

Master Thesis

Optimising Dispatch Strategies for Grid-Connected Green Hydrogen Systems with Wind PPAs, Battery Storage, and Day-Ahead Market Participation

I.E. Klein



Master Thesis

Optimising Dispatch Strategies for Grid-Connected Green Hydrogen Systems with Wind PPAs, Battery Storage, and Day-Ahead Market Participation

by

I.E. Klein

to obtain the degree of Master of Science
at the Delft University of Technology

Student number:	4931645	
Project duration:	September, 2024 – June, 2025	
Thesis committee:	Prof. dr. A. Papantoleon Dr. ir. G. F. Nane Dr. ir. J. T. van Essen Dr. R. Bhardwaj	TU Delft, Responsible supervisor TU Delft, Co-supervisor TU Delft, Committee member Power2X, Committee member
Additional Supervisor:	Dr. I. Bilibin	Power2X, Company supervisor
Study:	Master, Applied Mathematics	
Faculty:	Faculty of Electrical Engineering, Mathematics and Computer Science	
University:	Delft University of Technology	

Abstract

Green hydrogen is increasingly recognised as a key enabler of the energy transition, offering a carbon-free alternative for hard-to-abate sectors. However, its large-scale deployment remains economically challenging due to high capital expenditures (CAPEX) and production costs. The main cost driver is electricity procurement, given the reliance on affordable renewable energy and the challenges posed by its intermittent and uncertain nature. This underscores the importance of smart power sourcing and energy management strategies to reduce power-related expenses.

This thesis investigates the operational and economic performance of a hybrid renewable energy system designed for green hydrogen production. The system integrates an electrolyser with electricity procured via a wind Power Purchase Agreement (PPA), battery energy storage, and participation in the day-ahead electricity market. Hydrogen is assumed to be sold under a fixed-price offtake agreement, placing hydrogen production at the centre of operational decision-making. The primary objective is to determine the optimal dispatch strategy to maximise the long-term profitability of the hydrogen production plant, and to evaluate whether the addition of battery storage and market-based electricity sourcing enhances economic performance under different system configurations and external conditions.

To determine profit-maximising operational strategies, a mathematical optimisation framework is developed, comparing three solution methods with varying assumptions about foresight and uncertainty. The first method uses a deterministic formulation with perfect foresight, expressed as a mixed-integer linear programming (MILP) problem. This is extended into a rolling horizon approach with limited foresight, and a two-stage stochastic optimisation method that captures forecast uncertainty through correlated scenarios of wind generation and electricity prices. These methods are used to analyse how electricity sourcing strategies, system configuration and uncertainty treatment influence operational and economic indicators such as electrolyser load factor, system profitability, and the levelised cost of hydrogen (LCOH).

The results highlight the potential role of battery storage in reducing electricity costs and unlocking additional revenue by strategically charging and discharging in response to market price fluctuations. However, the associated cost savings are not sufficient to outweigh the high investment costs, making battery storage an unprofitable option under the assumed conditions. Another key finding is that allowing for supplementary electricity sourcing from the day-ahead market, in addition to the PPA, enables cost savings by taking advantage of low or even negative market prices. Nonetheless, to comply with green hydrogen certification requirements, the system must maintain a sufficiently large PPA to ensure a predominantly renewable electricity supply.

Regarding modeling approaches, the deterministic perfect foresight model achieves the highest system profitability, as it is able to optimise operations over the entire planning horizon with full knowledge of future conditions. In addition, it has the lowest computational burden, making it an attractive option for long-term planning and design studies. Although the rolling horizon and stochastic models provide a more realistic representation of uncertainty, the difference in profitability outcomes remains modest, with a maximum deviation of approximately 15 percent. For long-term profitability assessment and system design studies, the deterministic model proves to be sufficiently accurate while offering the advantage of lower computational effort.

By comparing different system configurations and optimisation approaches, this thesis provides insights into the trade-offs between model accuracy, economic performance, and operational robustness in integrated hydrogen systems, offering practical guidance for the design and operation of future green hydrogen facilities.

Acknowledgements

This thesis was written as the final part of the degree of Master of Science in Applied Mathematics at TU Delft. After starting the Bachelor's programme six years ago, it now marks the end of my academic journey. Looking back, I would not have chosen a different path. TU Delft has taken me to many inspiring places, both academically and personally, and has played a key role in shaping the interests and ambitions I hold today. One such experience was joining the Eco-Runner Team Delft during a gap year between my Bachelor and Master studies. With a dedicated team, we worked on developing a highly efficient hydrogen-powered city car, an opportunity that introduced me to the world of hydrogen and its role in the energy transition. That moment laid the foundation for my interest in the field of sustainable energy, and ever since, I have been eager to combine my mathematical education with meaningful real-world applications. This motivation guided my search for a thesis project and I feel incredibly fortunate to have found one that truly matched this ambition.

This project could not have been successful if it were not for several people. I would like to express my sincere gratitude to them here.

To start I would like to express my gratitude to Power2X for taking on this thesis project together with me. In particular, I would like to thank Rajat Bhardwaj for offering me the opportunity to be part of Power2X and to explore a topic that so closely aligns with my interests. Thank you for all the inspiring discussions, for trusting me with the freedom to shape my project and for helping me bridge the gap between theory and practice. Moreover, to everyone at Power2X: thank you for making me feel welcome from the very beginning and for your constant willingness to help in every way possible over the past ten months.

A special thank you goes to my company supervisor Ilya Bilibin. Your sharp insights, patience and consistent support have been invaluable. Our many sessions helped me stay on track while continuously learning from your depth of knowledge about power markets and modeling, and the clarity you brought to complex topics. Equally, I want to thank Fleur Theulen, with whom I had frequent check-ins and who became closely involved after the first couple of months. Thank you for the many hours you spent brainstorming with me, reviewing model outputs and challenging assumptions. I learned a great deal from working with you, and your feedback undoubtedly strengthened the quality of this work. Beyond that, you also became a friend and I truly value how you made working on my thesis more collaborative and fun than I could have hoped for.

I also want to sincerely thank my academic supervisors at TU Delft without whom this project would not have been possible. Firstly to Antonis Papapantoleon, thank you for your guidance and for keeping me focused whenever I strayed too far off track. I appreciate your encouragement early on, when the project was still taking shape, and your full trust and freedom in allowing me to pursue a topic that inspired me. Secondly to Tina Nane, although the final project ended up somewhat removed from your own research area, I'm incredibly thankful that you stayed involved throughout. I truly enjoyed our meetings and your calm and thoughtful presence always left me feeling more confident and at ease.

Finally, I want to thank my family for their unwavering support. Your encouragement throughout my studies has meant the world to me. A special thank you to my sister Anna, who provided valuable feedback on this report. I look forward to returning the favour when your own TU Delft thesis comes around.

In conclusion, I feel proud and grateful to have completed this thesis at the intersection of mathematics and the real-world challenges of the energy transition. The experience has only deepened my enthusiasm for working in this field and I look forward to continuing the journey ahead.

*I. E. Klein
Amsterdam, June 2025*

Contents

Abstract	i
Acknowledgements	ii
Nomenclature	vi
List of Figures	viii
List of Tables	ix
1 Introduction	1
1.1 Hydrogen as a clean and versatile energy carrier	1
1.2 Improving the green hydrogen business case	1
1.2.1 Challenges for large-scale green hydrogen deployment	2
1.2.2 Combining long-term contracts and spot market	2
1.2.3 Integrating battery storage	3
1.3 Modeling challenges in integrated renewable energy systems	3
1.4 Research objective	4
1.4.1 Cooperation with Power2X	5
1.4.2 Research questions	6
1.5 Report outline	6
2 Background Information: Energy Technologies and Markets	7
2.1 Hydrogen	7
2.1.1 Grey, blue and green hydrogen	7
2.1.2 Hydrogen production from water electrolysis	8
2.1.3 Hydrogen storage and transport infrastructure	10
2.1.4 Use cases for hydrogen	12
2.1.5 Dutch hydrogen infrastructure: national backbone	12
2.1.6 Hydrogen spot market and long-term contracting	13
2.1.7 Subsidies and regulations	14
2.2 Electricity Markets	14
2.2.1 Forward market	15
2.2.2 Day-ahead market	15
2.2.3 Intraday market	16
2.2.4 Balancing market	16
2.2.5 Pricing mechanism and merit order effect	16
2.2.6 Long-term contracting through PPAs	17
2.3 Energy Storage Systems	18
2.3.1 Battery Energy Storage Systems (BESS)	19
3 Literature Review	20
3.1 Mathematical techniques for optimisation under uncertainty	20
3.1.1 Optimisation of multi-period decision problems	20
3.1.2 Mixed-Integer Linear Programming	21
3.1.3 Dynamic Programming	22
3.1.4 Rolling Horizon Optimisation	24
3.1.5 Stochastic Programming	24
3.1.6 Scenario generation methods for stochastic optimisation	26
3.1.7 Multivariate Monte Carlo and correlated scenario generation	27

3.2	Classifying literature on combined energy systems management	28
3.3	Key related works on optimising integrated energy systems.	32
3.3.1	Multi-energy configurations using linear programming	32
3.3.2	Applications of rolling horizon and two-stage stochastic programming	33
3.3.3	Application of Markov decision process formulation.	33
3.4	Classifying this research and identifying the literature gap	34
3.4.1	Literature gap	36
4	Research Approach	37
4.1	System description	37
4.2	System configuration test cases	38
4.3	Model workflow.	39
4.4	Key Performance Indicators.	40
4.4.1	Operational KPIs	40
4.4.2	Financial KPIs	41
4.5	Modeling assumptions and scope clarifications.	42
4.5.1	Technical and operational modeling.	42
4.5.2	Market participation	43
4.5.3	Hydrogen offtake and regulatory green labeling	43
4.5.4	Economic and financial evaluation	44
4.5.5	Forecasting and uncertainty modeling.	44
4.6	Data for external information	44
4.6.1	Day-ahead electricity price data	45
4.6.2	Wind load factor data	48
4.7	Model parameters.	49
4.7.1	Temporal resolution and modeling period	49
4.7.2	Electrolyser parameters	50
4.7.3	Battery parameters.	51
4.7.4	Commercial parameters	51
4.8	Model implementation	53
5	Model Formulation	54
5.1	Objective function	54
5.1.1	Time horizon and resolution.	55
5.1.2	Time-dependent external information: day-ahead prices and wind load factors	55
5.1.3	Fixed Costs.	55
5.1.4	Variable costs	55
5.1.5	Revenue	56
5.1.6	Total objective function	57
5.2	Variables	57
5.3	Constraints	58
5.4	Model characteristics and problem classification	61
5.4.1	Linear reformulation.	61
5.4.2	Look-ahead period and uncertainty incorporation.	62
5.5	Final formulation MILP	64
6	Advanced Solution Methods	65
6.1	Rolling horizon optimisation formulation.	65
6.1.1	Temporal structure and solution logic	65
6.1.2	Modified constraints.	66
6.2	Two-stage stochastic optimisation with rolling horizon.	67
6.2.1	Model structure and decision stages	68
6.2.2	Objective function	68
6.2.3	Constraints	68

6.3	Construction of day-ahead price and wind load factor scenarios	70
6.3.1	Forecast error modeling and scenario generation method	70
6.3.2	Estimating mean and standard deviation for forecast errors	72
6.3.3	Illustration and validation of scenario generation method	73
6.3.4	Scenario probabilities	74
7	Results	75
7.1	Base case results for deterministic model	75
7.1.1	System dynamics under varying market and wind conditions	75
7.1.2	Operational and financial KPI results	79
7.1.3	Sensitivities	83
7.2	Comparing different optimisation approaches	84
7.2.1	Base case PPA size: minimal differences between DO, RHO and SORH solutions	85
7.2.2	Sensitivity analysis PPA size	85
7.2.3	Smaller PPA sizes: increased variation between DO, RHO and SORH solutions	86
7.3	Impact of differentiated hydrogen sales prices (green vs grey) on system operation and profitability	90
7.4	Analysis battery profitability	93
7.4.1	Optimal battery size	93
7.4.2	Assessment economic impact of battery inclusion	93
7.4.3	Performance across different years.	93
7.4.4	Test case: no wind PPA	95
8	Conclusion and Discussion	96
8.1	Optimisation method hybrid energy systems	96
8.2	Profitability implications of market access and battery storage	98
8.2.1	Trade-offs between wind PPA and grid electricity sourcing.	98
8.2.2	Profitability of the battery	99
9	Recommendations for Future Research	101
	Appendices	104
A	Appendix	105
A.1	Correlation Structure Scenario Generation Method	105
B	Appendix	107
B.1	Battery and electrolyser behaviour for the differentiated hydrogen price model	107
C	Appendix	109
C.1	Results for No Battery Model for 100 MW PPA and 50 MW Electrolyser Setup	109
C.2	Results for Electrolyser only Model for different Historical Day-Ahead Price Data.	110
	Bibliography	112

Nomenclature

List of Abbreviations

Abbreviation	Definition
AE	Alkaline Electrolyser
BESS	Battery Energy Storage System
CAPEX	Capital Expenditure
CVaR	Conditional Value at Risk
DA	Day-Ahead
DO	Deterministic Optimisation
DP	Dynamic Programming
EES	Electrical Energy Storage System
FC	Fuel Cell
GHG	Green House Gas
HPA	Hydrogen Purchase Agreement
ID	Intraday
IRR	Internal Rate of Return
LCOH	Levelised Cost of Hydrogen
LCOE	Levelised Cost of Energy
LHV	Lower Heating Value
LP	Linear Programming
MCP	Market Clearing Price
MCV	Market Clearing Volume
MCS	Monte Carlo Simulation
MDP	Markov Decision Process
MILP	Mixed Integer Linear Programming
MNILP	Mixed Non-Integer Linear Programming
NPV	Net Present Value
OPEX	Operational Expenditure
OTC	Over-the-Counter
PEM	Proton Exchange Membrane (Electrolyser)
PPA	Power Purchase Agreement
PSD	Positive Semi-Definite
PV	Photovoltaic (Solar)
RES	Renewable Energy Sources
RFNBO	Renewable Fuels of Non-Biological Origin
RH	Rolling Horizon
RHO	Rolling Horizon Optimisation
RM	Reserves Market
RO	Robust Optimisation
SMR	Steam Methane Reforming
SO	Stochastic Optimisation
SOC	State of Charge
SOE	Solid Oxide Electrolyser
SORH	Stochastic Optimisation with Rolling Horizon
SP	Stochastic Programming
TSO	Transmission System Operator
WT	Wind Turbine

List of Symbols

Symbol	Description	Units	Type
Decision Variables			
$p_t^{M \rightarrow E}$	Electricity from market to electrolyser at time t	MWh	Continuous
$p_t^{M \rightarrow B}$	Electricity from market to battery at time t	MWh	Continuous
$p_t^{B \rightarrow E}$	Electricity from battery to electrolyser at time t	MWh	Continuous
$p_t^{B \rightarrow M}$	Electricity from battery to market at time t	MWh	Continuous
$p_t^{W \rightarrow E}$	Wind PPA electricity used by electrolyser at time t	MWh	Continuous
$p_t^{W \rightarrow B}$	Wind PPA electricity stored in battery at time t	MWh	Continuous
$p_t^{W \rightarrow M}$	Wind PPA electricity sold to market at time t	MWh	Continuous
$p_t^{W, \text{overcap}}$	Wind PPA power not utilised and subject to penalty	MWh	Continuous
z_t^{state}	Binary on/off state of the electrolyser	{0, 1}	Binary
Total Power Flows			
$p_t^{M, \text{buy}}$	Total electricity bought from market at time t	MWh	Continuous
$p_t^{M, \text{sell}}$	Total electricity sold to market at time t	MWh	Continuous
$p_t^{B, \text{charge}}$	Power used to charge battery at time t	MWh	Continuous
$p_t^{B, \text{discharge}}$	Power discharged from battery at time t	MWh	Continuous
$p_t^{E, \text{in}}$	Power used by electrolyser at time t	MWh	Continuous
State Variables			
$p_t^{W, \text{PPA}}$	Wind power available under PPA at time t	MWh	Continuous
SOC_t	Battery state of charge at time t	%	Continuous
h_t	Hydrogen produced at time t	MWh (LHV)	Continuous
U_t	Utilisation rate of the electrolyser at time t	%	Continuous
$z_t^{\text{su count}}$	Binary indicating electrolyser shutdown at time t	{0, 1}	Binary
Auxillary variables / linearisation variables			
z_t^{grid}	Binary indicator for grid direction (buy/sell)	{0, 1}	Binary
z_t^{battery}	Binary indicator for battery mode (charge/discharge)	{0, 1}	Binary
Exogenous Information (time series)			
α_t	Wind load factor at time t	-	Continuous
λ_t	Day-ahead electricity market price at time t	€/MWh	Continuous
Model Parameters			
D^{target}	Annual hydrogen production target	MWh	Fixed
π^{H_2}	Hydrogen selling price (HPA)	€/MWh	Fixed
π^{PPA}	PPA electricity price	€/MWh	Fixed
$\tilde{\lambda}$	Small artificial penalty to guide dispatch behaviour	-	Fixed
K^{PPA}	Contracted wind PPA capacity	MW	Fixed
K^{Elec}	Installed capacity of electrolyser	MW	Fixed
K^{Batt}	Battery energy capacity	MWh	Fixed
K^{grid}	Grid connection capacity	MW	Fixed
C^{bal}	Penalty cost for unused wind PPA power	€/MWh	Fixed
C^{su}	Cost of shutdown for electrolyser	€	Fixed
η^e	Efficiency of the electrolyser	%	Fixed
p^{min}	Minimum capacity of the electrolyser	MWh	Fixed
N^{off}	Maximum number of shutdowns allowed per year	-	Fixed
η^C	Charging efficiency of battery	%	Fixed
η^D	Discharging efficiency of battery	%	Fixed
$p^{C, \text{max}}$	Maximum charging power	MWh	Fixed
$p^{D, \text{max}}$	Maximum discharging power	MWh	Fixed
$\text{SOC}^{\text{min}}, \text{SOC}^{\text{max}}$	Min/max state of charge	%	Fixed

List of Figures

1.1	Simplified schematic of the base case hydrogen production system setup.	5
2.1	Schematic representation of the water electrolysis process for hydrogen production [94].	8
2.2	Illustration of non-constant electrolyser DC and system efficiencies at varying load levels. Adapted from Siemens Energy [48].	10
2.3	Overview of planned Dutch hydrogen infrastructure development [38].	13
2.4	Overview of the sequential structure of electricity markets based on trading timeframes.	15
2.5	Market clearing mechanism in the day-ahead electricity market.	16
2.6	Illustration of the merit order curve and the market clearing mechanism.	17
2.7	Visual representation of the merit order effect.	17
2.8	Schematic representation of a lithium-ion battery module.	19
3.1	Schematic of the Rolling Horizon Optimisation (RHO) framework.	24
3.2	Conceptual framework for classifying literature for optimisation of energy management systems. 29	29
4.1	Schematic overview of the hybrid renewable energy and hydrogen production system with a wind Power Purchase Agreement (PPA), market interaction, a battery, an electrolyser and hydrogen offtake.	38
4.2	Schematic overview of the research approach. The figure illustrates the full workflow of the study, including the model inputs, pre-processing steps, optimisation model, and post-processing activities leading to final outputs.	39
4.3	Hourly electricity prices and offshore wind capacity factors over a selected period in 2023. The plot illustrates the inverse relationship between wind availability and market prices.	45
4.4	Monthly distribution of day-ahead electricity prices in 2023 to illustrate seasonal trends of price levels and volatility.	46
4.5	Distribution of hourly changes in day-ahead electricity prices (€/MWh) in 2023 which illustrates the short-term volatility of day-ahead prices.	46
4.6	Distribution of hourly day-ahead electricity prices in the Netherlands from 2019 to 2024 illustrating differences between historical price years.	47
4.7	Monthly distribution of hourly wind load factors for the Netherlands in 2023 illustrating strong seasonal trends in wind energy generation.	48
4.8	Distribution of hourly wind load factors in the Netherlands for the years 2019–2024.	49
6.1	Illustration of scenario generation method for electricity prices and wind load factors, using 2023 as base case year.	73
7.1	Base case results: dispatch behavior during a generally high-wind, low-price period.	76
7.2	Base case results: dispatch strategy under low-wind, high-price conditions	77
7.3	Hourly revenue, costs and profit during the low-wind, high-price scenario for base case results. 77	77
7.4	Base case dispatch results: system operation during wind overcapacity and ultra-low price conditions.	78
7.5	Base case dispatch results: system behavior under high wind and price variability.	79
7.6	Sensitivity analysis of the LCOH and IRR with respect to key economic drivers: hydrogen sales price and PPA price.	84
7.7	Sensitivity analysis results of hydrogen system under varying PPA sizes.	86
7.8	Comparison of dispatch patterns under different hydrogen sales price assumptions.	91
7.9	Sensitivity of LCOH and IRR to PPA size under a differentiated hydrogen pricing model.	92
A.1	Heat maps for the smoothed cross-correlation matrices R_{wp} for different exponential decay factors $\delta \in \{0.6, 0.9\}$	106

List of Tables

2.1	Simplified comparison of water electrolysis technologies: Alkaline Electrolysis (AE), Proton Exchange Membrane Electrolysis (PEM), and Solid Oxide Electrolysis (SOE).	9
3.1	Overview of several optimisation studies in integrated energy and hydrogen systems analysed for this thesis.	31
4.1	Base case system configuration.	38
4.2	Overview of system test cases.	39
4.3	Overview of operational Key Performance Indicators (KPIs).	40
4.4	Summary statistics of hourly day-ahead electricity prices in the Netherlands for the years 2019–2024 illustrating differences between historical price years.	48
4.5	Summary statistics of hourly wind load factors in the Netherlands for the years 2019–2024.	49
4.6	Temporal parameters for rolling horizon models.	50
4.7	Electrolyser’s technical and economical input parameters.	50
4.8	Battery storage technical and economical input parameters.	51
4.9	Commercial input parameters regarding the grid connection, power purchase agreement (PPA) and hydrogen sales agreement (HPA).	52
4.10	Historical wind capture prices and calculated PPA prices (2019–2024).	52
5.1	Overview of modeling solution methods compared in this thesis.	63
6.1	Estimated coefficients for the temporal autocorrelation and cross-correlation of wind generation and electricity prices, based on 2023 historical data.	71
6.2	Statistical comparison of historical and simulated electricity price and wind load factor data for the base year 2023.	73
7.1	Summary of base case operational results: total system power and hydrogen flows for the optimised dispatch.	80
7.2	Summary of base case financial results: breakdown of all revenue streams and costs.	80
7.3	Summary of electrolyser utilisation metrics for base case results.	81
7.4	Shut-off cycle characteristics of the electrolyser for base case results.	81
7.5	Battery operation summary, including charging and discharging behavior for base case results.	81
7.6	Wind PPA electricity allocation and electrolyser input breakdown for base case results.	82
7.7	Battery charging and discharging distribution for base case results.	82
7.8	Electricity market buy and sell distributions for base case results.	82
7.9	KPI results for base case model including the projects NPV, IRR and LCOH.	82
7.10	LOCH cost component breakdown for base case results.	83
7.11	Comparison of results from the different modeling approaches: DO, RHO and SORH. Results generated for base case system (oversized PPA).	85
7.12	Comparison of results from the different modeling approaches: DO, RHO and SORH. Results generated for smaller wind PPA size and lowered hydrogen sales price.	87
7.13	Relative difference in KPI’s between models: DO, RHO, SORH. Results for smaller wind PPA size and lowered hydrogen sales price concept.	87
7.14	Comparison of results from the different modeling approaches: DO, RHO and SORH. Results generated for smaller wind PPA size and lowest hydrogen sales price.	87
7.15	Relative difference in KPI’s between models: DO, RHO, SORH. Results for smaller wind PPA size and lowest hydrogen sales price concept.	87
7.16	Comparison of results from the different modeling approaches: DO, RHO and SORH. Results generated for the no-wind PPA system concept.	88

7.17	Relative difference in KPI's between models: DO, RHO, SORH. Results for no-wind system concept.	88
7.18	KPI results for the differentiated hydrogen sales price model, including the projects NPV, IRR and LCOH.	91
7.19	Annual average market and wind conditions for different years.	94
7.20	Summary of annual performance metrics for different historical years.	94
7.21	Breakdown of LCOH components by year, expressed in €/kg.	95
A.1	Estimated lag-1 autocorrelation from historical 2023 data	105
A.2	Empirical Pearson correlation between wind and electricity prices	105
B.1	Battery operation results under the differentiated hydrogen sales price model.	108
B.2	Distribution of power sources to the electrolyser and utilisation summary results for the differentiated hydrogen sales price model.	108
C.1	Financial KPI results for model set-up without battery.	109
C.2	Summary of annual power flows results for model set-up without battery.	109
C.3	Electrolyser utilisation behaviour results for model set-up without battery.	109
C.4	Distribution of power to electrolyser for model set-up without battery	109
C.5	Summary of costs and revenues for model set-up without battery.	110
C.6	Wind power distribution for model set-up without battery.	110
C.7	Performance metrics for electrolyser and battery (no PPA) model	111
C.8	Performance metrics for electrolyser only (no PPA, no battery) model	111

1

Introduction

The global energy transition towards a carbon-neutral society calls for a fundamental transformation in how energy is produced, consumed and managed. Deep decarbonisation across sectors, especially in hard-to-abate industries such as steel, chemicals and heavy-duty transport, demands innovative solutions beyond direct electrification [42]. In this context, hydrogen is increasingly recognised as a critical energy carrier. It has become a key priority for achieving climate goals, including those outlined in the European Green Deal and national net-zero targets [25, 40, 61]. Hydrogen complements electrification by offering a flexible, low-carbon solution where direct electrification is not feasible. By 2050, the European Union envisions the use of 400 million tonnes of hydrogen, aiming to cover 8% of total energy demand [39, 40]. This ambition underscores hydrogen's strategic importance in the future energy system.

1.1. Hydrogen as a clean and versatile energy carrier

Hydrogen is the lightest and most abundant element in the universe, found naturally in compounds such as water and fossil fuels [40]. Hydrogen can be produced through several methods, each with different environmental impacts. The most common route today is steam methane reforming of natural gas, resulting in so-called grey hydrogen, which emits substantial CO₂ [58]. Incorporating carbon capture and storage (CCS) into this process yields blue hydrogen, a lower-carbon alternative. Green hydrogen, by contrast, is produced by splitting water via electrolysis using renewable electricity, making it a carbon-free energy carrier [40, 41]. This thesis focuses on green hydrogen, specifically in combination with renewable electricity from wind, due to its potential to support long-term decarbonisation goals.

A key strength of hydrogen lies in its versatility. It can serve as a fuel, feedstock, energy carrier or medium for energy storage, with potential applications across sectors such as industry, transport and power [25]. Historically, hydrogen has been used primarily as a feedstock in industrial processes such as ammonia production and oil refining. However, these uses have largely relied on fossil-based hydrogen and are associated with high CO₂ emissions [58]. In the coming years, a shift toward blue and green hydrogen is expected, along with its expanded applications in mobility, heating and chemical manufacturing [40]. Hydrogen also enables the production of low-carbon molecules such as ammonia, bio-methanol and synthetic fuels, which are critical to decarbonising the chemical and transport sectors [87]. In addition, hydrogen can facilitate the integration of renewable energy by enabling long-duration energy storage. During periods of excess electricity generation, surplus power can be converted into hydrogen. This stored hydrogen can be used directly or reconverted into electricity, enhancing grid flexibility, improving energy security, and supporting the integration of variable renewable energy sources [61].

1.2. Improving the green hydrogen business case

Despite its promise, the green hydrogen market faces several obstacles that hinder its large-scale deployment. Overall project costs are turning out much higher than initially projected, leading to delays and cancellations across the sector [14, 78]. As a result, only a limited number of large commercial facilities in Europe have reached Final Investment Decision (FID), and Shell's Holland Hydrogen 1 project in the Netherlands stands out as one of the few examples currently moving forward [82].

1.2.1. Challenges for large-scale green hydrogen deployment

The hydrogen market remains immature, with uncertainty surrounding future supply and demand, limited trading mechanisms and underdeveloped infrastructure for storage and transport [7, 42]. Key preconditions such as dedicated pipelines, clear regulatory frameworks and sufficient technical expertise are still lacking, undermining investor confidence and slowing large-scale deployment [40, 61]. The evolving regulatory framework at the European level adds further complexity, especially with ongoing negotiations around certification requirements for Renewable Fuels of Non-Biological Origin (RFNBOs), adding regulatory uncertainty for developers and off-takers [39].

The main challenge is that the cost of green hydrogen remains significantly higher than that of grey hydrogen and other fossil-based energy sources [25, 40]. This is primarily due to high capital investment (CAPEX) requirements for electrolyzers and the cost of securing renewable electricity [7, 41, 42, 58, 94]. Although renewable electricity costs have declined significantly over the past decade, ensuring access to low-cost renewable power with high utilisation rates for dedicated hydrogen production remains a key challenge. This constraint continues to hinder the large-scale deployment of green hydrogen [41]. However, recent studies indicate that falling renewable electricity prices, combined with expected reductions in electrolyser CAPEX, could lead to a substantial drop in green hydrogen production costs by 2030 [13].

The Levelised Cost of Hydrogen (LCOH) is a key metric for evaluating the economic viability of hydrogen projects and comparing different production pathways. It reflects the average cost of producing one kilogram of hydrogen over the full lifetime of an asset, incorporating CAPEX, operational costs, and electricity-related expenses. A detailed cost breakdown from a Dutch case study by TNO (2024) estimates the LCOH in the range of €10 to €14 per kilogram, with electricity-related costs, including power procurement, grid fees and energy taxes, accounting for more than half of the total [19]. These findings emphasise the importance of smart electricity sourcing strategies, as reducing power costs can significantly improve the economic viability of green hydrogen production. Several approaches can help reduce these costs, including combining long-term contracts with spot market sourcing, discussed in Section 1.2.2, and integrating battery storage to increase operational flexibility, discussed in Section 1.2.3.

In addition to high costs, green hydrogen projects also face significant operational challenges. Specifically, integrating hydrogen production with variable renewable energy sources introduces significant system design and dispatch challenges. Electrolysers must be able to respond to fluctuating generation profiles and market signals, which adds technical and economic complexity. The inherent intermittency of renewables, combined with increasing volatility and uncertainty in electricity market conditions, complicates a steady power procurement strategy and predictable operations cashflow. These dynamics can increase the risk of underutilisation or unprofitable operation, further weakening the economic case for green hydrogen deployment [41].

These operational challenges underscore the need for models that can capture system-wide dynamics and determine the optimal operational strategy. Such models must optimise energy dispatch across all assets in response to fluctuating renewable generation and electricity prices. By adopting a forward-looking approach, they enable cost-effective scheduling of electrolyser operation, battery charging and discharging, and power market transactions. The modeling challenges associated with these optimisation methods are further discussed in Section 1.3.

1.2.2. Combining long-term contracts and spot market

To address the economic and operational challenges of large-scale green hydrogen deployment, several strategies have been proposed to improve the viability of such projects. A widely discussed approach includes the use of long-term fixed contracts to mitigate risks related to market volatility and future uncertainty, thereby strengthening investor confidence. For electricity procurement, Power Purchase Agreements (PPAs) can supply renewable power, such as wind or solar, to electrolyzers at a stable and predictable cost [29, 58]. These contracts reduce operational cost uncertainty, as electricity accounts for a large share of hydrogen production costs. Similarly, Hydrogen Purchase Agreements (HPAs) can secure future revenue by contractually fixing both the purchase price and offtake volume of the produced hydrogen [7, 79]. These contractual arrangements benefit both energy suppliers and offtakers by reducing exposure to price fluctuations and supply insecurity [20, 29]. In doing so, they support more predictable business models, which are essential for enabling investment and facilitating project development in an uncertain market landscape.

Beyond contractual mechanisms, a flexible and adaptive electricity sourcing strategy offers another critical lever to improve the economics of green hydrogen. One effective approach is to build a diversified power procurement portfolio that combines long-term fixed contracts with participation in the spot market [58].

While fixed contracts such as PPAs offer stability and supply security, they often come with a price premium in exchange for guaranteed volumes and long-term commitments [57]. In contrast, sourcing from the spot market allows producers to take advantage of periods with low electricity prices, though it introduces higher exposure to price volatility and uncertainty. A strategic combination of the two can allow producers to balance cost minimisation with risk management, ensuring supply while capturing value from favourable market conditions.

Another important benefit of sourcing electricity through wind or solar PPAs is that it supports compliance with emerging hydrogen certification schemes. In particular, the European Union's RFNBO framework requires hydrogen producers to demonstrate that their electricity input comes from renewable sources. Excessive reliance on grid power, which may include fossil-based generation, could disqualify the hydrogen from being sold at a green premium. As a result, PPAs with clearly traceable renewable sources are becoming increasingly important in the commercial viability of green hydrogen projects [85].

When combining a variable renewable supply with a fixed-capacity electrolyser, determining the optimal sizing ratio between the two becomes a key design challenge. If the PPA capacity is too small, hydrogen production may be constrained during low-generation periods and require greater reliance on grid electricity to meet demand targets. Conversely, if the PPA is oversized, surplus electricity may need to be sold back to the market at unfavourable prices, lowering overall project returns. These dynamics are further influenced by the inclusion of a battery energy storage system, which can buffer variability and enhance operational flexibility. Finding the optimal configuration among renewables, electrolyser, and battery capacity requires detailed modeling and scenario testing.

1.2.3. Integrating battery storage

To further address renewable intermittency and enhance system performance, energy storage can play a key role. Battery energy storage systems (BESS) can significantly enhance operational flexibility [2]. Batteries allow producers to store excess renewable electricity when supply exceeds demand and use it during periods of shortfall. Beyond storing surplus generation, BESS can also be used to purchase and store electricity from the grid during periods of low or even negative prices, typically associated with high renewable penetration, and supply it to the electrolyser when additional power is needed. This can lower the average electricity procurement cost of hydrogen production. Moreover, by enabling a more stable and controllable power input, batteries can smooth the load profile of the electrolyser, improving its utilisation and potentially extending its lifetime.

In addition to supplying the electrolyser, BESS opens up new revenue opportunities through participation in power markets [2, 58]. By anticipating future price developments, stored electricity can be strategically discharged and sold back to the grid during high-price periods. In this way, batteries can simultaneously support cost savings and value generation, strengthening the economic case for integrated hydrogen systems. However, despite declining costs, the capital investment required for large-scale battery systems remains substantial [47]. Furthermore, the ability to realise stable revenues from combined market trading and electrolyser optimisation remains uncertain and highly dependent on project-specific and market-specific factors [47]. This calls for research to quantify the potential benefits and assess under which conditions the additional investment in BESS can be economically justified.

1.3. Modeling challenges in integrated renewable energy systems

The shift towards low-carbon energy systems introduces a number of modeling challenges, particularly for integrated systems combining renewable electricity, hydrogen production, and storage technologies [89]. As variable renewable energy sources such as wind and solar become more prominent their intermittency and limited predictability increase the need for operational flexibility. This requires models that can capture time-dependent dynamics and identify optimal operational schedules in anticipation of changing conditions. In particular, forward-looking optimisation becomes essential to coordinate when and how assets are operated under fluctuating generation and price signals.

A key modeling challenge lies in the increasing interdependence between system components. Electrolysers, batteries, and grid connections are no longer independent units, but parts of a tightly coupled energy system. As a result, modeling these components in isolation may overlook critical interactions. Decisions such as when to run the electrolyser or store electricity in the battery depend on system-wide conditions, including relative asset sizes and contractual arrangements. This calls for integrated optimisation approaches that capture these interdependencies within a unified system framework. However, such integrated models

can become computationally expensive, particularly when exploring multiple configurations or time horizons. A careful balance must be found between model realism and tractability. This balance also concerns the level of detail used to represent individual components. While real-time operation may require detailed models of electrolyser process dynamics, including variable efficiency and startup behaviour, long-term strategic assessments benefit more from a simplified but representative system-level model. In these cases the focus is on capturing the key interactions between components to evaluate long-term profitability under different configurations and market conditions, rather than simulating real-time operational constraints in detail.

Finally, uncertainty plays a central role in modeling energy systems dominated by renewables. Key variables such as electricity prices and renewable generation from wind and solar fluctuate considerably over time. Many traditional energy models assume perfect foresight, meaning all future values of these variables are known in advance. While this simplifies optimisation and reduces computational burden, it becomes increasingly unrealistic in volatile and uncertain markets. Methods capable of representing forecast errors and external variability can support more robust operational strategies and better-informed investment decisions [80]. One such approach is stochastic optimisation, which considers multiple possible futures and determines operational strategies that perform well across a range of scenarios. However, these methods increase model complexity and require accurate quantification of uncertainty, which remains a challenging task.

Some researchers argue that, for long-term investment and planning studies, it may be more practical to rely on deterministic models under perfect foresight, combined with extensive scenario analysis to explore uncertainty in outcomes [2, 63]. This approach maintains computational efficiency, while still providing insights into how system performance might vary under different conditions. These models are powerful in guiding investment decisions but tend to overestimate profitability by assuming full knowledge of future conditions. In contrast, short-term operational models, often used for real-time scheduling, incorporate uncertainty more explicitly but are typically not used for long-term strategic planning. Determining the appropriate balance between model realism, computational cost, and robustness of results remains an open research question. This thesis contributes to this discussion by comparing different optimisation approaches and assessing their suitability for long-term profitability assessment of integrated green hydrogen systems under uncertainty.

1.4. Research objective

This thesis investigates the operational and economic performance of a hybrid renewable energy system designed for green hydrogen production. The system consists of an electrolyser powered by offshore wind electricity procured through a PPA, supported by battery energy storage and participation in the day-ahead electricity market. Hydrogen is assumed to be sold under a fixed-price offtake contract, placing hydrogen production at the centre of operational decision-making. A simplified schematic overview of the base case system setup is presented in Figure 1.1. The primary objective of this research is to develop and evaluate optimisation models that determine the optimal electricity procurement, storage dispatch, and hydrogen production strategy to maximise long-term system profitability. A key novelty of this study lies in its integrated treatment of all system components and the co-optimisation of dispatch decisions under fluctuating electricity prices and variable wind generation.

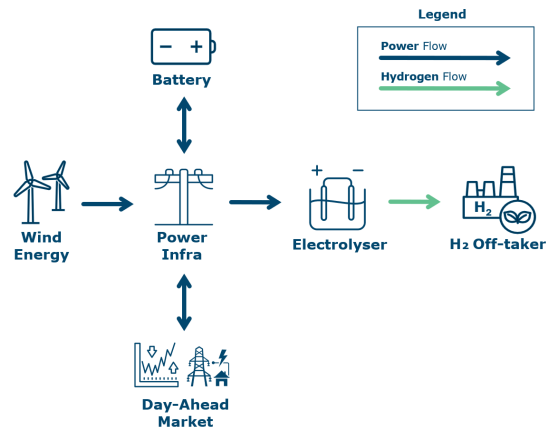


Figure 1.1: Simplified schematic of the base case hydrogen production system setup. Electricity is sourced from wind power via a PPA and the day-ahead market, with a battery enabling flexible dispatch to the electrolyser and the power market. The produced hydrogen is sold to an off-taker under a fixed-price contract.

The research focuses in particular on the role of short-term energy storage and market interaction in improving the business case for green hydrogen. It addresses fundamental questions regarding the contribution of battery storage and electricity market access to overall system profitability. Financial performance is evaluated using key performance indicators such as the LCOH, while operational outcomes are assessed through electrolyser utilisation, battery behaviour, and market trading patterns.

To address these questions, a mathematical optimisation model is formulated as a mixed-integer linear program (MILP). The model determines the optimal hourly dispatch of electricity between system components over a full year of operation, accounting for the dynamic interactions between wind availability, electricity prices, battery charging and hydrogen production. A central contribution of this thesis is the comparative analysis of three solution methods that differ in their treatment of foresight and uncertainty. The first is a deterministic optimisation (DO) approach with perfect foresight over the full year of wind generation and electricity prices. The second is a rolling horizon (RHO) method that limits foresight to a one-week look-ahead window. The third extends this rolling framework by incorporating forecast uncertainty through scenario-based stochastic optimisation (SORH). In this method, multiple future scenarios of wind and electricity prices are generated by introducing correlated forecast errors, enabling more robust decision-making under uncertainty. This comparison allows for the evaluation of how different levels of foresight and uncertainty handling affect long-term system profitability and model performance.

A novel aspect of this work is the application of rolling horizon and stochastic optimisation methods to evaluate long-term system profitability, an approach more commonly used in short-term dispatch or real-time operations. Long-term planning models typically rely on deterministic formulations due to lower computational burden, despite the risk of overestimating profitability. This thesis compares the proposed models in terms of financial performance and computational efficiency, offering practical insights into which modeling approach is best suited for long-term green hydrogen investment analysis.

1.4.1. Cooperation with Power2X

This thesis is conducted in collaboration with Power2X, a company based in Amsterdam that specialises in developing and operating large-scale sustainable energy projects. Founded in 2020, Power2X focuses on accelerating the global energy transition by facilitating the production of green and low-carbon molecules, including hydrogen, ammonia, methanol and sustainable aviation fuels. As a scale-up active in both project development and consultancy, Power2X also helps industrial partners design, develop and implement clean energy solutions at an industrial scale.

Their projects are designed to transform the traditional energy landscape by promoting sector coupling, innovation and the deployment of alternative energy technologies. In particular, Power2X aims to reduce carbon emissions and enhance the flexibility and resilience of the future energy system by integrating renewable generation, storage technologies and hydrogen-based solutions. The company's collaboration with various stakeholders positions it at the forefront of advancing sustainable energy systems worldwide [73].

The cooperation with Power2X provides valuable industry context for this research, ensuring that the

optimisation and modeling work developed in this thesis aligns closely with the real-world challenges and opportunities faced by emerging green hydrogen projects.

1.4.2. Research questions

The research is guided by the following two central questions:

- How can one determine the optimal dispatch strategy of a hybrid energy system that combines renewable electricity, energy storage, hydrogen production and participation in the day-ahead market, for the purpose of long-term profitability assessment?
- Does adding a battery and participation in the day-ahead market improve the hydrogen production plant's long-term profitability, and to what extent is this influenced by the system configuration and external conditions?

To systematically address these main questions, the following sub-questions are formulated:

Optimal dispatch strategy modeling

- What optimisation approaches can be applied to determine profit-maximising dispatch decisions under different assumptions regarding foresight and uncertainty?
- How can uncertainty in wind generation and electricity prices be modeled and incorporated into the system's dispatch strategy?

Sensitivity analysis and evaluation of system configurations

- How do different system configurations, defined by the inclusion or exclusion of battery storage, market participation and fixed renewable power contracting, affect the operational flexibility and economic outcomes of the hybrid energy system?
- How does the system's operational and economic performance respond to changes in key techno-economic parameters?
- How do variations in external market conditions, based on historical electricity price profiles and wind generation data from multiple years, influence system performance and profitability?

1.5. Report outline

To address the research questions, this thesis is structured as follows. Chapter 2 provides the necessary background on the energy technologies and markets relevant to this study, including hydrogen production and use cases, battery energy storage systems, electricity wholesale markets and long-term renewable power contracts. Chapter 3 reviews existing literature on optimal energy management models. It begins with an overview of key mathematical techniques for optimisation under uncertainty, followed by a classification framework to position this thesis within the broader literature. The chapter then discusses a selection of key related works and concludes by clarifying the scope and contribution of this research. Chapter 4 outlines the methodological framework, describing the system configurations used for the case studies, the model workflow, key assumptions, input data and model parameters. Chapter 5 presents the mathematical formulation of the optimisation model, defining the objective function, decision and state variables and the set of system constraints. The problem is classified as a MILP and three solution techniques are introduced to reflect different levels of foresight and uncertainty incorporation. Chapter 6 elaborates on these advanced solution methods, including rolling horizon and two-stage stochastic optimisation. It also describes the construction, implementation and validation of the scenario generation procedure used to model uncertainty in electricity prices and wind generation. Chapter 7 presents and analyses the model results across different system configurations and optimisation approaches. It compares operational and financial performance, evaluates the role of battery storage and market interaction, and includes sensitivity analyses on key input parameters. Finally, Chapter 8 presents the conclusions and discussion, summarising the main findings, addressing the limitations of the research. This is followed by Chapter 9, which provides recommendations for future research.

2

Background Information: Energy Technologies and Markets

This chapter provides the necessary background information on the energy technologies and market structures relevant to this thesis. Section 2.1 outlines the key components of the hydrogen value chain, including production, storage, transport and end-use applications. It also highlights recent developments in the Dutch hydrogen infrastructure, policy support mechanisms and market initiatives. Section 2.2 presents an overview of the European electricity markets, with particular focus on day-ahead market operations and the role of Power Purchase Agreements (PPAs) in renewable energy contracting. Finally, Section 2.3 introduces the concept of Battery Energy Storage Systems (BESS), discussing their technical characteristics, operating principles and relevance within integrated energy systems.

2.1. Hydrogen

Hydrogen (H_2) is the simplest and lightest element in the periodic table, and it is also the most abundant element in the universe [40]. Despite its abundance, hydrogen does not occur naturally in its pure form and cannot be captured directly from the atmosphere. Instead, it is found in compounds such as water (H_2O) and fossil fuels like methane (CH_4), from which it must be extracted. Releasing hydrogen from these compounds requires energy [40].

Hydrogen is a versatile energy carrier that is expected to play an increasingly important role in the decarbonisation of the global energy system [25]. To understand its role, it is useful to consider the structure of the hydrogen value chain. This can broadly be divided into three segments. The upstream segment involves hydrogen production, either from fossil fuels or through water electrolysis using renewable electricity. The midstream segment encompasses storage and transportation, which are essential to connect production sites with demand centres. Finally, the downstream segment covers end-use applications across sectors such as industry, transport and power. The following sections of this chapter discuss each of these stages in more detail.

2.1.1. Grey, blue and green hydrogen

Depending on the process and energy input, hydrogen production is labeled with different colours, each signifying a unique technological and environmental profile [40]. Today, the vast majority of global hydrogen production comes from fossil fuels [25, 42, 87]. About 95% of the current global hydrogen supply is “grey” hydrogen, produced mainly through steam methane reforming (SMR) of natural gas or gasification of coal [87]. These fossil fuel-based methods result in a substantial CO_2 footprint and are not compatible with achieving net-zero targets [40].

To replace grey hydrogen with a cleaner form of production, two main routes are under consideration: green hydrogen and blue hydrogen [40]. There are also other, less common and less mature hydrogen production routes. These include pink hydrogen from nuclear power, turquoise hydrogen from methane pyrolysis, white hydrogen from naturally occurring underground sources and hydrogen from biomass-based processes such as gasification or fermentation [40, 42, 87]. These pathways are still in early stages of development and are not further discussed in this thesis.


Blue hydrogen is a low-carbon variant of grey hydrogen that incorporates carbon capture and storage (CCS) technologies. Like grey hydrogen, it is produced by reforming fossil fuels but it seeks to mitigate carbon dioxide emissions by capturing and storing the CO₂, often in geological formations such as depleted oil and gas reservoirs or saline aquifers [40]. This approach enables continued use of existing fossil fuel assets and infrastructure, while reducing overall greenhouse gas (GHG) emissions compared to conventional hydrogen production without emission mitigation [40].

However, the environmental benefits of blue hydrogen are highly dependent on two critical factors: the extent of methane leakage across the natural gas supply chain and the effectiveness of the CO₂ capture process. Significant emissions reductions are only possible if methane leakage is extremely low (below 0.2%) and CO₂ capture rates are close to 100% [40]. In addition, blue hydrogen faces economic and technical challenges, including exposure to volatile natural gas prices and the development of CCS infrastructure [42]. Despite these challenges, blue hydrogen is often regarded as a transitional solution to help establish clean hydrogen markets while green hydrogen production scales up [40, 75].

A long-term, sustainable alternative is green hydrogen, which is produced via the electrolysis of water using electricity generated from renewable sources such as wind or solar energy. Because no fossil fuels are involved and the electricity input is carbon-free, green hydrogen is considered a zero-emission option at the point of production [40, 41]. While the underlying electrolysis technologies are well understood, large-scale green hydrogen production remains limited today due to high costs and constraints in renewable energy availability [40]. The next section discusses the production of green hydrogen in more detail, focusing on the electrolysis process and the technologies involved.

2.1.2. Hydrogen production from water electrolysis

Water electrolysis is a widely used method for producing hydrogen using electricity to split water molecules into hydrogen and oxygen [87]. As illustrated in Figure 2.1, the process takes place in an electrolyser, where an electric current is applied across two electrodes submerged in an electrolyte solution. At the cathode, water is reduced to form hydrogen gas (H₂), while at the anode, water is oxidised to produce oxygen gas (O₂) [94]. The overall reaction is:



The performance of the electrolysis process can be improved through the use of electrolytes and electrocatalysts, which enhance reaction kinetics and increase energy efficiency [87].




Figure 2.1: Schematic representation of the water electrolysis process for hydrogen production [94].

In practice, an electrolyser system consists of two main parts: the electrolyser stack, where the electrochemical splitting of water into hydrogen and oxygen occurs, and the balance of plant (BoP), which includes supporting equipment. The BoP typically comprises systems for water treatment, cooling, gas purification and compression, power conversion, and sometimes hydrogen storage. Since most power sources such as offshore wind farms supply alternating current (AC), a rectifier is used to convert AC to direct current (DC). These rectifiers must provide a stable and controllable current to optimise energy use. Additionally, electrolysers require ultra-pure deionised water, as impurities can significantly reduce system lifetime and performance.

While the fundamental chemistry of electrolysis is consistent, electrolyser technologies vary significantly in terms of materials, operating temperatures, efficiencies, cost structures, flexibility and maturity levels. Table 2.1 provides a simplified comparison of three main electrolyser types: Alkaline Electrolysis (AE), Proton

Exchange Membrane Electrolysis (PEM), and Solid Oxide Electrolysis (SOE), based on multiple sources [32, 39, 87, 93, 94].

Alkaline Electrolysers (AE): AE is the most mature and widely deployed electrolysis technology. It operates at moderate temperatures (typically 30–80°C) using a liquid alkaline electrolyte. AE systems are characterised by low capital and operating costs and are well-suited for continuous operation with steady power input. They are known for their long lifetime and robust durability. However, AE systems exhibit limited operational flexibility. They are slower to start up and less responsive to load changes, making them less ideal for dynamic operation with intermittent renewable energy. Furthermore, the hydrogen they produce typically requires additional purification to meet high-purity standards.

Proton Exchange Membrane Electrolysers (PEM): PEM electrolysers use a solid polymer membrane as the electrolyte and operate at lower temperatures (70–90°C). PEM systems offer several advantages, including high efficiency, compact design, and rapid response times, which make them well-suited for coupling with variable renewable energy sources (RES) such as wind or solar. They can operate across a wide range of loads, with fast start-up and shut-down capabilities, and produce very high-purity hydrogen. However, these benefits come at the cost of higher capital and operational expenses, largely due to the use of scarce and expensive noble metal catalysts such as platinum and iridium. PEM stacks also tend to have shorter lifespans compared to AE. However, significant reductions in both investment costs and operational expenses of PEM electrolysers are expected as the technology continues to mature.

Solid Oxide Electrolysers (SOE): SOE operates at significantly higher temperatures (typically above 700°C), using water in the form of steam rather than liquid water. At these conditions, electrochemical reactions can reach very high efficiencies, theoretically around 100%, and the systems can potentially make use of waste heat from industrial processes. However, they are currently less mature and not as widely commercialised as AEs or PEMs, and their high operating temperature makes them less suitable for direct coupling with intermittent renewable sources. Despite its high efficiency potential, the low technology readiness level and durability concerns currently limit its deployment.

