

Modeling of Battery Degradation for Integrated Design of Storage Systems

Thesis MSc Sustainable Energy Technology
Thodoris Tsonopoulos

Delft University of Technology

Modeling of Battery Degradation for Integrated Design of Storage Systems

by

Thodoris Tsonopoulos

to obtain the degree of Master of Science in Sustainable Energy Technology
at the Delft University of Technology,
to be defended publicly on Friday August 28, 2026 at 13:00 PM.

Student number: 6175902
Project duration: December 1, 2025 – August 28, 2026
Thesis committee: Dr. R. Schmehl, TU Delft, Chair
Ir. A. j. Lagugna, TU Delft, Second member
Dr. ir. J. Iori, TU Delft, supervisor

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Preface

This thesis marks the completion of my MSc degree in Sustainable Energy Technology at Delft University of Technology. The research was carried out within the Wind Energy section of the Faculty of Aerospace Engineering, and it examines how battery degradation influences the operation and the design of hybrid wind and storage power plants.

I would first like to express my sincere gratitude to my supervisor, Dr. Jenna Iori, for her guidance and support throughout this project. Our weekly meetings kept the work on track and gave me a place to test an idea before it went into the report. Her written feedback was detailed down to the level of individual sentences, and being asked to justify every claim was demanding at the time, but the report is a great deal better for it. I am equally grateful for how we approached the subject itself. Neither of us started with expertise in battery degradation modeling, and working through it together left us both understanding it far better than when we began, along with an extension to the SHIPP framework that will outlast this thesis.

I would also like to thank my friends back home, Kostas, Neria and Charis, whose support carried over the distance across this master and this thesis. My thanks go equally to the friends I made here in Delft, Greek and international alike, among them Dimitris, Ilia and Ahmed, to name only a few, who put up with me and made this city feel like a place to live rather than a place to study. I am thankful too for the people who were part of this journey for a time and are no longer part of my life, who shaped it just as much.

To my parents, thank you for the encouragement and the support, in every form, that made these studies possible. This master would not have happened without your immense support, financial and emotional, and without your constant encouragement to be better and to improve every day.

This work closes a path that began at the Department of Mechanical Engineering and Aeronautics in Patras, where an interest in renewable energy first took hold, and continued here in Delft. I am writing this at twenty-seven, no more certain about what comes next than I was at eighteen starting the bachelor, or at twenty-two finishing it. What I do believe is that both of them would be glad about the leap this took: not to a new city this time but to a new country, starting over again. I am grateful for these two years abroad and for everything they taught me, and to everyone who was part of them.

*Thodoris Tsonopoulos
Delft, August 2026*

Summary

Hybrid power plants that pair wind generation with battery storage are commonly designed with the battery treated as an ideal asset, or with degradation reduced to a fixed cost per megawatt-hour of throughput. Both simplifications remove the dependence of aging on how the battery is cycled. They cannot show which features of operation drive capacity fade, and they pass no degradation sensitivity to the sizing decision. This thesis asks to what extent explicitly accounting for battery degradation influences optimal operational strategies and design decisions in hybrid power plants.

The framework couples a linear dispatch program to a semi-empirical degradation model through rain-flow cycle counting, the cycle-extraction method used for fatigue analysis in mechanical systems. The dispatch maximizes day-ahead arbitrage revenue at a fixed battery size. Cycles are extracted from the resulting state-of-charge trajectory, converted to a degradation increment, and used to fade the capacity for the following project year. The battery is replaced when its state of health reaches the 70% end-of-life threshold. All reported degradation uses the physical reference model of Xu et al. A convex polynomial surrogate after Shi et al. is fitted to the same cycle-amplitude response and run as a second branch through the sizing study, so that the dependence of the design conclusion on the degradation model can be measured. The case study is the IEA Wind Task 50 onshore reference site in western Denmark: a 325 MW wind farm of 65 NREL 5 MW turbines and a 150 MW / 300 MWh lithium manganese oxide battery behind a 325 MW grid connection. It is valued over a 20-year horizon at a 3% real discount rate on DK1 day-ahead prices, using 2022 for the design study and 2019 as a price-regime sensitivity.

Two operational levers are identified. Cycle amplitude produces the larger change in total degradation, but it acts on cycle aging alone and is tied to revenue. Mean state of charge produces a smaller change, acts on both the cycle and the calendar component, and does so at no revenue cost, which makes it the more useful lever. At the degradation-aware optimum and the 10–90% window, calendar aging is 54% of the annual total, so a measure aimed at cycle amplitude alone captures less than half of the degradation. The two window parameters act differently: width trades revenue against degradation, while the window center reduces degradation at constant revenue.

Pricing degradation reduces the optimal battery from 800 MWh and 250 MW to a broad plateau near 550 MWh and 175 MW, a reduction of about 30%, driven by replacement cost rather than by lost revenue. The result does not depend on which degradation model is used: the physical reference model and the polynomial surrogate select the same design. Price volatility has a larger effect than degradation on whether the battery is worth building at all: at the reference design its battery-marginal net present value is –89.7 MEUR under 2019 prices and +42.9 MEUR under 2022 prices.

An attempt to embed the degradation cost directly in the dispatch, as a single nonlinear program, fails for a structural reason. Rainflow cycle counting has a discontinuous gradient where cycles merge or split, so a smooth-gradient solver fails to certify an optimum on nearly half the days attempted, returns no improvement on the linear program where it does, and most of the apparent improvement across the year comes from solves that did not converge. The alternative keeps the dispatch linear and carries the degradation feedback through a sub-gradient, which the polynomial surrogate admits because it is convex and the physical reference model is not. Its two degradation derivatives are implemented and verified here. Closing the outer loop additionally requires the sensitivity of the optimal dispatch to capacity, which remains future work.

No measured cell or pack data is used, so the framework is verified against its own specification rather than validated against a physical battery, and the dispatch assumes perfect price foresight on a single market.

Contents

Preface	i
Summary	ii
Nomenclature	xii
1 Introduction	1
1.1 Context and Motivation	1
1.2 Hybrid Power Plants and Storage Dispatch	1
1.2.1 Hybrid Power Plants	1
1.2.2 Storage System Dispatch	2
1.3 System Integration and Socio-Economic Context	2
1.4 Literature Overview and Research Gaps	3
1.5 Research Question and Objectives	6
2 Battery Degradation in Lithium-Ion Systems	7
2.1 Operating Principle of Lithium-Ion Batteries	7
2.2 Degradation Mechanisms	7
2.2.1 Calendar Aging	8
2.2.2 Cycle Aging	8
2.2.3 Relative Contribution	8
2.3 Quantifying Degradation: State of Health	9
2.4 Cycle Counting Under Irregular Operation	9
2.5 Degradation Modeling Approaches	10
2.6 The Physical Reference Model (Xu 2016)	11
2.6.1 Model Structure	11
2.6.2 Stress Factor Models	12
2.6.3 End-of-Life Definition	15
2.7 Polynomial Degradation Model (Shi 2018)	15
2.7.1 Polynomial Stress Function	15
2.7.2 Scope of the Polynomial Model	16
2.7.3 Convexity of the Rainflow Degradation Cost	16
2.8 Sub-gradient of the Rainflow Degradation Cost	17
2.8.1 Degradation Cost and the Dispatch Variables	17
2.8.2 Sub-gradient with Respect to Power	18
2.8.3 Efficiency Ratio Identity	18
2.8.4 Differentiation with Respect to Cycle Amplitude Only	19
3 Methodology	20
3.1 Modeling Framework	20
3.1.1 Integrated Evaluation of a Battery Design	21
3.1.2 Investigative Approach	21
3.2 Case Study and Input Data	23
3.2.1 Reference Site and Plant	23
3.2.2 Electricity Price Data and Economic Parameters	23
3.3 Wind Resource & Power Generation Model	23
3.3.1 Site and Wind Farm Configuration	24
3.3.2 Wind Resource Data	24
3.3.3 Wake and Turbulence Model	26
3.4 Battery Dispatch Optimization	27
3.4.1 Optimization Problem	27

3.4.2	Objective Function	28
3.4.3	Storage Model	28
3.4.4	Storage Constraints	30
3.4.5	Dispatch Constraints	30
3.4.6	Periodicity and inter-year continuity.	30
3.4.7	Implementation	31
3.5	Degradation Evaluation	31
3.5.1	The Two Degradation Models	31
3.5.2	Fitting the Polynomial Stress Function	33
3.5.3	Degradation Sub-gradient and the Capacity Gradient	36
3.5.4	Gradient Verification	36
3.6	Capacity Reduction	37
3.6.1	Multi-Year Degradation and Battery Replacement	37
3.7	Design Space Exploration via Parameter Sweep	39
3.7.1	Stage 1: Energy and Power Capacity Sweep	40
3.7.2	Stage 2: SoC Operating Window Sweep	41
3.8	Towards a Gradient-Based Degradation-Aware Sizing	41
3.8.1	The Monolithic NLP Formulation	42
3.8.2	Non-Smoothness of the Rainflow Degradation Cost	44
3.8.3	A Gradient-Based Sizing Method	44
3.9	Modeling Limitations	45
4	Results	48
4.1	Electricity Price Regimes	48
4.2	Model Verification	50
4.2.1	Verification Scope and Data	50
4.2.2	Dispatch and Cycle Counting	51
4.2.3	Response to the Operating Window	53
4.3	Baseline Behavior at the Reference Design	56
4.3.1	State-of-Health Trajectory	56
4.3.2	Degradation Components: Cycle and Calendar Aging	57
4.3.3	Price Sensitivity of Lifetime, Revenue and Net Present Value	60
4.4	Parameter Sweep	61
4.4.1	Energy and Power Capacity Sweep	61
4.4.2	State-of-Charge Operating Window Sweep	65
4.4.3	Fidelity of the Polynomial Surrogate	68
4.5	Monolithic Nonlinear Program	69
4.5.1	Convergence Behavior	69
4.5.2	Dispatch Outcome	70
5	Conclusions and Recommendations	73
5.1	Link to the Research Gaps	73
5.2	Research Objective	74
5.3	Answers to the Research Questions	74
5.4	Impact	75
5.5	Limitations	76
5.6	Recommendations for Future Research	77
	References	78
A	Justification for the Use of AI Tools	82
B	Directional Wind Resource	83
C	Gradient Verification	85
C.1	Test 1: Frozen-Dispatch Capacity Gradient	85
C.2	Test 2: Per-Timestep Dispatch Sub-gradient	88
C.3	Fused Per-Year Sensitivity	92

D	Verification on Synthetic Input	95
D.1	Rainflow Counter on Synthetic Traces	95
D.2	Efficiency Threshold in the Dispatch	96
E	Monolithic NLP Supporting Material	98
E.1	Per-Day Exemplar Dispatch Panels	98
E.2	Distribution of Outcomes	102
E.3	Representative Sample of Per-Day Results	103

List of Figures

1.1	Three approaches to coupling a rainflow-based degradation model to a dispatch optimization. The five blocks and their positions are identical in each panel. Under post-processing (a) the degradation result never returns to the optimization; under the iterative-sequential approach (b) it returns between solves; under the embedded approach (c) it is evaluated at every solver iteration rather than between solves. The inputs are the wind and price time series and the battery energy and power capacity, and the outputs are the capacity fade, the end of life and the net present value.	5
2.1	Rainflow half-cycle identification on a synthetic SoC trajectory. Red circles mark lower turning points, green circles upper turning points. The solid bracket spans a charge half-cycle, the dashed bracket the discharge half-cycle that follows it. Dotted lines link the third half-cycle to its cycle amplitude δ_i and mean state of charge σ_i	10
2.2	Cycle amplitude stress factor $S_\delta(\delta)$ and mean SoC stress factor $S_\sigma(\sigma)$ of the physical reference model, over the amplitude and mean SoC ranges permitted by a 10–90% operating window.	13
2.3	Isolated effect of S_σ . Two cycles of equal amplitude ($\delta = 0.30$) at mean SoC $\sigma = 0.25$ (left) and $\sigma = 0.75$ (right), each shown in color within its surrounding dispatch trajectory in grey. Both are taken from a simulated annual dispatch produced by the model of this thesis.	14
3.1	Three-way coupling between battery size, operational strategy and degradation rate. . .	21
3.2	The shared evaluation pipeline and the first two investigative approaches. In the baseline evaluation (a), the dispatch LP (Section 3.4) is solved for a fixed design and its SoC trajectory is passed to the physical (Xu) and convex (Shi) degradation models (Section 3.5), which return lifetime NPV and end-of-life through the capacity reduction step (Section 3.6). The parameter sweep (b, Section 3.7) repeats this evaluation over a grid of designs and selects the maximum-NPV point. The monolithic NLP (c, Section 3.8.1) instead places the degradation cost inside the optimization objective. The third approach, gradient-based sizing, is shown separately in Figure 3.13 (Section 3.8.3).	22
3.3	Omnidirectional wind-speed distribution at 90 m hub height for the site, 2022. Histogram: hourly ERA5 wind speeds; curve: two-parameter Weibull fit ($A = 9.24$ m/s, $k = 2.39$). The wake model uses the directional twelve-sector fit instead (Table B.1, Figure B.1). Data are taken from the nearest ERA5 grid cell and therefore approximate the resource near the site.	25
3.4	Wind rose at 90 m hub height for the site, 2022, with wind speed binned by turbine operating state (below cut-in, partial load, approaching rated, at or above rated, above cut-out). Data are taken from the nearest ERA5 grid cell and therefore approximate the resource near the site.	26
3.5	XDSM of the dispatch LP. The fixed design parameters $\bar{E}, \bar{P}$ enter above the optimizer and the wind and price series above the objective; the optimizer iterates between the revenue objective (Equation 3.2a) and the constraint set (Equations 3.2b–3.2i) to optimality, returning the optimal dispatch. The reading convention and data exchanges are described in the text above.	29
3.6	The two computation paths, from the shared rainflow cycles to the simulation outputs (top) and the optimization input (bottom). The Shi accumulation of Equation 3.8 also maps to State of Health, for the model comparison reported in Chapter 4; that path is not drawn here.	33

3.7	Second derivative of $S_\delta(\delta)$. The shaded region ($\delta < 0.144$) is non-convex; the sign change is marked by the dashed line. The vertical axis is truncated at 12, and S''_δ continues to rise beyond the frame, reaching about 35 at $\delta = 0.80$	34
3.8	Φ_{shi} (dashed) fitted to S_δ (solid) over $\delta \in [0.15, 0.80]$. The lower panel confirms $\Phi''_{\text{shi}} > 0$ everywhere while S''_δ goes negative below $\delta_c = 0.1437$. The vertical axis of the lower panel is truncated at 12; both second derivatives continue beyond the frame.	35
3.9	$\Phi_{\text{shi}}(\delta)$ and its second derivative, for the coefficients fitted over $\delta \in [0.15, 0.80]$ under the 10–90% SoC window. The second derivative is positive everywhere, confirming global convexity.	35
3.10	Multi-year simulation loop. The effective energy capacity $\bar{E}_k = \bar{E} \cdot \text{SoH}_{k-1}$ is the only dispatch input that changes from year to year, as marked in the figure; the wind and price series and the nominal power $\bar{P}$ are held fixed.	38
3.11	Two-stage parameter sweep structure. Stage 1 (left): lifetime NPV is evaluated across an $(\bar{E}, \bar{P})$ grid under three degradation scenarios to identify the optimal battery sizing. Stage 2 (right): the SoC operating window is varied at the identified optimum to quantify the effect of window width and window center on degradation and NPV.	40
3.12	XDSM of the monolithic NLP. Inputs and the dispatch output follow the dispatch LP of Figure 3.5; the additional rainflow analysis and gradient blocks, and the discontinuity marked in the figure, are described in the text above and in Section 3.8.2.	43
3.13	Conceptual gradient-based sizing method. The sizing search would set the energy capacity $\bar{E}$ and pass it to the dispatch LP, which builds the SoC trajectory, counts rainflow cycles and accumulates degradation. The total degradation would then be used to compute the capacity gradient $d\text{NPV}/d\bar{E}$, which measures how lifetime NPV responds to a change in capacity; that value would drive the next update of $\bar{E}$, and the process would repeat.	45
4.1	Hourly DK1 day-ahead prices for 2019 and 2022. The dashed line marks zero price.	49
4.2	Empirical cumulative distribution of the within-day price spread for DK1 2019 and 2022. The dotted lines mark the median of each year (Table 4.1).	50
4.3	One week of the solved year-1 dispatch, 300 MWh / 150 MW battery, 10–90% window, DK1 2022 prices. Top: day-ahead price. Middle: wind output, with battery discharge positive and charge negative. Bottom: state of charge, with the minimum and maximum allowed values dashed. Week 20, hours 3192 to 3360, all 168 hours shown.	52
4.4	Rainflow cycles of week 20, 10–90% window, all 168 hours. Top: state-of-charge trajectory, with the minimum and maximum allowed values dashed and the counted turning points marked. Bottom: one bar per counted cycle, spanning the state-of-charge levels its two turning points join and placed at their midpoint in time. Cycles of zero amplitude appear as a single dot. Three cycles with their midpoint in an adjacent week are not shown. The scale bar at the right gives $\delta_c = 0.1437$	53
4.5	Effect of the state-of-charge window width, at fixed center 0.50, 300 MWh / 150 MW battery. Left: state-of-charge trajectory for each window, with that window's bounds dashed in the same color, first 72 hours. Right: summed multiplicity $\sum_i n_i$ of the cycles counted over the full week, binned by amplitude, with the three windows drawn side by side within each bin so that coincident values stay visible. The lower panel magnifies the range below $\delta = 0.42$, where every bar carries $\sum_i n_i \leq 1$; triangles mark the three amplitudes at which all three windows agree. Dotted line: $\delta_c = 0.1437$	54
4.6	State-of-health trajectory over the 20-year project for the 150 MW / 300 MWh reference battery, under DK1 2019 and DK1 2022 prices. The vertical segment in each line is the battery replacement, where the state of health returns to 100%. The dashed lines mark the 80% (IEC/EV) and 70% (warranty) end-of-life levels.	57
4.7	Annual degradation rate f_d split into cycle and calendar components over the 20-year project, for the 150 MW / 300 MWh reference battery. For each project year the left bar is DK1 2019 and the right bar is DK1 2022.	59
4.8	Calendar share of the annual degradation rate f_d over the 20-year project, for the 150 MW / 300 MWh reference battery under DK1 2019 and DK1 2022 prices. The cycle share is the complement of the plotted value. The dashed line marks an equal split.	59

4.9	Lifetime NPV across the full $(\bar{E}, \bar{P})$ grid, coarse pass, DK1 2022 prices. The star marks the NPV-maximizing design in each panel and the diamond marks the no-degradation optimum; the two coincide in (a). Gray dots are the evaluated grid points and the dashed rectangle marks the region covered by the refined pass. The uncolored corner regions are designs removed by the duration and value filters of Section 3.7.1 and not evaluated. The two panels use independent color scales, so a given color does not carry the same NPV across panels.	62
4.10	Lifetime NPV over the refined grid under the two degradation models, DK1 2022 prices. In each panel the star marks the grid optimum, the open circle the continuous quadratic optimum (Equation 3.16) and the gray dots the evaluated grid points. The diamond marks the best no-degradation design among the refined-grid points; it lies at the corner because the no-degradation optimum is outside this window (Figure 4.9).	62
4.11	Lifetime NPV as a function of energy capacity at fixed $\bar{P} = 175$ MW, DK1 2022 prices. Markers are the evaluated designs of both passes, so the spacing is 50 MWh between 450 and 750 MWh and wider outside that range. The dotted line marks the Xu optimum at 550 MWh.	63
4.12	Lifetime NPV as a function of power capacity at fixed $\bar{E} = 550$ MWh, DK1 2022 prices. The dotted line marks the Xu optimum at 175 MW.	64
4.13	Battery revenue, cost and their rates of change against energy capacity at fixed $\bar{P} = 175$ MW, DK1 2022 prices. Quantities in (a) are discounted to year 0. Rates in (b) are forward differences between adjacent grid points, plotted at the interval midpoint; the series starts at 450 MWh because the grid spacing below that is 150 MWh. The dotted line marks the Xu optimum at 550 MWh.	65
4.14	Lifetime NPV against year-1 annual degradation rate for the two window series, battery fixed at 550 MWh / 175 MW, DK1 2022 prices, Xu branch. The width series varies $\delta_{\max}$ at fixed center 0.50; the center series varies the center at fixed width 0.80. The two share the 10–90% window, circled, and are drawn on the same axes. Only the endpoints of the center series are labeled; Table 4.8 lists all five. Degradation is shown here in percent per year and in the tables as a fraction.	67
4.15	Objective difference of the nonlinear dispatch relative to the LP, against final solver optimality, for the 356 active days at $\bar{E} = 550$ MWh and $\bar{P} = 175$ MW under DK1 2022 prices. Navy points are converged solves (optimality below 10^{-6}) and dark red points are solves that did not converge. The dashed line marks the convergence threshold and the dotted line the divergence threshold at 10^{-2} . The difference is clipped to +150 EUR for display, which affects six days, all of them non-converged. The four example days of Table 4.11 are labeled.	71
B.1	Per-sector Weibull distributions at 90 m hub height, 2022. The twelve colored curves are the directional fits (solid: sectors 15° – 195° ; dashed: 225° – 345°); the black curve is the omnidirectional fit. Legend entries give the sector center and its frequency. Data are taken from the nearest ERA5 grid cell.	84
C.1	Finite-difference error of the cycle-degradation slope against the analytical slope of Equation C.2, over the full range of step sizes h , for the central and the forward difference. On the descent, truncation error falls at second order in h for the central difference and at first order for the forward difference, marked by the guide lines. On the rise, rounding error grows as the step shrinks. The central-difference minimum reaches order 10^{-12} near $\epsilon^{1/3} = 6 \times 10^{-6}$, which is agreement to about eleven digits.	87
C.2	Schematic of a rainflow re-pairing. Left: a state-of-charge trace with one ambiguous excursion, and the two readings the counter can give it. Right: the degradation cost against p_t^{ch} at step t . The two readings meet where the comparison is tied, so the cost is continuous, but they arrive with the different slopes D^- and D^+ of Equation C.6. The interval between those slopes is the bracket of Equation C.7. The figure carries no data from a run.	89

C.3	Sub-differential bracket at 48 sampled corner time steps from each price year, ordered by the position of the analytic value between the two quotients of Equation C.6. The grey bar spans the two quotients and its ends are those two values. Over the full populations the analytic sub-gradient lies inside the bar at every time step, at an endpoint in 85 of 97 for DK 2019 and 106 of 123 for DK 2022. The central difference is the mean of the two quotients and lies at the middle of the bar throughout.	91
C.4	Fused per-year sensitivity g_k , DK1 2022 prices, years 1 to 12, normalized to its annual mean. The dotted line marks the annual mean.	93
C.5	The magnitude and alignment terms overlaid on the g_k bars of Figure C.4 at reduced opacity, DK1 2022 prices. The magnitude term follows the bar heights; the alignment term does not.	94
D.1	Rainflow read-back on the two synthetic state-of-charge traces. Top: three periods of the constant-amplitude trace, with the 10% and 90% window bounds shown as dashed lines. Bottom: the nested trace, with each counted cycle drawn as a vertical bar spanning its amplitude, placed at the midpoint of the interval it was extracted from. The four ripples at $\delta = 0.10$ are resolved separately from the excursion at $\delta = 0.80$	96
D.2	Energy discharged over three days against the price ratio between the expensive and the cheap block. Left: the full sweep. Right: detail around the transition, markers only, since no ratio between the two was solved. The battery is idle at 1.0975 and saturates at 1.0995, bracketing the predicted threshold $1/(\eta_{in}\eta_{out}) = 1.0989$ within 0.18%. Six of the thirteen sampled ratios lie between the idle and the saturated band and are shown but not asserted, because the linear program is close to degenerate in a narrow interval around the threshold.	97
E.1	Type A, converged gain: day D049 (Low regime). Converged solve, final optimality 9.5×10^{-7} (gtol). Revenue lower by 28 EUR, degradation cost lower by 115 EUR, objective change +87 EUR.	99
E.2	Type B, converged loss: day D179 (High regime). Converged solve, final optimality 2.2×10^{-7} (gtol). Revenue lower by 22 EUR, degradation cost higher by 24 EUR, objective change -45 EUR.	100
E.3	Type C, stalled: day D084 (Mid regime). Final optimality 3.9×10^{-3} (xtol, step collapse), above the convergence threshold. Revenue lower by 31 EUR, degradation cost lower by 138 EUR, objective change +107 EUR.	101
E.4	Type D, diverged: day D363 (Low regime). Final optimality 8.6×10^9 (maxiter). Reported degradation cost lower by 986 EUR on a trajectory that does not change over the day, with about 0.5 EUR of revenue separating the two dispatches; objective change +986 EUR.	102

List of Tables

2.1	Cycle amplitude stress function coefficients from Xu et al. [63], repeated with the full parameter set in Table 2.2.	13
2.2	Parameters of the physical reference model for LMO cells, reproduced without modification from Table I of Xu et al. [63].	15
3.1	Inputs and outputs of the integrated evaluation of a single battery design.	21
3.2	Economic parameters used in the NPV calculation. Battery cost components follow the Danish Energy Agency technology-data catalogue for energy storage, lithium-ion technology, real EUR ₂₀₂₀ [10]. Initial costs use the 2025 technology year; replacement costs use the 2035–2040 expansion-cost figures.	23
3.3	Key parameters of the NREL 5 MW reference turbine used in this study.	24
3.4	PyWake configuration used to generate the annual power time series.	27
3.5	Dispatch LP constraints: standard SHIPP compared with the extensions made for this thesis. The modified SHIPP source is available at [61].	31
3.6	Composition of the two degradation models as applied in this thesis. Both use the same rainflow cycles, the same calendar term and the same capacity reduction, so the two accumulations differ in the cycle-amplitude stress function alone. The last two rows show where the branches part company: what each is used for, and how much of the accumulation is carried into the derivative.	33
4.1	DK1 day-ahead price statistics for 2019 and 2022. The daily spread is the within-day maximum minus minimum price.	49
4.2	Response of the implementation to a change in the operating window. Predictions are fixed before the runs. Width-series entries are given in the order 20–80, 10–90, 0–100%; center-series entries in the order 20–100, 0–80% against the 10–90% reference. The entry “none” marks the widest window, which buys rather than sells at hour 95.	55
4.3	Week statistics for the width series, at fixed center 0.50, battery 300 MWh / 150 MW, week 20 of DK1 2022. $\sum_i n_i$ is the summed cycle multiplicity. EFC is the discharge throughput divided by nominal capacity, $EFC = \sum_t \max(p_t^s, 0) \Delta t / \bar{E}$, as used for the annual figures elsewhere in this chapter. $\bar{\delta}$ is the mean amplitude over the counted cycles. Degradation columns are the week’s contribution on the Xu reference branch.	56
4.4	Year-1 degradation components and dispatch statistics, 10–90% state-of-charge window, 300 MWh LMO battery. Cycle quantities are defined in Table 4.3. The two components are rounded independently and need not sum to the rounded total.	58
4.5	End-of-life years and project economics under DK1 2019 and 2022 prices, 10–90% window, 300 MWh LMO battery. End-of-life years are the crossings of the simulated state-of-health trajectory, interpolated between year-end values. The 70% year quoted for 2022 is the year the level would have been crossed had the battery remained in service, since replacement is triggered at 70.5%.	60
4.6	Grid and continuous optima from the energy and power capacity sweep, DK1 2022 prices. A dash denotes a value not computed.	63
4.7	Width series: window swept at fixed center 0.50, battery fixed at 550 MWh / 175 MW, DK1 2022 prices. Degradation quantities are Year-1 values on the Xu reference branch; t_{repl} is the replacement year; SoH_{20} is the state of health at year 20. The 10–90% window is the reference shared with Table 4.8.	66
4.8	Center series: window swept at fixed width $\delta_{\text{max}} = 0.80$, battery fixed at 550 MWh / 175 MW, DK1 2022 prices. Degradation quantities are Year-1 values on the Xu branch; t_{repl} is the replacement year; SoH_{20} is the state of health at year 20. The 10–90% window is the reference shared with Table 4.7.	66

4.9	Xu and Shi branches across the width series, at fixed center 0.50, battery 550 MWh / 175 MW, DK1 2022 prices. f_d is the year-1 annual degradation rate, t_{repl} the replacement year and R^2 the fit quality of the power law over the amplitude range of the window. The last column is the share of year-1 cycle degradation, on the Xu branch, carried by cycles below $\delta_c = 0.1437$, where the surrogate extrapolates outside its fitting range.	68
4.10	Monolithic solver outcome by price-variability regime, 356 active days, 550 MWh / 175 MW, DK1 2022. The objective difference is positive where the nonlinear dispatch is worth more than the LP dispatch.	70
4.11	The four solver behaviors across the 356 active days, 550 MWh / 175 MW, DK1 2022.	71
B.1	Per-sector Weibull parameters and frequencies, 90 m hub height, 2022 (nearest ERA5 grid cell).	83
C.1	Test 1 quantities for the two price years. The cycle and calendar degradations are model outputs. The cycle slope is Equation C.2, checked by the finite difference of Figure C.1. The calendar coupling is Equation C.3. Their sum is the bracket of Equation C.5.	87
C.2	Test 2 quantities on the solved Year 1 trajectories of the two price years. The two populations are separated at the reference perturbation of 10^{-6} fractional SoC. Relative disagreements are normalized by the larger of the analytic value and the central difference, with a floor at one thousandth of the median analytic magnitude, so that time steps where the sensitivity passes through zero do not inflate the ratio.	90
E.1	Distribution of the change in objective value across the 356 active days, split by solver convergence. Most days cluster just below zero, and the positive bins are dominated by non-converged solves.	102
E.2	Sign of the objective change and positive mass by convergence band. The positive mass is the sum of all positive changes in the band. Converged days are the majority of the sample but carry 13% of the positive mass; the remainder comes from solves that did not reach an optimum.	103
E.3	Representative per-day results: the ten largest objective gains and six typical converged losses. Reg: regime (H/M/L). Spr: daily price spread (EUR/MWh). ΔRev : revenue difference, positive when the reported dispatch earns less (EUR). ΔDeg : degradation cost difference, positive when it degrades less (EUR). Rat: ratio of the two. ΔObj : change in objective value against the LP (EUR). It: iterations. Opt: final optimality; a day is converged when this is below 10^{-6} . The full dataset is summarized in Tables E.1 and E.2.	103

Nomenclature

Abbreviations

Abbreviation	Definition
ASTM	American Society for Testing and Materials
BESS	Battery Energy Storage System
BPA	Bastankhah–Porté-Agel (Gaussian wake model)
CAPEX	Capital Expenditure
DK1	Danish day-ahead electricity bidding zone (western Denmark)
DoD	Depth of Discharge: maximum cycle depth set by the SoC operating window, parameter $d = 1 - \sigma_{\min}$ (distinct from the cycle amplitude δ)
DST	Dynamic Stress Test
DTU	Technical University of Denmark
EFC	Equivalent Full Cycle
ENTSO-E	European Network of Transmission System Operators for Electricity
EoL	End of Life
ERA5	ECMWF Reanalysis v5 (wind resource dataset)
HPP	Hybrid Power Plant
IEA	International Energy Agency
IEC	International Electrotechnical Commission
LFP	Lithium Iron Phosphate (battery chemistry)
LMO	Lithium Manganese Oxide (battery chemistry)
LP	Linear Program
NCA	Nickel Cobalt Aluminum (battery chemistry)
NLP	Nonlinear Program
NMC	Nickel Manganese Cobalt (battery chemistry)
NPV	Net Present Value
NREL	National Renewable Energy Laboratory
P2D	Pseudo-Two-Dimensional (electrochemical model)
SEI	Solid Electrolyte Interphase
SHIPP	Sizing optimization for Hybrid Power Plants (optimization framework)
SoC	State of Charge
SoH	State of Health
UTC	Coordinated Universal Time
XDSM	eXtended Design Structure Matrix

Symbols

Decorations follow a fixed convention. Bold face denotes the vector of a quantity over the horizon, so $\mathbf{p}^{\text{ch}}$ is the vector of p_t^{ch} . A bar denotes a rated design value ($\bar{E}$, $\bar{P}$) or a mean over the stated set ($\bar{\sigma}$, $\bar{T}_c$). A hat denotes a fitted or pinned value and a star an optimal value. On degradation quantities, the superscripts X_u and Sh_i select the degradation branch. Composites of declared symbols are not listed separately.

Symbol	Definition	Unit
A	Weibull scale parameter (omnidirectional wind-speed fit)	[m/s]
a_i	Opening turning-point (time-step) index of rainflow cycle i	[–]
$b_0, \dots, b_5$	Coefficients of the quadratic NPV surrogate	[varies]

Symbol	Definition	Unit
b_i	Closing turning-point (time-step) index of rainflow cycle i	[-]
$C(t)$	Usable battery capacity at time t	[MWh]
C_0	Nominal capacity of a new cell	[MWh]
C_p	Turbine power coefficient	[-]
C_{repl}	Battery replacement capital cost	[€]
D	Rotor diameter	[m]
d	Depth of discharge (SoC operating-window depth), $d = 1 - \sigma_{\min}$	[-]
D_h	Finite-difference estimate of a derivative at step h (distinct from the rotor diameter D)	[varies]
D^+, D^-	Forward and backward difference quotients of the degradation cost at one time step	[€/MWh]
e_0	Battery stored energy before the first time step, used where the trajectory is written as a running sum	[MWh]
e_t	Battery stored energy at time step t	[MWh]
$e_{\min}, e_{\max}$	Minimum and maximum allowable stored energy ($= \sigma_{\min} \bar{E}$, $= \sigma_{\max} \bar{E}$)	[MWh]
$\bar{E}$	Nominal battery energy capacity	[MWh]
$\bar{E}_k$	Effective energy capacity in year k , $\bar{E}_k = \bar{E} \cdot \text{SoH}_{k-1}$	[MWh]
$\bar{E}^*$	Continuous (quadratic-surrogate) optimal energy capacity	[MWh]
EFC	Equivalent full cycles over a period, $\sum_t \max(p_t^i, 0) \Delta t / \bar{E}$	[-]
f	Rainflow degradation cost, $f = \lambda_E^{\text{repl}} \bar{E} \sum_i n_i \Phi(\delta_i)$; carries the cycle-amplitude stress alone, with no mean-SoC, temperature or calendar term	[€]
f_d	Linearized degradation fraction, mapped to capacity loss by the SEI fade curve $L(f_d)$	[-]
$f_d^{\text{cyc}}, f_d^{\text{cal}}$	Cycle-aging and calendar-aging contributions to f_d	[-]
$f_{d,i}$	Life fraction lost by rainflow cycle i (per-cycle loss)	[-]
$f_{d,k}$	Annual degradation increment in year k	[-]
$f_{d,\text{cum},k}$	Cumulative degradation at the end of year k , $\sum_{j \leq k} f_{d,j}$, reset on replacement	[-]
g_k	Fused per-year sensitivity signal, $-\bar{E}_k \langle s_k, \mu_k \rangle$ (examined and not used)	[€/MWh]
h	Finite-difference step size; a fraction of the capacity in the sub-gradient check and a relative step in the capacity check	[-]
I	Ambient turbulence intensity	[-]
J	Objective of the monolithic NLP (Equation 3.17)	[€]
k	Weibull shape parameter (wind-speed fit); see also the year index in Subscripts	[-]
k_3	Shi polynomial stress coefficient	[-]
k_4	Shi polynomial exponent ($k_4 > 1$ for convexity)	[-]
k^*	Wake expansion coefficient (BPA model)	[-]
k_T	Temperature stress coefficient (Xu model)	[-]
$k_{\delta 1} - k_{\delta 3}$	Cycle-amplitude stress coefficients (Xu model)	[-]
k_{σ}	Mean-SoC stress coefficient (Xu model)	[-]
k_{τ}	Calendar-aging (time) rate coefficient (Xu model)	[s ⁻¹]
L	Battery capacity loss (fraction), $L = 1 - \text{SoH}$, via the SEI fade curve $L(f_d)$	[-]
N	Project lifetime	[years]
n_i	Cycle-count weight of rainflow record i ($= 1$ for a full cycle, $\frac{1}{2}$ for a half cycle)	[-]
$\text{NPV}_{\text{revenue}}$	Discounted revenue over the project horizon, before replacement cost	[€]
NPV_{net}	Lifetime NPV after subtracting discounted replacement cost	[€]

Symbol	Definition	Unit
$\widehat{\text{NPV}}$	Quadratic surrogate of lifetime NPV over the refined grid	[€]
$P(v)$	Turbine power curve as a function of wind speed	[MW]
$P_{\max}$	Grid connection power capacity	[MW]
$P_{\text{wind}}(t)$	Wind-farm power output at time step t	[MW]
$\bar{P}$	Battery power capacity	[MW]
$\bar{P}^*$	Continuous (quadratic-surrogate) optimal power capacity	[MW]
p_t^s	Signed storage power at time step t ($p_t^s < 0$ charging, $p_t^s > 0$ discharging)	[MW]
p_t^c	Curtailed wind power at time step t	[MW]
$p_t^{\text{ch}}, p_t^{\text{dis}}$	Charging and discharging power at time step t (sub-gradient derivation, ≥ 0)	[MW]
r	Discount rate	[-]
R^2	Coefficient of determination of a fit	[-]
s	Uniform scale factor applied to the SoC trajectory in the frozen-dispatch check	[-]
s_t	Dispatch degradation sub-gradient at time step t , $\partial f / \partial p_t$, where f is the degradation cost in euros	[€/MW]
s_k	Vector of $\partial f_d / \partial e_t$ over year k , used only in the fused signal	[MWh ⁻¹]
g_k		
S_δ	Cycle-amplitude stress factor (Xu model)	[-]
S_σ	Mean-SoC stress factor (Xu model)	[-]
S_T	Temperature stress factor (Xu model)	[-]
S_τ	Time (calendar) stress factor within a cycle, $S_\tau(t) = k_\tau t$ (Xu model)	[-]
$\mathcal{S}(t)$	Set of rainflow cycles spanning time step t	[-]
$\text{SoC}_{\min,i}$, $\text{SoC}_{\max,i}$	Lower and upper state of charge reached by rainflow half-cycle i	[-]
SoH_k	State of health at the end of project year k	[-]
$\text{SoH}_{\text{eol}}, \text{SoH}_{\text{repl}}$	End-of-life state of health, 0.70, and the replacement trigger, 0.705, half a percentage point above it	[-]
T	Number of time steps in the horizon; 8,760 for the annual dispatch, 24 for the monolithic NLP	[-]
T_c	Battery cell temperature (held constant at 25 °C)	[K]
T_{ref}	Reference temperature for the Arrhenius stress model	[K]
t	Time (also the simulation period duration in calendar aging)	[s] or [h]
t_{repl}	Project year in which battery replacement occurs	[years]
$u(z)$	Wind speed at height z (power-law shear profile)	[m/s]
u_{ref}	ERA5 wind speed at reference height z_{ref}	[m/s]
v	Wind speed (turbine power-curve argument)	[m/s]
$\mathbf{x}$	State-of-charge profile (vector over the horizon)	[-]
z, z_{ref}	Height above ground and the ERA5 reference height	[m]
α	Wind shear exponent (power-law profile)	[-]
$\alpha_{\text{sei}}, \beta_{\text{sei}}$	SEI fade-curve coefficients (Xu model)	[-]
γ	Annuity factor converting one simulated year to lifetime NPV	[-]
δ, δ_i	Cycle amplitude: fraction of capacity cycled in a single rain-flow event, cycle i where indexed	[-]
$\bar{\delta}$	Mean cycle amplitude over all identified cycles in a year	[-]
δ_c	Inflection amplitude of S_δ , $\delta_c = 0.1437$; S_δ is convex above it and concave below	[-]
$\delta_{\min}$	Lower bound of the Φ_{shi} fitting range	[-]
$\delta_{\max}$	Maximum achievable cycle amplitude, $\delta_{\max} = \sigma_{\max} - \sigma_{\min}$	[-]
Δf_d	Difference in annual f_d between the Xu and Shi branches	[-]
Δt	Simulation time-step duration	[h]

Symbol	Definition	Unit
ε	Curtailment regularization weight in the dispatch objective ($\varepsilon = 1 - 10^{-6}$)	[-]
ϵ	Machine epsilon of double precision, 2.2×10^{-16} (distinct from the curtailment weight ε)	[-]
η	One-way efficiency where charging and discharging are symmetric, $\eta = \eta_{\text{in}} = \eta_{\text{out}}$	[-]
$\eta_{\text{in}}, \eta_{\text{out}}$	Battery charging and discharging efficiency (SHIP convention, $\in (0, 1]$)	[-]
η_d	Discharge efficiency in the Shi convention, $\eta_d = 1/\eta_{\text{out}} \geq 1$	[-]
λ_t	Day-ahead electricity price at time step t	[€/MWh]
λ_E, λ_P	Battery energy and power capital cost	[€/MWh], [€/MW]
$\lambda_E^{\text{repl}}, \lambda_P^{\text{repl}}$	Battery replacement energy and power cost	[€/MWh], [€/MW]
μ_k	Vector of LP dual prices on the lower SoC bound over year k ; μ_t its entry at time step t	[€/MWh]
ρ	Air density	[kg/m ³]
σ, σ_i	Mean state of charge of a rainflow half-cycle, cycle i where indexed	[-]
$\bar{\sigma}$	Mean SoC over the simulation period (calendar aging); also the SoC window center	[-]
$\hat{\sigma}$	Fixed SoC fraction for inter-year continuity pinning	[-]
σ_{ref}	Reference mean SoC at which $S_{\sigma} = 1$	[-]
$\sigma_{\text{min}}, \sigma_{\text{max}}$	Lower and upper SoC operating-window bound	[-]
σ_{Φ}	Φ -weighted mean SoC, $\sigma_{\Phi} = \sum_i f_{d,i} \sigma_i / \sum_i f_{d,i}$	[-]
$\tau_{c,i}$	Time spent in rainflow cycle i (Xu general form)	[s]
Φ	Cycle-amplitude stress function (general)	[-]
Φ_{shi}	Shi convex polynomial stress function (gradient computation)	[-]

In the quoted Theorem 1, $\lambda \in [0, 1]$ denotes a convex-combination weight, unrelated to the prices and costs above. The operator $\langle \cdot, \cdot \rangle$ is the inner product, and $|\bar{s}_k|, |\bar{\mu}_k|$ denote the mean absolute values of the vectors s_k and μ_k . The set $\mathcal{K}_{\text{repl}}$ is the set of project years in which battery replacement occurs.

