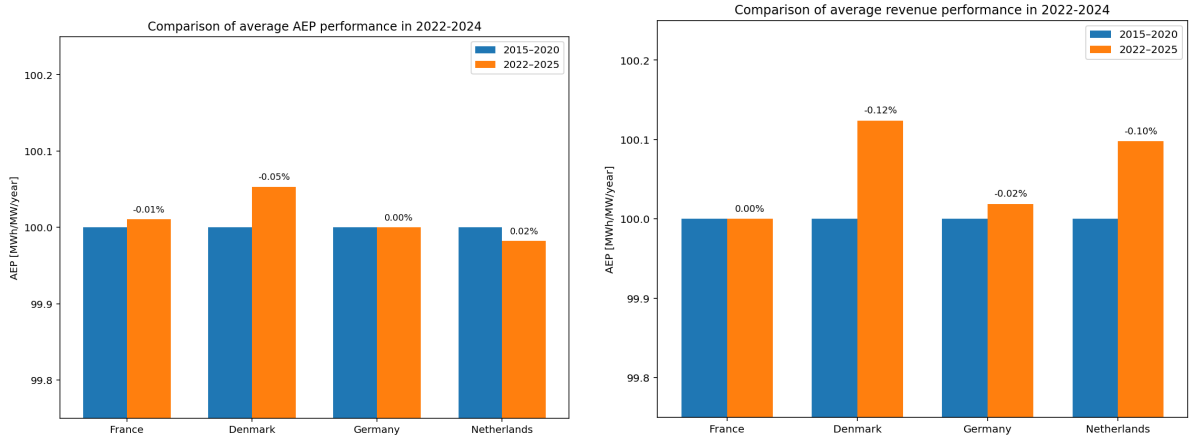

(a) Comparison of average AEP performance per location and per interval of optimization

(b) Comparison of average revenue performance per location and per interval of optimization

Figure 4.16: France - comparison of performance in 2022-2024 across different optimization intervals

4.2.3. Results Summary

Case Study 2 confirms that layouts optimized on historical data remain effective when evaluated under slightly different future conditions. Across all sites, both revenue and AEP changes stayed below about 0.12%, showing that small updates in wind and market data do not significantly degrade performance. This means that, at least over short time horizons, the optimization framework is robust: a layout optimized using several recent years of data remains nearly optimal when exposed to new but similar conditions.

It is important to note, however, that this stability is observed because the optimization and test periods are close to one another, and the underlying market structure appears to have changed little between them. Over longer time spans, where structural changes such as stronger cannibalization effects could emerge, larger performance deviations could be expected.

Another practical implication is that revenue optimization, even though it requires additional price data, did not prove more fragile than AEP optimization. Both approaches maintained their performance, suggesting that integrating revenue considerations does not make layouts overly sensitive to moderate uncertainty in market inputs. This is encouraging for projects exposed to merchant markets: a revenue-based approach appears at least as stable as the traditional AEP-driven design while offering small but consistent economic advantages.

4.3. Case Study 3

This case study imagines the thought process of a wind farm developer aiming to design an offshore wind farm optimized for future revenue generation. The objective is to assess how predictions about

future market conditions influence the expected revenues of the farm.

4.3.1. Market Scenario Creation

To build the necessary hourly market price scenarios, the methodology from 2.5 is applied.

This study focuses on the impact of wind energy penetration on electricity prices and revenues, particularly the cannibalization effect. Therefore, the primary variable for scenario creation is the correlation factor ρ. The mean electricity price (μ) and coefficient of variation (CV) are held constant across scenarios: CV is set to 0.25, based on Mehta et al. [17], and the yearly mean electricity price is set at 60€/MWh. This simplification avoids complexities related to inflation or Net Present Value (NPV) considerations.

Three correlation factors are selected to represent different scenarios:

- $\rho = 0$: No cannibalization effect, with no correlation between electricity prices and wind speed.
- $\rho = -0.3$: Moderate correlation, reflecting a scenario where the electricity grid significantly depends on wind energy (e.g., Denmark's current yearly averaged correlation is approximately -0.37) [17].
- $\rho = -0.6$: High correlation, indicating a future scenario with substantial wind energy penetration.

The choice was made to use the Danish location for this Case Study because of the good performance of revenue-based layouts at this location, and also the standard shape of the wind profile. Indeed, while the French location seemed to raise higher differences between AEP and revenues, the singularity of the energy and revenue roses did not make it the most suitable choice for a more general analysis

The period 01/01/2035–31/12/2040 is arbitrarily chosen as an objective period for this case study, the aim being to have a relatively large time interval, corresponding to an operation period of a future wind farm project. The actual interval was selected randomly, and this doesn't impact the study. Wind data from Denmark (2015–2020) are used, assuming similar wind resource conditions for 2035–2040. This choice is based on the consistent performance of AEP-based layouts in Denmark and the standard wind profile observed.

The constructed scenarios, based on different values of the correlation coefficient ρ, are displayed in Figures 4.17a and 4.17b.

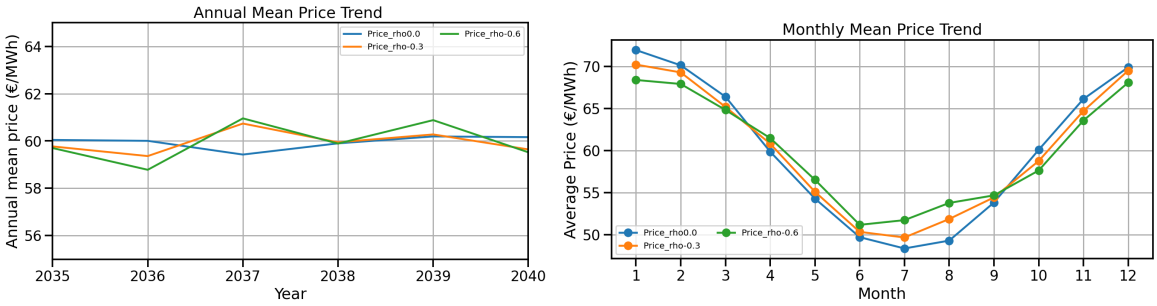

(a) Yearly average and seasonal variation of electricity prices.