Table 2.1: Simplified comparison of water electrolysis technologies: Alkaline Electrolysis (AE), Proton Exchange Membrane Electrolysis (PEM), and Solid Oxide Electrolysis (SOE). The table compares key performance indicators including efficiency, operating temperature, capital expenditure (CAPEX), operational expenditure (OPEX), system lifetime, operational flexibility (i.e., ability to ramp up/down and perform frequent start/stop cycles), hydrogen purity and technology maturity level. Based on information from [32, 39, 87, 93, 94].

	Eff. [%]	Temp [°C]	CAPEX	OPEX	Lifetime	Flexibility	Purity	Maturity
AE	60–80	30–80	Low	Low	Long	Low	Low	High
PEM	80–90	70–90	Medium	Medium	Medium	High	High	Med
SOE	90–100	700–1000	High	High	Short	Low	Med	low

This thesis focuses on green hydrogen production using AE, as they represent the most commercially mature and cost-effective option currently available for large-scale applications [93].

Electrolyser efficiency

Electrolyser efficiency is a key performance metric, especially in systems powered by renewable energy sources, which often supply variable rather than continuous power. Under such conditions, the efficiency of electrolysers at partial load becomes particularly relevant [48]. Efficiency is typically defined as the ratio between the energy content of the produced hydrogen and the electrical energy input. This is usually expressed based on either the Lower Heating Value (LHV) or Higher Heating Value (HHV) of hydrogen. HHV includes the latent heat of vaporisation of water, while LHV excludes it. However, for system-level evaluations that compare process chains or fuels, LHV is generally preferred because it reflects the usable energy when hydrogen is combusted in its gaseous state [8].

Electrolyser performance varies with operating conditions such as current density, temperature and pressure [3]. Typically, system-level efficiency declines as current density increases, while lower temperatures and higher pressures may further reduce performance. Both AE and PEM electrolysers exhibit a distinction between their DC efficiency, linked directly to the electrochemical reaction, and system efficiency, which accounts for auxiliary components such as cooling, water purification and power electronics. At part-load

operation, auxiliary consumption becomes proportionally larger, leading to reduced system efficiency even if DC efficiency remains relatively stable. This behaviour is illustrated in Figure 2.2, where the divergence between DC and system efficiency is shown to grow at low operating loads [48].

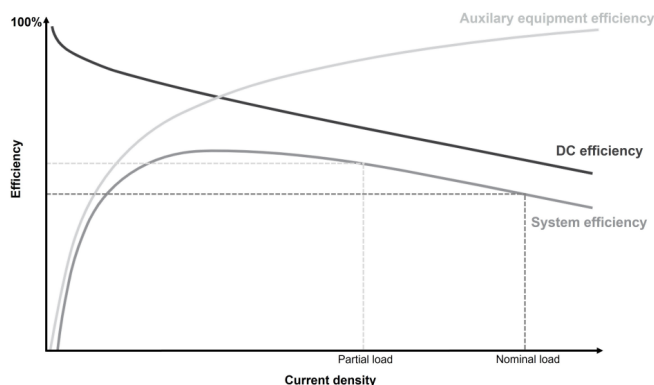


Figure 2.2: Illustration of non-constant electrolyser DC and system efficiencies at varying load levels. Adapted from Siemens Energy [48].

In this thesis, electrolyser efficiency is modeled with a fixed conversion efficiency, defined as the ratio of energy stored in hydrogen to the electrical energy input. Although real-world electrolyser performance depends nonlinearly on several factors, assuming a constant efficiency is a widely adopted simplification in system-level studies to maintain tractability and avoid nonlinear formulations [2]. Nevertheless, to approximate the real-world behaviour where efficiency drops significantly at very low power levels, the model enforces a minimum loading threshold (typically around 30% of the electrolyser's rated capacity) below which the unit is shut-down. This emulates the concave relationship between power input and efficiency without increasing model complexity [1, 51].

2.1.3. Hydrogen storage and transport infrastructure

In this thesis, hydrogen storage and transport infrastructure are not modeled explicitly. It is assumed that hydrogen offtakers are responsible for all downstream logistics, including compression, purification and delivery. However, it is important to acknowledge that these components represent a significant share of total system costs and will play a crucial role in the development and operation of future hydrogen markets. Hydrogen storage and transport are essential for enabling a flexible and large-scale hydrogen economy, particularly when production occurs in locations far from demand centres. Despite hydrogen's high energy content by mass, its low volumetric energy density poses technical and economic challenges for both storage and distribution [42]. Several storage technologies exist, including compressed hydrogen, liquefied hydrogen, and chemical storage.

Compressed gas storage

Compressed gas storage is currently the most widely used due to its simplicity and rapid discharge capability [87]. Hydrogen is compressed to high pressures, often up to 700 bar, to increase its density. At 700 bar, it reaches a volumetric density of 36 kg/m^3 [42, 94]. It offers advantages like simple operation, rapid charge/discharge cycles, and relatively low energy requirements compared to liquefaction technologies [87]. Hydrogen can be stored in pressure vessels, like spherical or pipe vessels, or geologically in formations such as salt or lined rock caverns [38].

An obstacle, however, is the requirement for big systems due to hydrogen's low density as a gas, needing over twice the volume of natural gas for the same energy output [40, 87]. Compressed hydrogen is also highly flammable, and leaks can be explosive [87]. For large-scale and long-duration applications, geological storage, particularly in salt caverns, is considered promising due to its low cost and high stability [40, 94]. However, their development is limited by technical aspects such as the tightness of boreholes and the transfer capacity of surface installations [94]. Scaling up geological storage requires careful planning, as some suitable sites may already be allocated to the storage of methane, biomethane or CO_2 during the energy transition [40].

Liquid hydrogen storage

For liquid hydrogen storage, hydrogen is converted into liquefied form by cooling it to very low temperatures (-253°C) [94]. This method achieves a higher volumetric density than compressed gas [42]. Liquefaction is an energy-intensive process, consuming a significant portion of the hydrogen's potential energy (25-35% lost) [40, 42, 87, 94]. The technique to liquefy is much more difficult and consumes more energy than compression [92], but it is a relatively well-established technology at a small scale and it does not require complex re-conversion compared to chemical hydrogen carriers [88]. Currently, liquid hydrogen storage is primarily used for special high-tech applications like space travel and is not widely commercialised [94]. Ports are also largely unprepared for liquid hydrogen trade [42].

Chemical storage

Chemical storage involves binding hydrogen to other materials or compounds, allowing it to be stored and transported in more stable and often denser forms. One such method uses metal hydrides, where hydrogen is stored in solid form by forming reversible chemical bonds with metals [94]. This approach is considered relatively safe and offers high volumetric storage efficiency [87]. However, it is still an emerging technology, with ongoing research focused on improving adsorption properties, reducing material costs and overcoming the challenge that some hydride systems require high temperatures for charging and discharging [87].

Another well-established option is ammonia (NH₃), which has a high hydrogen density and can be stored at near-ambient conditions (around 1 bar and 25°C) [87]. Ammonia supports long-term, stable storage and is especially suitable for large-scale export and import [42]. However, it poses safety risks due to its toxicity, may lead to nitrogen oxide emissions if leaked and requires energy-intensive processes to extract pure hydrogen for end use [40]. Methanol (CH₃OH) is another hydrogen carrier that offers a higher energy density than ammonia, but it contains less hydrogen by weight and volume [42, 87]. Finally, Liquid Organic Hydrogen Carriers (LOHCs) and other solid hydrogen carriers are being actively explored. These systems allow hydrogen to be chemically bound and released through reversible reactions, but they currently face barriers related to round-trip efficiency, cost and scalability, limiting their near-term commercial viability [42, 94].

Transport

In general, transport costs are currently high but are expected to decrease through economies of scale, reduced project risks and technological advancements [40, 42]. Depending on the volume of hydrogen, delivery distances and local circumstances, hydrogen can be transported from the production site and distributed to the end consumer in multiple ways. There are three main categories for hydrogen transport based on currently existing and mature transport systems: road transportation, pipeline transportation, and shipping [92].

Pipelines are the most cost-effective option for transporting hydrogen over short to medium distances and at large volumes, particularly when existing natural gas infrastructure can be repurposed with minimal upgrades [25, 40, 42]. Repurposing pipelines can reduce costs by 65–94% compared to new builds, but requires technical modifications to address challenges such as hydrogen embrittlement, increased material wear and the need for additional compression [25, 40, 42]. Additionally, safety, environmental, and regulatory considerations must all be evaluated before repurposing infrastructure [25, 40, 42]. Nonetheless, planning is underway for medium-distance and backbone transmission infrastructure, with an EU-wide hydrogen network envisioned after 2030 [25].

Where pipelines are impractical, alternative transport modes such as trucks and tube trailers are used, typically for short distances or last-mile distribution [42, 94]. These methods can involve hydrogen in either liquefied or compressed form, but are generally less efficient and more costly due to hydrogen's low volumetric energy density compared to conventional fuels like diesel [42, 87].

On the other hand, for long-distance or intercontinental trade, shipping becomes the most viable option. Due to hydrogen's low density, it must be transported in the form of liquefied hydrogen, ammonia or another hydrogen carrier. Among these, ammonia is particularly attractive due to its relatively high hydrogen density and the availability of existing infrastructure, although toxicity and re-conversion energy losses remain key challenges [40, 87]. As the green hydrogen economy becomes increasingly global, transporting hydrogen by sea is expected to play a crucial role. This is particularly significant for long-distance trade between renewable-rich regions producing hydrogen and countries with high demand for hydrogen imports [68].

Beyond transport considerations, safety remains a critical concern across all hydrogen infrastructure due to its unique physical properties, such as high flammability and its tendency to cause metal embrittlement [40, 42]. While these risks are well understood, the implementation of appropriate national and international

standards is essential to ensure safe deployment [40]. Countries such as the Netherlands are already developing dedicated guidelines for the safe handling and use of gaseous hydrogen [61].

2.1.4. Use cases for hydrogen

Hydrogen serves as a versatile energy carrier [40, 42]. It has a high specific energy density, providing three times more energy than gasoline combustion per unit mass [87, 94]. Its use produces no direct CO₂ emissions and almost no air pollution, offering significant potential for decarbonisation across various sectors. [25]. When combusted, hydrogen can generate high-temperature heat exceeding 1000°C, which is useful for industrial applications, although it may lead to the formation of nitrogen oxides (NO_x) under certain conditions [40]. Alternatively, hydrogen can be used in fuel cells to generate electricity through an electrochemical reaction with oxygen, emitting only water as a by-product [40, 42].

In addition to its clean end use, hydrogen also offers a solution to one of the key challenges of renewable energy: variability. Green hydrogen, produced through electrolysis powered by renewable electricity, enables the conversion of excess wind or solar power into a storable, transportable gas [41]. It can be stored over various time scales, from hours to seasons, and later reconverted to electricity when needed. This makes hydrogen a valuable tool for grid balancing, reducing curtailment, enhancing flexibility and improving the utilisation of renewable generation capacity [41].

Beyond its role in balancing electricity supply and demand, hydrogen's versatility is reflected in its wide range of possible end uses. It can be reconverted into electricity or used directly in industrial processes, transport, or the production of chemicals and fuels. This flexibility is often captured by the concept of Power-to-X (PtX), which refers to the conversion of renewable electricity ('power') into other energy carriers or products ('X'), such as hydrogen, methane, methanol, ammonia or synthetic fuels [40]. Currently, hydrogen accounts for a small share of global and EU energy consumption [25]. Global demand stands just below 100 million tonnes per year, which is already five times higher than in 1975 but still representing only about 3% of global final energy use [42]. Today, hydrogen is used predominantly as a feedstock in the (petro)chemical sector, notably in oil refining, ammonia production for fertilisers and methanol synthesis [40, 42]. Looking forward, hydrogen is expected to play a central role in decarbonising hard-to-abate sectors such as heavy industry, (green) steelmaking, chemicals and long-distance transport, including shipping and aviation [25, 42]. Its use in the built environment is also possible, for instance in heating, but is generally considered less cost-effective than alternatives like electric heat pumps. Early large-scale demand for low-carbon hydrogen is anticipated to be driven primarily by industrial decarbonisation efforts [13].

2.1.5. Dutch hydrogen infrastructure: national backbone

The Netherlands is actively developing a national hydrogen backbone, led by HyNetwork, a subsidiary of Gasunie, the Dutch state-owned natural gas infrastructure company [38]. The backbone will form the core of the country's hydrogen transport system and is expected to span 1400 km by 2030, primarily using repurposed natural gas pipelines, which will make up around 85% of the total network [38]. The remaining segments will consist of newly built hydrogen-dedicated pipelines to ensure full connectivity across key regions. As shown in Figure 2.3, the network will link five major industrial clusters: Eemshaven (Northern Netherlands), Amsterdam/IJmuiden (North Sea Canal Area), Rotterdam/Moerdijk, Westerscheldemond (Zeeland), and Chemelot (Limburg). These clusters serve as both significant demand centres and potential supply hubs, given their proximity to renewable resources and import terminals.

In addition to enabling domestic hydrogen transport, the network will provide cross-border connections with Belgium and Germany by 2030. The coastal location of several clusters allows integration with offshore wind farms and future hydrogen imports, reinforcing the Netherlands' ambition to become a central hydrogen hub for Northwest Europe [61].

To complement the transport network, the Netherlands is also investing in large-scale hydrogen storage. The country has favourable geological conditions for salt cavern development, and HyStock, also a Gasunie initiative, is currently developing four hydrogen storage caverns with a total planned capacity of 216 GWh [38]. By integrating storage with the national backbone, the Netherlands aims to balance supply and demand fluctuations and strengthen its role in the regional hydrogen market.

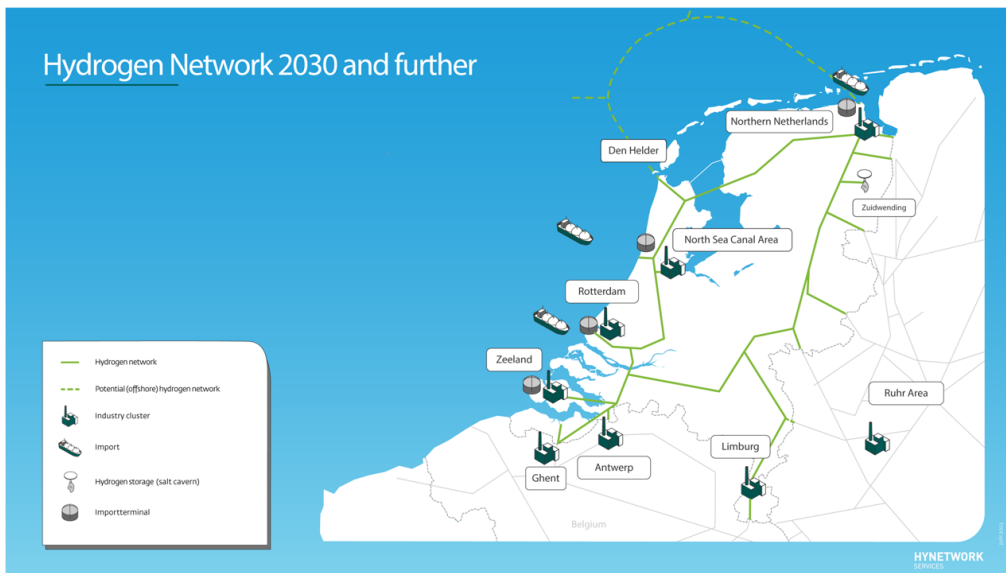


Figure 2.3: Overview of planned Dutch hydrogen infrastructure development [38].

2.1.6. Hydrogen spot market and long-term contracting

As hydrogen production and consumption scale up, the development of a liquid spot market is increasingly seen as a critical step toward a mature and efficient hydrogen economy [17, 25]. A well-functioning spot market would enable transparent pricing, short-term trading and efficient dispatch based on marginal costs [17]. For flexible producers, it opens opportunities to respond to price volatility and optimise operations, while also improving market access and supporting day-to-day balancing [17]. The Netherlands is well-positioned to develop such a market. It benefits from vast potential for green hydrogen production from offshore wind, established industrial clusters with significant hydrogen demand and key port infrastructure capable of receiving hydrogen imports [38]. The national hydrogen backbone plays an essential role here, as it will connect large demand and supply centres. Combined with the country's natural gas legacy, this positions the Netherlands as a logical gateway to a broader northwest European hydrogen economy [17].

However, several conditions must be in place before a fully liquid hydrogen spot market can emerge. These include a reliable transport network, a sufficient level of supply and demand to ensure liquidity and a standardised digital trading platform [17, 64]. At present, limited liquidity and a lack of market standardisation remain key barriers [16]. Nonetheless, the Netherlands is actively advancing on all these fronts. Alongside the rollout of the national hydrogen backbone, the HyXchange initiative is developing a dedicated trading platform. Through regional spot market simulations and the creation of the HyClix price index, HyXchange aims to lay the foundation for a robust and transparent hydrogen market [17].

Until the spot market becomes sufficiently liquid and standardised, hydrogen transactions will primarily rely on long-term bilateral contracts [64]. These contracts are negotiated directly between a hydrogen producer and a buyer without involving a centralised market platform. Such agreements, commonly referred to as Hydrogen Purchase Agreements (HPAs), typically span 10 to 15 years and provide fixed terms regarding volumes, prices and delivery conditions [16]. By offering long-term revenue stability, HPAs play a crucial role in de-risking investment and securing financing for early-stage hydrogen projects, particularly in the absence of a fully liberalised and liquid spot market [79]. Over time, as trading platforms develop and infrastructure becomes more interconnected, the market is expected to evolve toward a hybrid model in which long-term bilateral contracts coexist with flexible and transparent spot market mechanisms [64].

This thesis assumes a take-as-produced HPA, in which the offtaker accepts hydrogen as it becomes available, typically following the variable production profile of the renewable power source. This structure closely resembles Power Purchase Agreements (PPAs) in the electricity sector and is further discussed in Section 2.2.6. An alternative is the baseload contract, where the offtaker expects a continuous and constant supply of hydrogen regardless of production variability. Such contracts are likely to be favoured by industrial consumers with steady, inflexible demand. However, meeting this demand poses additional operational challenges for producers who rely on intermittent renewable electricity, which does not naturally provide a stable, baseload output.

2.1.7. Subsidies and regulations

The development of the hydrogen economy is strongly influenced by public policy support. Regulatory frameworks and financial incentives play a vital role in bridging the cost gap between green or low-carbon hydrogen and fossil-based alternatives. Understanding these evolving policy frameworks is key to assessing the technical and economic feasibility of hydrogen projects, particularly in a rapidly changing regulatory environment. In the European Union, the Renewable Energy Directive (RED II) provides the overarching legal framework for promoting renewable energy, including hydrogen [85]. To operationalise specific parts of this directive, the European Commission has adopted two key Delegated Acts, which set out the detailed rules for certifying hydrogen as a Renewable Fuel of Non-Biological Origin (RFNBO) [85]. These acts define how producers must prove that the electricity used for electrolysis is truly renewable. Since green and grey hydrogen molecules are chemically identical, such certification schemes are essential to verify the origin of hydrogen and provide transparency to governments, regulators and end users [40].

To qualify as renewable, hydrogen production must comply with several criteria outlined in the Delegated Acts governing the electricity used for electrolysis. This includes demonstrating a verifiable link to renewable electricity, either through a direct connection to a renewable source or via a PPA. In addition, producers must meet requirements for additionality, temporal correlation and geographical correlation [85]. Additionality ensures that the electricity is sourced from newly built renewable capacity; temporal correlation requires that hydrogen is produced at the same time as renewable electricity is generated; and geographical correlation mandates proximity or grid-based linkage between the hydrogen production site and the renewable source [20, 85]. These certification rules are intended not only to support EU climate goals but also to facilitate international trade by ensuring a clear definition of renewable hydrogen.

Alongside regulatory instruments, financial support mechanisms are critical to enable early-stage projects and accelerate market growth [25]. These include instruments such as the EU Innovation Fund, which provides grants to large-scale demonstration projects aimed at low-carbon technologies; national hydrogen strategies, which often allocate public funding to support hydrogen infrastructure and production pilots (as seen in countries like the Netherlands and Germany); and contracts for difference (CfDs), which offer investors a guaranteed strike price by compensating the difference between a pre-agreed price for renewable hydrogen and the prevailing market price [85]. These instruments can offer price guarantees or bridge the cost gap between renewable and fossil hydrogen, making them essential tools in advancing the commercial deployment of green hydrogen.

2.2. Electricity Markets

Since the liberalisation and restructuring of electricity markets that started in the 1990s, electricity has transitioned from a centrally dispatched public service to a commodity traded on competitive markets [23, 56]. In this liberalised framework, electricity is bought and sold across different wholesale markets depending on the timing of delivery. These include the forward market, day-ahead (DA) and intraday (ID) spot markets, and a real-time balancing market. Together, these interconnected markets ensure that electricity supply and demand are aligned continuously, from years in advance down to real-time. Electricity trading is facilitated by Transmission System Operators (TSOs), which maintain grid stability and ensure the secure transmission of power. In the Netherlands, this role is fulfilled by TenneT, a government-owned entity responsible for operating the high-voltage grid and organising a stable and transparent electricity market [84].

Unlike many other commodities, electricity must be consumed at the moment it is produced, since it cannot be stored efficiently at scale. This physical constraint introduces operational complexity and results in high price volatility, especially during periods of fluctuating demand or intermittent renewable energy supply. A crucial aspect of electricity systems is the concept of dispatchability, which reflects a generator's ability to adjust output on short notice. Dispatchable plants, such as gas-fired or coal-fired units, can be ramped up or down to meet variations in demand. In contrast, intermittent RES are non-dispatchable and depend on the availability of wind or solar irradiation. Their variable output complicates real-time balancing and increases the system's reliance on flexible trading mechanisms and ancillary services.

Figure 2.4 illustrates the sequence of the main electricity markets and their progression over time. Trading starts in the forward market, with contracts typically agreed upon between five years and one month before delivery. Closer to real-time, participants trade electricity in the DA market, where prices are determined via daily auctions for each hour of the following day. The ID market allows further adjustments in response to updated forecasts and the balancing market ensures that any real-time mismatches between generation and consumption are resolved. Together, these markets form the wholesale electricity market, where large

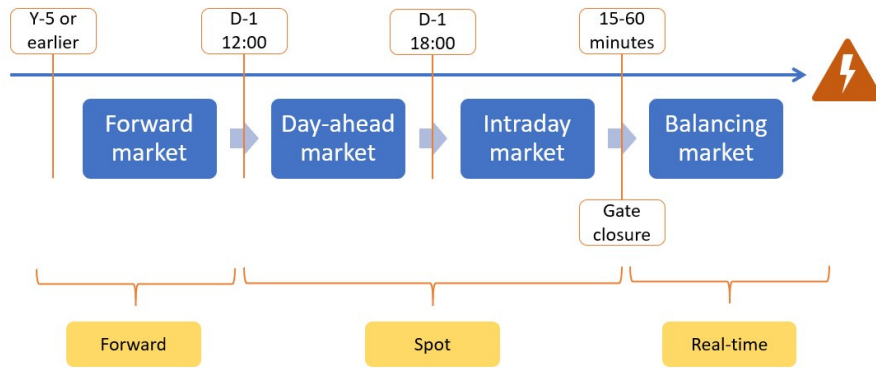


Figure 2.4: Overview of the sequential structure of electricity markets based on trading timeframes. Electricity is first traded in the forward market, followed by the spot markets comprising the day-ahead and intraday markets, and finally in the real-time balancing market shortly before delivery. Adapted from KYOS [46].

generators, retailers and industrial consumers engage in continuous trading.

2.2.1. Forward market

The forward market is a financial marketplace where electricity is traded through contracts that specify delivery at a future date. These forward contracts, often also referred to as futures when standardised, typically span delivery periods from five years up to one month in advance of actual consumption [24]. Forward markets allow participants to hedge against price risks by securing electricity prices and delivery volumes well in advance, thereby reducing exposure to market volatility. Participants in the forward market include electricity producers, retailers, industrial consumers and financial traders. By locking in prices in advance, these actors can secure revenue streams or cost forecasts, which is particularly valuable in capital-intensive sectors like energy generation. As such, the forward market plays a key role in supporting investment decisions and long-term planning in the electricity sector.

Forwards are generally traded bilaterally, or over-the-counter (OTC). In OTC trading, two parties negotiate the terms of the contract directly without using a centralised exchange. These bilateral deals offer flexibility in contract size, duration and pricing structure, but lack transparency and are not publicly reported [23]. As a result, it is difficult to obtain comprehensive price signals or volume data from OTC activity.

2.2.2. Day-ahead market

The day-ahead market is the most prominent segment of the spot market, where electricity is traded for delivery on the following day in hourly intervals [24]. Each day, market participants submit their bids and offers until 12:00 Central European Time (CET), after which the market clearing process is conducted through a uniform-price auction [23]. In this auction, market participants submit bids specifying the quantity of electricity they wish to buy or sell and the corresponding prices. The collected bids form aggregated supply and demand curves, and their intersection determines the market clearing price (MCP) and market clearing volume (MCV) for each hourly delivery period. Figure 2.5 illustrates this mechanism.

All accepted supply and demand bids are settled at the MCP, meaning that all sellers receive and all buyers pay the same price, regardless of their individual bid. This marginal pricing mechanism ensures that the price is set by the most expensive unit required to meet demand in that hour, commonly referred to as the “marginal plant”. As a result, lower-cost generators are dispatched first, promoting cost-efficient system operation. A more detailed discussion of marginal pricing and the role of the merit order is provided in Section 2.2.5. Trading is carried out through the European Power Exchange (EPEX), and the DA prices used are therefore also referred to as EPEX spot prices [84]. The prices established in the DA market serve as a benchmark for the value of electricity in each hour and form the basis for financial hedging and planning by market participants.

For the purposes of this thesis, only trading on the DA market is considered, and the hourly DA spot price is used as the electricity price input. The DA prices tend to show the lowest level of variation and the clearest seasonal patterns among all markets. Research by Mulder [60] has shown that prices in this market follow recognisable patterns throughout the year, such as higher prices in winter and lower prices in summer. These trends are driven by predictable changes in electricity supply and demand, as well as by the way the market is structured and how participants behave strategically within it.




Figure 2.5: Market clearing mechanism in the day-ahead electricity market. The intersection of the aggregated supply and demand curves determines the Market Clearing Price (MCP) and Market Clearing Volume (MCV). Adapted from Nmeou et al. [65].

2.2.3. Intraday market

After the DA market closes, trading continues on the intraday market. This market enables market participants to update their positions closer to the time of delivery, responding to forecast errors or unexpected changes in demand and generation. Trading on the ID market typically occurs in hourly and quarter-hourly intervals, with contracts being traded up to five minutes before delivery. Electricity is traded continuously among market participants and trades are executed immediately when buy and sell orders match. This price-setting mechanism is known as the pay-as-bid principle [23]. Although trading volumes on the ID market are currently lower than those on the DA market, they are expected to grow in response to the increasing penetration of RES [24]. The inherent uncertainty and variability of RES generation make it necessary for market participants to adjust their positions closer to delivery [12].

As delivery time nears, demand forecasts become more accurate due to the availability of real-time information, such as weather updates and power plant availability. Nonetheless, the ID market remains more volatile than the DA market. This volatility stems from last-minute shifts in consumption or unexpected changes in supply, and is further intensified by the shorter time available for market participants to respond to such fluctuations.

2.2.4. Balancing market

Even after the ID market closes, imbalances may persist between scheduled generation and actual consumption. These mismatches are resolved in the balancing market, which is operated in real-time by the TSO [84]. Market participants such as generators and demand-response providers submit flexible bids indicating how much power they can adjust on short notice. The TSO then activates the most cost-effective bids to stabilise the high-voltage grid and maintain its frequency close to 50 Hz [84].

The balancing market is essential for ensuring short-term grid security, particularly in systems with limited storage and high shares of RES [84]. As RES penetration increases, so does the potential for sudden and unforeseen imbalances, making real-time flexibility even more critical. In addition, the balancing market facilitates cross-border cooperation by enabling real-time electricity exchange between countries, helping balance supply and demand across regions when needed [23].

2.2.5. Pricing mechanism and merit order effect

In the DA market, electricity prices are determined by two key principles: the marginal pricing mechanism and the merit order model [95]. Under marginal pricing, each electricity supplier submits bids that reflect their marginal cost of production, which refers to the cost of generating one additional unit of electricity. These bids are arranged from lowest to highest, resulting in what is known as the merit order curve [95]. This curve represents the economic ranking of available generation technologies based on their marginal costs.

To determine the MCP, the aggregated demand curve is placed on top of the merit order curve. The point at which supply meets demand defines both the MCV and the MCP, as illustrated in Figures 2.5 and 2.6. The price is set by the marginal plant, which is the most expensive generating unit required to meet total demand [95]. All generators whose bids are accepted receive this same price per unit of electricity, regardless

of their original bid. This uniform pricing mechanism ensures that cheaper plants are dispatched first and encourages cost-efficient system operation. As a result, most generators earn more than their marginal operating costs, with the difference between the MCP and their production cost allowing them to recover capital expenditures and earn a profit [6].

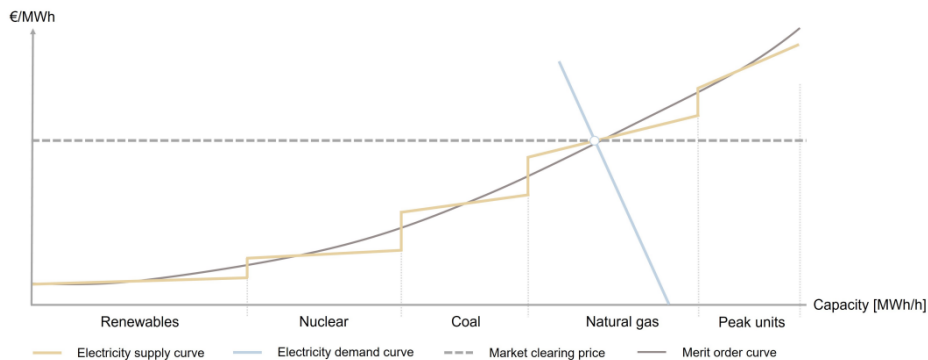


Figure 2.6: Illustration of the merit order curve, where electricity generation technologies are ranked by marginal cost. The market clearing price is set by the most expensive unit needed to meet demand. Adapted from van den Berkmortel [86].

In the Netherlands, gas-fired power plants often play a significant role in setting the MCP, particularly during periods of moderate demand. Renewable energy sources (RES) such as wind and solar typically appear on the far left of the curve because of their near-zero variable costs and priority dispatch status. A key consequence of this system is the so-called merit order effect. As the share of low-cost RES in the system increases, the merit order curve shifts to the right. This results in conventional, higher-cost generators being displaced and a lower overall market clearing price, particularly during periods of high wind or solar output. While this effect benefits consumers through reduced electricity prices, it also places economic pressure on conventional generators, whose operating hours and revenues may decline as a result [81]. This effect can be seen in Figure 2.7 and shows the changing dynamic of the power generation mix.

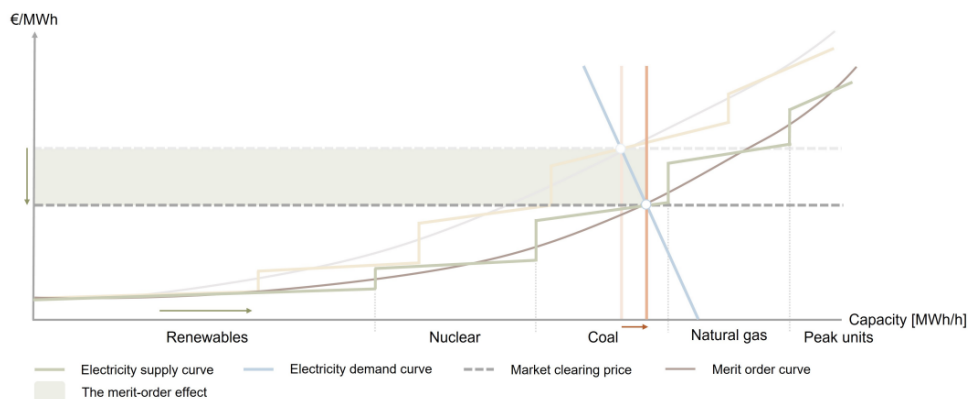


Figure 2.7: Demonstration of the merit order effect. The inclusion of low-cost renewable energy sources shifts the supply curve to the right, displacing higher-cost generation and lowering the market clearing price. Adapted from van den Berkmortel [86].

2.2.6. Long-term contracting through PPAs

The increasing penetration of renewable energy sources is fundamentally reshaping the structure and operation of electricity markets. Unlike conventional generation, these technologies are variable and non-dispatchable, making their output dependent on weather conditions and difficult to predict with precision. To remain profitable in such a setting, renewable producers must find ways to manage both supply-side uncertainty and growing exposure to price volatility. Historically, many renewable projects operated under fixed feed-in tariffs or subsidy-backed contracts, receiving guaranteed payments for every unit of electricity produced. These schemes de-risked market participation but are increasingly being phased out in favour of

more market-based frameworks. As a result, renewable energy producers are now expected to participate actively in competitive wholesale markets, including the day-ahead and intraday markets, often without perfect foresight of their output.

To complement spot market participation, Power Purchase Agreements (PPAs) are a form of bilateral long-term contract used in the power sector to reduce investment risk [36]. They are increasingly popular for renewable generation projects, such as wind and solar, because they offer revenue stability in highly volatile electricity markets [36, 57]. By locking in a price and offtake volume in advance, PPAs help derisk projects and improve their bankability.

PPAs can be classified along several dimensions, including the location of power delivery (on-site vs. off-site) and the form of settlement (physical vs. financial) [36]. On-site PPAs involve direct delivery of electricity to the buyer, often behind the meter, and avoid the need for grid connection. Off-site PPAs rely on transmission through the public grid and allow for greater geographic flexibility. In physical PPAs, electricity is delivered to the consumer either directly or via the grid. In financial (or virtual) PPAs, there is no physical delivery. Instead, the buyer and seller settle the financial difference between a fixed PPA price and the market price, typically using a Contract for Difference (CfD) structure [57]. In a CfD structure, the producer continues to sell electricity on the spot market, but if the market price falls below the agreed PPA price, the buyer pays the producer the difference. Conversely, if the market price exceeds the PPA price, the producer pays the difference back to the buyer. This ensures a fixed net revenue for the producer while allowing both parties to benefit from market exposure.

One of the most important aspects of PPA design is the choice of offtake profile, which defines how contracted volumes relate to actual generation or consumption. This also determines which party carries the volume and price risk. Various delivery structures have been identified by Mittler et al. [57], and although not exhaustive, a few common examples help illustrate the diversity in risk allocation:

- **As-produced:** The buyer takes electricity as it is generated. The buyer carries the risk of variability, curtailment and forecasting errors.
- **As-forecasted:** The buyer agrees to purchase a forecasted generation profile. The producer bears the forecasting risk and is responsible for deviations.
- **As-consumed:** The buyer only pays for electricity actually consumed at the time of use, shifting all volume risk to the producer.
- **Baseload:** A fixed quantity is delivered continuously, typically over each hour of the contract term. The seller ensures this supply, absorbing generation risk.

These profiles directly influence PPA pricing. Contracts with greater uncertainty for the seller (e.g. baseload or as-consumed) typically command a higher price, while profiles that place more risk on the buyer (e.g. as-produced) will be priced differently [57]. Selecting an appropriate offtake structure is especially important for renewable generators like wind farms, whose variable output and exposure to imbalance costs must be carefully managed through contract design.

2.3. Energy Storage Systems

Energy storage systems play a crucial role in enhancing power system flexibility and reliability, and help mitigate the mismatches between electricity generation and consumption. A wide range of storage technologies exists, typically classified into mechanical and electrochemical systems [11].

Among mechanical storage technologies, Pumped Hydro Storage (PHS) and Compressed Air Energy Storage (CAES) are the most well-known. In the case of PHS, excess electricity is used to pump water from a lower to a higher reservoir, storing energy in the form of gravitational potential. When needed, water is released to flow back down through turbines, generating electricity [11, 18]. CAES operates by compressing air into underground caverns or high-pressure tanks during periods of low demand or surplus generation, and later releasing the air to drive a turbine [11, 62].

Both PHS and CAES are well suited for large-scale and long-duration energy storage. However, their applicability is limited by geographic and technical constraints [11]. PHS requires substantial elevation differences and significant land use, while CAES depends on the availability of suitable underground caverns. In terms of performance, PHS typically offers discharge durations of 6–10 hours, while CAES systems tend to have shorter durations of around 2 hours but suffer from lower efficiency, typically around 50% [11]. In contrast, Battery Energy Storage Systems (BESS) are highly efficient and better suited for short-duration applications.

2.3.1. Battery Energy Storage Systems (BESS)

BESS store electricity in electrochemical form, relying on reversible electrochemical reactions to store and release energy. Several battery chemistries are used in practice, including lead-acid, nickel-cadmium, sodium-sulfur, and lithium-ion (Li-ion), each with distinct performance, cost and efficiency characteristics [11]. Among these, Li-ion batteries dominate utility-scale application due to their high energy and power density, low self-discharge rate and high round-trip efficiency [66]. This thesis therefore focuses specifically on Li-ion batteries. Charging and discharging efficiencies of Li-ion cells are typically around 95%, resulting in an overall round-trip efficiency of approximately 90% [66].

Figure 2.8 shows a schematic of a Li-ion battery cell. In practice, a BESS consists of multiple modules composed of such cells. Li-ion batteries operate via reversible electrochemical reactions between a positive electrode (cathode), a negative electrode (anode), and an electrolyte. During charging, lithium ions (Li^+) move from the cathode to the anode via the electrolyte, while electrons move through an external circuit. During discharging, the process is reversed: electrons flow from the anode to the cathode, generating usable electrical power [11]. Operationally, batteries are characterised by three key parameters: their power rating (MW), indicating the maximum (dis)charge rate, their energy capacity (MWh), determined by power rating times duration, and their state-of-charge (SOC), which tracks the remaining energy relative to total capacity.




Figure 2.8: Schematic representation of a lithium-ion battery, illustrating the movement of lithium ions (Li^+) and electrons (e^-) during charging and discharging. During charging, Li^+ ions migrate from the cathode to the anode through the electrolyte, while electrons flow externally to the anode. The process reverses during discharge. Adapted from Chakraborty et al. [11].

Despite their advantages, BESS face limitations compared to mechanical storage technologies. Their storage capacity is typically smaller, limiting their usefulness for long-duration applications [11]. Furthermore, they cannot charge and discharge simultaneously and must operate within fixed limits. System efficiency and battery lifespan are influenced by operating conditions, including cycling behaviour, depth of discharge, and temperature. Although Li-ion batteries can handle between 1,000 and 10,000 cycles, degradation is inevitable over time [67]. A charging cycle, or simply a cycle, refers to one full sequence of charging and discharging, typically measured as a complete round-trip of energy into and out of the battery. This does not necessarily mean charging and discharging from 0% to 100%. For example, two partial cycles of 50% depth of discharge would also count as one full cycle. The number of cycles a battery can complete before experiencing a notable decline in capacity is a key factor in determining its overall lifespan [47]. To mitigate wear and preserve battery life, the SOC is typically constrained within a 20%–90% range [47, 66, 67].

3

Literature Review

This chapter reviews existing literature on the optimisation of energy management problems for integrated electricity and hydrogen systems. Section 3.1 begins with an overview of key mathematical techniques for optimisation under uncertainty. Section 3.2 introduces a classification framework for structuring the literature on energy management optimisation. In Section 3.3, selected studies are discussed using this framework, with a focus on system configurations, problem formulation methods, and the optimisation techniques used to solve them. Finally, Section 3.4 concludes the chapter by positioning this thesis within the broader literature and identifying the research gaps it aims to address.

3.1. Mathematical techniques for optimisation under uncertainty

This section provides an overview of key mathematical techniques and solution approaches that underpin the optimisation methods applied in this thesis for decision-making under uncertainty. These techniques are particularly relevant in the context of energy systems optimisation, where future market prices, renewable energy availability and system states are uncertain. These include linear programming, stochastic programming with a two-stage structure, a rolling horizon approach and Monte Carlo simulation. These concepts lay the foundation for the solution methods detailed in later sections and chapters.

3.1.1. Optimisation of multi-period decision problems

Optimisation is the process of selecting the best solution from a set of feasible options by maximising or minimizing an objective function. A general optimisation problem can be written as:

$$\min_{x \in \mathbb{R}^n} f(x) \quad \text{subject to} \quad x \in X, \quad (3.1)$$

where $x \in \mathbb{R}^n$ is a vector of decision variables, $X \subseteq \mathbb{R}^n$ is the feasible set defined by a set of constraints, and $f(x)$ is the objective function to be minimized. A maximisation problem can be formulated analogously.

A widely used subclass of optimisation problems is linear programming (LP), in which both the objective function and all constraints are linear [5]. A linear program takes the form:

$$\min_{x \in \mathbb{R}^n} c^\top x \quad \text{subject to} \quad Ax \leq b, \quad (3.2)$$

where $c \in \mathbb{R}^n$ is a cost vector, $A \in \mathbb{R}^{m \times n}$ is a constraint matrix, and $b \in \mathbb{R}^m$ is a constraint bound vector.

In energy systems, optimisation often involves decisions made sequentially over time. This leads to a multi-period decision problem, typically formulated over a finite time horizon $t = 0, 1, \dots, T$. These models capture system evolution, such as battery charge or hydrogen storage, and respond to external factors like prices or weather. The optimisation problem can be formulated as a linear program of the following general form:

$$\max / \min_x \sum_t f(x_t, s_t), \quad (3.3)$$

$$\text{s.t.: } s_{t+1} = U(s_t, x_t, d_t) \quad \forall t, \quad (3.4)$$

$$H(x_t, s_t) = 0 \quad \forall t, \quad (3.5)$$

$$G(x_t, s_t) \leq 0 \quad \forall t, \quad (3.6)$$

$$x_{\min} \leq x_t \leq x_{\max} \quad \forall t, \quad (3.7)$$

$$s_{\min} \leq s_t \leq s_{\max} \quad \forall t. \quad (3.8)$$

Here, x_t denotes the decision variables at time t , such as power flows, hydrogen production rates, or battery charging levels. The system's state at time t is captured by s_t , for example the state of charge (SOC) of a battery or the hydrogen storage level. External inputs or disturbances, such as electricity market prices or renewable energy availability, are represented by d_t . The function $f(x_t, s_t)$ defines the objective, which typically involves minimising total system costs or maximising profit over the time horizon.

The system dynamics are described by the update function $U(s_t, x_t, d_t)$, which determines the next system state, s_{t+1} , based on current decisions, states, and external inputs. For instance, in a battery system, this could model how the SOC evolves over time depending on the charge/discharges and efficiency losses.

The constraints $H(x_t, s_t) = 0$ and $G(x_t, s_t) \leq 0$ ensure that system operation remains physically and operationally feasible. Equality constraints H may include energy balance equations or coupling between subsystems, such as requiring that total power output matches demand or market dispatch levels. Inequality constraints, G , can capture limitations such as maximum ramping rates, storage capacity limits, minimum up/down times, or electrolyser operating ranges.

3.1.2. Mixed-Integer Linear Programming

Building on the general multi-period framework introduced above, an important class of problems emerges when the optimisation involves both continuous and discrete decisions. In energy systems, discrete decisions are particularly common. For example, the model may need to determine whether or not to activate an electrolyser at a given time. These types of decisions cannot be captured using purely continuous variables and require the incorporation of discrete variables into the model. Mixed-Integer Linear Programming (MILP) extends the linear programming framework by allowing some decision variables to take on discrete values, typically integers or binary (0–1) values. Despite this added complexity, MILP retains the linearity of the objective function and constraints, enabling it to model a wide range of real-world systems while remaining compatible with efficient commercial solvers such as CPLEX and Gurobi [33].

In the context of energy system optimisation, MILP provides the flexibility to model start-up and shut-down decisions, minimum on/off times and logical interdependencies between components. It is particularly well-suited for modeling the commitment status of generation units, the activation of market bidding strategies or operational constraints that depend on discrete switching. Various studies have leveraged the use of MILP to achieve precise control over complex hybrid energy systems. For instance, Alvarez et al. [2] employ MILP to optimise the day-ahead and reserve market scheduling of a hydrogen-based virtual power plant, incorporating both continuous variables (e.g., electricity purchased, sold, production levels) and binary variables (e.g., electrolyser commitment and startup states). Similarly, Xiao et al. [91] optimise the operation of a wind-electrolytic hydrogen storage system using MILP, effectively balancing the trade-off between selling wind power directly to the electricity market or producing hydrogen for the hydrogen market. The use of MILP in these studies allows for precise control over the complex interactions between the system components, including the battery energy storage system, electrolyser, and hydrogen storage.

The MILP formulation builds on the problem defined in equations (3.3)-(3.8), extending it to include binary variables. A simplified MILP structure is presented below:

$$\max / \min_{x_t, z_t} \sum_{t=1}^T f(x_t, s_t, z_t), \quad (3.9)$$

$$\text{s.t. } s_{t+1} = U(s_t, x_t, d_t) \quad \forall t, \quad (3.10)$$

$$H(x_t, s_t, z_t) = 0 \quad \forall t, \quad (3.11)$$

$$G(x_t, s_t, z_t) \leq 0 \quad \forall t, \quad (3.12)$$

$$x_{\min} \cdot z_t \leq x_t \leq x_{\max} \cdot z_t \quad \forall t, \quad (3.13)$$

$$z_t \in \{0, 1\} \quad \forall t, \quad (3.14)$$

$$s_{\min} \leq s_t \leq s_{\max} \quad \forall t. \quad (3.15)$$

Here, z_t is a binary variable indicating, for example, the activation status of a discrete component such as the electrolyser. The conditional bound (3.13) ensures that x_t can only take positive values if $z_t = 1$, effectively linking the continuous and binary decisions. This formulation allows the model to reflect realistic operational constraints while maintaining a tractable linear structure.

Nonlinear programming

MILP offers a strong balance between model expressiveness and computational efficiency, particularly when the problem structure remains linear. However, many real-world energy systems exhibit inherently nonlinear behaviour. Examples include quadratic cost functions in generation units, nonlinear efficiency curves in storage systems and thermodynamic relationships in hydrogen production processes [26, 62, 63]. These systems are better represented using Mixed-Integer Nonlinear Programming (MINLP), which permits nonlinear expressions in both the objective function and constraints.

While MINLP expands modeling flexibility, it also significantly increases computational complexity, especially in long-horizon or high-resolution problems. In addition, the availability and maturity of MINLP solvers is limited compared to those for MILP. Only a few commercial solvers, such as Gurobi, support specific convex MINLP classes, including quadratic or second-order cone constraints [33]. However, general nonconvex MINLP problems remain computationally challenging to solve. Due to these limitations, MINLP is frequently avoided in large-scale system-level studies. In many practical cases, effective linearisation techniques are applied to approximate nonlinear relationships within the MILP framework. These approximations enable the model to remain tractable while still capturing the essential system dynamics.

From deterministic to stochastic optimisation

The MILP formulation defined in equations (3.9)–(3.15) represents a deterministic case, where uncertainty in exogenous inputs is not considered. This setup is often referred to as a *perfect foresight* model, in which all external information, such as electricity prices and wind availability (denoted by d_t), is assumed to be known in advance for the entire modeling horizon. While this assumption simplifies the optimisation problem and enables the generation of high-resolution operational strategies that respect detailed system constraints, it represents a strong simplification of reality. In practice, future inputs are subject to uncertainty, and forecast errors can significantly affect system performance, particularly in flexible energy systems with market exposure [5].

A common way to address uncertainty is through parameter estimation and sensitivity analysis, which examine how optimal decisions respond to variations in key inputs. However, these techniques may be insufficient when the optimisation outcome is highly sensitive to uncertain parameters. In such cases, more robust methods are required that explicitly model uncertainty through random variables and integrate them directly into the optimisation framework. This leads to the class of stochastic optimisation methods, which are better suited for decision-making under uncertainty. The next subsection introduces several of these techniques.

3.1.3. Dynamic Programming

Dynamic Programming (DP) is a powerful method for solving multi-stage decision problems, where decisions taken at one point in time influence the future state of the system. It is well-suited for sequential decision-making under uncertainty and relies on the principle of optimality and Bellman's equation [72].

A widely used formalism within DP is the Markov Decision Process (MDP), which provides a structured approach for modeling systems where the future state depends on the current state, the action taken, and

stochastic exogenous information [37, 74]. MDP theory was originally introduced by Bellman [4] and is a cornerstone in dynamic optimisation. MDPs are commonly used in stochastic control problems and real-time operational strategies, where decisions must be updated frequently in response to new information [12, 79].

MDPs are defined over a finite decision horizon $t = 0, 1, \dots, T$, where the system evolves according to a stochastic transition function. The core elements of an MDP include a system state S_t , a decision or action A_t , and a stochastic process W_t , capturing exogenous uncertainty (e.g., electricity prices, wind availability). The system transitions from state S_t to S_{t+1} according to a known function, $S_{t+1} = S^M(S_t, A_t, W_t)$, and accumulates rewards based on a function $r(S_t, A_t)$. Powell [72] highlights that the transition function $S^M(S_t, A_t, W_t)$ plays a critical role in modeling the dynamics of the system. The random variable W_t introduces significant complexity, as it requires accurately characterising how external information impacts state transitions.

The objective is to identify a policy π , which maps each state S_t to an action $A_t^\pi(S_t)$, that maximises the expected cumulative reward over the horizon:

$$\max_{\pi} \mathbb{E} \left[\sum_{t=0}^T r(S_t, A_t^\pi(S_t)) \mid S_0 \right]. \quad (3.16)$$

This problem is solved recursively through the Bellman equation, which characterises the value of being in a particular state:

$$V_t(S_t) = \max_{a \in A_t} [r(S_t, a) + \gamma \mathbb{E}[V_{t+1}(S_{t+1}) \mid S_t, a]]. \quad (3.17)$$

Here, $V_t(S_t)$ denotes the value function at time t , $\gamma \in [0, 1]$ is the discount factor which adjusts the importance of future rewards, and $P(S_{t+1} \mid S_t, A_t)$ defines the transition probabilities. The transition probabilities encapsulate the stochastic nature of the system, describing the likelihood of moving from state S_t to S_{t+1} after taking action A_t . The goal is to find an optimal policy $\hat{\pi}$ that selects the action A_t in each state S_t to maximise the expected return over time.

MDPs have found various applications in energy systems for problems such as storage dispatch, market bidding, and hybrid renewable energy control [12, 79, 83]. A particularly relevant study is by Schrottenboer et al. [79], who apply an MDP framework to optimise a green hydrogen system with wind PPA's, hydrogen production and storage, market interaction, and hydrogen offtake contracts. Their model involves daily decisions on how much electricity to allocate to hydrogen production or to sell on the market, under stochastic prices and wind conditions, closely resembling the operational setting considered in this thesis.

To be able to solve the MDP, the authors discretise the entire decision and state space, including price levels, storage states, and production rates, and apply backward dynamic programming. This discretisation enables tractable computation but introduces certain limitations. First, discretisation introduces approximation error since continuous variables must be rounded to fixed levels, which leads to information loss and reduced modeling accuracy. Second, as Powell [71] notes, MDPs suffer from the well-known *curse of dimensionality*: the size of the state space, action space and stochastic information space each grow exponentially with problem complexity. In large-scale systems with finer temporal resolution and multiple interacting assets, such as the hourly, year-long hybrid energy system studied in this thesis, this exponential growth makes exact solution via classical dynamic programming computationally challenging.