Subscripts and Indices

Index	Meaning
t	Time step (hour); t' is a second time index in the same sum
i, j	Rainflow half-cycle index (charging and discharging, respectively)
k	Project year; also the annuity and quadratic-surrogate summation index

Introduction

1.1. Context and Motivation

Hybrid Power Plants (HPPs), which combine renewable energy sources with Battery Energy Storage Systems (BESS), are increasingly used to improve flexibility, support market participation and reduce grid congestion [49]. Recent road-map studies indicate that HPPs will play a growing role in future power systems where the majority of energy will be generated from wind and solar sources [11]. The storage component allows energy to be shifted over time, which improves the dispatchability of an otherwise variable wind and solar resource [25].

In many studies, BESS are modeled as ideal and fully flexible assets. In practice, however, the performance of the battery degrades over time due to electrochemical aging [21]. This battery degradation depends strongly on operational decisions such as cycle depth, cycling frequency and charge-discharge rates. Aggressive operation strategies can increase short-term revenues but accelerate degradation, leading to reduced battery lifetime and higher long-term costs [24].

Despite the importance of degradation, many system-level optimization frameworks either neglect it entirely [13] or represent it using simplified power-throughput-based cost models. These approaches are computationally convenient but fail to capture the coupling between operational strategy, degradation behavior and optimal system design. As a result, they provide limited insight into how operational decisions can affect battery lifetime.

Cycle-based degradation models based on rainflow counting address this limitation. They link degradation directly to the battery state-of-charge trajectory, so aging can be evaluated under realistic and irregular operating conditions [63]. They are also harder to place inside an optimization, because the rainflow operator is non-smooth and the degradation it returns depends on the whole trajectory rather than on any single time step [8]. How tightly a framework couples degradation to the operating and sizing decisions is therefore a modeling choice and a computational one at the same time. This thesis investigates how accounting for battery degradation affects the operational strategies and the design decisions of a wind-storage hybrid power plant.

1.2. Hybrid Power Plants and Storage Dispatch

This section introduces the two concepts that frame the rest of the thesis: the hybrid power plant as a system and the dispatch of its storage component as an optimization problem.

1.2.1. Hybrid Power Plants

A hybrid power plant combines two or more generation or storage technologies behind a single grid connection point and operates them as one unit under a single operator [11]. The most common configurations pair a variable renewable source, such as wind or solar, with an energy storage system. Co-location lets the technologies share the cost of the grid connection and its associated electrical infrastructure. It also lets them use the connection capacity more fully than a single technology could [11],

49].

For wind energy, the limitation addressed by a hybrid configuration is the mismatch between generation and demand. Wind output follows the resource rather than the electricity price or the system load. Adding storage allows energy to be shifted from periods of high wind availability to periods of high demand or high price, which improves the dispatchability of the plant [30, 3]. A specific case is base-load operation, in which the plant is required to deliver at least a constant power level at all times. This target can be met with wind and storage, but the required storage capacity and its cost are sensitive to the design assumptions [30].

The single connection point introduces a design coupling that does not exist for standalone plants. The generation capacity is often over-sized relative to the grid connection, so that the storage can absorb the surplus and the connection capacity is used a larger share of the time [44, 11]. The sizes of the generation and storage components, the grid connection capacity and the operating strategy are therefore interdependent. Selecting them is a design optimization problem rather than a sequence of independent choices [44, 17].

1.2.2. Storage System Dispatch

Dispatch is the decision of when and how much to charge or discharge the storage system over time. In a market setting, this decision is driven by the electricity price and the available renewable generation. Dispatch is commonly formulated as an optimization problem. The decision variables are the charge and discharge power at each time step. The state of charge evolves through an energy balance with fixed charge and discharge efficiencies, and the power and energy ratings impose linear bounds [43, 30]. The objective is usually to maximize revenue or net present value, or to minimize the cost of meeting a demand target [17, 54]. When the storage model is linear, the problem is a linear program; when binary variables are added, for example to prevent simultaneous charging and discharging, it becomes a mixed-integer linear program [43].

The simplest value stream is energy arbitrage, in which the storage charges when the price is low and discharges when the price is high [40]. The revenue from arbitrage depends on the spread between high and low prices, so it varies with the market and with the year. It can be small for two distinct reasons. In wind-storage plants operated for base-load, arbitrage contributes only a small share of the business case, because most of the value comes from meeting the dispatch target [30]. In day-ahead markets with narrow price spreads, the arbitrage revenue is low in absolute terms, whether or not another value stream competes with it [40]. Where neither condition holds, arbitrage sets both the revenue and the cycling of the battery, and therefore the rate at which it ages.

Dispatch strategies differ in their objective. Some maximize revenue, some maximize the share of renewable energy delivered and some target a reliability level for the delivered power [3]. The dispatch also depends on forecasts of wind and price. Robust formulations optimize against a conservative forecast, improving reliability at a small cost in revenue [31]. Formulations that assume perfect foresight sit at the optimistic end of that range. Strategies differ further in how the design itself is searched, from full-year simulation with heuristic search [5] to convex formulations applied to storage sizing for offshore wind [37].

The dispatch strategy sets the sequence of charge and discharge cycles that the battery experiences and therefore the rate at which it ages. More accurate degradation modeling has been shown to change the operating decision and improve the long-term outcome [52]. How long the battery lasts under a given dispatch is therefore a plant-level question with consequences beyond just the plant.

1.3. System Integration and Socio-Economic Context

As the share of wind and solar generation rises, the difficulty shifts from producing renewable energy to matching it to demand, and a hybrid power plant with storage addresses that match rather than generation alone. By delivering dispatchable power, providing grid services and using the grid connection more fully, it reduces the need for additional network capacity and for backup from conventional plants [11]. The over-sizing of generation behind a single connection is itself a form of integration efficiency, since it delivers more renewable energy through a given amount of grid infrastructure [44]. Producing base-load power from wind and storage is one example of this role, trading a higher storage

cost for a firm and predictable supply [30]. The value of such a plant is therefore partly a system value that a plant-level revenue figure does not capture.

The economic case for storage depends on its capital cost, which has fallen substantially and is projected to continue declining along an experience curve as cumulative capacity grows [56]. Cell and pack costs fall fastest, at learning rates of 15% and 23% respectively, while the balance of system, the grid connection and the engineering and construction costs decline slowly or not at all [10]. A replacement therefore becomes cheaper over time relative to the initial installation, but it remains a lump sum committed part-way through the project life, and it is the timing of that commitment rather than its unit cost that a degradation-aware design can influence. Accounting for degradation is therefore not only a matter of one plant's profitability, but of whether storage can be deployed cost-effectively at the scale that a renewable-dominated system requires. That cost is ultimately carried by electricity consumers through wholesale prices and network tariffs, so cost-effective sizing is also a question of affordability.

Battery replacement is a material decision as well as a financial one. Each replacement requires a new cell inventory, with the associated extraction of lithium, manganese and graphite and the emissions embodied in cell manufacturing. Supply of these materials is geographically concentrated, and refining capacity for several of them sits with a single dominant supplier [28]. Demand for them therefore carries supply-security concerns beyond price. Any design or operating choice that reduces the installed energy capacity, or that extends the time to replacement, reduces material demand as well as lifetime cost. The case for accounting for degradation runs in the same direction for materials as it does for economics.

The value a storage system can earn from arbitrage also changes as renewable penetration rises. Increasing wind and solar generation lowers wholesale electricity prices and narrows the spread between high and low prices, an effect known as the merit-order effect [57]. A narrower spread lowers the revenue available from arbitrage. System-level studies indicate that the dominant value of storage comes instead from deferring investment in generation and transmission capacity, with arbitrage and reserve provision secondary, and that this deferral value increases with the share of wind and solar in the system [38]. Both value streams depend on the battery remaining available across the project lifetime, and therefore on how quickly it ages.

Degradation therefore sits at the intersection of a plant-level decision and a system-level cost. The degradation models and optimization approaches available to represent it are reviewed in the following section.

1.4. Literature Overview and Research Gaps

Battery degradation models used in energy system studies are commonly classified into two main categories: power-based and cycle-based models [21].

Power-based models approximate degradation as a function of charging and discharging power or total energy throughput. Because of their simple mathematical structure, they are widely used in system-level optimization problems [13]. However, these models rely on aggregated stress measures and do not capture the nonlinear relationship between cycling behavior, depth of discharge and battery aging [14]. Shi et al. [59] report that an NMC battery experiences roughly ten times more degradation when cycled at a depth of 100% than at a depth of 10%, for the same energy throughput. A power-based model treats both cases identically.

Cycle-based degradation models explicitly relate battery aging to charge-discharge cycles and their characteristics. By accounting for characteristics such as the cycle depth, frequency and amplitude, these models provide a more physically meaningful representation of degradation mechanisms in lithium-ion batteries. As a result, cycle-based models are commonly used in battery lifetime assessment studies [55]. The improved accuracy comes at a computational cost: cycle identification requires processing the full state of charge trajectory, making these models harder to embed directly in optimization solvers. Reniers et al. [52] quantify what that accuracy is worth, comparing three battery models of increasing fidelity against experimental degradation data. The root-mean-square capacity prediction error decreases from 11% to 5% when moving from the simplest to the most detailed model. Used inside an optimal control algorithm for energy arbitrage, the most detailed model gives a best-case

profit improvement of 175% over the simplest, mostly through reduced degradation rather than higher revenue. Their fidelity ladder is electrochemical rather than cycle-based, but the choice of degradation model changes the operating decision and not only the lifetime estimate.

In practical applications, battery operation rarely follows regular or repeated cycling patterns. Instead state of charge trajectories are irregular and depend on market signals, renewable generation and operational constraints. Rainflow counting, originally developed for fatigue analysis in mechanical systems, provides a systematic method to extract equivalent charge-discharge cycles from such irregular trajectories. By expressing degradation in terms of identified cycles rather than raw time-series data, rainflow counting enables cycle-based aging models to be applied under realistic operating conditions [63].

The degree to which degradation is coupled to operational and design decisions has direct consequences for the quality of the solution. Operation affects degradation, degradation determines replacement timing and cost, and replacement cost feeds back into the optimal storage sizing. A framework that treats these links loosely may find an operational strategy that appears profitable in the short term but leads to premature replacement, or a storage size that is optimal only under the assumption of degradation-free operation [30, 52]. The strength of the coupling therefore determines whether the framework can identify the design that maximizes lifetime value, or only the dispatch that maximizes revenue at a fixed design.

Based on the reviewed literature, three distinct approaches can be identified for integrating rainflow-based degradation models into energy system optimization. These approaches differ primarily in how strongly degradation is coupled to operational decision-making. Figure 1.1 shows the three on a common structure, so that the only difference between them is the path taken by the degradation result.

- *Post-processing approach.* In this approach, degradation is evaluated after the operational or sizing optimization has been completed. Battery operation is optimized without explicitly accounting for aging and rainflow-based cycle counting is applied afterwards to the resulting state-of-charge trajectory. While this method provides insight into expected battery lifetime under different operating profiles, degradation does not influence operational decisions. As a result, this approach is mainly used for techno-economic evaluation rather than decision making [63, 23].
- *Iterative-sequential approach.* In this approach, operational optimization and degradation evaluation are coupled through an iterative feedback loop. After an initial optimization, rainflow-based degradation metrics, such as state of health or expected lifetime, are computed from the resulting state-of-charge trajectory and fed back into the optimization as updated costs, constraints, or sizing parameters. This process is repeated until convergence. Compared with post-processing methods, this approach allows degradation estimates to influence subsequent optimization runs. However, degradation is still treated as an external update rather than as a decision-dependent term within the optimization problem. As a result, the final dispatch may be suboptimal for the fully coupled operation-degradation problem. This limitation is common to iterative-sequential architectures used in storage market bidding and hybrid plant co-optimization studies, where degradation estimates are updated between successive optimization runs [54, 18].
- *Embedded approach.* In this approach, the rainflow-based degradation model is embedded directly in the optimization problem, so that degradation influences the operating decision within a single solve rather than through an external update. The main obstacle is that the rainflow operator has no closed form and is non-smooth, which prevents its direct use in gradient-based or linear solvers. Shi et al. [59] address this by proving that the rainflow-based degradation cost is convex in the state-of-charge profile when the cycle-depth stress function is convex and by providing a sub-gradient method to solve the resulting problem. The journal version of this work adds an online control policy whose distance from the full-information optimum is bounded [60]. Lee and Kim [36] reach a comparable result by a different route, introducing an auxiliary state-of-charge variable that tracks discharge only, which allows the per-cycle degradation cost to be written in a form compatible with deterministic scheduling. These formulations make cycle-based degradation usable inside an optimization, but they have been applied mainly to operation or planning for a fixed storage size and rarely to a framework that links the operating decision to the system design.




Figure 1.1: Three approaches to coupling a rainflow-based degradation model to a dispatch optimization. The five blocks and their positions are identical in each panel. Under post-processing (a) the degradation result never returns to the optimization; under the iterative-sequential approach (b) it returns between solves; under the embedded approach (c) it is evaluated at every solver iteration rather than between solves. The inputs are the wind and price time series and the battery energy and power capacity, and the outputs are the capacity fade, the end of life and the net present value.

Two limitations run across all three approaches. The first concerns the fidelity of the degradation term. In many studies, degradation is reduced to a single aggregated quantity in the objective function or constraints. Murcia Leon et al. [44] represent degradation cost as proportional to total energy throughput, assigning the same marginal cost to every megawatt-hour cycled regardless of cycle depth. He et al. [23] include a per-cycle penalty derived from an assumed cycle-life curve and total regulation throughput, without distinguishing between deep and shallow cycles in the actual state-of-charge trajectory. González-Garrido et al. [18] constrain the maximum number of equivalent full cycles per year, so that the optimizer receives no marginal cost signal for individual cycles. In all three cases, the formulation eliminates the dependence of degradation cost on cycle-specific characteristics such as cycle amplitude and mean state of charge. The optimizer therefore has no gradient information connecting individual dispatch decisions to capacity fade. Two things therefore remain unknown: which characteristics of the operation drive capacity fade, and which of them can be used to reduce it. This holds for the dispatch trajectory itself and for the parameters that bound it, such as the state-of-charge operating window.

The second limitation concerns the structure of the coupling rather than the fidelity of the model. Approaches that rely on post-processing or iterative feedback typically depend on exhaustive parameter sweeps or repeated re-optimization. Saez-de-Ibarra et al. [54], for instance, co-optimize storage sizing and control strategy by re-running the full dispatch problem for each candidate size, while González-Garrido et al. [18] update degradation constraints between successive bidding optimization runs without passing sensitivity information to the sizing level. Martínez-Rico et al. [41] follow the same simulation-based logic for a hybrid renewable plant, simulating each candidate battery size to end of life with cycling and calendar aging and reporting that energy-to-power ratios above three to four hours are less cost-effective under the studied market scenarios. These simulation-based methods give reliable sizing results and are well suited to studying how degradation affects the design, which is why they are widely used. Their main cost is computational. The dispatch problem is solved once for every candi-

date size, so the effort grows with the number of design variables. A second cost is structural. The degradation sensitivities computed inside the operational problem are not passed to the sizing level. Dykes et al. [13] identify this weak structuring between operation and design optimization as an open challenge for hybrid power plant frameworks.

Together, these two limitations mark what a degradation-aware framework for a wind-storage plant has yet to establish. The first leaves open how operational choices govern degradation and whether accounting for it changes the storage size that is selected. The second leaves open whether that link can be optimized over rather than enumerated around. The following section develops the first into the research question and its sub-questions, and the second into a methodological objective.

1.5. Research Question and Objectives

The two limitations share a single root. Degradation is caused by the operating decision, and it is a cost the design must carry. Existing frameworks represent it too coarsely to trace the first link, or too loosely to act on the second. This motivates the main research question of this thesis:

To what extent does explicitly accounting for battery degradation influence optimal operational strategies and design decisions in hybrid power plants?

The question has an operational side and a design side, both following from the first limitation. The operational side comes first: the effect of degradation cannot be isolated without knowing which features of operation drive it. This gives two sub-questions, one for the dispatch trajectory and one for the state-of-charge window that bounds it. The design side gives the third, which compares the storage size selected with and without a degradation cost.

1. Which characteristics of battery operation, such as cycle amplitude and mean state of charge, act as the effective levers for reducing degradation?
2. How does the choice of state-of-charge operating window trade revenue against battery degradation and lifetime?
3. How does accounting for degradation change the optimal storage sizing, compared with a design obtained under degradation-free operation?

The first two sub-questions are answered with a cycle-based degradation model. The model is evaluated on the full dispatch trajectory and reports degradation per identified cycle, not per megawatt-hour. The third is answered with a parameter sweep that re-solves the dispatch at each candidate size, structure (b) of Figure 1.1.

This sweep is reliable, but it does not remove the structural limitation. It enumerates rather than searches, and it does not pass degradation sensitivities to the design level. The embedded approach would remove that limitation, by placing the degradation cost inside the dispatch. The reviewed formulations show this is possible, but only at a fixed storage size and with a purpose-built solver. This thesis formulates the same problem for a general-purpose gradient-based solver and reports whether it can be solved that way.

Research Objectives

The first objective builds and verifies the model on which the rest depend. The next three answer the sub-questions in turn. The fifth addresses the second limitation.

- Implement a rainflow-based cycle-counting degradation model inside a dispatch optimization framework for a wind-battery hybrid power plant, and verify the implementation against analytical checks.
- Identify the operational levers for reducing degradation (sub-question 1).
- Quantify the revenue and degradation trade across state-of-charge windows (sub-question 2).
- Quantify the change in optimal storage sizing against a degradation-free design (sub-question 3).
- Formulate the embedded degradation-aware dispatch problem and establish whether a general-purpose gradient-based solver can solve it.

2

Battery Degradation in Lithium-Ion Systems

This chapter establishes the theoretical foundations for the rest of the thesis. It has two parts. The first (Sections 2.1–2.4) is background on battery degradation for readers without prior knowledge in electrochemical systems: how a lithium-ion cell ages, how that aging is quantified, and how individual cycles are extracted from an irregular operating profile. The second (Sections 2.5–2.8) introduces the two degradation models this thesis uses and derives the mathematical results they rest on: the physical reference model used for reporting, the polynomial model used for the gradient, and the sub-gradient of the degradation cost. The methodology of Chapter 3 applies these tools; the present chapter defines them.

2.1. Operating Principle of Lithium-Ion Batteries

A lithium-ion battery stores and releases energy through the reversible intercalation of lithium ions between a positive electrode (cathode) and a negative electrode (anode), separated by an electrolyte. During discharge, lithium ions migrate from the anode to the cathode while electrons travel through the external circuit and perform work. Charging reverses this process. The state of charge (SoC) is the fraction of available capacity currently stored in the battery, ranging from fully depleted ($\text{SoC} = 0$) to fully charged ($\text{SoC} = 1$).

The degradation model used in this thesis is parameterized for Lithium Manganese Oxide (LMO), a lithium-ion chemistry with a manganese-oxide spinel cathode and a graphite anode. LMO is not the dominant chemistry in modern grid-scale storage; LFP and NMC are more common. It is used here because Xu et al. [63] provide a complete, experimentally validated semi-empirical degradation parameter set for LMO cells, which allows the framework to be implemented without introducing additional uncertainty from parameter fitting.

The use of LMO therefore affects the numerical degradation rates and the projected end-of-life values, but not the structure of the methodology. Controlled cycling of commercial LFP, NCA and NMC cells under matched conditions confirms that capacity-fade rates differ substantially across chemistries [51].

2.2. Degradation Mechanisms

With the operating principle and the state of charge established, the question is how a cell loses capacity over time.

Battery degradation is the irreversible loss of electrochemical performance over time. It appears primarily in two forms. The first is capacity fade: the maximum amount of energy the battery can store decreases, so a cell rated at 300 MWh when new may only be able to store 240 MWh after several years of use. The second is power fade: the battery's internal resistance increases, which means a larger fraction of energy is lost as heat during each charge and discharge cycle, reducing both the

power the battery can deliver and the round-trip efficiency [14].

This thesis models capacity fade only. For arbitrage duty it is the binding constraint, since it limits the energy available to shift, and it is the quantity on which end of life is defined here. Power fade is not represented, so the round-trip efficiency is held constant across the project lifetime. Two physically distinct mechanisms drive capacity fade simultaneously: calendar aging, which proceeds whether or not the battery is used, and cycle aging, which is caused by charging and discharging. Only the second is directly set by the dispatch, and both are affected by the state of charge at which the battery operates.

2.2.1. Calendar Aging

Calendar aging occurs continuously, even when the battery is idle. The dominant mechanism is growth of the Solid Electrolyte Interphase (SEI) layer on the graphite anode. The SEI is a thin film of reduction products that forms when the electrolyte reacts with the anode surface. An initial SEI layer is necessary to stabilize the electrode, but continued film growth is parasitic: it consumes cyclable lithium, which reduces the amount of charge that can be transferred between the electrodes and therefore reduces the usable capacity of the cell. The growing film also increases the resistance to ion transport at the electrode surface, contributing to power fade [21, 14].

SEI growth is limited by how fast reactant species can diffuse through the existing film to reach the electrode surface [14]. As the layer thickens, this diffusion path lengthens and the growth rate slows. The resulting capacity loss follows a characteristic pattern: degradation accumulates rapidly in a new cell and then decelerates over time. Quantitatively, capacity loss is commonly modeled as proportional to the square root of storage time, reflecting the diffusion-limited nature of the process [48]. Growth rate is sensitive to temperature and to the SoC at which the battery rests. Elevated temperature and high resting SoC both accelerate SEI formation [14]. Calendar aging therefore represents a baseline level of degradation that accumulates regardless of how the battery is operated.

2.2.2. Cycle Aging

Cycle aging is caused by the mechanical and electrochemical stresses of charge-discharge events. Each time lithium ions intercalate into and de-intercalate from the electrode lattice, the host material expands and contracts. Repeated volumetric changes cause microscopic fatigue: electrode particles crack, active material becomes electrically isolated and contact resistance increases [14]. In LMO cathodes, manganese ion dissolution into the electrolyte provides an additional degradation pathway, activated at elevated temperature and high SoC.

The primary operational driver of cycle aging is the cycle amplitude δ , defined as the fraction of total capacity cycled in a single charge-discharge event. This quantity is the depth of an individual cycle, as identified by the rainflow counting algorithm. It should not be confused with the depth-of-discharge parameter d , which sets the SoC operating window (Section 3.4.4). Deeper cycles impose larger volumetric strains and dissipate more energy within the cell, accelerating fatigue at a rate that is nonlinear in the depth of discharge. A full discharge from SoC = 1 to SoC = 0 causes more degradation than two half-depth cycles covering the same net energy throughput. This nonlinearity is the reason that simple energy-throughput models cannot accurately predict degradation under realistic operating conditions.

2.2.3. Relative Contribution

Total degradation is commonly modeled with separate calendar and cycle components that are summed to give the overall capacity loss. This additive decomposition is a modeling choice, not a physical law. It is justified when the two damage pathways operate through largely independent mechanisms, as in LMO cells where SEI growth and particle cracking affect different parts of the cell [63]. In chemistries where the mechanisms interact more strongly (for example, where cycling accelerates SEI reformation), the additive assumption becomes less accurate. The relative contribution of each component depends on operating conditions. Frequent deep cycling accumulates cycle aging rapidly. Extended rest at high SoC accumulates calendar aging at a comparable rate. In storage systems that are dispatched against electricity prices, the battery typically cycles daily during periods of high price volatility and rests at intermediate SoC during low-activity periods. The balance between the two components therefore depends on the market regime and dispatch strategy, not only on the battery chemistry. Field measurements from a deployed utility-scale system confirm this [19]. Operating strategy and cell temperature are the

primary drivers of capacity loss, and day-ahead and intraday market operation age the battery faster than frequency regulation. The share carried by each pathway is therefore a property of the operating regime, and it has to be measured for a given site and market rather than assumed. Where calendar aging carries a large share, operational measures directed at cycle amplitude alone address only part of the total.

2.3. Quantifying Degradation: State of Health

These mechanisms explain why a cell ages. To manage that aging, sizing and dispatch decisions need a single number that tracks the capacity lost so far. That number is the State of Health (SoH), defined as the ratio of current usable capacity to the original rated capacity,

$$\text{SoH}(t) = \frac{C(t)}{C_0}, \quad (2.1)$$

where $C(t)$ is the usable capacity at time t and C_0 is the nominal capacity of a new cell. A new battery has $\text{SoH} = 1$. SoH declines monotonically as degradation accumulates.

End-of-life (EoL) is the SoH value below which the battery is no longer fit for its intended purpose. It is a convention set by the application rather than a property of the cell, and both the criterion and its level vary between applications [21]. This thesis adopts a state of health of 70% as the replacement criterion and reports the 80% crossing for the baseline case alongside it; Section 2.6.3 sets out the choice.

State of health is the conventional metric in experimental studies. The model of Section 2.6 does not compute it directly: it accumulates damage in a separate variable whose increments add, and recovers the state of health from that variable through a nonlinear map.

2.4. Cycle Counting Under Irregular Operation

Cycle aging depends on the amplitude of each individual charge-discharge cycle. This assumes the individual cycles can be identified. Under realistic operation they cannot. Battery dispatch is driven by electricity prices and renewable generation, so the SoC trajectory is irregular: cycle depth, duration and mean level vary continuously, and no two days are alike. Standard cycle-counting methods assume regular, fixed-amplitude cycles, so they cannot be applied here without losing information.

Rainflow counting addresses this. The algorithm was developed for fatigue analysis in mechanical engineering, where structural components experience variable-amplitude stress reversals [12]. It was later reformulated with an explicit analytical definition of the cycle amplitude [53] and standardized as a counting procedure for fatigue analysis [1]. The analogy to battery aging is direct. An electrode particle accumulates electrochemical fatigue under irregular charge-discharge cycles, just as a structural component accumulates damage under variable mechanical loading. In both cases, the damage depends on the amplitude and position of individual cycles, not the total energy dissipated. Several studies have adopted rainflow counting for battery degradation modeling. Xu et al. [63] use the algorithm to extract cycles from irregular SoC trajectories and evaluate capacity fade under realistic operating profiles. Shi et al. [59] build on this by formulating a convex degradation cost function over rainflow-identified cycles. He et al. [23] apply cycle counting to evaluate battery aging in market bidding applications. The algorithm follows the standard variable-amplitude fatigue cycle counting procedure defined in ASTM E1049 [2] and is used here without modification.

The algorithm identifies all charge and discharge half-cycles embedded within an irregular SoC trajectory, regardless of how they overlap or nest. Each half-cycle is bounded by two turning points: the local extrema where the battery switches between charging and discharging. A full charge-discharge cycle is a pair of half-cycles, one running from a lower turning point to an upper one and one running back. Each half-cycle is characterized by two quantities that appear throughout the degradation model, defined in the two subsections below. Figure 2.1 shows a synthetic trajectory whose turning points bound four half-cycles. One charge half-cycle and the discharge half-cycle that follows it are marked, and a third is annotated with its cycle amplitude and mean state of charge.




Figure 2.1: Rainflow half-cycle identification on a synthetic SoC trajectory. Red circles mark lower turning points, green circles upper turning points. The solid bracket spans a charge half-cycle, the dashed bracket the discharge half-cycle that follows it. Dotted lines link the third half-cycle to its cycle amplitude δ_i and mean state of charge σ_i .

Cycle amplitude δ .

The cycle amplitude is the absolute SoC change between the turning points that bound the half-cycle,

$$\delta_i = |\text{SoC}_{\max,i} - \text{SoC}_{\min,i}|, \quad (2.2)$$

where $\text{SoC}_{\max,i}$ and $\text{SoC}_{\min,i}$ are the upper and lower SoC turning points of half-cycle i . A full discharge from $\text{SoC} = 1$ to $\text{SoC} = 0$ gives $\delta = 1$; a cycle between $\text{SoC} = 0.4$ and $\text{SoC} = 0.7$ gives $\delta = 0.3$.

Terminology. In the battery literature, the symbol δ and the term depth of discharge (DoD) are sometimes used interchangeably to refer to the amplitude of a single cycle. This thesis uses distinct terms to avoid confusion: *cycle amplitude* δ_i denotes the actual SoC range of a specific rainflow half-cycle i , while *depth of discharge* d denotes the SHIPP operating-constraint parameter that sets the lower state-of-charge bound ($d = 1 - \sigma_{\min}$, Section 3.4.4). The cycle amplitude varies from cycle to cycle and is bounded above by the width of the operating window, $\delta_i \leq \sigma_{\max} - \sigma_{\min}$, which equals d only when $\sigma_{\max} = 1$.

Mean state of charge σ .

The mean SoC is the midpoint of the SoC range traversed by the half-cycle,

$$\sigma_i = \frac{\text{SoC}_{\max,i} + \text{SoC}_{\min,i}}{2}. \quad (2.3)$$

Two cycles with the same depth but different mean SoC levels cause different degradation. Higher mean SoC accelerates the electrochemical side reactions that drive aging. This is why σ appears as a separate stress factor in degradation models that account for mean SoC, as described in Section 2.6.

2.5. Degradation Modeling Approaches

The preceding sections established the two aging mechanisms, the SoH metric, and the per-cycle parameters δ and σ . A degradation model is what combines these into a prediction of capacity loss. This section surveys the available classes, selects the semi-empirical class, and names the two formulations this thesis uses.

Degradation models range from high-fidelity electrochemical formulations to compact empirical expressions. The choice involves a trade-off between physical fidelity and computational tractability.

Electrochemical models resolve lithium concentration, electric potential and reaction rates within the cell using coupled partial differential equations. The Pseudo-Two-Dimensional (P2D) model, for example,

solves for the solid-phase lithium concentration $c_s(r, t)$ at each electrode particle [14]:

$$\frac{\partial c_s}{\partial t} = \frac{D_s}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial c_s}{\partial r} \right), \quad (2.4)$$

where D_s is the solid-phase diffusion coefficient and r is the radial coordinate within the particle. Degradation is predicted from the evolution of side reaction rates and film growth at the particle surface.

P2D models can achieve high physical fidelity but require detailed cell-level parameters (electrode thickness, porosity, exchange current densities) that are rarely available for commercial cells. A single year of hourly simulation at the P2D level is computationally prohibitive for the 8,760-step dispatch problems considered in this thesis.

Semi-empirical models represent the dominant aging mechanisms through algebraic expressions calibrated to experimental cycling and storage data. They express the per-cycle degradation as a product of stress factors, each capturing one physical driver. The general structure common to most semi-empirical formulations is

$$f_{d,i} = [S_\delta(\delta_i) + S_\tau(\tau_{c,i})] S_\sigma(\sigma_i) S_T(T_{c,i}), \quad (2.5)$$

where S_δ is the cycle amplitude stress factor, S_σ is the mean SoC stress factor, S_T is the temperature stress factor and S_τ is a time stress term.

The calendar term S_τ can be applied inside each cycle, as written above, or accumulated once over the operating period as a single term. This thesis uses the second form, in which the per-cycle term carries only S_δ and calendar aging is a separate bulk term (Section 2.6.1).

Data-driven models fit degradation trajectories directly from measured capacity data without imposing physical structure. Machine learning approaches, such as the early-life prediction model of Severson et al. [58], can achieve high accuracy within their training domain. However, they are specific to the chemistry, cell format and cycling conditions of their training data, with limited ability to extrapolate to operating regimes not represented in the dataset. This makes them less suitable for the variable dispatch profiles produced by market-driven optimization.

For system-level optimization of hybrid power plants over multi-year horizons with irregular cycling, the semi-empirical class is the suitable one: it represents the dominant aging mechanisms at a cost compatible with multi-year hourly simulation. Xu et al. [63] and He et al. [23] use semi-empirical models within dispatch optimization frameworks for grid-scale storage, while Shi et al. [59] build a convex formulation on top of the same class.

This thesis uses two models from that class. The physical reference model of Xu et al. [63] is used for degradation reporting. It specifies empirical functional forms for each stress factor, calibrated against LMO cycling data. Its equations, calibration accuracy and parameter values are given in Section 2.6.

The polynomial model of Shi et al. [59] replaces the cycle amplitude stress function with a power-law form whose convexity is enforced during fitting, and proves that the resulting rainflow degradation cost is convex in the SoC profile. That convexity admits an analytic sub-gradient. Shi et al. do not propose a new aging model; they provide the mathematical framework that makes cycle-based degradation compatible with gradient-based optimization. The stress function and the convexity guarantee are given in Section 2.7, the sub-gradient in Section 2.8.

2.6. The Physical Reference Model (Xu 2016)

This model is presented first because it is the reference against which all degradation in this study is reported. It separates total capacity loss into a calendar aging component and a cycle aging component accumulated through the rainflow counting of Section 2.4. It is parameterized for LMO cells and validated against Dynamic Stress Test (DST) cycling data with a capacity estimation error of $\pm 3\%$ at the 80% EoL threshold [63].

2.6.1. Model Structure

The total linearized degradation fraction f_d is the sum of a calendar aging term and a cycle aging term,

$$f_d = f_d^{\text{cal}} + f_d^{\text{cyc}}, \quad f_d^{\text{cyc}} = \sum_i n_i f_{d,i}, \quad (2.6)$$

where $f_{d,i}$ is the linearized life fraction lost in cycle i (Equation 2.9), n_i is the cycle multiplicity returned by the rainflow algorithm, equal to 0.5 for a half-cycle and 1 for a full cycle, and f_d^{cal} is the linearized calendar life loss over the simulation period (Equation 2.15). The two contributions are accumulated and reported separately throughout.

The quantity f_d is a linearized accumulation variable, not a capacity loss. The physical capacity loss L follows the SEI two-exponential fade curve (Equation (12) in [63]),

$$L(f_d) = 1 - \alpha_{\text{sei}} e^{-\beta_{\text{sei}} f_d} - (1 - \alpha_{\text{sei}}) e^{-f_d}, \quad (2.7)$$

where α_{sei} and β_{sei} are empirical coefficients fitted to LMO degradation test data. Here L is the fraction of capacity lost, with $L = 0$ for a new cell.

The two-exponential structure captures the fast initial capacity drop from SEI film formation followed by a slower linear degradation phase. The State of Health is the capacity retained, equal to the complement of the loss,

$$\text{SoH}(f_d) = 1 - L(f_d) = \alpha_{\text{sei}} e^{-\beta_{\text{sei}} f_d} + (1 - \alpha_{\text{sei}}) e^{-f_d}, \quad (2.8)$$

with $\text{SoH} = 1$ for a new cell. End-of-life is reached when SoH falls to the replacement threshold of Section 2.6.3.

2.6.2. Stress Factor Models

Equation 2.6 left the per-cycle loss $f_{d,i}$ that builds f_d^{cyc} unspecified. The physical reference model gives it as the product of three operational stress factors (Equation (18) in [63]),

$$f_{d,i} = S_\delta(\delta_i) S_\sigma(\sigma_i) S_T(T_{c,i}), \quad (2.9)$$

where S_δ is the cycle amplitude stress factor, S_σ is the mean SoC stress factor and S_T is the temperature stress factor.

Calendar aging is not included here. It is accumulated once over the operating period as the separate term f_d^{cal} (Equation 2.15), following the decomposition in Equation (3) of [63]. Xu's combined per-cycle form (Equation (20) in [63]) is used for parameter fitting from cycling tests, not for operational accumulation, which would otherwise count calendar aging twice.