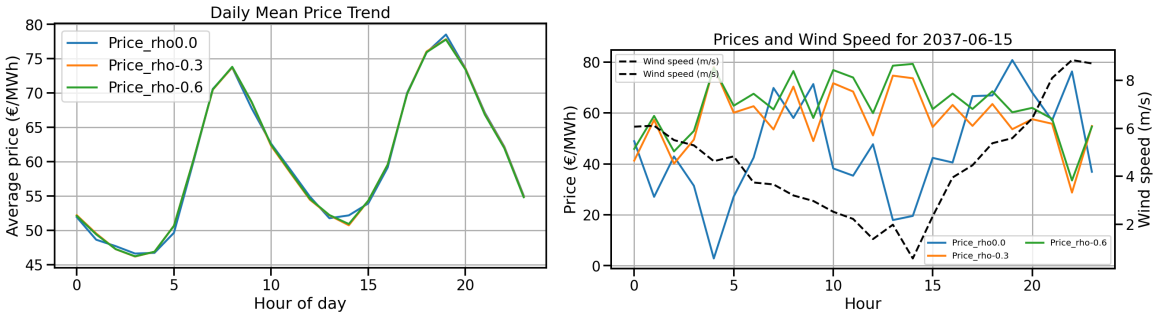

(b) Average wind prices per hour and wind speed-price relationship for each scenario.

Figure 4.17: Constructed market scenarios based on varying correlation coefficients between wind speed and electricity prices.

4.3.2. Results and Discussion

Once the three market scenarios are defined, wind farm layouts are optimized for the scenario with a correlation coefficient of -0.3 (the baseline scenario). The other two scenarios (zero correlation and a stronger -0.6 correlation) serve as boundary cases: a 0 correlation scenario corresponding to a case with no cannibalization (wind generation has no impact on prices), and a -0.6 correlation scenario representing a future with roughly double the current cannibalization effect. Optimizing the layout for the -0.3 scenario and then evaluating its performance under these different scenarios aims at mimicking the strategy of a wind farm developer who designs for the most likely conditions but wants to assess how the design would perform if market dynamics turn out less favorably or more favorably than expected.

In addition, an extra test scenario is generated using the same -0.3 correlation but with a different random seed for the stochastic price perturbations (denoted as the -0.3 correlation (test) scenario). This test scenario represents a situation where the expected long-term cannibalization effect is the same as in the baseline, but the exact hourly price fluctuations differ. By comparing the baseline and test scenarios, we can gauge the impact of unpredictable short-term price variations on the optimized layout's revenue. (A detailed description of this test scenario is provided in Appendix C, similar to the format of Figure 4.17.) The optimized layout from the baseline -0.3 scenario is evaluated across all these scenarios, and the resulting average annual revenues per unit capacity are summarized in Table 4.7.

The first observation from Table 4.7 is that the revenues under all scenarios are very close to one another (within roughly 0.5%). In particular, the baseline -0.3 optimized layout achieves 257.34 k€/MW/year in its own design scenario, and in the perturbed -0.3 correlation (test) scenario it achieves 257.11 k€/MW/year. This difference of about 0.09% in revenue is negligible, indicating that the inability to precisely predict hourly prices has virtually no impact on the optimized layout's annual revenue. This analysis is also supported by the previous section and Figure 4.16 In other words, small random deviations in hourly electricity prices do not meaningfully reduce the revenue of a layout that was optimized for the correct long-term price–wind correlation. This confirms that when performing wind farm layout optimization for day-ahead market revenue, perfect foresight of hourly price patterns is not essential. A layout optimized using the expected statistical relationship between wind and price remains near-

Table 4.7: Average annual revenues of the layout optimized for the -0.3 correlation scenario, evaluated under different scenarios.

Evaluation Scenario	Average Revenue (k€/MW/year) of the -0.3 optimized layouts
Baseline -0.3 correlation scenario	257.34
- 0.3 correlation (test)	257.11
0 correlation scenario	257.67
-0.6 correlation scenario	256.51

optimal even if real hourly prices fluctuate unpredictably around that expectation.

More noticeable are the revenue changes when the wind–price correlation differs from the design assumption of -0.3. As expected, the layout optimized for the -0.3 scenario yields the highest revenue in the 0 correlation scenario, at 257.67 k€/MW/year, which is about 0.13% higher than in the baseline case. This slight increase occurs because a zero correlation scenario imposes no cannibalization penalty on high wind output. Under these conditions, times of strong wind are not associated with lower electricity prices. The optimized layout was originally tuned to favor wind conditions that produce good revenue under moderate cannibalization; when that penalty is removed (correlation = 0), those same wind conditions become even more profitable, hence a modest increase in revenue is observed.

Conversely, under the -0.6 correlation scenario (stronger cannibalization), the layout's revenue drops to 256.51 k€/MW/year, about a 0.32% decrease relative to the baseline. In this scenario, high wind periods incur even steeper price drops than assumed in the design, so the layout's emphasis on those periods yields slightly lower returns. These results demonstrate that the assumed correlation between wind generation and price is an important factor for revenue optimization: changing the correlation has a more pronounced impact on revenue than the random price noise does.