To overcome these limitations and implement MDPs in practice, researchers have proposed approximation techniques such as Approximate Dynamic Programming (ADP) and reinforcement learning methods that relax the need for exact value function computation. For example, Shuai et al. [83] develop an ADP method for microgrid optimisation by solving single-period subproblems sequentially using Bellman's recursion. Similarly, Zhou et al. [96] apply deep reinforcement learning to dynamically dispatch a multi-energy system under uncertainty from renewable sources.

Despite these advances, the full MDP formulation remains computationally challenging for year-long optimisation of energy systems with hourly resolution and multiple interacting components. While they provide powerful theoretical structure and have appealing properties (e.g., optimal policies), the requirement to discretise the state and action space would lead to loss of granularity and significantly increased computational burden. Therefore, we looked into two alternative modeling frameworks that provide a balance between tractability and realism: Rolling Horizon Optimisation (RHO), also known as Model Predictive Control, and

Scenario-Based Stochastic Programming. These approaches allow sequential planning and incorporate uncertainty without requiring full discretisation of the solution space. These methods are further explained in the following sections.

3.1.4. Rolling Horizon Optimisation

Rolling Horizon Optimisation (RHO), also known as Model Predictive Control (MPC), is a sequential solution method for optimisation problems that evolve over time. Rather than solving the full decision problem across the entire time horizon at once, RHO breaks the problem into a series of smaller, overlapping segments and repeatedly solves this smaller subproblem.

Figure 3.1 provides a schematic representation of the RHO framework. At each solution step, an optimisation model is solved over a finite forward-looking window, known as the *prediction horizon*. Only the first period's decisions, known as the *control period*, are implemented. After each control step, the system state is updated with the most recent information, and the optimisation is repeated over a shifted time window. This iterative procedure continues until the full horizon has been covered.




Figure 3.1: Schematic of the Rolling Horizon Optimisation (RHO) framework, adapted from Ding et al. [50]. At each stage t , the optimisation problem is solved over a finite prediction horizon (orange bar), but only the first part of the solution, called the control period (blue bar), is implemented. After each stage, the system state is updated and the horizon shifts forward, forming a sequence of overlapping look-ahead windows. This iterative structure allows the model to incorporate updated forecasts and respond to changing system conditions in a computationally efficient manner.

The RHO approach is particularly attractive for large-scale energy systems where frequent re-optimisation is necessary due to evolving forecasts and operational states. Compared to classical dynamic programming, RHO does not require explicit knowledge of transition probabilities or full discretisation of the decision space. Instead, it uses deterministic optimisation formulations, often as a MILP, to re-plan based on updated forecasts, while maintaining computational efficiency [18].

While RHO enables adaptive decision-making over time, it does not directly model uncertainty beyond point forecasts or forecast scenarios. Because of its modular structure and ability to react to new information, RHO is frequently combined with stochastic optimisation methods [10, 18, 27]. For example, scenario-based stochastic programming can be embedded within each rolling optimisation step, allowing the system to account for future uncertainties using scenario generation techniques while still maintaining computational tractability. These stochastic methods are introduced in detail in the next section.

3.1.5. Stochastic Programming

In many real-world optimisation settings, model inputs are not known with certainty at the time decisions must be made. Stochastic programming extends linear programming by incorporating uncertainty, treating certain inputs as random variables [34]. Early developments in this field, such as the seminal work of Dantzig [15], introduced randomness through scenario-based formulations that extend deterministic models. In a scenario-based formulation, uncertainty is represented by a set of possible future outcomes, or scenarios. These scenarios are typically generated by introducing random forecast errors to baseline predictions for uncertain parameters, such as electricity prices and wind generation. For example, forecast errors may be sampled from a normal distribution and added to expected values to construct a range of plausible realisations [91]. The resulting scenario tree enables the optimisation model to evaluate decisions across multiple futures and identify strategies that are robust to variability in input conditions. This modeling approach is

particularly relevant in energy systems, where uncertainty in electricity prices, renewable production and demand forecasts plays a central role in operational planning [1, 18, 26, 62].

In addition to standard scenario generation methods, some studies have adopted a more advanced approach by modeling wind power and electricity prices as stochastic processes, directly capturing their dynamic nature over time. For example, Tankov [12] models wind power production using a truncated log-normal distribution derived from a stylised wind speed process, while electricity prices are represented by a stochastic differential equation with brownian motion. This approach allows for continuous-time modeling of both wind generation and price dynamics, enabling a more accurate representation of their evolution and volatility. Such methods are best suited for applications where variables evolve continuously, and the objective is to determine optimal control policies in real time, such as dynamic trading or continuous portfolio management. However, they are less compatible with scenario-based stochastic programming, which is inherently a discrete-time approach relying on a finite set of scenarios. Moreover, solving stochastic control problems often requires solving Hamilton-Jacobi-Bellman equations [12], which are computationally intensive, whereas stochastic programming typically involves solving MILP programs over a reduced set of scenarios, offering greater computational efficiency for large-scale problems.

Two-stage stochastic optimisation

A widely adopted structure within stochastic programming is the two-stage formulation, which separates decisions into those made before and after uncertainty is realised. This is particularly appropriate in energy system applications, where advance planning must be followed by real-time adjustments. Two-stage stochastic programming divides the decision-making process into a first “here-and-now” stage and a second “wait-and-see” stage. First-stage decisions are made prior to the realisation of uncertainty, typically based on forecasts or historical estimates. After the uncertainty is revealed, second-stage or recourse decisions are made to adjust the initial plan [5]. This structure is particularly relevant for problems such as day-ahead market bidding, electrolyser scheduling, or reserve allocation in power systems, where adjustments are made once real-time conditions deviate from forecasts [18, 28].

An example of a combined RHO and two-stage stochastic approach is demonstrated by Ding et al. [18], where they optimise a wind farm and storage system participating in the day-ahead and intraday markets, under uncertain wind forecasts. Their model adopts a two-stage stochastic programming framework for the day-ahead market, where the first stage determines the energy bids, and the second stage adjusts these decisions based on updated wind forecasts. The recourse actions include adjusting wind power generation, charging and discharging the storage, and managing penalties for deviations from the day-ahead commitments. Although they consider a rolling optimisation for the intraday market, the two-stage structure in the day-ahead market remains static, emphasising the impact of forecast uncertainty on bidding strategies.

Ghasemi [28] further expands on two-stage stochastic programming by applying it to an energy hub, which includes multiple energy carriers and interactions with the electricity grid. The model minimises the hub’s total cost by optimising energy exchanges, storage, and unit commitments while accounting for uncertainty in renewable generation (solar energy) and energy demand. Ghasemi’s work highlights the flexibility of two-stage stochastic optimisation in handling multi-carrier energy systems, where first-stage decisions (e.g., energy exchange with the grid) are made before the realisation of uncertainty, and second-stage adjustments are made once the uncertain parameters are known [28].

Using the notation from Ghasemi [28], let x denote the vector of first-stage decisions and ξ a random vector capturing uncertainty, such as forecast errors in electricity prices or wind production. Upon realisation of ξ , the second-stage decisions y are taken to minimise system costs under the realised scenario. The total expected cost includes both the first-stage cost and the expected cost of the recourse decisions. The problem can be expressed as:

$$\min_{x \in X} (f(x) + \mathbb{E}_{\xi} [Q(x, \xi)]), \quad (3.18)$$

where $Q(x, \xi)$ denotes the optimal value of the second-stage problem:

$$Q(x, \xi) = \min_{y \in Y(x, \xi)} g(y, x, \xi), \quad (3.19)$$

with $Y(x, \xi)$ representing the feasible set of second-stage decisions conditional on the first-stage plan and scenario ξ , and $g(y, x, \xi)$ denoting the second-stage cost function.

In the context of this thesis, the vector ξ includes forecast uncertainty in wind generation α and electricity prices λ , such that $\xi = (\lambda, \alpha)$. As computing the expectation over a continuous distribution is generally

intractable, a standard practice is to discretise the uncertainty space using a finite set of representative scenarios. These scenarios form the basis for a deterministic equivalent formulation of the stochastic program. Assuming that ξ has a finite set of realisations $\{\xi_1, \dots, \xi_S\}$ with associated probabilities ϕ_s , the expectation can be approximated as:

$$\mathbb{E}_\xi[Q(x, \xi)] \approx \sum_{s \in S} \phi_s Q(x, \xi_s), \quad (3.20)$$

yielding the following deterministic equivalent problem:

$$\min_{x \in X} \left(f(x) + \sum_{s \in S} \phi_s Q(x, \xi_s) \right). \quad (3.21)$$

The associated probabilities ϕ_s of each scenario must sum to one, determining the relative contribution of each scenario to the calculated expected value. While equal weighting implies that all scenarios are equally likely, more sophisticated approaches assign weights based on historical frequency or statistical similarity to a central forecast. For instance, inverse distance weighting can prioritise scenarios that are closer to the expected trajectory, providing a more accurate representation of likely outcomes.

Alternatively, robust optimisation (RO) can be used to find solutions that perform well across worst-case scenarios. Mansoor-Saatloo et al. [54] developed a robust technique for day-ahead scheduling of a hydrogen-based smart micro-energy hub integrated with demand response and fuel cell-based hydrogen storage system. RO was employed to incorporate power price uncertainty where the optimisation is performed under the worst-case scenario to achieve risk-averse outcomes.

3.1.6. Scenario generation methods for stochastic optimisation

In stochastic optimisation, the accuracy and effectiveness of the model depend heavily on the quality of the scenario set used to represent uncertainty. This motivates the use of structured scenario generation and reduction methods. In energy systems, where key inputs such as electricity prices and renewable output are inherently uncertain, carefully constructed scenarios are essential. These scenarios must capture a realistic range of possible outcomes to ensure that expected value calculations are meaningful and that the resulting optimisation decisions are robust and reliable.

Seljom et al. [80] provide a comparative evaluation of several scenario generation methods, including random sampling, iterative sampling, k-means clustering and moment-based optimisation. Their study highlights trade-offs between computational efficiency and statistical representativeness. More advanced techniques offer improved performance in capturing statistical properties and ensuring model stability but often at an increased computational cost. The choice of method must therefore balance accuracy with feasibility, especially in large-scale stochastic programs.

Monte Carlo simulation

One of the most widely used techniques for scenario generation in stochastic optimisation is Monte Carlo simulation (MCS) [72]. MCS generates a large number of random scenarios by sampling from predefined probability distributions for each uncertain variable, such as wind power, electricity prices or demand [31]. This approach is highly flexible, allowing for the modeling of a wide range of distribution types, including normal, log-normal and beta distributions. For instance, load demand uncertainty can be modeled using a normal distribution [1], while wind speed may follow a Weibull distribution due to its skewed nature [79].

While MCS is a widely used and flexible method for scenario generation, its effectiveness depends on the accuracy of the assumed probability distributions for forecast errors. MCS requires that these distributions are known and explicitly defined. This reliance on predefined distributions can be limiting, particularly when the true distribution is unknown or deviates from standard forms.

Empirical and model-based scenario generation

As a non-parametric alternative, bootstrapping can be used for scenario generation. Instead of sampling from a predefined distribution, bootstrapping resamples historical data directly, generating new scenarios that maintain the empirical characteristics of the observed data without assuming any specific statistical form [21, 35]. However, because it relies entirely on historical data, bootstrapping cannot generate novel patterns beyond what has been observed. This means it may fail to capture extreme events or unprecedented conditions, such as market shocks or rare weather events affecting renewable energy generation. Consequently,

while bootstrapping provides a robust alternative to MCS when distributional assumptions are uncertain, it is constrained by the scope of the available data.

Alternatively, rather than relying on forecasted values provided by external data sources, researchers can construct their own forecasting models. This is commonly achieved through time series forecasting methods, such as Auto-Regressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH). These models are particularly effective for directly forecasting electricity prices or wind power due to their ability to account for temporal dependencies in the data. ARIMA models can be suitable for both stationary and non-stationary time series, as the integration component allows for differencing to address non-stationarity. GARCH models, on the other hand, are particularly effective for modeling time-varying volatility, and are well suited to markets like electricity spot markets, where prices exhibit volatility clustering and heavy-tailed behaviour.

However, the use of time series forecasting methods is not without challenges. These models require careful calibration and extensive validation to ensure accuracy, as their performance is highly sensitive to model specifications and parameter settings. Moreover, they may struggle to account for structural breaks, such as sudden changes in market dynamics or unexpected shifts in renewable generation patterns [90]. Such complexities make time series forecasting a resource-intensive approach, often justifying the preference for simpler scenario generation techniques based on forecast errors.

Modeling forecast errors instead of direct values

A common approach in the literature is to generate scenarios based on the forecast errors of key variables, rather than the forecasted values themselves [1, 18, 62, 91]. This method enables the direct incorporation of forecast uncertainty without the need to build complex forecasting models. For instance, Xiao et al. [91] model the uncertainty in wind power outputs and electricity prices using normally distributed forecast errors. The forecast errors are assumed to follow a normal distribution with a mean of zero, reflecting unbiased forecasts, and a standard deviation estimated from historical data. This approach allows them to capture the inherent variability in the forecasts while maintaining their expected values.

Similarly, Aftab et al. [1] and Nease et al. [62] adopt the same methodology for demand forecasts, where the forecast errors are also assumed to follow a normal distribution. The use of normally distributed forecast errors is a practical choice, as it simplifies the scenario generation process and aligns with the Central Limit Theorem, which suggests that the sum of independent random variables tends to be normally distributed.

Scenario reduction techniques

While MCS can generate hundreds or thousands of scenarios, solving optimisation problems over such large sets is computationally demanding. Scenario reduction techniques aim to approximate the full scenario set using a smaller, representative subset, while preserving key statistical properties such as means, variances and correlations.

Methods include forward selection, backward elimination, k-means clustering and Kantorovich distance minimization, which seeks to minimize the Wasserstein distance between the original and reduced distributions [1, 26]. Although scenario reduction improves computational feasibility, it may reduce model robustness, especially in capturing extreme events or rare but high-impact conditions.

3.1.7. Multivariate Monte Carlo and correlated scenario generation

In its basic form, MCS assumes independent sampling from marginal distributions of uncertain parameters. However, this independence assumption is often unrealistic in energy systems, where temporal and inter-variable correlations are significant, such as the strong link between renewable generation availability and electricity market prices.

To preserve these dependencies, *multivariate* MCS using Cholesky decomposition is commonly applied in literature [30]. This method allows for the transformation of uncorrelated random samples into samples that follow a desired correlation structure. Given a correlation matrix R , Cholesky decomposition produces a lower triangular matrix L such that $R = LL^T$. Uncorrelated samples Z are transformed into correlated ones via $Y = LZ$, which are then used to add noise to the forecast trajectories. This ensures that the simulated scenarios maintain the empirical correlations observed in historical data.

To illustrate this method, consider a case where the goal is to generate scenarios for wind forecast errors $w^{(s)}$ and electricity price errors $z^{(s)}$, assuming these are normally distributed and exhibit an empirical correlation coefficient ρ . The procedure begins by sampling:

$$\mathbf{z}^{(s)} = \begin{bmatrix} w^{(s)} \\ z^{(s)} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, I).$$

To induce the correlation, we define the correlation matrix:

$$R = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \quad \text{with} \quad L = \begin{bmatrix} 1 & 0 \\ \rho & \sqrt{1-\rho^2} \end{bmatrix},$$

where the lower triangular matrix L is derived by the Cholesky decomposition such that $R = LL^\top$. The transformation $\mathbf{y}^{(s)} = L\mathbf{z}^{(s)}$ yields correlated samples:

$$\mathbf{y}^{(s)} = \begin{bmatrix} w^{(s)} \\ \rho w^{(s)} + \sqrt{1-\rho^2} z^{(s)} \end{bmatrix}.$$

This guarantees that the generated scenarios preserve the target linear correlation structure. The resulting samples can then be scaled by historical standard deviations and added to base forecasts to construct realistic scenario trajectories.

Beyond inter-variable correlation, Cholesky decomposition can also be extended to account for temporal dependencies by constructing a full covariance matrix across time. This involves estimating autocorrelation coefficients and embedding them into block covariance matrices. As a result, the method enables the generation of time series scenarios that capture both intra-variable (temporal) and inter-variable (cross-dependence) correlation structures.

While Cholesky-based MCS is computationally efficient and analytically tractable, it assumes multivariate normality and linear correlation. These assumptions may oversimplify real-world dependencies, especially under regime shifts or extreme events. As an alternative, copula-based methods have been developed to allow for more flexible dependency modeling. Copulas decompose joint distributions into marginal components and a dependence structure, enabling non-Gaussian, non-linear correlations to be modeled independently of the marginals [45]. Gaussian and Archimedean copulas are widely used in energy systems to model joint behavior between renewable generation, load and market prices [69]. Alternatively, empirical copula approaches can be used to reproduce observed rank correlations without assuming a specific functional form [77]. These methods offer a richer but often more computationally intensive alternative to Cholesky-based MCS.

3.2. Classifying literature on combined energy systems management

Optimal energy management is an essential topic in modern energy systems research, driven by the need for economic viability, efficient integration of renewable energy sources, and the challenges posed by energy market dynamics. To systematically explore and categorise optimisation models for energy management, Weitzel and Glock [89] developed a conceptual framework tailored for electrical energy storage systems (EES). This framework has proven valuable in structuring research on energy management by providing a clear and consistent method for classifying existing studies. While their framework was initially developed for stationary electrical storage, its adaptability allows for its application to hybrid systems including energy storage systems, hydrogen production, and sales, as explored in this research.

The framework is built on three primary components that guide the classification of optimisation models: *contextual characteristics*, *mathematical formulation*, and *optimisation techniques*. Each component addresses different aspects of energy management problems and provides a systematic way to evaluate and position existing research. An overview of the framework by Weitzel and Glock [89] is presented in Figure 3.2.

Contextual characteristics

The first component of the framework investigates the context of the energy management problem and answers the question, “*What is considered?*” It encompasses the system scope, system characteristics, time horizon, and control architecture.

The system scope defines the type of participants and their interactions. This can range from simple configurations involving only storage, to more complex hybrid systems or microgrids that include producers, consumers and multiple forms of energy storage. Depending on the configuration, the overarching model objectives may be economic (e.g., profit maximisation), technical (e.g., grid stability or base-load production), environmental (e.g., carbon reduction), or user-driven (e.g., self-sufficiency). The system characteris-

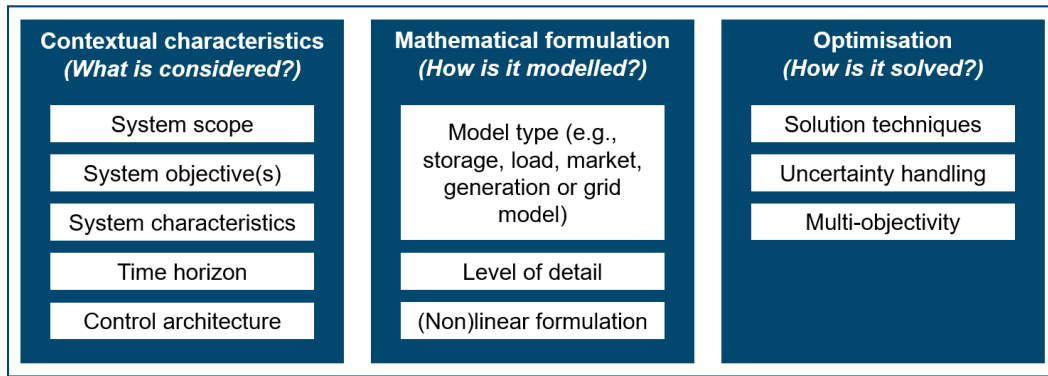


Figure 3.2: Conceptual framework for classifying literature for optimisation of energy management systems, based on the framework developed by Weitzel and Flock [89]. The framework categorises studies based on three key dimensions: contextual characteristics, mathematical formulation, and optimisation techniques. This structured approach enables consistent comparison and positioning of optimisation approaches and is adapted in this thesis to evaluate hybrid electricity-hydrogen systems.

tics detail the specific components included in the system, such as the type of storage (e.g., battery or hydrogen), production (e.g., electricity from wind or solar, hydrogen from electrolysis) and market interaction (e.g., participation in the day-ahead market). These characteristics shape both the operational behavior and the optimisation complexity of the system.

The time horizon defines the temporal scope of the analysis. Depending on the objectives, energy systems can be evaluated over short-term periods (hours to days) or extended durations (weeks to years). Using a longer time frame enables a more accurate representation of seasonal dynamics, long-term pricing trends, and component degradation. The control architecture specifies how decision-making is organised across system components. In a centralised architecture, a single decision-maker oversees the entire system. A decentralised approach involves multiple autonomous agents making independent decisions, while a distributed setup enables collaborative control among components.

An important consideration related to both the time horizon and control architecture is the intended purpose of the model. Specifically, whether it is designed for real-time control or long-term profitability evaluation. A model aimed at real-time control focuses on dynamic decision-making, continuously adapting to changing conditions such as electricity prices, renewable energy availability, or demand fluctuations. These models often require fast computation and efficient handling of operational constraints, making them suitable for automated dispatch and short-term planning. Conversely, models designed for long-term profitability assessment emphasise strategic decision-making over extended periods. These models account for economic metrics, investment costs, operational revenues, and risk management. They are used to evaluate the financial viability of energy systems, optimise asset sizing, or design market participation strategies. The distinction between these two model types is crucial, as it directly influences the level of detail in the mathematical formulation, the optimisation techniques applied, and the computational requirements.

Mathematical formulation

The second pillar of the framework addresses the question, “*How is the problem modeled?*” It covers the mathematical representation of the system and includes five potential modeling components: storage, load, market, generation and grid model.

The storage model forms a fundamental part of most energy system studies, capturing the physical and operational behaviour of storage technologies. Key parameters include charging and discharging efficiencies, storage capacity, energy losses and state-of-charge dynamics. The load model represents the system’s electricity demand, which can be modeled as either deterministic, using for example fixed demand profiles, or stochastic to account for uncertainty in consumption patterns. The level of demand detail affects the realism and adaptability of the dispatch strategy. The market model describes how the system interacts with electricity markets, including day-ahead, real-time, or ancillary service markets. This model may incorporate pricing schemes, bidding strategies, and regulatory constraints, though not all studies model these in full detail. The generation model is particularly important when renewable generation sources such as wind or solar are part of the system. It defines the production characteristics, availability profiles, and intermittency of these resources. Lastly, the grid model, though often abstracted in system-level studies, represents network

constraints, power flows and grid-related costs.

In addition to the selection of components, the level of detail and abstraction in each model varies significantly depending on the study objective. For instance, process-level models may include thermodynamic representations of hydrogen production via electrolysis, accounting for temperature, pressure, and degradation [1, 3, 62]. However, such complexity is often avoided in system-level optimisation where the focus lies on high-level dispatch strategies and energy flows [2, 43, 63]. Similarly, while detailed market rules and participation structures can be included, many models simplify them by representing electricity prices as a time series and assuming simplified bidding behaviour

The chosen level of detail has direct implications for the type of mathematical formulation employed. Simplified and linear representations often allow the use of MILP, which is computationally tractable and widely applied in energy systems optimisation [89]. When models include nonlinearities, such as those introduced by component degradation or nonconstant conversion efficiencies, more advanced formulations like MINLP may be required [83]. In highly stochastic environments, MDPs are often used to capture sequential uncertainty and adaptive decision-making [12, 54, 79]. The choice of formulation depends not only on what is being modeled but also on how the system and its dynamics are represented

Optimisation technique

The third component of the classification framework answers the question, “*How is it solved?*” It refers to the computational methods used to find optimal or near-optimal solutions to the formulated models, the way uncertainty is accounted for in the decision-making process and the approach taken when multiple conflicting objectives are present.

Solution techniques vary widely depending on the problem formulation, level of complexity and computational requirements. LP and MILP are among the most common techniques in the literature due to their tractability and compatibility with existing solvers such as CPLEX and Gurobi [2, 3, 18, 27, 28]. These solvers are specialised software tools that implement advanced mathematical algorithms to efficiently compute optimal solutions for structured optimisation problems. When models include nonlinear relationships, nonlinear programming may be applied, although this typically incurs significantly higher computational costs. These approaches aim to find exact solutions, meaning that they identify a mathematically optimal result within a defined solution space. In addition to exact solution methods, in some cases, heuristic or metaheuristic methods are employed to solve complex or non-convex problems, particularly when a global optimum is not strictly required or feasible to compute.

When the energy management problem involves uncertain inputs such as renewable generation, electricity prices or demand levels, different solution techniques can be applied to handle this uncertainty. A common approach is to ignore uncertainty entirely and use deterministic models with perfect foresight, assuming full knowledge of all inputs over the planning horizon. While computationally efficient, this assumption may lead to unrealistic or overly optimistic results. To address this limitation, stochastic optimisation methods are used. Uncertainty can be handled through scenario-based stochastic programming or dynamic programming for problems formulated as MDPs. Stochastic programs can often still be formulated as MILPs and solved using exact methods, though the inclusion of scenarios increases model size and complexity. For MDPs, or when problem complexity becomes intractable, approximate methods such as reinforcement learning or heuristics are required [10]. A detailed description of these solution techniques and their mathematical foundations is provided in the previous Section 3.1.

Lastly, in many energy system applications, multiple objectives must be considered. For example, an operator may aim to minimize cost while also reducing emissions or maintaining high reliability. These trade-offs are typically handled using multi-objective optimisation techniques, such as Pareto front analysis, weighted sum methods or ϵ -constraint approaches [55]. The chosen technique determines how trade-offs are navigated and how final decisions are selected from competing alternatives.

Table 3.1: Overview of optimisation studies in integrated energy and hydrogen systems analysed for this thesis. Each study is analysed considering the following categories: the type of renewable energy source (RES), the presence or absence of grid connection, inclusion of hydrogen production and sales, the use of fuel cells, the electricity markets involved, and any storage components. Moreover, each optimisation model is characterised by its objective function, modeling horizon, mathematical formulation, solution method, and whether uncertainty was incorporated. For notation purposes, the following abbreviations are used: Wind Turbine (WT), Photovoltaic Solar Energy (PV), Hydrogen Storage (H₂), Battery Energy Storage System (BESS), Fuel Cell (FC), Combined Heat and Power unit (CHP), Compressed Air Energy Storage (CAES), Pumped Hydro Storage (PHS), Day-Ahead Market (DA), Intraday Market (ID), Reserves Market (RM), Mixed Integer Linear Programming (MILP), Mixed Integer Non-Linear Programming (MINLP), Conditional Value at Risk (CVaR), Load Following (LF), Linear Programming (LP), Dynamic Programming (DP), Rolling Horizon (RH), Robust Optimisation (RO), Stochastic Optimisation (SO), Scenario-Based (SB), Two-Stage (2S).

Ref	RES (Type)	Grid Conn.	H ₂ Prod.	H ₂ Sale	Power Markets	Storage/ Assets	Math Form.	Obj. Function	Modeling Horizon	Solution Method	Uncertainty
[79]	✓ (WT)	✓	✓	✓	DA	H ₂ , FC	MDP	Max profit	1 year	DP	✓ (prices, wind)
[12]	✓ (WT)	✓	-	-	DA, ID	BESS	MDP	Max profit	1 week	DP	✓ (prices, wind)
[2]	✓ (PV)	✓	✓	✓	DA, RM	H ₂ , BESS	MILP	Min cost	1 week	LP	-
[54]	✓ (WT)	✓	✓	✓	-	H ₂ , BESS, FC, CHP	MILP	Min cost	24 hrs	RO	✓ (prices)
[91]	✓ (WT)	✓	✓	✓	DA	H ₂	MILP	Max profit + CVaR	24 hrs	LP, SB-SO	✓ (prices, wind)
[18]	✓ (WT)	✓	-	-	DA, ID	PHS	MILP	Max profit	24 hrs	LP, RH, 2S-SO	✓ (wind)
[1]	✓ (PV)	✓	✓	-	-	H ₂ , BESS	MINLP	Min cost	24 hrs	NLP, SB-SO	✓ (PV, load)
[26]	-	✓	✓	-	DA	H ₂ , BESS, FC	MINLP	Min cost	24 hrs	NLP, SB-SO	✓ (demand)
[63]	✓ (WT, PV)	-	✓	-	-	H ₂ , BESS	MINLP	Min LCOH	1 year	NLP	-
[43]	✓ (PV)	✓	✓	✓	DA	H ₂ , FC	MILP	Max profit + CVaR	24 hrs	LP, 2S-SO	✓ (prices, PV, demand)
[62]	-	✓	-	-	DA	FC, CAES	MILNP	Max profit + LF	1 week	NLP, RH	✓ (demand)
[28]	✓ (PV)	✓	-	-	DA	BESS, CHP	MILP	Min cost	24 hrs	LP, 2S-SO	✓ (PV, load)

3.3. Key related works on optimising integrated energy systems

This section highlights the most relevant studies that align with the objectives of this thesis. As the previous section has outlined the key mathematical techniques used for optimisation under uncertainty, the focus here is on how these techniques are applied in the context of integrated energy systems. Particular attention is given to studies that explore hydrogen production, energy storage and interaction with energy markets.

Table 3.1 provides an overview of the literature analysed. Each study is selected based on its relevance to the scope of this thesis, including the role of hydrogen as a central system component, the integration of renewable energy sources, electricity and hydrogen market participation and the use of storage technologies. In addition, studies are characterised by their modeling objectives, mathematical formulation, solution approach, and the treatment of uncertainty.

3.3.1. Multi-energy configurations using linear programming

Recent studies have explored optimisation-based approaches for the operation of integrated energy systems that combine electricity and hydrogen production. One such example is the work by Alvarez et al. [2], who develop a deterministic optimisation model for the self-scheduling of a hydrogen-based virtual power plant (VPP) participating in both the day-ahead electricity and reserve markets. Their system comprises solar PV, battery storage, electrolysers, and hydrogen storage, forming a comprehensive multi-energy configuration. The model is formulated as a MILP problem, with the objective of minimising total system costs by optimally dispatching the various system components. The optimisation is performed over a one-week horizon with hourly resolution, and focuses on computing the optimal 24-hour day-ahead schedule for each trading day. It is solved using the solver Gurobi, achieving efficient computation times. However, the model does not incorporate uncertainty in renewable generation or market prices, which limits its ability to capture the variability inherent in such systems. The authors acknowledge this limitation and suggest that a rolling horizon approach could improve the model's responsiveness to dynamic operating conditions.

In contrast, Xiao et al. [91] present a stochastic optimisation framework for the operation of a wind-powered hydrogen storage system with access to both electricity and hydrogen markets. The system includes a wind farm, an electrolyser, and hydrogen storage, and seeks to maximise expected revenue by determining the optimal allocation of wind power between electricity sales and hydrogen production. Like Alvarez et al. [2], the authors use a MILP formulation and focuses on short-term operation, with their modeling horizon limited to 24 hours. However, their model extends the deterministic framework by incorporating uncertainty through a scenario-based stochastic approach. Random forecast errors in wind generation and electricity prices are modeled using a normal distribution with parameters estimated from historical data. Additionally, financial risk is addressed using Conditional Value-at-Risk (CVaR), which allows the model to weigh expected returns against potential losses under unfavourable conditions. This study demonstrates how hydrogen production can help mitigate wind power curtailment, enhancing the overall business case for wind farm operators by providing an additional revenue stream through hydrogen sales.

Integration of multiple energy components for enhanced profitability

A key conclusion of both the study by Xiao et al. [91] and by Alvarez et al. [2] is that the integration of multiple components, such as renewable energy sources, battery storage, and hydrogen storage, significantly enhances the profitability of the combined hydrogen and electricity systems. Alvarez et al. [2] explains how the battery storage provides rapid response to demand fluctuations, while the hydrogen storage tank offers long-term storage, ensuring supply during periods of low renewable generation. Without these components, the hydrogen VPP's ability to adapt to market conditions and seasonal demand is severely limited. This finding underscores the importance of combined energy systems, and the need to model them together. Despite Xiao et al. [91] and Alvarez et al. [2] including many elements that we also consider in this thesis, their studies lack consideration of long-term contracts, such as power purchase agreements (PPAs) or hydrogen offtake agreements.

Long-term optimisation and capacity sizing

One of the few studies that adopts a long-term perspective is presented by Niaz et al. [63], who investigate the optimal sizing of a battery system in a renewable-powered hydrogen production setup. Their model aims to minimise the levelised cost of hydrogen (LCOH) by optimising the BESS capacity in a standalone hybrid system composed of solar PV, wind, an electrolyser and battery storage. The system is simulated over a full year with hourly resolution, capturing seasonal variation in renewable availability. Unlike many studies

focused on short-term operation, this work does not consider uncertainty in renewable generation or system operation, nor does it model market participation. While this makes the approach less applicable for uncertainty-aware planning, it reflects a different modeling focus aimed at design and sizing decisions under deterministic conditions. Furthermore, although the model is described as a mixed-integer dynamic optimisation problem, the solution method is not specified. This lack of methodological transparency presents challenges for reproducibility, particularly due to the structure of the objective function. When LCOH is used as the objective, total system costs are divided by the amount of hydrogen produced. However, this quantity is not fixed as it depends on decision variables such as the dispatch profile and battery capacity. As a result, the objective involves dividing by a function that is itself endogenous to the optimisation, introducing non-linearity. Moreover, battery capacity directly constrains the feasible space of hourly dispatch decisions, meaning that sizing and operational strategies must be optimised jointly. This interdependence significantly increases the complexity of the problem and makes solving it computationally demanding.

Still, the study offers valuable insights into cost-driven system design. The results show a strong sensitivity of battery size to capital cost assumptions, with the lowest LCOH of €10.33/kg achieved using a 6575 kW battery paired with a 4.5 MW alkaline electrolyser. While uncertainty and market-driven dynamics are not included, this work is notable for addressing investment decision-making over an extended horizon, which is a perspective still underrepresented in the current literature.

3.3.2. Applications of rolling horizon and two-stage stochastic programming

Building on the narrative by Xiao et al. [91] of enhanced profitability for wind farms through storage and market interaction, Ding et al. [18] explore a similar concept, focusing on the coordinated operation of a wind farm and a pumped hydro storage (PHS) plant active on both the day-ahead (DA) and intraday (ID) electricity market. Their primary objective is to maximise profit by optimising bidding and dispatch strategies while accounting for the stochastic nature of wind power generation. It is a short-term operational model, with a modeling horizon of 24 hours and high temporal resolution (hourly for the DA market and 10-minute intervals for ID operations). The model is formulated as a two-stage stochastic optimisation problem combined with rolling horizon approach. In the first stage, the system determines day-ahead bids, while in the second stage, these bids are adjusted in real-time as updated wind forecasts become available. Uncertainty in wind generation is modeled using a scenario-based approach, with scenarios generated through Monte Carlo simulation and reduced using a backward reduction method for computational efficiency. The solution employs a rolling horizon strategy, where the optimisation is periodically re-optimised with the latest forecast data. This approach demonstrates how rolling horizon optimisation can effectively manage forecast uncertainty in renewable energy systems. While Ding et al. [18] focus on intraday adjustments in response to updated forecasts, our setup does not explicitly model intraday markets.

Moreover, applications of two-stage frameworks specifically for hydrogen production systems are relatively scarce. One notable exception is Rezaee et al. [43] who present a two-stage stochastic programming approach for optimal electricity procurement of hydrogen fuel stations. Their model integrates an electrolyser, hydrogen storage and participation in multiple electricity procurement channels, including the DA market, balancing market and bilateral contracts. The first stage involves scenario-independent decisions, such as the volume of electricity procured through bilateral contracts, while the second stage captures scenario-dependent decisions, including transactions in the DA and balancing markets. This setup effectively captures the trade-off between securing electricity through fixed contracts and responding to real-time price fluctuations. However, Rezaee et al. [43] focus on minimising electricity procurement costs for hydrogen fuel stations rather than optimising a broader hydrogen production system with renewable energy sources. Their reliance on the balancing market also introduces elements that are outside the scope of this research. Furthermore, their model is a static two-stage optimisation over a fixed period, without the continuous re-optimisation seen in rolling horizon approaches like that of Ding et al. [18].

3.3.3. Application of Markov decision process formulation

Schrotenboer et al. [79] explore an advanced control strategy for a Green Hydrogen Energy System (GHES), which integrates renewable energy from wind, hydrogen production via electrolysis, hydrogen storage, and fuel cell technology for converting hydrogen back to electricity. The model adopts a MDP framework to optimise the long-term profitability of the system, balancing hydrogen production, storage and sales. The model accounts for uncertainty in wind generation (modeled as a Weibull distribution) and electricity prices (modeled as an AR(1) process), as well as hydrogen spot market prices following a similar AR(1) structure. It explores different strategies for selling hydrogen, including fixed contracts and flexible market-based sales. The

GHES operates under both spot market interactions and long-term contracts, including PPAs and hydrogen offtake agreements. This structure allows the model to capture the complex interplay between market interactions and contractual obligations.

This study by Schrottenboer et al. [79] employs a case study in the Netherlands using 2019 electricity price data, demonstrating that optimal management of hydrogen production and storage can significantly enhance system profitability. The results suggest that operational revenues could increase by up to 51% through effective contract and market strategies. However, the findings rely on 2019 price data and assume a functioning hydrogen spot market, highlighting the need for validation with more recent data and consideration of current market limitations.

The model was applied over a full-year time horizon with a daily resolution. While this long-term perspective enables strategic planning and the evaluation of contract-based decisions, the coarse temporal resolution limits the model's ability to capture the short-term dynamics of variable renewable generation and electricity prices. This trade-off is likely necessary to ensure computational tractability within the MDP framework, which becomes increasingly complex as the temporal granularity and number of state variables increase. The MDP approach remains computationally intensive due to the curse of dimensionality inherent in dynamic programming, requiring substantial computing resources. In contrast, MILP-based approaches, such as those proposed by Alvarez et al. [2], Xiao et al. [91] and Rezaee et al. [43] provide a more tractable and scalable alternative.

Oversizing renewable energy contracts

A critical factor influencing hydrogen production costs within the study by Schrottenboer et al. [79] is the amount of renewable power procured through long-term PPA contracts, raising the question of how to best size the PPA capacity relative to the electrolyser capacity. Schrottenboer et al. [79] address this by conducting a sensitivity analysis on different PPA sizes, demonstrating how variations in PPA capacity impact the system's profitability and operational efficiency.

In a different study Moran et al. [59] advocate for oversizing PPA capacity relative to the electrolyser size, allowing surplus renewable power to be sold on the spot market, while accepting some degree of curtailment. This approach effectively increases the electrolyser's full-load hours, improving its economic performance. Similarly, Ali Khan et al. [44] argue that even in an on-site PPA configuration, oversizing can be advantageous, as it boosts full-load hours and enhances the cost-efficiency of electrolyser operation. However, excessive oversizing can lead to significant curtailment costs or force the sale of surplus electricity at a loss.

3.4. Classifying this research and identifying the literature gap

The studies reviewed in the previous section consistently highlight the value of integrated energy systems, where the combination of renewable energy sources, grid connection, energy storage, and hydrogen production can significantly improve system profitability. These findings underscore the importance of modeling and optimising the system as a whole, rather than analysing components in isolation. Several studies also point to the benefits of oversizing renewable generation capacity, allowing for increased self-sufficiency and better utilisation of surplus energy through storage or hydrogen production. Moreover, combining short-term market participation with long-term contractual arrangements has been shown to enhance both economic performance and system robustness.

In terms of optimisation methodology, the literature review revealed a wide range of approaches. On one end of the spectrum are deterministic models that rely on perfect foresight of electricity prices and renewable availability. Fully deterministic models, such as that of Alvarez et al. [2], offer the fastest solving times and lowest computational burden, making them suitable for rapid evaluations or applications with limited computational resources. However, this efficiency comes at the expense of realism, as these models do not capture the uncertainty inherent in renewable generation or electricity market dynamics. On the other end are stochastic models that incorporate uncertainty through scenarios like those of Xiao et al. [91] and Rezaee et al. [43], though at the cost of increased computational complexity. These contrasting methods illustrate a fundamental trade-off between computational efficiency and modeling realism in the optimisation of integrated energy systems. Motivated by these insights, this thesis explores and compares optimisation approaches with varying levels of look-ahead capability and uncertainty handling, aiming to identify effective strategies for the design and operation of hybrid systems that include wind energy, battery storage, and hydrogen production under realistic market and weather conditions.

Following the classification scheme by Weitzel and Glock [89] presented in Section 3.2, this research is classified in the following way.

Contextual characteristics

In this thesis, we consider a hybrid energy system that has hydrogen production via electrolysis at the heart of the system, integrated with a battery for energy storage. For electricity procurement we have a bilateral PPA with a wind park and additionally interact with the electricity grid through the DA market. The key components therefore include wind power generation, an electrolyser, a battery system, and market interaction for electricity trading.

The primary objective is to maximise the system's profit over a full-year horizon (8760 hours) with hourly resolution, capturing both short-term operational dynamics and long-term seasonal trends. Unlike real-time dispatch models, which focus on immediate operational decisions, this thesis aims to evaluate the system's long-term profitability. The focus is on strategic decision-making that balances electricity procurement, storage operation, and hydrogen production to maximise economic returns. A centralised control approach is employed, where all components are managed under a unified optimisation framework to ensure coordinated operation.

Mathematical formulation

The optimisation model is formulated as a simplified, system-level representation of the hybrid energy system. The objective is to determine the optimal dispatch strategy while satisfying all technical and contractual constraints. Specifically, the model determines the volumes of electricity to buy or sell on the day-ahead market, the battery's charging and discharging schedule, and the operation of the electrolyser, including its on-off status.

The optimisation problem is formulated as a MILP, ensuring that it can be efficiently solved while maintaining the flexibility to incorporate binary decision variables for operational constraints. This model adopts several simplifications to maintain computational tractability. It avoids detailed process models of hydrogen production or battery degradation, instead assuming constant efficiencies for the electrolyser and battery. The market model is also simplified, representing electricity prices as exogenous information drawn from historical data or forecasted values. The system is assumed to be a price-taker, meaning that its bids on the day-ahead market do not influence market prices. Similarly, wind power generation is treated as an exogenous input, determined by wind load factor data, without accounting for operational interactions with the wind farm. Hydrogen demand is represented as a fixed annual target that must be met, but the production can be flexibly scheduled throughout the year to maximise profitability. This approach allows the model to focus on optimising the balance between electricity procurement, battery storage, and hydrogen production, rather than modeling the entire supply-demand dynamics of the electricity market.

Optimisation

The base case of this model is a fully deterministic MILP with perfect foresight. This deterministic model outcome serves as a best-case benchmark, establishing an upper bound on system profitability. Leveraging state-of-the-art commercial solvers Gurobi, ensures fast and reliable computation, making this approach practical even for the full-year, hourly-resolution problem.

However, to bridge the gap between theoretical best-case scenarios and real-world operations, this thesis first extends the deterministic MILP with a rolling horizon framework and reduces the look-ahead window from one year to one week. In rolling horizon optimisation (RHO), the optimisation problem is solved over a moving, finite planning horizon, while looking ahead in the forecast window. Specifically, at each 24-hour planning step, the model is re-optimised using the latest available information for electricity prices and wind availability. The immediate decisions for the first 24-hours are implemented, and the model is then re-solved in the next solution step as new data becomes available. This adaptive approach allows the system to respond to changing conditions while maintaining computational tractability. Nevertheless, RHO by itself is still deterministic, relying on point forecasts without explicitly accounting for forecast uncertainty.

To introduce uncertainty directly into the model, this thesis further develops a Scenario-Based Stochastic Optimisation approach within the RHO framework. In this extended model, a two-stage stochastic programming structure is applied. At each solution step, we assume perfect foresight for the 24-hour planning period (when actual day-ahead prices are known), while scenarios are generated for the remaining 144 hours of the 168-hour (one week) look-ahead window. These scenarios incorporate uncertainty in both wind generation and electricity prices and are used to optimise expected profit over the planning horizon. Specifically, the first-stage decisions are the dispatch decisions for the first 24-hour planning period and the second-stage decisions are the dispatch decision for the rest of the week, which are scenario dependent.

To ensure that the generated scenarios accurately capture the statistical properties of wind and electricity price uncertainties, this thesis employs a Monte Carlo Simulation approach. Specifically, scenarios are generated for the forecasting errors of the wind load factor and day-ahead electricity prices using historical data as a basis. To preserve the negative correlation typically observed between wind generation and electricity prices, a Cholesky decomposition is applied, transforming independently generated scenarios into correlated ones. This multivariate Monte Carlo method ensures that the scenarios reflect realistic joint behavior of the two variables.

Lastly, while the optimisation model is primarily single-objective, it is inherently flexible. In future work, this model could be extended to include multiple objectives, such as minimizing emissions, reducing grid interactions, or enhancing system reliability. For the scope of this thesis, however, the focus remains on profit maximisation, providing a clear and quantifiable performance metric.

3.4.1. Literature gap

To the best of our knowledge, several important gaps remain in the existing literature. This thesis aims to address several of these gaps through an innovative system set up and solution approach and contributes to the existing literature in the following ways:

- **Unique system configuration with hydrogen production as central objective:** Existing studies rarely focus on a hybrid energy system where hydrogen production and sales are the primary objective. Our model uniquely combines a wind PPA, a battery for energy storage and active participation in the day-ahead electricity market. The coordinated operation of these components is aimed at improving the hydrogen business case, and their individual contributions are evaluated to understand their respective impacts on overall system performance.
- **Long-term profitability modeling under uncertainty:** In contrast to most existing studies that focus on short-term dispatch (e.g., daily or weekly), this evaluates system performance over a full-year horizon (8,760 hours) with hourly resolution, providing a comprehensive view of long-term profitability. It incorporates uncertainty in wind generation and electricity prices, linking short-term market strategies to long-term profitability outcomes through the use of stochastic optimisation techniques.
- **Innovative application of rolling horizon and two-stage stochastic optimisation:** The thesis proposes a unique solution approach that integrates two-stage stochastic programming with a rolling horizon framework. While rolling horizon methods and stochastic programming are typically applied in short-term operational contexts, this work extends their application to long-term planning for hydrogen production systems under uncertainty.
- **Correlation-aware scenario generation:** A key innovation in this thesis is the use of a scenario generation method that accurately captures the observed negative correlation between wind generation and electricity prices. This approach ensures that the generated scenarios reflect the statistical dependencies of these variables, reducing the risk of generating unrealistic scenarios and improving the robustness of the optimisation results.

4

Research Approach

This chapter outlines the methodological framework and modelling choices that form the foundation of the analysis. It provides a structured overview of how the hydrogen production system is represented, how key input parameters are selected and how performance is evaluated across different system configurations. The chapter begins with a high-level description of the base case setup in Section 4.1, introducing the core system components and their interactions. Section 4.2 then presents the specific test cases designed to answer the research questions, explaining the rationale behind different configurations. The modeling structure and workflow are described in Section 4.8, outlining the logic used to optimise system operation and defining model output. This is followed by Section 4.4, which defines the key performance indicators (KPIs) used to assess the economic and operational outcomes of each test case. To clarify the scope of the analysis, Section 4.5 summarises the main assumptions, simplifications and modeling limitations. Afterwards, the external data sources for electricity prices and wind generation are discussed in Section 4.6, while Section 4.7 provides a complete overview of the technical and economic parameters used throughout the model. Lastly, Section 4.8 outlines the model implementation, including the programming tools, solver configuration and computational environment used to develop and run the optimisation models.

4.1. System description

This thesis focuses on the optimisation of a hybrid renewable energy system located in the Netherlands that integrates wind power through a Power Purchase Agreement (PPA), a battery energy storage system (BESS), an electrolyser for hydrogen production and market participation through the day-ahead electricity market. The goal is to maximise system profitability by jointly optimising energy dispatch and market interaction decisions, while respecting technical and contractual constraints.

Figure 4.1 provides a schematic overview of the system architecture. Electricity can be sourced from two channels: (1) a fixed-price wind PPA and (2) the day-ahead electricity market. In addition, the system has the ability to sell electricity back to the market. All power flows in and out of the system occur through a single bidirectional grid connection.

The wind PPA is modeled as a *take-or-pay* contract: at each hour, the system operator pays a fixed price per MWh for the full amount of wind power available, regardless of whether or not it is utilised. The availability of wind power is determined by hourly wind load factors, which are treated as exogenous inputs to the model. In cases where the available PPA power is curtailed, the system incurs balancing penalties, which reflect redispatching and imbalance costs borne by the grid operator, as observed in practice [84].

The system also participates in the day-ahead electricity market, where hourly electricity prices are known with certainty one day in advance and serve as model inputs. The system may submit bids to buy electricity (e.g., for storage or hydrogen production) or to sell surplus electricity back to the grid. We assume the system behaves as a *price taker*, meaning its operations do not influence market clearing prices. This is justified by the relatively small size of the electrolyser and battery relative to the scale of the competitive and liberalised national electricity market [39].

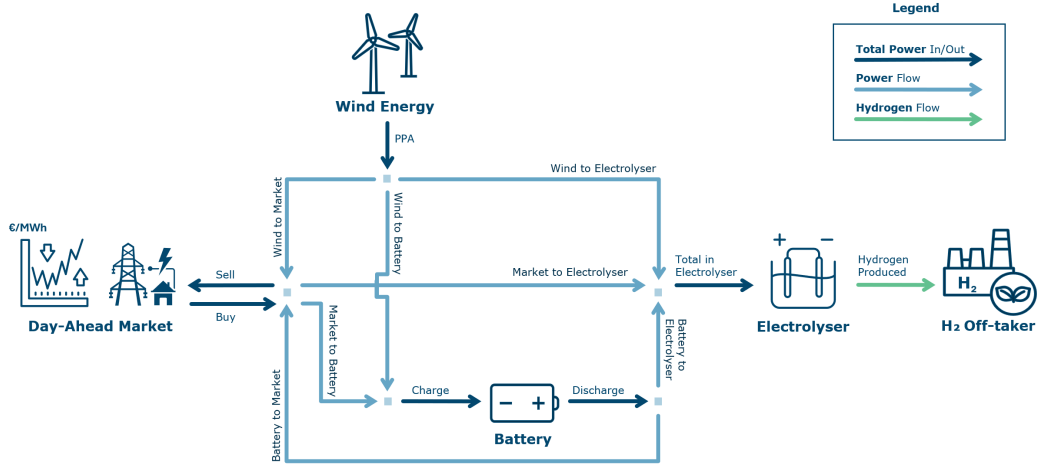


Figure 4.1: Schematic overview of the hybrid renewable energy and hydrogen production system with a wind Power Purchase Agreement (PPA), market interaction, a battery, an electrolyser and hydrogen offtake.

The sourced electricity, whether purchased from the market or sourced through the PPA, can be routed to the electrolyser to produce hydrogen, or stored temporarily in the battery. The battery can discharge its stored energy either to the electrolyser or directly to the market, depending on prevailing electricity prices and system needs. No hydrogen storage is modeled in this research. Instead, we assume that the hydrogen offtaker has sufficient on-site storage capacity to accept all hydrogen produced.

Hydrogen is sold under a long-term, fixed-price Hydrogen Purchase Agreement (HPA) governed by a *take-as-produced* contract. Under this agreement, all hydrogen produced can be sold to the offtaker at the agreed price. The contract includes an annual minimum hydrogen demand target, which the producer is obligated to meet. However, there is no maximum sales cap, meaning any excess hydrogen production beyond the target may also be sold at the same price if it is economically viable to do so.

Table 4.1: Base case system configuration.