Figure 2.2 shows the two factors over the ranges the operating window permits. Their sensitivities differ by an order of magnitude. Over the amplitude range $0.02 \leq \delta \leq 0.80$, S_δ rises from 0.11×10^{-5} to 2.98×10^{-5} , a factor of 26. Over the mean SoC range $0.10 \leq \sigma \leq 0.90$, S_σ rises from 0.66 to 1.52, a factor of 2.3. Cycle amplitude is therefore the stronger of the two stress factors in the physical reference model.




Figure 2.2: Cycle amplitude stress factor $S_\delta(\delta)$ and mean SoC stress factor $S_\sigma(\sigma)$ of the physical reference model, over the amplitude and mean SoC ranges permitted by a 10–90% operating window.

Cycle amplitude stress model.

Xu et al. propose an empirical expression for LMO cells,

$$S_\delta(\delta) = (k_{\delta 1} \delta^{k_{\delta 2}} + k_{\delta 3})^{-1}, \quad (2.10)$$

where $k_{\delta 1}$, $k_{\delta 2}$ and $k_{\delta 3}$ are empirical coefficients fitted to LMO cycle life data. Xu et al. report that this form attains an R^2 of 0.993 on their cycle-life dataset, a closer fit than the exponential and power-law alternatives they also tested [63]. The fitted values are listed in Table 2.1.

Table 2.1: Cycle amplitude stress function coefficients from Xu et al. [63], repeated with the full parameter set in Table 2.2.

Symbol	Value	Description
$k_{\delta 1}$	1.40×10^5	Scale coefficient
$k_{\delta 2}$	-5.01×10^{-1}	Exponent
$k_{\delta 3}$	-1.23×10^5	Offset

This form is not convex over the full amplitude range. Differentiating Equation 2.10 twice and setting $S''_\delta(\delta) = 0$ gives a single interior inflection point,

$$\delta_c = \left[\frac{k_{\delta 3} (k_{\delta 2} - 1)}{k_{\delta 1} (k_{\delta 2} + 1)} \right]^{1/k_{\delta 2}}, \quad (2.11)$$

which for the coefficients of Table 2.1 evaluates to $\delta_c = 0.1437$. Below this amplitude S_δ is concave and above it convex, so a degradation cost built on S_δ falls outside the convexity result of Section 2.7.3, which requires a stress function that is convex over the whole amplitude range.

Mean SoC stress model.

The mean SoC stress model follows an exponential form,

$$S_\sigma(\sigma) = e^{k_\sigma(\sigma - \sigma_{\text{ref}})}, \quad (2.12)$$

where $\sigma_{\text{ref}} = 0.50$ is the reference SoC at which $S_\sigma = 1$. Higher mean SoC levels accelerate degradation. Figure 2.3 isolates this effect on two cycles of equal amplitude, so that S_σ is their only difference. With $k_\sigma = 1.04$, the cycle at the higher mean SoC carries a stress factor larger by $e^{1.04 \times 0.50} = 1.68$.




Figure 2.3: Isolated effect of S_σ . Two cycles of equal amplitude ($\delta = 0.30$) at mean SoC $\sigma = 0.25$ (left) and $\sigma = 0.75$ (right), each shown in color within its surrounding dispatch trajectory in grey. Both are taken from a simulated annual dispatch produced by the model of this thesis.

Temperature stress model.

The temperature stress model is derived from the Arrhenius relation [46],

$$S_T(T_c) = e^{k_T(T_c - T_{\text{ref}}) \cdot T_{\text{ref}}/T_c}, \quad (2.13)$$

where k_T is the temperature stress coefficient and T_{ref} is the reference temperature in Kelvin. Cell temperature is assumed constant at $T_c = 25^\circ\text{C}$ throughout the simulation, consistent with active thermal management in grid-scale installations. Under this assumption $S_T = 1$ for all cycles and the temperature factor drops out of Equation 2.9.

Time stress and calendar aging.

Calendar aging is accumulated once over the operating period, not per cycle. The time stress factor is linear in elapsed time (Equation (27) in [63]),

$$S_\tau(t) = k_\tau \cdot t, \quad (2.14)$$

where k_τ is the calendar aging rate coefficient and t is the elapsed time in seconds. The calendar aging contribution over the simulation period is the product of the time, mean SoC and temperature stress factors (Equation (19) in [63]),

$$f_d^{\text{cal}} = S_\tau(t) S_\sigma(\bar{\sigma}) S_T(\bar{T}_c) = k_\tau \cdot t \cdot S_\sigma(\bar{\sigma}) S_T(\bar{T}_c), \quad (2.15)$$

where $\bar{\sigma}$ is the mean SoC and $\bar{T}_c$ is the mean cell temperature over the full simulation period. Under the constant temperature assumption $S_T = 1$ and this reduces to $f_d^{\text{cal}} = k_\tau \cdot t \cdot S_\sigma(\bar{\sigma})$.

The coefficients of all four stress factors, together with those of the fade curve of Equation 2.7, are listed in Table 2.2.

Table 2.2: Parameters of the physical reference model for LMO cells, reproduced without modification from Table I of Xu et al. [63].

Parameter	Symbol	Value	Unit
SEI coefficient	α_{sei}	5.75×10^{-2}	—
SEI coefficient	β_{sei}	121	—
Cycle-amplitude scale	$k_{\delta 1}$	1.40×10^5	—
Cycle-amplitude exponent	$k_{\delta 2}$	-5.01×10^{-1}	—
Cycle-amplitude offset	$k_{\delta 3}$	-1.23×10^5	—
SoC stress coefficient	k_{σ}	1.04	—
Reference SoC	σ_{ref}	0.50	—
Temperature coefficient	k_T	6.93×10^{-2}	—
Reference temperature	T_{ref}	298.15	K
Calendar aging rate	k_{τ}	4.14×10^{-10}	s^{-1}

2.6.3. End-of-Life Definition

The replacement criterion adopted in Section 2.3 is set out here. Battery end-of-life is defined as the point at which the State of Health falls below a threshold. No single standard applies across all use cases, and the criterion itself differs between them: for high-energy cells in vehicle service it is set on capacity, at the point where 80% of the initial capacity remains, because below that the cell no longer meets the requirements of the vehicle; for high-power cells it is set on available power, at the point where power has halved [21]. The 80% level is the value manufacturers most often quote in specification sheets and the level adopted by Xu et al. [63], but it originates in a requirement that stationary service does not impose, and cells in grid applications are commonly operated beyond it [51]. Stationary operation of degraded cells is taken considerably lower, and the literature does not agree on how far: second-life installations are variously taken to end at a state of health of 60% [19] or to extend to 40% [45].

This thesis adopts a state of health of 70% as the replacement criterion. The level reflects two considerations that pull in opposite directions. Stationary practice permits operation well below 80%, while the cycle-life data underlying the degradation model, and the accuracy reported with it, both correspond to 80% capacity retention, so confidence in the simulated trajectory decreases as the threshold is lowered. A 70% threshold sits between the two. The 80% crossing is reported for the baseline case, so the effect of the choice can be checked against the calibrated range rather than assumed away.

2.7. Polynomial Degradation Model (Shi 2018)

The physical reference model captures the two aging mechanisms, but its cycle amplitude stress function S_{δ} is not convex (Section 2.6.2). Convexity gives the degradation cost properties that make it efficient to optimize, set out in Section 2.7.3. Shi et al. [59] work with a power-law stress function whose convexity can be enforced during fitting, and prove that the rainflow degradation cost built from any convex stress function is convex in the state-of-charge profile. The sub-gradient that this cost admits is derived separately in Section 2.8, because it applies to any stress function of cycle amplitude and not only to the polynomial one.

2.7.1. Polynomial Stress Function

Shi et al. [59] model the cycle amplitude stress with a power-law function, which they term the polynomial stress model,

$$\Phi_{\text{shi}}(\delta) = k_3 \delta^{k_4}, \quad (2.16)$$

where $k_3 > 0$ is a scale coefficient and k_4 is a real exponent. Shi et al. attribute this form to Koller et al. [34] rather than proposing it. Differentiating twice with respect to the cycle amplitude gives

$$\Phi''_{\text{shi}}(\delta) = k_3 k_4 (k_4 - 1) \delta^{k_4-2}. \quad (2.17)$$

Convexity is not a property of the functional form. From Equation 2.17, $\Phi''_{\text{shi}}(\delta) \geq 0$ for all $\delta > 0$ if and only if $k_4 \geq 1$, with strict convexity for $k_4 > 1$. The exponent is therefore constrained to $k_4 > 1$ during

fitting rather than assumed. The coefficients are not universal: they are fitted over the range of cycle amplitudes that the operating window permits, so they change with the window.

In this model the stress function is the whole of the per-cycle life loss. Accumulating it over the rainflow cycles of a state-of-charge profile gives the total cycle degradation $\sum_i n_i \Phi_{\text{shi}}(\delta_i)$, where δ_i is the amplitude of cycle i and n_i is its multiplicity, 0.5 for a half-cycle and 1 for a full cycle (Equation (6) in [59]). Shi et al. convert this to a cost by scaling it with the energy capacity and the replacement cost per unit energy. That cost is the object of Section 2.8.

2.7.2. Scope of the Polynomial Model

The polynomial model is a stress function of cycle amplitude and an accumulation rule, and nothing else. Three differences from the physical reference model follow, and all three matter for how the model can be used.

First, the per-cycle loss carries no mean SoC factor and no temperature factor. The quantity $\Phi_{\text{shi}}(\delta_i)$ is therefore the counterpart of the three-factor product $f_{d,i}$ of Equation 2.9, not of the amplitude factor S_δ alone. Two cycles of equal amplitude receive the same per-cycle loss in this model regardless of the state of charge at which they occur.

Second, the model has no calendar aging term. What it accumulates is cycle aging only, so it predicts no degradation for a battery that is held at rest.

Third, the model defines no mapping from accumulated degradation to capacity loss or to State of Health. The accumulated quantity enters an optimization objective directly as a cost, scaled by the energy capacity and the replacement cost per unit energy. There is no counterpart to the SEI fade curve of Equation 2.7.

Shi et al. [59] state the first of these restrictions and give their reasoning. In grid applications the cell temperature is held within a narrow band by thermal management, and the other factors they name, high C-rate and cell imbalance, are either absent or regulated by the battery management system, so cycle amplitude is the factor they retain. They do not discuss mean state of charge. The restriction is also the scope of their convexity result: Section 2.7.3 is stated for a stress function of cycle amplitude alone, and a per-cycle loss carrying a mean SoC factor would fall outside it, because the cycle mean is itself a function of the state-of-charge profile.

The consequence is that the two models of this chapter do not produce comparable quantities as published. Using them together requires supplying the mean SoC and temperature factors, supplying a calendar term, and choosing a mapping from accumulated degradation to State of Health. Those are modeling decisions rather than properties of either published model.

2.7.3. Convexity of the Rainflow Degradation Cost

Section 2.7.1 established when the stress function itself is convex. This subsection gives the result that makes that property useful: convexity of the stress function transfers to the cost accumulated over rainflow cycles. Constraining $k_4 > 1$ therefore buys convexity of the object the dispatch actually sees.

A convex dispatch cost has two properties that a non-convex cost does not. A sub-gradient method applied to it converges to the global optimum rather than to a local one, and the gap between any intermediate solution and that optimum can be bounded at each iteration. These are the reasons a convex stress function is preferred when degradation enters the dispatch as a cost.

Shi et al. [59] prove the following. Let $\mathbf{x} \in \mathbb{R}^T$ be a state-of-charge profile, and let $\text{Rainflow}(\mathbf{x})$ produce cycle amplitudes $\{\delta_i\}$ with multiplicities $\{n_i\}$. If the stress function Φ is convex, then the accumulated stress $\sum_i n_i \Phi(\delta_i)$ is convex in $\mathbf{x}$, and hence so is the degradation cost $f = \lambda_E^{\text{repl}} \bar{E} \sum_i n_i \Phi(\delta_i)$. The proof decomposes an arbitrary SoC profile into unit steps and proceeds by induction on the number of steps [59]. The converse is not claimed: convexity of Φ is sufficient, not necessary.

One point about the scope of this result matters for what follows. It establishes convexity with respect to the state-of-charge profile, that is, with respect to the dispatch. It does not establish convexity with respect to the energy capacity $\bar{E}$; the sizing problem is not convex in $\bar{E}$, and no such claim is made here or in the sizing method of Chapter 3.

The reference stress function S_δ fails the premise, for the reason given in Section 2.6.2. Using it directly would lose the guarantee; discarding it would lose the experimentally validated accuracy. The power-law function Φ_{shi} resolves this: it is fitted to S_δ over the operating range while its exponent is constrained to keep it convex. The two functions serve complementary purposes, S_δ for degradation reporting and Φ_{shi} for the gradient, as set out in the methodology (Section 3.5).

The framework itself is not specific to LMO. The rainflow procedure and the separation between cycle and calendar aging apply to any lithium-ion chemistry. Transferring the model would require refitting the stress-factor coefficients to chemistry-specific cycling and calendar-aging data, and re-verifying that the fitted amplitude stress function is convex before the polynomial branch could be used.

2.8. Sub-gradient of the Rainflow Degradation Cost

Convexity establishes that the degradation cost admits a sub-gradient but does not construct one. This section derives it, since a gradient-based sizing method requires the cost to be differentiable with respect to the charging and discharging power. The derivation below applies to any stress function of cycle amplitude. The convex power law of Section 2.7.1 is one such function, and is the one used here, for the reasons given in Section 2.7.3.

2.8.1. Degradation Cost and the Dispatch Variables

Differentiating the cost with respect to the dispatch first requires writing it in terms of the dispatch variables. Section 2.7.1 expressed it through the cycle amplitudes; the amplitudes must now be traced back to the charging and discharging power. Following Shi et al. [59], the degradation cost over a dispatch horizon is the accumulated cycle stress scaled by the energy capacity and the replacement cost per unit energy,

$$f(\mathbf{p}^{\text{ch}}, \mathbf{p}^{\text{dis}}) = \lambda_E^{\text{repl}} \bar{E} \sum_i n_i \Phi_{\text{shi}}(\delta_i), \quad (2.18)$$

where $\mathbf{p}^{\text{ch}}$ and $\mathbf{p}^{\text{dis}}$ are the vectors of charging and discharging power over the horizon, with entries p_t^{ch} and p_t^{dis} at each time step, $\bar{E}$ is the energy capacity and λ_E^{repl} is the replacement energy cost per unit energy.

The cost takes the whole vectors as arguments because the rainflow cycles depend on the entire trajectory rather than on any single time step. Its scale is set by the cost of a full replacement, $\lambda_E^{\text{repl}} \bar{E}$, of which the accumulated stress is the fraction consumed over the horizon. The sum carries the amplitude stress alone, for the reasons set out in Sections 2.7.2 and 2.8.4, and carries no calendar term, because calendar aging does not depend on the dispatch and contributes nothing to the derivatives that follow.

The link to the dispatch is through the cycle amplitude. Because the dispatch variables are powers, the derivation is carried out in stored energy e_t rather than in the normalized state of charge of Section 2.4, the two being related by $e_t = \bar{E} \text{SoC}_t$. The stored energy is the running sum of charge and discharge power,

$$e_t = e_0 + \Delta t \sum_{t' \leq t} (\eta_{\text{in}} p_{t'}^{\text{ch}} - \eta_{\text{out}}^{-1} p_{t'}^{\text{dis}}), \quad (2.19)$$

where $p_{t'}^{\text{ch}} \geq 0$ and $p_{t'}^{\text{dis}} \geq 0$ are the charging and discharging power at time step t' , $\eta_{\text{in}}, \eta_{\text{out}} \in (0, 1]$ are the one-way efficiencies in the SHIPP convention, and Δt is the time step. The amplitude of cycle i is the normalized difference between the battery energy at its two turning points a_i and b_i ,

$$\delta_i = \frac{|e_{b_i} - e_{a_i}|}{\bar{E}}, \quad a_i < b_i, \quad (2.20)$$

where a_i and b_i are the time-step indices of the turning points that open and close the cycle, as identified by the rainflow procedure of Section 2.4. Equation 2.20 is Equation 2.2 written in energy rather than state of charge; the division by $\bar{E}$ returns the normalized form.

Equations 2.18 to 2.20 form a chain. The cost depends on the cycle amplitudes, each amplitude depends on the stored energy at two turning points, and the stored energy depends on the power at every earlier time step. The next subsection differentiates along that chain, working from the last link back to the first.

2.8.2. Sub-gradient with Respect to Power

The derivative of the cost with respect to a single power variable follows from how a change in that variable moves the SoC trajectory. Raising the charging power p_t^{ch} at one time step t by a small amount raises the battery energy $e_{t'}$ at every later step $t' > t$ by the same amount, and leaves every earlier step unchanged, by Equation 2.19. Because the amplitude is a difference between the stored energy at two turning points, it changes only if the nudge moves one of them and not the other. If both lie after t , the cycle shifts bodily and its amplitude is unchanged. If both lie at or before t , the cycle is untouched. Only when $a_i \leq t < b_i$ does the later point move while the earlier one stays, so that the distance between them changes.

A cycle in this first case is said to span time step t : at t the battery is partway through it, with one turning point already reached and the other still ahead. Collect these cycles into the spanning set,

$$\mathcal{S}(t) = \{i : a_i \leq t < b_i\}. \quad (2.21)$$

For a spanning cycle, differentiating Equation 2.20 through Equation 2.19 gives

$$\frac{\partial \delta_i}{\partial p_t^{\text{ch}}} = \frac{\Delta t \eta_{\text{in}}}{\bar{E}} \text{sign}(e_{b_i} - e_{a_i}), \quad i \in \mathcal{S}(t), \quad (2.22)$$

where the sign records whether the cycle rises or falls across t : charging into a rising cycle deepens it and charging into a falling cycle makes it shallower. Substituting this into Equation 2.18 through the chain rule, the $1/\bar{E}$ cancels the $\bar{E}$ in the cost, and the derivative with respect to the charging power is

$$\frac{\partial f}{\partial p_t^{\text{ch}}} = \lambda_E^{\text{repl}} \Delta t \eta_{\text{in}} \sum_{i \in \mathcal{S}(t)} n_i \Phi'(\delta_i) \text{sign}(e_{b_i} - e_{a_i}), \quad (2.23)$$

and with respect to the discharging power,

$$\frac{\partial f}{\partial p_t^{\text{dis}}} = -\frac{\lambda_E^{\text{repl}} \Delta t}{\eta_{\text{out}}} \sum_{i \in \mathcal{S}(t)} n_i \Phi'(\delta_i) \text{sign}(e_{b_i} - e_{a_i}). \quad (2.24)$$

The two expressions share the sum $\sum_{i \in \mathcal{S}(t)} n_i \Phi'(\delta_i) \text{sign}(e_{b_i} - e_{a_i})$; only the leading efficiency coefficient differs.

The stress function used here is the power law of Section 2.7.1, so the derivative that Equations 2.23–2.24 require is

$$\Phi'_{\text{shi}}(\delta) = k_3 k_4 \delta^{k_4-1}. \quad (2.25)$$

With $k_4 > 1$ this derivative is positive and increasing on $\delta > 0$, so a deeper cycle both costs more and attracts a larger marginal penalty. It vanishes at $\delta = 0$, so time steps that lie inside no cycle contribute nothing to the sub-gradient.

Equations 2.23–2.24 are computed for the whole horizon in a single pass. The contribution of each cycle is added at its opening index a_i and removed at its closing index b_i , and one cumulative sum over the time steps then gives the summation for every t at once. The computation is linear in the number of cycles and time steps.

This construction differs from the index-set partition of Shi et al. [59], who assign each time step to a single half-cycle. Their assignment holds when cycles do not nest: exactly one cycle spans each time step, and the sum in Equations 2.23–2.24 reduces to a single term with $n_i = \frac{1}{2}$. Under irregular market dispatch the SoC trajectory nests shallow cycles inside deeper ones, so more than one cycle spans a given time step and the single-assignment partition no longer applies. Whether nesting occurs often enough to matter is a property of the dispatch rather than of the method, and it is measured on the trajectories of this study in Appendix C.2. The spanning-set form is exact in both cases.

2.8.3. Efficiency Ratio Identity

The two derivatives share a structure that yields an exact check on any implementation of them. Because the summation in Equations 2.23–2.24 is common to the charge and discharge derivatives, their

ratio at any time step is independent of the stress function and the cycle amplitudes,

$$\left| \frac{\partial f / \partial p_t^{\text{dis}}}{\partial f / \partial p_t^{\text{ch}}} \right| = \frac{1}{\eta_{\text{in}} \eta_{\text{out}}}, \quad (2.26)$$

the inverse of the round-trip efficiency. This is an exact pointwise identity, and it provides a consistency check on any implementation: at every time step where the sub-gradient is non-zero, the ratio of the discharge to the charge derivative must equal $1/(\eta_{\text{in}} \eta_{\text{out}})$. The numerical verification is reported in Section 3.5.4.

2.8.4. Differentiation with Respect to Cycle Amplitude Only

The stress function carried into the sub-gradient is the cycle-amplitude function $\Phi(\delta)$ alone, without the mean-SoC factor S_σ or the temperature factor S_T of the physical reference model. Two reasons require this.

The first is the structure of the derivation. Raising p_t^{ch} adds the same constant to the stored energy at every step after t . An amplitude is a difference between two stored-energy values, so that constant cancels whenever both turning points move together, which is why only the cycles spanning t contribute to Equations 2.23–2.24. A cycle mean is an average rather than a difference, so the constant does not cancel: every cycle lying after t has its mean displaced. A cost carrying $S_\sigma(\sigma_i)$ would therefore collect a contribution at each time step from every later cycle in the horizon, not from the few that span it. On the annual dispatch of this study that additional term exceeds the amplitude term by more than two orders of magnitude, so it would dominate the sub-gradient rather than refine it (Appendix C.2).

The second is the convexity result of Section 2.7.3, which is stated for a stress function of cycle amplitude alone. A differentiated cost carrying S_σ falls outside it and forfeits the convergence guarantee that motivates using a sub-gradient at all.

The mean-SoC and temperature factors are therefore excluded from the dispatch sub-gradient, while the physical reference model of Section 2.6 retains them. The exclusion applies only to differentiation with respect to the dispatch. Differentiating with respect to the battery capacity is a different operation: scaling $\bar{E}$ stretches the trajectory rather than shifting it, so the cycle means rescale along with the amplitudes and S_σ can be differentiated in the usual way.

That capacity derivative is the second of two things this chapter leaves open. The first, in Section 2.7.2, is that the two models do not measure the same quantity as published. Neither is a defect of either model. Both are choices to be made in assembling a framework from them.

3

Methodology

This chapter builds the framework that evaluates a degradation-aware hybrid wind and battery design and describes how the design space is searched with it.

Section 3.1 introduces the modeling framework, the three-way coupling it has to represent and the three approaches through which the research question is examined. Section 3.2 then defines the case study with its reference plant, price data and economic parameters, and Section 3.3 describes the wind resource model and the hourly power time series it produces.

Sections 3.4 to 3.6 build the evaluation pipeline that the three approaches share: the battery dispatch optimization, the application of the two degradation models of Chapter 2 and the fitting of the polynomial one, and the reduction of capacity from the accumulated degradation across the project lifetime, with replacement.

Sections 3.7 and 3.8 develop the approaches themselves, which differ in how the sizing decision and the degradation model couple to that pipeline. Section 3.9 collects the modeling limitations.

3.1. Modeling Framework

Battery size, operational strategy and degradation rate are connected through distinct mechanisms that create a three-way coupling. The operational strategy depends directly on battery size: a larger energy capacity allows deeper arbitrage cycles and the time series of battery power scales roughly in proportion to the power capacity. This is a direct dependency, not a feedback loop.

The connection between battery size and degradation is economic. Degradation leads to earlier battery replacement, which adds cost. Sizing optimization must account for replacement costs that depend on how fast the battery degrades, so the optimal battery size is coupled to the degradation rate through the lifetime cost of replacements.

The degradation rate, in turn, depends on the operational strategy. Deep cycles, high states of charge and frequent cycling all accelerate capacity fade. The dispatch that maximizes short-term revenue is rarely the one that minimizes long-term degradation. Figure 3.1 illustrates these dependencies.

This coupling presents a methodological challenge. A sizing tool that ignores degradation may identify an optimal battery size that is economically inferior once replacement costs are accounted for. Conversely, a degradation model evaluated independently of the dispatch schedule cannot capture how operational choices alter the degradation trajectory. Integrating the two requires a framework in which dispatch, degradation and sizing are evaluated jointly.

Integrating dispatch, degradation and sizing requires three models evaluated together: a dispatch optimization, a degradation model and an economic model. Two degradation models are used for two purposes. The physical reference model of Xu et al. [63] provides experimentally validated accuracy and is used for all degradation reporting. The polynomial model [59] provides a differentiable cost




Figure 3.1: Three-way coupling between battery size, operational strategy and degradation rate.

and is used only where a gradient is required. Both are defined in Chapter 2; how each is applied is described in Section 3.5.

3.1.1. Integrated Evaluation of a Battery Design

Evaluating a single battery design requires three models applied in sequence. A dispatch optimization determines how the battery is operated. A degradation model translates that operation into capacity loss. An economic model converts revenue and replacement costs into a lifetime net present value. The three are applied in sequence rather than solved jointly, because degradation is computed from the dispatch solution rather than embedded in it (Section 3.4.7).

The dispatch optimization is a linear program that maximizes battery revenue over one year for a fixed energy capacity $\bar{E}$ and power capacity $\bar{P}$ (Section 3.4). It takes the hourly wind power series, the day-ahead price series, the SoC operating window and the round-trip efficiency, and returns the hourly stored-energy trajectory e_t . It has no representation of degradation.

The degradation evaluation takes the stored-energy trajectory, extracts rainflow cycles and accumulates annual capacity loss with the physical reference model (Section 3.5). Over the project horizon the accumulated loss is mapped to State of Health, and battery replacement is triggered when SoH reaches its threshold (Section 3.6).

The economic evaluation discounts annual revenue over the project lifetime and subtracts capital and replacement costs to give the lifetime NPV (Section 3.2.2).

Table 3.1 lists the inputs and outputs. A single design $(\bar{E}, \bar{P})$ maps to a single lifetime NPV. The design study of Section 3.7 evaluates this map over a grid of designs to locate the degradation-aware optimum.

Table 3.1: Inputs and outputs of the integrated evaluation of a single battery design.


Inputs	Outputs
Wind power series $P_{\text{wind}}(t)$	Stored-energy trajectory e_t
Day-ahead price series	Rainflow cycles (δ_i, σ_i)
Energy capacity $\bar{E}$, power capacity $\bar{P}$	Annual degradation f_d
SoC window $[\sigma_{\min}, \sigma_{\max}]$	State of Health trajectory
Round-trip efficiency	Battery replacement years
Cost parameters (Table 3.2)	Lifetime NPV
Degradation parameters (Table 2.2)	

3.1.2. Investigative Approach


The research question is examined through three approaches, each examining the coupling between dispatch and degradation at a different level of analytical integration. They do not represent a progression of successes: the first is completed and provides the main design result; the second is attempted and runs into a limitation of the degradation model it embeds; the third is developed analytically but not closed into a working implementation.

The three approaches share the same underlying evaluation pipeline and differ only in how the sizing decision and the degradation model are coupled to it. Figure 3.2 shows the baseline evaluation pipeline (a) together with the first two approaches, the parameter sweep (b) and the monolithic NLP (c); the third, gradient-based sizing, is shown in Figure 3.13.

(a) Baseline evaluation



(b) Parameter sweep



(c) Monolithic NLP




Figure 3.2: The shared evaluation pipeline and the first two investigative approaches. In the baseline evaluation (a), the dispatch LP (Section 3.4) is solved for a fixed design and its SoC trajectory is passed to the physical (Xu) and convex (Shi) degradation models (Section 3.5), which return lifetime NPV and end-of-life through the capacity reduction step (Section 3.6). The parameter sweep (b, Section 3.7) repeats this evaluation over a grid of designs and selects the maximum-NPV point. The monolithic NLP (c, Section 3.8.1) instead places the degradation cost inside the optimization objective. The third approach, gradient-based sizing, is shown separately in Figure 3.13 (Section 3.8.3).

The first approach is a *parameter sweep* (Section 3.7). For each point in a two-dimensional grid of energy capacity $\bar{E}$ and power capacity $\bar{P}$, the dispatch LP is solved to optimality and the resulting SoC trajectory is passed to the physical (Xu) degradation model, which accumulates annual capacity loss across a 20-year horizon and evaluates lifetime NPV. Repeating this over the grid evaluates NPV across the full design space and identifies where degradation-awareness shifts the optimal design relative to a degradation-free baseline. The sweep is exhaustive rather than directed and is the primary delivered contribution of this thesis.

The second approach is a *monolithic nonlinear program* (Section 3.8). It embeds the degradation cost directly inside the optimization objective, replacing exhaustive grid search with a single directed solve. The approach carries a known limitation: rainfall cycle counting is continuous but not always differentiable, so small changes to the dispatch can cause cycles to merge or split and the gradient a solver relies on to jump. Whether that limitation prevents a solution in practice is tested in Chapter 4, which reports that the solver certifies an optimum on roughly half the days attempted. Documenting that outcome is itself a finding, and it is what motivates the decomposition described next.

The third approach, *gradient-based sizing* (Section 3.8.3), is arranged so that no solver has to differentiate the rainfall map, and is the intended long-term direction of this work. It separates the dispatch LP from the sizing decision: the dispatch LP is solved for a fixed battery size, and the sizing is updated using a degradation sub-gradient computed analytically on the polynomial model. The analytical ingredients for the capacity gradient, the frozen-dispatch sub-gradient and the LP dual prices, are implemented and verified in this thesis. The sizing search itself is not closed and is identified as the principal direction for future work in Chapter 5.

Sections 3.2 to 3.6 develop each component of the shared pipeline in turn. Sections 3.7 and 3.8 then develop the three approaches built on it.

3.2. Case Study and Input Data

The methodology of Section 3.1 is independent of any particular site or market. This section defines the case study used to demonstrate it: the reference site and plant, the electricity price data and the economic parameters. The wind resource modeling that converts this site definition into the hourly power series $P_{\text{wind}}(t)$ is described in Section 3.3, as it forms part of the methodology rather than the case study.

3.2.1. Reference Site and Plant

This study uses the onshore reference site defined by IEA Wind Task 50 [26], located in western Denmark (56.20° N, 8.54° E, 86 m above ground level). The reference hybrid power plant at this site comprises a 325 MW wind farm, a 400 MW photovoltaic plant, a 150 MW/300 MWh battery and a 325 MW grid connection. Only the wind farm and the battery are modeled here; the photovoltaic plant is outside the scope of this study. The battery is a 150 MW/300 MWh lithium manganese oxide (LMO) system operated within a 10–90% SoC window (Section 3.4.4).

3.2.2. Electricity Price Data and Economic Parameters

Electricity prices are taken from the ENTSO-E Transparency Platform for the Danish day-ahead bidding zone DK1 [15]. Two price years are considered: 2019, representing moderate and stable market conditions and 2022, representing the high price volatility of the European energy crisis. Comparing the two years isolates the sensitivity of the optimal design and degradation trajectory to the price regime while the physical model parameters are held fixed. A discount rate of $r = 3\%$ and a project lifetime of $N = 20$ years are used throughout the NPV calculation. The battery is modeled as participating only in this day-ahead market; intraday, balancing and ancillary-service revenues are not represented. This understates the total revenue potential but keeps the dispatch problem tractable and the price signal unambiguous.

The battery cost assumptions follow the Danish Energy Agency technology-data catalogue for energy storage [10]. Initial costs use the 2025 technology year. Replacement costs use the catalogue's expansion-cost figures for 2035 to 2040. The catalogue reports constant values across that window, so the replacement cost is insensitive to the exact replacement year provided it falls inside it. The simulated replacement years reported in Chapter 4 satisfy this. For the initial build the catalogue reports an energy component of 95 EUR/kWh, a power component of 86 EUR/kW and a separate other-project-costs term of 150 EUR/kWh. The initial energy cost λ_E combines the energy component with the other-project-costs term, giving 245 EUR/kWh; the initial power cost λ_P is the power component alone. Table 3.2 lists these together with the remaining economic parameters.

Table 3.2: Economic parameters used in the NPV calculation. Battery cost components follow the Danish Energy Agency technology-data catalogue for energy storage, lithium-ion technology, real EUR₂₀₂₀ [10]. Initial costs use the 2025 technology year; replacement costs use the 2035–2040 expansion-cost figures.

Parameter	Value	Unit
Initial energy cost, λ_E (energy + other project)	245	EUR/kWh
Initial power cost, λ_P	86	EUR/kW
Replacement energy cost	72	EUR/kWh
Replacement power cost	96	EUR/kW
Discount rate, r (real)	3	%
Project lifetime, N	20	yr

3.3. Wind Resource & Power Generation Model

The wind power time series $P_{\text{wind}}(t)$ is the fixed exogenous input to the dispatch LP. Its shape determines the SoC trajectory that results from optimal dispatch, which in turn drives the rainflow cycle

statistics and the degradation output. The wind model must therefore be physically representative of a realistic onshore site.

3.3.1. Site and Wind Farm Configuration

The wind farm consists of 65 turbines arranged in a layout defined by IEA Wind Task 50, with a wind power density of 7 MW/km². The turbine coordinates and spatial layout are retained from the Task 50 site definition. The updated Task 50 repository specifies the NREL 5 MW reference turbine [33] for this site. Its key parameters are listed in Table 3.3.

Table 3.3: Key parameters of the NREL 5 MW reference turbine used in this study.

Parameter	Value	Unit
Rated power	5	MW
Rotor diameter	126	m
Hub height	90	m
Specific power	≈ 401	W/m ²
Cut-in wind speed	3	m/s
Rated wind speed	11.4	m/s
Cut-out wind speed	25	m/s

The power curve is reconstructed from the published C_p and C_t curves as $P(v) = C_p(v) \frac{1}{2} \rho (\pi D^2 / 4) v^3$, capped at rated power and set to zero outside the cut-in and cut-out range, using an air density of $\rho = 1.225 \text{ kg/m}^3$ (sea level, 15 °C). Because the dispatch LP operates on the normalized wind power series and battery sizing is expressed relative to the wind capacity, the degradation-relevant features of the SoC trajectory are not materially affected by the absolute power scale. The original Task 50 site specification implies a custom turbine with a rotor-specific power of 244 W/m², against approximately 401 W/m² for the NREL 5 MW machine adopted here. The substitution preserves the normalized capacity-factor profile and the wake structure. Both degradation models are driven by the same wind power series, so any scaling bias enters both identically and cancels in the comparison between them.

3.3.2. Wind Resource Data

Hourly wind speed and direction are obtained from the ERA5 reanalysis dataset, accessed through the Open-Meteo archive endpoint [64] with the model pinned to ERA5. ERA5 is provided on an approximately 31 km grid, so the series is extracted at the grid cell nearest to the site, 56.25° N, 8.50° E. The dataset covers 1 January to 31 December 2022 (8,760 hourly observations, UTC).

At this location the ERA5 archive provides wind speed at 10 m and 100 m. The 100 m series is scaled to the 90 m hub height using the power-law shear relation [39],

$$u(z) = u_{\text{ref}} \left(\frac{z}{z_{\text{ref}}} \right)^\alpha, \quad (3.1)$$

where u_{ref} is the wind speed at $z_{\text{ref}} = 100 \text{ m}$, $z = 90 \text{ m}$ is the hub height and $\alpha = 0.17$ is the shear exponent. The exponent is derived from the site's own ERA5 levels by fitting Equation 3.1 to the 10 m and 100 m wind speeds, rather than assumed from a terrain table. The resulting scaling factor is $(90/100)^{0.17} \approx 0.982$. A constant ambient turbulence intensity of $I = 0.14$ is applied across all directions, equal to the reference turbulence intensity I_{ref} of IEC turbulence class B [27], representative of medium-turbulence onshore conditions.

The shear exponent is held constant at $\alpha = 0.17$ throughout the year, although in reality it varies with atmospheric stability, time of day and season. The constant value is standard practice for the open, flat terrain of the Danish west coast, and the resulting bias in hub-height wind speed is small relative to other model uncertainties. Because ERA5 returns the nearest grid cell rather than the exact site

coordinates, the resulting resource is also an approximation of the wind climate near the planned wind farm; the spatial-averaging limitation of this reanalysis-driven approach is discussed in Section 3.9.

The resulting resource is summarized in Figures 3.3 and 3.4. The omnidirectional wind-speed distribution (Figure 3.3) is well described by a two-parameter Weibull distribution with scale $A = 9.24$ m/s and shape $k = 2.39$, giving an annual mean of 8.19 m/s. The wake model does not use this single pooled fit; it uses a directional fit in which a separate Weibull is fitted to each of twelve 30° sectors (frequency-weighted mean shape $k = 2.59$). The pooled fit yields a lower shape parameter than the sector-mean because pooling all directions mixes regimes of different strength into one broader distribution; the per-sector fits isolate each regime and recover a higher shape. The per-sector parameters are listed in Table B.1 and plotted in Figure B.1.




Figure 3.3: Omidirectional wind-speed distribution at 90 m hub height for the site, 2022. Histogram: hourly ERA5 wind speeds; curve: two-parameter Weibull fit ($A = 9.24$ m/s, $k = 2.39$). The wake model uses the directional twelve-sector fit instead (Table B.1, Figure B.1). Data are taken from the nearest ERA5 grid cell and therefore approximate the resource near the site.

The site is strongly directional. As shown in Figure 3.4, flow is dominated by the west, south-west and north-west sectors, which is consistent with the prevailing mid-latitude westerlies and the maritime exposure of the Danish west coast to the North Sea. Easterly sectors, which draw continental air from the European mainland, are both less frequent and lighter. The strongest winds, those at and above the rated speed of the turbine, arrive almost exclusively from the western half of the rose, so the directions that contribute most of the energy are also the most frequent. Wind speeds above the 25 m/s cut-out occur in only 6 hours of the entire year, so the turbines are effectively never shut down for high wind at this site. This directional structure is the reason the wake model uses the sector-wise Weibull fit: a single omnidirectional distribution would average over sectors whose scale parameter ranges from about 5 m/s in the rare north-easterly directions to over 10 m/s in the dominant westerlies (Appendix B).




Figure 3.4: Wind rose at 90 m hub height for the site, 2022, with wind speed binned by turbine operating state (below cut-in, partial load, approaching rated, at or above rated, above cut-out). Data are taken from the nearest ERA5 grid cell and therefore approximate the resource near the site.

3.3.3. Wake and Turbulence Model

Farm-level power production is computed using PyWake, the open-source wind farm flow modeling framework developed at the Department of Wind and Energy Systems of DTU [50]. Two aerodynamic models are used in combination: the Bastankhah–Porté-Agel (BPA) Gaussian wake model [4] for the velocity deficit and the Crespo–Hernández model [9] for the added turbulence intensity within the array. The full configuration is listed in Table 3.4.