Finally, we compare the revenue performance of a layout optimized for revenue (under the -0.3 scenario) with that of a layout optimized purely for Annual Energy Production (AEP), to highlight the benefit of including price effects in the optimization. Table 4.8 summarizes the average revenues for both the revenue-optimized and AEP-optimized layouts across the scenarios, along with the percentage revenue loss when using the AEP-optimized layout. Notably, the revenue-optimized layout achieves slightly higher revenues in every scenario. The advantage is small(on the order of 0.05%), but it is consistent. Even in scenarios that differ from the one used for optimization, the layout tailored for market revenue maintains a marginal lead. This indicates that accounting for the price–wind correlation in the design (even an uncertain future correlation) can consistently improve financial performance compared to a layout optimized only for AEP. In practical terms, a wind farm developer stands to gain a modest but reliable increase in expected revenue by using a revenue-optimization approach as opposed to a traditional AEP-maximization approach, even if the future market conditions deviate from the predictions.

Table 4.8: Performance of the revenue-optimized vs. AEP-optimized layouts across different scenarios. All values are expected annual revenues in k€/MW/year. The Loss column indicates how much less revenue (in percent) the AEP-optimized layout earns relative to the revenue-optimized layout for the given scenario.

Correlation	Revenue-Optimized	AEP-Optimized	Loss
0.3 (baseline)	257.34	257.21	0.05%
0.3 (test seed)	257.11	256.99	0.05%
0	257.67	257.55	0.05%
-0.6	256.51	256.38	0.05%

In summary, this case study suggests that a wind farm layout optimized for revenue under a realistic baseline scenario is highly robust to both short-term price volatility and to plausible shifts in the long-

term wind–price correlation. The extreme scenarios considered (one with zero cannibalization and one with a doubled cannibalization effect) bound the range of potential outcomes, yet the optimized layout's performance varies only marginally across this range. In real-world project development, one would expect more narrowly informed projections for the future cannibalization effect (for example, an expected correlation coefficient somewhere between 0 and -0.6 rather than these extremes). Such informed estimates would further reduce the uncertainty in revenue outcomes. Importantly, the results consistently show that explicitly optimizing for market revenue yields a small but consistent benefit in expected revenues compared to optimizing for AEP alone. Therefore, beyond the precise prediction of hourly prices, the key factor for revenue-maximizing layout design is anticipating the strength of the wind-price correlation in the market, and incorporating that expectation into the optimization process.

5

Conclusion

This study investigates how optimizing offshore wind farm layouts for expected day-ahead market revenues compares with the traditional approach based on AEP. The analysis combines historical wind and electricity price data with an advanced layout optimization framework to quantify potential performance differences between the two objectives. Through three complementary case studies, the research examines (1) the direct comparison of AEP- and revenue-based layouts under known conditions, (2) the robustness of revenue-optimized layouts under unseen wind and price scenarios, and (3) the sensitivity of results to future market predictions and cannibalization effects. Together, these analyses provide a comprehensive understanding of whether revenue-aware optimization can improve financial outcomes and enhance the efficiency of offshore wind farms operating in merchant electricity markets.

5.1. Answers to the Research Questions

Does optimizing an offshore wind farm layout for expected day-ahead market revenue (instead of AEP) improve performance?

Yes — revenue-based layouts consistently achieve slightly higher revenues than AEP-based ones, while leaving energy production acceptable. Across all case studies, the revenue gains remain small (typically on the order of 0.05–0.35%), and the associated AEP trade-offs remain below 1%. These improvements are modest but systematic: the revenue objective provides a subtle but measurable economic advantage without compromising energy production.

How do market conditions influence revenue-based layout optimization, and how robust are revenue-optimized layouts to out-of-sample wind and market conditions?

Revenue-optimized layouts show strong robustness to both short-term and structural uncertainties in market conditions. When tested on unseen wind and price years, layouts retain their performance with objective-function losses below 0.15%. Even when tested against alternative future market assumptions regarding the wind–price correlation — a critical factor driving price cannibalization — the impact remains small. For example, layouts optimized under a baseline correlation $\rho = -0.3$ achieve 257.34 k€/MW/year in its design scenario and 257.11 k€/MW/year when hourly price fluctuations are perturbed with a random noise (a negligible 0.09% reduction). When the price–wind correlation itself is altered, the same layouts perform slightly better under no cannibalization ($\rho = 0$, 257.67 k€/MW/year, +0.13%) and slightly worse under stronger cannibalization ($\rho = -0.6$, 256.51 k€/MW/year, −0.32%). These small differences confirm that layouts optimized for revenue remain efficient across a wide range of plausible market conditions. It is also shown that these revenue-optimized layouts consistently outperform AEP-optimized layouts on future scenarios.

5.2. Practical Implications for Stakeholders

For developers and operators, revenue-aware layout optimization offers an enhancement to conventional design. Even though the revenue differences compared to AEP optimization are modest, they are consistently positive across all tested scenarios — including those deviating from the market conditions used in optimization. This means that a developer optimizing for revenue is unlikely to be worse off, and can expect slightly higher project income without sacrificing significant energy yield. Importantly, exact hourly price forecasts are not needed; robust layouts can be obtained by using statistical market relationships (such as expected wind–price correlation).

For policy makers and market designers, the findings emphasize that while layout optimization can marginally increase project revenue, long-term market structures remain the dominant driver of economic performance. Market frameworks that reduce uncertainty about the wind–price relationship would help developers plan and finance merchant-exposed projects with greater confidence. Encouraging the use of revenue objectives can also improve the overall resilience of new offshore wind developments.