Component	Specification
Electrolyser	50 MW capacity
Wind Power Purchase Agreement (PPA)	100 MW contracted capacity
Battery Energy Storage System	100 MWh energy capacity, 4-hour duration, 25 MW power rate
Grid Connection	100 MW bidirectional power capacity
Annual Hydrogen Production Target	180 GWh, equivalent to 5.4 kilo tonne H ₂ per annum (ktpa)

The system's base case configuration is presented in Table 4.1. At its core is a 50 MW alkaline electrolyser, chosen as a representative mid-scale deployment size. The other asset sizes are strategically selected to ensure balanced operational flexibility and enable meaningful system trade-offs. The wind PPA is deliberately oversized at 100 MW to prioritise renewable-sourced hydrogen. The lithium-ion BESS is configured with a 4-hour duration, providing 100 MWh of energy and a 25 MW charge/discharge rate. The grid connection is sized at 100 MW in both directions, supporting simultaneous electrolyser operation and market interaction. Finally, the annual hydrogen production target of 180 GWh (5.4 ktpa) ensures a minimum electrolyser utilisation rate of approximately 70%, reflecting a realistic commercial benchmark. Further technical details and comparisons of the alkaline electrolyser and lithium-ion BESS to other hydrogen production and energy storage technologies are provided in Chapter 2, in Sections 2.1.2 and 2.3, respectively.

4.2. System configuration test cases

To evaluate the impact of different system configurations on profitability and operational flexibility, four distinct test cases are modeled. These cases differ in three key dimensions: (1) *power sourcing strategy*: electricity can be procured from the market and/or through a fixed-price Wind Power Purchase Agreement (PPA), (2) *energy storage integration*: a battery energy storage system may be included or excluded, and (3) *market*

interaction ability: the system may be restricted to buying electricity, or allowed to both buy and sell power on the day-ahead market. This structured setup allows for a comparative analysis of the marginal value of individual system components (e.g., battery, market participation, PPA) under uncertainty in wind availability and electricity prices. An overview of the four test cases is presented in Table 4.2.

Table 4.2: Overview of system test cases.

Case	Battery	PPA	Grid Trading	System Description
Case 1	✓	✓	Buy and sell	Full flexibility for power sourcing and battery (<i>base case</i>)
Case 2	–	✓	Buy and sell	Renewable and grid power sourcing, no battery
Case 3	–	–	Buy-only	Grid-only power sourcing, no battery
Case 4	✓	–	Buy and sell	Grid-only power sourcing with battery

4.3. Model workflow

Figure 4.2 provides a schematic overview of the model workflow, illustrating the key model inputs, pre-processing steps, optimisation model, post-processing and resulting outputs.

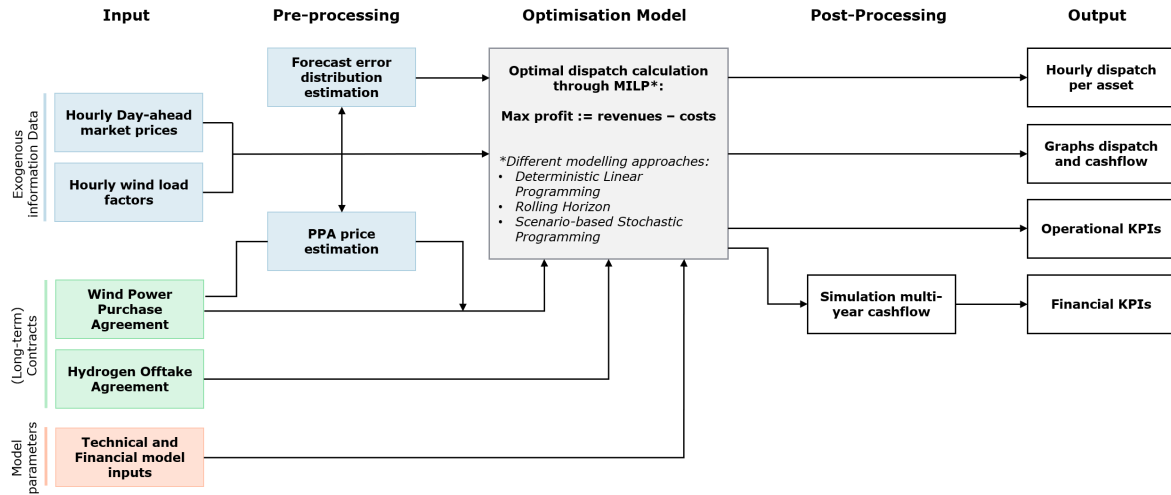


Figure 4.2: Schematic overview of the research approach. The figure illustrates the full workflow of the study, including the model inputs, pre-processing steps, optimisation model, and post-processing activities leading to final outputs.

Input: The model relies on several external input datasets, historical day-ahead electricity prices and wind load factors. Hourly day-ahead electricity price data for the Netherlands bidding zone is sourced from the ENTSO-E Transparency Platform [22]. Wind input is represented through historical wind load factors obtained from a representative location in the Netherlands, using data extracted from the certified and open-source Renewable Ninja platform [70]. Section 4.6 provides a detailed description and statistical analysis of these datasets. Section 4.7 presents an overview of the technological and financial model parameters used.

Pre-processing: Prior to optimisation, several steps are required to prepare the input data. The distribution of forecast errors for both wind generation and electricity prices is estimated to enable scenario generation. Due to the lack of available forecast data, the error distributions for both variables are approximated using the standard deviation calculated from historical time series data. Further details are provided in Section 6.3. Additionally, the PPA price is estimated based on the input datasets, with the methodology described in Section 4.7.4.

Optimisation model: The core optimisation problem determines the optimal dispatch strategy to maximise system profitability. The formulation of the optimisation problem, including the objective function, decision variables, and constraints, is detailed in Chapter 5. The selection of the modeling approach (e.g., deterministic or stochastic) determines whether scenario generation is applied and how long the optimisation look-ahead period is. The solution methods for deterministic and stochastic optimisation models, including rolling horizon and scenario-based approaches are explained in Chapter 6.

It is important to note which factors influence the optimisation decisions. In this research, the system component sizes (e.g., the installed capacity of the electrolyser, the PPA size, and the battery storage size) are assumed to be fixed inputs to the model. Consequently, parameters that depend solely on these fixed capacities, such as capital expenditures (CAPEX) and fixed operational expenditures (OPEX) do not impact the optimal dispatch decisions. These cost components are incorporated later in the post-processing phase for financial evaluation purposes. Instead, the optimisation is primarily driven by dynamic external inputs and variable operational costs. These include the time-varying day-ahead electricity prices, the hourly wind load factors, the hydrogen selling price, wind power redispatch costs and operational penalty costs associated with electrolyser shutdowns. In addition, technical constraints such as minimum load levels, maximum ramping rates, and battery state-of-charge limitations guide feasible operational decisions, although these are considered fixed constraint parameters rather than dynamic optimisation drivers.

Post-processing: After solving the optimisation problem, a post-processing step is conducted to evaluate the financial and operational performance of each system configuration and modeling approach. Based on the optimised dispatch over one operational year, a 20-year project lifetime is simulated to calculate financial key performance indicators (KPIs), namely the Net Present Value (NPV), Internal Rate of Return (IRR), and Levelised Cost of Hydrogen (LCOH). These calculations are detailed in Section 4.4.2. Operational KPIs are derived from the dispatch outputs, summarising energy flows and utilisation metrics as described in Section 4.4.1.

Outputs: The outputs of the model include the detailed hourly dispatch decisions, which can be exported to external files for further analysis if required. Graphical results are presented to illustrate key dynamics such as price evolution, wind load factors, hydrogen production, battery charging and discharging actions, market buying and selling behavior, and corresponding revenue and cost streams. These results are discussed in Chapter 7. In addition, the post-processed financial and operational KPIs provide a clear quantitative basis to compare different models and system setups.

4.4. Key Performance Indicators

To systematically evaluate the performance of the different system configurations and optimisation models developed in this thesis, a set of KPIs is defined. These KPIs allow for a comprehensive assessment of both the operational behavior and the financial viability of the hybrid hydrogen production system. The KPIs are categorised into two groups: operational and financial KPIs.

4.4.1. Operational KPIs

Operational KPIs provide insight into the technical performance and operational efficiency of the system over the one-year simulation period. These KPIs are crucial for understanding system flexibility, operational stability and the impact of different design choices on the dispatch strategy. The key operational KPIs considered are presented in Table 4.3.

Table 4.3: Overview of operational Key Performance Indicators (KPIs).

Category	Specific KPIs
Electrolyser Operation	<ul style="list-style-type: none"> - Electrolyser utilisation factor - Frequency and duration of start-up/shut-sown cycles - Distribution of power source (wind, market, battery)
Battery Usage	<ul style="list-style-type: none"> - Battery utilisation rate (charging/discharging hours) - Distribution of battery usage (wind, market, electrolyser)
Market Interaction	<ul style="list-style-type: none"> - Total volumes traded (buy/sell volumes) - Distribution market trades (direct buy/sale, battery storage)
Wind PPA	<ul style="list-style-type: none"> - Wind surplus volume and loss costs - Distribution wind PPA power (electrolyser, battery storage, direct sale)

The *electrolyser utilisation factor* measures the average utilisation of the electrolyser over the year. It is calculated as the ratio of hydrogen production to the installed maximum production capacity at each hour,

averaged over the entire simulation period. This metric provides insight into how intensively the electrolyser is used. The *frequency and duration of start-up and shut-down cycles* records the number of times the electrolyser is switched on or off, as well as the average duration of each start-up or shut-down event. This KPI is essential for assessing operational stability, as frequent cycling can lead to increased maintenance costs and reduced equipment lifespan. The *distribution of power source* tracks the share of electricity supplied to the electrolyser from wind (PPA), the market, or battery storage, providing insights into the system's reliance on different energy sources.

For renewable energy integration, the *wind surplus volume and loss costs* KPI measures the volume of surplus wind power available under the PPA that is not directly used by the electrolyser. This surplus can be directed towards battery storage, sold to the market, or curtailed (with a penalty cost). Since surplus power is often sold to the market, even when market prices are unfavorable, this KPI also tracks the associated annual loss costs incurred from selling at low prices or paying curtailment penalties. The *distribution of wind PPA power* captures how the accepted wind energy is allocated between electrolyser operation, battery storage, or direct sale to the market. Together, these operational KPIs provide a detailed understanding of the system's flexibility, efficiency, and dispatch strategy under varying market and renewable generation conditions.

4.4.2. Financial KPIs

The main financial outputs considered are revenues from hydrogen sales, revenues from electricity sales to the grid, electricity procurement costs, and shutdown penalty costs. In addition, the model accounts for fixed CAPEX and fixed annual OPEX costs. The annual profit is calculated based on these cost and revenue components and serves as a basis for comparing different system configurations. Additionally, three financial KPIs assess the economic feasibility and long-term profitability of the project under the assumption that the operational patterns observed in the simulated year are repeated over the full lifetime of the plant (20 years). The key financial KPIs calculated are the Net Present Value (NPV), the Internal Rate of Return (IRR), and the Levelised Cost of Hydrogen (LCOH).

To calculate these financial KPIs, we construct annual cash flow statements based on the operational results of the optimisation models. Several financial modeling assumptions and simplifications are applied. For the long-term financial assessment, annual cash flows are extrapolated over a 20-year project lifetime, assuming constant operational behaviour across years. While a more realistic approach would involve optimisation over the full project lifetime horizon, such a formulation would substantially increase model complexity. Therefore, this study adopts a one-year representative modeling period to focus on operational strategy decisions in response to hourly electricity prices and wind variability. The resulting financial metrics are used as indicative measures to assess the economic implications of different system setups and to support comparative analysis.

Secondly, depreciation and amortization (D&A) are included by depreciating the initial CAPEX linearly over the 20 years of plant life using a straight-line method. This means that each year, an amount equal to $\frac{1}{20}$ of the total CAPEX is deducted as a non-cash D&A expense. Although D&A is not an actual cash outflow, it is subtracted when calculating the taxable profit, thereby reducing the corporate tax burden. The annual taxable profit is computed as: *Revenue* minus *total operating costs* and minus *D&A* amount. A fixed corporate tax rate of 25.8% is applied, consistent with 2024 Dutch regulations [19]. Additionally, a discount rate of 10% is used to discount future cash flows to present value. It is noted that no inflation, equipment degradation, grid tariffs, or water consumption costs are included in the financial modeling. Furthermore, the analysis is conducted on an unleveraged basis, and interest rates or financing structures are not considered.

The resulting annual post-tax cash flows form the basis for the computation of NPV, IRR, and LCOH, providing a consistent and comparable framework for evaluating the profitability and economic feasibility across different optimisation models and system configurations.

Net present value

The NPV measures the cumulative net value of the project's future cash flows discounted back to present terms. A higher NPV indicates a more attractive investment opportunity. It is calculated as:

$$NPV = R_0 + \sum_{y=1}^N \frac{R_y}{(1+r)^y}, \quad (4.1)$$

where R_0 is the initial investment cost (negative value), R_y is the net cash flow in year y after tax, r is the discount rate (10%), and N is the total project lifetime (20 years).

Internal rate of return

The IRR is defined as the discount rate r that delivers an NPV of the project equal to zero: $NPV = 0$ at $r = \text{IRR}$. The IRR provides an intuitive measure of the expected annualised return on investment. A higher IRR reflects a more profitable and desirable project.

Levelised cost of hydrogen

The LCOH represents the discounted average cost per unit of hydrogen produced over the project's lifetime. It is a widely used benchmark metric in energy project evaluation. In this thesis, LCOH is calculated by taking the ratio of the discounted total project costs to the discounted total hydrogen production:

$$LCOH = \frac{NPV_{\text{Costs}}}{NPV_{\text{H}_2 \text{ Produced}}} = \frac{\sum_{y=1}^N \frac{C_y^{\text{CAPEX}} + C_y^{\text{O\&M}} + C_y^{\text{Shutoff}} + (C_y^{\text{Power}} - R_y^{\text{Grid}}) + C_y^{\text{Tax}}}{(1+r)^y}}{\sum_{y=1}^N \frac{H_y}{(1+r)^y}}, \quad (4.2)$$

where C_y^{CAPEX} represents the CAPEX cost, $C_y^{\text{O\&M}}$ are the fixed O&M costs, C_y^{Shutoff} the shutdown penalty costs, C_y^{Tax} the cost associated with tax, C_y^{Power} the electricity procurement costs, and R_y^{Grid} the electricity sales revenue. Moreover, H_y represents the total hydrogen production. Thus, the numerator represents the NPV of all project costs, while the denominator represents the NPV of hydrogen production over 20 years.

Higher hydrogen production typically leads to a lower LCOH by spreading fixed costs over more kilograms of hydrogen. However, maximising production can also increase variable costs (such as buying electricity at expensive hours), creating a trade-off that is evaluated through the optimisation models. Minimising the LCOH would ideally be part of the optimisation objective itself. However, implementing a minimisation over a variable denominator (total hydrogen production) makes the optimisation problem nonlinear and substantially more difficult to solve. Therefore, this thesis focuses on profit maximisation and uses LCOH only as a post-processing metric for performance evaluation.

4.5. Modeling assumptions and scope clarifications

To maintain tractability and focus, this thesis employs a number of modeling assumptions and simplifications. These are grouped into five categories: (i) technical and operational modeling assumptions, (ii) market participation and regulatory considerations, (iii) hydrogen offtake and green labeling, (iv) financial modeling scope, and (v) uncertainty representation. Where relevant, simplifications are justified based on the objectives and limitations of this research.

4.5.1. Technical and operational modeling

This research adopts a high-level optimisation approach to capture the operational flexibility and system-level dynamics of hybrid energy assets. Rather than modeling electrochemical processes in detail, the focus is on strategic dispatch and market interactions. Several simplifying assumptions are made to maintain computational tractability while preserving the key drivers of system behavior.

- **Electrolyser efficiency:** The electrolyser converts electricity into hydrogen at a fixed efficiency, assumed constant across all operating levels. Although real-world efficiency varies with factors such as temperature, pressure and load, a constant value is commonly used in system-level studies for tractability [2, 3]. In reality, efficiency is nonlinear and depends on load and electrolyser type [1, 51]. For alkaline electrolysers, it is lowest at very low loads due to auxiliary losses, peaks at moderate loads and declines at high loads. To retain linearity, this is approximated by enforcing a minimum load threshold below which the electrolyser shuts down. More details on electrolyser efficiency can be found in Section 2.1.2.
- **Degradation and maintenance:** The degradation of components such as the electrolyser and battery is not modeled explicitly over time. However, in reality degradation over time does occur and frequent on and off cycling increases this rate [9]. To reflect this technical wear and operational inefficiency, our model penalised start-ups with a fixed start-up cost and imposes a constraint on the annual number of shut-off events. These approximations are also included by [18, 28] and capture the operational trade-offs without requiring detailed lifetime degradation models.
- **Battery modeling:** The battery energy storage system is modeled with fixed charging and discharging efficiencies and a constant round-trip efficiency. In reality, these efficiencies can also vary with operational variables such as the state-of-charge (SOC), but we simplify by assuming constant values as seen

in existing research [27, 54]. Energy losses are only considered at the time of charging or discharging (no self-discharge or idle losses are modeled). The battery is subject to upper and lower SOC bounds to avoid extreme operating conditions that accelerate degradation [53].

- **Power losses and auxiliary consumption:** All power flows in the model are assumed to be lossless. Transmission losses, inverter/converter inefficiencies, and auxiliary loads (e.g. for control systems or balance-of-plant) are neglected. Additionally, electricity consumption for hydrogen compression is not explicitly modeled and is assumed to be included in the overall electrolyser efficiency.

4.5.2. Market participation

In this research, the primary focus lies on modeling strategic electricity procurement and hydrogen production scheduling, rather than developing a fully dynamic and detailed market trading model. Therefore, several simplifications are made regarding market behavior, participation mechanisms and grid-related considerations. These assumptions allow the model to concentrate on dispatch optimisation and operational flexibility, without introducing the complexities of endogenous market price formation and balancing or ancillary services participation.

- **Price-taking behaviour:** The system is modeled as a price taker in the electricity market, meaning its operations do not influence the market clearing price or the bidding behaviour of other participants. This assumption is justified by the relatively small scale of the electrolyser and battery compared to total system capacity [39]. All bids are assumed to be accepted at the prevailing market price, and trading decisions are made solely to maximise the system's economic performance, without consideration of strategic market interaction.
- **Market access:** The system participates solely in the day-ahead market. No revenue is modeled from intra-day, balancing, or ancillary service markets, although these are relevant in real-world operation.
- **No grid tariffs:** Grid connection and usage fees, such as capacity-based transmission tariffs or reactive power penalties, are not included in the profit or cost calculations.
- **No trading transaction costs:** No trading transaction costs, fees, or other market-related charges are included in the model. The system is assumed to transact electricity at the prevailing market-clearing price without incurring additional costs. This implies that electricity can be bought at the exact market price and sold at the same rate, enabling full potential in response to hourly price fluctuations.
- **No market failures:** Market mechanisms are assumed to function perfectly. Government interventions, subsidies, price caps, or scarcity pricing mechanisms are not included.

4.5.3. Hydrogen offtake and regulatory green labeling

This research aims to model hydrogen sales through a simplified and practical offtake structure, without delving into the detailed requirements of future certification frameworks. Several assumptions are made regarding the sales contract type, compliance with green hydrogen regulations and the evolving carbon intensity of grid electricity. These assumptions enable flexible system operation and allow the focus to remain on dispatch optimisation and system profitability.

- **Flexible load contract and no hydrogen storage:** All hydrogen is assumed to be sold under a fixed-price take-as-produced HPA which allows the electrolyser to operate flexibly and does not restrict the offtake of a baseload profile. We therefore assume that the offtaker has sufficient storage capacity and accepts any volume produced. Moreover, an annual minimum production target is enforced, but there is no upper bound, and any additional hydrogen can be sold for the contracted fixed price.
- **Green hydrogen assumption:** The base case model assumes that all hydrogen produced is eligible to be sold as green hydrogen, regardless of the electricity source. While the system occasionally draws power from the grid which may not be fully renewable, all hydrogen is assumed to be sold at a fixed price aligned with expected green or low-carbon hydrogen benchmarks. This is supported by the long-term assumption that grid carbon intensity in the Netherlands will decline due to increased renewable energy penetration, making grid electricity an increasingly low-carbon source [39]. Furthermore, the model does not account for compliance with Renewable Fuels of Non-Biological Origin (RFNBO) regulations; requirements such as hourly matching, additionality, and geographic correlation are not enforced.

4.5.4. Economic and financial evaluation

The primary focus of the optimisation model is on maximising operational profit during the simulation period, rather than conducting a financial modeling assessment. Accordingly, several simplifications are made concerning cost components and post-processing of financial indicators. This structure enables the research to provide meaningful financial insights while keeping the optimisation model computationally tractable.

- **Missing cost components:** The optimisation focuses on maximising operational profit (revenue minus energy costs and operational penalties). Costs such as water use, insurance, permitting, balancing of plant and grid connection fees are excluded. This might result in lower LCOH values than in reality. Recent studies by TNO mention that due to increase grid tariffs costs, nowadays grid tariffs make a significant contribution of more than €2/kg hydrogen to the LCOH value [19].
- **Financial modeling simplifications for NPV, IRR and LCOH:** Financial KPIs such as NPV, IRR and LCOH are calculated in post-processing based on simplified annual cash flow models. These metrics are not part of the optimisation itself but are used for comparative analysis. To approximate long-term project performance, annual results from the representative optimisation year are extrapolated over a 20-year lifetime, assuming constant operation each year.
- **Financing and support mechanisms:** In practice, hydrogen projects often benefit from public subsidies, concessional loans, tax credits or debt financing structures, which can substantially impact the economic viability and financial performance (improving NPV and IRR). In this study, such support mechanisms are not considered; a simplified, fully equity-financed project structure is assumed for the financial post-processing calculations. Consequently, the reported financial indicators may underestimate the attractiveness of projects under real-world conditions where support measures are available.

4.5.5. Forecasting and uncertainty modeling

This research does not focus on the development of a sophisticated forecasting model for electricity prices or renewable availability, as such modeling would require a separate and extensive effort. Instead, forecasts are based on historical data with simplified uncertainty modeling to realistically capture forecast errors without overcomplicating the optimisation framework. Different assumptions are made depending on the optimisation model variant, as detailed below.

- **Short-term forecasts as certain:** For all model variants, the 24-hour ahead forecast of electricity prices and wind availability is assumed to be accurate and known with certainty.
- **No endogenous learning:** Forecasts are treated as exogenous inputs. The model does not update or improve forecasts during the optimisation horizon and no learning mechanisms are included.
- **Simplified representation of uncertainty:** Uncertainty in prices and wind generation is modeled using normally distributed forecast errors. These errors are derived from historical data's standard deviation. This approach reflects a simplified uncertainty structure, despite the fact that real-world forecast errors result from a combination of sources and exhibit temporal and spatial dependencies. Due to the lack of detailed insight into forecasting algorithms and the unavailability of forecast datasets, this simplified, data-driven method was chosen to approximate uncertainty in a tractable manner.

4.6. Data for external information

This study incorporates two key sources of external information: hourly day-ahead electricity prices and hourly wind load factors. The year 2023 is used as the primary reference for input data, as it represents the most recent full year of operational data available at the start of this study. It reflects contemporary market conditions, regulatory developments, and renewable generation patterns, making it a relevant and representative basis for evaluating system performance under realistic assumptions. Figure 4.3 illustrates both variables over a selected period in February–March 2023. The plot clearly shows temporal variability in both wind availability and electricity prices, along with a negative correlation between the two. Periods of high offshore wind capacity factors often correspond with price dips, reflecting the impact of abundant low-marginal-cost renewable energy on market prices. Conversely, when wind availability is low, electricity prices tend to rise due to increased reliance on more expensive generation sources. However, this relationship is not deterministic and market prices are also influenced by other factors such as overall demand levels, fuel prices, generation mix, and intra-week dynamics.

Moreover, electricity prices exhibit a distinct 24-hour intraday pattern commonly referred to as the “duck curve”. This pattern features a midday dip in prices due to high solar PV output, followed by steep increases during the evening ramp-up period. The magnitude of this price swing can reach up to 46% in Europe, and is often even more pronounced in the Netherlands with high solar penetration [84]. Recognising these patterns in both inter-day and intra-day variability underscores the need to model external drivers with temporal granularity and inter-variable correlation.




Figure 4.3: Hourly electricity prices and offshore wind capacity factors over a selected period in 2023. The plot illustrates the inverse relationship between wind availability and market prices: as offshore wind capacity factors increase, electricity prices tend to decrease due to the increased supply of low-marginal-cost renewable energy. This dynamic is especially visible during mid-February and late February, where high wind output corresponds with noticeable price dips. Such patterns motivate the inclusion of inter-variable correlation in the scenario generation method.

4.6.1. Day-ahead electricity price data

This study relies on historical hourly day-ahead electricity price data sourced from the ENTSO-E Transparency Platform [22], which provides publicly available electricity market data across Europe. The data is used to inform the dispatch optimisation of the hybrid energy system, with the year 2023 serving as the base case for the model. In addition, sensitivity analyses are performed using historical price data from previous years (2019–2024) to assess model robustness under varying market conditions. To illustrate the temporal variability and price dynamics, several statistical and visual analyses are conducted. Figure 4.4 presents a distribution of hourly prices throughout 2023 using monthly boxplots. These plots clearly highlight the seasonal variability in electricity prices, with the highest volatility observed in the summer months. The monthly average prices vary between approximately €85/MWh and €135/MWh, while price extremes reach as low as -€480/MWh and as high as €400/MWh, reflecting occurrences of both scarcity and oversupply conditions.

Hourly price volatility is further analysed in Figure 4.5, which shows the histogram of hourly price changes. The distribution is centered around zero with a high peak, indicating frequent small changes and infrequent but extreme jumps, which is relevant for assessing the responsiveness requirements of the system.




Figure 4.4: Monthly distribution of day-ahead electricity prices in 2023. Each boxplot illustrates the spread, central tendency, and presence of outliers in hourly day-ahead prices for each month. The plots show clear seasonal trends, with lower median prices in winter and higher price volatility in the summer months. Extreme negative and positive price events are observed across the year, reflecting both oversupply and scarcity conditions. Source: ENTSO-E Transparency Platform [22].




Figure 4.5: Distribution of hourly changes in day-ahead electricity prices (€/MWh) in 2023. The histogram shows the frequency of hourly price differences (Δ Price), highlighting the high concentration of small changes around the mean and the presence of occasional large jumps. The red dashed line indicates the mean, while the orange dashed lines show one standard deviation above and below the mean. This distribution illustrates the short-term volatility of day-ahead prices, which is relevant for assessing the responsiveness and flexibility requirements of energy system components. Source: ENTSO-E Transparency Platform [22].

Market prices over the past years

Figure 4.6 illustrates the distribution of hourly day-ahead electricity prices in the Netherlands from 2019 to 2024, while Table 4.4 provides key statistical indicators for each year. The data highlights a substantial shift in both the average price level and its volatility during this period.

In the years 2019 and 2020, electricity prices were relatively stable, with median values around €40/MWh and interquartile ranges spanning approximately €34 to €47/MWh. However, from 2021 onward, a sharp increase in both the average price and its dispersion is observed. This shift coincides with the onset of the European energy crisis, triggered by a rapid rise in natural gas prices, reduced Russian gas flows to Europe, and broader geopolitical tensions. As a result, the mean day-ahead electricity price surged to €102.96/MWh in 2021 and further to €241.92/MWh in 2022, marking an unprecedented escalation.

The volatility of prices also increased dramatically during this period. The standard deviation, which remained around €11–15/MWh in the pre-crisis years, rose to €74.71/MWh in 2021 and peaked at €131.53/MWh in 2022. This reflects the extreme price fluctuations experienced during periods of tight supply-demand bal-

ance, particularly in winter months with limited renewable generation and high heating demand. Many hourly prices exceeded €500/MWh, and several instances in 2022 and 2024 saw peaks approaching or surpassing €800/MWh. These extreme outliers were typically linked to system stress conditions such as low renewable output, high peak demand, unplanned outages, or constrained cross-border capacity.

Although average prices began to decline in 2023 and 2024 (to €95.82/MWh and €77.29/MWh respectively) the interquartile range (Q25–Q75) remained significantly wider than in the pre-crisis period. For example, in 2024, the interquartile range spanned from €56.58 to €100.00/MWh, indicating persistent market uncertainty. The elevated spread suggests that while the worst of the crisis may have passed, electricity markets remain exposed to volatility due to renewable intermittency, fuel market instability, and ongoing regulatory reforms such as capacity mechanism designs and market structure changes.

In addition to higher volatility, the frequency of negative electricity prices has increased significantly in recent years. This trend is particularly pronounced in 2023 and 2024, which saw a substantial rise in the number of hours with negative day-ahead prices in the Netherlands. This development is primarily attributed to a combination of high solar PV output and limited export capacity during periods of low demand. For electrolyser systems, such events present valuable arbitrage opportunities, as producers can effectively be paid to consume electricity while generating hydrogen. This dynamic strengthens the economic case for allowing grid-sourced electricity in hydrogen production strategies, especially under flexible operating regimes.

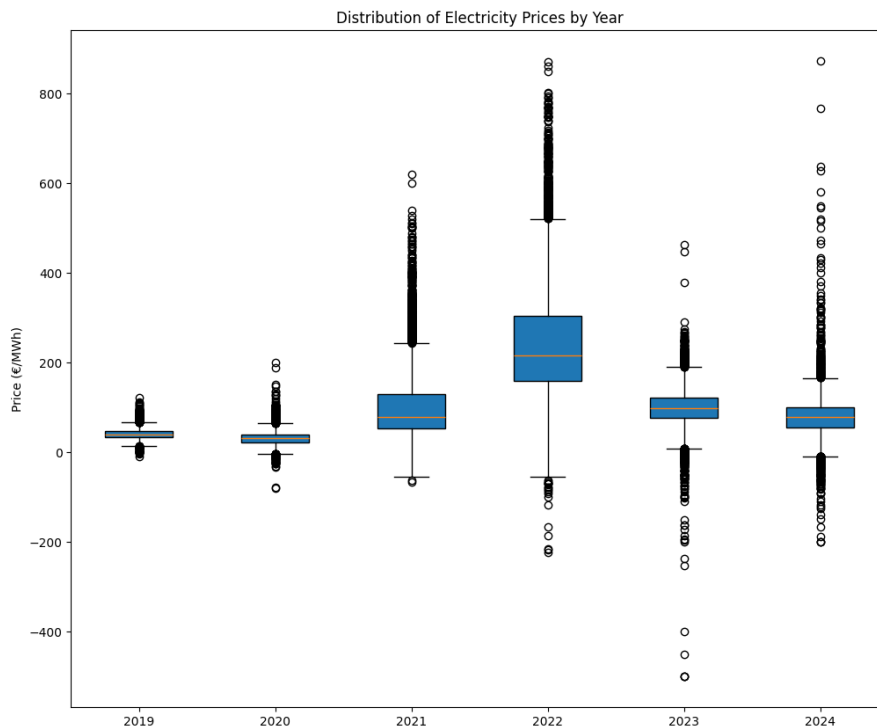


Figure 4.6: Distribution of hourly day-ahead electricity prices (€/MWh) in the Netherlands from 2019 to 2024. The figure shows increasing price levels and volatility during the 2021–2022 energy crisis, with a rise in both extreme values and interquartile spread. A higher frequency of negative prices is also visible in 2023 and 2024. Summary statistics are provided in Table 4.4. Source: ENTSO-E Transparency Platform [22].

Table 4.4: Summary statistics of hourly day-ahead electricity prices (€/MWh) in the Netherlands for the years 2019–2024. The table reports key indicators including the mean, median, standard deviation, variance, and interquartile range (Q25–Q75). These values complement the distributional trends illustrated in Figure 4.6, highlighting the significant rise in both average prices and volatility during the 2021–2022 energy crisis. The post-crisis years show some price stabilisation, though variability remains elevated compared to pre-crisis levels. Source: ENTSO-E Transparency Platform [22].

Year	Mean	Median	Std Dev	Variance	Q25	Q75
2019	41.20	39.70	11.28	127.20	34.10	47.39
2020	32.25	31.67	15.31	234.27	22.90	40.19
2021	102.96	78.30	74.71	5581.17	53.72	130.05
2022	241.92	216.99	131.53	17299.03	160.00	304.94
2023	95.82	99.22	49.05	2405.57	76.94	122.40
2024	77.29	80.00	49.49	2449.27	56.58	100.00

4.6.2. Wind load factor data

The wind load factor is a number bounded between 0 and 1 that represents the ratio of actual wind energy generation to the theoretical maximum generation capacity of a wind farm at any given moment. Mathematically, the load factor at time t is defined as:

$$\text{Load Factor}_t = \frac{P_{\text{actual},t}}{P_{\text{rated}}},$$

where $P_{\text{actual},t}$ is the wind power output at hour t , and P_{rated} is the installed capacity of the wind farm. This thesis uses hourly historical wind load factor time series for the Netherlands obtained from the Renewables.ninja platform¹ [70], which generates wind power profiles based on reanalysis weather data. Specifically, the wind data is derived from the MERRA-2 reanalysis dataset [76], calibrated to local conditions using satellite-based validation techniques and weather station data, including the Royal Netherlands Meteorological Institute (KNMI) data. These load factors are representative of average wind power output across the country and are suitable for use in modeling aggregated wind energy inputs to a national hydrogen production system.

Figure 4.7 presents the monthly distribution of wind load factors for 2023. The data reveals clear seasonal variation, with higher median values in winter months and lower levels during late summer, consistent with meteorological wind patterns over the North Sea.




Figure 4.7: Monthly distribution of hourly wind load factors for the Netherlands in 2023. The figure shows strong seasonal variation, with higher wind availability during winter and lower values during summer months. The spread in each boxplot reflects temporal variability in hourly wind conditions, relevant for modeling renewable generation potential. Source: Renewables.ninja [70].

Figure 4.8 displays the distribution of hourly wind load factors in the Netherlands for the years 2019 through 2024, with accompanying summary statistics provided in Table 4.5. The data reveals that wind con-

¹<https://www.renewables.ninja>

ditions have remained relatively consistent from 2020 to 2024, with mean load factors ranging between 0.45 and 0.51, and standard deviations clustered around 0.30. This suggests a stable wind resource profile over recent years, supporting its reliability for long-term energy planning and investment analysis.

In contrast, the year 2019 stands out as an outlier, with a significantly lower mean wind load factor of 0.33 and a narrower interquartile range. This deviation may be attributed to below-average wind speeds during that year, limitations in early wind fleet deployment and measurement coverage, or improvements in turbine technology and capacity factors that have occurred since 2019, leading to more efficient wind power generation in later years.




Figure 4.8: Distribution of hourly wind load factors in the Netherlands for the years 2019–2024. The figure highlights the relative consistency of wind resource availability from 2020 onwards, with mean values between 0.45–0.51. The year 2019 stands out with a lower average and narrower distribution. Summary statistics are provided in Table 4.5. Source: Renewables.ninja [70].

Table 4.5: Summary statistics of hourly wind load factors in the Netherlands for the years 2019–2024. The table provides key indicators including the mean, median, standard deviation, variance and interquartile range (Q25–Q75). The results confirm a relatively stable wind resource profile from 2020 onwards, while 2019 appears as an outlier with both a lower mean load factor and a narrower interquartile range. These values complement the distributional trends shown in Figure 4.8. Source: Renewables.ninja [70].

Year	Mean	Median	Std Dev	Variance	Q25	Q75
2019	0.332	0.284	0.239	0.057	0.145	0.471
2020	0.501	0.517	0.315	0.099	0.205	0.786
2021	0.454	0.437	0.296	0.088	0.184	0.706
2022	0.462	0.446	0.297	0.088	0.199	0.718
2023	0.510	0.519	0.302	0.091	0.240	0.776
2024	0.480	0.474	0.298	0.088	0.217	0.741

4.7. Model parameters

The model parameters can be categorised into three main groups: temporal parameters (Section 4.7.1), including time resolution and modeling horizon, technical parameters related to the electrolyser (Section 4.7.2) and battery systems (Section 4.7.3), and commercial parameters that define renewable electricity prices, hydrogen contracts, and other market-related inputs (Section 4.7.4).

4.7.1. Temporal resolution and modeling period

The system is modeled with an hourly time resolution, capturing short-term operational dynamics such as electricity price fluctuations, intermittent renewable generation, and energy storage behaviour. This level of granularity is necessary for accurately representing variability in wind power availability and market conditions on the day-ahead electricity market. The full modeling horizon spans one calendar year (8760 hours), allowing the model to incorporate seasonal trends, long-term contractual commitments (e.g., hydrogen offtake

agreements), and annual system performance metrics. All time-series inputs, including wind load factors and day-ahead electricity prices, are synchronised to this hourly resolution to ensure temporal consistency across the model.

Rolling horizon parameters

For the rolling horizon methods, the optimisation problem is solved repeatedly over a shorter forward-looking time window, referred to as the *planning horizon* or *look-ahead period*. After each optimisation run, only the decisions corresponding to a limited *planning period* are implemented, after which the horizon shifts forward by one planning period, and the process repeats. This iterative approach allows the model to incorporate new data or forecasts as they become available, enabling a more realistic representation of decision-making under uncertainty and limited foresight.

Table 4.6 summarises the key temporal parameters used in the model. While the full operation length is fixed at one year, the rolling horizon methods divide this horizon into overlapping segments defined by a planning horizon of 168 hours (one week) and a planning period of 24 hours (one day).

Table 4.6: Temporal parameters for rolling horizon models.

Parameter	Value (hours)	Symbol
Operation length	8760	T
Planning period	24	T'
Planning horizon	168	H

4.7.2. Electrolyser parameters

Table 4.7 summarises the main techno-economic input parameters used for the electrolyser component. These values are based on a combination of Lazard's industry reports [47], Power2X internal assumptions, and expert consultation. Due to confidentiality, exact CAPEX breakdowns are not disclosed.

The electrolyser converts electricity into hydrogen at a fixed energy conversion efficiency η_{el} , which determines the amount of hydrogen produced per unit of electricity consumed. In line with Lazard's 2024 benchmark for alkaline electrolysers [47], a constant efficiency of 60% is assumed. The electrolyser operates within a partial load range, constrained by a minimum load threshold of 30% and a maximum capacity of 100%. In our base case configuration, this corresponds to a 50 MW electrolyser operating between 15 MW and 50 MW.

Table 4.7: Electrolyser's technical and economical input parameters.

Parameter	Unit	Value	Symbol
Maximum capacity	% of rated power	100	P_{\max}
Minimum capacity	% of rated power	30	P_{\min}
Efficiency	%	60	η_{el}
Maximum # of shut-downs	per year	20	N^{off}
Cold start-up cost	€/cycle	8,000	C_{su}
Electrolyser CAPEX	€/MW	1,750,000	C_e^{capex}
Electrolyser OPEX	% of CAPEX/year	2.5	δ_e^{opex}

Cold start-up cost estimation

The fixed cost associated with a shut-down/start-up cycle represents the incremental wear imposed on the electrolyser stack during each cold start. Frequent cycling leads to mechanical stress and reduced stack lifetime, and is therefore penalised in the model using a fixed cold start-up cost C_{su} , as commonly adopted in the literature [18, 28]. This penalty cost is calculated using three variables: the total stack replacement cost C_{stack} , the expected operational lifetime T_{lifetime} , and the number of equivalent operating hours attributed to a single shut-down cycle $T_{\text{equiv shutdown}}$. The latter reflects the estimated degradation impact of a cold start, set to 30 hours based on expert judgment and representing a conservative "worst-case" assumption. The resulting shut-down cost is computed as:

$$\text{Shutdown Cost (€/cycle)} = C_{\text{stack}} \cdot \frac{T_{\text{equiv shutdown}}}{T_{\text{lifetime}}}. \quad (4.3)$$

Assuming a stack lifetime of 67,500 hours [47] and internal estimates for replacement cost, the derived cost is approximately €7,778. For implementation purposes, this value is rounded to €8,000.

4.7.3. Battery parameters

Table 4.8 summarises the technical and economic parameters used for the battery system in this thesis. The operational limits as well as the CAPEX and OPEX parameter values are primarily based on Lazard's 2024 Levelised Cost of Storage estimates for utility-scale, standalone lithium-ion systems with a 4-hour duration and power capacities between 50–100 MW [47]. In this model, we implement a 25 MW battery with a 4-hour duration, resulting in a total usable energy capacity of 100 MWh. This size was chosen to complement the 50 MW electrolyser, offering meaningful flexibility in system operation without significant oversizing.

The operational characteristics of the battery, including charging and discharging efficiency as well as the usable state-of-charge (SOC) range, are based on values commonly used in academic and applied research. Several studies adopt similar assumptions, typically setting round-trip efficiency at 90–95% and SOC limits between 10–20% minimum and 80–90% maximum [1, 2, 12, 26, 28, 54]. For this work, both charging and discharging efficiencies are set to 95% and the SOC range is restricted between 20% and 90% of the total capacity.

Table 4.8: Battery storage technical and economical input parameters.

Parameter	Unit	Value	Symbol
Charging efficiency	%	95	η^C
Discharging efficiency	%	95	η^D
Minimum SOC	% of capacity	20	SOC^{\min}
Maximum SOC	% of capacity	90	SOC^{\max}
Maximum charge power	MWh	25	$p^{C,\max}$
Maximum discharge power	MWh	25	$p^{D,\max}$
Battery CAPEX	€/MWh	282,512	C_b^{capex}
Battery OPEX	% of CAPEX/year	2.8	δ_b^{opex}

4.7.4. Commercial parameters

The model relies on a number of key commercial parameters, summarised in Table 4.9. These values define core assumptions around electricity sourcing, hydrogen sales, and system operation under commercial constraints. Electricity is primarily sourced from a dedicated wind PPA with a contracted capacity of 100 MW. The PPA price is estimated using a weighted average of the capture price and levelised cost of energy (LCOE), as further detailed in the next paragraph. In addition, a grid export limit of 100 MW is imposed, even though the combined capacity of wind generation and battery discharge could exceed this. In practice, grid connection capacity is often minimised to reduce infrastructure costs and avoid long permitting procedures, making 100 MW a realistic and conservative design choice.

The hydrogen production system is expected to meet an annual production target of 180 GWh. This value is based on a minimum target utilisation rate of approximately 70% of the 50 MW electrolyser's full capacity, ensuring the system operates at a realistic and economically viable level throughout the year.

The hydrogen selling price is a key parameter affecting system profitability. It is important to set this price above the LCOH to ensure economic viability, while remaining within realistic bounds for future green hydrogen markets. TNO's RHyCEET study [19] estimates a range of hydrogen production costs for different technologies and scenarios, generally between €200/MWh and €400/MWh. Based on these benchmarks, a hydrogen selling price of €300/MWh is adopted in this work to reflect a reasonable market-consistent value for green hydrogen.

Table 4.9: Commercial input parameters regarding the grid connection, power purchase agreement (PPA) and hydrogen sales agreement (HPA).

Parameter	Unit	Value	Symbol
Grid connection capacity	MW	100	K^{grid}
Contracted wind PPA capacity	MW	100	P^{PPA}
PPA electricity price	€/MWh	97	π^{PPA}
Redispatch penalty	€/MWh	150	C^{bal}
Annual hydrogen production target	GWh	180	D^{target}
Hydrogen selling price (HPA)	€/MWh	300	π^{H_2}

PPA price calculation

To estimate a realistic pricing level for the wind PPA, the *average capture price* is computed. This reflects the average market value of electricity generated by the wind installation, taking into account hourly market prices and wind production profiles. At each hour $t \in \{1, \dots, T\}$, where $T = 8760$ represents the total number of hours in a year, the wind power output is calculated as:

$$p_t^{\text{wind PPA}} = \alpha_t \cdot K_{\text{PPA}},$$

where $\alpha_t \in [0, 1]$ denotes the wind load factor, and K_{PPA} is the rated capacity of the wind park in MW. The electricity price at hour t is denoted by λ_t in euros per megawatt-hour (€/MWh). The revenue at hour t , r_t , is then calculated by:

$$r_t = p_t^{\text{wind PPA}} \cdot \lambda_t = \alpha_t \cdot K_{\text{PPA}} \cdot \lambda_t.$$

The total revenue R and energy E over the year are computed as:

$$R = \sum_{t=1}^T \alpha_t \cdot \lambda_t \cdot K_{\text{PPA}}, \quad E = \sum_{t=1}^T \alpha_t \cdot K_{\text{PPA}}, \quad \text{with } T = 8760 \text{ hours.}$$

The *average capture price* is defined as the revenue-weighted average price received per MWh of wind energy, which shows to be independent of the wind PPA size K_{PPA} :

$$\pi_{\text{cap}} = \frac{R}{E} = \frac{\sum_t \alpha_t \cdot \lambda_t}{\sum_t \alpha_t}. \quad (4.4)$$

The contractual PPA price π_{PPA} is then calculated as a weighted average² between the market-based capture price and a fixed LCOE π_{LCOE} :

$$\pi_{\text{PPA}} = 0.3 \cdot \pi_{\text{cap}} + 0.7 \cdot \pi_{\text{LCOE}}. \quad (4.5)$$

In this thesis we set $\pi_{\text{LCOE}} = 97\text{€/MWh}$, in line with recent literature values published by Lazard for offshore wind power [47]. The resulting capture prices and PPA contract prices calculated using historical hourly wind and electricity data are summarised for the years 2019–2024 in Table 4.10.

Table 4.10: Historical wind capture prices and calculated PPA prices (2019–2024) based on Equations 4.4 and 4.5 respectively.

Year	Average Capture Price (€/MWh)	Final PPA Price (€/MWh)
2019	40.07	79.92
2020	30.35	77.00
2021	97.39	97.12
2022	208.80	130.54
2023	85.10	93.43
2024	68.30	88.39

²This method for estimating PPA price was introduced by Power2X and is used for project evaluations. For comparison, TNO [19] estimated a PPA price for electricity from offshore wind in 2023 at 75 €/MWh making our estimation more conservative.

Redispatch costs

To reflect curtailment penalties for unused wind power under the PPA, a redispatch cost of €150/MWh is imposed whenever the model is unable to accept the full wind generation. This value is derived from national congestion cost statistics published by TenneT [84].

In the Netherlands, the total redispatch volume in 2023 amounted to 599 GWh, a significant reduction from the 972 GWh recorded in 2022. This corresponds to a 39% year-on-year decrease in redispatch volumes. In parallel, total congestion management costs which include redispatch, restriction contracts, and reactive power support, declined from €388 million in 2022 to €278 million in 2023, representing a 30% cost reduction. These reductions are largely attributed to lower wholesale electricity prices and increased reliance on restriction and reactive power contracts. However, the €278 million figure reflects the full scope of congestion management, not just redispatch. According to available breakdowns, redispatch typically accounts for approximately one third of total congestion management costs. Based on this, we estimate the redispatch-specific costs to be around €93 million.

To estimate a representative redispatch penalty for curtailment in the model, the total dispatch cost is divided by the total redispatch volume in 2023:

$$\frac{93 \text{ million euro}}{599 \text{ GWh}} \approx 155 \text{ €/MWh.}$$

We round this to an average redispatch penalty of €150/MWh, which is used in the model to represent the opportunity cost of curtailing wind power sourced from the PPA. This cost is incurred when available renewable energy cannot be utilised due to storage constraints, operational strategies, or market interactions. All data is sourced from the Annual Market Update 2023 published by TenneT, which provides an overview of key developments in the Dutch and German electricity markets [84].

4.8. Model implementation

The optimisation model was developed entirely from the ground up for the purpose of this thesis. All code used for formulating the MILP problem, implementing the rolling horizon and stochastic structure, and processing the results was written by the author, without relying on any pre-existing optimisation model implementations or template code. While standard modeling libraries such as Gurobi's Python API were used, the model logic and structure were designed and implemented independently.

The implementation is written in Python, using the Jupyter Notebook interface combined with Visual Studio Code for enhanced programming flexibility. Python was chosen for its extensive libraries and compatibility with optimisation solvers.

The optimisation problems were solved using the Gurobi Optimizer under a free academic license. Gurobi is renowned for its robust performance in solving large-scale linear and mixed-integer problems efficiently and accurately [33]. The specific solver version utilised was Gurobi Optimizer 11.0.2, build v11.0.2rc0, running on Windows 11 (OS Build 22621.2).

The computational experiments were conducted on a Microsoft Surface Laptop Studio equipped with an 11th Gen Intel(R) Core(TM) i7-11370H CPU operating at 3.30 GHz and 16 GB of installed RAM. The operating system is Windows 11 Business, Version 22H2.

5

Model Formulation

This chapter presents the formal mathematical formulation of the optimisation problem. The system is structured using a mathematical optimisation framework, consisting of three key components. First, the objective function to be maximised is described in Section 5.1. Second, the system's variables are defined in Section 5.2, including the decision variables that control system behaviour and the state variables that describe the system's evolution over time. Third, the constraints that define the set of feasible operations are formulated in Section 5.3. Finally, Section 5.4 outlines the main characteristics of the resulting optimisation model, including its classification as a Mixed-Integer Linear Program (MILP) following appropriate linearisation of certain constraints. This section also introduces model extensions that incorporate different foresight assumptions and approaches to handling uncertainty. Finally, Section 5.5 concludes the chapter by presenting the complete mathematical formulation of the optimisation model.

Notation For clarity, the notation used throughout this thesis is as follows: time-dependent variables are denoted by lower case letters, with power flows specifically indicated by p_t . Binary decision variables are represented by z_t . The electricity market prices are denoted by λ_t while the wind load factors are denoted by α_t . Abbreviations are used for each asset for the powerflow variables: E for the electrolyser, B for the battery, W for wind power, and M for the day-ahead market. Finally, most fixed model parameters are denoted by uppercase letter, where generally C is used for costs, and K for capacity values, with the exception of π used for fixed prices. For an overview of all symbols and their descriptions used throughout this thesis, refer to the List of Symbols Table in the Nomenclature.

5.1. Objective function

The objective of the optimisation problem is to maximise the total gross profit over a given time horizon. This profit is defined as the total revenue from selling hydrogen and electricity, minus all relevant operational and capital expenditures. In the mathematical formulation, the total profit is represented by the reward function \mathcal{R} , where all revenue components are denoted by S , referring to sales, and all cost components are grouped under F , referring to financial expenditures.

The revenue components consist of hydrogen sales revenue S_{H_2} and electricity sales revenue in the day-ahead market S_{DA} . Costs are divided into variable and fixed components. Variable costs include the electricity procurement cost F_{power} , which accounts for the cost of electricity drawn from the grid or purchased through the wind PPA, and the shut-off cost $F_{\text{shut-off}}$, representing the operational penalties or inefficiencies associated with switching the electrolyser on and off. Fixed costs include the annualised capital expenditure F_{Capex} , which reflects the investment costs associated with the system components, and the fixed operational expenditure F_{Opex} , which covers recurring costs such as maintenance and labour.

The overall objective function is thus given by:

$$\max \mathcal{R} := \underbrace{S_{H_2} + S_{DA}}_{\text{Revenue}} - \underbrace{F_{\text{power}} - F_{\text{shut-off}}}_{\text{Variable Costs}} - \underbrace{F_{\text{Capex}} - F_{\text{Opex}}}_{\text{Fixed Costs}}. \quad (5.1)$$

5.1.1. Time horizon and resolution

We define the discrete time horizon as the set $\mathcal{T} = \{1, \dots, T\}$, where each element $t \in \mathcal{T}$ represents a single hour. Throughout this thesis, we assume an hourly time resolution ($\Delta t = 1$ hour) over a full year, resulting in $T = 8760$ time steps. This choice allows the model to capture seasonal patterns and operational flexibility over long time scales.