The BPA model represents wake velocity deficits as Gaussian profiles with a constant wake expansion coefficient $k^* = 0.0572$. This value is not a default; it is set from the ambient turbulence intensity ($I = 0.14$) through the empirical relation $k^* = 0.38 I + 0.004$ of Niayifar and Porté-Agel [47], so that the assumed expansion rate is consistent with the site turbulence rather than an arbitrary constant. Bastankhah and Porté-Agel validated the deficit model against wind tunnel measurements and large-eddy simulation data for both single wakes and wind farm arrays [4]. The Crespo–Hernández model computes the added turbulence intensity generated within each wake, tracking the build-up of turbulence experienced by turbines in successive rows of the array as a function of downstream distance, ambient turbulence and thrust coefficient [9].

The BPA model was selected over simpler alternatives such as the N.O. Jensen (top-hat) model [32] because the Gaussian wake profile provides a more realistic representation of partial wake overlap in large arrays with irregular layouts [4]. At the other end of the fidelity spectrum, computational fluid dynamics or large-eddy simulation methods would provide higher physical accuracy but are computationally prohibitive for the 8,760 hourly time steps required by the dispatch optimization. The chosen combination offers a balance between physical accuracy and computational cost that is appropriate for generating the annual power time series needed by the dispatch LP.

Table 3.4: PyWake configuration used to generate the annual power time series.

Component	Model / setting	Value
Wake deficit model	Bastankhah–Porté-Agel (2014)	$k^* = 0.0572$ (constant)
Turbulence model	Crespo–Hernández	default coefficients
Wake superposition	Linear sum	–
Rotor averaging	Equidistant grid	3×3 (≈ 9 points)
Ambient turbulence	I	0.14 (IEC class B)
Shear exponent	α	0.17
Air density	ρ	1.225 kg/m ³

For each hourly time step, PyWake receives the hub-height wind speed and direction from the ERA5 time series. It computes the flow field through the full 65-turbine array and returns the output power of each turbine. Farm-level power is obtained by summing across all turbines, producing an 8,760-point annual time series $P_{\text{wind}}(t)$ in megawatts. Any generation that exceeds the grid connection limit and cannot be stored in the battery is curtailed by the LP. The resulting annual farm energy yield is approximately 1,230 GWh at a net capacity factor of 43%, after array wake losses of about 5% relative to the gross (wake-free) production. This capacity factor is high for an onshore site and sits above the Danish fleet average, which the 402 W/m² specific power of the NREL 5 MW turbine would ordinarily work against. It follows from the hub-height resource, which places the farm at or above rated power for 17.8% of the year. The resource is taken from reanalysis rather than from site measurement (Section 3.9). The absolute yield is not itself a result of this thesis: it sets the generation profile against which dispatch and degradation are evaluated in Chapter 4.

3.4. Battery Dispatch Optimization

The wind power time series of Section 3.3 and the day-ahead prices of Section 3.2.2 are the exogenous inputs to the dispatch problem, which determines the optimal charge and discharge schedule of the battery for a given energy capacity $\bar{E}$ and power capacity $\bar{P}$. The problem is a linear program (LP) solved once for each fixed capacity evaluated during the simulation. The solution produces the hourly SoC time series e_t that is subsequently passed to the degradation model.

The formulation follows the `lp_alt` variant of SHIPP [29], which uses a single signed power variable p_t^s for the storage system. The sign convention is: $p_t^s < 0$ denotes charging and $p_t^s > 0$ denotes discharging. The net power delivered to the grid from the battery is p_t^s .

3.4.1. Optimization Problem

The full LP is stated below. The objective function and each constraint group are explained in the subsections that follow.

$$\max_{p_t^s, p_t^c, e_t} \quad \gamma \sum_{t=1}^T \lambda_t (p_t^s - \varepsilon p_t^c) \Delta t - \lambda_P \bar{P} - \lambda_E \bar{E} \quad (3.2a)$$

$$\text{subject to} \quad e_{t+1} - e_t \leq -\eta_{\text{in}} p_t^s \Delta t, \quad \forall t, \quad (3.2b)$$

$$e_{t+1} - e_t \leq -\frac{p_t^s \Delta t}{\eta_{\text{out}}}, \quad \forall t, \quad (3.2c)$$

$$e_{\min} \leq e_t \leq e_{\max}, \quad \forall t, \quad (3.2d)$$

$$-\bar{P} \leq p_t^s \leq \bar{P}, \quad \forall t, \quad (3.2e)$$

$$0 \leq P_{\text{wind},t} + p_t^s - p_t^c \leq P_{\max}, \quad \forall t, \quad (3.2f)$$

$$p_t^s \leq \max(P_{\max} - P_{\text{wind},t}, 0), \quad \forall t, \quad (3.2g)$$

$$0 \leq p_t^c \leq P_{\text{wind},t}, \quad \forall t, \quad (3.2h)$$

$$e_1 = e_{T+1}. \quad (3.2i)$$

The decision variables are the storage power p_t^s , the curtailed power p_t^c and the stored energy e_t at each of the $T = 8,760$ hourly time steps. The exogenous inputs are the wind power $P_{\text{wind},t}$ and the day-ahead electricity price λ_t .

The LP is solved with full knowledge of the day-ahead prices over the whole horizon, so the dispatch has perfect price foresight. This yields an upper bound on achievable arbitrage revenue. Because both degradation models are evaluated on the same dispatch schedule, the overestimate affects them equally and does not bias the model comparison.

3.4.2. Objective Function

The objective (3.2a) maximizes the net present value (NPV) of the hybrid plant over a project lifetime of N years. The first term is the annualized revenue from electricity sales. The second and third terms are the capital costs of the battery power and energy capacity. The annuity factor γ converts the simulated year to a lifetime present value:

$$\gamma = \frac{365 \times 24}{T} \sum_{k=1}^{N-1} (1+r)^{-k}, \quad (3.3)$$

In this study, $r = 0.03$ and $N = 20$ years. The project horizon comprises one construction year (year 0), in which the capital expenditure is incurred and no revenue is generated, followed by $N - 1 = 19$ operating years over which revenue is discounted. This $k = 1 \dots N-1$ convention follows the SHIPP framework [29] and was confirmed as the intended behavior by the framework maintainer. The state-of-health trajectory is reported across all $N + 1$ year boundaries (years 0 to 20), as it is a capacity state rather than a discounted cash flow.

The regularization parameter $\varepsilon = 1 - 10^{-6}$ makes curtailed power marginally less valuable than stored power in the objective [29]. This ensures the LP prefers to store surplus wind rather than curtail it when both options are feasible. Without this term, the solver may produce degenerate solutions in which it curtails power that could have been stored at no cost. When the energy and power capacities are fixed parameters (as in each dispatch solve), the CAPEX terms $\lambda_P \bar{P}$ and $\lambda_E \bar{E}$ become constants and do not affect the dispatch. The LP then reduces to a pure revenue-maximization problem.

Figure 3.5 summarizes the dispatch LP as an extended design structure matrix (XDSTM) [35]. In this convention the component that drives the process sits on the diagonal, the data passed between components appear in the off-diagonal grey parallelograms, thick grey lines carry data and thin black lines carry the order of execution, numbered by step. The diagonal holds the LP optimizer, which coordinates two analysis blocks: the revenue objective (3.2a) and the constraint set (3.2b)–(3.2i). The exogenous inputs enter from the top: the fixed design parameters $\bar{E}$ and $\bar{P}$ above the optimizer, which holds them constant at each solve, and the wind and price time series above the objective, which reads them to value revenue. The optimizer proposes a dispatch, the two blocks evaluate it, and it iterates to optimality, returning the optimal dispatch: the storage power, curtailment and stored-energy trajectory $p_t^{s*}, p_t^{c*}, e_t^*$. Of the internal couplings, the objective reads the storage and curtailment power while the constraints additionally read the stored energy. The constraint groups shown here are detailed in the subsections that follow.

3.4.3. Storage Model

The stored energy e_t changes at each time step according to the charging and discharging efficiencies. When the battery charges ($p_t^s < 0$), only a fraction η_{in} of the input power reaches the cells. When the battery discharges ($p_t^s > 0$), the cells must supply more energy than is delivered, by a factor of $1/\eta_{\text{out}}$. The exact relationship is

$$e_{t+1} - e_t = \begin{cases} -\eta_{\text{in}} p_t^s \Delta t & \text{if } p_t^s \leq 0 \text{ (charging),} \\ -\frac{p_t^s \Delta t}{\eta_{\text{out}}} & \text{if } p_t^s > 0 \text{ (discharging),} \end{cases} \quad (3.4)$$

where $\eta_{\text{in}} \in (0, 1]$ is the charging efficiency and $\eta_{\text{out}} \in (0, 1]$ is the discharging efficiency. Both are set to the symmetric one-way value $\eta_{\text{in}} = \eta_{\text{out}} = 0.9539$. This value follows from the AC round-trip efficiency




Figure 3.5: XDSM of the dispatch LP. The fixed design parameters $\bar{E}, \bar{P}$ enter above the optimizer and the wind and price series above the objective; the optimizer iterates between the revenue objective (Equation 3.2a) and the constraint set (Equations 3.2b–3.2i) to optimality, returning the optimal dispatch. The reading convention and data exchanges are described in the text above.

of 0.91 reported for lithium-ion systems in the Danish Energy Agency technology-data catalogue [10], through $\eta_{\text{in}} = \eta_{\text{out}} = \sqrt{0.91} = 0.9539$. The same catalogue reports a DC round-trip efficiency of 0.95, which implies a power-conditioning efficiency of $\sqrt{0.91/0.95} = 0.979$ at each AC-DC conversion. The resulting AC round-trip efficiency is $\eta_{\text{in}} \cdot \eta_{\text{out}} = 0.91$. The efficiency assumption and the battery cost assumption of Section 3.2.2 are therefore drawn from the same source.

Equation (3.4) switches between two linear expressions depending on the sign of p_t^s . This cannot be used directly in a linear program without introducing binary variables. SHIPP avoids binary variables by replacing the equality (3.4) with two relaxed inequalities (3.2b)–(3.2c), restated here for reference:

$$e_{t+1} - e_t \leq -\eta_{\text{in}} p_t^s \Delta t, \quad (3.2b)$$

$$e_{t+1} - e_t \leq -\frac{p_t^s \Delta t}{\eta_{\text{out}}}. \quad (3.2c)$$

Both constraints hold for all t . At the LP optimum, exactly one inequality is active at each time step. The charging constraint (3.2b) binds when $p_t^s < 0$ and the discharging constraint (3.2c) binds when $p_t^s > 0$. The relaxation is therefore exact at the optimum. This exactness holds only where the objective actively drives the constraint. At time steps where the objective is indifferent to the stored energy, and at SoC-bound steps, one inequality can remain slack, which injects shallow phantom micro-cycles into the recorded SoC trace. At the operating points examined here these inflate the raw rainflow cycle count by up to about 16% at small energy capacities, but their effect on f_d is below 0.25%, because accumulation is dominated by deep cycles ($k_4 = 1.18 > 1$ in Φ_{shi}) and the phantom cycles are shallow. Revenue and NPV are unaffected, as they are computed from dispatch power rather than the energy trace. Reported raw cycle counts and equivalent full cycles should therefore be read as approximate upper bounds.

Efficiency convention Shi et al. [59] define the discharge efficiency as $\eta_d \geq 1$, the reciprocal of the SHIPP convention: $\eta_d = 1/\eta_{\text{out}}$. Under this convention, the discharge constraint (3.2c) can be written as $e_{t+1} - e_t \leq -\eta_d p_t^s \Delta t$. The two forms are algebraically equivalent. This thesis uses the SHIPP convention ($\eta_{\text{out}} \in (0, 1]$) in the LP formulation and notes the correspondence when referencing Shi’s sub-gradient derivation in Section 3.4.7.

3.4.4. Storage Constraints

Energy bounds and SoC operating window.

The stored energy at every time step is bounded by (3.2d), where

$$e_{\min} = \sigma_{\min} \cdot \bar{E}, \quad e_{\max} = \sigma_{\max} \cdot \bar{E}, \quad (3.5)$$

and $\sigma_{\min}$, $\sigma_{\max}$ are the minimum and maximum allowable SoC fractions. The baseline runs use $\sigma_{\min} = 0.10$ and $\sigma_{\max} = 0.90$, giving $e_{\min} = 0.10 \bar{E}$ and $e_{\max} = 0.90 \bar{E}$.

The SoC operating window is an extension added to the standard SHIPP formulation for this thesis. The standard SHIPP energy bounds use the depth-of-discharge parameter d for the lower limit and the full capacity $\bar{E}$ for the upper limit: $\bar{E}(1 - d) \leq e_t \leq \bar{E}$. Setting $d = 1 - \sigma_{\min}$ and $\sigma_{\max} = 1$ recovers the standard bounds. The explicit $\sigma_{\max}$ parameter allows restricting the upper SoC limit independently, which is relevant for degradation analysis because high SoC levels accelerate calendar aging [63].

The maximum achievable cycle depth under this window is $\delta_{\max} = \sigma_{\max} - \sigma_{\min} = 0.80$, which sets the upper bound of the Φ_{shi} fitting range in Section 3.5.2. The window is held at 10–90% throughout, except in the second sweep stage, where it becomes a design parameter (Section 3.7.2); its effect on degradation is reported in Section 4.4.2.

Power bounds.

The storage power is bounded symmetrically by the power capacity (3.2e). Because the `lp_alt` formulation uses a single signed variable, the lower bound $-\bar{P}$ corresponds to the maximum charging power and the upper bound $\bar{P}$ to the maximum discharging power.

3.4.5. Dispatch Constraints

Grid connection bounds.

The combined output of the wind farm and the storage must respect the grid connection capacity (3.2f). The upper bound $P_{\max}$ prevents the plant from injecting more power than the grid connection allows. The lower bound of zero prevents the battery from drawing power from the grid: the plant can only charge the battery from its own wind generation.

Storage headroom.

An additional constraint (3.2g) limits the storage discharge power to the available headroom between the grid connection capacity and the wind power. This prevents the storage from discharging into the grid beyond what the connection can deliver, independently of curtailment. In a single-storage configuration, this constraint is $p_t^s \leq \max(P_{\max} - P_{\text{wind},t}, 0)$.

Curtailment.

Curtailment is an explicit decision variable in SHIPP, bounded by the available wind power (3.2h). When net production exceeds the grid connection capacity and the battery cannot absorb the surplus, the optimizer curtails wind via p_t^c . The regularization parameter ε in the objective (3.2a) ensures curtailment is used only when the surplus cannot be stored.

3.4.6. Periodicity and inter-year continuity.

The periodicity constraint (3.2i) enforces that the battery ends the simulated year at the same stored energy at which it started. In the first year of each battery lifetime, the optimizer freely selects the starting SoC level. This is the standard SHIPP behavior.

From the second year onward, a pinning constraint fixes the starting level:

$$e_1 = \hat{\sigma} \cdot \bar{E}_k, \quad (3.6)$$

where $\hat{\sigma}$ is the SoC fraction selected by the optimizer in year 1 and $\bar{E}_k = \bar{E} \cdot \text{SoH}_{k-1}$ is the degraded capacity entering year k . Because $\hat{\sigma}$ is a fraction of current capacity (not an absolute energy level), degradation does not break inter-year continuity. The SoC fraction resets when a battery is replaced.

The pinning constraint (3.6) and the SoC fraction continuity logic are additions to the standard SHIPP formulation made for this thesis..

3.4.7. Implementation

The LP is implemented in the SHIPP framework [29] using Pyomo [22] and solved with Gurobi [20].

In the standard SHIPP formulation, energy capacity $\bar{E}$ and power capacity $\bar{P}$ are decision variables of the LP. The solver jointly optimizes capacity and dispatch in a single solve, sizing the battery to maximize revenue and minimize capital cost without any representation of degradation. This is the appropriate formulation when degradation is neglected entirely.

This thesis modifies that formulation: $\bar{E}$ and $\bar{P}$ are fixed parameters at each dispatch solve. In the parameter sweep, they are set by the grid coordinates. In gradient-based sizing, they would be updated between solves by the sizing search (Section 3.8.3). Degradation is not embedded in the LP objective or constraints. It is evaluated *after* each LP solve by passing the resulting SoC trajectory e_t to the degradation model described in Sections 2.6–3.5.

Dual prices.

Each dispatch LP solve also returns the dual prices of the energy balance constraints (3.2b)–(3.2c): the marginal revenue of stored energy at each time step. Together with the degradation sub-gradient, they form the capacity gradient used by the gradient-based sizing method, assembled in Section 3.5.3. Dual price extraction is an addition to the standard SHIPP formulation made for this thesis.

Summary of extensions.

Table 3.5 summarizes the standard SHIPP constraints and the additions made for this thesis.

Table 3.5: Dispatch LP constraints: standard SHIPP compared with the extensions made for this thesis. The modified SHIPP source is available at [61].

Constraint	Standard SHIPP	Thesis extension
Relaxed storage dynamics (3.2b)–(3.2c)	✓	
Power bounds (3.2e)	✓	
Energy lower bound ($e_{\min}$)	$\bar{E}(1 - d)$	$\sigma_{\min} \bar{E}$ ¹
Energy upper bound ($e_{\max}$)	$\bar{E}$	$\sigma_{\max} \bar{E}$
Periodicity (3.2i)	✓	
Inter-year SoC pinning (3.6)		✓
Dispatch bounds (3.2f)	✓	
Storage headroom (3.2g)	✓	
Curtailment bound (3.2h)	✓	
Curtailment regularization (ε)	✓	
Dual price extraction		✓
Sub-gradient coupling		✓

3.5. Degradation Evaluation

Chapter 2 defined two degradation models: the physical reference model (Section 2.6) and the polynomial model (Section 2.7). This section describes how they are applied: the role each plays, the fitting of the polynomial stress function to the operating window, the verification of the sub-gradient and the treatment of calendar aging.

3.5.1. The Two Degradation Models

The physical reference model and the polynomial model are not competing alternatives. The first supplies every degradation quantity reported in this thesis. The second supplies the gradient information a sizing search requires and is never used for reporting.

Neither is applied exactly as published. Chapter 2 presents each model as its authors define it. This subsection states how they are assembled into the two computation paths used here, referred to throughout

¹The SHIPP lower bound $\bar{E}(1 - d)$ with $d = 1 - \sigma_{\min}$ is mathematically identical to $\sigma_{\min} \bar{E}$. The difference is in parametrization: SHIPP exposes a depth-of-discharge parameter d on the Storage object, while this thesis uses explicit SoC fraction limits $\sigma_{\min}$, $\sigma_{\max}$ for clarity.

as the Xu branch and the Shi branch. Both receive the same rainflow-identified cycles $\{(\delta_i, \sigma_i, n_i)\}$ from the same SoC trajectory, and both are evaluated at every design point.

Xu branch.

The physical reference model is applied unchanged. Substituting the per-cycle loss of Equation 2.9 into the accumulation of Equation 2.6 gives the annual degradation

$$f_d^{\text{Xu}} = f_d^{\text{cal}} + \sum_i n_i S_\delta(\delta_i) S_\sigma(\sigma_i) S_T(T_{c,i}). \quad (3.7)$$

Shi branch.

The polynomial model as published carries a cycle-amplitude stress function and an accumulation rule only (Section 2.7.2). Three elements are supplied from the physical reference model so that the two branches return comparable quantities. The mean SoC and temperature factors are applied per cycle, so the per-cycle loss keeps the structure of Equation 2.9 with Φ_{shi} in place of S_δ . The calendar term of Equation 2.15 is added unchanged. The accumulation is therefore

$$f_d^{\text{Shi}} = f_d^{\text{cal}} + \sum_i n_i \Phi_{\text{shi}}(\delta_i) S_\sigma(\sigma_i) S_T(T_{c,i}), \quad (3.8)$$

identical to Equation 3.7 except in the amplitude term. The third supplied element is the mapping from f_d to State of Health, which is the SEI fade curve of Equation 2.7 applied to both branches; that choice is set out in Section 3.6. Because Φ_{shi} is convex by construction ($k_4 > 1$), Theorem 1 of Section 2.7.3 applies to the cost built from its amplitude term.

Calendar aging.

The polynomial model accumulates through cycle counting alone and has no calendar term of its own. Calendar aging can carry a large share of total degradation, and that share depends on the operating regime rather than on the chemistry, so omitting it would understate annual degradation. The Xu calendar term is therefore supplied to both branches. It is computed from the same elapsed time and the same period mean SoC in each case, so it takes the same value on both, and the difference between the two accumulations is confined to the cycle term.

The differentiated cost is narrower than the accumulation.

Equation 3.8 is what the Shi branch accumulates. It is not what the sub-gradient differentiates. The differentiated quantity is the dispatch cost of Equation 2.18, which carries the amplitude term $\Phi_{\text{shi}}(\delta_i)$ alone. The calendar term is excluded because it does not depend on the dispatch: $\partial f_d^{\text{cal}} / \partial p_t^{\text{ch}} = 0$ for all t , so including it would leave every sub-gradient value unchanged. The mean SoC and temperature factors are excluded for the reason given in Section 2.8.4. The two quantities are kept distinct throughout: Equation 3.8 is reported, Equation 2.18 is differentiated.

Temperature assumption.

The cell temperature is held constant at $T_c = 25^\circ\text{C}$, so $S_T = 1$ for every cycle on both branches. For a grid-scale system with active thermal management this is a reasonable approximation at the reference temperature for which the Xu parameters were calibrated. The Xu temperature model is validated only above 15°C [63], so the assumption also marks the lower bound of model validity in cold conditions, since wintertime ambient temperatures at the site can fall below this threshold.

Table 3.6 sets out the composition of the two branches element by element, and Figure 3.6 shows the two computation paths from rainflow extraction to their respective outputs.

Table 3.6: Composition of the two degradation models as applied in this thesis. Both use the same rainflow cycles, the same calendar term and the same capacity reduction, so the two accumulations differ in the cycle-amplitude stress function alone. The last two rows show where the branches part company: what each is used for, and how much of the accumulation is carried into the derivative.

Element	Physical reference model (Xu)	Polynomial model (Shi)
Cycle amplitude stress	S_δ (Eq. 2.10)	Φ_{shi} (Eq. 2.16)
Mean SoC stress	S_σ (Eq. 2.12)	S_σ , unchanged
Temperature stress	S_T (Eq. 2.13)	S_T , unchanged
Calendar term	Eq. 2.15	Eq. 2.15, unchanged
Accumulation	Eq. 3.7	Eq. 3.8
Fade mapping to SoH	Eq. 2.7	Eq. 2.7, unchanged
End-of-life prediction	reported throughout	model comparison only
Carried into the derivative	not differentiated	amplitude term only (Eq. 2.18)




Figure 3.6: The two computation paths, from the shared rainflow cycles to the simulation outputs (top) and the optimization input (bottom). The Shi accumulation of Equation 3.8 also maps to State of Health, for the model comparison reported in Chapter 4; that path is not drawn here.

3.5.2. Fitting the Polynomial Stress Function

The polynomial $\Phi_{shi}(\delta) = k_3 \delta^{k_4}$ and its second derivative are defined in Section 2.7.1. This subsection fits the coefficients k_3 and k_4 to S_δ over the operating range of cycle amplitudes for the window used in this study.

Fitting window.

The coefficients k_3 and k_4 are fitted to S_δ over $\delta \in [0.15, 0.80]$. Both bounds follow from the system's operating constraints.

The upper bound $\delta_{\max} = 0.80$ is set by the SoC operating window. With $\sigma_{\min} = 0.10$ and $\sigma_{\max} = 0.90$,

$$\delta_{\max} = \sigma_{\max} - \sigma_{\min} = 0.90 - 0.10 = 0.80. \quad (3.9)$$

Any cycle deeper than this is physically impossible under the enforced SoC constraints, so fitting Φ_{shi} beyond $\delta = 0.80$ has no effect on the gradient computation.

The lower bound $\delta_{\min} = 0.15$ is set by the concave region of S_δ , which extends up to the inflection point $\delta_c = 0.1437$ of Equation 2.11. Figure 3.7 shows the sign change. Fitting Φ_{shi} within this region would

either violate the $k_4 > 1$ requirement or produce a poor approximation of the gradient. The lower bound is placed just above the inflection point.




Figure 3.7: Second derivative of $S_\delta(\delta)$. The shaded region ($\delta < 0.144$) is non-convex; the sign change is marked by the dashed line. The vertical axis is truncated at 12, and S''_δ continues to rise beyond the frame, reaching about 35 at $\delta = 0.80$.

Fitted coefficients.

The coefficients are obtained by ordinary least squares in log-log space. Taking the logarithm of Equation 2.16 gives a linear relation, $\log \Phi_{\text{shi}}(\delta) = \log k_3 + k_4 \log \delta$, so $\log k_3$ and k_4 are the intercept and slope of a straight-line fit to the points $(\log \delta, \log S_\delta(\delta))$. The stress function S_δ is sampled on a uniform grid of 500 points over $\delta \in [0.15, 0.80]$. The convexity condition $k_4 > 1$ is confirmed after fitting rather than imposed as a constraint. The fitted values for the $[0.10, 0.90]$ SoC window are

$$k_3 = 3.24 \times 10^{-5}, \quad k_4 = 1.179, \quad (3.10)$$

with a coefficient of determination $R^2 = 0.963$, evaluated in linear space, for the fit of Φ_{shi} to S_δ over the fitting window. This is distinct from the $R^2 = 0.993$ reported in Section 2.6.2 for the fit of S_δ to the underlying LMO cycle-life data. The coefficients depend on the SoC window through the upper fitting bound $\delta_{\text{max}} = \sigma_{\text{max}} - \sigma_{\text{min}}$ and are recomputed whenever the window changes. Figure 3.8 shows the fitted Φ_{shi} against S_δ over the full range including the non-convex region.

The two functions agree closely across the middle of the fitting range and separate at its upper end, where S_δ grows faster than any fixed power of δ . The polynomial therefore under-represents the curvature of S_δ at the deepest cycles.




Figure 3.8: Φ_{shi} (dashed) fitted to S_δ (solid) over $\delta \in [0.15, 0.80]$. The lower panel confirms $\Phi''_{\text{shi}} > 0$ everywhere while S''_δ goes negative below $\delta_c = 0.1437$. The vertical axis of the lower panel is truncated at 12; both second derivatives continue beyond the frame.

Because $\Phi_{\text{shi}}(\delta) = k_3 \delta^{k_4}$ with $k_4 = 1.179 > 1$ is strictly convex, Theorem 1 (Section 2.7.3) applies to the fitted polynomial: the degradation cost accumulated with the polynomial model is a convex function of the SoC profile. Figure 3.9 confirms the convexity visually.




Figure 3.9: $\Phi_{\text{shi}}(\delta)$ and its second derivative, for the coefficients fitted over $\delta \in [0.15, 0.80]$ under the 10–90% SoC window. The second derivative is positive everywhere, confirming global convexity.

Cycles below the fitting window.

The lower fitting bound is set at $\delta = 0.15$ rather than at the inflection point $\delta_c = 0.1437$, so that the fitted polynomial lies wholly inside the convex region of S_δ with margin. Cycles shallower than the bound contribute to f_d through S_δ in the physical model and receive sub-gradient values extrapolated from Φ_{shi} in the convex model. Shallow cycles carry little of the accumulated damage, so the gradient signal in this region is small. At the reference design, cycles in the non-convex region below δ_c account for 1.9% of year-1 cycle degradation on the physical branch and 1.3% on the convex branch, while making up about a quarter of the counted rainflow records. The error the extrapolation introduces is quantified across the operating window series in Section 4.4.2.

3.5.3. Degradation Sub-gradient and the Capacity Gradient

Two derivatives of the degradation model are computed in this thesis, and they differentiate different quantities for different purposes. This subsection states both and the role each one plays.

Dispatch sub-gradient.

The first is the derivative of the dispatch degradation cost with respect to the charging and discharging power at each time step, derived in Section 2.8 and given by Equations 2.23 and 2.24. It is evaluated on the cycles of a solved dispatch and carries the amplitude term of the polynomial model alone. Its physical reading is the efficiency asymmetry of Equation 2.26: with $\eta_{in} = \eta_{out} = 0.9539$ the discharge sensitivity exceeds the charge sensitivity by $1/(\eta_{in}\eta_{out}) \approx 1.099$, because the cells must supply the extra energy that covers the conversion loss. That identity holds at every time step and is used as a consistency check on the implementation (Appendix C.2).

Frozen-dispatch capacity gradient.

The second is the derivative of the annual degradation with respect to the energy capacity at a fixed dispatch. It is the quantity a sizing search needs. Raising $\bar{E}$ while holding the energy dispatch fixed rescales the normalized state-of-charge trajectory, so every cycle amplitude δ_i , every cycle mean σ_i and the year's mean state of charge $\bar{\sigma}$ all fall as $1/\bar{E}$. Both the cycle term and the calendar term respond, and the derivative is available in closed form,

$$\left. \frac{df_d}{d\bar{E}} \right|_{\text{frozen}} = -\frac{1}{\bar{E}} \left[\sum_i f_{d,i}^{\text{Shi}} (k_4 + k_\sigma \sigma_i) + k_\sigma \bar{\sigma} f_d^{\text{cal}} \right], \quad (3.11)$$

where $f_{d,i}^{\text{Shi}} = n_i \Phi_{shi}(\delta_i) S_\sigma(\sigma_i) S_T$ is the per-cycle degradation of the Shi branch (Equation 3.8), k_4 is the fitted exponent, k_σ is the mean-SoC stress coefficient and f_d^{cal} is the calendar term.

The bracket has three parts. The factor k_4 is the response of the amplitude power law to the rescaling. The factor $k_\sigma \sigma_i$ is the response of the per-cycle mean-SoC stress, which changes because rescaling moves every cycle mean. The final term is the response of the calendar stress to the same shift in mean state of charge.

The mean-SoC coupling is not a correction: dropping it leaves the amplitude-only slope $k_4 f_d^{\text{cyc,Shi}}$, which is about 31% below the complete cycle slope in both price years, and 48% below the total slope once the calendar coupling is counted (Section C.1). It is the mechanism behind the sensitivity of degradation to the window center reported in Section 4.4.2.

Equation 3.11 is assembled from the per-cycle contributions and does not contract the dispatch sub-gradient against the LP dual prices. A per-year signal formed that way was examined and rejected, because it varies with the temporal alignment of the two vectors rather than with any change in the battery; Appendix C.3 gives the diagnosis. The dual prices enter the sizing gradient on the revenue side instead, as set out in Section 3.8.3, where the two components of $d\text{NPV}/d\bar{E}$ are described.

3.5.4. Gradient Verification

The two derivatives of Section 3.5.3 are verified by separate checks on the same solved Year 1 dispatch. Test 1 (Section C.1) verifies the frozen-dispatch capacity gradient of Equation 3.11 against a finite difference over a sweep of step sizes, and against the closed-form expression directly. Test 2 (Section C.2) verifies the per-timestep sub-gradient of Equations 2.23 and 2.24 by perturbing the charging power one

step at a time, and confirms the discharge-to-charge identity of Equation 2.26. Where rainflow re-pairs the cycles no derivative exists, so the criterion there is that the analytic value lies between the two one-sided difference quotients rather than equaling either.

3.6. Capacity Reduction

Both degradation models produce an annual fractional degradation rate f_d as their primary output. This quantity is a linearized measure of accumulated stress and is not directly interpretable as a capacity loss percentage. A capacity reduction function is required to translate f_d into the physical capacity loss L , whose complement is the State of Health against which battery end-of-life is defined.

The mapping is the second place where the two models could differ, after the accumulation rule. A linear mapping $L = 1 - c_{\text{deg}} f_d$ is the simplest, but its scaling constant has no anchor in experimental data. Calibrating that constant so the two models reach end of life together removes the arbitrariness at one point only: the models are then free to diverge at intermediate states of health, and the accumulation rule and the fade curve are conflated, so a difference in lifetime can no longer be attributed to either.

This study instead applies the Xu SEI two-exponential fade curve to both branches. Holding the mapping fixed and varying only the accumulation rule makes any difference in projected lifetime attributable solely to how f_d is accumulated from rainflow cycles. The capacity loss on both branches is then computed from Equation 2.7,

$$L(f_d) = 1 - \alpha_{\text{sei}} e^{-\beta_{\text{sei}} f_d} - (1 - \alpha_{\text{sei}}) e^{-f_d}, \quad (3.12)$$

where α_{sei} and β_{sei} are the LMO coefficients from Table 2.2. The State of Health is the complement, $\text{SoH} = 1 - L(f_d)$ (Equation 2.8), and replacement is triggered when SoH reaches the trigger applied in the annual loop (Section 3.6.1). For the physical reference model, $f_d = f_d^{\text{Xu}}$ of Equation 3.7. For the polynomial model, $f_d = f_d^{\text{Shi}}$ of Equation 3.8. Both include the calendar term, so the capacity reduction function receives the same kind of quantity in each case.

This embeds a working hypothesis: that the relationship between accumulated degradation and capacity retention is independent of how that degradation was accumulated. It is adopted for the controlled comparison rather than asserted as a physically validated property.

3.6.1. Multi-Year Degradation and Battery Replacement

The single-year EoL estimate assumes that the year-1 rate $f_{d,1}$ is constant across the project lifetime. This neglects two feedback mechanisms. First, as the battery degrades, its effective energy capacity decreases, which changes the SoC trajectory and therefore the cycle depth distribution and annual f_d . Second, a battery that has lost enough capacity to fall below the replacement threshold must be replaced before the project ends. A multi-year simulation loop addresses both. The loop can be run for any number of years; a 20-year horizon is adopted here. This duration is standard in hybrid power plant techno-economic assessments and is consistent with the HyDesign reference framework [44]. It is also long enough to observe at least one replacement event under both price scenarios while retaining several years of post-replacement operation, which allows the second battery generation to be compared against the first.

The loop repeats annually for the full 20-year project horizon. The wind power series, the day-ahead price series and the nominal power capacity $\bar{P}$ are identical in every year, so each year replays the same representative year of operation. The only dispatch input that changes between years is the effective energy capacity $\bar{E}_k$, which is reduced in proportion to the retained state of health as the battery fades. At each year the dispatch LP is re-solved with this degraded capacity, producing an updated SoC trajectory. The trajectory passes through rainflow counting and the degradation model to compute the annual $f_{d,k}$, which is accumulated and mapped through the SEI fade curve to update SoH. If SoH falls below the replacement trigger at the end of the year, replacement is triggered: SoH resets to 1.0, the replacement cost is subtracted from NPV and the SoC continuity fraction $\hat{\sigma}$ resets for the new battery. If SoH remains above it, the loop advances to the next year with the degraded capacity. Figure 3.10 shows the loop structure and marks which dispatch inputs are held fixed across years and which one updates.




Figure 3.10: Multi-year simulation loop. The effective energy capacity $\bar{E}_k = \bar{E} \cdot \text{SoH}_{k-1}$ is the only dispatch input that changes from year to year, as marked in the figure; the wind and price series and the nominal power $\bar{P}$ are held fixed.

Annual capacity feedback.

At the start of each year k , the effective energy capacity is

$$\bar{E}_k = \bar{E} \cdot \text{SoH}_{k-1}, \quad (3.13)$$

where SoH_{k-1} is the state of health at the end of the previous year. The dispatch LP is re-solved with $\bar{E}_k$ as the fixed capacity, producing a dispatch schedule adapted to the degraded battery.

The power capacity $\bar{P}$ is held at its nominal value throughout, on the assumption that the battery management system maintains the same rate limits regardless of capacity fade; a consequence is that the effective C-rate $\bar{P}/\bar{E}_k$ rises each year as SoH declines, which the C-rate-free Xu stress model does not reflect (Section 3.9).

The annual degradation $f_{d,k}$ is computed from this schedule. The cumulative degradation is updated as $f_{d,\text{cum},k} = f_{d,\text{cum},k-1} + f_{d,k}$ and the state of health from the fade curve,

$$\text{SoH}_k = 1 - L(f_{d,\text{cum},k}), \quad (3.14)$$

Inter-year SoC continuity.

Each annual LP solve enforces the within-year periodicity constraint $e_1 = e_{T+1}$ (Equation 3.2i). This ensures the battery returns to the same SoC fraction at the end of each simulated year as it started, which is the physically correct constraint when the same representative dataset is repeated every year. It does not, however, link the terminal SoC of year k to the initial SoC of year $k+1$.

To enforce physical continuity across year boundaries, a pinning constraint is added from year 2 onward. Year 1 optimizes freely and the optimal SoC fraction at which the battery settles is recorded as $\hat{\sigma} = e_{T+1}/\bar{E}_1$. For all subsequent years of the same battery lifetime, the starting SoC is fixed to

$$e_1^{(k)} = \hat{\sigma} \cdot \bar{E}_k, \quad k \geq 2, \quad (3.6)$$

where $\bar{E}_k$ is the degraded capacity of year k . The periodic constraint $e_1 = e_{T+1}$ remains active alongside the pinning constraint in all years, so both must hold simultaneously. This formulation enforces continuity in SoC *fraction* rather than absolute energy, which is physically consistent because SoH is defined relative to current capacity. Battery replacement resets $\hat{\sigma}$: the new battery's year-1 LP optimizes the SoC fraction freely again.

Battery replacement.

Replacement is triggered at the end of any year in which SoH_k falls below the replacement trigger $\text{SoH}_{\text{repl}} = 0.705$, which is the end-of-life threshold $\text{SoH}_{\text{eol}} = 0.70$ of Section 2.6.3 plus a tolerance of half a percentage point. The check is made once a year, so without the tolerance a battery that ended a year marginally above the threshold would operate a further full year and spend all of it below the level it is meant to respect. The tolerance can only advance a replacement, never delay one, so it does not extend any simulated lifetime. Upon replacement, the cumulative degradation accumulator resets to zero, SoH resets to 1.0 and the nominal capacity $\bar{E}$ is restored. The replacement cost $C_{\text{repl}} = \lambda_E^{\text{repl}} \bar{E} + \lambda_P^{\text{repl}} \bar{P}$ uses the replacement energy and power costs of Table 3.2 (72 EUR/kWh and 96 EUR/kW), not the initial-build costs λ_E and λ_P . It is discounted to project year k and subtracted from the project NPV,

$$\text{NPV}_{\text{net}} = \text{NPV}_{\text{revenue}} - \sum_{k \in \mathcal{K}_{\text{repl}}} \frac{\lambda_E^{\text{repl}} \bar{E} + \lambda_P^{\text{repl}} \bar{P}}{(1+r)^k}, \quad (3.15)$$

where $\mathcal{K}_{\text{repl}}$ is the set of years in which replacement occurs. The number of replacements over the 20-year project lifetime is a simulation output and varies with the price year and the accumulation model.

The end-of-life threshold itself is set in Section 2.6.3; this section defines only how it is applied inside the annual loop.

3.7. Design Space Exploration via Parameter Sweep

The degradation model and dispatch framework of Sections 2.6–3.6 quantify how a given battery size and operating strategy translate into lifetime NPV. To systematically explore how these quantities interact across the feasible design space, a two-stage parameter sweep was conducted. The first stage varies the battery size across an $(\bar{E}, \bar{P})$ grid to locate the NPV-optimal design point under degradation-aware and degradation-free assumptions. The second stage fixes the battery size at the identified optimum and varies the SoC operating window to identify which window parameters act as degradation levers. Figure 3.11 illustrates the two-stage structure.