5.3. Further Work

Several extensions could strengthen and expand the findings of this study. First, refining the version SGD optimization algorithm, including hyperparameter tuning and convergence diagnostics, could help ensure that layouts found are close to global optima. Indeed, the tuning of this algorithm was performed in this study with a limited number of test runs, due to time availability.

Second, future research could focus on predictive indicators of when revenue-based optimization is likely to yield meaningful gains over AEP. Identifying site and market characteristics that correlate with larger revenue improvements would help practitioners decide when this additional modeling step is justified. The 2D Wasserstein distance proposed in this work does not account for the symmetric component of the wind resource in layout optimization. Future work should first assess the significance of this symmetry and, if it proves relevant, adapt the 2D Wasserstein distance or develop an alternative metric to incorporate its effect.

Finally, future scenario analysis of Case Study 3 could be expanded. While this work evaluated layouts optimized for one correlation assumption across a few alternative futures, optimizing layouts for each scenario separately and cross-comparing them would clarify how wrong correlation assumptions translate to lost revenue. A study of this type could help conclude as to what would be the optimal design strategy for a wind farm developer facing uncertainties in future cannibalization (e.g., is the best overall performance found with low or high correlation coefficients optimizations). It could also be helpful to extend this case study to other locations, to evaluate the impact of the site on the expected revenues yielded by the layouts.

5.4. Closing Statement

Revenue-aware layout optimization is a low-risk enhancement to standard offshore wind farm design practice. Although the financial improvements are small in absolute terms, they are systematic and require little additional market foresight beyond an expected price–wind correlation. The approach produces layouts that remain efficient under realistic price and wind variability and even under plausible shifts in market structure. As offshore wind transitions further into competitive merchant markets, integrating market-based revenue optimization alongside energy production-based design can help developers reduce financial risk and improve long-term project resilience.

References

[1] Enerdata <OR> I. Surname, I. Surname, and I. Surname. *Global wind installations reached a record high in 2023 with 117 GW added*. 2024. URL: https://www.enerdata.net/publications/daily-energy-news/global-wind-installations-reached-record-high-2023-117-gw-added.html (visited on 02/14/2025).

[2] L. Hirth. "The market value of variable renewables, The effect of solar wind power variability on their relative price". In: *Energy Economics* 38 (2013). URL: https://doi.org/10.1016/j.eneco.2013.02.004.

[3] L. Reichenberg, T. Ekholm, and T. Boomsma. "Revenue and risk of variable renewable electricity investment: The cannibalization effect under high market penetration". In: *Energy* 284 (2023). URL: https://doi.org/10.1016/j.energy.2023.128419.

[4] T. Nguyen et al. "Offshore Wind Farm Layout Optimization Accounting for Participation to Secondary Reserve Markets". In: *Wind Energy Science* (2024). URL: https://doi.org/10.5194/wes-2024-131.

[5] B. Pérez, R. Mínguez, and R. Guanche. "Offshore wind farm layout optimization using mathematical programming techniques". In: *Renewable Energy* (2013). URL: https://doi.org/10.1016/j.renene.2012.12.007.

[6] P. Gebraad et al. "Maximization of the Annual Energy Production of Wind Power Plants by Optimization of Layout and Yaw-Based Wake Control". In: *BYU ScholarsArchive* (2016). URL: http://hdl.lib.byu.edu/1877/3679.

[7] A. P. J. Stanley et al. "Objective and algorithm considerations when optimizing the number and placement of turbines in a wind power plant". In: *Wind Energy Science* 6.5 (2021). URL: 10.5194/wes-6-1143-2021.

[8] G. Gualtieri. "A novel method for wind farm layout optimization based on wind turbine selection". In: *Energy Conversion and Management* 193 (2019). URL: https://doi.org/10.1016/j.enconman.2019.04.059.

[9] P. Ziyaei et al. "Minimizing the levelized cost of energy in an offshore wind farm with non-homogeneous turbines through layout optimization". In: *Ocean engineering* 249 (2022). URL: https://doi.org/10.1016/j.oceaneng.2022.110859.

[10] S. Cox. *Financiers must be reassured offshore wind can still guarantee a good return*. 2020. URL: https://www.rechargenews.com/circuit/financiers-must-be-reassured-offshore-wind-can-still-guarantee-a-good-return/2-1-778801 (visited on 04/23/2025).

[11] O. Catsaros. *Global Cost of Renewables to Continue Falling in 2025 as China Extends Manufacturing Lead: BloombergNEF*. 2025. URL: https://about.bnef.com/blog/global-cost-of-renewables-to-continue-falling-in-2025-as-china-extends-manufacturing-lead-bloombergnef/?utm_source=chatgpt.com (visited on 04/23/2025).

[12] B. H. Meijer. *Dutch test the water for subsidy-free offshore wind farms*. 2017. URL: https://www.reuters.com/article/business/environment/dutch-test-the-water-for-subsidy-free-offshore-wind-farms-idUSKBN1E91RH/?utm_source=chatgpt.com (visited on 04/23/2025).

[13] Offshore Energy. *Report: Ørsted Not Competing in Dutch Zero Subsidy Offshore Wind Tender*. 2017. URL: https://www.offshore-energy.biz/report-orsted-not-competing-in-dutch-zero-subsidy-offshore-wind-tender/?utm_source=chatgpt.com (visited on 04/23/2025).

[14] Reuters. *German agency awards offshore wind licences for 2026 at zero subsidy*. 2021. URL: https://www.reuters.com/business/energy/german-agency-awards-offshore-wind-licences-2026-zero-subsidy-2021-09-09/?utm_source=chatgpt.com (visited on 04/23/2025).

[15] E. Loth et al. “Why we must move beyond LCOE for renewable energy design”. In: *Advances in Applied Energy* 8 (2022). URL: 10.1016/j.adapen.2022.100112.