5.1.2. Time-dependent external information: day-ahead prices and wind load factors

At each time step, the system is exposed to exogenous information that varies over time. The two primary inputs are the electricity market price and the renewable energy availability via wind power. Based on the electricity price λ_t and the wind load factor α_t , the operator must decide at each time step how much electricity to buy or sell on the market, whether to charge or discharge the battery, and how much hydrogen to produce.

The hourly day-ahead electricity market prices are denoted by the sequence $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_T)$, where λ_t represents the market clearing price at time t , in €/MWh. This price determines the cost of purchasing electricity from the grid, whether it is used directly by the electrolyser or to charge the battery. When electricity is sold to the market, a transaction fee or spread $\bar{\lambda}$ is subtracted from the price received, reflecting trading costs or losses due to inefficiencies.

The operator also has access to renewable electricity through a PPA for wind energy. This long-term contract follows a take-or-pay structure: the operator pays a fixed price π^{PPA} per MWh of contracted capacity, regardless of actual usage. The wind availability at each time step is determined by the wind load factor sequence $A = (\alpha_1, \alpha_2, \dots, \alpha_T)$, where each $\alpha_t \in [0, 1]$ indicates the proportion of the maximum turbine capacity that is converted to usable electricity at time t . These values are influenced by external conditions such as wind speed and turbine technology, among others.

The actual wind energy available through the PPA is computed as $p_t^{W, \text{PPA}} = \alpha_t K^{\text{PPA}}$, where K^{PPA} is the contracted PPA capacity in MW. If the operator chooses not to utilise the full available wind power, for example, due to storage constraints or a low market price, the unused share must still be paid for, and a balancing cost C^{bal} is incurred for each MWh of excess power. This cost accounts for system-level penalties or redispatch needs when wind energy is curtailed.

5.1.3. Fixed Costs

Capital and fixed operational costs are incurred independently of the system's operational decisions, as the installed capacities are assumed to be fixed. Although they do not influence the dispatch decisions directly, these costs are included in the model to ensure a realistic representation of system profitability. The capital expenditure (CAPEX), denoted by $C_{e/b}^{\text{capex}}$, is calculated based on the installed capacities of the electrolyser and the battery, denoted by K^{Elec} and K^{Batt} , respectively. The unit CAPEX for the electrolyser is expressed in €/MW, and for the battery in €/MWh of usable energy. These are multiplied by an annualisation factor ξ^{annual} to account for the investment being spread over its economic lifetime.

Fixed operational costs (OPEX) are assumed to be a fixed percentage of the initial CAPEX and are incurred annually as part of maintenance and operational support. Here $\delta_{e/b}^{\text{opex}} \in [0, 1]$ is the fixed OPEX ratio expressed as a fraction of the CAPEX for the electrolyser and battery, respectively.

The fixed cost term is given by:

$$F_{\text{Capex}} + F_{\text{Opex}} = \xi^{\text{annual}} \left(C_e^{\text{capex}} K^{\text{Elec}} + C_b^{\text{capex}} K^{\text{Batt}} \right) + \delta_e^{\text{opex}} C_e^{\text{capex}} K^{\text{Elec}} + \delta_b^{\text{opex}} C_b^{\text{capex}} K^{\text{Batt}}. \quad (5.2)$$

5.1.4. Variable costs

Power costs

Variable power-related costs F_{power} consist of three components: electricity purchases from the day-ahead market ($p_t^{M, \text{buy}}$), electricity obtained from the wind PPA ($p_t^{W, \text{PPA}}$), and penalties for unused wind energy ($p_t^{W, \text{overcap}}$). These are combined as:

$$F_{\text{power}} = \sum_{t \in \mathcal{T}} \left(\lambda_t p_t^{M, \text{buy}} + \pi^{\text{PPA}} p_t^{W, \text{PPA}} + C^{\text{bal}} p_t^{W, \text{overcap}} \right), \quad (5.3)$$

where λ_t denotes the day-ahead market electricity price at time t , π^{PPA} is the fixed contractual price per MWh of electricity procured via the wind power purchase agreement, and C^{bal} represents the balancing penalty cost imposed for curtailing excess wind energy that cannot be utilised within the system.

Additional power penalty cost: preference for wind-powered hydrogen production

To encourage the use of behind-the-meter wind energy for hydrogen production, aligned with both economic and environmental objectives, an additional modeling penalty is introduced. Specifically, a small artificial cost is added to the electricity used from the market and from the battery for electrolyser operation. This penalty ensures that, when the model faces multiple dispatch combinations yielding identical profits, due to equal total hydrogen production and summed market buy/sell power flows, it preferentially selects configurations where the electrolyser is supplied mostly by wind power rather than market power or battery power.

The penalty term is implemented as:

$$\sum_{t \in \mathcal{T}} \bar{\lambda} (p_t^{M \rightarrow E} + p_t^{B \rightarrow E}), \quad (5.4)$$

where in this case $\bar{\lambda} \approx 0.001$ is a symbolic penalty coefficient, small enough not to distort economic decisions, but sufficient to break ties in favor of green hydrogen production. This adjustment addresses a modeling artifact where different power-routing configurations, such as using wind for battery charging instead of direct electrolyser supply (market is used for electrolyser powering), can yield the same objective value, despite differing in sustainability impact. The introduced term subtly prioritises configurations that maximise the use of available wind energy for direct hydrogen production.

Electrolyser shut-down costs

Frequent start-stop cycles accelerate the degradation of electrolyser stacks and increase maintenance needs. To discourage excessive shutdowns and start-ups, a cost C^{su} is applied each time the electrolyser is turned off. This is implemented using a binary variable $z_t^{\text{su count}}$, which equals 1 when the electrolyser shuts off at time t , and 0 otherwise. A similar approach was used by Ding et al. [18] and Ghasemi [28]. The annual shutdown cost term is then calculated as:

$$F_{\text{shut-off}} = \sum_{t \in \mathcal{T}} C^{\text{su}} z_t^{\text{su count}}. \quad (5.5)$$

5.1.5. Revenue

Hydrogen sales revenue

Hydrogen is sold through a fixed-price Hydrogen Purchase Agreement (HPA), in which the operator commits to supplying a predefined annual quantity of hydrogen at a fixed price π_{H_2} , expressed in €/MWh hydrogen. Although hydrogen is typically expressed in kilograms, we adopt an energy-based representation in this model to ensure consistency with the electricity system. Specifically, we convert kilograms of hydrogen to thermal MWh using the lower heating value (LHV) of hydrogen. All hydrogen production volumes h_t are thus expressed in thermal MWh throughout the model.

The total revenue from hydrogen sales is calculated as:

$$S_{H_2} = \sum_{t \in \mathcal{T}} \pi_{H_2} h_t. \quad (5.6)$$

Electricity sales revenue

In addition to hydrogen production, the operator has the option to generate additional revenue by selling electricity to the day-ahead market. This can occur in several ways. First, surplus electricity from the wind PPA that exceeds the electrolyser's consumption can be redirected to the market. Second, favorable market conditions, such as high spot prices, may make electricity sales more profitable than hydrogen production. Lastly, the battery contributes to economic optimisation by storing electricity during low-price hours and selling it during high-price periods. In this thesis, electricity trading is limited to the day-ahead market. Participation in other wholesale markets (e.g., intraday or balancing) is considered out of scope.

To avoid unrealistic dispatch behavior during periods of zero electricity prices, a small artificial penalty is applied to each unit of electricity sold to the market. This penalty, again denoted by $\bar{\lambda}$, reduces the effective sales price and discourages economically meaningless transactions. Without this adjustment, the model might engage in simultaneous buying and selling of electricity during consecutive zero-price hours without affecting the objective value, resulting in unrealistic outcomes.

Although this term could be interpreted as a symbolic transaction cost or market spread, its primary role in this thesis is to ensure modeling consistency. The same value for $\bar{\lambda}$ is used as in the wind power preference

penalty (e.g., $\bar{\lambda} = 0.001$), small enough to avoid distorting economic decisions, yet sufficient to guide the model away from nonphysical solutions. The resulting electricity sales revenue is calculated as:

$$S_{DA} = \sum_{t \in \mathcal{T}} (\lambda_t - \bar{\lambda}) \cdot p_t^{M, \text{sell}}. \quad (5.7)$$

5.1.6. Total objective function

Based on the revenue and cost components defined above, the full objective function can now be stated explicitly. The goal is to determine the control strategy that maximises the total annual profit \mathcal{R} , which includes electricity and hydrogen revenues, variable operational costs, and fixed annual costs. The complete objective function is given by:

$$\begin{aligned} \max \mathcal{R} := & \underbrace{S_{H_2} + S_{DA}}_{\text{Revenue (5.8)}} - \underbrace{F_{\text{power}} - F_{\text{shut-off}}}_{\text{Variable Costs (5.9)}} - \underbrace{F_{\text{Capex}} - F_{\text{Opex}}}_{\text{Fixed Costs (5.10)}} \\ = & \sum_{t \in \mathcal{T}} \left(\pi_{H_2} h_t + (\lambda_t - \bar{\lambda}) p_t^{M, \text{sell}} \right) \end{aligned} \quad (5.8)$$

$$- \sum_{t \in \mathcal{T}} \left(\lambda_t p_t^{M, \text{buy}} + \pi^{\text{PPA}} p_t^{W, \text{PPA}} + C^{\text{bal}} p_t^{W, \text{overcap}} + \bar{\lambda} (p_t^{M \rightarrow E} + p_t^{B \rightarrow E}) + C^{\text{su}} z_t^{\text{su count}} \right) \quad (5.9)$$

$$- \xi^{\text{annual}} \left(C_e^{\text{capex}} K^{\text{elec}} + C_b^{\text{capex}} K^{\text{Batt}} \right) - \delta_e^{\text{opex}} C_e^{\text{capex}} K^{\text{Elec}} - \delta_b^{\text{opex}} C_b^{\text{capex}} K^{\text{Batt}}. \quad (5.10)$$

This objective function forms the core of the optimisation model. In the following sections, we introduce the decision and state variables in more detail and formulate the system constraints that ensure feasibility with respect to technological and contractual limitations.

5.2. Variables

We distinguish between several categories of variables and parameters in the model: decision variables, state variables, power flow variables and fixed parameters.

Decision variables

The decision variables represent the controllable actions of the operator at each time step. The goal of the optimisation is to determine the optimal dispatch of these variables such that the overall objective function, total system profit, is maximised. These include decisions on how to allocate electricity from both the day-ahead market and the wind PPA to various components of the system: the electrolyser, the battery, or directly to the market.

Electricity from the market can either be sent to the electrolyser ($p_t^{M \rightarrow E}$) or used to charge the battery ($p_t^{M \rightarrow B}$). Likewise, electricity from the wind PPA ($p_t^{W, \text{PPA}}$) can be allocated to the electrolyser ($p_t^{W \rightarrow E}$), stored in the battery ($p_t^{W \rightarrow B}$), or sold directly to the market ($p_t^{W \rightarrow M}$). Any portion of the wind power that is not used incurs a balancing penalty and is captured by the state variable $p_t^{W, \text{overcap}}$.

The battery can be discharged to supply electricity to the electrolyser ($p_t^{B \rightarrow E}$) or to sell electricity back to the market ($p_t^{B \rightarrow M}$). In addition, the operator decides whether to activate the electrolyser at each time step. This is modeled with the binary variable z_t^{state} , where $z_t^{\text{state}} = 1$ indicates that the electrolyser is operational and $z_t^{\text{state}} = 0$ means it is turned off. Each time the electrolyser is shut down, the auxiliary binary variable $z_t^{\text{su count}}$ is triggered and set to 1, which accounts for shutdown-related degradation costs in the objective function.

State variables

State variables describe the internal condition of the system at each time step and evolve based on either exogenous information or the outcomes of past and current control decisions.

The battery state of charge SOC_t tracks the amount of energy stored in the battery. It is updated dynamically based on the charging and discharging decisions made at the current time step, as well as the SOC of the previous time step. Similarly, the hydrogen production h_t is determined endogenously, based on the total electricity input to the electrolyser and its efficiency η^e .

The binary variable $z_t^{\text{su count}}$, which was introduced in relation to electrolyser control, is considered a state variable. It is not directly controlled but depends on the transition in the decision variable z_t^{state} . Specifically, $z_t^{\text{su count}} = 1$ if the electrolyser is turned off at time t after being on at time $t-1$, thereby tracking shutdown events over time.

The utilisation rate of the electrolyser, denoted U_t , expresses the operating level relative to its maximum capacity. It is calculated as the ratio of the electricity input to the electrolyser at time t to its nominal installed capacity K^{Elec} , and is used to enforce minimum operating thresholds.

Lastly, the available wind power from the Power Purchase Agreement is given by $p_t^{W, \text{PPA}} = \alpha_t \cdot K^{\text{PPA}}$, where the wind load factor $\alpha_t \in [0, 1]$ is exogenous and varies hourly based on weather conditions. This makes $p_t^{W, \text{PPA}}$ an exogenously updated state variable.

Power flows

To model the internal energy dynamics of the system in a transparent and flexible manner, we adopt a flow-based formulation, as commonly used in energy systems modeling [89]. In this approach, electricity transfers between system components such as the market, battery and electrolyser are explicitly represented through distinct power flow variables. While this increases the number of control variables, which is typically undesirable in optimisation problems, it allows for more granular control, clearer constraint formulation, and ensures physical consistency across all subsystems.

The power flows capture key operational decisions: electricity purchased from the market ($p_t^{M, \text{buy}}$), electricity sold to the market ($p_t^{M, \text{sell}}$), charging ($p_t^{B, \text{charge}}$) and discharging ($p_t^{B, \text{discharge}}$) of the battery, and the electricity input to the electrolyser for hydrogen production ($p_t^{E, \text{in}}$). The allocation of wind PPA energy is modeled via flows from wind to each subsystem: to the market ($p_t^{W \rightarrow M}$), to the battery ($p_t^{W \rightarrow B}$), and to the electrolyser ($p_t^{W \rightarrow E}$).

Unutilised wind power under the PPA contract is accounted for by the variable $p_t^{W, \text{overcap}}$, which captures the excess generation beyond the sum of wind allocations to other components. The available wind power itself is an exogenous state variable given by $p_t^{W, \text{PPA}} = \alpha_t \cdot K^{\text{PPA}}$, where α_t represents the wind load factor at time t .

5.3. Constraints

The optimisation model incorporates a set of constraints that ensure the operational strategy complies with technical limitations of the electrolyser, battery system, and grid interactions, as well as economic considerations such as market transactions and wind PPA agreements.

Power flow constraints

The model uses a flow-based formulation to explicitly define how electricity moves between system components at each time step. The total power consumed by the electrolyser is the sum of electricity flows from the market, battery, and wind PPA. Likewise, charging, discharging, and market transactions are decomposed into their respective power sources. The complete set of flow definitions is:

$$p_t^{E, \text{in}} = p_t^{M \rightarrow E} + p_t^{B \rightarrow E} + p_t^{W \rightarrow E} \quad \forall t \in \mathcal{T}, \quad (5.11)$$

$$p_t^{M, \text{buy}} = p_t^{M \rightarrow B} + p_t^{M \rightarrow E} \quad \forall t \in \mathcal{T}, \quad (5.12)$$

$$p_t^{M, \text{sell}} = p_t^{B \rightarrow M} + p_t^{W \rightarrow M} \quad \forall t \in \mathcal{T}, \quad (5.13)$$

$$p_t^{B, \text{charge}} = p_t^{M \rightarrow B} + p_t^{W \rightarrow B} \quad \forall t \in \mathcal{T}, \quad (5.14)$$

$$p_t^{B, \text{discharge}} = p_t^{B \rightarrow E} + p_t^{B \rightarrow M} \quad \forall t \in \mathcal{T}, \quad (5.15)$$

$$p_t^{W, \text{overcap}} = p_t^{W, \text{PPA}} - (p_t^{W \rightarrow E} + p_t^{W \rightarrow B} + p_t^{W \rightarrow M}) \quad \forall t \in \mathcal{T}. \quad (5.16)$$

Wind power constraints

Wind energy is sourced through a long-term Power Purchase Agreement (PPA), which provides a fixed contracted capacity K^{PPA} . The actual available power at each time step is driven by the wind load factor $\alpha_t \in [0, 1]$. The resulting wind power available to the system is defined by:

$$p_t^{W, \text{PPA}} = \alpha_t \cdot K^{\text{PPA}} \quad \forall t \in \mathcal{T}. \quad (5.17)$$

This available wind power can be allocated to the electrolyser, battery, or directly sold to the market. The total allocation must not exceed the available generation:

$$p_t^{W \rightarrow E} + p_t^{W \rightarrow B} + p_t^{W \rightarrow M} \leq p_t^{W, \text{PPA}} \quad \forall t \in \mathcal{T}. \quad (5.18)$$

Any wind energy not allocated is considered overcapacity and is stored in $p_t^{W, \text{overcap}}$ (Eq. 5.16), which incurs a penalty cost.

Battery constraints

The battery is modeled as an energy storage unit with finite energy and power capacity. Its operation is governed by physical limitations related to charging and discharging rates, efficiency losses, and the evolution of its state of charge (SOC) [1, 27, 28, 53].

The total usable energy capacity of the battery is denoted by K^{Batt} (in MWh). The maximum charging and discharging power are denoted by $P^{C, \text{max}}$ and $P^{D, \text{max}}$, respectively. These power limits are directly related to the battery's duration t^{duration} , which indicates the number of hours the battery can charge or discharge at maximum power until fully (dis)charged:

$$P^{C, \text{max}} = P^{D, \text{max}} = \frac{K^{\text{Batt}}}{t^{\text{duration}}}.$$

For example, a battery with a total energy capacity of 60 MWh and a duration of 4 hours has:

$$P^{C, \text{max}} = P^{D, \text{max}} = \frac{60}{4} = 15 \text{ MW},$$

which means that in one time step (e.g., one hour), the battery can charge or discharge up to 25% of its full capacity. More generally, the power limit can be interpreted as the inverse of the energy-to-power ratio. A value of $C^{\text{max}} = 0.25$ (expressed as a fraction of the energy capacity) implies that four time steps are required to fully charge the battery (neglecting efficiency losses), assuming maximum allowable charging power is used in each period.

The battery must operate within these bounds, and cannot be charged and discharged simultaneously. These constraints are implemented as follows:

$$p_t^{B, \text{charge}} \cdot p_t^{B, \text{discharge}} = 0 \quad \forall t \in \mathcal{T}, \quad (5.19)$$

$$0 \leq p_t^{B, \text{charge}} \leq P^{C, \text{max}} \quad \forall t \in \mathcal{T}, \quad (5.20)$$

$$0 \leq p_t^{B, \text{discharge}} \leq P^{D, \text{max}} \quad \forall t \in \mathcal{T}. \quad (5.21)$$

The state of charge $\text{SOC}_t \in [0, 1]$ represents the relative energy level of the battery at time t , expressed as a fraction of K^{Batt} . For instance, $\text{SOC}_t = 0.8$ implies the battery holds 80% of its total energy capacity, or $0.8 \cdot K^{\text{Batt}}$ MWh. The SOC evolves based on the net charging and discharging decisions, adjusted for charging efficiency $\eta^C \in (0, 1]$ and discharging efficiency $\eta^D \in (0, 1]$:

$$\text{SOC}_t = \text{SOC}_{t-1} + \frac{p_t^{B, \text{charge}} \cdot \eta^C - \frac{p_t^{B, \text{discharge}}}{\eta^D}}{K^{\text{Batt}}} \quad \forall t \in \mathcal{T}, \quad (5.22)$$

$$\text{SOC}^{\text{min}} \leq \text{SOC}_t \leq \text{SOC}^{\text{max}} \quad \forall t \in \mathcal{T}. \quad (5.23)$$

To avoid unrealistic end-of-horizon behaviour, such as fully depleting or overcharging the battery in the final hour, we enforce that the state of charge at the end of the optimisation horizon matches its initial value: $\text{SOC}_0 = \text{SOC}_T = \text{SOC}_{\text{init}}$. In this model, both the initial and terminal SOC are fixed at the minimum allowable SOC level, which represents 20% of the total battery capacity. This is defined as $\text{SOC}_{\text{init}} = \text{SOC}_{\text{min}} = 0.2$.

Electrolyser constraints

The electrolyser converts electricity into hydrogen, operating between a predefined minimum and maximum load. In this model, we assume constant efficiency, denoted by $\eta^e \in (0, 1]$, meaning that the conversion rate from electricity to hydrogen is fixed across all power levels. Although in reality electrolyser efficiency depends on factors such as temperature, pressure, and input power, a constant efficiency is chosen here to keep the operational model simple. The focus of this thesis is not to model the electrochemical conversion process in detail, but rather to capture system-level dispatch and interaction dynamics across energy sources and storage. Hence the hydrogen production at each time step t is given by the total power input to the electrolyser, multiplied by its efficiency:

$$h_t = p_t^{E, \text{in}} \cdot \eta^e, \quad \forall t \in \mathcal{T}. \quad (5.24)$$

The electrolyser can operate at a level up to its installed capacity K^{Elec} , which defines the upper bound of its operating range. To avoid inefficient operation and accelerated degradation, a minimum power threshold P^{min} is enforced whenever the electrolyser is switched on. This threshold represents a minimum loading condition (e.g., $P^{\text{min}} = 0.3 \cdot K^{\text{Elec}}$) below which operation is considered technically suboptimal. These operational bounds are implemented using the binary decision variable z_t^{state} , which indicates whether the electrolyser is active at time t :

$$z_t^{\text{state}} \cdot P^{\text{min}} \leq p_t^{E, \text{in}} \leq z_t^{\text{state}} \cdot K^{\text{Elec}}, \quad \forall t \in \mathcal{T}. \quad (5.25)$$

In addition to the operational limits, the electrolyser must also undergo scheduled maintenance to ensure reliable performance and prevent excessive wear. This model incorporates a mandatory maintenance period of M^{time} hours, during which the electrolyser is not available for operation.

To enforce this, a maintenance constraint is included, ensuring that the total operating hours of the electrolyser over the planning horizon do not exceed the available operational time after accounting for maintenance:

$$\sum_{t \in \mathcal{T}} z_t^{\text{state}} \leq M^{\text{time}}. \quad (5.26)$$

To avoid artificial switching in the first hour, we fix the initial on/off state of the electrolyser. In this model, the electrolyser is assumed to be operational at the beginning of the planning horizon:

$$z_t^{\text{state}} = 1, \quad \text{for } t = 0. \quad (5.27)$$

Frequent on-off cycling of the electrolyser can accelerate wear and reduce its technical lifetime. To capture this effect, we introduce a state binary variable $z_t^{\text{su count}}$, which is triggered whenever the electrolyser shuts down between consecutive time steps:

$$z_t^{\text{su count}} \geq z_{t-1}^{\text{state}} - z_t^{\text{state}}, \quad \forall t \in \mathcal{T} \setminus \{0\}. \quad (5.28)$$

To limit degradation, the total number of shutdowns over the year is restricted by an upper bound N^{off} , enforced as:

$$\sum_{t \in \mathcal{T}} z_t^{\text{su count}} \leq N^{\text{off}}. \quad (5.29)$$

This constraint implicitly discourages unnecessary shutdowns, and by enforcing a minimum on-time between cycles, it also helps avoid rapid cycling behavior that would otherwise require warm-up/cool-down modeling.

Grid connection constraints

Electricity trading with the external grid is constrained by the physical capacity of the site's grid connection. This includes both electricity bought from the market and electricity sold back to the grid. Let K^{grid} denote the grid connection capacity (in MW), which applies as a symmetric limit to both import and export:

$$p_t^{M, \text{buy}} \leq K^{\text{grid}} \quad \forall t \in \mathcal{T}, \quad (5.30)$$

$$p_t^{M, \text{sell}} \leq K^{\text{grid}} \quad \forall t \in \mathcal{T}. \quad (5.31)$$

In addition to this limit, we impose a mutual exclusivity condition to ensure that electricity cannot be bought and sold simultaneously:

$$p_t^{M, \text{buy}} \cdot p_t^{M, \text{sell}} = 0 \quad \forall t \in \mathcal{T}. \quad (5.32)$$

This constraint ensures that the model must choose at each time step whether to import or export power, but not both. It prevents trading behavior driven purely by numerical structure in the objective function, and reflects the physical reality of using a single cable to interface with the electricity grid.

Hydrogen demand constraints

In commercial settings, hydrogen offtake is typically governed by long-term supply contracts such as a Hydrogen Purchase Agreement (HPA). These contracts often stipulate a minimum annual volume of hydrogen that the producer must deliver to the offtaker. To reflect this requirement, the model includes a minimum annual production constraint:

$$\sum_{t \in \mathcal{T}} h_t \geq D^{\text{target}}. \quad (5.33)$$

Here, D^{target} represents the minimum volume of hydrogen (in MWh thermal, based on LHV) contractually promised to the offtaker over the planning horizon. This ensures that the producer can fulfill baseline contractual obligations.

However, the model is not restricted to this minimum production level. It allows for additional hydrogen production whenever it is economically advantageous, under the assumption that the offtaker has sufficient storage capacity to accommodate any surplus hydrogen and is willing to purchase this additional volume at the same fixed price specified in the HPA. Moreover, the model assumes flexibility in delivery timing, meaning the producer is not bound to a fixed hourly delivery profile or baseload commitment. This flexibility enables the system to optimise hydrogen production in response to dynamic electricity prices and operational conditions, maximising profitability.

5.4. Model characteristics and problem classification

The model developed in this chapter is fundamentally designed as a Mixed-Integer Linear Program (MILP). This classification arises from the combination of continuous variables, which represent power flows within the system, and binary decision variables, such as z_t^{state} for the electrolyser's operational state. The use of binary variables allows the model to capture logical decisions (e.g., whether the electrolyser is active), while continuous variables describe energy quantities and power exchanges. This mixed-integer structure is essential for accurately representing the operation of a hybrid energy system.

A linear formulation is chosen because it offers a balanced trade-off between model expressiveness and computational efficiency. Linear models can be efficiently solved using commercial optimisation solvers such as Gurobi, making them suitable for large-scale, time-dependent problems like the one in this thesis. However, despite our objective to maintain a fully linear model, the initial formulation includes two non-linear constraints that must be addressed. Specifically, the model currently contains two bilinear constraints that introduce non-convexities:

$$p_t^{M, \text{buy}} \cdot p_t^{M, \text{sell}} = 0 \quad \forall t \in \mathcal{T}. \quad (5.34)$$

$$p_t^{B, \text{charge}} \cdot p_t^{B, \text{discharge}} = 0 \quad \forall t \in \mathcal{T}. \quad (5.35)$$

Although modern solvers such as Gurobi support certain classes of non-convex optimisation, MILPs are generally more robust and computationally efficient. To this end, these constraints are reformulated using standard linearisation techniques involving auxiliary binary variables.

5.4.1. Linear reformulation

To enable MILP solvability, two binary variables are introduced. The variable $z_t^{\text{grid}} \in \{0, 1\}$ governs market interaction: when set to 1, electricity may be purchased from the grid; when set to 0, only selling is allowed. Similarly, $z_t^{\text{battery}} \in \{0, 1\}$ controls battery operation: a value of 1 permits charging, while 0 permits discharging.

The mutual exclusivity of buying and selling is enforced as:

$$p_t^{M, \text{buy}} \leq K^{\text{grid}} \cdot z_t^{\text{grid}}, \quad (5.36)$$

$$p_t^{M, \text{sell}} \leq K^{\text{grid}} \cdot (1 - z_t^{\text{grid}}). \quad (5.37)$$

These linear constraints replace the original non-linear constraints (Eq. 5.30- 5.32), ensuring that electricity can either be bought or sold at each time step, but not both. This reformulation maintains the model's logical structure while preserving its computational efficiency.

Similarly, the linear constraints ensuring exclusive battery charging or discharging operation and thus replacing Eq. 5.19 - 5.21 are given by:

$$p_t^{B, \text{charge}} \leq K^{\text{Batt}} \cdot z_t^{\text{battery}}, \quad (5.38)$$

$$p_t^{B, \text{discharge}} \leq K^{\text{Batt}} \cdot (1 - z_t^{\text{battery}}). \quad (5.39)$$

These linear formulations preserve the logical structure of the original bilinear constraints while allowing the model to remain within the MILP framework. The resulting model is both expressive and computationally efficient, enabling solution with commercial optimisation solvers such as Gurobi [33]. The total MILP formulation is presented in Section 5.5.

5.4.2. Look-ahead period and uncertainty incorporation

The initial version of the optimisation model in this thesis is based on the assumption of *perfect foresight*, meaning that all input parameters, including wind availability α_t and electricity market prices λ_t are assumed to be known with complete certainty for the entire planning horizon. This setup allows for optimising the dispatch strategy with maximum precision, making it possible to perfectly plan when to produce hydrogen, when to purchase electricity, and when to sell excess power to the market. As a result, the perfect foresight model effectively serves as an upper bound on system profitability, representing the highest achievable profit under ideal conditions.

While the perfect foresight model offers computational simplicity and clarity, it represents an idealised scenario that does not reflect the realities of energy system operations. In the real-world, operators don't have perfect foresight an entire year ahead but instead must rely on forecasts to guide their scheduling and decision-making. These forecasts are inherently uncertain, and their accuracy decreases with longer time horizons. As a result, operators cannot achieve the theoretical upper bound of profitability defined by the perfect foresight model. Instead, their decisions are subject to the uncertainty of future conditions, which can lead to suboptimal outcomes.

This does not mean that the perfect foresight model is without value. On the contrary, it provides critical insights by establishing a benchmark for maximum possible profit. By comparing the performance of realistic models, which account for forecast uncertainty, against the perfect foresight scenario, we can assess how much profit is lost due to uncertainty. Such comparisons offer a clear understanding of the economic impact of imperfect information.

To investigate the impact of different look-ahead periods and the incorporation of uncertainty in input variable forecasts, this thesis explores two extensions to the modeling approach:

1. **Rolling horizon optimisation (RHO):** A deterministic solution method that limits the foresight window to a shorter period (e.g., one week) and re-optimises periodically (e.g., every 24 hours). This approach assumes no knowledge beyond the foresight window, but within the window, perfect foresight is maintained.
2. **Scenario-based two-stage stochastic optimisation:** A probabilistic formulation where future prices and wind profiles are represented by multiple scenarios. First-stage decisions are made before uncertainty is revealed, while second-stage decisions are scenario-dependent. This approach explicitly models forecast uncertainty and enables risk-aware planning.

These methods provide a more realistic representation of operational decision-making, allowing for a direct comparison between the idealised, perfectly informed scenario and the more practical, uncertainty-aware approaches. In the literature review Chapter 3 the theoretical background of these methods is explained

in detail. These modeling extensions still build upon the core structure presented in this chapter but require additional formulations and algorithmic approaches. Table 5.1 offers an overview of the three model types compared in this thesis. Chapter 6 elaborates on the implementation extensions, including the adapted mathematical formulation and the tools used to solve them.

Table 5.1: Overview of modeling solution methods compared in this thesis.

Name	Model type	Foresight window	Forecast uncertainty	Solution approach
DO	Deterministic optimisation with Perfect Foresight	1 year	None	Single MILP solve
RHO	Deterministic optimisation with Rolling Horizon	1 week	None	Receding horizon MILP
SORH	Two-Stage Stochastic optimisation with Rolling Horizon	1 week	None (Day 1) + Multiple random scenarios (Day 2-7)	Scenario-based + receding horizon MILP

5.5. Final formulation MILP

$$\begin{aligned}
\max \quad & \sum_{t \in \mathcal{T}} (\pi_{H_2} h_t + (\lambda_t - \bar{\lambda}) p_t^{M, \text{sell}}) \\
& - \sum_{t \in \mathcal{T}} (\lambda_t p_t^{M, \text{buy}} + \pi^{\text{PPA}} p_t^{W, \text{PPA}} + C^{\text{bal}} p_t^{W, \text{overcap}} + \bar{\lambda} (p_t^{M \rightarrow E} + p_t^{B \rightarrow E}) + C^{\text{su}} z_t^{\text{su count}}) \\
& - \xi^{\text{annual}} (C_e^{\text{capex}} K^{\text{Elec}} + C_b^{\text{capex}} K^{\text{Batt}}) \\
& - \delta_e^{\text{opex}} C_e^{\text{capex}} K^{\text{Elec}} - \delta_b^{\text{opex}} C_b^{\text{capex}} K^{\text{Batt}} \\
\text{s.t.} \quad & p_t^{E, \text{in}} = p_t^{M \rightarrow E} + p_t^{B \rightarrow E} + p_t^{W \rightarrow E} \quad \forall t \in \mathcal{T} \\
& p_t^{M, \text{buy}} = p_t^{M \rightarrow B} + p_t^{M \rightarrow E} \quad \forall t \in \mathcal{T} \\
& p_t^{M, \text{sell}} = p_t^{B \rightarrow M} + p_t^{W \rightarrow M} \quad \forall t \in \mathcal{T} \\
& p_t^{B, \text{charge}} = p_t^{M \rightarrow B} + p_t^{W \rightarrow B} \quad \forall t \in \mathcal{T} \\
& p_t^{B, \text{discharge}} = p_t^{B \rightarrow E} + p_t^{B \rightarrow M} \quad \forall t \in \mathcal{T} \\
& p_t^{W, \text{overcap}} = p_t^{W, \text{PPA}} - (p_t^{W \rightarrow E} + p_t^{W \rightarrow B} + p_t^{W \rightarrow M}) \quad \forall t \in \mathcal{T} \\
& p_t^{W, \text{PPA}} = \alpha_t \cdot K^{\text{PPA}} \quad \forall t \in \mathcal{T} \\
& p_t^{W \rightarrow E} + p_t^{W \rightarrow B} + p_t^{W \rightarrow M} \leq p_t^{W, \text{PPA}} \quad \forall t \in \mathcal{T} \\
& p_t^{B, \text{charge}} \leq K^{\text{Batt}} \cdot z_t^{\text{battery}} \quad \forall t \in \mathcal{T} \\
& p_t^{B, \text{discharge}} \leq K^{\text{Batt}} \cdot (1 - z_t^{\text{battery}}) \quad \forall t \in \mathcal{T} \\
& 0 \leq p_t^{B, \text{charge}} \leq P^{C, \text{max}} \quad \forall t \in \mathcal{T} \\
& 0 \leq p_t^{B, \text{discharge}} \leq P^{D, \text{max}} \quad \forall t \in \mathcal{T} \\
& \text{SOC}_t = \text{SOC}_{t-1} + \frac{p_t^{B, \text{charge}} \cdot \eta^C - p_t^{B, \text{discharge}}}{K^{\text{Batt}}} \quad \forall t \in \mathcal{T} \\
& \text{SOC}^{\min} \leq \text{SOC}_t \leq \text{SOC}^{\max} \quad \forall t \in \mathcal{T} \\
& \text{SOC}_0 = \text{SOC}^{\text{init}}, \quad \text{SOC}_T = \text{SOC}^{\text{init}} \\
& h_t = p_t^{E, \text{in}} \cdot \eta^e \quad \forall t \in \mathcal{T} \\
& z_t^{\text{state}} \cdot P^{\min} \leq p_t^{E, \text{in}} \leq z_t^{\text{state}} \cdot K^{\text{elec}} \quad \forall t \in \mathcal{T} \\
& \sum_{t \in \mathcal{T}} z_t^{\text{state}} \leq M^{\text{time}} \\
& z_0^{\text{state}} = 1 \\
& z_t^{\text{su count}} \geq z_{t-1}^{\text{state}} - z_t^{\text{state}} \quad \forall t \in \mathcal{T} \setminus \{0\} \\
& \sum_{t \in \mathcal{T}} z_t^{\text{su count}} \leq N^{\text{off}} \\
& p_t^{M, \text{buy}} \leq K^{\text{grid}} \cdot z_t^{\text{grid}} \quad \forall t \in \mathcal{T} \\
& p_t^{M, \text{sell}} \leq K^{\text{grid}} \cdot (1 - z_t^{\text{grid}}) \quad \forall t \in \mathcal{T} \\
& \sum_{t \in \mathcal{T}} h_t \geq D^{\text{target}} \\
\text{Variable types:} \quad & z_t^{\text{state}}, z_t^{\text{su count}}, z_t^{\text{grid}}, z_t^{\text{battery}} \in \{0, 1\}, \quad \forall t \in \mathcal{T} \\
& \text{All other variables} \in \mathbb{R}_{\geq 0}
\end{aligned}$$

6

Advanced Solution Methods

This chapter presents advanced solution methods developed to improve the realism and performance of the optimisation framework introduced earlier. The baseline model assumes perfect foresight over the entire optimisation horizon, which is often unrealistic in practical settings. To address this, we introduce two extensions that incorporate limited foresight and uncertainty in a structured manner. Section 6.1 describes the rolling horizon optimisation (RHO) formulation and solution algorithm, which extends the original model by solving sequential subproblems over a moving time window. Section 6.2 presents the two-stage stochastic optimisation approach with a rolling horizon (SORH), which builds upon the RHO formulation by incorporating uncertainty explicitly. In this formulation, a set of future scenarios is considered to optimise short-term decisions against multiple possible realisations of electricity prices and wind availability. Finally, Section 6.3 explains the scenario generation method used to produce representative samples of day-ahead electricity prices and wind generation load factors.

6.1. Rolling horizon optimisation formulation

Rolling Horizon optimisation (RHO) is a heuristic technique that extends the deterministic MILP model with limited foresight and iterative reoptimisation. Instead of solving the full-year dispatch problem in one shot with perfect foresight, RHO simulates real-world decision-making more realistically by solving a sequence of smaller sub-problems over shorter planning horizons. This section describes how the RHO method modifies the base MILP and outlines the structure, logic, and implementation details of the rolling solution algorithm.

For the RHO formulation we modify the deterministic MILP from Chapter 5 by restricting foresight to a shorter forecast horizon. The model is solved sequentially over a moving time window, and only the decisions for the immediate planning period are retained in each solve.

6.1.1. Temporal structure and solution logic

The full modeling period $\mathcal{T} = \{1, \dots, T\}$ (e.g. $T = 8760$ hours) is divided into overlapping planning windows of length T' (e.g. 168 hours = 1 week). We let $\mathcal{I} = \{1, 1 + \Delta i, 1 + 2\Delta i, \dots, T\}$ be the set of solution steps where $\Delta i = T'$ and the total number of solution steps is $I = \frac{T}{T'}$. At each step i , a sub-problem is solved over the planning period $\mathcal{H}^i = \{i, \dots, i + H\}$, with only the first T' hours of the resulting control trajectory implemented. After this step, the horizon is rolled forward by $\Delta i = T'$, and the process repeats until the end of \mathcal{T} is reached.

To summarise we let:

- Δt : time resolution (1 hour),
- T' : planning period (e.g., 24 hours),
- H : forecast horizon (e.g., 168 hours).
- At each step $\mathcal{I} = \{1, 1 + \Delta i, 1 + 2\Delta i, \dots, T\}$, we define:

$$\mathcal{T}^i = \{i, i + 1, \dots, i + T'\}, \quad \mathcal{H}^i = \{i, i + 1, \dots, i + H\}.$$

The MILP from Chapter 5 is solved over each window H^i . At each iteration $i \in \mathcal{I}$, the optimisation solves:

$$\max \mathcal{R}^i = \sum_{t \in \mathcal{H}^i} \pi_{H_2} h_t + (\lambda_t - \bar{\lambda}) p_t^{M, \text{sell}} - \left(\lambda_t p_t^{M, \text{buy}} + \pi^{\text{PPA}} p_t^{W, \text{PPA}} + C^{\text{bal}} p_t^{W, \text{overcap}} + \bar{\lambda} (p_t^{M \rightarrow E} + p_t^{B \rightarrow E}) + C^{\text{su}} z_t^{\text{su count}} \right). \quad (6.1)$$

6.1.2. Modified constraints

Most constraints are unchanged and applied for $t \in \mathcal{H}^i$, but key time-coupled constraints are modified to ensure correct state tracking and smooth transitions between planning steps.

The iterative procedure used in this thesis to implement the rolling horizon optimisation is summarised in Algorithm 1. At each step, the optimisation model is solved over a fixed planning period using the most recent forecast information. Only the decisions for the current planning period are retained, while the remaining horizon is re-optimised in the next iteration based on updated initial states.

Initial conditions and transition between solution steps

To ensure consistency between successive optimisation windows, initial conditions must be passed from one solution step to the next. Specifically the variables for the electrolyser on/off state, z^{state} , and the state of charge of the battery SOC . For each iteration i , the system is initialised using terminal values from the previous solve:

$$z_i^{\text{state}} = z^{\text{start}}, \quad SOC_i = SOC^{\text{start}}.$$

To model continuity and avoid artificial decisions at sub-horizon boundaries, the following adjustments are made to constraints:

For $t = i$:

$$z_t^{\text{su count}} \geq z^{\text{start}} - z_t^{\text{state}},$$

$$SOC_t = SOC^{\text{start}} - \frac{\Delta_t^{\text{charge}}}{K^{\text{Batt}}},$$

and for $t > i$:

$$z_t^{\text{su count}} \geq z_{t-1}^{\text{state}} - z_t^{\text{state}},$$

$$SOC_t = SOC_{t-1} + \frac{\Delta_t^{\text{charge}}}{K^{\text{Batt}}},$$

where $\Delta_t^{\text{charge}} = p_t^{B, \text{charge}} \cdot \eta^C - \frac{p_t^{B, \text{discharge}}}{\eta^D}$. At the end of each step we update the current *start* value of z^{state} and SOC to be the last value of the previous solution step:

$$z^{\text{start}} := z_{i+T'}^{\text{state}}, \quad SOC^{\text{start}} := SOC_{i+T'}.$$

This ensures that each solve starts from the physical end state of the previous solve.

Hydrogen production target and number of shut-offs per horizon

The hydrogen production constraint is adjusted proportionally to the horizon size, to ensure the annual target is met over all subproblems. To approximate the annual hydrogen production target across rolling steps, we enforce:

$$\sum_{t \in \mathcal{H}^i} h_t \geq D_{\text{horizon}}^{\text{target}} = \frac{D_{\text{annual}}^{\text{target}}}{T} \cdot T'.$$

Similarly we adjust the number of allowed shut-offs annually to a number more proportioned to the horizon period.

$$\sum_{t \in \mathcal{H}^i} z_t^{\text{off count}} \leq N_{\text{horizon}}^{\text{Off}}.$$

Maintenance period adjustment

In the modified RHO model, we introduce a pre-defined maintenance period to account for planned downtime of the electrolyser. This adjustment is necessary to ensure that the electrolyser undergoes scheduled maintenance without the optimisation model exploiting the flexibility to avoid maintenance entirely.

To account for the fact that operators have some control over when to schedule maintenance within the year, we assume that they make a rational choice by scheduling the maintenance during the period of highest average electricity prices. This approach minimises the economic impact of the downtime, as it avoids operating the electrolyser when production would be most expensive. Specifically, the maintenance period is defined as a consecutive block of hours, M^{time} , during which the electrolyser is forced to shut down.

The maintenance window is identified using perfect foresight, which is not entirely realistic in practice. Operators are assumed to perfectly predict the highest price period of the year, which allows the model to minimise the impact of the maintenance period. Mathematically, this window is determined using the following:

$$\text{Start} = \operatorname{argmax}_t \left(\frac{1}{M^{\text{time}}} \sum_{t'=t}^{t+M^{\text{time}}} P_{t'} \right).$$

The set of maintenance hours \mathcal{M} is then defined as the consecutive range of hours starting from this identified period: $\mathcal{M} = \{t : \text{Start} \leq t < \text{Start} + M^{\text{time}}\}$. To ensure that the maintenance period is strictly adhered to, the following constraint is added to the RHO model. For each hour within the maintenance window: $z_t^{\text{state}} = 0, \forall t \in \mathcal{M}$. These adjustments ensure that the electrolyser cannot operate, and thus cannot produce hydrogen, during the pre-defined maintenance period. The mathematical formulation of these constraints is expressed as:

$$\sum_{t \in \mathcal{M}} z_t^{\text{state}} = 0.$$

This adjustment ensures that the electrolyser's operating hours are reduced by the predefined maintenance period, effectively lowering the total available operating hours over the year. It also prevents the optimiser from strategically avoiding the maintenance period, maintaining the realism of the model.

Algorithm 1 Rolling horizon optimisation procedure

- 1: **Input:** Forecast data (wind, prices), model parameters, initial states $z^{\text{start}}, \text{SOC}^{\text{start}}$
 - 2: **Set:** Planning period T' , forecast horizon H , Modeling period T , time resolution $\Delta t = 1$
 - 3: **for** each solution step $i = 1, 1 + T', 1 + 2T', \dots$ until T **do**
 - 4: Define planning window $T'^i = \{i, \dots, i + T'\}$, forecast window $H^i = \{i, \dots, i + H\}$
 - 5: **initialise:** $z_i^{\text{state}} := z^{\text{start}}, \text{SOC}_i := \text{SOC}^{\text{start}}$
 - 6: Solve MILP over T'^i using forecasts on H^i
 - 7: Store decisions for $t \in T'^i$
 - 8: **Update:** $z^{\text{start}} := z_{i+T'}^{\text{state}}, \text{SOC}^{\text{start}} := \text{SOC}_{i+T'}$
 - 9: **end for**
 - 10: **Output:** Combined decisions over full horizon
-

6.2. Two-stage stochastic optimisation with rolling horizon

To explicitly capture forecast uncertainty in electricity prices and wind availability, we extend the deterministic MILP into a two-stage stochastic optimisation framework. In this formulation, decisions for the immediate planning period (first stage) are made with perfect knowledge of electricity prices and wind generation. In contrast, future decisions (second stage) are optimised across multiple forecast scenarios, each representing a possible realisation of future conditions. This structure enables risk-aware planning by balancing expected performance across these scenarios, ensuring a robust operational strategy under uncertainty.

The two-stage stochastic formulation builds directly upon the MILP from Chapter 5 and the rolling horizon structure introduced in Section 6.1, but introduces key changes in how the objective function, variables, and constraints are defined. This section outlines these changes and presents the complete structure of the two-stage model.

6.2.1. Model structure and decision stages

At each solution step $i \in \mathcal{I}$, we again define the planning period $\mathcal{T}^i = \{i, i+1, \dots, i+T'\}$, representing the hours over which *first-stage decisions* are made (e.g., 24 hours). Secondly we define the forecast horizon $\mathcal{H}^i = \{i, i+1, \dots, i+H\}$, representing the full scenario window (e.g., 168 hours), and the scenario set \mathcal{S} , representing discrete realisations of uncertainty in wind and electricity prices.

In this two-stage framework, the distinction between first-stage and second-stage variables is made purely based on time. The first-stage corresponds to the decisions made during the first 24 hours of the forecast horizon, i.e., the planning period $\mathcal{T}^i = \{i, i+1, \dots, i+T'-1\}$. The second-stage then covers the remaining hours of the forecast horizon, defined as $\mathcal{L}_2^i = \mathcal{H}^i \setminus \mathcal{T}^i = \{i+T', i+T'+1, \dots, i+H-1\}$.

This structure is motivated by operational reality: in the day-ahead market, the operator knows the exact electricity prices for the next 24 hours and can rely on relatively accurate wind forecasts over this period. As such, decisions in the planning period are treated as deterministic (first stage), while decisions beyond that are subject to uncertainty (second stage). Second-stage decisions are modeled per scenario $s \in \mathcal{S}$, which represents different realisations of future wind and price conditions. As a result, the same types of control variables (e.g., hydrogen production, power dispatch, battery SOC) appear in both stages, but with different labels to reflect their temporal and stochastic scenario. The total number of variables is thus theoretically doubled compared to the original MILP formulation. This allows the model to account for uncertainty beyond the planning period in a structured and risk-aware manner.

6.2.2. Objective function

The optimisation maximises the expected total reward across all scenarios. The objective function at each step $i \in \mathcal{I}$ becomes:

$$\max \mathcal{R}^i = \underbrace{\sum_{t=i}^{i+T'-1} \mathcal{R}_t^{\text{stage 1}}}_{\text{first stage}} + \underbrace{\sum_{s \in \mathcal{S}} \phi_s \sum_{t=i+T'}^{i+H-1} \mathcal{R}_{t,s}^{\text{stage 2}}}_{\text{expected second stage reward}}, \quad (6.2)$$

where ϕ_s is the probability of scenario s and the reward terms \mathcal{R}_t are defined as in the deterministic model (Eq. 5.8–5.10), but evaluated per scenario where applicable.

6.2.3. Constraints

First-stage constraints

All constraints for the battery, electrolyser, and grid are identical to those in the deterministic MILP and rolling horizon setup (see Chapter 5, Section 6.1). These are defined for $t \in \mathcal{T}^i$ and are scenario-independent. Since this model distinguishes between decision stages, we now explicitly label all first-stage variables with a superscript 1. For example, the hydrogen production constraint in the first stage becomes:

$$h_t^1 = \eta^e \cdot p_t^{1,E,\text{in}}, \quad \forall t \in \mathcal{T}^i. \quad (6.3)$$

Second-stage constraints

For the remaining hours $t \in \mathcal{H}^i \setminus \mathcal{T}^i$, constraints are defined per scenario $s \in \mathcal{S}$. Each scenario has its own state evolution, dispatch variables, and power flow decisions. These variables are labeled with a superscript 2 and subscript s to indicate second-stage and scenario-dependence. For example, the SOC constraint in the second stage becomes:

$$\text{SOC}_{t,s}^2 = \text{SOC}_{t-1,s}^2 + \frac{\Delta_{t,s}^{2,\text{charge}}}{K^{\text{Batt}}}, \quad \forall t \in \mathcal{H}^i \setminus \mathcal{T}^i, \forall s \in \mathcal{S}, \quad (6.4)$$

where $\Delta_{t,s}^{2,\text{charge}} = p_{t,s}^{2,\text{B,charge}} \cdot \eta^C - \frac{p_{t,s}^{2,\text{B,discharge}}}{\eta^D}$.

Hydrogen production and shutdown constraints across stages

Since total hydrogen production and shutdown events are contractual and must be tracked across the full horizon, we combine the first- and second-stage contributions. These constraints must hold for every scenario:

$$\sum_{t=i}^{i+T'-1} h_t^1 + \sum_{t=i+T'}^{i+H-1} h_{t,s}^2 \geq D_{\text{horizon}}^{\text{target}} \quad \forall s \in \mathcal{S}, \quad (6.5)$$

$$\sum_{t=i}^{i+T'-1} z_t^{1,\text{su count}} + \sum_{t=i+T'}^{i+H-1} z_{t,s}^{2,\text{su count}} \leq N_{\text{horizon}}^{\text{off}} \quad \forall s \in \mathcal{S}. \quad (6.6)$$

Transition constraints between stages

To ensure consistency between stages, the final states of stage 1 become the initial states of stage 2 for all scenarios:

$$z_{i+T'}^{2,\text{state,start}} := z_{i+T'-1}^{1,\text{state}} \quad (6.7)$$

$$\text{SOC}_{i+T'}^{2,\text{start}} := \text{SOC}_{i+T'-1}^1. \quad (6.8)$$

These are then used to initialise the second-stage constraints.

Transition between solution steps

As in the RHO model, we update states across solution steps:

$$z^{\text{start}} := z_{i+T'}^1, \quad \text{SOC}^{\text{start}} := \text{SOC}_{i+T'}^1.$$

The overall procedure of the scenario-based two-stage optimisation with rolling horizon is summarised in Algorithm 2, which outlines the iterative solution process, including scenario generation, decision variable partitioning, and state transitions between stages and across solution steps. The next section provides a detailed explanation of how the electricity price and wind availability scenarios used in this algorithm are constructed.