Figure 3.11: Two-stage parameter sweep structure. Stage 1 (left): lifetime NPV is evaluated across an $(\bar{E}, \bar{P})$ grid under three degradation scenarios to identify the optimal battery sizing. Stage 2 (right): the SoC operating window is varied at the identified optimum to quantify the effect of window width and window center on degradation and NPV.

3.7.1. Stage 1: Energy and Power Capacity Sweep

The sizing sweep evaluates 20-year lifetime NPV across a two-dimensional grid of energy capacity $\bar{E}$ and power capacity $\bar{P}$. For each grid point, the multi-year simulation loop of Section 3.6.1 is executed in full: the dispatch LP is solved for each of the 20 project years at the degraded effective capacity, rainflow counting and the degradation models accumulate annual f_d and battery replacement is triggered when SoH falls below the replacement trigger of Section 3.6.1. Three scenarios are evaluated at each point: no-degradation (dispatch LP only, no degradation cost), degradation with the physical reference model and degradation with the polynomial model.

The sweep was conducted in two passes designed to balance computational cost against resolution. A coarse pass over an 11 × 8 grid spanning $\bar{E} \in [150, 1,500]$ MWh and $\bar{P} \in [75, 325]$ MW (88 grid points) established the broad NPV landscape and located the region of highest NPV under degradation. A refined pass over a 7 × 6 grid spanning $\bar{E} \in [450, 750]$ MWh in 50 MWh steps and $\bar{P} \in [125, 250]$ MW in 25 MW steps (42 grid points) sampled the identified region at higher resolution, adding the 150 MW power level and the intervening energy capacities not present in the coarse grid. Where both passes cover the same $(\bar{E}, \bar{P})$ point, the refined run takes precedence. All runs use DK1 2022 prices, the 10–90% SoC window and the full 8,760-hour annual horizon with 20 project years.

Evaluating every grid point over the full 20-year horizon is expensive, so two filters reduce the coarse grid first. The first removes designs whose storage duration $\bar{E}/\bar{P}$ falls outside 1 to 8 hours. The upper bound is set well above the three to four hours that Martinez-Rico et al. [41] find to be the point beyond which added duration stops being cost effective in a day-ahead and reserve market, so it excludes only designs that are far outside the economically relevant range. The second filter runs before any

lifetime evaluation. A preliminary single-year sweep, costing one dispatch solve per point instead of a full 20-year loop, ranks the grid cheaply: a design is carried forward only if its single-year NPV reaches 65% of the highest single-year NPV on the grid, under at least one of the two degradation models. The single-year optimum under each scenario is kept regardless of both filters. The refined grid is not filtered.

The 65% threshold is chosen rather than derived, and the single-year NPV it screens on is an approximation of the lifetime NPV that is reported: it holds capacity constant and levelizes degradation instead of simulating replacement. The threshold is set wide so that only clearly uncompetitive designs are removed.

The optimal design point under each scenario is identified as the grid maximum. A quadratic fit,

$$\widehat{\text{NPV}}(\bar{E}, \bar{P}) = b_0 + b_1 \bar{E} + b_2 \bar{P} + b_3 \bar{E}^2 + b_4 \bar{E} \bar{P} + b_5 \bar{P}^2, \quad (3.16)$$

is fitted by nonlinear least squares over the refined grid region to refine the grid maximum to a continuous optimum $(\bar{E}^*, \bar{P}^*)$. The stationary point is accepted as a maximum if and only if the fitted Hessian is negative-definite. The quadratic fit is applied separately to the Xu and Shi degradation scenarios; it is not applied to the no-degradation scenario because the no-degradation optimum lies outside the refined grid boundary.

3.7.2. Stage 2: SoC Operating Window Sweep

Stage 1 fixes the size of the battery. Stage 2 fixes the size at the Xu grid optimum found there, whose coordinates are reported in Section 4.4, and varies the state-of-charge operating window to quantify its effect on degradation and NPV.

The window is parametrized by its width $\delta_{\max} = \sigma_{\max} - \sigma_{\min}$ and its center $(\sigma_{\max} + \sigma_{\min})/2$, both chosen before the year is solved. They are not the same as the cycle amplitude δ_i and cycle mean state of charge σ_i of Section 2.4, which the dispatch produces. The width caps the amplitude, and a cycle reaches that cap only if it spans the whole band. The center positions the band: shifting it at fixed width leaves the dispatch schedule unchanged and moves every state-of-charge level, and with it every σ_i and the profile mean $\bar{\sigma}$, by the amount of the shift.

The two parameters therefore act through different channels of the degradation model. Width acts on the amplitude stress $S_\delta(\delta_i)$, which appears in the cycle term alone. Center acts on the mean-SoC stress S_σ , which appears in both the cycle and the calendar term. Varying width and center independently across two series separates these two effects.

The *width series* fixes the center at 0.50 and varies the width, $\delta_{\max} \in \{0.40, 0.60, 0.80, 1.00\}$, corresponding to windows 30–70%, 20–80%, 10–90% and 0–100%. The *center series* fixes the width at $\delta_{\max} = 0.80$ and shifts the center through 0.60, 0.55, 0.50, 0.45 and 0.40, corresponding to windows 20–100%, 15–95%, 10–90%, 5–85% and 0–80%. The 10–90% window belongs to both series, at center 0.50 and width 0.80, and is the common reference linking them.

Changing the width changes the fitting range of the polynomial surrogate, so Φ_{shi} is refitted over $[\delta_{\min}, \delta_{\max}]$ before each simulation, with the lower bound $\delta_{\min}$ held at 0.15. The fit is accepted only if the fitted exponent satisfies $k_4 > 1$, the condition for Φ_{shi} to be convex (Section 2.7.1). A window with no convex fit is skipped, because the gradient-based method requires one. This excludes the 30–70% window, whose fit over $[0.15, 0.40]$ returns $k_4 = 0.908$, leaving seven simulated windows. The center series shares one fit across all five of its windows, because the width is fixed.

Each of the seven windows runs the full 20-year loop of Section 3.6.1 at the fixed sizing, producing the same outputs as Stage 1. All degradation quantities are reported with the physical reference model, so they are directly comparable across windows.

3.8. Towards a Gradient-Based Degradation-Aware Sizing

The parameter sweep of Section 3.7 finds the degradation-aware optimum by evaluating every point on a grid. Each point runs the full 20-year loop of Section 3.6.1, which needs 41 dispatch solves: one for the no-degradation case, reused across years because capacity does not fade, and one per year

for each of the two degradation branches. Evaluated without a screen, the 130 points of the two grids would cost about 5,300 linear programs. The filter of Section 3.7 reduces this by roughly a third, at the cost of one additional solve per coarse-grid point, but the order of magnitude is set by the grid rather than by the screen.

The limit is the grid rather than the cost of each design. A grid search evaluates m^q points for m levels in each of q design variables, so its cost grows exponentially with the number of design variables, while a gradient-based search grows roughly linearly when the gradient is available analytically [42]. Adding the two SoC window bounds to the two capacity variables, at roughly seven levels each, would take 2,401 designs and about 10^5 linear programs, against the 42 designs of the refined grid. A grid search also returns no sensitivity information, so it shows that degradation reduces the optimal size without showing which term is responsible.

The alternative is to search directly, using degradation sensitivity to guide the steps. Two routes are set out below. The first puts the degradation cost inside a single nonlinear program (NLP). The second keeps the dispatch LP separate and links it to a sizing search through an analytical degradation sub-gradient. The monolithic route comes first, because the property of the degradation model that makes it hard to solve is what motivates the second.

3.8.1. The Monolithic NLP Formulation

The monolithic formulation extends the dispatch LP of Section 3.4.1 by adding the degradation cost directly to the objective function. The decision variables remain the storage power p_t^s and stored energy e_t at each of the $T = 24$ hourly time steps of a single representative day, a horizon that keeps the problem small while preserving the structure of the degradation–dispatch coupling. The objective is

$$J(p^s, e) = \sum_{t=1}^T \lambda_t (p_t^s - \varepsilon p_t^c) \Delta t - \lambda_E^{\text{repl}} \bar{E} \sum_i n_i \Phi_{\text{shi}}(\delta_i) S_\sigma(\sigma_i), \quad (3.17)$$

which is maximized over p_t^s and e_t subject to the same storage dynamics, energy bounds, dispatch constraints and periodicity condition as the dispatch LP (Equations 3.2b–3.2i).

The degradation term accumulates the polynomial amplitude stress Φ_{shi} over all rainflow cycles $\{(\delta_i, \sigma_i)\}$ extracted from the current SoC trajectory e_t , weighted by the cycle multiplicity n_i and by the mean-SoC stress factor S_σ of Section 2.6.2. The temperature stress factor is unity under the constant-temperature assumption of that section and is omitted here, as it is from Equation 2.9.

The two terms of Equation 3.17 are in the same units and over the same horizon. The degradation sum is the fraction of battery life consumed during the day and $\lambda_E^{\text{repl}} \bar{E}$ is the cost of a full replacement, so their product is the euro value of the life consumed against which the day's revenue is traded.

The objective and its gradient are nonetheless both divided by the magnitude of the LP net revenue for the same day. The reason is how convergence is measured. The solver terminates on an absolute threshold applied to the norm of the Lagrangian gradient, so an objective expressed in euros carries its monetary scale into the reported optimality, and the same threshold is a stricter requirement on a low-revenue day than on a high-revenue one. Dividing by the LP net revenue brings the objective to order one on every day, which is what allows the single set of optimality thresholds used in Section 4.5 to mean the same thing throughout. The maximizer is unaffected, since the objective and the gradient are divided by the same positive constant.

The scaling does not change the relative size of the two terms. Across the days solved in Section 4.5 the degradation term is a median 2.3% of daily revenue, a property of the problem rather than of the formulation, and it bounds what the degradation term can contribute to the solution.

Figure 3.12 shows the monolithic NLP in the same XDSM convention as Figure 3.5, with two additions forced by the degradation term. A rainflow analysis block extracts the cycle set $\{(\delta_i, \sigma_i)\}$ from the stored-energy trajectory and feeds both the objective, which needs the cycles to evaluate the degradation cost, and the gradient. A gradient block then supplies the objective gradient to the optimizer, because a general-purpose NLP solver advances using derivatives rather than function values alone. That gradient is discontinuous wherever a small change in the dispatch reclassifies the rainflow cycles,

marked in the figure and explained in Section 3.8.2. This discontinuity is the one difference from the dispatch LP of Figure 3.5, whose objective and constraints are smooth throughout.




Figure 3.12: XDSM of the monolithic NLP. Inputs and the dispatch output follow the dispatch LP of Figure 3.5; the additional rainflow analysis and gradient blocks, and the discontinuity marked in the figure, are described in the text above and in Section 3.8.2.

The intended solution method is a general-purpose constrained NLP solver, here `scipy.optimize.minimize` [62] with the `trust-constr` method, which implements the trust-region interior-point algorithm of Byrd et al. [6]. No analytic Hessian is supplied, so the implementation builds a quasi-Newton approximation to it across iterations, and the linear constraints of Equations 3.2b–3.2i are handled through a barrier subproblem; both techniques are set out in Martins and Ning [42]. The convergence properties of such solvers are established for smooth objectives: the analysis assumes that the gradient varies continuously with the decision variables, so that curvature information accumulated across iterations remains valid near the current iterate. The solver is run with an initial trust radius of 0.1 rather than the default of 1.0, which limits how far the first steps move before curvature information has been built up, and with the step tolerance set to 10^{-12} and the iteration limit to 500, so that a solve which stops on step collapse has exhausted the trust region rather than met a loose tolerance.

The objective of Equation 3.17 carries the mean-SoC factor $S_\sigma(\sigma_i)$, but the gradient supplied to the solver is the amplitude sub-gradient of Section 2.8, which differentiates the amplitude-only cost of Equation 2.18 and holds S_σ fixed at its value for the current cycle set. The solver therefore receives a gradient that is exact in the amplitude term and omits the response of S_σ to a change in the cycle mean. The restriction is the one adopted in Section 2.8.4 for the sizing gradient, and for the same reason: it is the form for which the convexity result of Section 2.7.3 holds.

The omission is bounded, and the bound is a closed-form expression in the fitted parameters. The decision variables here are the stored energies e_t , so perturbing one moves only the cycles with a turning point at t , unlike the dispatch sub-gradient of Section 2.8.4, where the perturbed variable is a power and the shift propagates to every later step. For such a cycle the amplitude is a difference between the two turning points while the mean is their average, so the mean responds to the perturbation at half the rate; with the power law of Equation 2.16 the omitted term therefore stands to the retained one in the ratio

$$\frac{\Phi_{\text{shi}}(\delta_i) k_\sigma / 2}{\Phi'_{\text{shi}}(\delta_i)} = \frac{\delta_i k_\sigma}{2k_4}, \quad (3.18)$$

with $k_\sigma = 1.04$ (Table 2.2) and $k_4 = 1.179$ (Equation 3.10). It vanishes for shallow cycles and reaches 0.35 at $\delta = 0.80$, the widest cycle a 10–90% operating window admits. A spot comparison against a

central difference of Equation 3.17 on a solved day agrees with this: the two match to three decimals at shallow cycles and diverge at the turning points of deep ones by the amount the ratio predicts.

3.8.2. Non-Smoothness of the Rainflow Degradation Cost

The degradation term of Equation 3.17 is not differentiable in the decision variables. The property belongs to the cycle-counting step rather than to the way the problem is solved, so it is present in any formulation that places a rainflow-based degradation cost inside a solver.

Rainflow cycle counting maps a continuous SoC trajectory to a discrete set of half-cycles $\{(\delta_i, \sigma_i)\}$. This mapping is continuous but not always differentiable. When the dispatch trajectory is perturbed by a small amount, which turning points of the trajectory are paired into a cycle can change, so the cycle set before and after the perturbation is not the same set counted with slightly different amplitudes. This occurs when a local extremum in the trajectory crosses an adjacent turning point in the SoC sequence: two shallow half-cycles then merge into one deeper cycle, or a single deep cycle splits into two shallower ones. Wherever cycles are re-paired in this way the total degradation cost remains continuous, but its gradient with respect to e_t jumps, because the stress contribution of a cycle scales nonlinearly with its amplitude δ_i and the amplitudes themselves change abruptly when cycles are re-paired.

The convergence properties of gradient-based optimization algorithms are derived under the assumption that the objective and constraints are continuously differentiable [42]. That assumption does not hold here, so those properties do not apply. This does not mean a solver returns nothing on such a problem; gradient-based methods can make progress on non-smooth objectives in some cases. It means there is no basis for expecting them to, and no guarantee available to fall back on when they do not. Whether the limitation is material for this formulation is an empirical question, and Section 4.5 answers it.

3.8.3. A Gradient-Based Sizing Method

Two properties of the monolithic route argue for a different arrangement. The first is its size: one power and one energy variable per time step gives of order 10^4 variables at an annual horizon, in a nonlinear program rather than a linear one, and a nonlinear program of that size is considerably harder to solve. The second is the loss of differentiability of Section 3.8.2, which removes any guarantee that a gradient-based solve converges.

Separating the sizing decision from the dispatch removes the first difficulty entirely. The dispatch stays a linear program at full horizon length, and the search that responds to degradation acts on one or two sizing variables. It does not remove the second: rescaling the trajectory can re-pair the rainflow cycles, so the capacity gradient is a sub-gradient rather than a derivative. Shi et al. show that the degradation cost is convex in the SoC profile (Theorem 1, Section 2.7.3), which is what makes a sub-gradient method appropriate, but the theorem is stated for the SoC profile and not for the capacity, so convergence of the outer search would have to be established separately. The concept is shown in Figure 3.13.

At each sizing iteration the dispatch LP of Section 3.4.1 would be solved to optimality for a fixed battery size. It would produce the optimal SoC trajectory, from which rainflow counting and the degradation model return the accumulated degradation. The sizing search would set $\bar{E}$, receive the degradation outcome and update $\bar{E}$ in the direction that increases lifetime NPV.

The signal that drives the sizing search is the gradient of net present value with respect to the sizing variable, $dNPV/d\bar{E}$. This gradient has three contributions. The first is direct: a larger battery costs more to build, so capacity enters NPV linearly through capital expenditure. The second and third act through the dispatch, because changing $\bar{E}$ changes the optimal SoC trajectory. The second is the change in arbitrage revenue that follows. The third is the change in rainflow cycle statistics, and therefore in degradation and in the timing of battery replacement. Computing the gradient from the polynomial model avoids finite-differencing through the rainflow map, because the amplitude term is differentiated analytically. It does not make the gradient a derivative in the ordinary sense: the cycle set can change as the capacity changes, so what the expression returns is a sub-gradient, and the outer search would have to be one that accepts a sub-gradient rather than requiring a derivative.

Assembling the gradient cuts the same quantity a different way, by whether the dispatch is held fixed or




Figure 3.13: Conceptual gradient-based sizing method. The sizing search would set the energy capacity $\bar{E}$ and pass it to the dispatch LP, which builds the SoC trajectory, counts rainflow cycles and accumulates degradation. The total degradation would then be used to compute the capacity gradient $d \text{ NPV} / d \bar{E}$, which measures how lifetime NPV responds to a change in capacity; that value would drive the next update of $\bar{E}$, and the process would repeat.

allowed to respond. The *frozen-dispatch* sensitivity is the change in degradation cost with capacity at a fixed dispatch, with the corresponding revenue sensitivity given by the LP dual prices on the state-of-charge bounds, a direct by-product of the dispatch solve. The *re-dispatch* sensitivity is what remains: after a change in capacity the optimal dispatch shifts, and capturing that shift requires re-solving the dispatch LP and propagating the effect through the degradation model. Together the two constitute the full capacity gradient.

The frozen-dispatch component is implemented and verified here (Section C.1), as is the per-timestep degradation sub-gradient that the re-dispatch term would require (Section C.2). The re-dispatch component and a suitable step-size rule are the items that remain, and are identified as the main direction for future work in Chapter 5.

The same gradient extends to more than one design variable. The operating window enters the degradation model through the mean SoC factor S_σ and through the range of amplitudes it admits, both of which already appear in Equation 3.11, so a search over $\bar{E}$, $\sigma_{\min}$ and $\sigma_{\max}$ would require no further derivation.

In the absence of the completed sizing search, the parameter sweep of Section 3.7 approximates what such a search would return. Re-solving the dispatch LP at every grid point captures the re-dispatch effect numerically, which is the information the analytical re-dispatch component would supply, and locates the optimum to within the grid spacing. The sweep therefore also provides the reference against which a completed gradient-based method could be verified.

3.9. Modeling Limitations

The simplifying assumptions on which the methodology rests are stated in the sections where they are introduced, next to the models they qualify.

Interannual variability not captured.

The representative-year assumption excludes years with anomalously high or low wind resource or price volatility. A price year with larger spread between peak and off-peak hours would produce more

intensive cycling and higher annual degradation; a calmer year would produce the opposite. The sensitivity to price year is partially addressed by comparing 2019 and 2022 single-year results, but the multi-year loop uses only the 2022 tile.

ERA5 spatial resolution.

ERA5 provides wind speed at approximately 31 km horizontal resolution. The site coordinates are used to extract a single grid point, so sub-grid mesoscale variability and local flow features are not resolved. The hub-height extrapolation using the derived shear exponent ($\alpha = 0.17$) addresses the vertical wind speed profile, but does not compensate for spatial averaging within the ERA5 grid cell.

No mast or lidar measurement is available for this site, so the size and direction of the resulting bias cannot be established here. The consequence is visible in the yield: a hub-height mean of 8.19 m/s places the farm at or above rated power for 17.8% of the year and gives a net capacity factor of 43%, which is high for onshore operation and above the Danish fleet average. A resource error therefore shifts revenue levels throughout, and the absolute yield figures should be read as conditional on the reanalysis rather than as a site assessment. The design comparisons of Chapter 4 are less exposed, because degradation-aware and degradation-free cases are evaluated on the same generation profile.

C-rate effect on degradation not modeled.

The LP bounds charging and discharging power by the power capacity $\bar{P}$ but does not impose a separate C-rate limit. The physical reference model also neglects C-rate effects explicitly, on the basis that battery current is coupled to cell temperature and depth of discharge in the test conditions used for parameter calibration [63]. In practice, high C-rates accelerate degradation through lithium plating and increased internal resistance, mechanisms that neither the LP constraint set nor the Xu stress model captures. The dispatch schedule may therefore include high-rate events that a degradation-aware operational strategy would otherwise avoid.

Capacity fade only; efficiency and power fade not modeled.

Section 2.2 identifies two forms of degradation, capacity fade and power fade. This thesis models the first alone. The physical reference model tracks the loss of usable energy capacity and returns a State of Health defined on that basis [63], and the charge and discharge efficiencies are held constant at $\eta_{\text{in}} = \eta_{\text{out}} = 0.9539$ for the full battery lifetime. In a real cell the internal resistance grows with cycle count and elapsed time, so a few percentage points of round-trip efficiency are lost over a service life.

The omission works in one direction. A battery whose efficiency falls delivers less energy per unit charged, so revenue declines faster over the horizon than reported here, and the extra energy drawn to deliver the same output raises throughput and therefore cycle aging, a feedback the model does not carry. Every lifetime, revenue and net present value reported here is optimistic in that respect. The effect on the sizing comparison is smaller than the effect on the levels, because it applies to every design in the sweep rather than to one of them.

Average SoC as a passive dispatch output.

The SoC stress factor S_σ modifies both cycle and calendar aging rates as a function of the average SoC of each identified cycle and the average profile SoC over the operating period. In the current implementation, these averages are computed from the LP dispatch output and passed to the degradation model as inputs; they are not optimization variables. A degradation-aware operational strategy could in principle target a lower average SoC to reduce S_σ and slow aging. Within the framework built here that target can only be imposed from outside, by moving the bounds of the operating window, which is what Section 3.7 varies. The level the dispatch settles at inside those bounds is never a decision variable, so the mean SoC can be constrained but not optimized.

LMO chemistry.

The physical reference model parameters are calibrated for Lithium Manganese Oxide (LMO) cells. Modern grid-scale batteries predominantly use LFP or NMC chemistry, which have different aging mechanisms and longer cycle lives. The LMO parameterization is adopted because it is the only chemistry for which the full published parameter set is available from Xu et al. [63].

Chemistry enters only through the parameter values. Chemistry-dependent are the stress-factor coefficients, the calendar rate and the two fade-curve coefficients, and with them every absolute quantity

reported here: annual degradation, end-of-life year, optimal capacity and NPV. The split between calendar and cycle aging is chemistry-dependent for the same reason. A chemistry with longer cycle life would shift that split further toward calendar aging and reduce the degradation cost per unit of energy cycled, which would move the optimal capacity upward and shrink the inward shift reported in Chapter 4; the direction of that shift would be unchanged, its size would not.

Chemistry-independent is the structure built on those parameters. The accumulation rule, the dual-model architecture of Section 3.5.1, the convexity condition $k_4 > 1$ and both gradient derivations follow from the functional forms rather than the fitted values, and the non-differentiability of the rainflow map is a property of the counting algorithm alone. Applying the framework to LFP or NMC would require re-fitting the coefficients and the polynomial surrogate, and re-deriving nothing. The comparison between the Xu and Shi models also holds as it stands, because both use the same parameterization and the comparison isolates the accumulation rule rather than the parameter values.

Late-life capacity acceleration not modeled.

The SEI fade curve predicts approximately linear capacity loss at late life. It does not capture the capacity knee observed in lithium-ion cells at advanced degradation levels, where electrode particle cracking, lithium plating and electrolyte decomposition create positive feedback loops absent from the SEI growth model [48]. The cycle-life data underlying the amplitude stress function, and the $\pm 3\%$ capacity estimation error reported with the model, both correspond to 80% capacity retention [63]. Simulating to the 70% replacement threshold therefore applies the fade curve ten percentage points below the retention at which the model's accuracy is quantified. The curve is smooth and monotone over that interval, so the extension is a continuation rather than a change of regime, but no accuracy figure is available for it. The 80% crossing is reported alongside every lifetime for this reason. No published semi-empirical model structured compatibly with the rainflow framework used here exists for the post-knee regime in LMO or an equivalent grid-scale chemistry; physics-based models such as O'Kane et al. [48] capture particle cracking dynamics but require electrochemical parameters unavailable at the sizing stage.

Other hybrid-plant frameworks handle late life differently: HyDesign [44] switches to a second fade regime at a fixed capacity level, which at this site would take effect from the second simulated year. That regime carries no published calibration, so it is not adopted here.

No penalty for operating at the state-of-charge extremes.

The stress factor S_σ is evaluated at the mean state of charge of each counted cycle and at the mean of the profile, so a cycle is scored by where it sits on average and not by the levels it reaches. A cycle that touches a full or an empty cell therefore carries no additional degradation beyond what its mean implies. Xu et al. note that operation at low state of charge raises cell impedance and accelerates cycle aging, and that their model does not represent it [63]; the same holds at the upper bound, where lithium plating and electrolyte oxidation are not represented either. The consequence is one-directional at the lower bound: because S_σ decreases monotonically with mean state of charge and has no floor, the model reports a benefit from lowering the operating window that continues to grow as the window approaches the empty cell. Results at windows reaching 0% or 100% state of charge should be read with this in mind.

Single-cell parameters applied at pack scale.

The physical reference model is calibrated on single-cell cycling tests [63] and is applied here to a pack of several hundred megawatt-hours as though it aged as one large cell. Effects that appear only at pack scale are therefore absent: cell-to-cell variation in capacity and impedance, thermal gradients across a container, losses in balancing, and the failure of individual modules. Each of these makes an installed pack age faster and less uniformly than the model predicts, so the lifetimes reported here are optimistic in a direction the model cannot quantify. The framework itself is unaffected, since a pack-level parameterization would enter through the same coefficients.

4

Results

This chapter reports the degradation-aware analysis of the IEA Wind Task 50 hybrid wind and battery plant under the DK1 day-ahead electricity market. It is organized in five sections. Section 4.1 characterizes the two DK1 price years and states which one the sizing sweep uses. Section 4.2 verifies the dispatch and degradation implementation on one week of a solved year. Section 4.3 reports the behavior of the model at the reference battery over the 20-year horizon. Section 4.4 reports the parameter sweep, which delivers the sizing result. Section 4.5 reports the monolithic nonlinear program, which tests whether the degradation cost can be embedded directly in the dispatch and returns a negative result.

Unless a section states otherwise, every run uses an LMO battery at a 10–90% state-of-charge operating window, and replacement triggered at the end of any simulation year in which the state of health falls below 70.5%, that is, the 70% end-of-life threshold plus a 0.5 percentage point tolerance (Section 3.6.1). Deviations from this configuration, including the battery size, the price year and the operating window itself, are stated at the start of the section in which they occur. All degradation, and every lifetime, replacement year and value figure derived from it, is reported on the physical reference model of Xu et al. [63]. The polynomial model of Shi et al. [59] is the surrogate required by the gradient-based method of Chapter 3, which needs a convex degradation cost; it is run as a second branch in Section 4.4, and the two are referred to throughout as the Xu branch and the Shi branch, and the gradients built on it are verified in Appendix C.


4.1. Electricity Price Regimes

Every run in this chapter uses one of two DK1 day-ahead price years, 2019 and 2022. This section characterizes the two years as market inputs, before any model output is reported. Both years are used in the baseline comparison of Section 4.3; the parameter sweep of Section 4.4 uses DK1 2022 only, for the reason given in Section 4.3.3.

DK1 2019 is a year of low and stable prices. DK1 2022 is a year of high and volatile prices, driven by the reduction in European gas supply over 2021 and 2022 [16]. Table 4.1 lists the price statistics for both years, and Figures 4.1 and 4.2 show the annual price series and the distribution of the daily price spread. The mean price is 38.5 EUR/MWh in 2019 and 219.0 EUR/MWh in 2022, a factor of 5.7. The hourly standard deviation is 13.2 EUR/MWh in 2019 and 145.4 EUR/MWh in 2022, a factor of 11. The mean within-day price spread, the maximum minus the minimum price within a day, is 26.9 EUR/MWh in 2019 and 182.3 EUR/MWh in 2022, a factor of 6.8. The higher level and the wider swings persist across the whole year rather than arising from a few extreme days (Figure 4.1).

Table 4.1: DK1 day-ahead price statistics for 2019 and 2022. The daily spread is the within-day maximum minus minimum price.

Metric	DK1 2019	DK1 2022
Mean price (EUR/MWh)	38.5	219.0
Median price (EUR/MWh)	38.8	190.6
Standard deviation (EUR/MWh)	13.2	145.4
Minimum (EUR/MWh)	-48.3	-19.0
Maximum (EUR/MWh)	109.5	871.0
5th percentile (EUR/MWh)	15.8	21.4
95th percentile (EUR/MWh)	56.6	503.9
Negative-price hours	132.0	38.0
Mean daily spread (EUR/MWh)	26.9	182.3
Median daily spread (EUR/MWh)	24.0	165.0

**Figure 4.1:** Hourly DK1 day-ahead prices for 2019 and 2022. The dashed line marks zero price.

A battery earns arbitrage revenue by charging when the price is low and discharging when it is high within the same day. The margin available on each cycle is set by the within-day price spread, not by the price level. The spread is 6.8 times larger in 2022 than in 2019, and it is larger on nearly every individual day rather than only on average (Figure 4.2). Each cycle is therefore worth more in 2022, and the dispatch cycles more often and more deeply. More frequent and deeper cycling raises cycle aging (Section 4.3.2). The 2022 regime thus combines the largest arbitrage value per cycle with the largest cycle aging per cycle.

The wider spread in 2022 raises the value of moving energy across hours and drives the dispatch to cycle more, which in turn raises cycle aging and its cost. Under 2019 prices the spread is narrower, so added capacity earns less arbitrage revenue and is cycled less; storage remains useful in that regime, but at lower value. Repeating the full sweep under 2019 prices is listed in Section 3.9 as an extension of the numerical results.

One property of the 2022 data should be kept in view. It was an exceptional year: its spreads reflect a supply shock rather than typical operating conditions, while 2019 is closer to the range seen before 2021. Any result obtained with 2022 prices therefore describes a high-arbitrage environment and is not a forecast of expected revenue.




Figure 4.2: Empirical cumulative distribution of the within-day price spread for DK1 2019 and 2022. The dotted lines mark the median of each year (Table 4.1).

4.2. Model Verification

Chapter 3 specifies a model in which a linear dispatch program produces a state-of-charge trajectory, the rainflow algorithm converts that trajectory into a list of cycles, and a semi-empirical degradation model computes the degradation from that list. These three steps run in this order within each simulation year, and the simulation is repeated over a 20-year horizon (Section 3.6.1). A verification of the implementation is provided in this section.

The degradation model itself is not under test. The Xu formulation and its LMO coefficients are taken from the published work, where they were fitted and validated against cell measurements [63], and this thesis uses them as given. What this thesis adds is the application of that model to the output of a dispatch optimizer. Four properties of the implementation are verified here:

1. The state-of-charge trajectory analyzed below is the one stored by the baseline run of Section 4.3, so the checks apply to the results reported in the rest of this chapter and not to a separate solve;
2. The battery stays inside its energy and power limits and responds to the price signal;
3. The rainflow counting algorithm returns the cycles contained in the dispatched state-of-charge trajectory;
4. The degradation model correctly calculates degradation from the output of the rainflow counting algorithm as Equation 2.6 specifies, and the implementation responds to a change in the operating window as the equations predict.

The checks use the full solved year at five operating windows. One week of it is displayed, short enough that each counted cycle can be traced to the swing that produced it. The week is a slice of the first annual solve of the multi-year loop of Section 3.6, taken from the 2022 price year. Only that year is used below, and the properties verified concern the implementation and not the market data, so the choice of price year does not affect the conclusions. The complementary tests on synthetic input, where every expected value is available in closed form, are reported in Appendix D.

4.2.1. Verification Scope and Data

Year 1 of the DK1 2022 baseline was solved in full, over all 8,760 hours, at the reference battery of 300 MWh / 150 MW, on the ERA5 wind series and the DK1 day-ahead prices of Section 4.1. Five state-of-charge windows were solved: the 10–90% reference, the width series 20–80 / 10–90 / 0–100% at

center 0.50, and the center series 20–100 / 10–90 / 0–80% at width 0.80. The checks reported below therefore cover 43,800 hourly steps and the 4,883 cycles counted within them, of summed multiplicity $\sum_i n_i = 3,659.5$.

The year-1 trajectory analyzed here matches the trajectory stored by the baseline run element by element over all 8,760 steps, to within double precision, and a non-zero difference halts the procedure. Property 1 therefore holds: the trajectory analyzed in this section is the one the baseline run produced. The inter-year pinning of Equation 3.6 does not apply here, because this is the first year of the battery's life and the optimizer selects its starting level freely.

The degradation figures reported below are recomputed from the published cycle list rather than read from the run. For each window the cycle term is re-evaluated as $\sum_i n_i S_\delta(\delta_i) S_\sigma(\sigma_i)$ over the counted cycles, and the calendar term as $k_\tau t S_\sigma(\bar{\sigma})$. The stress coefficients are taken from the printed values in Xu et al. [63] rather than from the degradation modules, so that a wrong coefficient would appear as a disagreement instead of canceling on both sides. The two routes are independent implementations of the same equation and agree to a relative difference of 4×10^{-13} , the level at which double-precision arithmetic differs between two orderings of the same sum. This confirms that Equation 2.6 is correctly implemented in the degradation modules. The cycle list is published with the run, so this comparison can be repeated by a reader who has neither the solver nor the input data.

The displayed week is the one whose mean daily price spread is closest to the annual median, a rule fixed before the results were seen. It is week 20, hours 3,192 to 3,360, with a mean daily spread 2.0% below the annual median. Its cycling is slightly above average for the year, by 4% in summed cycle multiplicity and by 14% in energy moved, so the week is not an unusually quiet one for the counter or the degradation model.

Cycle counting is run once on the full year, exactly as Sections 4.3 and 4.4 count it. Each cycle is then assigned to the week containing its midpoint, so the weekly sets partition the annual set with nothing dropped and nothing counted twice. Counting a week in isolation is not equivalent, because it cuts cycles that straddle the week boundary and closes residue half-cycles at the edges that the annual count never produces.

4.2.2. Dispatch and Cycle Counting

Three checks are made on the displayed week: that the dispatch stays inside its limits and responds to the price, that the battery losses are applied on the correct side of the storage dynamics, and that the rainflow algorithm returns the cycles contained in the resulting trajectory.

Figure 4.3 shows the full 168 hours at the reference window: the price, the wind and battery power, and the state of charge. The battery starts the week at 0.85 state of charge and discharges within two hours. This week is part of the annual simulation, so its starting state of charge is set by the earlier hours of the year, which are not shown.

The dispatch follows the price and stays within its rated limits. The battery charges when the price is low and discharges when it is high: the arithmetic mean price over the 66 charging hours is 162.8 EUR/MWh and over the 25 discharging hours 235.5 EUR/MWh, a ratio of 1.45 against the round-trip threshold of 1.10 set by the efficiency of Section 3.4.3. The mean prices therefore differ by more than the round-trip loss, which is the response the objective of Equation 3.2a requires. The state of charge stays between 0.10 and 0.90 at every step and reaches both values, so the limits are enforced and both are reached. Peak power is 150 MW in both directions, with no step above the rating. Wind output peaks at 325 MW, the installed rating of the 65 turbines, which equals the grid connection capacity, so the export limit is not binding at any step. These results are the expected behavior for property 2 and confirm that the energy and power constraints of Section 3.4.3 are correctly implemented.

The limits and the price response are therefore correct. The next check is narrower and concerns the battery losses. One hour of charging at the rated 150 MW adds $\eta_{\text{in}} \bar{P} \Delta t = 150 \times 0.953939 = 143.1$ MWh to the stored energy, which is 0.4770 of the 300 MWh capacity. Applying the discharging efficiency instead would give 157.2 MWh, or 0.5241. Ahead of the price peak at hours 101 to 103 the state of charge remains at 0.4230 through hours 96 to 99 and then rises to the 0.90 ceiling in the single hour ending at hour 100, leaving a gap below the ceiling of 0.4770. The same gap appears at the other two




Figure 4.3: One week of the solved year-1 dispatch, 300 MWh / 150 MW battery, 10–90% window, DK1 2022 prices. Top: day-ahead price. Middle: wind output, with battery discharge positive and charge negative. Bottom: state of charge, with the minimum and maximum allowed values dashed. Week 20, hours 3192 to 3360, all 168 hours shown.

width-series windows, at resting levels of 0.3230 and 0.5230 below ceilings of 0.80 and 1.00. At all three windows the gap matches the charging value to within 6×10^{-17} and not the discharging value. This shows that Equation 3.4 is correctly implemented, and that the charging inequality of Equation 3.2b holds with equality at that hour, so the relaxation reproduces the exact dynamics there.

The dispatch is therefore the one Section 3.4.3 specifies. What remains is whether the cycles counted on that dispatch are the ones the trajectory contains. Figure 4.4 places the counted cycles on the same trajectory. Following Section 2.4, each cycle is bounded by two turning points and carries a multiplicity n_i , equal to 1 for a full cycle and 0.5 for a half-cycle, and it is n_i rather than the number of entries in the cycle list that enters the degradation sum of Equation 2.6. In the lower panel each cycle is drawn as a vertical bar spanning the two state-of-charge levels its turning points join, so the bar height is the amplitude δ_i and the dot is the mean state of charge $\bar{\sigma}_i$. These are the two stress inputs of the degradation model. Both panels share the state-of-charge axis, so a bar can be matched to the swing above it by its position in time and by the levels it spans.

No horizontal extent is drawn, because a cycle is a pair of turning points and not an interval: one half-cycle in this week joins hour 65 to hour 163, and a horizontal span would give it a duration it does not have. A shallow cycle extracted from inside a deeper one appears as a short bar at the state of charge at which it occurred, above or below a taller one.

The counter returns 19 cycles in the week at the reference window, 12 full and 7 half, so $\sum_i n_i = 15.5$. Twelve of the 19 span the full window, at an amplitude of 0.80, and those twelve are 5 full cycles together with all 7 half-cycles. Of the remaining seven, three have numerically zero amplitude and four lie between δ_c and 0.34, so 16 of the 19 are deeper than δ_c . A further 3 records have their midpoint in an




Figure 4.4: Rainflow cycles of week 20, 10–90% window, all 168 hours. Top: state-of-charge trajectory, with the minimum and maximum allowed values dashed and the counted turning points marked. Bottom: one bar per counted cycle, spanning the state-of-charge levels its two turning points join and placed at their midpoint in time. Cycles of zero amplitude appear as a single dot. Three cycles with their midpoint in an adjacent week are not shown. The scale bar at the right gives $\delta_c = 0.1437$.

adjacent week and are counted there. Applying the three-point rule of Section 2.4 to the published trajectory in a separate implementation reproduces all 19 records, with the same amplitudes, mean states of charge and multiplicities. This shows that the rainflow counting algorithm is correctly implemented, which is property 3.