[16] K. Dykes. “Optimization of Wind Farm Design for Objectives Beyond LCOE”. In: *Journal of physics conference series* (2020). URL: http://dx.doi.org/10.1088/1742-6596/1618/4/042039.

[17] M. K. Mehta, M. Zaaijer, and D. von Terzi. “Designing wind turbines for profitability in the day-ahead market”. In: *Wind Energy Science* 9.12 (2024). URL: 10.5194/wes-2024-43.

[18] J. Zhang and Y. Jiang. “Joint optimization of the number, type and layout of wind turbines for a new offshore wind farm”. In: *J. Renewable Sustainable Energy* 12.5 (2020). URL: https://doi.org/10.1063/5.0020204.

[19] T. Kunakote et al. “Comparative Performance of Twelve Metaheuristics for Wind Farm Layout Optimisation”. In: *Archives of Computational Methods in Engineering* 29 (2022). URL: https://doi.org/10.1007/s11831-021-09586-7.

[20] J. Feng and W. Z. Shen. “Solving the wind farm layout optimization problem using random search algorithm”. In: *Renewable Energy* 78 (2015). URL: http://dx.doi.org/10.1016/j.renene.2015.01.005.

[21] H. Long, P. Li, and W. Gu. “A data-driven evolutionary algorithm for wind farm layout optimization”. In: *Energy* 208 (2020). URL: https://doi.org/10.1016/j.energy.2020.118310.

[22] J. Quick et al. “Stochastic gradient descent for wind farm optimization”. In: *Wind Energy Science* 8.8 (2023). URL: 10.5194/wes-8-1235-2023.

[23] M. Baricchio, P. M. O. Gebraad, and J. W. van Wingerden. “Evaluating the potential of a wake steering co-design for wind farm layout optimization through a tailored genetic algorithm”. In: *Wind Energy Science* 9 (2024). URL: https://doi.org/10.5194/wes-9-2113-2024.

[24] J. J. Thomas et al. “A comparison of eight optimization methods applied to a wind farm layout optimization problem”. In: *Wind Energy Science* 8 (2023). URL: https://doi.org/10.5194/wes-8-865-2023.

[25] F. Paraschiv, D. Erni, and R. Pietsch. “The impact of renewable energies on EEX day-ahead electricity prices”. In: *Energy Policy* 73 (2014). URL: https://doi.org/10.1016/j.enpol.2014.05.004.

[26] S. Djørup, J. Z. Thellufsen, and P. Sorknæs. “The electricity market in a renewable energy system”. In: *Energy* 162 (2018). URL: https://doi.org/10.1016/j.energy.2018.07.100.

[27] L. Blickwedel et al. “Future economic perspective and potential revenue of non-subsidized wind turbines in Germany”. In: *Wind Energy Science* 6 (2021). URL: https://doi.org/10.5194/wes-6-177-2021.

[28] S. Busch et al. “The Development of Renewable Energy in the Electricity Market”. In: *European Commission, Economy and Finance* (2023). URL: https://economy-finance.ec.europa.eu/publications/development-renewable-energy-electricity-market_en.

[29] M. Jansen et al. “Offshore wind competitiveness in mature markets without subsidy”. In: *Nature Energy* 5 (2020). URL: https://doi.org/10.1038/s41560-020-0661-2.

[30] Centrales Next. *Qu’est-ce que le M0 ?* 2024. URL: https://www.centrales-next.fr/glossaire-energies-renouvelables/m0-mzero (visited on 02/21/2025).

[31] B. Mauch, P. M. S. Carvalho, and J. Apt. “Can a wind farm with CAES survive in the day-ahead market”. In: *Energy Policy* 48 (2012). URL: https://doi.org/10.1016/j.enpol.2012.05.061.

[32] A. S. Sætherø. “Hourly Price Forward Curves for Electricity Markets”. In: (2017). URL: https://duepublico2.uni-due.de/servlets/MCRFileNodeServlet/duepublico_derivate_00044584/DissASSaethroe.pdf.

[33] Kiesel R., Paraschiv F., and Sætherø A. “On the construction of hourly price forward curves for electricity prices”. In: *Computational Management Science* 16 (2019), pp. 345–369. URL: 10.1007/s10287-018-0300-6.

[34] F. Ziela and R. Steinert. “Probabilistic mid- and long-term electricity price forecasting”. In: *Renewable and Sustainable Energy Reviews* 94 (2018). URL: https://doi.org/10.1016/j.rser.2018.05.038.

[35] R. A. de Marcos, J. Reneses, and A. Bello. “Long-Term Spanish Electricity Market Price Forecasting with Cointegration and VEC Models”. In: *Institute of Electrical and Electronics Engineers* (2016). URL: https://doi.org/10.1109/PMAPS.2016.7764158.

[36] R. Weron. “Electricity price forecasting: A review of the state-of-the-art with a look into the future”. In: *International Journal of Forecasting* 30 (2014). URL: https://doi.org/10.1016/j.ijforecast.2014.08.008.

[37] P. Gabrielli et al. “Data-driven modeling for long-term electricity price forecasting”. In: *Energy* 244 (2022). URL: https://doi.org/10.1016/j.energy.2022.123107.

[38] B. Oh, D. Lee, and D. Lee. “Oil-Price Based Long-Term Hourly System Marginal Electricity Price Scenario Generation”. In: *Institute of Electrical and Electronics Engineers* 10 (2022). URL: https://ieeexplore.ieee.org/document/9724189.