Algorithm 2 Two-stage stochastic optimisation with rolling horizon

- 1: **Input:** Base forecast data (wind, prices), model parameters, initial states $z^{\text{start}}, \text{SOC}^{\text{start}}$
 - 2: **Set:** Planning period T' , forecast horizon H , modeling period T , time resolution $\Delta t = 1$
 - 3: **for** each solution step $i = 1, 1 + T', 1 + 2T', \dots$ **until** T **do**
 - 4: Define planning window $\mathcal{T}^i = \{i, \dots, i + T' - 1\}$, forecast horizon $\mathcal{H}^i = \{i, \dots, i + H - 1\}$
 - 5: **initialise:** $z_i^{1,\text{state}} := z^{\text{start}}, \quad \text{SOC}_i^1 := \text{SOC}^{\text{start}}$
 - 6: **Generate** scenarios \mathcal{S} for electricity prices and wind load factors over \mathcal{H}^i *{See Section 6.3 for details}*
 - 7: **Solve two-stage stochastic MILP:**
 - 8: First-stage decisions: x_t^1 for $t \in \mathcal{T}^i$ (common across all scenarios)
 - 9: Second-stage decisions: $x_{t,s}^2$ for $t \in \mathcal{H}^i \setminus \mathcal{T}^i, s \in \mathcal{S}$ (scenario-dependent)
 - 10: Subject to:
 - 11: First-stage constraints over \mathcal{T}^i
 e.g., $h_t^1 = \eta^e \cdot p_t^{1,\text{E},\text{in}}$
 - 12: Second-stage constraints over $\mathcal{H}^i \setminus \mathcal{T}^i$, per scenario s
 e.g., $h_{t,s}^2 = \eta^e \cdot p_{t,s}^{2,\text{E},\text{in}}$
 - 13: Transition constraints:
 - 14: $z_{i+T'}^{2,\text{state},s} = z_{i+T'-1}^{1,\text{state}}$
 - 15: $\text{SOC}_{i+T'}^{2,s} = \text{SOC}_{i+T'-1}^1$
 - 16: Combined hydrogen and shutdown constraints per scenario $s \in \mathcal{S}$:

$$\sum_{t \in \mathcal{T}^i} h_t^1 + \sum_{t \in \mathcal{H}^i \setminus \mathcal{T}^i} h_{t,s}^2 \geq D_{\text{horizon}}^{\text{target}}, \quad \sum_{t \in \mathcal{T}^i} z_t^{1,\text{su count}} + \sum_{t \in \mathcal{H}^i \setminus \mathcal{T}^i} z_{t,s}^{2,\text{su count}} \leq N_{\text{horizon}}^{\text{off}}$$
 - 17: **Store:** first-stage decisions x_t^1 for $t \in \mathcal{T}^i$
 - 18: **Update:** $z^{\text{start}} := z_{i+T'}^{1,\text{state}}, \quad \text{SOC}^{\text{start}} := \text{SOC}_{i+T'}^1$
 - 19: **end for**
 - 20: **Output:** Combined first-stage decisions over full horizon
-

6.3. Construction of day-ahead price and wind load factor scenarios

To capture uncertainty in day-ahead electricity prices and wind load factors within the scenario-based stochastic model, we employ Monte Carlo Simulation (MCS) combined with Cholesky decomposition. MCS allows for the generation of multiple realistic scenarios of future prices and wind load factors, while Cholesky decomposition ensures that these scenarios maintain the empirical characteristics of the data, including temporal autocorrelation and the typically negative correlation between wind and prices. By using MCS, we can flexibly generate any desired number of scenarios, each representing a plausible realisation of future conditions. The goal is to generate realistic and correlated trajectories of forecast errors that reflect empirical variability and dependencies. This section explains how these scenarios are constructed.

6.3.1. Forecast error modeling and scenario generation method

Rather than directly forecasting full trajectories of future electricity prices and wind availability, we adopt a forecast error modeling approach. At each solution step i , the operator is assumed to have access to a deterministic base forecast for electricity prices $\bar{\lambda}_t$ and wind load factors $\bar{\alpha}_t$ over the forecast horizon $\mathcal{H}^i = \{i_1, \dots, i_H\}$. This base forecast represents the best available estimate without accounting for uncertainty. To introduce uncertainty, the base forecast is perturbed by scenario-specific error terms, but only for the second-stage decision window $\mathcal{H}^i \setminus \mathcal{T}^i$, which excludes the initial planning period. This setup allows the model to maintain deterministic control over immediate decisions while incorporating uncertainty for future decisions. The length of the scenario generation window is defined as $H' := H - T'$, where H is the total forecast horizon, and T' is the planning period. For example, if the forecast horizon is one week (168 hours) and the planning period is one day (24 hours), then $H' = 144$ hours.

The forecast error terms for electricity prices and wind load factors are modeled as normally distributed random variables:

$$\varepsilon_{\lambda,t} \sim \mathcal{N}(0, \sigma_\lambda^2), \quad \varepsilon_{\alpha,t} \sim \mathcal{N}(0, \sigma_\alpha^2).$$

This assumption is consistent with commonly observed characteristics of forecast errors, where deviations from the forecast tend to be symmetric around zero but vary in magnitude. The standard deviations σ_λ and σ_α quantify the scale of uncertainty, representing the expected variability of forecast errors. These values are estimated from historical data to ensure that the scenarios reflect the observed volatility of prices and wind availability.

Now for each scenario $s \in \mathcal{S}$ we want to generate a sequence of correlated random errors. The scenarios are constructed by perturbing the base forecasts using the forecast errors, as follows:

$$\lambda_t^{(s)} = \bar{\lambda}_t + \sigma_\lambda \cdot \varepsilon_{\lambda,t}^{(s)} \tag{6.9}$$

$$\alpha_t^{(s)} = \bar{\alpha}_t + \sigma_\alpha \cdot \varepsilon_{\alpha,t}^{(s)} \tag{6.10}$$

Here, $s \in \mathcal{S}$ denotes the scenario index, and $\varepsilon_{\lambda,t}^{(s)}$ and $\varepsilon_{\alpha,t}^{(s)}$ are correlated standard normal random variables that preserve the spatio-temporal dependence structure of the uncertainties.

Spatio-temporal correlation modeling

To ensure realistic scenario dynamics, we model both the temporal autocorrelation within each process and the instantaneous cross-correlation between wind and electricity prices. Specifically, we assume that the forecast error terms for prices and wind exhibit the same autocorrelation and cross-variable correlation structure as the historical data. This assumption ensures that the generated scenarios accurately capture the persistence and co-movement of the two series.

The temporal autocorrelation of wind and price forecast errors is modeled using AR(1)-style correlations:

$$\text{Corr}[\varepsilon_{\alpha,t}, \varepsilon_{\alpha,t'}] = \phi_\alpha^{|t-t'|}, \quad \text{Corr}[\varepsilon_{\lambda,t}, \varepsilon_{\lambda,t'}] = \phi_\lambda^{|t-t'|},$$

where ϕ_α and ϕ_λ are empirically estimated from historical wind and price time series using lag-1 autocorrelations. This formulation reflects the behaviour of a first-order autoregressive process, AR(1), which is a special case of the general autoregressive model AR(p). In an AR(p) model, the current value of a time series is regressed on its p previous values, capturing temporal dependencies through a finite memory of past observations. By assuming an AR(1)-type decay in correlation, we incorporate the short-term persistence typically observed in wind and electricity price forecast errors.

The cross-correlation between wind and electricity prices is represented by a Pearson correlation coefficient ρ , also estimated from historical data. An example of these estimated correlation values for the 2023 price and wind data is shown in Table 6.1.

Table 6.1: Estimated coefficients for the temporal autocorrelation and cross-correlation of wind generation and electricity prices, based on 2023 historical data. The table reports the first-order autocorrelation coefficients ϕ_α and ϕ_λ for wind and price data, respectively, as well as the Pearson correlation coefficient ρ capturing the instantaneous cross-correlation between the two processes. These values are used to preserve the statistical dependence structure in the scenario generation model for the forecast errors.

	Wind (ϕ_α)	Price (ϕ_λ)	Cross-Correlation (ρ)
Estimated Coefficient	0.9896	0.9245	-0.369

The inter-variable cross-correlation between wind and electricity prices is directly applied to the error terms, ensuring that the negative relationship between the two variables, commonly observed in electricity markets, is preserved. Mathematically, this is defined as:

$$\text{Corr}[\varepsilon_{\alpha,t}, \varepsilon_{\lambda,t}] = \rho.$$

These correlation structures are preserved in the scenario generation process using Cholesky decomposition, which transforms independent random variables into correlated forecast errors that adhere to the specified autocorrelation and cross-correlation patterns.

To apply Cholesky we construct a full block correlation matrix $R \in \mathbb{R}^{2H' \times 2H'}$, capturing both autocorrelation and inter-variable correlation:

$$R = \begin{bmatrix} R_{\alpha\alpha} & R_{\alpha\lambda} \\ R_{\lambda\alpha} & R_{\lambda\lambda} \end{bmatrix}.$$

Here, $R_{\alpha\alpha}$ and $R_{\lambda\lambda}$ are $H' \times H'$ autocorrelation matrices where each entry is given by $\phi_\alpha^{|i-j|}$ and $\phi_\lambda^{|i-j|}$, respectively, and the off-diagonal blocks $R_{\alpha\lambda} = R_{\lambda\alpha}^\top = \rho \cdot I_{H'}$ capture the pointwise cross-variable dependence.

Cholesky-based scenario generation

We apply Cholesky decomposition to the matrix R , obtaining $R = LL^\top$, where $L \in \mathbb{R}^{2H' \times 2H'}$ is a lower triangular matrix. We generate standard normal samples $\mathbf{z}^{(s)} \sim \mathcal{N}(0, I_{2H'})$ for each scenario and transform them using:

$$\mathbf{y}^{(s)} = L \cdot \mathbf{z}^{(s)}.$$

The resulting vector $\mathbf{y}^{(s)} = \begin{bmatrix} \mathbf{y}_\alpha^{(s)} \\ \mathbf{y}_\lambda^{(s)} \end{bmatrix} \in \mathbb{R}^{2H'}$ contains both wind and price errors with the desired correlation structure. The scenario-specific realisations at time $t \in \mathcal{T}$ are then constructed as:

$$\alpha_t^{(s)} = \bar{\alpha}_t + \sigma_\alpha \cdot y_{\alpha,t}^{(s)}, \quad (6.11)$$

$$\lambda_t^{(s)} = \bar{\lambda}_t + \sigma_\lambda \cdot y_{\lambda,t}^{(s)}. \quad (6.12)$$

Note, however, that wind load factor values $\alpha_t^{(s)}$ must lie within the physical interval $[0, 1]$, as they represent the fraction of available wind power relative to the installed capacity. To ensure feasibility, any simulated values falling outside this range are capped at the nearest bound.

Implementing cholesky

Applying Cholesky decomposition requires the correlation matrix R to be symmetric and positive semi-definite. A matrix is positive semi-definite (PSD) if all of its eigenvalues are non-negative, which ensures that a valid lower-triangular factor L exists such that $R = LL^\top$. In our case, we encountered numerical failures during Cholesky decomposition due to the constructed R not being PSD. This was primarily caused by the combination of high temporal autocorrelations ($\phi_\alpha = 0.99$, $\phi_\lambda = 0.92$) and a strong negative cross-correlation ($\rho = -0.37$), which introduced internal inconsistencies in the matrix structure.

To address this issue, we modified the structure of the cross-correlation block $R_{\alpha\lambda}$ to reflect a temporally smoothed dependence. Instead of using a pointwise identity structure $\rho \cdot I_{H'}$, we introduced decaying cross-correlation across neighboring hours:

$$(R_{\alpha\lambda})_{i,j} = \rho \cdot \delta^{|i-j|},$$

where $\delta \in (0, 1)$ is a temporal decay parameter. This adjustment reflects the realistic assumption that high wind at time t may also influence prices at nearby hours (e.g., $t \pm 1, t \pm 2, \dots$), due to merit order and supply ramping effects. A value of δ close to 1 results in long-lasting correlation across time, while a value close to 0 limits interaction to the same hour only. Figure A.1 in the Appendix showcases the effect of different values of δ on the correlation structure. Choosing a decay factor $\delta \in (0.8, 0.9)$ retains strong hour-to-hour linkage, capturing realistic temporal spillover, while keeping the correlation matrix numerically stable and PSD for Cholesky decomposition.

In addition, we applied a standard regularisation step by symmetrising the matrix and adding a small diagonal perturbation $R \leftarrow R + \varepsilon I$ (with $\varepsilon \approx 10^{-10}$) to ensure all eigenvalues remain non-negative. The resulting matrix R satisfies the necessary conditions for Cholesky factorisation and supports the generation of temporally and spatially correlated stochastic scenarios.

6.3.2. Estimating mean and standard deviation for forecast errors

To generate realistic scenarios for electricity prices and wind load factors, we need to define the statistical properties of the forecast errors. Specifically, we must determine the mean and standard deviation of these errors, which capture the expected deviation from the base forecast.

In this thesis, we take a simplified but practical approach by assuming that the forecast mean for both electricity prices and wind load factors is equal to the historical observed values. This assumption can be interpreted as assuming a “perfect forecast” for the base values, with uncertainty introduced solely through the stochastic forecast errors. Mathematically, this is expressed as:

$$\bar{\lambda}_t = \mathbb{E}[\lambda_t] = \lambda_t^{\text{historic}},$$

and

$$\bar{\alpha}_t = \mathbb{E}[\alpha_t] = \alpha_t^{\text{historic}}.$$

In other words, the base forecast $\bar{\lambda}_t$ for electricity prices is simply assumed to be the historic price for the corresponding time period, and the base forecast $\bar{\alpha}_t$ for wind load factors is assumed to be the historic wind load factor. This approach provides a clear and intuitive baseline for scenario generation. Conceptually, this can be viewed as assuming that operators have access to highly accurate point forecasting models for the next week, but they still account for forecast uncertainty by adding normally distributed random errors. This strikes a balance between leveraging historical data for realistic baseline values and capturing the inherent uncertainty of future outcomes.

The uncertainty in these forecasts is captured through normally distributed random error terms with zero mean and a specified standard deviation σ_λ and σ_α . Due to the absence of historical forecast data, which prevents direct calculation of forecast errors, we estimate forecast uncertainty using the standard deviation of historical price and wind load factor data. Specifically, the standard deviations are calculated as the standard deviation of the observed prices and wind load factors over the entire historical year. This method provides a simple yet robust measure of forecast uncertainty. However, it also assumes that the forecast error distribution is constant throughout the year, which may not fully capture seasonal variations in volatility. Optionally, different $\sigma_{\lambda,t}$ and $\sigma_{\alpha,t}$ could be used for each month to reflect seasonal volatility. This allows the model to better capture periods of higher or lower price uncertainty.

By using the standard deviation calculated from the entire year, we effectively assume perfect knowledge of future volatility. If we would want to use this stochastic optimisation model for real-time operational planning, this assumption is not realistic since it requires future knowledge. In practical applications, operators would typically rely on their own forecasting models to generate point forecasts and add an uncertainty measure based on recent performance of those forecasts. Our approach, however, is a deliberate simplification. Rather than directly modeling a forecasting process, we focus solely on capturing the inherent uncertainty around the forecast values. If we were to adapt this model for real-time operation without relying on future values, an alternative approach would be to estimate the mean and standard deviation of forecast errors using only the most recent historical period, such as the past week or month. This would make the model more realistic by ensuring that the forecast error distribution is responsive to recent market conditions and avoids any reliance on future data.

6.3.3. Illustration and validation of scenario generation method

To ensure that our scenario generation method accurately captures the statistical properties of the underlying electricity price and wind load factor data, we conduct a validation analysis using historical 2023 data. Our primary objective is to confirm that the generated scenarios maintain the key characteristics of the original data, including the mean values, standard deviations, and the negative cross-correlation between wind and electricity prices.

Figure 6.1 provides an illustration of ten generated scenarios for both electricity prices and wind load factors, alongside the base (historical) values for comparison. Each scenario is a stochastic realisation based on the forecast error modeling approach described in this section.




Figure 6.1: Illustration of scenario generation method for electricity prices and wind load factors, using 2023 as base case year. The top plot shows 10 generated price scenarios alongside the base (historical) price trajectory, while the bottom plot shows 10 generated wind scenarios. The generated scenarios were constructed using Monte Carlo Simulation and Cholesky's Decomposition to capture the variability, temporal auto-correlation and negative cross-correlation between wind and prices.

To quantitatively validate the scenario generation method, we conduct a simulation test using the year of 2023 data and simulating a full year (8760 hours). Specifically, we generate 100 scenarios for both electricity prices and wind load factors using the Monte Carlo method, ensuring that each scenario preserves the spatial and temporal correlation structure specified in our model. The results are summarised in Table 6.2, showing the mean, standard deviation, and calculated cross-correlation between the generated scenarios.

Table 6.2: Statistical comparison of historical and simulated electricity price and wind load factor data for the base year 2023. The table reports the mean, standard deviation, and cross-correlation values used to evaluate the performance of the scenario generation method. While the simulated scenarios closely replicate the historical mean values and preserve the negative cross-correlation between wind and electricity prices, they exhibit higher standard deviations, reflecting the added variability introduced by the stochastic generation process.

	Wind	Price	Cross-Correlation
Base (Historical)	Mean: 0.4935 Std Dev: 0.2971	Mean: 97.73 Std Dev: 48.81	-0.369
Simulated	Mean: 0.4961 Std Dev: 0.3428	Mean: 98.12 Std Dev: 69.35	-0.348

These results demonstrate that our scenario generation method effectively replicates the mean values and

the negative cross-correlation between wind and prices observed in the historical data. However, we do observe a slightly higher standard deviation for the simulated scenarios compared to the original data. This may be attributed to the added stochasticity in the scenario generation process, which introduces additional variability.

6.3.4. Scenario probabilities

In the two-stage stochastic optimisation model, each generated scenario $s \in \mathcal{S}$ represents a possible realisation of future electricity prices and wind availability. To account for the inherent uncertainty, each scenario is assigned a probability ϕ_s that represents its likelihood of occurring. In this thesis, we adopt an equal probability approach, where each scenario is considered equally likely. Mathematically, this is expressed as:

$$\phi_s = \frac{1}{|\mathcal{S}|}, \quad \forall s \in \mathcal{S}.$$

This choice of equal scenario probabilities is a common and practical approach when no clear information exists to suggest that some scenarios are more probable than others. It ensures that the expected value of the objective function is a simple average of the outcomes across all scenarios, providing a balanced assessment of the possible future states.

While more complex probability assignment methods exist, such as using historical data to calibrate scenario probabilities or employing machine learning techniques to predict scenario likelihoods, our approach prioritises simplicity and computational efficiency. Such methods could be explored in future work to further enhance the model's realism.

7

Results

This chapter presents the optimisation results and evaluates system performance under different modeling approaches and system configurations. Section 7.1 presents the base case results from the deterministic model, which serves as a benchmark for subsequent comparisons, and includes a sensitivity analysis on key economic parameters. Section 7.2 compares the performance of different modeling approaches and examines how variations in PPA sizing and hydrogen sales prices influence outcomes. Section 7.3 introduces differentiated hydrogen sales prices based on electricity source, providing a more realistic link between system operation and economic incentives. Finally, Section 7.4 evaluates the role and profitability of battery storage under different system setups and across multiple historical years.

7.1. Base case results for deterministic model

To establish a benchmark for evaluating system performance, this section presents results from the base case model. The configuration analysed includes a 50 MW electrolyser, a 100 MWh battery storage system, a 100 MW wind PPA contract and a 100 MW grid connection. The model is run over a full year of operations, representing 8,760 hourly time steps. The solution method applied in this base case is a deterministic optimisation (DO) approach with perfect foresight for entire year. These results therefore reflect the best-case system behaviour, unaffected by forecast uncertainty or limited foresight.

7.1.1. System dynamics under varying market and wind conditions

To illustrate the operational behavior of the optimised hybrid energy system, we analyse four representative time windows that capture distinct combinations of wind generation and electricity price dynamics. These scenarios reveal how the model coordinates the interaction between the battery, electrolyser, and electricity market in response to external conditions. Figures 7.1 to 7.5 each present five vertically aligned subplots showing: (1) electricity prices and wind load factor, (2) market and PPA power flows, (3) battery activity, (4) electrolyser operation, and (5) wind power allocation. The power flow subplot explicitly shows the direction and magnitude of electricity flows between sources and sinks (e.g., wind to battery, market to electrolyser). Additionally, battery state-of-charge and electrolyser utilisation rate are included to provide further insight into asset-level performance. This integrated visualization enables a detailed understanding of the system's dispatch logic and the consistency of energy flows between components.

Scenario 1: high wind and low prices

Figure 7.1 presents a period from December 19 to December 28, 2023 (hours 8464–8675), marked by consistently high wind generation and generally low electricity prices. Except for a brief drop around hour 8640, wind availability remains near maximum levels, often driving prices to zero or even negative values. Throughout this interval, the electrolyser operates continuously at full capacity, powered almost entirely by wind PPA electricity. This behavior reflects the model's prioritisation of hydrogen production from renewable sources. Surplus wind energy is either sold to the market or stored in the battery. Notably, excess power is sold even during negative price hours, as this results in lower losses compared to paying redispatch penalties.

A temporary shift occurs around hour 8640, when wind output declines and prices surge. The system responds by supplementing electrolyser input with battery discharge and grid purchases. The battery is dis-

charged strategically during peak-price periods, enabling cost-efficient electricity sourcing and stable hydrogen production under volatile conditions.

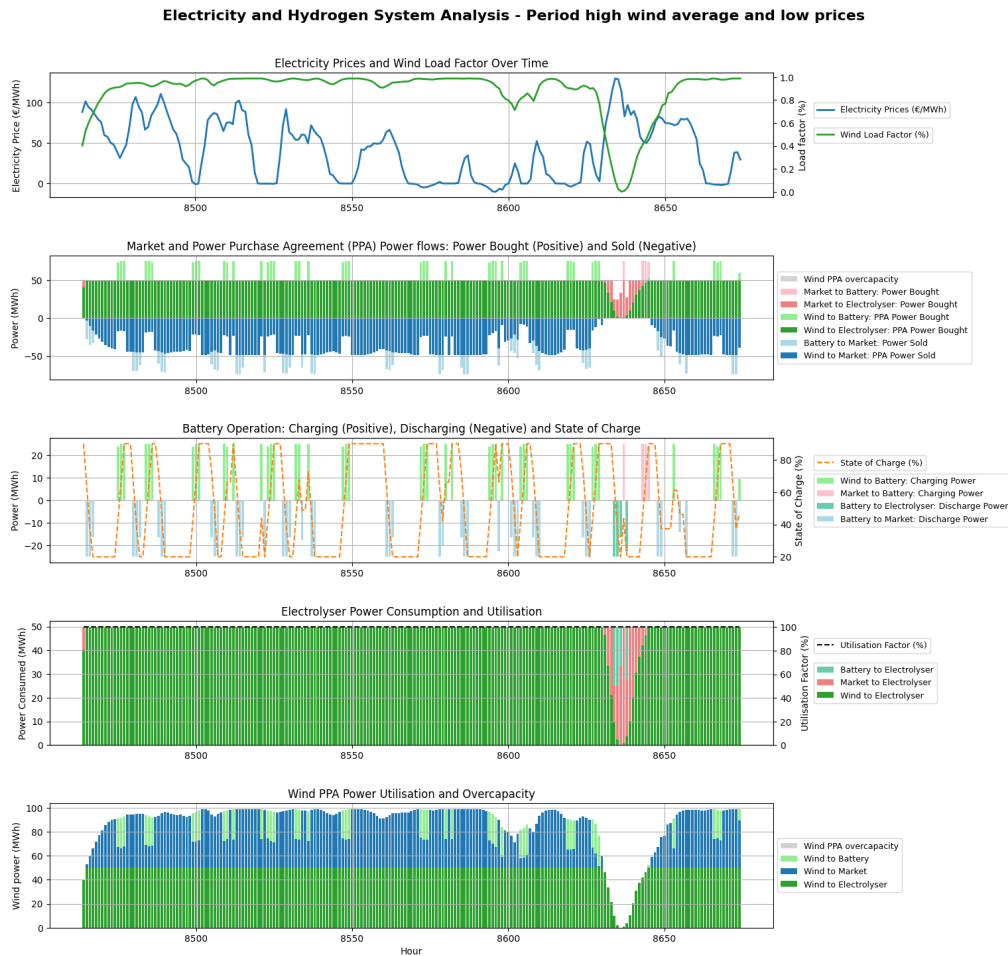


Figure 7.1: Base case results: dispatch behavior during a generally high-wind, low-price period, corresponding to December 19 to December 28, 2023 (hours 8464–8675). Electrolyser operation is sustained at full capacity with wind as the dominant input. Surplus wind energy is either sold to the market, even during negative pricing, or stored in the battery. A temporary low-wind episode near hour 8640 triggers price spikes and a shift to grid and battery support.

Scenario 2: low wind and high prices

Figure 7.2 depicts a period characterised by multiple episodes of low wind availability accompanied by elevated electricity prices. The inverse correlation between wind output and price is evident: prices rise during low wind periods and fall when generation increases. In response, the model shuts down the electrolyser for an extended period to avoid uneconomical hydrogen production. Instead, available wind energy is either sold to the market or used to charge the battery. Given that spot market prices frequently exceed the fixed PPA rate of €97/MWh, electricity resale yields attractive margins.

Battery usage intensifies to store wind energy and exploit price volatility, with frequent charging during price dips and discharging during peaks. Toward the end of the timeframe, as wind improves and prices moderate, the electrolyser resumes operation at full capacity. One exception is when it operates at minimum load, allowing the system to prioritise profitable electricity sales over hydrogen production. This strategy avoids triggering additional shutdown costs, while still reallocating most of the available wind power toward more profitable electricity sales. This behavior reflects a strategic trade-off between selling hydrogen versus electricity, and the importance of managing start-stop cycles to minimise operational penalties.

Figure 7.3 shows the hourly revenue and cost flows resulting from these dispatch activities, with the red line tracing net profit at each timestep. Profitability peaks when wind energy is resold during high-price periods, while losses emerge during grid dependence or electrolyser inactivity. The results confirm that the model's decisions are aligned with short-term economic optimisation.

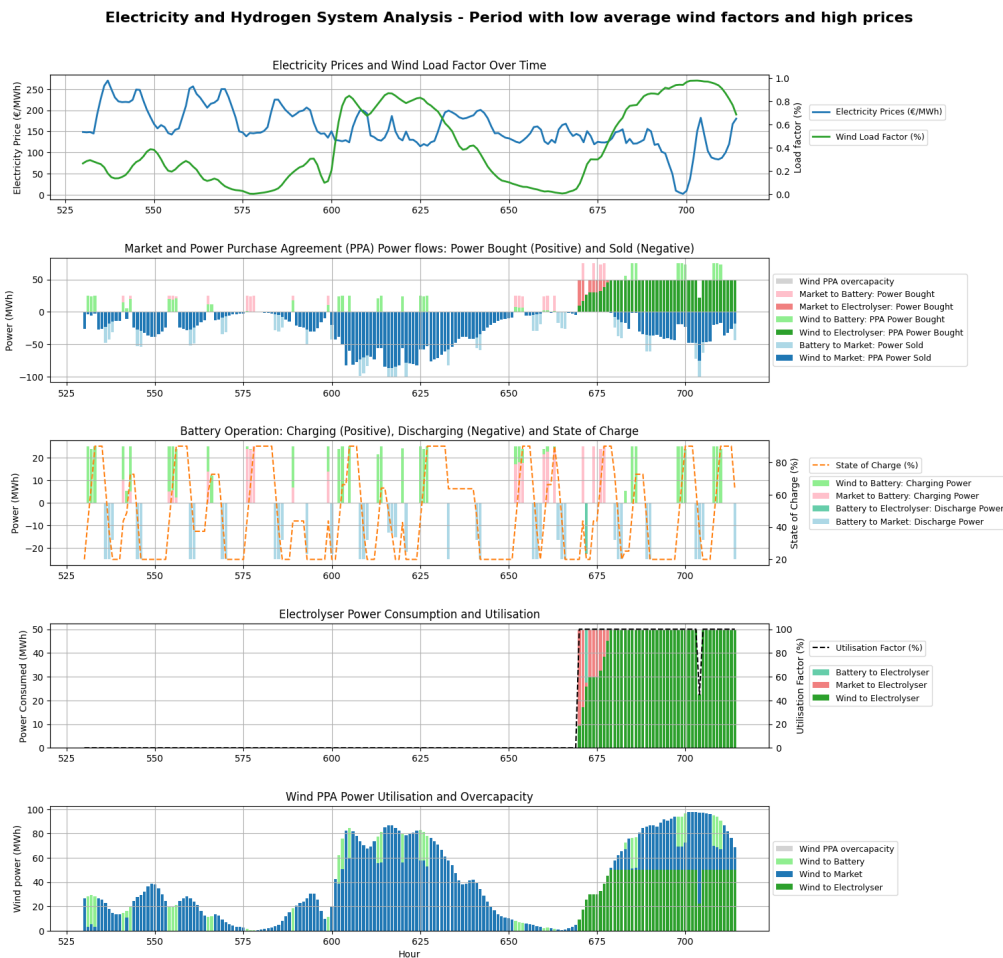


Figure 7.2: Base case results: dispatch strategy under low-wind, high-price conditions, corresponding to January 23 to January 30, 2023 (hours 530–715). The system shut-offs hydrogen production during low wind and expensive grid periods, reallocating wind power toward market sales and battery charging. High market prices enable profitable battery discharging to the market. Additionally, toward the end of the window we experience that the electrolyser is strategically ramped down to minimum load to avoid shutdown penalties.

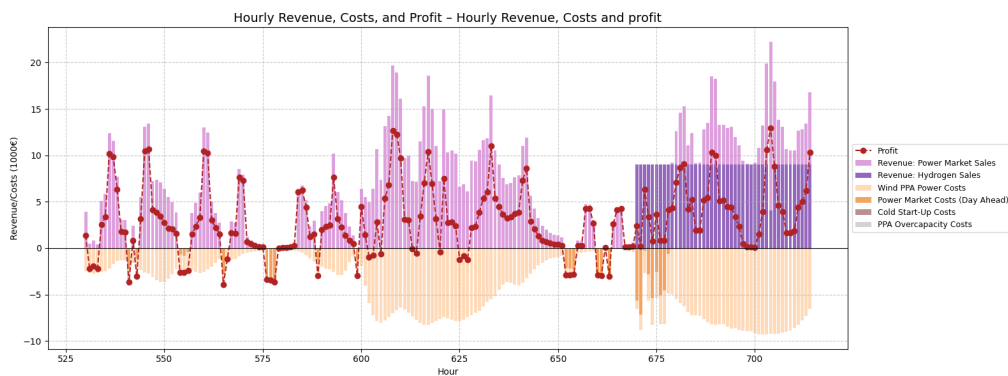


Figure 7.3: Hourly revenue, costs, and profit during the low-wind, high-price scenario: January 22–30, 2023 (hours 530–715) for base case results. Positive revenue from electricity market sales is observed when excess wind is sold during peak-price hours. Profitability declines during periods of low wind and high market prices, especially when the electrolyser is shut off or draws power from the grid and fixed PPA costs remain.

Scenario 3: wind overcapacity with ultra-low prices

Figure 7.4 captures a critical system behavior during a period of wind overcapacity and extremely low electricity prices, approximately between June 29 to July 3, 2023. Throughout this interval, wind generation far

exceeds the combined consumption capacity of the electrolyser and battery, both of which are operated at full capacity for much of the time.

Once storage and electrolyser limits are reached, the remaining surplus triggers curtailment and redispatch penalties, represented by the grey bars in the power flow subplot. Rather than incurring even larger losses by selling at highly negative market prices, the model accepts redispatch costs for excess wind. Interestingly, in some of these hours, the system purchases electricity from the market, effectively earning revenue due to the negative price signal.



Figure 7.4: Base case dispatch results: system operation during wind overcapacity and ultra-low price conditions corresponding to June 29 to July 3, 2023 (hours 4300–4400). Redispatch penalties occur when surplus wind exceeds storage and electrolyser capacity. Market purchases are used strategically when electricity prices are deeply negative.

Scenario 4: high wind variability

Figure 7.5 illustrates a highly volatile period from August 12 to August 17, 2023 (hours 6075–6252), with fluctuating wind availability and electricity prices. The persistently low and erratic wind levels prevent the electrolyser from relying solely on wind power. Instead, the system maintains high utilisation levels by flexibly sourcing electricity from the grid and the battery, in addition to whatever wind power is available.

Battery operation becomes particularly dynamic, with frequent charging and discharging cycles that help balance intermittent renewable input and capitalise on price deviations in the market. Moreover, we see the system frequently ramps the electrolyser up and down. Notably, a sharp price spike within this period triggers a brief electrolyser shutdown, illustrating the system's sensitivity to unprofitable operating conditions.

Summary dispatch dynamics When wind is abundant and electricity prices are low, the system prioritises hydrogen production and manages surplus power through battery storage or market sales, sometimes even at a loss, to avoid redispatch penalties. During low-wind, high-price periods, it shifts to market participation

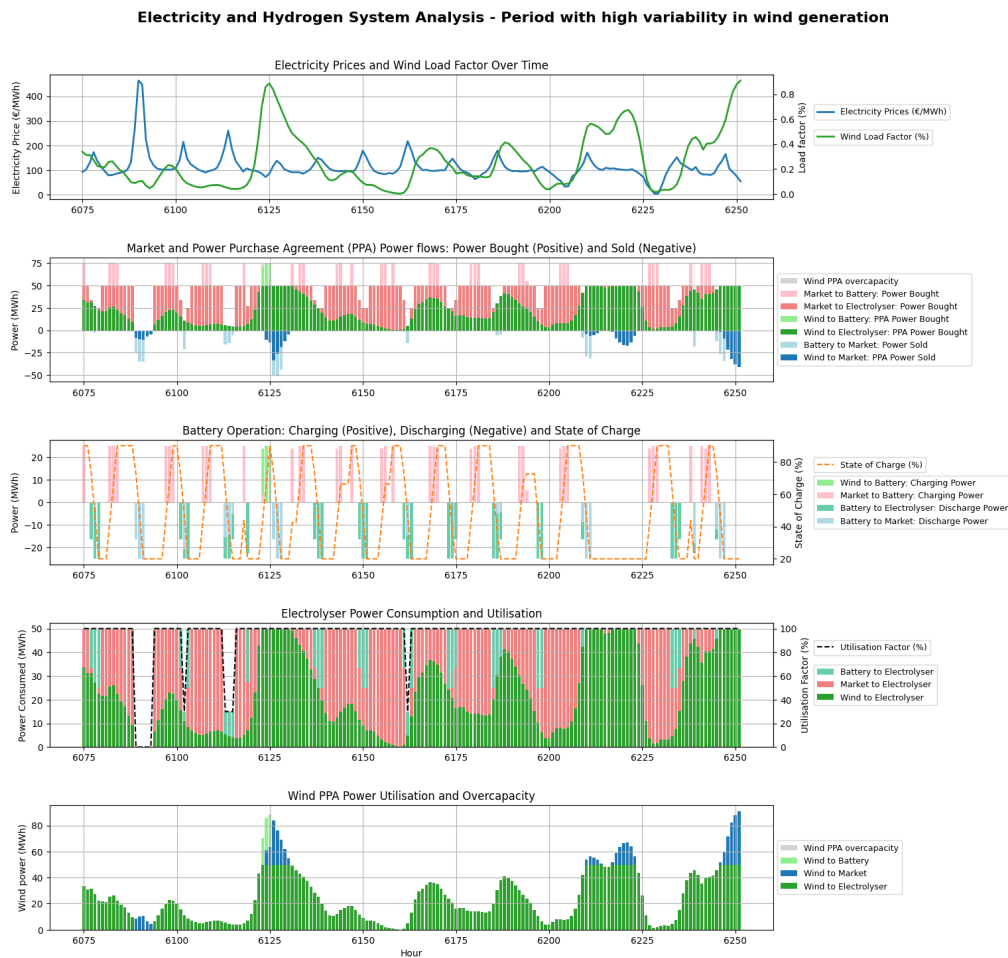


Figure 7.5: Base case dispatch results: system behavior under high wind and price variability, corresponding to August 12–17, 2023 (hours 6075–6252). Wind levels are generally low and irregular, leading to frequent cycling of the electrolyser and heavy reliance on the battery and market purchases.

and therefore frequently charges and discharges the battery to capitalise on market price differences, with the electrolyser strategically shut down or maintained at minimum load. Under overcapacity, redispatch penalty costs are accepted to avoid worse economic outcomes. In volatile conditions, the system combines all available sources to stabilise electrolyser operation, dynamically adjusting to short-term fluctuations. Collectively, these behaviors reflect an economically rational and flexible dispatch strategy.

7.1.2. Operational and financial KPI results

Tables 7.1 and 7.2 summarise the annual operational and financial performance of the optimised dispatch. Over the year, 419.1 GWh of electricity is supplied to the electrolyser, resulting in 7.54 kt of hydrogen production. The system draws most of its energy from the wind PPA (446.1 GWh), supplemented by 115.4 GWh from the market. Additionally, 48.2 GWh is cycled through the battery to support flexible operation.

The financial breakdown in Table 7.2 reports total annual revenue of €88.78 million, primarily from hydrogen sales (€75.43 million), with €13.35 million generated from electricity market transactions. On the cost side, wind PPA payments account for €42.48 million, and market electricity purchases for €10.84 million, summing to approximately €53 million in gross power costs. Including all other operating expenses, such as electrolyser and battery OPEX, cold start-ups, and redispatch penalties, the total operational costs (excluding CAPEX) reaches €56.57 million.

Electricity resales recover nearly 25% of total power costs, significantly improving cost efficiency. After accounting for €13.60 million in annualized CAPEX, the system achieves a net profit of €18.62 million. These results confirm the economic viability of the hybrid system and highlight the role of market participation and dispatch flexibility in enhancing profitability.

An additional metric tracks the annual loss from selling surplus wind electricity back to the market at prices lower than the fixed PPA rate. This “wind-to-market loss” is computed by multiplying the hourly surplus wind exported to the market by the difference between the market price and the PPA cost, then summing over the year. The result is a minor loss of €0.75 million, equivalent to just 4% of annual net profit. While this loss reflects the occasional inefficiency of forced wind acceptance, it is offset by the strategic advantage of overcapacity which enables an increased hydrogen production and revenue. Thus, although there is a marginal trade-off, maintaining slight overcapacity may strengthen the business case by enabling higher hydrogen output from wind power.

Table 7.1: Summary of base case operational results: total system power and hydrogen flows for the optimised dispatch.

Metric	Unit	Value
Avg wind load factor	-	0.51
Avg DA market price	€/MWh	95.8
Power bought DA market	GWh	115.4
Power sold DA market	GWh	137.7
Wind power from PPA	GWh	446.1
PPA overcapacity	GWh	1.01
Battery charge	GWh	48.2
Battery discharge	GWh	43.5
Power to electrolyser	GWh	419.1
Hydrogen produced	kt	7.54

Table 7.2: Summary of base case financial results: breakdown of all revenue streams and costs.

Metric	Unit	Value
<i>Revenues</i>		
Hydrogen revenue	M€	75.43
Power market revenue	M€	13.35
Total revenue	M€	88.78
<i>CAPEX (Annualized)</i>		
Electrolyser CAPEX	M€	10.28
Battery CAPEX	M€	3.32
Total annualized CAPEX	M€	13.60
<i>Operating Costs</i>		
Market power cost	M€	10.84
Wind PPA cost	M€	42.48
Redispatch penalty	M€	0.15
Cold start-up cost	M€	0.12
Electrolyser OPEX	M€	2.19
Battery OPEX	M€	0.79
Total annual OPEX	M€	2.98
Total annual cost (excl. CAPEX)	M€	56.57
<i>Profit</i>		
Annual profit (excl. CAPEX)	M€	32.22
Net profit (after all costs)	M€	18.62
<i>Wind power surplus loss</i>		
Estimated wind-to-market loss	M€	-0.75

Operational KPI's

Tables 7.3, 7.4, and 7.5 summarise the key operational behavior of the electrolyser and battery systems over the simulation period. The electrolyser exhibits consistently high utilisation, achieving a 95.7% average utilisation rate and running at full capacity for 95.3% of the time. Downtime totals 300 hours (3.4% of the year), which aligns exactly with the required maintenance time, indicating no economically driven idling beyond that constraint. A total of 15 shut-offs were observed, each lasting on average 20 hours, with a maximum duration of 79 hours. This suggests the model prefers fewer but longer interruptions, balancing the trade-off between shut-down penalties and low-revenue operating periods.

Battery usage is also active and consistent. Over the year, it charges for 2017 hours and discharges for 1957 hours, corresponding to approximately 46% of the year. Average charging and discharging powers are 23.9 MW and 22.2 MW respectively, in line with the system's round-trip efficiency of 90.2%. The battery operates at full capacity in over 66% of its active time, highlighting a strategy that maximises market trading value and storage utilisation during economically favorable periods.

Table 7.3: Summary of electrolyser utilisation metrics for base case results.

Metric	Unit	Value
Avg utilisation rate	%	95.7
Shut-offs (count)	-	15
Off status (hours)	hours	300
Off status (share)	%	3.4
Min capacity (hours)	hours	114
Min capacity (share)	%	1.3
Max capacity (hours)	hours	8346
Max capacity (share)	%	95.3

Table 7.4: Shut-off cycle characteristics of the electrolyser for base case results.

Metric	Unit	Value
Shut-offs (count)	-	15
Avg time per shut-off	hours	20.0
Max shut-off time	hours	79.0
Min shut-off time	hours	3.0

Table 7.5: Summary of annual battery charging and discharging behaviour in the base case, including energy throughput, average power per charge/discharge, operating intensity, share charged/discharged at full rated power and round-trip efficiency.

Metric	Unit	Charging	Discharging
Total energy throughput	GWh	48.2	43.5
Average power	MW	23.9	22.2
Operating hours	hours	2017	1957
Share at full rated power	% of op. hours	67.6	66.4
Round-trip efficiency	%	90.2	

Power flow distributions

As part of the analysis of system operation, we examine the distribution of power flows between components. Tables 7.6 through 7.8 provide a structured breakdown of energy flows between the wind PPA, electrolyser, battery and electricity market, capturing how energy is allocated, stored, and traded under the model assumptions.

Wind utilisation and electrolyser input: Table 7.6 shows that most wind power (70.9%) is directly routed to the electrolyser, followed by 24.3% sold to the market, and a small portion (4.6%) stored in the battery. Only 0.2% is curtailed due to overcapacity. On the input side, wind accounts for 75.7% of the electricity consumed by the electrolyser, while market power covers 20.9% and battery discharge contributes 3.4%. These shares highlight that wind-powered hydrogen production is prioritised wherever possible.

Battery operation: As shown in Table 7.7, battery charging is slightly skewed toward the market (57.5%) over wind (42.5%). On discharge, a clear majority (66.9%) is sold to the market rather than used for hydrogen production (33.1%). This behaviour suggests that the battery is mostly used to strategically trade with the market and make use of price differences over time, shifting excess electricity (especially from the market) to higher-priced periods for sale, rather than serving as a buffer for hydrogen production.

Market interaction: Of the electricity purchased from the market, 75.9% goes directly to the electrolyser, with the rest routed through the battery. On the sales side, 78.8% of sold electricity is directly from the wind PPA, and 21.2% is from battery discharge. These distributions confirm the system's logic: it uses wind energy as the primary driver of hydrogen production, leverages the market when prices allow, and employs the battery mainly to enhance revenue through market trading. Hydrogen production is prioritised over electricity export, and storage is used tactically rather than continuously. This balance highlights the hybrid model's economic and operational effectiveness. It's also important to recognise that these patterns are shaped by technical constraints, such as the electrolyser's 50 MW rating compared to the battery's 25 MW, which enables higher energy throughput and greater market participation on the electrolyser's side. Additionally, the contracted wind PPA has a capacity of 100 MW, and with an average load factor of approximately 50%, it often produces enough to fully power the electrolyser. Moreover, during high wind periods the PPA can even exceed this baseline, enabling simultaneous operation of both the electrolyser and battery, or even export to the market.

Table 7.6: Wind PPA electricity allocation and electrolyser input breakdown for base case results.

Wind PPA Output

Destination	GWh	Share (%)
Wind to Electrolyser	317.0	70.9
Wind to Battery	20.5	4.6
Wind to Market	108.6	24.3
Overcapacity	1.0	0.2
Total	447.1	100.0

Electrolyser Input

Source	GWh	Share (%)
Wind to Electrolyser	317.0	75.7
Battery to Electrolyser	14.4	3.4
Market to Electrolyser	87.6	20.9
Total	419.1	100.0

Table 7.7: Battery charging and discharging distribution for base case results.

Battery Charge

Source	GWh	Share (%)
Wind to Battery	20.5	42.5
Market to Battery	27.8	57.5
Total	48.2	100.0

Battery Discharge

Destination	GWh	Share (%)
Battery to Electrolyser	14.4	33.1
Battery to Market	29.1	66.9
Total	43.5	100.0

Table 7.8: Electricity market buy and sell distributions for base case results.

Market Buy

Usage	GWh	Share (%)
To Electrolyser	87.6	75.9
To Battery	27.8	24.1
Total	115.4	100.0

Market Sell

Source	GWh	Share (%)
Wind to Market	108.6	78.8
Battery to Market	29.1	21.2
Total	137.7	100.0

Financial KPI's

Table 7.9 summarises the high-level economic performance of the optimised hybrid energy system. The model yields a Net Present Value (NPV) of €97.14 million and an Internal Rate of Return (IRR) of 21.1%, indicating strong financial viability. The corresponding Levelised Cost of Hydrogen (LCOH) is €8.40 per kilogram, positioning the system competitively with other green hydrogen production pathways.

Table 7.10 details the breakdown of the LCOH. The cost is primarily driven by electricity purchases under the PPA, which accounts for over 70% of the LCOH. However, revenue from electricity sales offsets a significant portion of these costs, contributing a negative €2.11/kg, corresponding to a reduction of 20.2% of the total LCOH. CAPEX for the electrolyser and battery, as well as fixed OPEX, represent moderate shares. Taxes and minor costs such as redispatch and shut-offs contribute marginally. The final LCOH, inclusive of all adjustments, is €8.40/kg.

Table 7.9: KPI results for base case model including the projects NPV, IRR and LCOH.

Metric	Unit	Value
Net Present Value (NPV)	M€	97.14
Internal Rate of Return (IRR)	%	21.1
Levelized Cost of Hydrogen (LCOH)	€/kg	8.40

Table 7.10: LOCH cost component breakdown for base case results, including contributions of power, capital and operational expenditures.

Component	€/kg	Share (%)
Electricity PPA	5.937	70.6
Electricity Market	1.326	15.8
Electricity Sales	-2.113	-25.1
Capex Electrolyser	1.436	17.1
Capex Battery	0.464	5.5
Opex Electrolyser	0.306	3.6
Opex Battery	0.111	1.3
Shut-off costs	0.022	0.3
Redispatch / Overcapacity	0.020	0.2
Tax	0.895	10.7
LCOH (excl. power sales)	10.518	
Final LCOH (incl. power sales)	8.404	100.0

7.1.3. Sensitivities

To better understand the robustness of the model and to identify key economic drivers, we performed a series of sensitivity analyses on both technical and economic parameters. These include the hydrogen sales price, PPA price and size, electrolyser shut-off costs and lastly, battery size and unit CAPEX costs.

The analysis shows that some parameters affect only financial outcomes, such as fixed CAPEX and OPEX parameters, while others significantly influence both the dispatch strategy and profitability. For example, changes in CAPEX or OPEX of the electrolyser and battery do not alter the dispatch strategy in a fixed-size system, but do affect the overall financial feasibility. These will be revisited in more detail in the battery analysis where we discuss optimal battery size in Section 7.4.

PPA price and hydrogen sales price sensitivities

Figure 7.6 presents the results of the sensitivity analysis of the LCOH and IRR with respect to the hydrogen sales price and the PPA price parameters. The PPA price also does not influence operational behaviour in the current model due to the take-or-pay contract structure: the operator pays for all wind power generated regardless of actual use. As a result, dispatch decisions remain unaffected. Instead, the PPA price directly affects the cost structure, increasing the LCOH and reducing the IRR in a linear fashion, as shown in Figure 7.6b.

In contrast, parameters such as hydrogen sales price and PPA size do influence operational decisions. Figure 7.6a illustrates the effect of hydrogen price on both LCOH and IRR. As expected, higher hydrogen sales prices increase profitability, reflected in a rising IRR. The wider margin between electricity costs and hydrogen revenues incentivises greater electrolyser utilisation.

Interestingly, LCOH also increases with sales price, which may seem counterintuitive. The change in LOCH is due to three key effects. First, while increased hydrogen production spreads fixed CAPEX and OPEX over more output (reducing their contribution to LCOH), this is offset by a second effect: higher utilisation leads to production during less favourable (i.e., more expensive) hours, increasing electricity costs. Most importantly, as profit margins rise, taxes become a growing share of total costs, further raising the LCOH.

At lower hydrogen prices, utilisation drops significantly, and the system only produces the minimum required hydrogen. Below a certain threshold, profitability becomes negative, indicating that hydrogen sales price is a critical determinant of project viability.

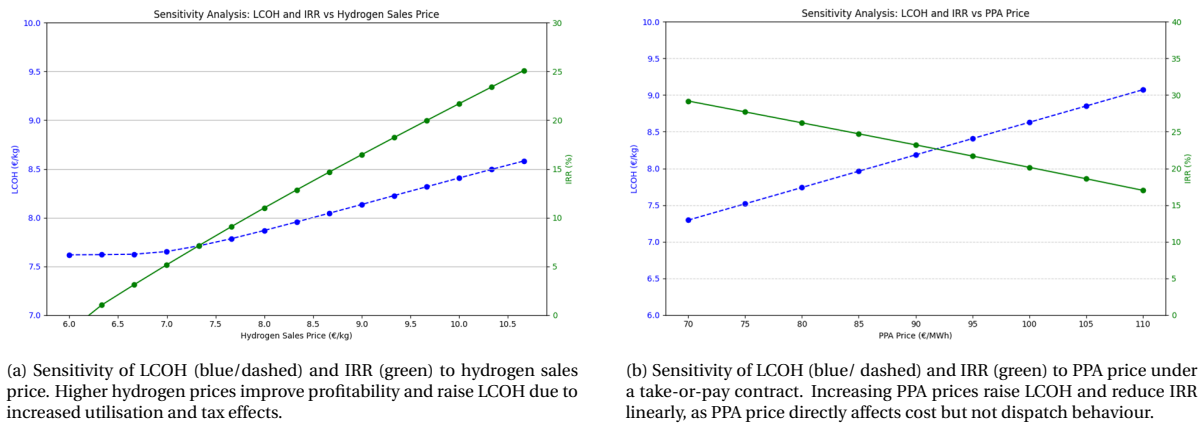


Figure 7.6: Sensitivity analysis of financial metrics LCOH and IRR with respect to key economic drivers. Subfigure (a) illustrates the impact of hydrogen sales price on the LCOH and the IRR, while subfigure (b) shows the effect of varying the PPA price under a take-or-pay contract. In both plots, the blue dashed line represents LCOH (left axis) and the green line represents IRR (right axis).

Shut-off costs sensitivity

The height of the shut-off cost influences the frequency of electrolyser start-up and shut-down cycles. Higher shut-off costs discourage frequent cycling, resulting in fewer annual transitions. However, the overall impact on annual profit is minimal, as shut-off costs represent only a small fraction of total system costs. As such, while these costs affect operational behaviour to some extent, they do not significantly alter the system's financial performance.