Whether a swing is recorded as a full cycle or as a half-cycle is set by the counting rule and not by its amplitude, which is why the twelve records at 0.80 split into 5 full and 7 half. Under the three-point rule a pair of turning points is removed as a full cycle when the trajectory returns to the level it left, and what cannot be paired in this way is reported as a half-cycle. Hours 21 to 55 show both outcomes in one event. The state of charge rises from 0.10 to 0.90, falls to 0.67, returns to 0.90 and falls to 0.10. The interruption at 0.67 returns to the level it left, so it is removed as one full cycle of amplitude 0.230 at a mean state of charge of 0.785, which is the short bar near 0.78 in Figure 4.4. Removing that pair leaves a single traverse from 0.10 to 0.90, reported as one half-cycle of amplitude 0.800, which is the tall bar beside it. The four monotonic segments therefore produce two records rather than four, as the three-point rule specifies.

A half-cycle, with $n_i = 0.5$, is a range traversed once rather than a full charge and discharge, so two half-cycles at the same amplitude contribute the same degradation as one full cycle. The number of cycles the counter returns is therefore not the number of charge and discharge cycles the battery performed, and only $\sum_i n_i$ enters the degradation model. Properties 1 to 3 therefore hold on the reference window, and the degradation sum itself was confirmed in Section 4.2.1. What remains of property 4 is the response to a change in the operating window, which the next subsection applies.

4.2.3. Response to the Operating Window

Property 4 requires the implementation to respond to a change in the operating window as the equations predict. The window is widened at fixed center and then shifted at fixed width. Each change fixes in advance what must stay the same and what must change, and by how much, which is what makes the response testable.

Figure 4.5 shows the battery dispatch and the counted cycles for operating windows of 20–80%, 10–

90% and 0–100%, on week 20 of DK1 2022. The first 72 hours of the trajectory are shown, because three overlaid trajectories are not readable at 168 hours. The cycle counts cover the full week.




Figure 4.5: Effect of the state-of-charge window width, at fixed center 0.50, 300 MWh / 150 MW battery. Left: state-of-charge trajectory for each window, with that window's bounds dashed in the same color, first 72 hours. Right: summed multiplicity $\sum_i n_i$ of the cycles counted over the full week, binned by amplitude, with the three windows drawn side by side within each bin so that coincident values stay visible. The lower panel magnifies the range below $\delta = 0.42$, where every bar carries $\sum_i n_i \leq 1$; triangles mark the three amplitudes at which all three windows agree. Dotted line: $\delta_c = 0.1437$.

The formulation predicts three things about widening the window, one for each way a discharge can end.

1. A discharge that ends on the lower bound can go no further, so its amplitude must equal the window width. These are the tall bars of the histogram, 13, 12 and 13 records, and each sits at its own window width.
2. A discharge that ends on the price stops while energy remains, so no bound enters the amount moved and the same discharge must have the same amplitude in every window. Three occur in the week, marked with triangles in the magnified panel, and their amplitudes agree across all three windows to six decimal places. Their mean state of charge does move, by 0.10 for each widening, because widening raises the ceiling and lowers the floor by that amount and an event anchored at one bound follows it. An identical swing therefore carries different degradation in each window, through the mean-state-of-charge stress factor alone.
3. A discharge limited by the upper bound has a closed form. Ahead of the price peak at hours 101 and 102 the model buys a full rated hour at hour 99, so it must leave one rated hour of space. Two rated hours are η_{in} of capacity at this two-hour battery, where $2\bar{P}\Delta t = \bar{E}$, so a window of width δ_{max} must sell $|\delta_{max} - \eta_{in}|$ at hour 95, or buy if it has room to spare. The two narrower windows sell that amount and the widest buys. This is the strongest of the three checks, because the predicted amount combines the width, the power rating and the charging efficiency, and an error in any one of them would break the match. The case arises only because the dispatch has perfect price foresight (Section 3.4).

Two further groups of records complete the week. The cluster at zero amplitude, one to three records per window, holds the phantom micro-cycles of Section 3.4.3, produced by slack in the state-of-charge

balance and carrying no degradation. The two records near 0.90 belong to the 0–100% window alone, where the floor is out of reach: one begins at a full battery and stops on the price, the other ends on the floor. Together with the three predictions these account for every record counted in the week, 20 at 20–80%, 19 at 10–90% and 19 at 0–100%, so the response is verified across the whole week and not only at the hours named above.

Shifting the window at fixed width must not change the dispatch at all. Neither the objective of Equation 3.2a nor the power constraints contains an absolute state of charge, so the formulation requires an unchanged power schedule and a trajectory offset by exactly the shift. Nothing in the implementation enforces this, which makes it an independent check on the code.

The schedules agree at every step for the upward shift and at 8,758 of 8,760 steps for the downward one, and total energy moved is identical in all three at 231,638.15 MWh. The two steps that differ lie at hours 8,527 to 8,530, near the end of the year, where the periodicity constraint of Equation 3.2i ties the final state of charge back to the first and the solver returned a different schedule with the same objective value. This is expected, because the formulation fixes the optimal value and not the choice among the schedules that achieve it.

The degradation must then scale by $\exp(k_\sigma \Delta \bar{\sigma})$, since the mean state of charge is the only stress input that moves, and it does to six significant figures. The residual comes from the zero-amplitude records, whose number changes with the center and which carry no degradation. This confirms that the mean-state-of-charge stress factor of Equation 2.6 is correctly implemented.

Table 4.2 collects the six requirements the formulation imposes on a change of window, together with the values the runs returned.

Table 4.2: Response of the implementation to a change in the operating window. Predictions are fixed before the runs. Width-series entries are given in the order 20–80, 10–90, 0–100%; center-series entries in the order 20–100, 0–80% against the 10–90% reference. The entry “none” marks the widest window, which buys rather than sells at hour 95.

Requirement	Predicted	Observed
Amplitude of a bound-limited discharge equals the window width	$\delta_{\max}$	0.60, 0.80, 1.00
Amplitude of a price-limited discharge does not depend on the width	identical	identical to six decimals
Amount sold at hour 95 equals $ \delta_{\max} - \eta_{\text{in}} $	0.354, 0.154, none	0.354, 0.154, none
Power schedule unchanged by a center shift	unchanged	8,760 and 8,758 of 8,760 steps
State-of-charge offset equals the shift	+30, –30 MWh	exact except hours 8,527–8,530
Degradation ratio equals $\exp(k_\sigma \Delta \bar{\sigma})$	1.10960, 0.90123	1.10960, 0.90121

The checks above establish that the response is correct. Table 4.3 gives its size. Widening from 20–80% to 0–100% raises the energy moved by a factor of 1.47 and the week’s cycle degradation by a factor of 2.77, while lowering the summed multiplicity to 0.81 of its value. Per equivalent full cycle the widest window’s cycle degradation is 1.89 times that of the narrowest. That factor isolates the superlinearity of the amplitude stress function from the amount of energy moved. The calendar term changes little, from 2.437×10^{-4} to 2.381×10^{-4} , because the center is fixed. These are week values, and scaling them to a year overstates the annual degradation by 3.3%, so the comparisons made here are between windows on the same week.

Table 4.3: Week statistics for the width series, at fixed center 0.50, battery 300 MWh / 150 MW, week 20 of DK1 2022. $\sum_i n_i$ is the summed cycle multiplicity. EFC is the discharge throughput divided by nominal capacity, $EFC = \sum_t \max(p_t^s, 0) \Delta t / \bar{E}$, as used for the annual figures elsewhere in this chapter. $\bar{\delta}$ is the mean amplitude over the counted cycles. Degradation columns are the week's contribution on the Xu reference branch.

Window	Records (full + half)	$\sum_i n_i$	EFC	$\bar{\delta}$	f_d^{cyc} (10^{-4})	f_d^{cal} (10^{-4})	f_d (10^{-4})
20–80%	11 + 9	15.5	6.332	0.409	1.772	2.437	4.209
10–90%	12 + 7	15.5	7.859	0.506	2.783	2.416	5.198
0–100%	6 + 13	12.5	9.274	0.738	4.906	2.381	7.287

The check was made on a week rather than on a year because errors in individual cycles can cancel in an annual total. None of the responses verified above is visible in an annual figure: the amplitude of a bound-limited discharge, the invariance of a price-limited one, and the rated hour of space left below the ceiling. At week scale each cycle can be matched to the dispatch that produced it.

4.3. Baseline Behavior at the Reference Design

The baseline runs use the IEA Wind Task 50 reference battery of 150 MW / 300 MWh at a fixed 10–90% state-of-charge window. Both price years, 2019 and 2022, are compared (Section 4.1). The purpose is to establish the magnitude of degradation at a fixed design, and its split between cycle and calendar aging, before the design space is explored.

4.3.1. State-of-Health Trajectory

The 20-year state-of-health trajectory is built by the multi-year simulation loop of Section 3.6.1. At the start of each project year k the inner LP is re-solved at the effective capacity $\bar{E}_k = \bar{E} \cdot \text{SoH}_{k-1}$. The annual degradation increment $f_{d,k}$ is computed from the resulting dispatch and used to calculate the SEI capacity loss and update the state of health.

Figure 4.6 shows the evolution of the state of health over the project lifetime. Under DK1 2022 prices the battery loses 7.95 percentage points of capacity in the first year, from 100% to a year-1 state of health of 92.05%, at a degradation rate $f_{d,1} = 0.02624$ (51% cycle, 49% calendar). After the first year the state of health declines steadily, and the annual loss eases from 2.6 percentage points in the second year to 2.0 in the year before replacement. The trajectory reaches 80% at year 6.14. It does not reach 70% while the first battery is installed: replacement is triggered at 70.5%, so the battery is replaced at the end of year 11 at a state of health of 70.05%. The second battery repeats the same trajectory and reaches year 20 at 74.0%.

Under DK1 2019 prices the first-year loss is 7.84 percentage points, to a year-1 state of health of 92.16%, at $f_{d,1} = 0.02534$ (50% cycle, 50% calendar), a total rate 3.4% below the 2022 case. The later decline has the same shape at a slightly lower rate. The trajectory reaches 80% at year 6.37 and 70% at year 11.44, so replacement falls at the end of year 12, at 68.95%. The second battery reaches year 20 at 76.7%.

Both features of the trajectory follow from the SEI fade function of Equation 2.7, which maps the cumulative degradation f_d to capacity loss through two exponential terms with $\alpha_{\text{sei}} = 0.0575$ and $\beta_{\text{sei}} = 121$ for the LMO cell used here. Over one battery life the input f_d runs from 0 to 0.297 under 2022 prices and from 0 to 0.313 under 2019 prices, so the trajectory is set by the shape of the function over that interval.

The two terms act on different scales. The first decays over $1/\beta_{\text{sei}} = 0.0083$, which is about one third of the first-year increment $f_{d,1} = 0.0262$. This term has lost 96% of its value by the end of year 1 and contributes 5.5 of the 7.95 percentage points of capacity lost in that year, 0.2 points in year 2, and less than 0.01 points from year 4 onward. The large first-year loss therefore comes from this term, which represents the initial SEI film formation. The framework reproduces this effect, so the first-year loss is expected model behavior rather than a numerical artifact. After the first year only the second term is active. Its slope $-(1 - \alpha_{\text{sei}})e^{-f_d}$ falls from -0.894 at the end of year 2 to -0.701 at the end of year 11, a reduction of 22% over the operating range of f_d , while the annual increment $f_{d,k}$ grows by 3.6%

over the same span. The combination gives the observed decline in annual capacity loss from 2.6 to 2.0 percentage points, so within each battery life the trajectory is steepest just after the first year and shallowest just before replacement.

For this site and duty the battery reaches the replacement threshold in 11 to 12 years, so a single replacement is required within the 20-year project under both price regimes.

The two price years differ in degradation because they differ in dispatch. The wider 2022 spreads make more arbitrage trades profitable, so the battery completes more cycles and holds a higher time-averaged state of charge; Section 4.3.2 separates these two contributions and shows that 81% of the difference comes from the cycle component. The size of the difference is small: the annual degradation rate is 3.4% higher under 2022 prices, the year-1 state of health differs by 0.1 percentage points, and replacement moves by one year. The one-year shift is a threshold effect rather than a physical gap. Under 2019 prices the state of health at the end of year 11 is 70.82%, 0.32 percentage points above the 70.5% replacement trigger, and 1.6% more cumulative degradation would take it below. That margin is smaller than the 3.4% difference in annual rate between the two price years, so replacement timing at this design is sensitive to inputs at the level of the difference the price regime itself produces.

The comparison is therefore asymmetric. The price regime changes the battery-marginal net present value by 132.6 MEUR and reverses its sign (Section 4.3.3), while changing the annual degradation rate by 3.4%. At a fixed design the price regime is a first-order driver of revenue and a second-order driver of degradation, because the dispatch can only respond to a wider spread within a capacity it cannot change. Whether this holds once capacity is free to vary is the question the sizing sweep of Section 4.4 addresses.




Figure 4.6: State-of-health trajectory over the 20-year project for the 150 MW / 300 MWh reference battery, under DK1 2019 and DK1 2022 prices. The vertical segment in each line is the battery replacement, where the state of health returns to 100%. The dashed lines mark the 80% (IEC/EV) and 70% (warranty) end-of-life levels.

4.3.2. Degradation Components: Cycle and Calendar Aging

Table 4.4 reports the year-1 degradation components and dispatch statistics for both price years, at the 150 MW / 300 MWh reference battery and the 10–90% state-of-charge window.

The table reports three different cycle counts, each measuring a different quantity. Under 2019 prices the rainflow counter returns 916 records, 518 of them full cycles and 398 half cycles. Weighting a full cycle 1 and a half cycle 0.5 gives the summed multiplicity $\sum_i n_i = 717.0$, which is what enters the degradation model. The equivalent full cycles, 339, are the discharge throughput divided by the nominal capacity, so they measure energy moved rather than the number of excursions. The mean cycle amplitude is 50.0%, so each counted cycle moves on average half the energy of a full charge

and discharge. This is why the equivalent full cycles are near half the summed multiplicity. The two counts are computed independently and do not reconcile term by term. Table 4.3 shows the same three quantities at week scale, and Appendix D sets out the full and half cycle convention.

The two components are close to equal in the first year. Under 2019 prices the split is 50.0% calendar and 50.0% cycle; under 2022 prices the cycle component is already the larger of the two, at 51.1% against 48.9%. Calendar aging reaches this share because the battery moves little energy, 339 equivalent full cycles per year under 2019 prices and 361 under 2022, in both cases fewer than one full cycle per day. Calendar aging in this model is a function of the age of the battery, the mean state of charge and the cell temperature (Equation 2.15), so at a fixed operating window and a fixed temperature it takes a similar value in both years by construction. The cycle component then grows with capacity fade while the calendar component does not, so cycle aging becomes the larger of the two from year 2 onward under 2019 prices and from year 1 under 2022 prices.

The two price years differ mainly in how much the battery cycles. The 2022 dispatch runs 8.3% more counted cycles and 6.3% more equivalent full cycles than the 2019 dispatch. Both follow from the higher price volatility of that year (Section 4.1). The two percentages differ because the 2022 cycles are on average shallower. This difference has a limited effect on the remaining statistics. The mean cycle amplitude changes by 1.2 percentage points, from 50.0% to 48.8%, and the time-averaged mean state of charge by about 1.3 percentage points, from 47.2% to 48.4%. The higher mean state of charge under 2022 prices raises the calendar component, but only by 0.00017, or 1.3%, against 5.8% for the cycle component. Of the difference in the annual degradation rate between the two years, 0.02624 against 0.02534, 81% comes from the cycle component.

Table 4.4: Year-1 degradation components and dispatch statistics, 10–90% state-of-charge window, 300 MWh LMO battery. Cycle quantities are defined in Table 4.3. The two components are rounded independently and need not sum to the rounded total.

Metric	DK1 2019	DK1 2022
f_d	0.02534	0.02624
f_d^{cyc}	0.01267 (50.0%)	0.01340 (51.1%)
f_d^{cal}	0.01268 (50.0%)	0.01284 (48.9%)
Year-1 state of health	92.2%	92.0%
Cycles (full + half cycles)	518 + 398	573 + 407
Summed multiplicity $\sum_i n_i$	717.0	776.5
Equivalent full cycles	339	361
Mean cycle amplitude $\bar{\delta}$	50.0%	48.8%
Mean state of charge $\bar{\sigma}$ (time average)	47.2%	48.4%

Figure 4.7 shows the two components in absolute terms for each battery during the project lifetime. The calendar component is flat, changing by less than 0.5% between the first and last year of a battery life, while the cycle component grows by about 10% over the same span. Both return to their year-1 values at replacement, in year 11 under 2022 prices and year 12 under 2019 prices.




Figure 4.7: Annual degradation rate f_d split into cycle and calendar components over the 20-year project, for the 150 MW / 300 MWh reference battery. For each project year the left bar is DK1 2019 and the right bar is DK1 2022.

Figure 4.8 expresses the same result as a share of the annual total, which is where the trend is visible. The share is not constant. It starts at 48.9% under 2022 prices and 50.0% under 2019 prices and falls by about 2.5 percentage points over the life of a battery, so the calendar share is below an equal split from year 1 under 2022 prices and from year 2 under 2019 prices. The decline comes entirely from the growth in the cycle component. That growth follows from capacity fade rather than from any change in dispatch: as the effective capacity falls, the same energy throughput corresponds to a larger cycle amplitude δ , and the per-cycle Xu stress $S_\delta(\delta)$ increases with amplitude.




Figure 4.8: Calendar share of the annual degradation rate f_d over the 20-year project, for the 150 MW / 300 MWh reference battery under DK1 2019 and DK1 2022 prices. The cycle share is the complement of the plotted value. The dashed line marks an equal split.

These results carry two implications for hybrid power plants that operate storage on a day-ahead market. First, reducing the number of cycles has a limited effect on total degradation for this duty. A 6.3% difference in equivalent full cycles between the two price years changes the annual degradation rate

by 3.5%, because the change acts on only one of two components of comparable size. Levers that act on the calendar component, chiefly the mean state of charge, have at least comparable reach, and the window sweep of Section 4.4.2 measures that reach directly. Second, the balance between the two components is a property of how the battery is operated rather than of the battery itself. A plant that completes more cycles per year would move the balance toward the cycle component, and the values reported here apply to the low-cycle end of that range. One condition bounds both statements: the cell temperature is fixed at 25 °C in these runs. A warmer installation would raise the calendar component and increase its share further.

4.3.3. Price Sensitivity of Lifetime, Revenue and Net Present Value

Table 4.5 summarizes the end-of-life years and the project-level economics under both price regimes. The 2019 scenario gives a longer battery lifetime at both levels: 0.23 years longer at 80% and 0.41 years longer at 70%. Replacement occurs one year later, at year 12 rather than year 11.

All net present values and revenues reported here and in Sections 4.4 and 4.5 are battery-marginal. The battery-marginal revenue is the plant revenue with the battery less the plant revenue without it. The battery-marginal net present value is that revenue stream discounted over the 20-year horizon, less the battery capital cost and the discounted replacement cost. It is therefore a net figure and not a capital cost. Plant revenue appears only in Table 4.5, and only to show the scale the battery works against: adding the battery raises annual plant revenue by 3.5% in 2019 and by 6.3% in 2022. Two net present values are reported for each price year. The first holds capacity at its initial value and charges no replacement; the second carries the simulated fade and the replacement it triggers. Their difference is the lifetime cost of degradation, and that is the term used for this quantity throughout the chapter.

On this basis the battery-marginal net present value is –89.7 MEUR under 2019 prices and +42.9 MEUR under 2022 prices, a swing of 132.6 MEUR and a change of sign. At this design the battery does not recover its cost in the 2019 regime and does in the 2022 regime.

Table 4.5: End-of-life years and project economics under DK1 2019 and 2022 prices, 10–90% window, 300 MWh LMO battery. End-of-life years are the crossings of the simulated state-of-health trajectory, interpolated between year-end values. The 70% year quoted for 2022 is the year the level would have been crossed had the battery remained in service, since replacement is triggered at 70.5%.

Metric	DK1 2019	DK1 2022
End of life at 80% state of health	6.37 yr	6.14 yr
End of life at 70% state of health	11.44 yr	11.03 yr
Replacement year	yr 12	yr 11
Final state of health at year 20	76.7%	74.0%
Present value of replacement cost	25.25 MEUR	26.01 MEUR
Year-1 plant revenue	50.3 MEUR	208.1 MEUR
Year-1 battery revenue	1.71 MEUR	12.26 MEUR
Battery net present value (NPV), degradation excluded	–62.0 MEUR	89.2 MEUR
Battery NPV, degradation included	–89.7 MEUR	42.9 MEUR
Lifetime cost of degradation	27.8 MEUR	46.3 MEUR

The one-year shift in replacement timing has a modest economic consequence. Discounting at a 3% real rate, a replacement at year 11 carries a present-value factor of about 0.722 and a replacement at year 12 about 0.701, so the discounted replacement cost is 0.76 MEUR higher under 2022 prices. The dominant difference between the two regimes is not degradation but revenue: the wide 2022 spreads let the LP execute more profitable arbitrage, which raises both revenue and the number of cycles. A negative battery net present value means the arbitrage revenue never recovers the capital. Under 2019 prices the discounted battery revenue over the horizon is 21.9 MEUR against 86.4 MEUR of capital, so the battery returns about a quarter of what it costs to build, before the 25.3 MEUR of discounted replacement cost is counted. Under 2022 prices the same battery returns 155.3 MEUR, or 1.8 times its capital cost. Storage at this design is therefore not a viable investment in the 2019 regime.

This result determines the price year used for the sizing study. The parameter sweep of Section 4.4

uses DK1 2022 for two reasons. First, it is the only one of the two price years in which the battery recovers its cost at the reference design, so it is the only one in which a cost-recovering optimal size exists. Second, it is the demanding case for the research question: the wider spread drives more cycling, so the degradation cost attached to each unit of added capacity is largest there, and any effect of degradation on the optimal size is most visible.

This comparison shows the tension that the degradation-aware framework addresses. The dispatch that maximizes short-term revenue is also the one that deepens cycling and advances the replacement date. In the reference case that tension is mild, because one price year moves the replacement date by a single year. It becomes the governing effect once the battery is sized larger.

One qualification carries beyond this site. The same hardware, the same site and the same degradation model give an investment case of opposite sign in the two price years, so the sizing results that follow are conditional on the price regime assumed and are not a property of the battery. Two conditions bound them. The 2022 price series describes an exceptional year rather than typical operating conditions (Section 4.1), and the dispatch is solved with perfect price foresight (Section 3.4), so the arbitrage revenue is an upper bound. Both act in the same direction: under narrower spreads or imperfect foresight, the storage capacity that recovers its cost would be no larger than the size reported in Section 4.4, and probably smaller.

4.4. Parameter Sweep

This section reports the two-stage parameter sweep of Section 3.7. Stage 1 (Section 4.4.1) evaluates 20-year lifetime net present value across a two-dimensional grid of energy capacity $\bar{E}$ and power capacity $\bar{P}$, under three scenarios: no degradation, the physical reference model and the polynomial model. Stage 2 (Section 4.4.2) fixes the battery at the Stage 1 optimum and varies the state-of-charge window to separate the effect of window width from that of window center. Both stages use the full 8,760-hour annual horizon repeated over 20 project years, with the multi-year replacement loop of Section 3.6.1, and use DK1 2022 prices (Section 4.1). The two stages share a common design point, the 550 MWh / 175 MW battery at a 10–90% window, which is the Stage 1 Xu optimum and the center point of the Stage 2 window series.

4.4.1. Energy and Power Capacity Sweep

Figure 4.9 reports lifetime NPV over the coarse $(\bar{E}, \bar{P})$ grid, with energy capacity on the vertical axis, power capacity on the horizontal axis and NPV as the contour value. Panel (a) is the no-degradation scenario and panel (b) the Xu branch. Without degradation the grid optimum is 800 MWh and 250 MW, at 134.2 MEUR; on the Xu branch it is 600 MWh and 175 MW, at 55.0 MEUR. At 1,300 MWh the Xu NPV falls to between –2.2 and 3.5 MEUR depending on power capacity, so a battery of that size returns no more than it costs.

The coarse grid is too widely spaced to fix the degradation-aware optimum. Energy capacity is sampled at 150 MWh and power at 25 MW, so a design 50 MWh from the reported optimum could not be distinguished from it. The dashed rectangle in Figure 4.9 marks the region resampled at 50 MWh to settle this.




Figure 4.9: Lifetime NPV across the full $(\bar{E}, \bar{P})$ grid, coarse pass, DK1 2022 prices. The star marks the NPV-maximizing design in each panel and the diamond marks the no-degradation optimum; the two coincide in (a). Gray dots are the evaluated grid points and the dashed rectangle marks the region covered by the refined pass. The uncolored corner regions are designs removed by the duration and value filters of Section 3.7.1 and not evaluated. The two panels use independent color scales, so a given color does not carry the same NPV across panels.

Figure 4.10 reports lifetime NPV over the refined grid, which samples this region at 50 MWh spacing, on the Xu branch (a) and the Shi branch (b). On both branches the grid optimum is 550 MWh / 175 MW, at 55.37 MEUR on the Xu branch and 56.48 MEUR on the Shi branch. A quadratic fit (Equation 3.16) over the refined region gives a continuous optimum of 560 MWh / 182 MW on the Xu branch ($R^2 = 0.970$) and 560 MWh / 183 MW on the Shi branch ($R^2 = 0.971$), with a negative-definite Hessian in both cases. Table 4.6 lists the grid and continuous optima for all three scenarios.




Figure 4.10: Lifetime NPV over the refined grid under the two degradation models, DK1 2022 prices. In each panel the star marks the grid optimum, the open circle the continuous quadratic optimum (Equation 3.16) and the gray dots the evaluated grid points. The diamond marks the best no-degradation design among the refined-grid points; it lies at the corner because the no-degradation optimum is outside this window (Figure 4.9).

The two branches select the same design, differing by 1 MW in power, which is below the 25 MW grid spacing. This checks the surrogate rather than the degradation models. The polynomial was fitted to the physical reference model over the amplitude range the 10–90% window admits (Section 3.5.2), so agreement there is expected; what it establishes is that the substitution does not move the sizing result, which is the condition the surrogate must satisfy to drive the gradient-based method of Section 3.8.3. The branches differ in value rather than in design. Across the refined grid the Shi branch returns 0.67

to 1.22 MEUR more than the Xu branch, or 1.3 to 2.5%, because the polynomial assigns a slightly lower cost to this amplitude distribution. NPV magnitudes in this chapter are therefore taken from the Xu branch.

Table 4.6: Grid and continuous optima from the energy and power capacity sweep, DK1 2022 prices. A dash denotes a value not computed.

Scenario	Grid optimum			Continuous optimum	
	$\bar{E}$ (MWh)	$\bar{P}$ (MW)	NPV (MEUR)	$\bar{E}^*$ (MWh)	$\bar{P}^*$ (MW)
No degradation	800	250	134.2	—	—
Xu [63]	550	175	55.37	560	182
Shi [59]	550	175	56.48	560	183

Two one-dimensional sections through the NPV surface show the two design variables separately, for both degradation models and for no degradation. Figure 4.11 holds power at 175 MW and varies energy capacity: the no-degradation curve peaks between 700 and 750 MWh and declines slowly, while the Xu and Shi curves peak near 550 MWh and fall more steeply above the peak. The vertical distance between the no-degradation and Xu curves is the lifetime cost of degradation, which rises at every sampled capacity from 48 MEUR at 300 MWh to 93 MEUR at 1,100 MWh. Figure 4.12 holds energy at 550 MWh and varies power capacity: the no-degradation curve rises to a shallow maximum near 225 MW while the Xu and Shi curves peak at 175 MW.




Figure 4.11: Lifetime NPV as a function of energy capacity at fixed $\bar{P} = 175$ MW, DK1 2022 prices. Markers are the evaluated designs of both passes, so the spacing is 50 MWh between 450 and 750 MWh and wider outside that range. The dotted line marks the Xu optimum at 550 MWh.




Figure 4.12: Lifetime NPV as a function of power capacity at fixed $\bar{E} = 550$ MWh, DK1 2022 prices. The dotted line marks the Xu optimum at 175 MW.

Accounting for degradation reduces the optimal size by about 30%, from 800 to 550 MWh in energy and from 250 to 175 MW in power. Figure 4.13 separates lifetime NPV into the revenue and the cost that produce it, at fixed 175 MW. Panel (a) plots four discounted quantities against energy capacity. The two navy curves are battery revenue, without degradation and on the Xu branch; their vertical gap is the revenue lost to capacity fade. The two dark red curves are cost: the capital paid at year 0, and that capital plus the discounted replacement. Their gap, shaded, is the cost of the one replacement each design requires. Lifetime NPV is the vertical distance from a revenue curve to the matching cost curve. Panel (b) plots how fast each of these changes as capacity is added, one MWh at a time.

The cause is a balance of rates rather than a change in battery lifetime. Every design in the sweep is replaced exactly once within the 20-year horizon, and a larger battery is replaced later rather than sooner: at 175 MW the replacement year rises from 11 at 300 MWh to 15 at 1,100 MWh, because a larger battery cycles at smaller amplitude for the same throughput. What grows with capacity is the cost of that single replacement, the shaded band in Figure 4.13(a), together with the revenue lost to fade, the gap between the two revenue curves. Their sum is the lifetime cost of degradation, and panel (b) of the same figure shows it rising at a nearly constant 0.05 MEUR per added MWh, three quarters of it discounted replacement capital. The gain in NPV per added MWh without degradation moves the other way, falling from 0.064 MEUR/MWh between 500 and 550 MWh to zero between 700 and 750 MWh. The two curves cross between 550 and 600 MWh, which is where the degradation-aware optimum lies. The shift is large because the NPV surface is flat over this range, not because the cost of degradation is large.

The breadth of the maximum and the size of the shift are separate properties. Of the 42 refined-grid designs, 3 lie within 1% of the Xu optimum and 8 within 2%, spanning 500 to 650 MWh and 150 to 200 MW, so the sizing decision can accommodate manufacturer block sizes or grid connection limits at little cost. The shift is larger than that region: evaluated on the Xu branch, the design that maximizes NPV without degradation returns 47.22 MEUR against 55.37 MEUR at the degradation-aware optimum, 8.15 MEUR or 15% less. Sizing without degradation therefore does not make the plant unprofitable at this site; it oversizes the battery and gives up a sixth of its lifetime NPV.

The mechanism behind the shift is more general than the model that produced it. Any degradation model that assigns a positive and increasing cost to added capacity produces a shift in the same direction. What the result does depend on is the revenue side: it holds for a battery whose only revenue is day-ahead arbitrage, and a revenue stack that pays for installed capacity rather than for energy cycled would move the optimum upward.




Figure 4.13: Battery revenue, cost and their rates of change against energy capacity at fixed $\bar{P} = 175 \text{ MW}$, DK1 2022 prices. Quantities in (a) are discounted to year 0. Rates in (b) are forward differences between adjacent grid points, plotted at the interval midpoint; the series starts at 450 MWh because the grid spacing below that is 150 MWh. The dotted line marks the Xu optimum at 550 MWh.

4.4.2. State-of-Charge Operating Window Sweep

Section 4.4.1 determined the size of the battery. With the energy and power capacity held at that optimum, 550 MWh / 175 MW, this section examines how the width and the center of the operating window affect degradation and lifetime NPV. The two parameters are varied one at a time (Section 3.7.2): a width series at fixed center 0.50, and a center series at fixed width 0.80. The 10–90% window belongs to both series and serves as the reference.¹ Seven windows were simulated; the 30–70% window admits no convex polynomial fit and was not run.

All degradation quantities in this section are on the Xu branch.

Two columns in the tables that follow are measured over different horizons, for different purposes. Revenue is the year-1 value, before any capacity fade, and isolates what the window does to the dispatch. Lifetime NPV is discounted over the full 20-year horizon and is the quantity the design decision rests on.

The width is examined first. Table 4.7 reports the three windows of that series. Widening the window

¹The reference-window value (55.48 MEUR) differs by 0.11 MEUR from the Stage 1 value (55.37 MEUR); this is run-to-run numerical noise, not an effect of the operating window.

from 20–80% to 0–100% raises the annual degradation rate from 0.02036 to 0.02998, an increase of 47%. Almost all of that increase is cycle aging, which rises from 0.00742 to 0.01714. Calendar aging stays near 0.0129 at all three widths. The cycle share of total degradation therefore grows from 36% to 57%.

Revenue rises with width as well, from 15.67 to 21.62 MEUR in year 1, because a wider window frees more of the installed energy for arbitrage. The battery also wears out sooner: replacement moves from year 14 to year 10 and the state of health at year 20 falls from 85.0% to 70.7%. Lifetime NPV rises across the series, from 13.36 to 86.50 MEUR, so over the range tested the revenue gained outweighs the cost of replacing the battery earlier.

Table 4.7: Width series: window swept at fixed center 0.50, battery fixed at 550 MWh / 175 MW, DK1 2022 prices. Degradation quantities are Year-1 values on the Xu reference branch; t_{repl} is the replacement year; SoH_{20} is the state of health at year 20. The 10–90% window is the reference shared with Table 4.8.

Window	δ_{max}	f_d^{cyc}	f_d^{cal}	f_d	Rev (MEUR)	t_{repl} (yr)	SoH_{20} (%)	NPV (MEUR)
20–80%	0.60	0.00742	0.01294	0.02036	15.67	14	85.0	13.36
10–90%	0.80	0.01079	0.01286	0.02365	18.88	12	79.5	55.48
0–100%	1.00	0.01714	0.01285	0.02998	21.62	10	70.7	86.50

The center is examined next. Table 4.8 reports the five windows of that series, all of the same width. Lowering the center from 0.60 to 0.40 reduces the annual degradation rate from 0.02624 to 0.02131, a fall of 18.8%. Cycle and calendar aging both fall by that same 18.8%, so the cycle share stays at 46% across the series.

Revenue does not change across the center series. It is 18.88 MEUR in year 1 at every center, because the width is fixed, so the usable energy and the optimal dispatch are the same in all five runs and only the level at which the battery sits differs. The battery lasts longer: replacement moves from year 11 to year 14 and the state of health at year 20 rises from 75.9% to 84.5%. Lifetime NPV rises from 53.40 to 56.94 MEUR, and since revenue is constant, the whole 3.55 MEUR comes from reduced degradation.

The cycle counts do not follow the degradation. The share of counted cycles below δ_c varies by 11.9 percentage points across the five windows, while the degradation those cycles carry varies by 0.006 percentage points. Because the dispatch schedule is identical across this series (Section 4.2.3), the cycles that differ are the phantom micro-cycles of Section 3.4.3, produced by slack in the state-of-charge balance, and they carry no measurable damage. Cycle counts should therefore be read as upper bounds, and the degradation share compared instead.

Table 4.8: Center series: window swept at fixed width $\delta_{\text{max}} = 0.80$, battery fixed at 550 MWh / 175 MW, DK1 2022 prices. Degradation quantities are Year-1 values on the Xu branch; t_{repl} is the replacement year; SoH_{20} is the state of health at year 20. The 10–90% window is the reference shared with Table 4.7.

Window	Center	f_d^{cyc}	f_d^{cal}	f_d	Rev (MEUR)	t_{repl} (yr)	SoH_{20} (%)	NPV (MEUR)
20–100%	0.60	0.01197	0.01427	0.02624	18.88	11	75.9	53.40
15–95%	0.55	0.01137	0.01355	0.02491	18.88	12	78.8	54.53
10–90%	0.50	0.01079	0.01286	0.02365	18.88	12	79.5	55.48
5–85%	0.45	0.01024	0.01221	0.02245	18.88	13	82.1	56.32
0–80%	0.40	0.00972	0.01159	0.02131	18.88	14	84.5	56.94

The two series can now be compared. Figure 4.14 plots both on the same axes, lifetime NPV against annual degradation rate. The two levers rank differently depending on what is measured. On degradation they are close: the width series spans 0.0096 in annual degradation rate and the center series 0.0049, a factor of two. On lifetime NPV they are not, 73.1 MEUR across the width series against

3.55 MEUR across the center series, a factor of 21. Revenue is the difference. Widening raises revenue and degradation together; lowering the center reduces degradation at unchanged revenue. Width is a trade between revenue and degradation, the center is not.




Figure 4.14: Lifetime NPV against year-1 annual degradation rate for the two window series, battery fixed at 550 MWh / 175 MW, DK1 2022 prices, Xu branch. The width series varies $\delta_{\max}$ at fixed center 0.50; the center series varies the center at fixed width 0.80. The two share the 10–90% window, circled, and are drawn on the same axes. Only the endpoints of the center series are labeled; Table 4.8 lists all five. Degradation is shown here in percent per year and in the tables as a fraction.

The two act through different parts of the degradation model, and the structure of the model explains why. Cycle amplitude enters only the cycle term (Equation 2.9), so width acts on cycle aging alone, through the range of amplitudes the window admits, which is why the cycle share rises from 36% to 57% across that series while calendar aging stays flat. Mean state of charge enters both the cycle term and the calendar term (Equations 2.9 and 2.15), so the center scales the two together through S_σ rather than trading one against the other, and the cycle share holds at 46% while both fall.

At this site the operating window affects lifetime NPV more than the sizing decision does. Widening from the reference window to 0–100%, at a fixed battery, raises lifetime NPV by 31.0 MEUR. Sizing the battery without degradation costs 8.15 MEUR by comparison (Section 4.4.1). How the battery is operated therefore matters at least as much as how it is sized.

That comparison also exposes a limitation of the two-stage sweep. Stage 1 sized the battery at a fixed 10–90% window and Stage 2 varied the window at the Stage 1 size, so size and window are chosen in sequence rather than together. Width changes both revenue and degradation, so the size that is best at one width need not be best at another, and a sequential sweep cannot detect that. Finding the best combination requires searching energy capacity, power capacity and the two window bounds at once, which is the four-variable problem the gradient-based method of Section 3.8 is intended to address.

That limitation does not affect the result reported here, because the center is the parameter the recommendation changes and the center does not enter the sizing trade through revenue. Year-1 revenue is 18.88 MEUR at every center, so moving the window from 10–90% to 5–85% leaves the revenue side of the sizing trade untouched and lowers the annual degradation rate by 5.1%. Applying that reduction to the cost of degradation in Figure 4.13(b) moves the crossing point by about 8 MWh, which is below the 50 MWh spacing of the refined grid and well inside the plateau of Section 4.4.1. Sizing at the recommended window would return the same design.

The better operating choice is to keep the width and to center the window low, within limits that come from the degradation model rather than from the results. The stress factor scores each cycle at its

mean state of charge and charges nothing extra for reaching a full or an empty cell (Section 3.9), so the model neither penalizes the 0–100% window for operating at both extremes nor limits the benefit of lowering the center. At the wide end this matters most at 0–100%, where the second battery reaches year 20 at the replacement threshold. At the low end, lowering the center to 5–85% raises lifetime NPV by 0.84 MEUR over the reference and the further step to 0–80% adds 0.62 MEUR, but that step places the lower bound at the empty cell, where the reported benefit is least reliable. A low but not extreme center, such as 5–85%, is therefore the compromise these results support.