[39] J. Gea-Bermúdez et al. “Day-Ahead Market Modelling of Large-Scale Highly-Renewable Multi-Energy Systems: Analysis of the North Sea Region towards 2050”. In: *Energies* 14 (2021). URL: https://doi.org/10.3390/en14010088.

[40] R. Riva et al. *Welcome to TOPFARM*. URL: https://topfarm.pages.windenergy.dtu.dk/TopFarm2/index.html.

[41] IEA Wind TCP. *IEA Wind TCP Task 37*. 2021. URL: https://iea-wind.org/task37/ (visited on 05/05/2025).

[42] IEA Wind TCP. *IEA 10 MW Trubine Github*. 2020. URL: https://nrel.github.io/turbine-models/IEA_10MW_198_RWT.html (visited on 05/05/2025).

[43] EMD. *EMD mesoscale time series*. 2025. URL: https://www.emd-international.com/data-services/mesoscale-time-series/ (visited on 05/05/2025).

[44] Ember. *European Wholesale Electricity Price Data*. 2025. URL: https://ember-energy.org/data/european-wholesale-electricity-price-data/ (visited on 05/05/2025).

[45] Eurostat. *Inflation rate*. 2025. URL: https://ec.europa.eu/eurostat/databrowser/view/tec00118/default/table?lang=en (visited on 05/05/2025).

[46] DTU. *Ørsted’s TurbOPark implementation in PyWake*. 2025. URL: https://topfarm.pages.windenergy.dtu.dk/PyWake/notebooks/literature_verification/TurbOPark.html (visited on 05/05/2025).

[47] D. Van Binsbergen et al. “Performance comparison of analytical wake models calibrated on a large offshore wind cluster”. In: *Journal of Physics: Conference Series* 2767 (2024). URL: https://researchportal.vub.be/en/publications/performance-comparison-of-analytical-wake-models-calibrated-on-a-?utm_source=chatgpt.com.

[48] Wikipedia. *Earth mover’s distance*. 2025. URL: https://en.wikipedia.org/wiki/Earth_mover%27s_distance#:~:text=In%20computer%20science%2C%20the%20earth,over%20which%20it%20is%20moved (visited on 05/05/2025).

[49] Joseph C. Y. Lee and M. Jason Fields. “An overview of wind-energy-production prediction bias, losses, and uncertainties”. In: *Wind Energy Science* (2021). URL: https://doi.org/10.5194/wes-6-311-2021.

A

KDE plots

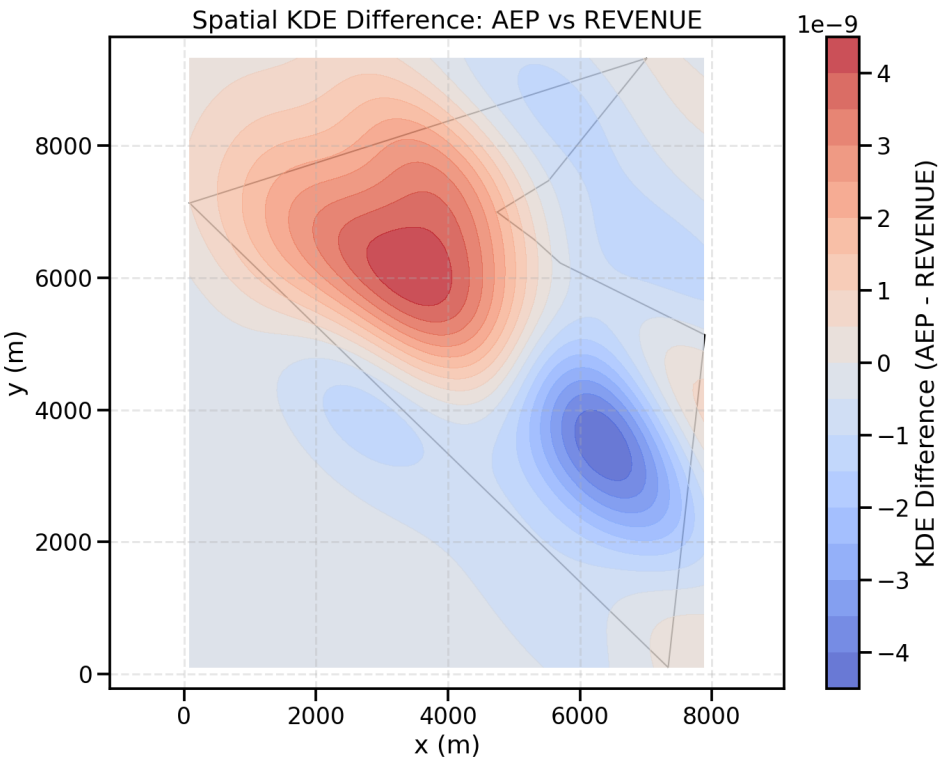

Figure A.1: Kernel Density Estimation of the French location, estimating spatial positioning preferences per objective function, red zones corresponding to areas favored by Revenue-based optimizations, and blue zones favored by AEP-based optimizations