The number of electrolyser shut-downs per year depends on the assumed cost per shut-off cycle. Recall that in the base case, a shut-off cost of €8,000 results in 15 shut-offs per year. At lower shut-off costs, the optimiser allows more frequent cycling (up to 20 events per year), while beyond a certain threshold, €14,000, the number stabilises at four annual shut-downs. Even with very high penalties (€40,000 per event), the optimiser still finds it beneficial to shut down 3 times per year. However, the overall impact on profitability is negligible; shutdown-related costs remain below 0.5% of the LCOH, and the effect on IRR and annual profit is minimal.

7.2. Comparing different optimisation approaches

This section synthesises the results obtained from comparing three optimisation approaches for the operation of the hydrogen energy system:

1. **Deterministic Optimisation (DO):** Assumes perfect foresight over the entire year. The model is solved once for the full 8760-hour horizon, optimising dispatch decisions with full knowledge of future electricity prices and wind availability.
2. **Rolling Horizon Optimisation (RHO):** Introduces a reduced foresight window of one week. The model is solved iteratively in daily planning periods (365 rolling steps), each time optimising a 168-hour look-ahead with perfect forecasts. Only the decisions for the first 24 hours of each step are implemented.
3. **Scenario-based Stochastic Optimisation with Rolling Horizon (SORH):** Extends the RHO approach by incorporating uncertainty in forecasts beyond the first 24 hours. For each daily solve, multiple future scenarios are constructed by adding correlated random errors to forecasted electricity prices and wind generation. The model then optimises dispatch decisions under uncertainty using a two-stage stochastic framework.

All models, including SORH, assume perfect knowledge of the first 24 hours of each planning period. This reflects real-world conditions, where day-ahead electricity prices are published one day in advance and wind forecasts are relatively accurate over short horizons. The key modeling differences lie in the treatment of foresight and uncertainty beyond the first 24 hours. The DO model benefits from full-year knowledge, while RHO operates with 1 week perfect foresight, and SORH captures the impact of forecast uncertainty beyond the day-ahead horizon. The objective is to evaluate how these assumptions affect dispatch strategy, electrolyser utilisation, and the overall financial performance of the hydrogen system.

7.2.1. Base case PPA size: minimal differences between DO, RHO and SORH solutions

When comparing the DO, RHO, and SORH models under the base case configuration (which includes a oversized wind PPA of 100 MW for a 50 MW electrolyser) we observe only minimal differences in model outcomes. As shown in Table 7.11, hydrogen production varies by less than 5% across all models. Profit differences are also limited, ranging from 1.1% for RHO to 5.9% for the 10-scenario SORH model, compared to the DO benchmark. Electrolyser utilisation remains consistently high (above 95%), and differences in power flows (market buying, selling, battery charging) are present but do not significantly impact overall profitability.

These limited variations are explained by the high availability of wind energy under the large PPA. Since the electrolyser can operate at full capacity most of the time using wind power, the advantage of perfect foresight in the DO model is largely neutralised. There is little opportunity for dynamic optimisation based on varying market prices, meaning that even models with limited foresight (RHO and SORH) achieve near-optimal dispatch strategies.

As observed throughout our analysis in changing the PPA size which will be further discussed later on, the main effect of moving away from perfect foresight in the DO model is a change in the electrolyser's utilisation pattern. Specifically, determining when and how frequently the electrolyser is turned on, off, or operated at minimum load. Without full knowledge of future prices and wind availability, the RHO and SORH models cannot perfectly schedule hydrogen production in the most cost-effective hours, and must instead optimise within the reduced 1-week horizon window.

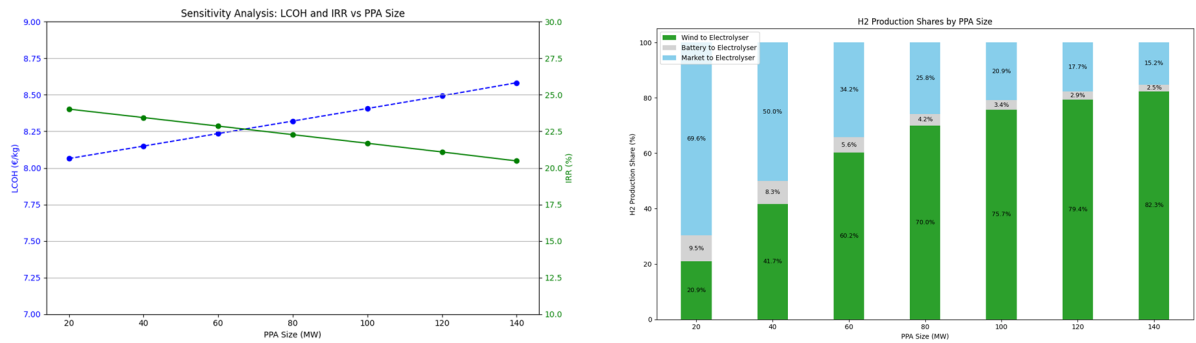
However, when the system is supplied with abundant renewable power which is the case under the 100 MW PPA, the electrolyser operates at or near full capacity across most of the year. In this scenario, optimisation decisions are dominated by the availability of wind power rather than fluctuations in electricity prices in the future. As a result, all three models converge on similar dispatch strategies and achieve comparable electrolyser utilisation rates (above 95%). This leaves little room for the advantages of perfect foresight, and thus the performance gap between the deterministic and limited-foresight models remains small.

Table 7.11: Comparison of results from the different modeling approaches: DO, RHO and SORH. Results generated for base case system: 100 MW wind PPA, 50 MW electrolyser and 100 MWh battery. Under these system conditions, the differences in operational behaviour and profitability between the various modeling approaches remain relatively limited.

Model	# Scenarios	H_2 Produced (kt)	Util. Rate (%)	Profit (M€)	$p^{M,buy}$ (GWh)	$p^{M,sell}$ (GWh)	$p^{B,charge}$ (GWh)	Runtime (s)
DO	-	7.15	95.7	18.62	105.53	149.52	48.20	19
RHO	-	7.11	95.7	18.46	108.31	154.02	48.24	26
SORH	5	7.51	95.3	18.45	115.45	139.54	47.89	112
SORH	10	7.50	95.2	18.39	115.43	139.71	50.63	231
SORH	15	7.50	95.2	18.45	115.26	139.72	48.86	356
SORH	20	7.50	95.2	18.46	115.34	139.86	48.39	678

7.2.2. Sensitivity analysis PPA size

To better understand when modeling assumptions begin to significantly impact system outcomes, we first performed a sensitivity analysis on PPA size using the DO model. Before comparing the different optimisation approaches (DO, RHO, SORH), this preliminary analysis helped reveal how the system responds to variations in renewable electricity availability and raised important questions about the validity of the assumed hydrogen sales price under different conditions.



(a) Sensitivity of LCOH and IRR to PPA size. The blue, dashed line shows LCOH (left axis), while the green line shows IRR (right axis). As PPA size increases, excess wind power is harder to monetise, leading to higher LCOH and lower IRR.

(b) Hydrogen production shares from wind, battery, and market electricity as a function of PPA size. Larger PPA capacities increase the share of hydrogen produced directly from wind, while smaller PPAs result in greater reliance on grid electricity and battery support.

Figure 7.7: Sensitivity of hydrogen system performance to PPA size. Subfigure (a) illustrates the impact of PPA size on the LCOH and the IRR, with LCOH shown in blue (left axis) and IRR in green (right axis). Subfigure (b) shows how the share of hydrogen produced using wind, battery, and market electricity varies with PPA size. As PPA size increases, wind usage dominates, while smaller PPAs rely more heavily on battery support and grid electricity.

One key observation was that reducing the PPA size leads to a corresponding decrease in the share of hydrogen produced using wind power. As shown in Figure 7.7b, market electricity begins to dominate the input mix for hydrogen production when the PPA is small, and the battery becomes more critical in supporting operation. As PPA size increases, the wind contribution to hydrogen production grows rapidly before levelling off around 80%, suggesting the optimiser still preserves some flexibility to benefit from low-cost market prices.

Tracking these shares is essential to determine the "greenness" of the hydrogen. A high share of renewable-based electricity is required for the hydrogen to qualify as renewable or low-carbon under various certification schemes, such as RFNBO. While grid electricity can be partially renewable, particularly during periods of high solar or wind output, it is not consistently green. The carbon intensity of grid-sourced electricity varies significantly depending on the generation mix at any given time. As a result, hydrogen produced using market electricity cannot automatically be assumed to meet green hydrogen standards unless time-matching or grid carbon-intensity tracking is explicitly incorporated. Therefore when market-sourced electricity becomes dominant, as is the case with smaller PPAs, the hydrogen would no longer meet these green criteria and would likely not be worth the green hydrogen sales price.

From an economic perspective, Figure 7.7a shows that increasing PPA size results in a higher LCOH and a lower IRR. This is because a large PPA often leads to excess wind power that cannot be used directly and must be sold back to the grid at unfavourable rates, decreasing overall system profitability. In contrast, smaller PPAs allow the system to rely more on market electricity and strategically optimise hydrogen production around low-cost hours, which boosts IRR.

However, this outcome highlights a mismatch between model assumptions and market reality. The model applies a uniform hydrogen sales price regardless of how the hydrogen is produced, but in practice, hydrogen made from predominantly grid electricity would not be sold at the same premium price as green hydrogen. As such, while the model outputs are correct under the assumed price structure, they likely overestimate the real-world profitability of systems with small PPAs and high market electricity use. To correct for this, we adjust the hydrogen sales price in later analyses when reducing the PPA size.

7.2.3. Smaller PPA sizes: increased variation between DO, RHO and SORH solutions

To better understand how modeling assumptions around foresight and uncertainty influence outcomes, we conducted additional analyses for reduced PPA capacities. Specifically, we evaluated two new test cases: a 50 MW wind PPA and a zero-PPA configuration where all electricity is sourced from the market. In line with earlier findings, a reduced PPA leads to a lower share of renewable-sourced hydrogen. To reflect this, we adjusted the hydrogen sales price downward to more realistic levels.

50 MW wind PPA and €250/MWh hydrogen sales price

We first ran simulations for the 50 MW PPA case using a hydrogen sales price of €250/MWh. The results for this configuration are presented in Tables 7.12 and 7.13. However, due to the still relatively favourable

margins, electrolyser utilisation remained high, and the differences between the modeling approaches were limited.

Table 7.12: Comparison of results from the different modeling approaches: DO, RHO and SORH. Results are generated for a 50 MW wind PPA, a 50 MW electrolyser, and a 100 MWh battery. Hydrogen sales price is €250/MWh (\approx €8.30/kg). Compared to the base case, the reduction in renewable supply leads to more pronounced variability in outcomes between modeling approaches, though the overall differences in performance remain moderate.

Model	# Scen.	H_2 prod. (kt)	Util. Rate (%)	# shut-offs	Market Rev. (M€)	Profit (M€)	LCOH (€/kg)	IRR (%)	Runtime (s)
DO	-	7.25	92.0	8	9.29	8.40	7.73	14.46	11
RHO	-	7.20	91.3	7	4.78	8.32	7.34	14.40	23
SORH	5	7.27	92.2	3	4.21	8.31	7.75	14.9	116
SORH	10	7.22	91.5	5	4.98	8.24	7.75	14.3	239
SORH	20	7.20	91.3	5	4.48	8.32	7.79	14.4	867

Table 7.13: Relative difference in KPI's between models: DO, RHO, SORH. Results for smaller wind PPA size and lowered hydrogen sales price concept: for a PPA size of 50 MW and hydrogen sales price of €250/kg (\approx €8.30/kg). The values under RHO and SORH indicate their relative differences to the DO base case.

Metric	DO (base)	Δ RHO	Δ SORH
H_2 produced (kt)	7.25	-0.69 %	-0.69%
Profit (M€)	8.4	-0.95%	-1.9%
Market revenue (M€)	4.29	+4.24%	+1.86%
IRR (%)	14.46	<-0.5%	-1.1%

50 MW wind PPA and €210/MWh hydrogen sales price

To test whether more constrained conditions would reveal larger deviations, we repeated the analysis using a lower hydrogen sales price of €210/MWh. As shown in Tables 7.14 and 7.15, this led to more noticeable variation between the DO, RHO, and SORH results.

Table 7.14: Comparison of results from the different modeling approaches: DO, RHO and SORH. Results are generated for a 50 MW wind PPA, 50 MW electrolyser, and 100 MWh battery concept. Hydrogen sales price is set at €210/kg (\approx €7/kg). Lowering the hydrogen sales price reduces the profit margin, resulting in lower utilisation rates and in turn more variability between solution approaches.

Model	# Scen.	H_2 prod. (kt)	Util. Rate (%)	# shut-offs	Market Rev. (M€)	Power Cost (M€)	Profit (M€)	LCOH (€/kg)	IRR (%)
DO	-	6.38	80.9	14	7.53	28.7	-0.8	7.40	7.13
RHO	-	6.48	82.2	10	7.22	29.4	-0.85	7.40	7.10
SORH	5	6.80	86.2	6	6.14	31.9	-0.96	7.39	7.00
SORH	10	6.41	81.2	18	7.47	29.57	-0.92	7.41	7.04

Table 7.15: Relative difference in KPI's between models: DO, RHO, SORH. Results for smaller wind PPA size and lowest hydrogen sales price concept: a PPA size of 50 MW and hydrogen sales price of €210/MWh. The values under RHO and SORH indicate percentage changes relative to the DO base case.

Metric	DO (base)	Δ RHO	Δ SORH
H_2 produced (kt)	6.38	+1.6%	+6.6%
Profit (M€)	-0.8	-6.2%	-20%
Market revenue (M€)	7.5	-4%	-18.7%
IRR (%)	7.13	<-0.5%	-1.8%

No wind PPA and €210/MWh hydrogen sales price

Finally, we tested the most market-exposed scenario by removing the wind PPA entirely (0 MW), while maintaining the hydrogen price at €210/MWh. This setup, summarised in Tables 7.16 and 7.17, showed the largest differences in both operational strategy and financial performance across the modeling approaches.

Table 7.16: Comparison of results from the different modeling approaches: DO, RHO and SORH. Results generated for the no-wind PPA system concept, where all power is sourced from the market and the system includes a 50 MW electrolyser and a 100 MWh battery. Hydrogen sales price is set to €210/MWh (\approx €7/kg). This case results in largest difference in both operational strategy and financial performance across the modeling approaches.

Model	# Scen.	H ₂ prod. (kt)	Util. Rate (%)	# shut-offs	Market Rev. (M€)	Profit (M€)	LCOH (€/kg)	IRR (%)	Runtime (s)
DO	-	6.37	80.8	14	1.94	7.34	7.15	9.0	10
RHO	-	6.56	83.7	8	1.69	1.24	7.16	8.9	20
SORH	5	6.79	86.2	6	1.50	1.19	7.16	8.84	60
SORH	10	6.40	82.1	17	1.99	1.23	7.16	8.9	127
SORH	20	6.29	79.8	14	2.04	1.31	7.15	8.94	228

Table 7.17: Relative difference in KPI's between models: DO, RHO, SORH. Results for no-wind system concept. This case results in largest difference in both operational strategy and financial performance across the modeling approaches.

Metric	DO (base)	Δ RHO	Δ SORH
H ₂ produced (kt)	6.37	+2.9%	+6.2%
Profit (M€)	1.34	-8.2%	-15%
Market revenue (M€)	2.94	-12.9%	-22.7%
IRR (%)	9.0	-1.1%	-1.8%

Effect of reduced foresight and added uncertainty

The variation between optimisation models becomes most pronounced in test cases with limited or no wind PPA. As observed in Tables 7.15 and 7.17, the difference in hydrogen production, profit and market revenue increases significantly as the system becomes more exposed to electricity market prices. This effect is particularly evident in the no-PPA configuration, where all electricity must be sourced from the market.

This can be explained by the fact that, in the presence of a large PPA, the system benefits from a steady and fixed-price renewable electricity supply. This reduces exposure to market price volatility and diminishes the value of precise forecasting and scheduling. The fixed-price wind PPAs hedge against future price volatility. Since PPA-sourced power is paid for regardless of whether it is used, the optimiser is incentivised to always consume, store or sell it in real-time. As such, foresight plays a limited role, especially when wind availability is high. Conversely, when relying on market power alone, dispatch decisions depend more on future price expectations. Under such conditions, reduced foresight (RHO) and uncertainty (SORH) have a higher impact on dispatch strategy and profitability. The SORH model shows up to a 15% profit drop compared to DO in the no-PPA case, despite a 6% increase in hydrogen production. This suggests that planning under uncertainty leads to overproduction to hedge against unfavourable outcomes, often at the cost of economic efficiency.

Operational differences: The most notable operational difference between the modeling approaches lies in the scheduling of the electrolyser. The DO model, with perfect foresight over the entire year, can optimise hydrogen production by concentrating operations during the most cost-effective hours and minimising start-stop cycles. In contrast, both the RHO and SORH models operate within a reduced one-week planning window, requiring them to meet weekly hydrogen production targets based on imperfect or uncertain information about future prices and wind availability. As a result, these models tend to exhibit slightly higher annual utilisation rates, but also increased variability in start-stop behaviour. For instance, the number of annual shut-offs varies from 14 in the DO model to as few as 3 or as many as 18 in SORH, depending on the specific scenario set used (Table 7.14).

Although the total downtime remains nearly identical across models, the timing of the shut-downs (mainly for the scheduled maintenance hours) differs. The DO model can optimally distribute the required 300 hours

of maintenance throughout the year, selecting periods of low profitability. In contrast, the RHO and SORH models apply a fixed downtime window of 300 consecutive hours, which is pre-scheduled based on a heuristic that targets the most expensive hours of the year. While this ensures comparability between approaches, it does implicitly use year-ahead price information to approximate the outcome achievable with perfect foresight.

Interestingly, while the total amount of power charged and discharged by the battery remains relatively consistent across all modeling approaches, we observe clear differences in market revenue. For instance, in the no-PPA case (Tables 7.16 and 7.17), market revenues differ by up to 22.7% between the DO and SORH models. This suggests that although the battery's overall activity (in terms of energy throughput) is largely unaffected by forecast uncertainty, the strategic timing of discharging (particularly when to sell electricity back to the market) is influenced by the modeling assumptions. In other words, the value extracted from the battery depends not just on how much is stored or released, but also on when these actions occur. These timing decisions, especially under uncertain market conditions, have a direct impact on realised market revenues.

However, despite these operational differences, the overall impact on long-term financial KPIs such as IRR and LCOH remains modest. This suggests that while operational behaviour is affected by the modeling approach, the long-term profitability is relatively stable.

Value of rolling horizon vs perfect foresight (DO vs RHO): When comparing DO and RHO, we isolate the effect of reducing the foresight window without introducing uncertainty. The DO model can schedule hydrogen production based on global optimisation across the full year, while RHO optimises week by week. As a result, RHO models show slightly conservative behaviour and tend to slightly overproduce hydrogen to ensure that weekly demand targets are met, leading to marginally higher utilisation and production levels but also higher power costs. In our analysis, the profit difference between DO and RHO was generally within 1–6% depending on the size of the PPA contract.

Impact of forecast uncertainty (RHO vs SORH): The SORH model further adds uncertainty by incorporating random errors into price and wind forecasts beyond the first 24 hours. This forces the optimiser to take a more cautious approach, often producing hydrogen earlier in the planning horizon to guard against low-wind or high-price scenarios later in the week. While this strategy increases hydrogen output and electrolyser utilisation, it often leads to higher production costs and lower market revenues. The profit loss relative to RHO is typically 1–3%, although it can reach up to 7–15% in the no-PPA case. These findings suggest that uncertainty-aware optimisation introduces realism into the operational strategy, but also reduces system profitability due to increased conservatism.

Effect of scenario number in SORH: We also tested the effect of increasing the number of scenarios used in the SORH model. Results indicate that using more scenarios generally brings the SORH outcome closer to that of the RHO model. This convergence is likely because with more scenarios, the average outcome approximates the expected value more closely. Since the distribution of the errors terms has zero mean, the theoretical expected value of the forecast scenarios are equal to the actual values. Moreover, since in the optimisation solution the decisions are averaged across all scenarios (which have equal weights), additional scenarios reduce the chance that a single extreme case distorts the initial 24-hour dispatch decisions. Usually we saw the run with least scenario's, namely 5, showed most deviation. Notably, increasing the number of scenarios also significantly increases runtime, suggesting a trade-off between accuracy and computational effort.

Long-term profitability implications

Although the DO, RHO, and SORH models differ in operational strategies, particularly in the scheduling of hydrogen production, their impact on long-term financial performance remains modest. Across all test cases, IRR values rarely differ by more than two percentage points, and the LCOH stays within a narrow range of approximately €0.10/kg. These findings suggest that for long-term investment evaluations, the deterministic DO model provides a sufficiently accurate approximation of system viability, particularly in configurations with a large, fixed-price renewable PPA.

Nevertheless, the DO model may slightly overestimate annual profits in setups exposed to high market price uncertainty. In such cases, adjusting DO-based financial outputs downward by approximately 5–10% can provide a more realistic benchmark. In addition, while DO captures the strategic essence of hydrogen

production under stable conditions, RHO and SORH offer a more operationally robust view, especially for short-term dispatch decisions in uncertain or volatile environments.

It is important to acknowledge, however, that the stochastic scenarios used in the SORH model were based on full-year historical distributions of wind and price volatility. In practice, this level of information would not be available in real time. Therefore, the results of the SORH model are most meaningful in a retrospective or planning context, rather than as a live operational tool. For real-time dispatch applications, a more dynamic model would be required. We would need one that continuously updates forecasts based on incoming data and adapts to rapidly changing market conditions. This would necessitate the development of a dedicated forecasting module capable of capturing near-term variability in electricity prices and wind availability, and integrating these updated predictions into the optimisation framework.

In summary, while the deterministic DO model remains a practical and computationally efficient choice for long-term planning and concept testing, future work aimed at operational implementation should consider more sophisticated stochastic control models paired with real-time data assimilation and forecasting capabilities.

7.3. Impact of differentiated hydrogen sales prices (green vs grey) on system operation and profitability


This section evaluates the effect of introducing differentiated hydrogen sales prices based on the electricity source: wind (green) versus grid (grey). This extension was developed to resolve the previously identified mismatch between the assumed hydrogen price and the electricity mix used for its production. By explicitly modeling separate revenue streams for green and grey hydrogen, the optimisation model now better reflects market realities and makes rational decisions on whether or not to use grid electricity for the electrolyser.

Although grid-sourced hydrogen can be cheaper to produce during periods of low electricity prices, it generally commands a lower market price due to its higher carbon footprint. However, utilising grid electricity can still be economically attractive in certain situations, such as when wind availability is insufficient or when minimising the number of start-up and shut-down cycles is prioritised. In these cases, supplementing wind power with grid electricity allows for smoother electrolyser operation and more consistent hydrogen production. To examine the impact of source-dependent hydrogen pricing, we conducted an exploratory optimisation run with this model extension. The model now includes differentiated hydrogen prices in the objective function and dynamically optimises production based on prevailing market conditions. Figure 7.8 illustrates the change in system dynamics over time. We compare these new outcomes with the results from the old model presented in Section 7.1.2. The detailed results of this analysis are presented in the Appendix B, where Tables B.1 and B.2 present the battery and electrolyser operation characteristics that support the discussion in this section.


Electrolyser operation: The results show a clear shift in both operational and financial performance. In terms of electrolyser operation, the average utilisation rate decreased slightly by 1%, while annual hydrogen output dropped from 7.5 kt to 7.3 kt. Although the total number of operational hours remained nearly unchanged, the hourly dispatch pattern became notably more variable as can be seen in Figure 7.8. Previously, the electrolyser operated predominantly at full capacity when active. With the revised pricing model, it now more frequently runs at intermediate load levels, as the model avoids full utilisation during hours when only grey electricity is available and profit margins are insufficient. This change in dispatch strategy also resulted in fewer start-stop cycles: 10 in the current run compared to 15 in the previous case.

Battery operation: Battery usage increased considerably in the updated model. Both the total energy charged and discharged grew by over 30%, and the share of battery output directed to the electrolyser rose substantially. This shift reflects the battery's role as a buffer that enables the system to store low-cost electricity and use it strategically for hydrogen production during hours of lower renewable availability or less favourable market conditions. Notably, we see more power from the grid being stored in the battery than previously. Importantly, in the current modeling setup, electricity discharged from the battery to the electrolyser is assumed to qualify as green, an assumption based on the premise that battery charging occurs primarily with wind power or during periods of low grid carbon intensity. However, this may represent a modeling simplification, as it allows the system to route market power through the battery to capture the green hydrogen premium. The observed increase in battery use appears to exploit this feature, and suggests a potential area for refinement in future certification logic.

7.3. Impact of differentiated hydrogen sales prices (green vs grey) on system operation and profitability Results



(a) Operational results for the old model with a single hydrogen sales price. The electrolyser operates at more stable utilisation levels (mostly 100%, 30% or off), and the battery system is used less frequently. Market power is used extensively to run the electrolyser, as all hydrogen is sold at the same fixed price.



(b) Operational results for the updated model with differentiated hydrogen sales prices. The electrolyser shows more variable utilisation, responding to price signals, and there is a notable increase in battery usage to enable market sales and support the flexible dispatch strategy. Compared to the old model, we observe reduced use of market electricity for the electrolyser since hydrogen produced from market power is now sold at a lower price. The system hence prefers to use battery power or ramp down the electrolyser instead.

Figure 7.8: Comparison of dispatch patterns under different hydrogen sales price assumptions. Introducing multiple H₂ prices in the model based on power source (grid or wind source) leads to a more dynamic operational strategy: electrolyser utilisation becomes more variable, and the battery plays a more active role in storing and shifting electricity. Market-sourced power is used more selectively for hydrogen production, as its associated hydrogen is sold at a lower price. This incentivises a shift towards maximising the use of renewable energy and battery storage in the system's dispatch decisions.

Table 7.18: KPI results for the differentiated hydrogen sales price model, including the projects NPV, IRR and LCOH. Relative to the base case results in Table 7.9, we observe a decrease in profitability: NPV drops from €97.2M to €87.6M, and IRR decreases from 21.1% to 20.1%. This is expected, as hydrogen produced using market electricity is now sold at a lower price, reducing revenues. Interestingly, the LCOH slightly decreases from €8.40/kg to €8.37/kg, indicating that the cost per unit of hydrogen remains nearly unchanged despite the drop in profits.

Metric	Unit	Value
Net Present Value (NPV)	M€	87.57
Internal Rate of Return (IRR)	%	20.10
Levelized Cost of Hydrogen (LCOH)	€/kg	8.37

The financial impact of introducing differentiated hydrogen sales prices is notable. Annual profit drops from €18.62 million in the base case, where all hydrogen was sold at a single fixed price, to €14.48 million under the new pricing scheme, representing a reduction of approximately 22%. This change is primarily driven by a €5 million decrease in hydrogen revenue, partially offset by a €1.1 million increase in electricity procurement costs. Overall, the updated model provides a more realistic estimate of profitability by reflecting that not all hydrogen can be sold at the premium green price.

Table 7.18 presents the new models NPV, IRR and LOCH results. Despite this change in revenue, the LCOH remains relatively stable at around €8.37/kg, indicating that the average production cost per kilogram of hydrogen is unaffected by the pricing strategy. The IRR, however, declines by 2.8 percentage points, further underlining the impact on project attractiveness. The observed profit reduction is primarily due to the share of hydrogen now sold at a lower grey price, which directly lowers total revenue.

To further investigate the implications of this updated pricing model, we repeated the sensitivity analysis on PPA size using the new dual-pricing setup. The results, presented in Figure 7.9, reveal a notable shift in the optimal PPA capacity compared to the previous analysis. Under the earlier model, which assumed a uniform hydrogen sales price, the IRR declined consistently with increasing PPA size due to diminishing marginal returns from surplus renewable energy. However, the revised model produces a more intuitive outcome: the IRR exhibits a clear maximum between 80 and 100 MW, reflecting the improved balance between green hydrogen production, renewable power utilisation, and economic efficiency. This concave trend indicates that smaller PPAs underperform due to over-reliance on lower-priced grey hydrogen, while excessively large PPAs introduce curtailment losses and reduce returns. The LCOH curve, meanwhile, continues to rise with PPA size as expected, due to the increasing volume of unused or inefficiently used renewable electricity. The optimal PPA size, resulting in highest IRR, is seen to be between 80-100 MW.

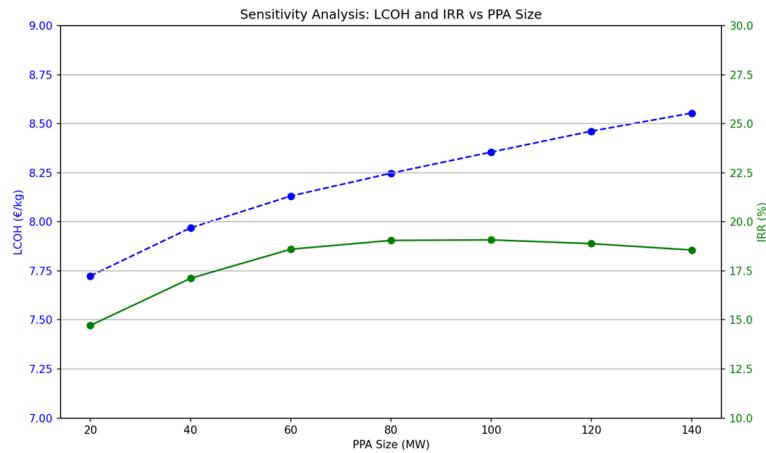


Figure 7.9: Sensitivity of the LCOH and IRR to PPA size under a differentiated hydrogen pricing model. The blue, dashed line (left axis) represents the LCOH in €/kg, while the green line (right axis) shows the IRR as a percentage. Compared to the base case with uniform pricing, the IRR curve now exhibits a clear maximum between 80 and 100 MW, reflecting an optimal trade-off between renewable utilisation and economic efficiency. The LCOH increases steadily with PPA size due to greater volumes of unused or low-value renewable electricity.

These findings underscore the importance of incorporating differentiated hydrogen sales prices into the model. By introducing separate prices for grid- and wind-sourced hydrogen, the model now better captures the incentives and trade-offs faced by a real-world operator. Under this revised framework, the battery plays a more strategic role. Rather than primarily arbitraging electricity prices, it now acts as a critical enabler of low-carbon hydrogen production. By storing renewable electricity and discharging it to the electrolyser when wind availability is low or market prices are otherwise unfavourable, the battery helps maintain hydrogen output while preserving green certification status. This operational flexibility becomes particularly valuable in the context of constrained renewable supply.

Given this shift in the battery's functional role and economic relevance, the updated model with differentiated hydrogen pricing is used in all subsequent analyses that evaluate the battery's impact on long-term profitability. This adjustment ensures a more accurate and realistic assessment of system performance and investment value in scenarios aligned with emerging hydrogen market regulations. Based on the revised sensitivity, we keep a 100 MW PPA as the representative base case for the subsequent battery value analysis, as it captures a realistic and financially robust system configuration under differentiated hydrogen pricing conditions.

Adjusted MIP gap for solving

Introducing differentiated hydrogen prices significantly increased the model's runtime, particularly when grey hydrogen prices were set to lower levels. To ensure computational efficiency during multiple test runs, we relaxed the MIP gap in gurobi from 0% to 0.1%, which substantially reduced the solution time without noticeably impacting result quality.

7.4. Analysis battery profitability

This section evaluates the role and profitability of battery storage under different system setups and across multiple historical years.

7.4.1. Optimal battery size

The capacity of the battery plays an important role in determining its economic value. Larger batteries offer increased flexibility for power trading, allowing more electricity to be stored and sold at more favourable times. However, as battery size increases, so do capital and operational costs. This creates a trade-off between the operational value of increased storage and its associated costs.

To identify the optimal balance, we initially attempted to extend the optimisation model by treating the battery capacity as a free variable and including unit capital costs in the objective function. However, using the base case cost parameters, the model consistently returned an optimal battery size of zero. This outcome indicated that, under current assumptions, battery investment is not economically viable. Only when the unit capital costs of the battery were reduced by approximately 50% did the model suggest a positive battery capacity. At that point, the optimal size jumped directly to the upper limit set in the model.

Moreover, due to the combination of sizing decision variables as well as hourly dispatch decision variables, this optimisation proved computationally intensive. Since all operational dispatch decisions (such as charging, discharging, and market interaction) are conditional on the installed battery size, allowing the model to optimise this size introduces significant additional complexity. Solve times increased substantially, making the process impractical for iterative testing.

Consequently, battery sizing was explored through a simplified, trial-and-error approach. Several discrete battery capacities were tested, and the corresponding models were solved individually. This effectively constituted a basic grid search to identify the capacity yielding the highest annual profit. While this approach lacks efficiency and precision, it was deemed appropriate given that optimal battery sizing was not the primary focus of this thesis. A more refined analysis, for instance using Bayesian optimisation or an adaptive grid search method, could have yielded faster and more accurate results but was outside the scope of this research.

Additionally, combining battery sizing with sensitivity to other structural parameters (such as wind PPA capacity) would further increase model complexity. Simultaneously treating both battery and PPA size as free optimisation variables would significantly expand the solution space, resulting in excessive computational burden. To maintain tractability and focus on the key research questions, the analysis proceeded with a fixed battery configuration of 100 MWh (four hours of duration, 25 megawatt charge and discharge rating). This size was selected based on the exploratory profit comparison and was used for all subsequent analyses involving battery integration.

7.4.2. Assessment economic impact of battery inclusion

To assess the economic value of battery storage under this configuration, we compared results from two model runs: one with the 100 MWh battery included, and another where battery capacity was set to zero. Operational and financial results for the no-battery case are presented in Appendix C.1.

The results show that adding a battery reduces total power costs by increasing electricity sales to the market, which offsets power purchases. However, this benefit is outweighed by the additional capital and operational costs of the battery system. Specifically, the no-battery model results in a €0.30/kg lower LCOH and a three percentage point higher IRR. This suggests that, despite higher market revenues, battery deployment is not profitable under baseline assumptions.

A sensitivity analysis on battery capital costs confirms this finding. The battery CAPEX would need to decrease by approximately 55 % to achieve the same LCOH as the no-battery configuration. This aligns with the earlier sizing experiment, which also found that battery integration becomes economically viable only after significant cost reductions.

7.4.3. Performance across different years

Several factors influence these results, including market power price levels and volatility. Since power market conditions vary across years, we evaluated the profitability of the battery under different historical wind and electricity price profiles. Table 7.19 presents the average electricity market prices and wind load factors for the years 2019 to 2024, which form the basis of the comparative analysis. The corresponding operational and economic performance results of the base case model are summarised in Table 7.20. Chapter 4 further

examines the distributions of electricity prices and wind generation across these years. In summary, the years 2019 and 2020 are characterised by relatively low electricity prices and low price volatility. From 2021 onwards, both price levels and volatility increase, with 2022 standing out due to the energy crisis that led to extreme market conditions. These dynamics are reflected in the model results: in high-price, high-volatility years such as 2022, market revenues rise substantially, and electrolyser utilisation decreases as the system engages more actively in electricity trading.

Nevertheless, the highest overall profitability is not observed in these volatile years, but rather in the early years (2019 and 2020), where low electricity prices lead to more favourable margins between electricity costs and hydrogen sales revenues. Notably, the wind PPA price is adjusted annually based on the market capture price, as discussed in Section 4.7.4, to ensure fair comparison across years. Interestingly, the base case year 2023 performs the worst in terms of profitability, with the highest LCOH and lowest IRR. This is likely due to elevated electricity prices without sufficient volatility to generate the high market trading revenues observed in 2022.

Table 7.19: Annual average market and wind conditions for different years.

Year	Average Electricity Market Price (€/MWh)	Average Wind Load Factor (%)
2019	41.20	33.18
2020	32.24	50.09
2021	102.96	45.36
2022	241.92	46.19
2023	95.82	51.04
2024	77.29	48.00

Table 7.20: Summary of annual performance metrics for different historical years. Columns include the KPIs LCOH, NPV, IRR. Additionally, the table reports the total market revenue from electricity trading, the average utilisation rate of the electrolyser and the total number of full shut-off of the electrolyser throughout the year.

Year	LCOH (€/kg)	NPV (M€)	IRR (%)	Profit (M€)	Market Revenue (M€)	Utilisation (%)	# shut-off
2019	7.02	144.54	26.16	25.24	1.80	96.58	13
2020	7.54	127.74	24.39	22.61	4.21	96.57	12
2021	7.56	105.12	21.99	19.07	21.01	83.23	20
2022	7.26	106.45	22.13	19.28	50.76	68.49	20
2023	8.36	78.15	19.08	14.48	12.45	92.74	10
2024	7.92	101.47	21.60	18.49	8.82	95.4	10

Battery deployment remains only marginally profitable under typical market conditions but becomes significantly more valuable in years with heightened price volatility. In 2022, characterised by extreme price spikes, the battery-inclusive model outperforms the no-battery configuration. In this scenario, the system benefits from increased electricity trading, resulting in higher market revenue and lower total power costs (see Table 7.21). As a result, the LCOH decreases by €0.40/kg, substantially improving the business case for battery integration. A similar improvement is observed in 2021.

However, 2021 and 2022 are the only years in which battery integration results in improved profitability. In all other years, the no-battery configuration outperforms the battery-inclusive model. Although battery usage improves dispatch flexibility and lowers electricity costs through better market timing, the high capital and operational costs associated with battery systems outweigh these benefits under most market conditions.

Overall, these findings suggest that battery systems can improve the economic performance of hydrogen production facilities, but only under extreme and volatile market conditions. Unless such conditions become more frequent, battery investments are unlikely to significantly enhance the long-term profitability of these systems.

Table 7.21: Breakdown of LCOH components by year, expressed in €/kg. The table presents individual cost components contributing to the final LCOH, including electricity costs from the PPA and market, electricity revenues from power sales, capital and operational expenditures for the electrolyser and battery system, shut-off related costs, redispatch and overcapacity penalties and tax contributions. Notably, electricity-related costs (both from procurement and revenues from trading) vary the most across years. Years with high electricity prices and volatility, such as 2022 and 2021, show the highest electricity procurement costs but also the highest electricity market revenues, partially offsetting those costs in the final LCOH.

Component	2019	2020	2021	2022	2023	2024
Electricity PPA	3.05	4.45	5.87	9.74	5.81	4.91
Electricity Market	0.93	0.46	1.31	2.56	1.18	1.09
Electricity Sales	-0.24	-0.55	-3.20	-4.00	-1.70	-1.17
Capex Electrolyser	1.35	1.35	1.57	1.90	1.41	1.37
Capex Battery	0.44	0.44	0.51	0.61	0.45	0.44
Opex Electrolyser	0.29	0.29	0.33	0.41	0.30	0.29
Opex Battery	0.10	0.10	0.12	0.15	0.11	0.11
Shut-off Costs	0.01	0.01	0.02	0.03	0.01	0.01
Redispatch / Overcapacity	0.00	0.00	0.00	0.01	0.02	0.00
Tax	1.09	1.00	1.02	1.25	0.77	0.88
LCOH (post-tax, excl. power sales)	7.26	8.09	10.76	16.66	10.06	9.09
Final LCOH (incl. power sales)	7.02	7.54	7.56	7.26	8.36	7.92

7.4.4. Test case: no wind PPA

The base case analysis assumes a 100 MW wind PPA, providing a steady supply of renewable electricity to power the electrolyser. To evaluate whether the value of battery storage increases when the system is more exposed to market price fluctuations, a comparative analysis was conducted using a test case without a PPA. In this configuration, the system relies solely on electricity procured from the day-ahead market, increasing its sensitivity to price volatility.

The results, however, mirror those observed in the base case: the no-battery configuration outperforms the battery-inclusive model in all years except 2022. This indicates that even under greater exposure to market dynamics, battery deployment remains economically unfavourable under typical conditions.

Moreover, the no-PPA configuration results in significantly lower overall system performance. The year 2022, in particular, stands out as a catastrophic case for the hydrogen business model when relying exclusively on market electricity. Extremely high electricity prices lead to severely negative profitability, with the NPV dropping to approximately €-200 million. This loss would more than offset the cumulative profits of the preceding three years, had this system been deployed in real-world conditions.

These results highlight the critical role of long-term fixed-price contracts, such as wind PPAs, in stabilising electricity costs and shielding the system from extreme market price spikes. Without such contracts, the hydrogen production business case becomes highly vulnerable to market risk. The full set of results supporting this analysis is provided in Appendix C.2.

8

Conclusion and Discussion

This thesis investigated the operational and economic performance of a hybrid renewable energy system for green hydrogen production via water electrolysis. The system combines power supplied through a long-term Power Purchase Agreement (PPA) for wind energy, battery energy storage and access to the day-ahead electricity market. The central objective was to determine profit-maximising dispatch strategies and assess the long-term profitability of different system configurations under varying external conditions and modeling assumptions. To evaluate long-term profitability, the two key financial metrics used were: the Levelised Cost of Hydrogen (LCOH), which reflects the average cost to produce one kilogram of hydrogen over the system's lifetime, and the Internal Rate of Return (IRR), which indicates the overall profitability of the investment.

The conclusions are structured in two parts. First, in Section 8.1 the key findings are summarised regarding the choices for optimisation approaches. Second, in 8.2, the impact of adding a battery and participation in the day-ahead market is evaluated.

8.1. Optimisation method hybrid energy systems

This section refers back to the first main research question presented in the introduction:

Main research question 1: *How can one determine the optimal dispatch strategy of a hybrid energy system that combines renewable electricity, energy storage, hydrogen production, and participation in the day-ahead market, for the purpose of long-term profitability assessment?*

To address this question, a mathematical modeling framework was developed to simulate and optimise the dispatch of the hybrid system responding to variable electricity market prices and wind generation. A system-level approach was essential in order to capture the complex interdependencies between the electrolyser, wind turbine, battery as well as the impact of external market signals on overall system performance.

The model was formulated as a Mixed-Integer Linear Program (MILP), allowing both continuous decision variables (such as power flows) and binary variables (such as the on/off status of the electrolyser). The MILP formulation enabled efficient and scalable solution methods, making use of its compatibility with the solver Gurobi. The optimisation was performed over a full-year horizon with hourly resolution, which allowed the model to represent temporal dynamics in wind availability, electricity prices and operational behaviour in a detailed way. The objective function maximised annual system profit by optimising the hourly dispatch of the electrolyser, battery and grid trading. The model included constraints on power flows, component capacities, operating limits of the electrolyser and battery, and hydrogen production targets. While certain technical behaviours, such as nonlinear efficiency curves and component degradation for the electrolyser and battery, were simplified or omitted, these choices were made to prioritise computational efficiency and enable strategic insights at the system level. These simplifications are, however, a limitation for studies that require more detailed modeling of individual components.

A detailed discussion of the optimised dispatch behaviour can be found in Section 7.1 of the Results chapter and can be summarised as follows: the optimised dispatch strategies reveal a flexible and economically driven response to variable market and wind conditions. When wind is abundant and electricity prices are low, the system prioritises hydrogen production and manages surplus power through battery storage or mar-

ket sales, sometimes even at a loss, to avoid high electricity redispatch penalties. During low-wind, high-price periods, it shifts to market participation and therefore frequently charges and discharges the battery to capitalise on market price differences, with the electrolyser strategically shut down or maintained at minimum load. When large wind overcapacity coincides with high market prices, redispatch penalty costs are occasionally accepted to avoid worse economic outcomes. In volatile conditions, the system combines all available sources to stabilise electrolyser operation, dynamically adjusting to short-term fluctuations.

Different optimisation approaches outcomes

To examine the influence of uncertainty and planning horizon on dispatch strategy and financial performance, three optimisation approaches were developed and compared: a deterministic optimisation model with perfect foresight (DO), a deterministic model solved using a rolling horizon with limited foresight (RHO), and a scenario-based stochastic optimisation model with a rolling horizon and two-stage structure (SORH).

All three models produced operationally feasible dispatch strategies. However, differences emerged in terms of profitability, operational patterns and computational complexity. As expected, the deterministic DO model yielded the highest annual profits due to full foresight, which enabled ideal scheduling of hydrogen production during low-price hours. The RHO and SORH models, achieved lower profitability and demonstrated more conservative operational behaviour.

From an operational perspective, the main effect of limiting foresight or introducing uncertainty was observed in the scheduling of the electrolyser. Without full-year perfect foresight, the model could no longer concentrate hydrogen production in the most cost-optimal hours across the entire year. As a result, the rolling horizon models (RHO and SORH) displayed notably different operational behaviour, including higher average utilisation rates and more variable shut-off patterns. These models tended to slightly overproduce hydrogen on a weekly basis to ensure that demand targets were met within the restricted planning window, leading to a more conservative yet robust dispatch strategy.

In terms of complexity and computational performance, the DO model was the most efficient, solving the full-year optimisation problem in a single run with the shortest overall computation time. The RHO model also maintained relatively fast solve times per iteration; however, it introduced complexity in other ways. Instead of solving one comprehensive annual problem, RHO required repeated optimisation over shorter planning windows, which necessitated adjustments to several model constraints. For example, the annual hydrogen production target had to be reformulated as a weekly target, and yearly constraints such as the maximum number of electrolyser shut-offs were also applied on a per-week basis. These adaptations reduced model flexibility and introduced additional tuning challenges. Moreover, this segmented approach complicates potential integration with system design optimisation, where a single annual optimisation is typically preferred. Lastly, the SORH model imposed a significantly higher computational burden than the DO and RHO models, primarily due to the need to solve multiple scenarios at each optimisation step. In particular, the solve time increases sharply as the number of scenarios grows.

To assess whether this additional effort yields meaningful improvements, the impact of number of scenario's on model behaviour was analysed. We observed that increasing the number of scenario's leads to convergence toward the RHO result. This is in line with theoretical expectations, since the forecast scenarios in the SORH model are constructed by adding zero-mean, correlated random errors to the historical values used in the RHO model. When the number of scenarios is low, for example 5, a single outlier can skew the expected value and lead to different short-term decisions. As the number increases, these deviations are averaged out since the expected value across a large number of scenarios should theoretically approximate the deterministic input of the RHO model.

DO as the preferred tool for long-term assessments

Although the DO, RHO and SORH models differed in operational strategies, the resulting differences in long-term financial performance were modest, as shown in Section 7.2. Across all test cases, the IRR values varied by no more than two percentage points and the LCOH remained within a narrow range of approximately €0.10/kg. These results indicate that for long-term investment evaluations, the DO model offers a sufficiently accurate approximation of system profitability. This is particularly true for system configurations with a large fixed-price PPA, where there is limited reliance on grid power and therefore less exposure to price uncertainty. We conclude that the DO model remains a valid and efficient tool for strategic planning and long-term profitability assessment of hybrid hydrogen systems. Nevertheless, the DO model may slightly overestimate annual profits in setups exposed to high market price uncertainty.

RHO and SORH for real-time operations modeling

While the DO model is well suited for long-term investment evaluation, the RHO and SORH approaches offer a more operationally realistic framework for short-term system control under uncertainty. These models are better aligned with real-time operations, where only limited foresight is available and decisions must adapt to forecast variability in electricity prices and renewable output. In system configurations with greater exposure to market dynamics or tighter operational constraints, these models provide valuable insight into how uncertainty and planning horizon limitations influence dispatch decisions.

However, their implementation in this thesis involved several simplifying assumptions posing limitations. Uncertainty was incorporated in the SORH model using a scenario-based approach. Forecast errors for wind and electricity prices were modeled as zero-mean, normally distributed variables with empirically derived standard deviations and correlation structures. These stochastic deviations were added to historical datasets to construct multiple plausible trajectories for prices and wind generation over a 168-hour horizon. The main limitation is that the statistical properties of the forecast errors were derived from full-year data, information that would not be available in a real-time setting. Moreover, this approach still relies on having a suitable baseline forecast method, which is not explicitly modeled in this thesis but would be essential for real-world implementation.

Lastly, both RHO and SORH models rely on perfect foresight for the first 24 hours, which is a reasonable assumption for day-ahead electricity prices but less accurate for wind generation. Conversely, they assume absolutely no forecast information beyond the weekly planning window, whereas in practice, operators can often rely on medium-term trends or probabilistic forecasts that extend beyond a few days.

8.2. Profitability implications of market access and battery storage

This section addresses the second main research question:

Main research question 2: *How does adding a battery and participation in the day-ahead market affect the hydrogen production plant's long-term profitability?*

The analysis and discussion of this question are structured around two core themes that emerged from the modeling results. First, the role of the day-ahead electricity market as an additional power source is examined, focusing on the trade-offs between wind PPA supply and opportunistic grid sourcing. This is discussed in detail in Section 8.2.1 and reflects how market participation introduces flexibility, cost-saving potential but enhanced exposure to price volatility. Second, the specific contribution of the battery to system profitability is evaluated. While the battery enhances operational flexibility and supports more strategic electricity use, its financial viability depends on the balance between cost savings and capital investment. This trade-off is explored further in Section 8.2.2.

8.2.1. Trade-offs between wind PPA and grid electricity sourcing

This section concludes how additional grid electricity sourcing, alongside a fixed wind PPA, influences the profitability and operational flexibility of green hydrogen production systems. This analysis also considered the implications of hydrogen pricing mechanisms that differentiate between renewable and grid electricity inputs.

The findings highlight that long-term renewable contracts, such as fixed wind PPAs, provide price stability and significantly reduce exposure to market risk. Configurations without a PPA showed drastically reduced profitability in volatile years like 2022, leading to substantial financial losses. The analysis also revealed, however, that an oversized PPA can also undermine overall system performance. When the contracted wind capacity exceeds the combined capacity of the electrolyser and storage, surplus electricity must be curtailed or sold into the market, sometimes at highly negative prices and lead to significant financial losses. This underscores the importance of carefully balancing PPA capacity with the system's ability to absorb, store or profitably dispatch renewable electricity.

Importantly, wind PPAs also ensure that a large share of hydrogen is produced using renewable electricity. This is a critical factor in determining whether the hydrogen qualifies as "green" under certification schemes for clean energy. In our analysis, however, grid electricity sourcing introduced greater flexibility and allowed the system to reduce power costs by capitalising on low or even negative market prices. The results showed that profitability improved as the share of grid-sourced electricity increased and the reliance on the wind PPA decreased. However, these configurations resulted in a higher share of hydrogen produced from

non-renewable sources, reducing its market value. Consequently, despite the lower production costs, overall profits would reduce when green hydrogen premiums could not be applied. This underlines the role of PPAs in both stabilising costs and securing access to higher-value revenue streams through certified green hydrogen.

Dual pricing scheme for green and grey hydrogen

To address the limitation of assuming that hydrogen produced from grid electricity could always be sold at a green premium, a dual pricing scheme was introduced. This approach differentiated hydrogen sales prices based on the origin of the electricity used in production. As a result, the model could weigh the economic benefit of sourcing low-cost grid power against the potential drawback of having to sell the hydrogen at a lower, grey price. As expected, the share of grid electricity used for hydrogen production decreased under this more realistic pricing regime, but it remained substantial, demonstrating that grid sourcing can still be economically attractive during low-price periods, even if the resulting hydrogen cannot always qualify for green certification. The detailed results supporting this conclusion are presented in Section 7.3.

Nonetheless, this approach still involves simplifications. The model did not incorporate detailed green certification criteria such as greenhouse gas intensity thresholds or hourly matching requirements between renewable generation and hydrogen production. In addition, all hydrogen produced using battery-stored electricity was assumed to qualify for the green premium. This introduces a potential modeling loophole, as the source of the battery's charge was not explicitly tracked. Although the majority of battery charging occurred during periods of surplus wind or very low market prices suggesting a likely renewable origin, this was not enforced through formal certification logic. As a result, the battery became increasingly valuable as a buffer that enabled the system to maintain green hydrogen status while taking advantage of favourable market dynamics.