4.4.3. Fidelity of the Polynomial Surrogate

Chapter 3 uses two degradation models for two purposes: the polynomial surrogate wherever a convex cost is required, and the physical reference model wherever degradation is reported (Section 3.5). That division was justified on structural grounds. This section tests it against data.

The surrogate is fitted over the amplitude range each window admits, so the width series tests it directly: each width gives it a different range to follow with a single power law. Table 4.9 compares the two branches across the three widths.

Table 4.9: Xu and Shi branches across the width series, at fixed center 0.50, battery 550 MWh / 175 MW, DK1 2022 prices. f_d is the year-1 annual degradation rate, t_{repl} the replacement year and R^2 the fit quality of the power law over the amplitude range of the window. The last column is the share of year-1 cycle degradation, on the Xu branch, carried by cycles below $\delta_c = 0.1437$, where the surrogate extrapolates outside its fitting range.

Window	f_d^{Xu}	f_d^{Shi}	Δf_d	$t_{\text{repl}}^{\text{Xu}}$ (yr)	$t_{\text{repl}}^{\text{Shi}}$ (yr)	R^2	f_d^{cyc} below δ_c
20–80%	0.02036	0.01989	–2.3%	14	15	0.988	2.66%
10–90%	0.02365	0.02234	–5.5%	12	13	0.963	1.87%
0–100%	0.02998	0.02609	–13.0%	10	11	0.896	1.05%

The surrogate underestimates the annual degradation rate at every width, and the error grows as the window widens: 2.3% at 20–80%, 5.5% at the 10–90% reference and 13.0% at 0–100%. It grows because the fit degrades, from $R^2 = 0.988$ to 0.896, as the fitted exponent rises from $k_4 = 1.03$ to 1.37. A single power law must follow S_δ over a longer amplitude range as the window widens, and the fitted exponent is a compromise across that range.

The error sits entirely in the cycle term. The polynomial model has no calendar term of its own, so the Xu calendar term is applied on both branches (Section 3.5.1) and the two accumulations differ only in how they treat cycle amplitude. Measured on the cycle term alone the underestimate reaches 22.7% at 0–100%, against 13.0% on the total.

The ordering of the windows is preserved: both branches rank the three identically, with degradation rising and replacement advancing as the window widens. The magnitudes are not: the surrogate reports replacement one year later at every width, leaves the battery 2 to 5 percentage points higher in state of health at year 20, and reports a lifetime NPV higher by 0.13, 1.00 and 3.55 MEUR. The surrogate is therefore adequate for computing the sensitivity of NPV to the operating window inside a sizing optimization, while the Xu branch should be used to compute the degradation of the design that optimization returns. That is why every absolute lifetime, replacement and NPV figure in this thesis is reported on the Xu branch.

The one-year offsets in Table 4.9 are not equally meaningful. Replacement is checked once a year, so a reported difference of one year can stand for anything between almost nothing and almost two years. Interpolating each state-of-health trajectory through the replacement trigger separates the two cases. At 20–80% the Xu branch crosses at 13.98 years and the Shi branch at 14.33, a real difference of 0.35 years reported as one; the Xu battery ends year 14 at 70.476% state of health, 0.024 percentage points below the trigger, so a marginally slower degradation rate would have removed the offset. At the 10–90% reference the crossings are 11.81 and 12.57 years, a difference of 0.76 years. At 0–100% they are 9.09 and 10.62 years, a difference of 1.53 years, so there the reported single year understates the gap. The offset at the narrow window is largely an artifact of the annual check; the offset at the widest window is not.

Two things could produce that error: the power law may fit poorly inside the amplitude range it was fitted to, or it may be inaccurate below δ_c , where it is applied outside that range. Evaluating both branches on the same cycles separates the two. Cycles below δ_c carry 2.66%, 1.87% and 1.05% of year-1 cycle degradation on the Xu branch as the window widens, against 2.04%, 1.28% and 0.64% on the Shi branch. The difference between each pair is the error contributed by the extrapolation, and it accounts for 11.7%, 6.1% and 2.4% of the total cycle-term error. Between 88 and 98% of the error therefore comes from cycles inside the fitted range, and the extrapolation is not the main source at any width.

4.5. Monolithic Nonlinear Program

Section 3.8.2 argues that a general-purpose optimization solver will have difficulty with a dispatch problem that includes a degradation cost built on rainflow cycle counting, because that cost is not differentiable in the dispatch variables. This section tests that argument against data.

The test compares two dispatches of the same day. The first is a revenue-only LP over the 24 hours of that day, using the formulation of Section 3.4.1. The second adds the degradation cost to the objective and solves the resulting nonlinear program (Equation 3.17) over the same 24 hours, using the `trust-constr` implementation and the solver settings of Section 3.8.1. Both use the same prices, the same wind and the same battery, held at 550 MWh / 175 MW, and both start and end the day at the same stored energy.

A single day is the smallest problem that still couples the dispatch to the degradation cost, so it is the setting in which the formulation is most likely to succeed. It is not a slice of the annual dispatch reported elsewhere in this chapter. The annual solve optimizes all 8,760 hours together, so a day within it opens at the state of charge the previous day ended at and is chosen jointly with the rest of the year. Here each day stands alone, which also means that a solve fails on its own trajectory rather than on error carried over from earlier days.

Three quantities describe each dispatch, all stated as the nonlinear result minus the LP result. The revenue difference is positive when the nonlinear dispatch earns less. The degradation cost difference is positive when it degrades the battery less. The objective difference is the degradation cost difference less the revenue difference. It is the objective of Equation 3.17 evaluated at the nonlinear dispatch less the same objective evaluated at the LP dispatch, so both dispatches are scored on the same terms and a positive value means the nonlinear dispatch is worth more once degradation is priced. A fourth quantity describes the solver rather than the dispatch: the final optimality measures how far the first-order optimality conditions are from being satisfied when it stops (Section 3.8.1). A day is converged when that value falls below 10^{-6} , stalled between 10^{-6} and 10^{-2} , and diverged above 10^{-2} . A stalled or diverged day has not satisfied the optimality conditions, so its dispatch is not a local optimum and its objective difference is not a result the solver established.

If the missing derivative matters in practice, two things should follow. The solver should stop without reaching an optimum on a large share of the days attempted, which Section 4.5.1 tests. And where the nonlinear dispatch appears to beat the LP, most of that gain should come from days on which the solver stopped early rather than from days on which it found a better dispatch, which Section 4.5.2 tests.

4.5.1. Convergence Behavior

The comparison was run for each of the 365 days of the DK1 2022 price series. On nine of them the price spread was too small to justify a charge cycle and the battery stayed idle in both dispatches, so those days are excluded and 356 remain. Days are grouped by price variability, using the tertiles of the daily price spread within 2022: Low at or below 127 EUR/MWh, Mid from 127 to 210 EUR/MWh, and High above 210 EUR/MWh. The cut values rank days within the year and are specific to 2022.

Of the 356 active days, 185 (52.0%) converge, 95 (26.7%) stall and 76 (21.3%) diverge. Table 4.10 reports the same outcome split by price-variability regime. The convergence rate does not order with the price spread: it is 61% in the Low regime, 44% in the Mid regime and 51% in the High regime, so the lowest rate occurs at neither extreme. The regime split establishes that the outcome is not confined to one part of the year, but it does not identify a cause.

The price regime does affect the size of the objective difference. The median difference is least negative

in the Low regime at -2.24 EUR and most negative in the High regime at -6.94 EUR, because a wider spread means a larger revenue reduction for the same change in dispatch. The positive days concentrate in the Low regime, 35 of 63, against 20 in the Mid and 8 in the High, because on a low-spread day the revenue attached to a single shallow cycle is small enough that removing it improves the objective.

Table 4.10: Monolithic solver outcome by price-variability regime, 356 active days, 550 MWh / 175 MW, DK1 2022. The objective difference is positive where the nonlinear dispatch is worth more than the LP dispatch.

Regime	Days	Converged	Stalled	Diverged	Median difference (EUR)	Positive days
High	122	62 (51%)	33 (27%)	27 (22%)	-6.94	8
Mid	120	53 (44%)	34 (28%)	33 (28%)	-4.08	20
Low	114	70 (61%)	28 (25%)	16 (14%)	-2.24	35

Figure 4.15 plots, for every day, the objective difference against the final optimality on a logarithmic axis. The converged days form a dense column at the left, most of them below the zero line. The largest positive differences appear as the optimality grows, to the right of the convergence threshold, and the very largest lie inside the diverged region beyond 10^{-2} .

Just under half of all days leave the solver unable to certify a local optimum, which is the first of the two expected outcomes. The scatter adds a second result: the largest reported improvements sit on days that did not converge, so no reported gain can be attributed to a degradation-aware dispatch until the trajectory behind it is inspected. Section 4.5.2 performs that inspection.

The cause lies in the objective, and it can be measured directly. Appendix C.2 perturbs the charging power at 200 sampled time steps of a solved annual dispatch and compares the two one-sided difference quotients. At 44% of those steps under DK1 2019 prices and 59% under DK1 2022 prices the quotients disagree by more than one percent, which curvature cannot produce and which indicates that the cycle set re-pairs and no derivative exists at that step. Points without a derivative are therefore common along the trajectory, not isolated.

Numerical conditioning is not the explanation. Every day reported here was solved with the objective and its gradient divided by the LP net revenue for that day (Section 3.8.1). Final optimality nonetheless spans from 9.4×10^{-8} to 4.4×10^{12} across the 356 active days and ends above unity on 38 of them. A first-order optimality measure reaching that magnitude on an objective normalized to order one is not a consequence of the scaling choice.

The inexactness of the supplied gradient is not the explanation either. The gradient handed to the solver omits the response of the mean-SoC factor S_σ to a change in the cycle mean (Section 3.8.1), an omission bounded by $\delta_i k_\sigma / 2k_4$ that vanishes for shallow cycles. An inexact but continuous gradient slows a solver or shifts the point it converges to; it does not remove the derivative. The measurement above concerns the objective itself rather than the gradient supplied for it, so the diagnosis does not depend on that choice.

The way the failed solves terminate points at the same cause. Of the 171 days that do not converge, 156 stop because the trust-region radius falls below the step tolerance of 10^{-12} rather than because the iteration limit of 500 is reached, and on 64 of those the first-order optimality is still above 10^{-2} . A trust region contracting through twelve orders of magnitude while the optimality measure does not fall is what happens when no step of any length improves the local model, which is the behavior a discontinuous gradient produces. Every converged day stops on the gradient criterion instead.

4.5.2. Dispatch Outcome

For the median day the nonlinear dispatch is worth less than the LP dispatch once degradation is priced. The median objective difference across the 356 active days is -4.39 EUR, and only 63 days show a positive difference. The mean is positive, at $+5.77$ EUR, but only because of a few large outliers, so the median is the statistic reported here.

Those 63 days do not support the conclusion they suggest. They account for 5,246 EUR of gain in




Figure 4.15: Objective difference of the nonlinear dispatch relative to the LP, against final solver optimality, for the 356 active days at $\bar{E} = 550$ MWh and $\bar{P} = 175$ MW under DK1 2022 prices. Navy points are converged solves (optimality below 10^{-6}) and dark red points are solves that did not converge. The dashed line marks the convergence threshold and the dotted line the divergence threshold at 10^{-2} . The difference is clipped to ± 150 EUR for display, which affects six days, all of them non-converged. The four example days of Table 4.11 are labeled.

total, and 87% of that comes from days on which the solver never reached an optimum. Nine of the ten largest single-day gains are among them. Of the 185 days the solver did solve to an optimum, 168 match or fall short of the LP, and the 17 that beat it account for the remaining 666 EUR. This is the second expected outcome: most of the apparent improvement is a product of the failure to converge rather than of a better dispatch.

The solver shows four behaviors, and every active day falls into one of them. Table 4.11 gives the share of each and names one example day; Appendix E.1 plots those four days against their LP counterparts.

Table 4.11: The four solver behaviors across the 356 active days, 550 MWh / 175 MW, DK1 2022.

Type	Solver outcome	Days	Share	Example
A	Converged, improves on the LP	17	4.8%	D049
B	Converged, matches or falls short	168	47.2%	D179
C	Stalled	95	26.7%	D084
D	Diverged	76	21.3%	D363

Types A and B are the outcomes the solver established, and together they are the 185 converged days. In Type A the nonlinear dispatch drops one shallow cycle that the LP performed and leaves the rest of the day unchanged, so the reduction in degradation cost exceeds the reduction in revenue; the 17 days of this type account for 666 EUR. Type B is the larger group by far, and its days divide between two patterns. In one the two trajectories coincide and revenue is given up for a reduction in degradation near zero. In the other the nonlinear trajectory sits above the LP for the whole day, earning less while degrading more, because a higher mean state of charge raises the stress factor. The median objective difference across Type B is -8.00 EUR.

Types C and D are the outcomes the solver did not establish, and they carry most of the reported gain. In Type C the state of charge is held flat over part of the day and a large reduction in degradation is reported, but the solver stops between 10^{-6} and 10^{-2} ; 24 of the 95 stalled days report a positive difference, totaling 1,088 EUR. In Type D the state of charge does not move at all, so the reported reduction is an accounting result for a battery that did nothing; 22 of the 76 diverged days report a

positive difference, totaling 3,492 EUR. Between them these two types account for 87% of all positive objective difference in the year.

These four behaviors follow from the shape of the degradation cost. Lowering degradation means changing which cycles the counter finds in the trajectory, and that requires crossing a point where the cycle set changes. The gradient gives the solver nothing to follow toward such a point, because it is flat on the approach and jumps at the crossing. The solver therefore stops short of the crossing (Type C), steps across it into a region where the accumulated curvature no longer describes the objective and diverges (Type D), or settles at a point it can certify that is no better than the LP, because the only directions in which the gradient varies leave the cycle set unchanged (Type B). The one improvement it does find, removing a whole shallow cycle (Type A), requires no crossing, and a revenue-only LP makes the same change once the revenue from that cycle falls below its degradation cost.

The monolithic formulation does not deliver a usable degradation-aware dispatch. Under the conditions most favorable to it, a one-day horizon at the reference design with the objective scaled, the solver reaches an optimum on about half the days attempted, and that rate does not order with the price regime. Where it does reach an optimum, the dispatch it returns is worse than the LP on five days out of six.

A degradation-aware dispatch is correspondingly difficult to characterize against a revenue-only one. The improvements the solver does certify come from removing a single low-value cycle rather than from a different way of operating the battery. The operational value of accounting for degradation is measured instead by the window sweep of Section 4.4.2, which requires no optimizer inside the dispatch. The route that keeps the dispatch linear and moves the degradation feedback outside the solver, through a per-cycle sub-gradient on the Shi branch, is verified in Appendix C; closing it into a full sizing optimizer is left as future work (Section 3.8.3). The sensitivity of degradation to the inputs of the dispatch LP can be calculated for that loop: combining the degradation sub-gradient with the LP dual price vector gives the fused per-year sensitivity, and Appendix C.3 reports the method and why the sizing gradient is assembled differently.

The test covers one design point and one price year, so the convergence rate itself would likely differ elsewhere. What generalizes is the pattern of behavior and not the 52.0% figure.

Conclusions and Recommendations

This thesis examined to what extent explicitly accounting for battery degradation changes the optimal operation and design of a hybrid wind and battery power plant. The analysis and the interpretation of the results are given in Chapter 4. This chapter provides the final synthesis, linking those findings to the research gaps identified in Chapter 1, to the research objective, and to the research question and its sub-questions. It introduces no new analysis.

5.1. Link to the Research Gaps

The literature review identified two limitations in degradation-aware frameworks for wind and storage plants. The first concerns the fidelity of the degradation term. Degradation is commonly reduced to a single aggregated quantity, most often a cost proportional to energy throughput, which charges the same amount for every megawatt-hour cycled regardless of how deeply the battery was cycled or at what state of charge. A representation of that form cannot show which features of operation drive degradation, because it does not distinguish between them. The second concerns the structure of the coupling rather than the model. Frameworks that evaluate degradation after the dispatch, or that feed it back between solves, rely on exhaustive enumeration and pass no degradation sensitivity to the design level.

The two limitations share a root. Degradation is produced by the operating decision and carried as a cost by the design decision. A framework that represents it coarsely cannot trace the first link, and one that couples it loosely cannot act on the second.

This thesis contributes to closing these gaps in five ways. Against the first gap, on the fidelity of the degradation term:

- It implements a rainflow cycle-counting degradation model inside the dispatch optimization of a wind and battery plant, evaluated on the full solved trajectory and reporting degradation per identified cycle rather than per megawatt-hour. Cycle amplitude, mean state of charge and calendar aging are accumulated and reported separately, which is what allows their contributions to be told apart.
- It measures the two operational characteristics against each other at a fixed design, separating the effect of the width of the state-of-charge window from that of its center, and establishes how much of the total degradation each can reach.
- It carries the degradation result into the sizing decision through a multi-year loop with capacity fade and replacement, and quantifies the change in the selected storage size against a degradation-free design.

Against the second gap, on the structure of the coupling:

- It formulates the embedded degradation-aware dispatch problem and establishes, against a full year of daily solves, that a general-purpose gradient-based solver cannot reliably solve it. The

cause is identified as a property of rainflow cycle counting rather than of the solver, which makes the result transferable to any formulation built the same way.

- It derives and verifies the analytical degradation sensitivities that a directed search would require, as the route that avoids that obstacle: a per-timestep sub-gradient of the degradation cost with respect to the dispatch, and a capacity gradient at fixed dispatch. Both are verified numerically, and the outer loop that would use them is left as future work.

Underlying all five is a two-model architecture. The physical reference model is used wherever degradation is reported, and the polynomial surrogate only where the gradient path requires a convex cost. The division is justified on structural grounds in Chapter 3 and tested against data in Section 4.4.3, where the surrogate is shown to preserve the ordering of designs while underestimating their degradation.

5.2. Research Objective

The objective of this thesis was to establish to what extent accounting for battery degradation changes the operating strategy and the storage size selected for a hybrid power plant. The objective was addressed by applying one reference plant, one wind resource and one electricity market to every scenario, and varying only the treatment of degradation and the design parameters under study.

Holding the plant and the market fixed is what makes the comparison attributable. The same hourly dispatch optimization over a full year, the same twenty-year horizon and the same replacement rule are used throughout. Between the no-degradation, physical and polynomial scenarios, only the degradation term changes; between the design points of the sweep, only the energy and power capacity change; between the windows of the second stage, only the state-of-charge bounds change. Differences in the reported net present value therefore follow from the quantity under study rather than from a change in the model.

All three coupling structures of Figure 1.1 appear in this work: degradation evaluated after each solve, returned between solves through the multi-year loop, and placed inside the dispatch objective. Which of them supported a usable result is the subject of Section 5.3.

5.3. Answers to the Research Questions

Sub-question 1: Which characteristics of battery operation, such as cycle amplitude and mean state of charge, act as the effective levers for reducing degradation?

Cycle amplitude is the stronger single driver of degradation, but the window center is the more useful lever. Widening the operating window from 20–80% to 0–100% raises the annual degradation rate by 47%, almost all of it cycle aging, because amplitude enters only the cycle term of the degradation model. Lowering the window center from 0.60 to 0.40 reduces the annual rate by 18.8%, a smaller change, but it acts on cycle and calendar aging together through the mean-SoC stress factor, and it costs no revenue. What either lever can achieve is bounded by calendar aging, which carries 54% of annual degradation at the selected design and the reference window and does not depend on the dispatch at all. An operating rule that regulates cycle amplitude alone therefore addresses less than half of the degradation at this site.

Sub-question 2: How does the choice of state-of-charge operating window trade revenue against battery degradation and lifetime?

Window width is a trade and window center is not. Across the width series, year-1 revenue rises from 15.67 to 21.62 MEUR, the replacement year advances from 14 to 10, and lifetime net present value rises from 13.36 to 86.50 MEUR. Over the range tested the revenue gained outweighs the earlier replacement, so a wider window is worth more despite the faster wear. Across the center series, year-1 revenue is 18.88 MEUR at every center because the usable energy and the optimal dispatch do not change; the replacement year moves from 11 to 14 and lifetime net present value rises by 3.55 MEUR, all of it from reduced degradation. The two levers differ by a factor of two on degradation and a factor of 21 on lifetime value, and the difference is revenue. The operating choice these results support is to keep the manufacturer-recommended width and to center the window low, at 5–85% rather than at

either extreme, since the widest and lowest windows are the points at which the degradation model is least reliable.

Sub-question 3: How does accounting for degradation change the optimal storage sizing, compared with a design obtained under degradation-free operation?

Accounting for degradation reduces the optimal battery by about 30%, from 800 MWh and 250 MW to 550 MWh and 175 MW, and both degradation models select the same design. Evaluated on the physical reference model, the design that maximizes net present value without degradation returns 47.22 MEUR against 55.37 MEUR at the degradation-aware optimum, so sizing without degradation costs 8.15 MEUR, or 15% of lifetime value. It does not make the plant unprofitable at this site; it oversizes the battery. The optimum is a plateau rather than a peak: 8 of the 42 refined designs lie within 2% of it, spanning 500 to 650 MWh and 150 to 200 MW, so the sizing decision can accommodate manufacturer block sizes or grid connection limits at little cost. The shift is large because the value surface is flat over this range, not because the cost of degradation is large.

Main research question: To what extent does explicitly accounting for battery degradation influence optimal operational strategies and design decisions in hybrid power plants?

Degradation changes the design decision more than the operating decision. It removes 250 MWh of energy capacity and 75 MW of power from the selected design. Carrying the larger design instead costs 8.15 MEUR of lifetime net present value: the extra capacity earns more revenue, but not enough to cover the capital and the replacement it adds, and the difference between the two net present values is what remains. On the operating side the influence of degradation is real but smaller. Window width is chosen the same way with or without a degradation cost, because lifetime net present value rises with width in both cases, so degradation does not change that choice. It changes only the center, worth 3.55 MEUR of lifetime net present value across the full range tested and 0.84 MEUR for the step from the reference window to the recommended one.

That influence cannot be delegated to a solver. Placing the degradation cost inside the dispatch objective produces a problem that a general-purpose gradient-based solver certifies on only 52% of the days attempted, even at the shortest horizon and the reference design. Of the 171 days that fail, 156 stop because the trust region collapses rather than because the iteration limit is reached, which is the behavior expected when no step improves the local model. Of the 185 days that do converge, 168 return a dispatch no better than the revenue-only linear program, and 87% of the apparent gain across the year comes from solves that never reached an optimum. The result is a property of rainflow cycle counting rather than of the solver, so it applies to any formulation that places a rainflow-based degradation cost inside a gradient-based optimization. What this thesis establishes about degradation therefore rests on evaluating a designed operating window and a designed size, not on an optimizer that discovers them.

5.4. Impact

The results bear on system integration through the size of the asset rather than through its dispatch. A design that accounts for degradation is about 30% smaller in energy capacity for the same plant role, so it commits less capital and less cell material to deliver the same firm output behind a given grid connection. Because the dominant value of co-located storage is capacity and integration value rather than arbitrage value (Section 1.3), and because that value depends on the battery remaining available, sizing and operating for slower aging serves it directly and without any change in market design.

The direction of these results matches how the storage industry itself approaches cost reduction. The Danish Energy Agency technology catalogue identifies two routes to a lower levelized cost of storage: reducing the capital cost of the cells, or extending the battery lifetime so that the capital already committed is spread over more cycles [10]. This thesis acts on the second. The window center extends the time to replacement by up to three years at no loss of revenue, and it requires no change in cell technology, only a change in the operating limits, which can be applied through operating rules or warranty terms rather than through capital expenditure.

The material consequence runs in the same direction as the economic one. Each replacement requires

a new cell inventory, with the extraction and manufacturing emissions that entails. The same catalogue reports that every lithium-ion chemistry considered for grid storage depends on critical raw materials whose production is concentrated in a few countries, and that end-of-life recycling currently supplies a small share of demand for most of them and almost none for lithium and graphite [10]. Recycling is not expected to close that gap as demand rises [28]. A smaller battery and a later replacement both reduce primary material demand, so the case for accounting for degradation runs the same way for materials as it does for cost.

The way degradation is modeled also bears on how retired batteries are valued. The catalogue reports that second-life systems, which reuse packs retired from electric vehicles, reduce capital cost by 21 to 36% but raise the levelized cost of storage by 11 to 32% over a fifteen-year period, because the remaining life of a used pack is uncertain and its performance is degraded [10]. A model that resolves degradation per cycle rather than per megawatt-hour addresses the first of those obstacles directly, since what a second-life valuation requires is the remaining life of a specific pack under a specific operating pattern, which is what a cycle-resolved model returns. That application is not pursued here.

The dependence on the price regime bounds all of this. At the reference design the battery-marginal net present value is -89.7 MEUR under DK1 2019 prices and +42.9 MEUR under DK1 2022 prices, a swing of 132.6 MEUR and a change of sign at the same site with the same wind resource (Section 4.3.3). Degradation separates the two years by far less: the annual rate differs by 3.4% and replacement moves by one year. The value of storage at this site is set by price volatility more than by degradation, and both figures value the plant through day-ahead arbitrage alone. As renewable penetration rises and the merit-order effect narrows price spreads, the arbitrage case weakens while the integration case does not, which is an argument for valuing such plants on capacity and system services rather than on energy arbitrage alone.

5.5. Limitations

Of the modeling limitations collected in Section 3.9, the one that bears most on the answers above is that the framework is verified but not validated. The checks of Section 4.2 and Appendix C, against the equations on a solved week, against closed-form answers on synthetic traces and against finite differences for the two derivatives, establish that the implementation computes the model it is intended to compute; they do not establish that the model predicts the behavior of an installed battery. Every absolute lifetime, replacement year and net present value reported here is conditional on the coefficients of Xu et al. as published. The relative results, the direction and approximate size of the sizing shift and the ranking of the two operating levers, follow from the structure of the model rather than from its fitted values and are correspondingly more robust.

The chemistry is the second such limitation. The model is parameterized for LMO, while grid-scale installations predominantly use LFP or NMC, which have longer cycle lives. Section 3.9 separates what this affects from what it does not: the parameter values carry every absolute quantity, while the accumulation rule, the two-model architecture and both gradient derivations follow from the functional forms. A longer-lived chemistry would shift the balance further toward calendar aging and reduce the cost per unit of energy cycled, which would move the optimum upward and shrink the reported shift without reversing it.

Three model properties bound the operating answer. The stress factor scores each cycle at its mean state of charge and charges nothing additional for reaching a full or an empty cell (Section 3.9), so the 0–100% window is not penalized for operating at both extremes and the benefit of lowering the window center is not capped; the two ends of the range tested are therefore the least reliable points in it. The late-life acceleration of capacity fade is not represented (Section 3.9), so lifetimes at high utilization are likely optimistic. And no C-rate stress term is included, so the model responds to how much energy is cycled but not to how fast, which means the answers say less about the choice of power capacity than about the choice of energy capacity.

Four properties of the case study bound the design answer. The results rest on one site, one chemistry and one revenue stream: the battery earns from day-ahead arbitrage only, and a revenue stack that pays for installed capacity rather than for energy cycled would move the optimum upward. The dispatch has perfect price foresight, so the revenue reported is an upper bound on what a plant operating against

forecasts could achieve. The replacement cost is held at its present value across the twenty-year horizon, whereas the catalogue it is drawn from projects cell costs to fall by roughly half by 2040 [10]; a declining replacement cost would reduce the penalty for degradation, so the reported shift is an upper bound in that respect as well. And the two-stage sweep chose size and window in sequence rather than together, so it cannot detect coupling between them; the effect of that sequence on the sizing result was estimated and found to be smaller than the grid spacing (Section 4.4.2), but it was not measured directly.

The polynomial surrogate is a limitation of the method rather than of the results. It underestimates the annual degradation rate by 2.3 to 13.0% depending on window width (Section 4.4.3), which is why every degradation figure reported in this thesis is taken from the physical reference model and the surrogate is used only where the gradient path requires a convex cost.

5.6. Recommendations for Future Research

Closing the gradient-based sizing loop is the first direction. The frozen-dispatch component of the capacity gradient is implemented and verified here, and the per-timestep degradation sub-gradient it needs is verified alongside it (Appendix C). What remains is the re-dispatch term, the change in the optimal dispatch that follows a change in capacity, and a step-size rule. The perturbation-based sensitivity framework of Castillo et al. [7] provides the formal route to the re-dispatch term, and because the inner problem is a linear program its Lagrangian Hessian is null, so the second-order form is available without the computational cost that a nonlinear inner problem would carry. Lighter alternatives are the parametric sensitivity of the linear program basis and finite differencing of the stored-energy trajectory across outer iterates. With the loop closed, energy capacity, power capacity and the two window bounds could be searched together, which is the four-variable problem the two-stage sweep of this thesis cannot address.

Extending the degradation model would sharpen both operating answers, and three additions matter most. The first is efficiency and power fade. This thesis models capacity fade alone and holds the charge and discharge efficiencies constant across the battery lifetime (Section 3.9), so a battery that loses round-trip efficiency as its internal resistance grows would deliver less revenue in later years than reported here, and would draw more energy to deliver the same output, raising cycle aging in turn. The second is a C-rate stress term, which would make the model sensitive to how fast energy is cycled and not only to how much, and would therefore give the choice of power capacity a degradation cost that the present formulation does not assign it. The third is a stress term for the state-of-charge extremes. Low-SoC impedance rise and high-SoC lithium plating are both absent from the model (Section 3.9), which is why the benefit of lowering the window center has no floor in these results; adding either term would bound the recommendation and test whether a low center survives a model that charges for approaching an empty cell.

Two further model extensions would test the boundaries of the reported results rather than sharpen them. Representing post-knee capacity fade is expected to increase the sizing shift; repeating the study on a longer-lived chemistry such as LFP or NMC is expected to reduce it, for the reasons given in Section 5.5. Neither would require re-deriving the framework, only refitting the coefficients and the polynomial surrogate.

Extending the case study would establish how far the results generalize. Repeating the sweep across additional price years and across sites with different wind resources and market conditions would show whether the sizing shift of about 30% is a property of this configuration or of the mechanism behind it, since the mechanism, a degradation cost that grows with capacity acting on a flat value surface, does not depend on the site. Replacing perfect price foresight with a forecast-based dispatch would quantify how much of the reported revenue survives realistic operation, and would also change the cycling pattern and therefore the degradation. The most consequential extension is to the revenue structure itself. This thesis values the battery through day-ahead arbitrage alone. Adding balancing and reserve markets, or a capacity payment that rewards installed energy rather than energy cycled, would change the sign of the trade this thesis measures, because a revenue stream that pays for availability rather than for throughput rewards a larger and less heavily cycled battery. That is the single condition most likely to change the answer to the third sub-question.


References

- [1] C. Amzallag et al. "Standardization of the rainflow counting method for fatigue analysis". In: *International Journal of Fatigue* 16.4 (June 1994), pp. 287–293. ISSN: 0142-1123. DOI: 10.1016/0142-1123(94)90343-3.
- [2] ASTM International. *Standard Practices for Cycle Counting in Fatigue Analysis*. Standard ASTM E1049-85(2017). West Conshohocken, PA: ASTM International, 2017. URL: <https://astm.org>.
- [3] Mostafa Bakhtvar et al. "A vision of flexible dispatchable hybrid solar-wind-energy storage power plant". In: *IET Renewable Power Generation* 15.13 (2021), pp. 2983–2996. DOI: <https://doi.org/10.1049/rpg2.12234>. URL: <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/rpg2.12234>.
- [4] Majid Bastankhah and Fernando Porté-Agel. "A new analytical model for wind-turbine wakes". In: *Renewable Energy* 70 (2014). Special issue on aerodynamics of offshore wind energy systems and wakes, pp. 116–123. ISSN: 0960-1481. DOI: <https://doi.org/10.1016/j.renene.2014.01.002>. URL: <https://www.sciencedirect.com/science/article/pii/S0960148114000317>.
- [5] Ted K. A. Brekken et al. "Optimal Energy Storage Sizing and Control for Wind Power Applications". In: *IEEE Transactions on Sustainable Energy* 2.1 (2011), pp. 69–77. DOI: 10.1109/TSTE.2010.2066294.
- [6] Richard H. Byrd, Mary E. Hribar, and Jorge Nocedal. "An Interior Point Algorithm for Large-Scale Nonlinear Programming". In: *SIAM Journal on Optimization* 9.4 (1999), pp. 877–900. DOI: 10.1137/S1052623497325107.
- [7] Enrique Castillo, Roberto Mínguez, and Carmen Castillo. "Sensitivity analysis in optimization and reliability problems". In: *Reliability Engineering & System Safety* 93.12 (2008). 17th European Safety and Reliability Conference, pp. 1788–1800. ISSN: 0951-8320. DOI: <https://doi.org/10.1016/j.ress.2008.03.010>. URL: <https://www.sciencedirect.com/science/article/pii/S0951832008000999>.
- [8] Kok-Hon Chew et al. "Analytical gradient-based optimization of offshore wind turbine substructures under fatigue and extreme loads". In: *Marine Structures* 47 (2016), pp. 23–41. ISSN: 0951-8339. DOI: <https://doi.org/10.1016/j.marstruc.2016.03.002>. URL: <https://www.sciencedirect.com/science/article/pii/S095183391630017X>.
- [9] A. Crespo and J. Hernández. "Turbulence characteristics in wind-turbine wakes". In: *Journal of Wind Engineering and Industrial Aerodynamics* 61.1 (1996), pp. 71–85. ISSN: 0167-6105. DOI: [https://doi.org/10.1016/0167-6105\(95\)00033-X](https://doi.org/10.1016/0167-6105(95)00033-X). URL: <https://www.sciencedirect.com/science/article/pii/016761059500033X>.
- [10] Danish Energy Agency and Energinet. *Technology Data for Energy Storage*. <https://ens.dk/en/analyses-and-statistics/technology-data-energy-storage>. First published October 2018; chapter 180, Lithium-ion Batteries for Grid-Scale Storage, amended May 2025. Accessed: 2026-07-01. 2018.
- [11] Kaushik Das et al. "Research Challenges and Opportunities of Utility-Scale Hybrid Power Plants". In: *WIREs Energy and Environment* 14.1 (2025), e70001. DOI: <https://doi.org/10.1002/wene.70001>. URL: <https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wene.70001>.
- [12] S. D. Downing and D. F. Socie. "Simple rainflow counting algorithms". In: *International Journal of Fatigue* 4.1 (Jan. 1982), pp. 31–40. ISSN: 0142-1123. DOI: 10.1016/0142-1123(82)90018-4. URL: <https://www.sciencedirect.com/science/article/pii/0142112382900184> (visited on 01/12/2026).
- [13] K. Dykes et al. *Opportunities for Research and Development of Hybrid Power Plants*. Tech. rep. NREL/TP-5000-75026. National Renewable Energy Laboratory (NREL), 2020. URL: <https://docs.nrel.gov/docs/fy20osti/75026.pdf>.

- [14] Jacqueline S. Edge et al. "Lithium ion battery degradation: what you need to know". In: *Phys. Chem. Chem. Phys.* 23 (14 2021), pp. 8200–8221. DOI: 10.1039/D1CP00359C. URL: <http://dx.doi.org/10.1039/D1CP00359C>.
- [15] ENTSO-E. *ENTSO-E Transparency Platform*. <https://transparency.entsoe.eu>. European Network of Transmission System Operators for Electricity. Accessed: 2026-03-16. 2024.
- [16] European Union Agency for the Cooperation of Energy Regulators. *Gas Wholesale Market Monitoring Report shows market expectations of high gas prices until 2024*. News Release. July 2022. URL: <https://acer.europa.eu/news/gas-wholesale-market-monitoring-report-shows-market-expectations-high-gas-prices-until-2024> (visited on 08/11/2026).
- [17] Iratxe Gonzalez-Aparicio et al. "Optimal Sizing and Operation of Hybrid Renewable Power Plants Participating in Coupled Power Markets with Different Execution Times". In: *Energies* 16.8 (2023), p. 3432. DOI: 10.3390/en16083432.
- [18] A. González-Garrido et al. "Electricity and reserve market bidding strategy including sizing evaluation and a novel renewable complementarity-based centralized control for storage lifetime enhancement". In: *Applied Energy* 262 (2020), p. 114591. ISSN: 0306-2619. DOI: <https://doi.org/10.1016/j.apenergy.2020.114591>. URL: <https://www.sciencedirect.com/science/article/pii/S0306261920301033>.
- [19] David Gräf et al. "What drives capacity degradation in utility-scale battery energy storage systems? The impact of operating strategy and temperature in different grid applications". In: *Journal of Energy Storage* 47 (Mar. 2022), p. 103533. ISSN: 2352-152X. DOI: 10.1016/j.est.2021.103533.
- [20] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. Accessed: 2026-03-16. 2024. URL: <https://docs.gurobi.com/projects/optimizer/en/current/index.html>.
- [21] Xuebing Han et al. "A review on the key issues of the lithium ion battery degradation among the whole life cycle". In: *eTransportation* 1 (2019), p. 100005. ISSN: 2590-1168. DOI: <https://doi.org/10.1016/j.etrans.2019.100005>. URL: <https://www.sciencedirect.com/science/article/pii/S2590116819300050>.
- [22] William E. Hart, Jean-Paul Watson, and David L. Woodruff. "Pyomo: modeling and solving mathematical programs in Python". In: *Mathematical Programming Computation* 3.3 (Sept. 2011), pp. 219–260. ISSN: 1867-2957. DOI: 10.1007/s12532-011-0026-8. URL: <https://doi.org/10.1007/s12532-011-0026-8>.
- [23] Guannan He et al. "Optimal Bidding Strategy of Battery Storage in Power Markets Considering Performance-Based Regulation and Battery Cycle Life". In: *IEEE Transactions on Smart Grid* 7.5 (2016), pp. 2359–2367. DOI: 10.1109/TSG.2015.2424314.
- [24] Ehsan Hosseini et al. "Optimal energy management system for grid-connected hybrid power plant and battery integrated into multilevel configuration". In: *Energy* 294 (2024), p. 130765. ISSN: 0360-5442. DOI: <https://doi.org/10.1016/j.energy.2024.130765>. URL: <https://www.sciencedirect.com/science/article/pii/S0360544224005371>.
- [25] IEA Wind TCP. *Annual Report 2021 – Task 50*. Tech. rep. International Energy Agency, 2021. URL: https://iea-wind.org/wp-content/uploads/2022/12/IEA_Wind_TCP_AR2021_Task50.pdf.
- [26] IEA Wind TCP. *Task 50: Hybrid Power Plants*. IEA Wind Technology Collaboration Programme. Accessed: 2026-03-16. 2024. URL: <https://iea-wind.org/task50/>.
- [27] International Electrotechnical Commission. *IEC 61400-1: Wind turbines – Part 1: Design requirements*. International Standard IEC 61400-1:2019. Geneva, Switzerland: International Electrotechnical Commission, 2019.
- [28] International Energy Agency. *Global Critical Minerals Outlook 2026*. Licence: CC BY 4.0. Accessed: 2026-08-04. Paris, France, 2026. URL: <https://www.iea.org/reports/global-critical-minerals-outlook-2026>.
- [29] Jenna Iori. *SHIPP - A design optimization software for hybrid power plants*. 2025. DOI: 10.4121/2EE36148-369F-4E1F-B770-C86752D7DCA4. URL: <https://github.com/jennaiori/shipp>.