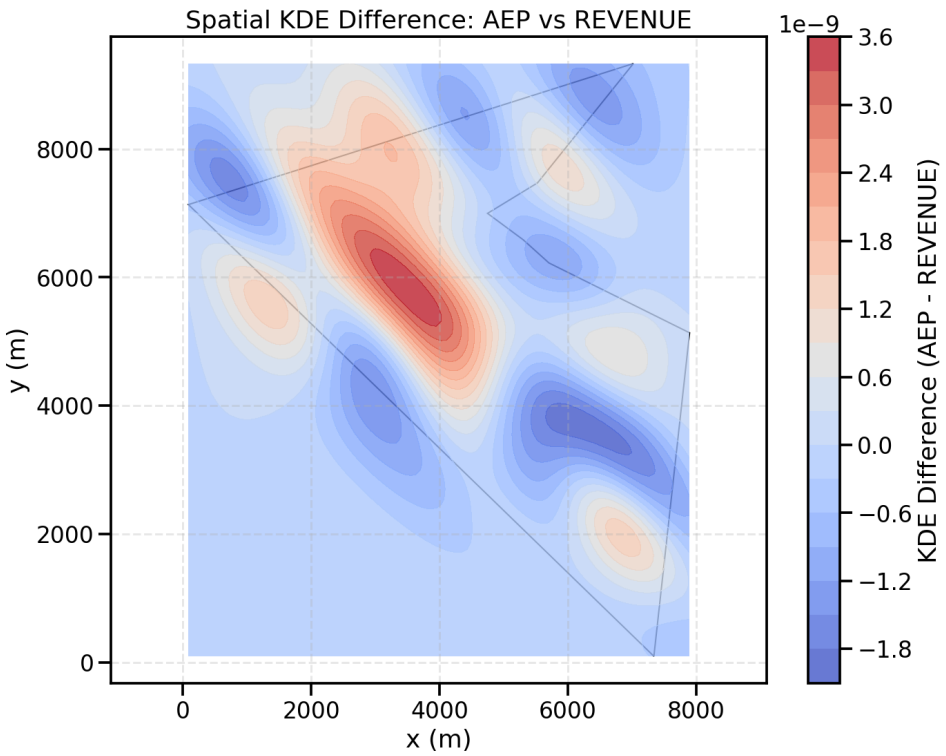

Figure A.2: Kernel Density Estimation of the Danish location, estimating spatial positioning preferences per objective function, red zones corresponding to areas favored by Revenue-based optimizations, and blue zones favored by AEP-based optimizations

B

Wind and market conditions for Case Study 2

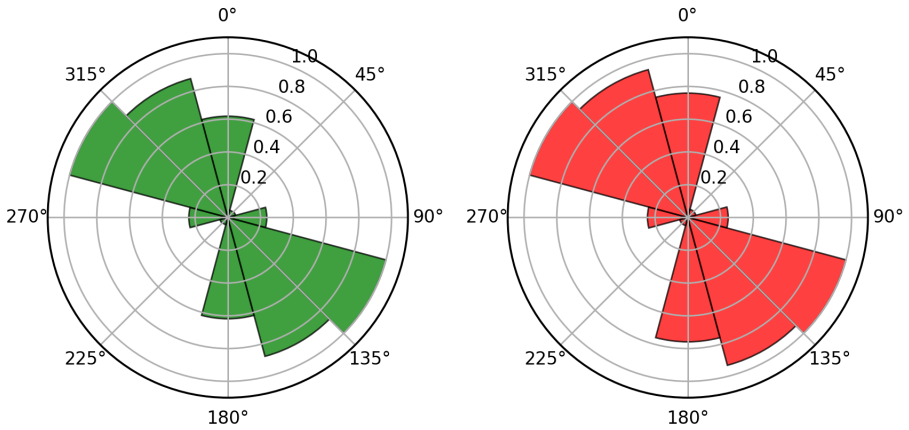

(a) France - Symmetric directional annual potential Energy (left) and Revenue (right)

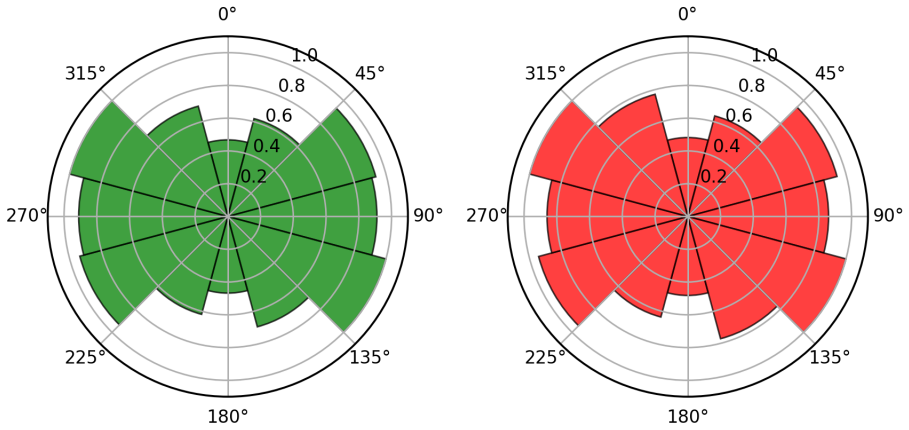

(b) Denmark - Symmetric directional annual potential Energy (left) and Revenue (right)

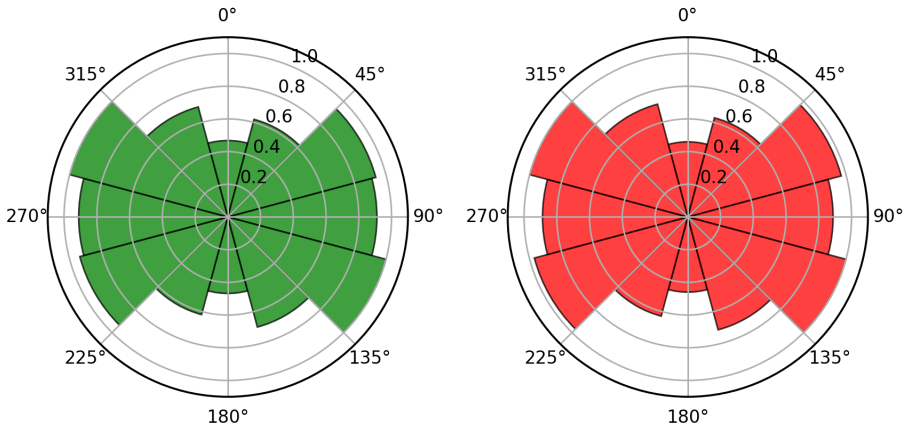

(c) Germany - Symmetric directional annual potential Energy (left) and Revenue (right)