The limitations of LCOH as standalone metric for hydrogen business case

Furthermore, our findings reveal that relying on LCOH as a standalone metric to assess system performance can be misleading. Although LCOH is widely used in hydrogen project evaluation, it does not fully capture the profitability of a project. In particular, our differentiated hydrogen pricing analysis showed that configurations with lower LCOH (especially those heavily reliant on low-cost market electricity) did not necessarily yield higher profits. This was because a substantial portion of the hydrogen produced did not qualify as renewable and therefore had to be sold at a lower "grey" price. The actual margin between LCOH and the hydrogen sales price proved to be the more relevant determinant of profitability. As such, future optimisation efforts should focus not only on reducing production costs, but also on maximising the spread between cost and attainable market price, which in turn depends on the hydrogen's carbon intensity.

8.2.2. Profitability of the battery

Battery storage improves operational flexibility by enabling strategic shifting of electricity consumption and helping to smooth the electrolyser's operation. This reduces overall power costs and allows the system to respond more effectively to market fluctuations. However, under baseline assumptions, the financial benefits of this added flexibility are insufficient to outweigh the battery's investment and operational costs. Compared to the no-battery case, the inclusion of a battery led to an increase in LCOH of approximately €0.35 per kilogram and a decrease in the IRR by about 2.5 percentage points. These results indicate that battery deployment, while beneficial for dispatch optimisation, does not consistently improve the long-term profitability of the hydrogen system. A sensitivity analysis on battery capital expenditure confirmed that costs would need to fall by more than 50 percent before battery integration becomes economically neutral compared to the no-battery case.

Testing the model on different historical years revealed that battery deployment only improved profitability under exceptional market conditions. In particular, during 2022, when electricity prices and volatility surged due to the energy crisis, the battery-enhanced configuration outperformed the no-battery case. In this year, increased electricity trading led to higher market revenues and reduced overall power costs, lowering the LCOH by €0.40 per kilogram. While this highlights the battery's potential in volatile markets, such conditions are not typical. Without more frequent periods of high price volatility or substantial reductions in battery capital costs, the financial benefits observed in 2022 are unlikely to hold in other years. Under current baseline assumptions, the battery remains too expensive to be profitable in this model set up.

Nevertheless, battery systems may offer strategic value that extends beyond this economic assessment. Even if they do not reliably improve long-term returns under present cost structures and market conditions,

batteries contribute to system resilience and operational flexibility. For example, they can help hedge against forecast errors and real-time price spikes, allowing for short-term adjustments that are particularly valuable in increasingly volatile electricity markets. Batteries also provide a buffer during grid congestion events, enabling better utilisation of surplus wind power under a PPA. Moreover, energy storage enhances autonomy and security, which may be essential for off-grid applications or regions with limited grid infrastructure.

Looking ahead, the value of batteries is likely to increase as power systems evolve. Higher renewable penetration and more dynamic electricity pricing will increase the need for flexible, fast-responding assets. In parallel, evolving regulatory frameworks may open access to additional revenue streams, including ancillary services such as frequency regulation, capacity reserves, or demand response. Moreover, batteries could participate in intraday trading to exploit short-term price volatility and provide balancing services closer to real time.

All these additional capabilities and potential revenue streams were not considered in the current model. These developments suggest that while battery investments may not yet be financially justified under baseline assumptions, their future role in system flexibility and market integration should be explored in further research.

9

Recommendations for Future Research

The following section highlights key areas where further research would be valuable and presents recommendations to support future work.

Enhanced modeling of uncertainty

The accuracy of optimisation outcomes is fundamentally tied to the quality of the input data. In this context, improving the realism of uncertainty modeling for electricity prices and wind generation could significantly enhance the reliability of model outcomes. This thesis modeled uncertainty by adding normally distributed forecast errors to historical values, using a fixed standard deviation estimated from the full year of data. This approach can be improved by refining the distributional assumptions. For instance, the forecast errors could follow a different distribution than the normal distribution, particularly if empirical data suggests skewness or fat tails. Additionally, rather than applying a constant mean and variance across the year, the error distributions could be made time dependent to reflect observed seasonality and time-varying volatility. As shown in Section 4.6, both wind and electricity price volatility exhibit clear seasonal patterns. Another enhancement would be to incorporate the increasing uncertainty associated with longer forecast horizons, by widening the error bands over time and narrowing them for short-term forecasts.

Ideally, however, future work should move beyond modeling uncertainty purely as deviations from point forecasts. A more advanced approach would use forecasting techniques that generate full scenario paths based on the underlying drivers of variability. By explicitly modeling factors such as weather conditions, time of day and seasonal effects, it is possible to create more structurally realistic and temporally coherent scenarios. These scenarios could then be combined with more advanced methods for modeling the correlation between wind generation and electricity prices. One promising technique is copula theory, which allows for flexible modeling of nonlinear dependencies and extreme joint events. This provides a richer and more accurate representation than standard linear correlation method applied in this thesis.

Finally, the scope of uncertainty could be expanded to include additional external factors such as fluctuations in hydrogen demand, variability in hydrogen sales prices or evolving policy and regulatory frameworks. Including these dimensions would further enhance the realism and practical applicability of the optimisation framework in real-world decision-making contexts.

Incorporating risk-aware optimisation through CVaR

Another direction for future research is to incorporate risk-aware decision-making into the dispatch optimisation model. A commonly used approach is to integrate Conditional Value-at-Risk (CVaR) into the objective function to manage downside risk arising from electricity price volatility or uncertainty in hydrogen revenues. CVaR is a risk metric that captures the expected loss in the worst-performing fraction of outcomes beyond a specified confidence level, such as the worst five percent of cases. It provides a more informative measure of tail risk than Value-at-Risk, as it reflects the average severity of the worst outcomes rather than just a threshold. CVaR can be incorporated either as part of a weighted multi-objective formulation that balances expected profit and downside risk, or as a constraint that guarantees a minimum acceptable performance under adverse conditions [43, 91]. Including such risk measures would enable the model to better account

for uncertainty, support more robust operational strategies and align with the risk preferences of investors or operators facing volatile market conditions.

Advanced sizing optimisation

This study generally assumed fixed capacity sizes for the system components, based on prior assumptions and validated through sensitivity analyses. While some exploratory optimisation was conducted for battery sizing, solving for optimal capacities was not the main focus of the thesis and was constrained by the computational burden of large-scale MILP formulations that include sizing variables as decision variables alongside dispatch optimisation. However, in integrated energy systems such as the one considered here, the relative sizing of components plays a crucial role in overall system performance. Each component, including the wind turbine, electrolyser, battery and grid connection, requires substantial capital investment. Jointly optimising these sizes could significantly reduce costs and improve system profitability. Further research into methods that enable combined optimisation of system sizing and operational strategy is therefore recommended.

Allowing all capacities to be free variables in an exact MILP framework, while simultaneously optimising dispatch decisions over a full year, is likely to be computationally infeasible. Similarly, exhaustive grid search across multiple sizing dimensions quickly becomes prohibitively expensive. A more promising direction is to apply advanced techniques such as Bayesian Optimisation (BO), a sequential and data-efficient method designed to identify global optima for computationally expensive black-box functions [52]. BO constructs a probabilistic surrogate model, often based on Gaussian Processes, of the objective function and uses an acquisition function to determine the most promising design points to evaluate next. This approach is well-suited to problems where each evaluation, such as a year-long dispatch optimisation, is time-consuming. Previous studies, including Lisicki et al. [52] and Li [49], have demonstrated the effectiveness of BO in energy system sizing. Applying BO to hybrid hydrogen systems could lead to more cost-effective system designs while requiring far fewer model evaluations than traditional grid search methods.

Hydrogen storage inclusion and sales contract structures

This thesis assumed a simplified hydrogen sales setup in which all produced hydrogen could be sold immediately, at a fixed price and flexible delivery profile, eliminating the urgency for hydrogen storage. However, this assumption does not reflect the likely structure of real-world offtake agreements. Commercial hydrogen contracts are likely to involve some kind of fixed-load delivery obligations. Offtakers are unlikely to accept highly variable delivery profiles and may instead demand a continuous and reliable flow of hydrogen. Meeting such obligations would make hydrogen storage not just beneficial but likely essential for producers.

Future research should incorporate hydrogen storage into the optimisation model to more accurately reflect the operational requirements and economic trade-offs of green hydrogen systems. Unlike batteries, hydrogen storage can offer longer-duration storage, making it better suited to capture energy mismatches that span multiple hours or even days. Including hydrogen storage would also enable a meaningful comparison between the roles of battery and hydrogen storage in managing fluctuations in renewable generation, electricity prices and delivery obligations. Optimising the size and operation of hydrogen storage alongside the current system setup is therefore a logical and valuable next step.

Moreover, baseload hydrogen delivery may command a price premium, due to its higher reliability and contractual value. This opens an important research opportunity: extending the current model to estimate the minimum premium price required to justify investment in hydrogen storage and flexible operation. By incorporating hydrogen storage and alternative offtake structures, future work could assess the commercial value of delivery flexibility, support optimal contract design, and contribute to more realistic hydrogen project development strategies.

RFNBO compliance and temporal matching

As regulatory frameworks evolve, compliance with certification schemes such as the Renewable Fuels of Non-Biological Origin (RFNBO) is expected to become a critical factor in determining the market value of green hydrogen. In particular, adherence to RFNBO rules will influence both the eligibility of hydrogen for green certification and the price premium it can command in offtake agreements. One of the most complex aspects of RFNBO compliance is the requirement for strict temporal matching between renewable electricity generation and hydrogen production, often at an hourly level.

The current model does not account for these regulatory constraints, focusing instead on economic and

technical optimisation. However, once grid connection and electricity market participation are introduced, ensuring RFNBO compliance becomes significantly more challenging. Producers must be able to demonstrate that electricity used for electrolysis is not only renewable, but also time-synchronised with hydrogen production. Moreover, RFNBO guidelines consider the greenhouse gas (GHG) intensity of grid electricity. If the average GHG intensity of the grid falls below a certain threshold at a given time, hydrogen produced using that electricity may still qualify as green. This adds a second layer of complexity to certification tracking that must also be addressed.

Therefore, future research should extend the current modeling framework to track both the temporal matching of renewable electricity and the carbon intensity of grid-sourced electricity. This would allow for evaluation of trade-offs between operational optimisation and regulatory compliance, and support the development of dispatch strategies that maximise profitability while preserving eligibility for green certification.

Short-term operational modeling and participation in additional markets

This thesis focused on long-term profitability using a dispatch model with hourly resolution over a one-year horizon. While suitable for strategic planning and system design, this approach does not capture short-term operational dynamics that are increasingly relevant in real-world electricity markets. To assess the system's performance under real-time conditions, future research should develop short-term operational models with higher temporal resolution, such as 5- or 15-minute intervals. These models would be better equipped to represent intra-day variability in electricity prices, wind generation and battery behaviour, as well as to handle forecasting errors and short-term balancing needs.

In parallel, the current modeling framework only includes participation in the day-ahead electricity market. However, energy systems may also participate in intra-day markets, reserve markets or ancillary service markets such as frequency regulation or spinning reserve. These markets offer additional revenue streams, particularly for fast-responding assets like batteries. Capturing these opportunities would require extending the existing model to reflect the specific operational and timing requirements of these markets, which in turn necessitates short-term optimisation frameworks. Future research should therefore explore the development of high-resolution dispatch models in combination with expanded market participation, as both are essential to more accurately assess the operational flexibility and economic potential of hybrid energy systems in increasingly dynamic electricity markets.

Appendices

A

Appendix

A.1. Correlation Structure Scenario Generation Method

To evaluate and support the realism of the correlation structure embedded in the scenario generator, we analysed key statistical properties from historical data. These include autocorrelation within the wind and price time series, cross-correlation between the two series, and the smoothed cross-correlation matrix structure used in scenario generation. This appendix provides empirical validation of the temporal and joint dynamics that underpin the generation of realistic scenarios for the stochastic optimisation model.

Autocorrelation Coefficients

The autoregressive behaviour of both wind load factors and day-ahead electricity prices is captured using an AR(1) structure. Table A.1 reports the estimated lag-1 autocorrelation coefficients from 2023 hourly data:

	Wind (ϕ_α)	Price (ϕ_λ)
Estimated AR(1) Coefficient	0.9896	0.9245

Table A.1: Estimated lag-1 autocorrelation from historical 2023 data

These values confirm the strong persistence in both time series, justifying the use of autoregressive sampling methods during scenario construction.

Cross-Correlation Between Series

To capture the interaction between wind availability and electricity prices, we compute the empirical Pearson correlation between the two series. Table A.2 shows the result:

Correlation Type	Value
Wind–Price Correlation (ρ)	−0.369

Table A.2: Empirical Pearson correlation between wind and electricity prices

This negative correlation reflects the merit-order effect, where high wind availability tends to lower electricity prices due to the near-zero marginal cost of wind generation. Capturing this negative relationship is crucial for scenario realism, especially when evaluating strategies that rely on market arbitrage.

Smoothed Cross-Correlation Matrix

To model temporal spillover effects in the wind–price relationship, we construct a smoothed cross-correlation matrix R_{wp} using exponentially decaying weights across time lags. Figure A.1 visualises the structure of R_{wp} for two decay factors, $\delta = 0.6$ and $\delta = 0.9$.

This smoothing strategy reflects the realistic assumption that wind availability at a given hour can affect electricity prices in nearby hours (e.g. $t \pm 1, t \pm 2, \dots$), due to market and ramping dynamics. A decay factor

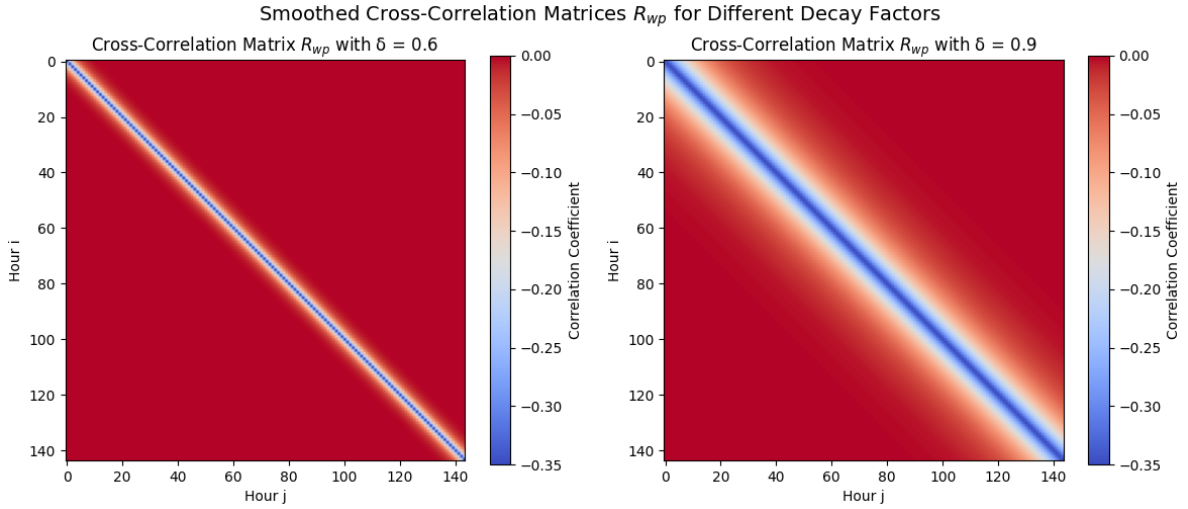


Figure A.1: Heat maps for the smoothed cross-correlation matrices R_{wp} for different exponential decay factors $\delta \in \{0.6, 0.9\}$. A higher δ preserves stronger temporal spillover in the wind–price correlation across hours, while a lower δ concentrates correlation near the diagonal. Both matrices reflect negative contemporaneous correlation, and their smoothed structure ensures that R_{wp} is positive semi-definite, allowing valid Cholesky decomposition for scenario generation.

δ close to 1 results in more persistent correlation across time, while lower values limit interaction to the immediate hour. Based on numerical experiments and stability considerations, we select $\delta \in (0.8, 0.9)$, which balances physical realism with the requirement for a positive semi-definite (PSD) correlation matrix.

As a final step, we ensure numerical stability by applying regularisation, symmetrising the matrix and adding a small diagonal perturbation:

$$R \leftarrow \frac{1}{2}(R + R^\top) + \varepsilon I \quad \text{with} \quad \varepsilon \approx 10^{-6}.$$

This guarantees that Cholesky decomposition can be applied to generate valid correlated samples.

Together, the results in this appendix validate the structure used in the scenario generation method, confirming that both auto- and cross-correlation dynamics are faithfully represented. This supports the model's ability to simulate realistic joint scenarios of wind availability and electricity prices for use in stochastic dispatch optimisation.

B

Appendix

B.1. Battery and electrolyser behaviour for the differentiated hydrogen price model

This appendix provides additional system-level insights based on an extended model run where hydrogen sales prices vary depending on the source of electricity, supporting the discussion in Section 7.3. In this model, a premium is awarded for green hydrogen produced using renewable electricity, while a lower price applies to grey hydrogen produced from grid power. This price differentiation alters dispatch decisions, especially regarding the role of the battery and the operation of the electrolyser.

Table B.1 summarises the impact on battery usage. Subtable (a) reveals a significant increase in battery activity and suggests that the battery is used strategically to buffer renewable energy and support the electrolyser under favourable pricing conditions. Subtable (b) shows the distribution of charging and discharging flows and indicates that stored energy is now used more deliberately to supply the electrolyser during periods when the value of green hydrogen is higher.

Table B.2 complements this analysis by detailing the sources of electricity used by the electrolyser and summarising its operational behaviour. The share of market electricity used drops, while battery power increases. Most power still comes directly from wind, but the system clearly prioritises low-carbon sources in response to pricing incentives. The operational profile of the electrolyser also adapts: The electrolyser now spends more time operating at partial load rather than switching between full operation and complete shut-down. This behaviour reflects a more gradual, dynamic dispatch strategy aligned with time-varying electricity and hydrogen prices.

Table B.1: Battery operation results under the differentiated hydrogen sales price model. Subtable (a) present the battery's operating hours which indicate a substantial increase in battery activity. Subtable (b) provides the breakdown of charging and discharging flows, showing how the battery is increasingly used to provide power to the electrolyser.

(a) Compared to the base case results in Table 7.5, the total number of battery operating hours increases substantially, from a combined total of 3,974 hours to 8,118 hours. This reflects a significantly more active role for the battery in the updated model. The system increasingly relies on the battery to shift wind energy and to limit hydrogen production from market electricity, which is now associated with a lower sales price. Additionally, the number of discharging hours is nearly double that of charging, suggesting a strategy of rapid charging followed by more gradual discharging, primarily to supply the electrolyser.

Metric	Unit	Charging	Discharging
Total energy throughput	GWh	65.22	58.86
Average power	MW	22.86	11.18
Operating hours	hours	2853	5265
Share at full rated power	% of oper. hours	68	0.84
Round-trip efficiency	%	90.25	-

(b) Compared to the base case in Table 7.7, the share of discharging directed to the electrolyser rises from 33.1% to 57.7%, reflecting a strategy to store electricity for hydrogen production during favourable pricing periods. Additionally, charging from the market increases, reinforcing the battery's role in the updated dispatch strategy.

Battery charge

Source	GWh	Share (%)
Wind to Battery	21.3	31.2
Market to Battery	46.8	68.8
Total	68.1	100.0

Battery discharge

Destination	GWh	Share (%)
Battery to Electrolyser	35.4	57.7
Battery to Market	26.0	42.3
Total	61.5	100.0

Table B.2: Distribution of power sources to the electrolyser and utilisation summary results for the differentiated hydrogen sales price model. Compared to the base case in Table 7.6, the share of market electricity used by the electrolyser decreases from 20.9% to 12.7%, while battery-supplied power increases from 3.4% to 8.7%. This indicates a strategic shift away from producing hydrogen using market power, which now yields a lower sales price, towards using more renewable and battery-stored energy. The average utilisation rate decreases slightly from 96% to 93%, while the number of shut-offs drops from 15 to 10. This reduction in shut-offs may be partly explained by an increase in intermediate-load operation: instead of switching between fully on, minimum load, or off, the electrolyser now ramps up and down more gradually in response to dynamic pricing, spending more time at partial utilisation levels.

Power sources share of the electrolyser

Source	GWh	Share (%)
Wind to Electrolyser	319.8	78.6
Battery to Electrolyser	35.5	8.7
Market to Electrolyser	51.5	12.7
Total	406.8	100.0

Electrolyser utilisation summary

Metric	Unit	Value
Avg utilisation rate	%	92.9
Shut-offs (count)	-	10
Off status (hours)	hours	300
Min capacity (hours)	hours	222
intermediate-load (hours)	hours	505
Max capacity (hours)	hours	7730

C

Appendix

C.1. Results for No Battery Model for 100 MW PPA and 50 MW Electrolyser Setup

This appendix presents detailed results for the base case configuration without a battery. The system includes a 100 MW Power Purchase Agreement (PPA) for offshore wind and a 50 MW electrolyser. All model outcomes are based on hourly operational optimisation for a single representative year with hydrogen sales prices set at €10/kg for green hydrogen and €7/kg for grey hydrogen. The purpose of this configuration is to establish a baseline reference case for system performance without a battery and excluding any price-driven energy shifting capabilities. The results include key financial indicators, power flows, electrolyser utilisation and cost and revenue breakdowns. All power not consumed by the electrolyser is either curtailed or sold on the day-ahead market.

Table C.2: Summary of annual power flows results for model set-up without battery.

Metric	Value
Avg wind load factor	0.510
Avg DA market price (€/MWh)	95.82
Power bought DA market (GWh)	76.71
Power sold DA market (GWh)	126.47
Wind power from PPA (GWh)	446.47
PPA overcapacity (GWh)	0.66
Battery charge (GWh)	0.00
Battery discharge (GWh)	0.00
Power to electrolyser (GWh)	396.77
Hydrogen produced (GWh)	238.02
Hydrogen produced (kt)	7.14

Table C.1: Financial KPI results for model set-up without battery.

Metric	Value
Net Present Value (NPV) (M€)	84.07
Internal Rate of Return (IRR) (%)	22.62
Levelized Cost of Hydrogen (LCOH) (€/kg)	8.036
Avg utilisation rate (%)	90.57

Table C.3: Electrolyser utilisation behaviour results for model set-up without battery.

Metric	Value
Avg utilisation rate (%)	90.6
Shut-offs (count)	8
Off status (hours)	300
Min capacity (hours)	433
Overig (hours)	589
Max capacity (hours)	7436

Table C.4: Distribution of power to electrolyser for model set-up without battery

Source	GWh	Share (%)
Wind to Electrolyser	320.0	80.7
Battery to Electrolyser	0.0	0.0
Market to Electrolyser	76.7	19.3
Total	396.7	100.0

Table C.5: Summary of costs and revenues for model set-up without battery.

Metric	Unit	Value
Hydrogen revenue	M€	67.26
Power market revenue	M€	10.13
Total revenue	M€	77.39
Market power cost	M€	7.15
Wind PPA cost	M€	42.48
Redispatch penalty	M€	0.10
Cold start-up cost	M€	0.06
Electrolyser OPEX	M€	2.19
Battery OPEX	M€	0.00
Total OPEX	M€	2.19
Electrolyser CAPEX (annual)	M€	10.28
Battery CAPEX (annual)	M€	0.00
Total CAPEX (annual)	M€	10.28
Total cost (excl. CAPEX)	M€	51.98
Annual profit (excl. CAPEX)	M€	25.41
Net profit (after all costs)	M€	15.13

Table C.6: Wind power distribution for model set-up without battery.

Destination	GWh	Share (%)
Wind to Electrolyser	320.0	71.6
Wind to Battery	0.0	0.0
wind to Market	126.6	28.3
Wind Overcapacity	0.7	0.1
Total	447.1	100.0
Wind to Electrolyser	320.0	80.7
Battery to Electrolyser	0.0	0.0
Market to Electrolyser	126.5	28.3
Total	396.7	100.0

C.2. Results for Electrolyser only Model for different Historical Day-Ahead Price Data

Price Data

This section presents annual results for two alternative configurations that rely solely on electricity sourced from the day-ahead market, without a wind PPA. These models are tested using real historical price data for six different years (2019–2024) to evaluate how price volatility affects system economics and the potential role of battery storage. The first configuration includes both an electrolyser and battery, allowing energy shifting based on market price dynamics. The second configuration includes only the electrolyser, with no flexibility beyond real-time market response. This analysis supports the discussion in Section 7.4.4, which highlights the impact of price volatility in the absence of long-term electricity contracts.

Results model: electrolyser and battery (no PPA)

In this configuration, the battery enables temporal optimisation by storing electricity during low-price hours and discharging during periods of high prices. Table C.7 reports the financial and technical performance metrics across six historical years. Despite the added operational flexibility, the battery-inclusive setup generally underperforms compared to the minimal electrolyser-only model in most years. The additional capital and operational expenditure of the battery is not offset by increased revenue from market sales under typical price conditions. Only in years of high volatility, most notably 2022, does battery operation offer high revenue potential.

Table C.7: Performance metrics for electrolyser and battery (no PPA) model

Year	NPV (M€)	IRR (%)	LCOH (€/kg)	Util. Factor (%)	Avg Power Cost (€/MWh)	Revenue Share Power Sales (%)
2019	92.77	21.7	5.33	91.3	37.87	0.5
2020	115.59	24.3	4.92	91.3	28.19	0.6
2021	10.32	11.4	6.76	73.3	66.54	7.4
2022	-200.82	–	11.79	68.5	169.32	16.2
2023	-7.33	9.0	7.14	80.9	77.02	4.2
2024	33.39	14.4	6.37	88.0	62.52	2.7

Results model: electrolyser only (no PPA, no battery)

This configuration represents the most exposed case, with electricity purchased exclusively from the day-ahead market and no ability to shift demand in time. Table C.8 summarises the annual results under this setup. Interestingly, this simpler configuration performs better than the battery-inclusive model in most years, due to lower capital costs and the limited added value of battery flexibility. However, it is also more vulnerable to price shocks. In 2022, extreme price volatility leads to catastrophic financial outcomes, with the NPV dropping to approximately –€208 million. This single year of adverse prices erases the cumulative profits from the preceding years, illustrating the fragility of the business case in the absence of price stability.

Table C.8: Performance metrics for electrolyser only (no PPA, no battery) model

Year	NPV (M€)	IRR (%)	LCOH (€/kg)	Util. Factor (%)	Avg Power Cost (€/MWh)
2019	116.97	28.8	4.90	91.3	39.47
2020	138.80	32.1	4.50	91.3	29.97
2021	26.16	14.6	6.41	73.3	66.54
2022	-208.02	0.0	11.97	68.5	173.51
2023	5.17	10.9	6.89	80.9	82.44
2024	45.35	17.7	6.15	88.0	68.04

These findings reinforce the broader conclusion that relying solely on market-based electricity procurement introduces substantial risk to the hydrogen production business case. Without long-term fixed-price contracts such as PPAs, even a technically viable system becomes financially fragile. Moreover, battery deployment, although useful in extreme conditions, is insufficient on its own to ensure profitability under current high investment costs.

Bibliography

- [1] M. A. Aftab, V. C. Pandey, S. G. Krishnan, F. Mir, G. Rolofs, E. Chukwureh, S. Ahmed, and C. Konstantinou. Demand flexibility in hydrogen production by incorporating electrical and physical parameters. *Electric Power Systems Research*, 239:111213, 2025. doi:10.1016/j.epsr.2024.111213.
- [2] E. F. Alvarez, P. Sanchez-Martin, and A. Ramos. Self-scheduling for a hydrogen-based virtual power plant in day-ahead energy and reserve electricity markets. *20th International Conference on the European Energy Market (EEM)*, pages 1–6, 06 2024. doi:10.1109/EEM60825.2024.10608848.
- [3] F. J. Baader, A. Bardow, and M. Dahmen. MILP formulation for dynamic demand response of electrolyzers. In Y. Yamashita and M. Kano, editors, *14th International Symposium on Process Systems Engineering*, volume 49 of *Computer Aided Chemical Engineering*, pages 391–396. Elsevier, 2022. doi:10.1016/B978-0-323-85159-6.50065-8.
- [4] R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [5] J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer, 1997. ISBN 0387982175.
- [6] C. Bowden. Understanding power markets: Merit order and marginal pricing, 2023. Retrieved June 2, 2025, from <https://www.squeaky.energy/blog/understanding-power-markets-merit-order-and-marginal-pricing>.
- [7] J. Burchardt, E. Hegnholt, M. Holm, F. Klose, D. Ritter, and S. Schönberger. Turning the European green hydrogen dream into reality: A call to action, October 2023. White Paper. Retrieved April 18, 2025, from <https://media-publications.bcg.com/Turning-the-European-Green-H2-Dream-into-Reality.pdf>.
- [8] A. Buttler and H. Spliethoff. Current status of water electrolysis for energy storage, grid balancing and sector coupling via power-to-gas and power-to-liquids: A review. *Renewable and Sustainable Energy Reviews*, 82:2440–2454, 2018. ISSN 1364-0321. doi:<https://doi.org/10.1016/j.rser.2017.09.003>.
- [9] C. Campbell-Stanway, V. Becerra, and S. Prabhu. Techno-economic analysis with electrolyser degradation modelling in green hydrogen production scenarios. *International Journal of Hydrogen Energy*, 106: 80–95, 2025. doi:10.1016/j.ijhydene.2025.01.359.
- [10] G. Ceusters, R. C. Rodríguez, A. B. García, R. Franke, G. Deconinck, L. Helsen, A. Nowé, M. Messagie, and L. R. Camargo. Model-predictive control and reinforcement learning in multi-energy system case studies. *Applied Energy*, 303:117634, 2021. doi:10.1016/j.apenergy.2021.117634.
- [11] M. R. Chakraborty, S. Dawn, P. K. Saha, J. B. Basu, and T. S. Ustun. A comparative review on energy storage systems and their application in deregulated systems. *Batteries*, 8(9), 2022. doi:10.3390/batteries8090124.
- [12] J. Collet, O. Féron, and P. Tankov. Optimal management of a wind power plant with storage capacity, 2017. Retrieved January 20, 2025, from <https://hal.science/hal-01627593v1>.
- [13] Hydrogen Council. Hydrogen insights 2023, 2023. Retrieved May 12, 2025, from <https://hydrogencouncil.com>.
- [14] A. Čučuk. Westwood: Over a fifth of all European hydrogen projects stalled or cancelled, December 2024. Retrieved June 10, 2025, from <https://www.offshore-energy.biz/westwood-over-a-fifth-of-all-european-hydrogen-projects-stalled-or-cancelled>.
- [15] G. B. Dantzig. Linear programming under uncertainty. *Management Science*, 1(3–4):197–206, 1955. doi:10.1287/mnsc.1.3-4.197.

- [16] C. A. del Pozo, S. Cloete, and Á. J. Álvaro. Techno-economic assessment of long-term methanol production from natural gas and renewables. *Energy Conversion and Management*, 266, 2022. doi:10.1016/j.enconman.2022.115785.
- [17] B. den Ouden, S. Hers, J. Schellekens, S. Blom, P. Verstraten, A. van der Veen, T. Verboon, J. Wolda, and W. Coenraads. Blueprinting the hydrogen market – hydrogen spot market simulation (H2SMS) technical report, 2024. Retrieved May 18, 2025, from <https://www.hyxchange.nl/document-library/>.
- [18] H. Ding, Z. Hu, and Y. Song. Rolling optimization of wind farm and energy storage system in electricity markets. *IEEE Transactions on Power Systems*, 30(5):2676–2684, 2015. doi:10.1109/TPWRS.2014.2364272.
- [19] L. Eblé and M. Weeda. Evaluation of the levelised cost of hydrogen based on proposed electrolyser projects in the netherlands. renewable hydrogen cost element evaluation tool (RHyCEET). Technical Report TNO 2024 R10766, TNO, May 2024. Commissioned by the Dutch Ministry of Economic Affairs and Climate Policy. Retrieved July 2, 2024, from <https://resolver.tno.nl/uuid:e5e1ab2e-ff69-48fb-8564-75f56282378c>.
- [20] L. F. J. Eblé. Temporal-correlation requirements for industrial onsite electrolysis: On the impact of temporal-correlation requirements and downstream industrial flexibility on the optimal design and costs for onsite electrolytic hydrogen production. Master's thesis, Delft University of Technology, September 2023. Retrieved June 18, 2025, from <http://repository.tudelft.nl/>.
- [21] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. CRC Press, New York, NY, 1994. ISBN 978-0412042317.
- [22] ENTSO-E. Entso-e transparency platform: Day-ahead market prices, 2024. Retrieved April 23, 2025, from <https://transparency.entsoe.eu>.
- [23] EPEX SPOT. Basics of the power market, 2022. Retrieved June 1st, 2025, from <https://www.epexspot.com/en/basicspowermarket>.
- [24] EPEX SPOT. Trading products, 2023. Retrieved June 1st, 2025, from <https://www.epexspot.com/en/tradingproduct>.
- [25] European Commission. A hydrogen strategy for a climate-neutral europe, July 2020. Final. Retrieved May 12, 2025, from <https://ec.europa.eu/commission/presscorner/api/files/attachment/865942>.
- [26] M. Feili and M. Aameli. The P2P energy management scheme for integrated energy microgrid considering P2G and electricity network fee. *International Journal of Industrial Electronics Control and Optimization*, 8(1):1–23, 2025. ISSN 2645-3517. doi:10.22111/ieco.2024.49044.1583.
- [27] H.-C. Gao, J.-H. Choi, S.-Y. Yun, H.-J. Lee, and S.-J. Ahn. Optimal scheduling and real-time control schemes of battery energy storage system for microgrids considering contract demand and forecast uncertainty. *Energies*, 11:1371, 2018. doi:10.3390/en11061371.
- [28] A. Ghasemi, M. Banejad, and M. Rahimiyan. Integrated energy scheduling under uncertainty in a micro energy grid. *IET Generation, Transmission & Distribution*, 12(12):2887–2896, 2018. doi:10.1049/iet-gtd.2017.1631.
- [29] Y. Ghiassi-Farrokhfal, W. Ketter, and J. Collins. Making green power purchase agreements more predictable and reliable for companies. *Decision Support Systems*, 144:113514, 2021. doi:10.1016/j.dss.2021.113514.
- [30] P. Glasserman. *Monte Carlo Methods in Financial Engineering*, volume 53 of *Applications of Mathematics*. Springer, 2003.
- [31] C. Graham and D. Talay. *Stochastic Simulation and Monte Carlo Methods: Mathematical Foundations of Stochastic Simulation*, volume 68 of *Stochastic Modelling and Applied Probability*. Springer, Berlin, 2013. ISBN 9783642393631. doi:10.1007/978-3-642-39363-1. Available online via TU Delft Library. Retrieved June 18, 2025, from <https://public.ebookcentral.proquest.com/choice/publicfullrecord.aspx?p=6312527>.

- [32] S. A. Grigoriev, V. N. Fateev, D. G. Bessarabov, and P. Millet. Current status, research trends, and challenges in water electrolysis science and technology. *International Journal of Hydrogen Energy*, 45(49): 26036–26058, 2020. ISSN 0360-3199. doi:<https://doi.org/10.1016/j.ijhydene.2020.03.109>.
- [33] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2024. Retrieved January 18, 2025, from <https://www.gurobi.com>.
- [34] W. K. K. Haneveld, M. H. van der Vlerk, and W. Romeijnders. *Stochastic Programming: Modeling Decision Problems*. Graduate Texts in Operations Research. Springer, Enschede, The Netherlands, 2024.
- [35] B. M. Hodge, D. Lew, M. Milligan, H. Holttinen, S. Sillanpaa, E. Gomez-Lazaro, R. Scharff, L. Soder, X. G. Larsen, G. Giebel, et al. Wind power forecasting error distributions: An international comparison; preprint. Technical report, Golden, CO, United States, September 2012. Retrieved January 20, 2025, from <https://www.osti.gov/biblio/1051129>.
- [36] S. Hollmén, F. Levihn, and G. Martinsson. When markets don't deliver: Bilateral hedging by means of pfas in managing intertemporal price risks in power generation investments. In *2022 18th International Conference on the European Energy Market (EEM)*, pages 1–6, 2022. doi:10.1109/EEM54602.2022.9921036.
- [37] R. A. Howard. *Dynamic Programming and Markov Processes*. MIT Press, 1960.
- [38] HyNetwork. Waterstofnetwerk Nederland, 2024. URL <https://www.hynetwork.nl/over-hynetwork/waterstofnetwerk-nederland>. Retrieved June 2, 2025, from <https://www.hynetwork.nl/over-hynetwork/waterstofnetwerk-nederland>.
- [39] IEA. Global hydrogen review 2023. 2023. Retrieved April 10, 2025, from <https://www.iea.org/reports/global-hydrogen-review-2023>.
- [40] International Renewable Energy Agency (IRENA). Geopolitics of the energy transformation: The hydrogen factor, 2022. International Renewable Energy Agency, Abu Dhabi. Retrieved June 18, 2025, from <https://www.irena.org/publications/2022/Jan/Geopolitics-of-the-Energy-Transformation-Hydrogen>.
- [41] B. J. Jesse, G. J. Kramer, V. Koning, S. Vögele, and W. Kuckshinrichs. Stakeholder perspectives on the scale-up of green hydrogen and electrolyzers. *Energy Reports*, 11:208–217, 2024. doi:10.1016/j.egy.2023.11.046.
- [42] N. Johnson, M. Liebreich, D. M. Kammen, P. Ekins, R. McKenna, and I. Staffell. Realistic roles for hydrogen in the future energy transition. *Nature Reviews Clean Technology*, 1:1–12, 2025. doi:10.1038/s44359-025-00050-4.
- [43] A. R. Jordehi, S. A. Mansouri, M. Tostado-Véliz, M. Carrión, M. J. Hossain, and F. Jurado. A risk-averse two-stage stochastic model for optimal participation of hydrogen fuel stations in electricity markets. *International Journal of Hydrogen Energy*, 49, 2024. doi:10.1016/j.ijhydene.2023.07.197.
- [44] M. H. A. Khan, R. Daiyan, et al. Designing optimal integrated electricity supply configurations for renewable hydrogen generation in Australia. *iScience*, 24:102539, 2023. doi:10.1016/j.isci.2021.102539.
- [45] D. Kurowicka and R. M. Cooke. *Uncertainty Analysis with High Dimensional Dependence Modelling*. Wiley, Chichester, UK, 2006. ISBN 978-0470017394.
- [46] KYOS. What is spot power trading?, 2023. Retrieved June 1st, 2025, from <https://www.kyos.com/faq/what-is-spot-power-trading/>.
- [47] Lazard. Levelized cost of energy+ 2024 (version 17.0), 2024. Retrieved April 23, 2025, from <https://www.lazard.com/research-insights/levelized-cost-of-energyplus-lcoeplus/>.
- [48] P. Lettenmeier. Efficiency - electrolysis. Technical report, Siemens Energy, 2021. Retrieved June 2, 2025, from <https://www.siemens-energy.com/>.

- [49] B. Li and J. Li. Sizing large numbers of grid-connected microgrids based on bayesian optimization. In *2023 8th Asia Conference on Power and Electrical Engineering (ACPEE)*, pages 1954–1958, 2023. doi:[10.1109/ACPEE56931.2023.10135641](https://doi.org/10.1109/ACPEE56931.2023.10135641).
- [50] Y. Li and Y. Liu. Optimizing flexible one-to-two matching in ride-hailing systems with boundedly rational users. *Transportation Research Part E: Logistics and Transportation Review*, 150:102329, 2021. doi:[10.1016/j.tre.2021.102329](https://doi.org/10.1016/j.tre.2021.102329).
- [51] A. Liponi, A. Baccioli, L. Ferrari, and U. Desideri. Techno-economic analysis of hydrogen production from PV plants. *E3S Web of Conferences*, 334:01001, 2022. doi:[10.1051/e3sconf/202233401001](https://doi.org/10.1051/e3sconf/202233401001).
- [52] M. Lisicki, W. Lubitz, and G. W. Taylor. Optimal design and operation of archimedes screw turbines using bayesian optimization. *Applied Energy*, 183:1404–1417, 2016. doi:[10.1016/j.apenergy.2016.09.084](https://doi.org/10.1016/j.apenergy.2016.09.084).
- [53] J. Liu, X. Chen, H. Yang, and Y. Li. Energy storage and management system design optimization for a photovoltaic integrated low-energy building. *Energy*, 190:116424, 2020. doi:[10.1016/j.energy.2019.116424](https://doi.org/10.1016/j.energy.2019.116424).
- [54] A. Mansour-Saatloo, M. Agabalaye-Rahvar, M. A. Mirzaei, B. Mohammadi-Ivatloo, M. Abapour, and K. Zare. Robust scheduling of hydrogen based smart micro energy hub with integrated demand response. *Journal of Cleaner Production*, 267:122041, 2020. doi:[10.1016/j.jclepro.2020.122041](https://doi.org/10.1016/j.jclepro.2020.122041).
- [55] R. T. Marler and J. S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26(6):369–395, 2004. doi:[10.1007/s00158-003-0368-6](https://doi.org/10.1007/s00158-003-0368-6).
- [56] K. Mayer and S. Trück. Electricity markets around the world. *Journal of Commodity Markets*, 9:77–100, 2018. ISSN 2405-8513. doi:[10.1016/j.jcomm.2018.02.001](https://doi.org/10.1016/j.jcomm.2018.02.001).
- [57] C. Mittler, M. Bucksteeg, and P. Staudt. Review of renewable power purchasing agreement types and classification through morphological analysis, 2023. Available at SSRN. Retrieved October 22, 2023, from <https://ssrn.com/abstract=4626108>.
- [58] I. Moradpoor, S. Syri, and A. Santasalo-Aarnio. Green hydrogen production for oil refining – Finnish case. *Renewable and Sustainable Energy Reviews*, 175:113159, 2023. doi:[10.1016/j.rser.2023.113159](https://doi.org/10.1016/j.rser.2023.113159).
- [59] C. Moran, P. Deane, S. Yousefian, and R. F. D. Monaghan. The hydrogen storage challenge: Does storage method and size affect the cost and operational flexibility of hydrogen supply chains? *International Journal of Hydrogen Energy*, 52:1090–1100, 2024. ISSN 0360-3199. doi:[10.1016/j.ijhydene.2023.06.269](https://doi.org/10.1016/j.ijhydene.2023.06.269).
- [60] M. Mulder and B. Willems. The dutch retail electricity market. *Energy Policy*, 127:228–239, 2019. ISSN 0301-4215. doi:doi.org/10.1016/j.enpol.2018.12.010.
- [61] Nationaal Waterstofprogramma (NWP). Hydrogen roadmap for the Netherlands, 2022. A publication commissioned by the Ministry of Economic Affairs and Climate Policy. Retrieved June 18, 2025, from <https://nationaalwaterstofprogramma.nl/documenten/handlerdownloadfiles.ashx?idnv=2339011>.
- [62] J. Nease and T. A. Adams. Application of rolling horizon optimization to an integrated solid-oxide fuel cell and compressed air energy storage plant for zero-emissions peaking power under uncertainty. *Computers & Chemical Engineering*, 68:203–219, 2014. doi:[10.1016/j.compchemeng.2014.06.001](https://doi.org/10.1016/j.compchemeng.2014.06.001).
- [63] H. Niaz and J. J. Liu. A mixed integer dynamic optimization approach for a hybrid-stand alone solar and wind powered alkaline water electrolyser for renewable hydrogen. In M. Türkay and R. Gani, editors, *31st European Symposium on Computer Aided Process Engineering*, volume 50 of *Computer Aided Chemical Engineering*, pages 1285–1291. Elsevier, 2021. doi:[10.1016/B978-0-323-88506-5.50198-4](https://doi.org/10.1016/B978-0-323-88506-5.50198-4).
- [64] N. Niedrig, J. Giehl, P. Jahnke, and J. Müller-Kirchenbauer. Market design options for a hydrogen market. Technical Report 04-2024, Copenhagen Business School, Denmark, 2024.
- [65] E. Ntmeou, F. Ioannidis, K. Kosmidou, and K. Andriosopoulos. Navigating electricity market design of greece: Challenges and reform initiatives. *Energies*, 18(10), 2025. doi:[10.3390/en18102575](https://doi.org/10.3390/en18102575).

- [66] E. Oh and H. Wang. Reinforcement-learning-based energy storage system operation strategies to manage wind power forecast uncertainty. *IEEE Access*, 8:20965–20976, 2020. doi:10.1109/ACCESS.2020.2968841.
- [67] S. Oh, J. Kong, Y. Yang, J. Jung, and C.-H. Lee. A multi-use framework of energy storage systems using reinforcement learning for both price-based and incentive-based demand response programs. *International Journal of Electrical Power & Energy Systems*, 144:108519, 2023. doi:10.1016/j.ijepes.2022.108519.
- [68] R. Ortiz-Cebolla, F. Dolci, and E. Weidner. Assessment of hydrogen delivery options, 2021. Joint Research Centre report. Retrieved April 10, 2025, from https://joint-research-centre.ec.europa.eu/document/download/5cdbd6f7-7ab4-447b-be0a-dde0a25198ab_en.
- [69] G. Papaefthymiou and D. Kurowicka. Using copulas for modeling stochastic dependence in power system uncertainty analysis. *IEEE Transactions on Power Systems*, 24(1):40–49, 2009. doi:10.1109/TPWRS.2008.2006984.
- [70] S. Pfenninger and I. Staffell. Long-term patterns of european PV output using 30 years of validated hourly reanalysis and satellite data. *Energy*, 114:1251–1265, 2016. doi:10.1016/j.energy.2016.08.060.
- [71] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, 2007.
- [72] W. B. Powell. A unified framework for stochastic optimization. *European Journal of Operational Research*, 275(3):795–821, 2019. doi:10.1016/j.ejor.2018.07.014.
- [73] Power2X. Power2x company website, 2024. Retrieved July 2, 2024, from <https://www.power2x.com/>.
- [74] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2005.
- [75] M. Rezaei, A. Akimov, and E. M. Gray. Economics of renewable hydrogen production using wind and solar energy: A case study for queensland, australia. *Journal of Cleaner Production*, 435:140476, 2024. doi:10.1016/j.jclepro.2023.140476.
- [76] M. M. Rienecker, M. J. Suarez, R. Gelaro, R. Todling, et al. MERRA: NASA’s modern-era retrospective analysis for research and applications. *Journal of Climate*, 24(14):3624–3648, 2011. doi:10.1175/JCLI-D-11-00015.1.
- [77] W. Römisch. Scenario reduction techniques in stochastic programming. In H. Niederreiter and D. Talay, editors, *Stochastic Algorithms: Foundations and Applications*, volume 2740 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2003. doi:10.1007/978-3-540-45199-5_1.
- [78] S. Schrapp, M. Caspersen, M. Holm, N. Ge, S. Jovanovic, L. Schoeller, and M. Egenhöfer. De-risking low-carbon hydrogen: A guide for machinery makers to navigate the global dynamics of hydrogen, October 2024. White Paper. Retrieved May 22, 2025, from <https://media-publications.bcg.com/BCG-De-Risking-Low-Carbon-Hydrogen.pdf>.
- [79] A. H. Schrottenboer, A. A. T. Veenstra, M. A. J. uit het Broek, and E. Ursavas. A green hydrogen energy system: Optimal control strategies for integrated hydrogen storage and power generation with wind energy. *Renewable and Sustainable Energy Reviews*, 168:112744, 2022. doi:10.1016/j.rser.2022.112744.
- [80] P. Seljom, L. Kvalbein, L. Hellemo, M. Kaut, and M. Muñoz Ortiz. Stochastic modelling of variable renewables in long-term energy models: Dataset, scenario generation & quality of results. *Energy*, 236:121415, 2021. doi:10.1016/j.energy.2021.121415.
- [81] F. Sensfuß, M. Ragwitz, and M. Genoese. The merit-order effect: A detailed analysis of the price effect of renewable electricity generation on spot market prices in germany. Working Paper S 7/2007, Fraunhofer Institute for Systems and Innovation Research (ISI), Karlsruhe, Germany, 2007.
- [82] Shell Nederland B.V. Shell start bouw van Europa’s grootste groene waterstoffabriek in Rotterdam, July 2022. Retrieved April 30, 2025, from <https://www.shell.nl/over-ons/nieuws/nieuwsberichten-2022/holland-hydrogen-1.html>.

- [83] H. Shuai, J. Fang, X. Ai, J. Wen, and H. He. Optimal real-time operation strategy for microgrid: An adp-based stochastic nonlinear optimization approach. *IEEE Transactions on Sustainable Energy*, 10(2):931–942, 2019. doi:10.1109/TSTE.2018.2855039.
- [84] TenneT. Annual market update 2023: Electricity market review focused on the netherlands and germany, 2024. Retrieved April 23, 2025, from <https://www.tennet.eu/nl/bedrijf/publicaties/technische-publicaties/>.
- [85] European Union. Directive (EU) 2018/2001 on the promotion of the use of energy from renewable sources (RED II), 2021. Retrieved May 12, 2025, from <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32018L2001>.
- [86] P. E. van den Berkmortel. Coupling of economic and engineering modeling approaches in an integrated framework for dynamic energy systems modeling: A proposal for utilizing economic engineering agent-based modeling for hydrogen market design. Literature Survey, Faculty of Mechanical, Maritime and Materials Engineering (3mE), Delft University of Technology, June 2023.
- [87] L. Vidas and R. Castro. Recent developments on hydrogen production technologies: State-of-the-art review with a focus on green-electrolysis. *Applied Sciences*, 11(23):11363, 2021. doi:10.3390/app112311363.
- [88] U. Weichenhain. Hydrogen transportation. the key to unlocking the clean hydrogen economy, 2021. Retrieved June 2, 2025, from https://www.rolandberger.com/publications/publication_pdf/roland_berger_hydrogen_transport.pdf.
- [89] T. Weitzel and C. H. Glock. Energy management for stationary electric energy storage systems: A systematic literature review. *European Journal of Operational Research*, 264(2):582–606, 2018. doi:10.1016/j.ejor.2017.06.052.
- [90] R. Weron. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4):1030–1081, 2014. doi:10.1016/j.ijforecast.2014.08.008.
- [91] P. Xiao, W. Hu, X. Xu, W. Liu, Q. Huang, and Z. Chen. Optimal operation of a wind-electrolytic hydrogen storage system in the electricity/hydrogen markets. *International Journal of Hydrogen Energy*, 45(46):24412–24423, 2020. ISSN 0360-3199. doi:10.1016/j.ijhydene.2020.06.302.
- [92] M. Yang, R. Hunger, S. Berrettoni, B. Sprecher, and B. Wang. A review of hydrogen storage and transport technologies. *Clean Energy*, 7, 2023. doi:10.1093/ce/zkad021.
- [93] M. Younas, S. Shafique, A. Hafeez, F. Javed, and F. Rehman. An overview of hydrogen production: Current status, potential, and challenges. *Fuel*, 316:123317, 2022. ISSN 0016-2361. doi:<https://doi.org/10.1016/j.fuel.2022.123317>.
- [94] M. Yue, H. Lambert, E. Pahon, R. Roche, S. Jemei, and D. Hissel. Hydrogen energy systems: A critical review of technologies, applications, trends and challenges. *Renewable and Sustainable Energy Reviews*, 146:111180, 2021. doi:10.1016/j.rser.2021.111180.
- [95] I. Zeier. How is the electricity price calculated?, 2023. Retrieved June 10, 2025, from <https://www.nordpoolgroup.com/en/the-power-market/how-is-the-electricity-price-calculated/>.
- [96] Y. Zhou, Z. Ma, J. Zhang, and S. Zou. Data-driven stochastic energy management of multi energy system using deep reinforcement learning. *Energy*, 261:125187, 2022. doi:10.1016/j.energy.2022.125187.