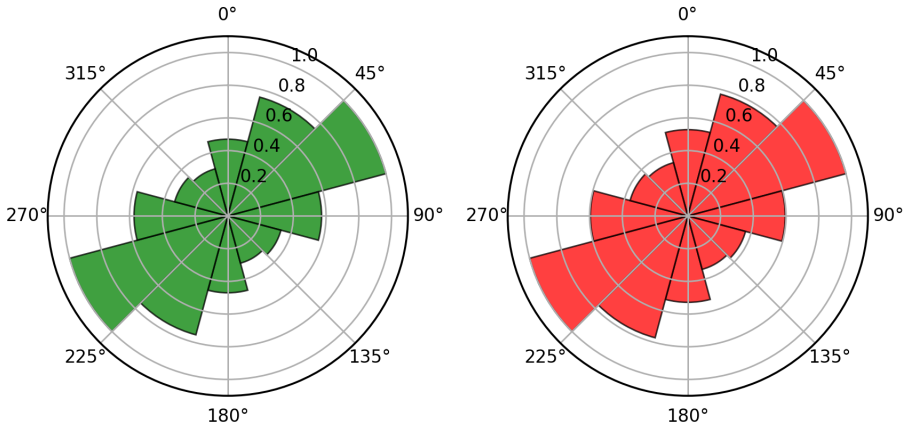

(d) Netherlands - Symmetric directional annual potential Energy (left) and Revenue (right)

Figure B.1: Comparison of potential Energy and Revenue roses for each location considered, for the time period 2015-2020, used to optimize layouts for case study 2

C

Noisy Scenario with -0.3 correlation between wind speed and electricity prices

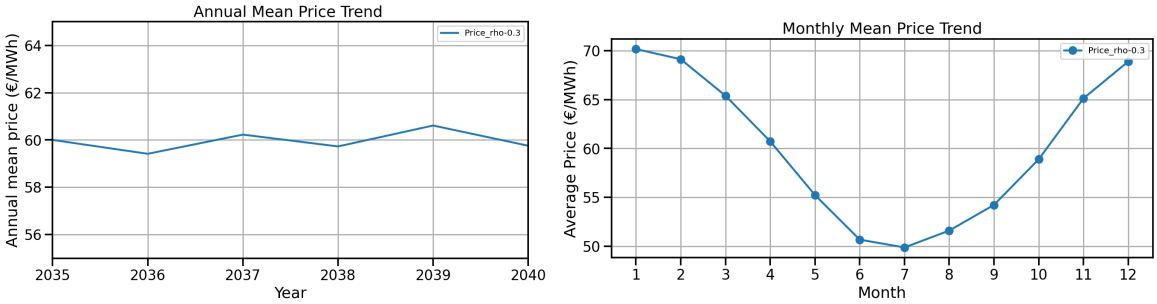

(a) On the left, yearly average of electricity prices. Small variations around 60 correspond to yearly wind resource variation. On the right, seasonal variation of the electricity prices.

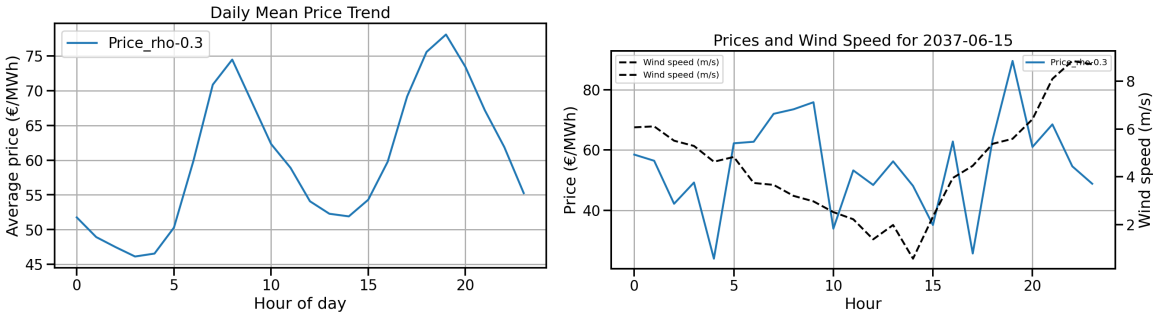

(b) On the left, average wind prices per hour of the day. On the right, wind speed and price for each scenario, during a specific day.

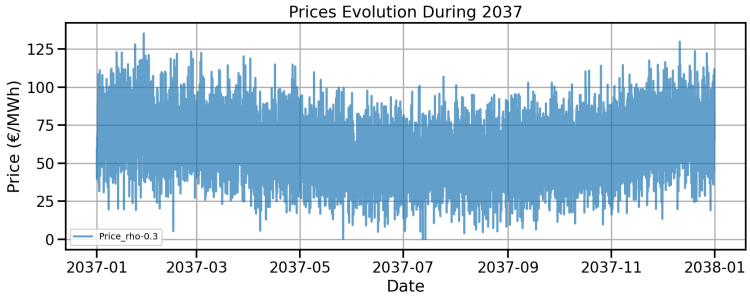

(c) Evolution of hourly electricity prices during the simulated 2037 year.

Figure C.1: Constructed scenario based on a correlation coefficient between wind speed and electricity prices of -0.3, but with a different random seed than the original used for layout optimization