- [30] Jenna Iori et al. "Design drivers for the storage system of baseload hybrid power plants". In: *NET Conference Proceedings 2024.2* (2024), pp. 245–250. DOI: 10.1049/icp.2024.1844.
- [31] Jenna Iori et al. "Reliable operation of wind-storage systems for baseload power production". In: *Journal of Physics: Conference Series* 3025.1 (2025), p. 012023. DOI: 10.1088/1742-6596/3025/1/012023.
- [32] N. O. Jensen. *A Note on Wind Generator Interaction*. Tech. rep. Risø-M-2411. Roskilde, Denmark: Risø National Laboratory, 1983, p. 16.
- [33] Jonkman JM et al. "Definition of a 5MW Reference Wind Turbine for Offshore System Development". In: *National Renewable Energy Laboratory (NREL)* (Jan. 2009). DOI: 10.2172/947422.
- [34] Michael Koller et al. "Defining a Degradation Cost Function for Optimal Control of a Battery Energy Storage System". In: *2013 IEEE Grenoble PowerTech*. Grenoble, France: IEEE, 2013, pp. 1–6. ISBN: 978-1-4673-5669-5. DOI: 10.1109/PTC.2013.6652329.
- [35] Andrew B. Lambe and Joaquim R. R. A. Martins. "Extensions to the Design Structure Matrix for the Description of Multidisciplinary Design, Analysis, and Optimization Processes". In: *Structural and Multidisciplinary Optimization* 46 (2012), pp. 273–284. DOI: 10.1007/s00158-012-0763-y.
- [36] Jin-Oh Lee and Yun-Su Kim. "Novel battery degradation cost formulation for optimal scheduling of battery energy storage systems". In: *International Journal of Electrical Power & Energy Systems* 137 (2022), p. 107795. DOI: 10.1016/j.ijepes.2021.107795.
- [37] Zetao Ma et al. "Rapid sizing of a hydrogen-battery storage for an offshore wind farm using convex programming". In: *International Journal of Hydrogen Energy* 48.58 (2023), pp. 21946–21958. DOI: 10.1016/j.ijhydene.2023.03.037.
- [38] Dharik S. Mallapragada, Nestor A. Sepulveda, and Jesse D. Jenkins. "Long-run system value of battery energy storage in future grids with increasing wind and solar generation". In: *Applied Energy* 275 (2020), p. 115390. DOI: 10.1016/j.apenergy.2020.115390.
- [39] J.F. Manwell, J.G. McGowan, and A.L. Rogers. *Wind Energy Explained: Theory, Design and Application*. 2nd. Wiley, 2009.
- [40] Jon Martinez-Rico et al. "Multi-objective Optimization of Production Scheduling Using Particle Swarm Optimization Algorithm for Hybrid Renewable Power Plants with Battery Energy Storage System". In: *Journal of Modern Power Systems and Clean Energy* 9.2 (2021), pp. 285–294. DOI: 10.35833/MPCE.2019.000021.
- [41] Jon Martinez-Rico et al. "Sizing a Battery Energy Storage System for Hybrid Renewable Power Plants Based on Optimal Market Participation Under Different Market Scenarios". In: *IEEE Transactions on Industry Applications* 58.5 (2022), pp. 5624–5634. DOI: 10.1109/TIA.2022.3189331.
- [42] Joaquim R. R. A. Martins and Andrew Ning. *Engineering Design Optimization*. Available: <https://mdobook.github.io/>. Cambridge University Press, 2021.
- [43] Luca Moretti et al. "A design and dispatch optimization algorithm based on mixed integer linear programming for rural electrification". In: *Applied Energy* 233–234 (2019), pp. 1104–1121. DOI: 10.1016/j.apenergy.2018.09.194.
- [44] J. P. Murcia Leon et al. "HyDesign: a tool for sizing optimization of grid-connected hybrid power plants including wind, solar photovoltaic, and lithium-ion batteries". In: *Wind Energy Science* 9.4 (2024), pp. 759–776. DOI: 10.5194/wes-9-759-2024. URL: <https://wes.copernicus.org/articles/9/759/2024/>.
- [45] Lucas Murphy and Curran A. Crawford. "Data-driven classification of lithium-ion batteries for second-life applications". In: *Journal of Energy Storage* (2025).
- [46] Wayne B. Nelson. *Accelerated Testing: Statistical Models, Test Plans, and Data Analysis*. Pbk. Wiley Series in Probability and Statistics. Wiley-Interscience, 2004.
- [47] Amin Niayifar and Fernando Porté-Agel. "Analytical Modeling of Wind Farms: A New Approach for Power Prediction". In: *Energies* 9.9 (2016), p. 741. DOI: 10.3390/en9090741.
- [48] Simon E. J. O’Kane et al. "Lithium-ion battery degradation: how to model it". In: *Phys. Chem. Chem. Phys.* 24 (13 2022), pp. 7909–7922. DOI: 10.1039/D2CP00417H. URL: <http://dx.doi.org/10.1039/D2CP00417H>.

- [49] Józef Paska, Piotr Biczal, and Mariusz Kłos. “Hybrid power systems – An effective way of utilising primary energy sources”. In: *Renewable Energy* 34.11 (2009), pp. 2414–2421. ISSN: 0960-1481. DOI: <https://doi.org/10.1016/j.renene.2009.02.018>. URL: <https://www.sciencedirect.com/science/article/pii/S096014810900086X>.
- [50] Mads M. Pedersen et al. *PyWake 2.5.0: An open-source wind farm simulation tool*. Feb. 2023. URL: <https://gitlab.windenergy.dtu.dk/TOPFARM/PyWake>.
- [51] Yuliya Preger et al. “Degradation of Commercial Lithium-Ion Cells as a Function of Chemistry and Cycling Conditions”. In: *Journal of the Electrochemical Society* 167 (Sept. 2020). DOI: 10.1149/1945-7111/abae37.
- [52] Jorn M. Reniers, Grietus Mulder, and David A. Howey. “Improving optimal control of grid-connected lithium-ion batteries through more accurate battery and degradation modelling”. In: *Journal of Power Sources* 379 (2018), pp. 91–102. DOI: 10.1016/j.jpowsour.2018.01.004.
- [53] I. Rychlik. “A new definition of the rainflow cycle counting method”. In: *International Journal of Fatigue* 9.2 (Apr. 1987), pp. 119–121. ISSN: 0142-1123. DOI: 10.1016/0142-1123(87)90054-5.
- [54] Andoni Saez-de-Ibarra et al. “Co-Optimization of Storage System Sizing and Control Strategy for Intelligent Photovoltaic Power Plants Market Integration”. In: *IEEE Transactions on Sustainable Energy* 7.4 (2016), pp. 1749–1761. DOI: 10.1109/TSTE.2016.2555704.
- [55] Dirk Uwe Sauer and Heinz Wenzl. “Comparison of different approaches for lifetime prediction of electrochemical systems—Using lead-acid batteries as example”. In: *Journal of Power Sources* 176 (Feb. 2008), pp. 534–546. DOI: 10.1016/j.jpowsour.2007.08.057.
- [56] Oliver Schmidt et al. “The future cost of electrical energy storage based on experience rates”. In: *Nature Energy* 2.8 (2017), p. 17110. DOI: 10.1038/nenergy.2017.110.
- [57] Frank Sensfuß, Mario Ragwitz, and Massimo Genoese. “The merit-order effect: A detailed analysis of the price effect of renewable electricity generation on spot market prices in Germany”. In: *Energy Policy* 36.8 (2008), pp. 3086–3094. DOI: 10.1016/j.enpol.2008.03.035.
- [58] Kristen A. Severson et al. “Data-driven prediction of battery cycle life before capacity degradation”. In: *Nature Energy* 4.5 (2019), pp. 383–391. DOI: 10.1038/s41560-019-0356-8. URL: <https://doi.org/10.1038/s41560-019-0356-8>.
- [59] Yuanyuan Shi et al. “A Convex Cycle-based Degradation Model for Battery Energy Storage Planning and Operation”. In: *2018 Annual American Control Conference (ACC)*. 2018, pp. 4590–4596. DOI: 10.23919/ACC.2018.8431814. URL: <https://ieeexplore.ieee.org/document/8431814>.
- [60] Yuanyuan Shi et al. “Optimal Battery Control Under Cycle Aging Mechanisms in Pay for Performance Settings”. In: *IEEE Transactions on Automatic Control* 64.6 (2019), pp. 2324–2339. DOI: 10.1109/TAC.2018.2867507.
- [61] Thodoris Tsonopoulos. *Modeling of Battery Degradation for Integrated Design of Storage Systems: source code*. <https://github.com/ThodorisTson/battery-degradation-thesis>. 2026.
- [62] Pauli Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.
- [63] Bolun Xu et al. “Modeling of Lithium-Ion Battery Degradation for Cell Life Assessment”. In: *IEEE Transactions on Smart Grid* 99 (2016), pp. 1–1. DOI: 10.1109/TSG.2016.2578950. URL: https://www.researchgate.net/publication/303890624_Modeling_of_Lithium-Ion_Battery_Degradation_for_Cell_Life_Assessment.
- [64] Patrick Zippenfenig. *Open-Meteo.com Weather API*. Version 1.4.0. Dec. 2024. DOI: 10.5281/zenodo.14582479. URL: <https://doi.org/10.5281/zenodo.14582479>.



Justification for the Use of AI Tools

AI tools were used selectively throughout the preparation of this thesis, under the principle that they support but do not replace the author's own technical reasoning, methodological decisions, or analytical work. All degradation modeling concepts, optimization architecture decisions, implementation choices, and interpretations of results are the author's own. AI tools were not used to generate original engineering content or methodological ideas.

Their use falls into four categories.

Writing assistance. AI tools assisted in improving the clarity, conciseness, and academic tone of technically complex sections across all chapters. This included restructuring sentences for directness, identifying passive constructions, and applying the MIT effective writing principles used as the style guide throughout the thesis. The technical content, arguments, and conclusions in all sections were authored independently.

Code implementation. When implementing Python routines for the degradation pipeline, including the `build_half_cycle_map` function, the dual- Φ accumulation branches, the sub-gradient assembly, and the multi-year simulation loop, AI tools were used as an interactive assistant to convert well-defined algorithmic ideas into clean and syntactically correct code. Every computational concept and implementation decision originated from the author's own design. AI did not propose algorithms or methodological approaches.

LaTeX formatting. AI tools assisted with LaTeX compilation issues, equation formatting, figure placement, table structure, and consistent styling across chapters. This included resolving Overleaf-specific formatting errors and maintaining consistency in notation throughout the document.

Verification of understanding. AI tools were occasionally used to confirm understanding of published methods, such as the Shi et al. (2018) sub-gradient formulas and the Xu et al. (2016) stress function conventions, by rephrasing or explaining them in different terms. This use was limited to clarification and did not contribute new technical content.

All engineering insights, modeling decisions, validation strategies, and interpretations presented in this thesis are the author's own work. The use of AI tools was limited to refinement and implementation support within a framework entirely conceived and validated by the author.

B

Directional Wind Resource

The wake model fits a separate two-parameter Weibull distribution to each of twelve 30° direction sectors. The fitted scale (A) and shape (k) parameters and the sector frequencies are listed in Table B.1 and plotted in Figure B.1. The scale parameter varies by roughly a factor of two across sectors, from about 5 m/s in the rare north-easterly directions to over 10 m/s in the dominant south-westerly and westerly directions, while the shape parameter stays within a narrow band (2.16 to 3.00). The sectors therefore differ mainly in wind strength rather than in distribution shape.

Table B.1: Per-sector Weibull parameters and frequencies, 90 m hub height, 2022 (nearest ERA5 grid cell).

Sector center [°]	Frequency [%]	A [m/s]	k [-]
15	2.2	5.12	2.55
45	4.7	6.69	2.77
75	4.3	7.28	2.89
105	8.6	8.85	2.55
135	7.9	8.91	3.00
165	5.8	8.49	2.62
195	8.8	10.13	2.67
225	11.9	10.20	2.67
255	13.2	9.53	2.62
285	15.5	9.64	2.43
315	13.1	10.61	2.39
345	4.1	6.85	2.16
<i>Frequency-weighted mean</i>		9.22	2.59





Figure B.1: Per-sector Weibull distributions at 90 m hub height, 2022. The twelve colored curves are the directional fits (solid: sectors 15°–195°; dashed: 225°–345°); the black curve is the omnidirectional fit. Legend entries give the sector center and its frequency. Data are taken from the nearest ERA5 grid cell.



Gradient Verification

This appendix gives the derivation, figures, and numerical results for the gradient work summarized in Section 3.5.4. It has three parts. Section C.1 verifies the capacity gradient, the derivative of annual degradation with respect to pack size, which is the quantity a sizing loop would use. It is checked against a finite difference over a sweep of step sizes. Section C.2 verifies the per-timestep dispatch sub-gradient on the solved Year 1 trajectory. The two derivatives answer different questions: the first is a single number per year and is blind to how the sensitivity is distributed over the horizon, while the second is the distribution itself. Section C.3 examines a per-year signal that combines the sub-gradient with the dispatch shadow prices, and sets out why it is not used as a capacity-update signal.

C.1. Test 1: Frozen-Dispatch Capacity Gradient

Test 1 checks the capacity gradient of the annual degradation against a finite difference. It is the aggregate counterpart to the per-timestep check of Section C.2. The gradient has a cycle part and a calendar part, and both are checked by finite difference. The cycle part needs it, because rainflow counting makes its slope hard to predict in advance. The calendar part is a closed-form derivative, and the same sweep confirms it.

Method

The Year 1 dispatch is held fixed. Every value in the 8,760-step normalized SoC trajectory is multiplied by a factor s near one. Each cycle amplitude δ_i and each cycle mean SoC σ_i is an energy divided by the capacity, so scaling the trajectory by s scales every δ_i and every σ_i by s . This is the same change produced by dividing the capacity by s at a fixed dispatch. The scaling is therefore the capacity direction, and its derivative is the frozen-dispatch capacity gradient.

The cycle degradation is:

$$f_d^{\text{cyc}} = \sum_i n_i \Phi(\delta_i) S_\sigma(\sigma_i), \quad \Phi(\delta) = k_3 \delta^{k_4}, \quad S_\sigma(\sigma) = \exp(k_\sigma (\sigma - \sigma_{\text{ref}})), \quad (\text{C.1})$$

where Φ is the amplitude power law and S_σ is the mean-SoC stress factor. The temperature factor equals one at the fixed cell temperature of 25 °C and is omitted. Differentiating with respect to s at $s = 1$, with $\delta_i \rightarrow s \delta_i$ and $\sigma_i \rightarrow s \sigma_i$, gives the cycle slope

$$\left. \frac{df_d^{\text{cyc}}}{ds} \right|_{s=1} = \sum_i f_{d,i} (k_4 + k_\sigma \sigma_i), \quad f_{d,i} = n_i \Phi(\delta_i) S_\sigma(\sigma_i). \quad (\text{C.2})$$

The two factors have a direct reading. The term k_4 is the response of the power law δ^{k_4} to the scaling. The term $k_\sigma \sigma_i$ is the response of the stress factor S_σ , which changes because the scaling moves each cycle mean.

The values $f_{d,i}$ are read from the rainflow cycles of the model, so the quantity checked is the degradation the model computes and not a separate reconstruction. The implemented gradient is compared against this expression directly, and the two agree to eleven digits.

The calendar degradation is $f_d^{\text{cal}} = S_\tau(t) S_\sigma(\bar{\sigma})$, where $S_\tau(t) = k_\tau t$ grows in proportion to time and $\bar{\sigma}$ is the mean SoC over the year. Only $\bar{\sigma}$ depends on the scaling. With $\bar{\sigma} \rightarrow s \bar{\sigma}$ and $dS_\sigma/d\bar{\sigma} = k_\sigma S_\sigma(\bar{\sigma})$, the calendar slope is

$$\left. \frac{df_d^{\text{cal}}}{ds} \right|_{s=1} = S_\tau(t) k_\sigma S_\sigma(\bar{\sigma}) \bar{\sigma} = k_\sigma \bar{\sigma} f_d^{\text{cal}}. \quad (\text{C.3})$$

The same sweep confirms it, as reported below..

The finite difference is evaluated over a wide range of step sizes h , from 10^{-1} to 10^{-13} , using the central form

$$D_h = \frac{f_d^{\text{cyc}}(1+h) - f_d^{\text{cyc}}(1-h)}{2h} \quad (\text{C.4})$$

and the one-sided forward form $D_h^+ = [f_d^{\text{cyc}}(1+h) - f_d^{\text{cyc}}(1)]/h$. The trajectory is rescaled and the rainflow count is repeated at each step, so the check runs from end to end.

Sweeping h separates two error sources. On the descent, at larger h , the estimate carries a truncation error that falls with the step, at second order for the central difference and at first order for the forward difference. On the rise, at smaller h , it carries a rounding error that grows, because the two close values in the numerator lose significant digits when subtracted and the small divisor magnifies the loss. The best step lies between the two.

The frozen-dispatch capacity gradient combines the two slopes. Scaling the normalized trajectory by s is equivalent to dividing the capacity by s , so $d/d\bar{E} = -(1/\bar{E}) d/ds$. The gradient is therefore

$$\left. \frac{df_d}{d\bar{E}} \right|_{\text{frozen}} = -\frac{1}{\bar{E}} \left[\sum_i f_{d,i} (k_4 + k_\sigma \sigma_i) + k_\sigma \bar{\sigma} f_d^{\text{cal}} \right]. \quad (\text{C.5})$$

The implemented gradient `dDegCost_dEcap` evaluates this expression.

Result

Figure C.1 shows the finite-difference error against the cycle slope of Equation C.2, on logarithmic axes, for the two price years and for both the central and the forward difference. Both curves have the same overall shape, a descent then a rise, which is the standard form for a numerical gradient. On the descent the two guide lines confirm the expected convergence: the central difference falls at second order in h and the forward difference at first order, so the central difference is the more accurate at every step. On the rise, rounding error grows as the step shrinks further. The central difference reaches its minimum at $h = 1.0 \times 10^{-5}$ for DK 2019 and $h = 2.4 \times 10^{-6}$ for DK 2022, at relative errors of 3.5×10^{-12} and 1.7×10^{-12} . Both sit near $\epsilon^{1/3} = 6 \times 10^{-6}$, the step at which the two error sources balance for a central difference in double precision. The forward difference reaches its minimum at $h = 1.7 \times 10^{-8}$ in both years, against $\sqrt{\epsilon} = 1.5 \times 10^{-8}$. The sweep therefore reproduces the predicted optimal steps as well as the predicted slopes, which confirms that Equation C.2 is the exact slope of the cycle degradation.

The descent and the minimum test the gradient; the rise does not. On the descent the error is the truncation error of the difference formula, so its slope tests whether Equation C.2 is the true derivative. On the rise the error comes from subtracting two nearly equal numbers and dividing by a small step, which is a limit of double precision rather than a property of the derivative. The full step range and both difference forms are shown so that the two regimes and the transition between them are visible.

The remaining quantities are evaluated directly. For DK 2019 the model gives a cycle degradation $f_d^{\text{cyc}} = 0.01093$, a calendar degradation $f_d^{\text{cal}} = 0.01268$, and a mean SoC $\bar{\sigma} = 0.472$. Evaluating the cycle slope of Equation C.2 on the rainflow cycles gives 0.01857, the value the finite difference confirms. The calendar slope of Equation C.3 follows from the same numbers, $k_\sigma \bar{\sigma} f_d^{\text{cal}} = 1.04 \times 0.472 \times 0.01268 = 0.00622$. Their sum is the bracket of Equation C.5, the total frozen-dispatch slope,




Figure C.1: Finite-difference error of the cycle-degradation slope against the analytical slope of Equation C.2, over the full range of step sizes h , for the central and the forward difference. On the descent, truncation error falls at second order in h for the central difference and at first order for the forward difference, marked by the guide lines. On the rise, rounding error grows as the step shrinks. The central-difference minimum reaches order 10^{-12} near $\epsilon^{1/3} = 6 \times 10^{-6}$, which is agreement to about eleven digits.

$0.01857 + 0.00622 = 0.02479$. Dividing by the capacity and changing sign gives the capacity gradient. DK 2022 follows the same steps, and both years are collected in Table C.1.

Table C.1: Test 1 quantities for the two price years. The cycle and calendar degradations are model outputs. The cycle slope is Equation C.2, checked by the finite difference of Figure C.1. The calendar coupling is Equation C.3. Their sum is the bracket of Equation C.5.

Quantity	DK 2019	DK 2022
Rainflow records	916	980
Records of zero amplitude	158	171
Summed multiplicity $\sum_i n_i$	717.0	776.5
Cycle degradation f_d^{cyc}	0.01093	0.01160
Calendar degradation f_d^{cal}	0.01268	0.01284
Mean SoC $\bar{\sigma}$	0.472	0.484
Φ -weighted mean SoC σ_Φ	0.5008	0.5069
Cycle slope, Equation C.2	0.01857	0.01978
Depth-only cycle slope $k_4 f_d^{cyc}$	0.01288	0.01367
Calendar coupling, Equation C.3	0.00622	0.00647
Total frozen-dispatch slope, Equation C.5	0.02479	0.02625
Central-difference minimum relative error	3.5×10^{-12}	1.7×10^{-12}
Step size at that minimum	1.0×10^{-5}	2.4×10^{-6}
Forward-difference minimum relative error	7.3×10^{-9}	1.0×10^{-8}
Step size at that minimum	1.7×10^{-8}	1.7×10^{-8}
Minimum with the calendar term included	2.1×10^{-12}	4.7×10^{-12}
Implemented gradient against Equation C.5	9.4×10^{-12}	9.2×10^{-12}

Interpretation

The finite difference verifies the two parts of the cycle slope together, and the calendar slope adds a third; the three form the capacity gradient of Equation C.5, which the implemented `dDegCost_dEcap` evaluates.

The mean-SoC coupling is not a small correction. Dropping the $k_\sigma \sigma_i$ term leaves the depth-only slope $k_4 f_d^{cyc}$, which is 0.01288 for DK 2019 and 0.01367 for DK 2022, or 30.6% and 30.9% below the complete

slope. The finite difference converges to the complete slope and not to the depth-only one, which is the evidence that the coupling belongs in the capacity gradient. Written out, the share is $k_\sigma \sigma_\Phi / (k_4 + k_\sigma \sigma_\Phi)$, where σ_Φ is the Φ -weighted mean SoC. It therefore depends only on the fitted exponent, the stress coefficient, and where the weighted mean SoC sits. For the two price years σ_Φ is 0.5008 and 0.5069. Both lie close to the center of the 10 to 90% window. That is why the share holds in both years although the two dispatches differ in cycle count and in amplitude distribution.

The calendar slope is also checked and not only stated. Repeating the sweep on $f_d^{\text{cyc}} + f_d^{\text{cal}}$ against the full bracket of Equation C.5 gives a central-difference minimum of 2.1×10^{-12} for DK 2019 and 4.7×10^{-12} for DK 2022, at the same step sizes and the same order as the cycle-only sweep. The calendar slope is a quarter of the total slope, so a one percent error in it would appear at 2.5×10^{-3} . The measured value is nine orders below that. Equation C.3 is the exact response of the calendar term to the scaling, evaluated on the production calendar function.

The implemented gradient evaluates Equation C.5 rather than something close to it. Comparing the two directly gives a relative difference of 9.4×10^{-12} for DK 2019 and 9.2×10^{-12} for DK 2022. The residual is fully accounted for. The two sides use equivalent expressions for the amplitude term, $k_4 \Phi(\delta_i)$ here and $\Phi'(\delta_i) \delta_i$ in the implementation, and the polynomial surrogate Φ clips its amplitude argument at 10^{-9} . This is a different clip from the one on the amplitude stress function S_δ , which acts at 10^{-6} (Appendix D). The 158 and 171 records of zero amplitude therefore carry a small cost in the first form and none in the second, and that contribution accounts for the reported differences to four digits. The two expressions agree to within one unit in the last place on every record that carries wear.

The scatter within the rise comes from the cycle counter rather than from the gradient. The rainflow records are recounted at four points of the sweep. The record count changes by up to 35 of 916 and 31 of 980. The summed multiplicity, the amplitude-weighted throughput, and the amplitude-weighted wear are unchanged to 3×10^{-16} . A floating-point tie changes whether the counter reports one full cycle or two half cycles of the same amplitude, and since two half cycles of weight $\frac{1}{2}$ carry the same wear as one full cycle, no degradation moves. The counting is therefore stable in every quantity the model uses, and unstable only in how many records it emits.

C.2. Test 2: Per-Timestep Dispatch Sub-gradient

Test 2 checks the per-timestep dispatch sub-gradient of Equations 2.23 and 2.24. It differs from Test 1 in three ways. Test 1 perturbs the energy capacity $\bar{E}$, a single scalar, and returns one number per year; Test 2 perturbs the charging power at one time step out of 8,760 and returns one number per time step. The two are different partial derivatives of the same cost, one with respect to a design variable and one with respect to the operating variables. Test 1 sums the second over the horizon, so it cannot separate two sensitivity vectors that share a sum, and the inner loop of the nested architecture uses the vector rather than the sum. The third difference decides the criterion. Scaling the whole trajectory keeps the rainflow pairing intact, so the cost is smooth in the direction Test 1 perturbs. Perturbing one time step does not, and at roughly half of them no derivative exists. Test 2 therefore asks whether the analytic value lies between two one-sided derivatives, not whether it equals one. Both price years are run, as in Test 1.

Method

The test is a one-sided perturbation test, applied one time step at a time. For a chosen charging time step t , the charging power p_t^{ch} is raised slightly and the annual degradation cost is recomputed from the start, rainflow counting included. The same is then done with p_t^{ch} lowered. Dividing each change in cost by the size of the perturbation gives two difference quotients, one from each side:

$$D^+ = \frac{f(e + \Delta e) - f(e)}{h}, \quad D^- = \frac{f(e) - f(e - \Delta e)}{h}, \quad \Delta e_{t'} = \begin{cases} \Delta t \eta_{\text{in}} h, & t' > t \\ 0, & t' \leq t \end{cases} \quad (\text{C.6})$$

Raising p_t^{ch} raises the battery energy at every later time step by the same amount and leaves every earlier one unchanged, which is what the third expression states and what defines the spanning set $\mathcal{S}(t)$

of Section 2.8.2. The perturbation h is expressed as a fraction of the pack size so that it is dimensionless. The reference perturbation is 10^{-6} , which is 3×10^{-4} MWh on the 300 MWh pack.

Two hundred charging time steps are sampled at regular intervals from the solved Year 1 trajectory. At each one the two quotients are computed and compared, first with each other and then with the analytic value. Comparing them with each other splits the sample into two populations, and the two populations are judged by different criteria.

Where the two quotients agree, the cost is differentiable at that time step and the analytic value must equal them, which is the ordinary case.

Where they disagree, the cost has a corner at that time step and no single derivative exists. A central difference returns the mean of the two quotients. At a corner that is the midpoint of an interval whose ends are the two answers, so it cannot equal either of them however small the perturbation is made. Equality is therefore the wrong criterion, and the criterion becomes membership. The amplitude-only cost is the form proved convex in Section 2.7.3, and for a convex function of one variable the sub-differential at a corner is the closed interval between the left and right derivatives. Every element of that interval is a valid sub-gradient, so the analytic value s_t must satisfy

$$\min(D^-, D^+) \leq s_t \leq \max(D^-, D^+). \quad (\text{C.7})$$

At a finite perturbation this interval contains the sub-differential rather than equaling it, so Equation C.7 is a necessary condition that tightens as the perturbation shrinks.

Where the corners come from.

A corner is a topology-change event in the sense of Section 3.8.2, where two shallow half-cycles merge into one deeper cycle or a deep cycle splits into two shallower ones. It is a change in how the counter groups the excursions, not a change in the battery. Because the perturbation shifts every value after t by the same amount, every excursion lying entirely after t keeps its amplitude exactly, and the excursion from t to $t+1$ is the only one that moves. The counter decides whether a small excursion is a cycle of its own or part of a larger one by comparing the amplitudes of adjacent excursions. A change in that one amplitude can flip the comparison and reassign the excursion. Figure C.2 shows the two readings. At the size where the comparison is tied, both readings give the same total wear, which is why the cost stays continuous. They do not give the same rate of change, because Φ' grows with amplitude and the two readings assign the perturbation to cycles of different amplitude. The two rates are D^- and D^+ .




Figure C.2: Schematic of a rainflow re-pairing. Left: a state-of-charge trace with one ambiguous excursion, and the two readings the counter can give it. Right: the degradation cost against p_t^{ch} at step t . The two readings meet where the comparison is tied, so the cost is continuous, but they arrive with the different slopes D^- and D^+ of Equation C.6. The interval between those slopes is the bracket of Equation C.7. The figure carries no data from a run.

Why no step-size sweep.

Test 1 sweeps the step size and reports the minimum of the resulting curve. That approach does not transfer, because at a corner there is no derivative for the difference to approach and the sweep has no target. On the differentiable time steps a sweep is possible but shows nothing. The two quotients never agree to better than about 6×10^{-7} relative at any perturbation. Curvature sets a floor above 10^{-6} and the arithmetic of the subtraction sets one below it, so the curve has no descending branch. The single comparison at the reference perturbation is reported instead, together with the arithmetic floor of the quotient at that perturbation, which is what decides whether a residual belongs to the analytic value or to the finite difference.

Result

All quantities come from the solved Year 1 trajectories at the sizing point, 150 MW and 300 MWh over 8,760 steps at the 10 to 90% window. The replacement cost is $\lambda_E^{\text{repl}} = 72$ EUR/kWh, so the sub-gradient has units of EUR/MW. The annual degradation cost being differentiated is 236,408 EUR under DK 2019 prices and 249,451 EUR under DK 2022 prices. Table C.2 collects both years.

Table C.2: Test 2 quantities on the solved Year 1 trajectories of the two price years. The two populations are separated at the reference perturbation of 10^{-6} fractional SoC. Relative disagreements are normalized by the larger of the analytic value and the central difference, with a floor at one thousandth of the median analytic magnitude, so that time steps where the sensitivity passes through zero do not inflate the ratio.

Quantity	DK 2019	DK 2022
Rainflow records	916	980
Median analytic magnitude	1.261 EUR/MW	1.261 EUR/MW
Sampled charging time steps	200	200
Differentiable at the reference perturbation	103	77
Flagged non-differentiable	97	123
Disagreement above 1%	87	118
Differentiable steps, median relative disagreement	1.79×10^{-7}	1.42×10^{-7}
Differentiable steps, largest	2.16×10^{-7}	1.79×10^{-7}
Arithmetic floor of the quotient at that perturbation	1.32×10^{-7}	1.40×10^{-7}
Corner steps, median against a central difference	0.33	0.24
Corner steps, median against the nearer quotient	2.25×10^{-7}	1.51×10^{-7}
Analytic value inside the bracket, Equation C.7	97 of 97	123 of 123
Analytic value at the lower quotient	25	30
Analytic value at the upper quotient	60	76
Analytic value strictly inside the bracket	12	17
Largest deviation from $1/(\eta_{\text{in}}\eta_{\text{out}})$, all 8,759 steps	4.4×10^{-16}	2.2×10^{-16}
Spanning set size $ \mathcal{S}(t) $	0 to 4, median 1	0 to 4, median 1
Steps spanned by exactly one cycle	58.8%	60.8%
Local term against mean-SoC-weighted derivative	factor 293	factor 305

On the differentiable time steps the analytic value agrees with the central difference to a median of 1.79×10^{-7} for DK 2019 and 1.42×10^{-7} for DK 2022, and no time step in either year disagrees by more than one percent. The arithmetic floor of the quotient at the same perturbation is 1.32×10^{-7} and 1.40×10^{-7} , the same order as those medians. The residual is therefore the precision of the finite difference and not of the analytic value.

On the corner time steps the analytic value lies inside the bracket of Equation C.7 at every one, 97 of 97 for DK 2019 and 123 of 123 for DK 2022. It coincides with one of the two quotients at 85 and at 106 of them, to within one thousandth of the bracket width. That is 87.6% and 86.2% of each population. The remaining 12 and 17 lie strictly inside. Measured against the nearer of the two quotients the corner steps sit at a median of 2.25×10^{-7} and 1.51×10^{-7} , the same order as the differentiable population. Measured against a central difference they sit at 0.33 and 0.24, six decades higher.

Figure C.3 shows 48 corner time steps from each year, ordered by where the analytic value sits between the two quotients, so the lower-endpoint steps are on the left and the upper-endpoint steps on the right. The central difference sits at the middle of every bracket. Where the two quotients are equal and opposite the central difference is zero, while the sub-gradient is close to the median magnitude of 1.261 EUR/MW.




Figure C.3: Sub-differential bracket at 48 sampled corner time steps from each price year, ordered by the position of the analytic value between the two quotients of Equation C.6. The grey bar spans the two quotients and its ends are those two values. Over the full populations the analytic sub-gradient lies inside the bar at every time step, at an endpoint in 85 of 97 for DK 2019 and 106 of 123 for DK 2022. The central difference is the mean of the two quotients and lies at the middle of the bar throughout.

The efficiency identity of Equation 2.26 holds at every time step, not on average. The ratio of the discharge sensitivity to the charge sensitivity equals $1/(\eta_{\text{in}}\eta_{\text{out}}) = 1.0989010312$ at all 8,759 time steps where the sensitivity is non-zero. The largest deviation is 4.4×10^{-16} for DK 2019 and 2.2×10^{-16} for DK 2022.

Two further quantities come from the same runs. The spanning set $\mathcal{S}(t)$ holds between zero and four cycles, with a median of one, and holds exactly one at 58.8% of the horizon under DK 2019 prices and 60.8% under DK 2022 prices. The local amplitude term, compared against a central difference of the cost with the mean-SoC stress factor applied on the differentiable time steps only, is smaller by a factor of 293 and 305

Interpretation

The two populations answer different questions and both answers are needed. Where the cost is differentiable the analytic value is the derivative, to the precision the finite difference can resolve. Where it is not, the analytic value is an element of the sub-differential, and coincides with one of the two one-sided derivatives in roughly seven cases out of eight. Both hold in both price years, on dispatches that differ in cycle count and in amplitude distribution, so neither is an artifact of one price series. A single measure pooled over the two populations would report the first result and hide the second.

The size of the corner population is a property of the problem rather than of the test, and it is not the same in the two years. Under DK 2019 prices, 87 of the 200 sampled time steps, or 44%, disagree with a central difference by more than one percent, which curvature cannot produce; under DK 2022 prices it is 118, or 59%. The reason the count is high in both is that the dispatch is a linear program. An optimal schedule saturates its limits: the state of charge rests at a window bound for long stretches, and charging and discharging run at the rated 150 MW for whole time steps, one of which moves 0.4770 of nominal capacity. Excursions of exactly equal amplitude are therefore common, and an exactly tied comparison is the condition under which any perturbation flips the counter. The same mechanism produces the 158 and 171 records of zero amplitude in Table C.1, where consecutive turning points land on the same value at a bound.

The gap between the two years follows from the same mechanism. The more volatile 2022 prices produce more cycles in the same horizon, 776.5 against 717.0, from 980 records against 916, so there are more comparisons between adjacent excursions and more of them fall on a tie. On these two years the problem therefore becomes less smooth as the market becomes more volatile, which is the regime in which storage arbitrage is most valuable. Two price years do not establish that as a general trend, but the direction follows from the mechanism. Either way, between roughly a half and three fifths of the charging time steps in one year of dispatch sit where the degradation cost is not differentiable. This is the per-timestep form of the non-smoothness reported in Section 3.8.2, and it is why a smooth nonlinear solver does not converge on the monolithic formulation.

The span count justifies the spanning-set construction of Section 2.8.2. The index-set partition of Shi et al. [59] assigns each time step to a single half-cycle, which is exact only where one cycle spans the step. That holds at 58.8% and 60.8% of the horizon here. At the remaining time steps either no cycle spans the step, and the sensitivity is zero, or several nested cycles do, and the sum has more than one term. A construction that assigns each time step to a single owning cycle cannot represent the second case. The spanning-set form reduces to Shi's expression where their assumption holds and stays correct where it does not.

The factor of 293 and 305 is the measured basis for excluding the mean-SoC stress from the gradient path, as set out in Section 2.8.4. Cycle amplitudes are invariant to a uniform shift of the SoC trajectory, so the amplitude term is local to the cycles spanning the perturbed time step. Cycle means are not invariant, so adding the mean-SoC factor gives the derivative a contribution from every cycle downstream, and on an annual dispatch that non-local contribution is two orders of magnitude larger than the local one. Keeping the gradient path on cycle amplitude alone is what keeps it local, and what keeps it inside the convexity result the construction relies on. The mean-SoC and temperature factors remain in the reporting path and in the capacity gradient of Section C.1, neither of which carries a convexity claim.

C.3. Fused Per-Year Sensitivity

This section examines a single per-year quantity that combines the degradation sub-gradient with the dispatch shadow prices. It was an early candidate signal for updating the capacity between years. It is reported here because it explains why the capacity gradient of Section 3.8.3 is assembled from separate analytical terms rather than from this quantity.

Definition

At each simulation year the sub-gradient machinery of Section 2.8 produces a fused per-year sensitivity,

$$g_k = -\bar{E}_k \langle s_k, \mu_k \rangle, \quad (\text{C.8})$$

in which s_k is the degradation sub-gradient over year k , μ_k is the vector of LP dual prices on the lower state-of-charge bound, and $\bar{E}_k$ is the effective energy capacity.

The two vectors carry different information. The sub-gradient s_k is large at time steps inside deep cycles. The dual μ_k is large at time steps where the lower state-of-charge limit binds. Their inner product is large only when the two coincide in time, so g_k measures whether the hours that matter most for degradation are also the hours that matter most for revenue, scaled by capacity. It was examined as a capacity-update signal on that basis. The analysis uses DK1 2022, the more volatile year, where the effect is largest.

Result

Figure C.4 shows g_k for each year of the first battery under DK1 2022 prices. The sub-gradient is computed on the polynomial model, so the figures use the baseline run of Section 4.3 evaluated on the Shi branch, where the first battery lasts twelve years rather than the eleven the Xu branch gives.

Its composition is found by writing g_k as a product of a magnitude term and an alignment term,

$$g_k = \underbrace{\bar{E}_k |\bar{s}_k| |\bar{\mu}_k|}_{\text{magnitude}} \times \underbrace{\frac{\langle s_k, \mu_k \rangle}{|\bar{s}_k| |\bar{\mu}_k|}}_{\text{alignment}}, \quad (\text{C.9})$$

where $|\bar{s}_k|$ and $|\bar{\mu}_k|$ are the mean absolute values of the two vectors. The first term measures the size of the two vectors and the second measures how their large entries line up in time.

Normalized to their annual means, the two terms behave differently. The magnitude term declines as capacity fades, from 1.18 to 0.90 times its mean, and its correlation with g_k is +0.92. The alignment term stays within 6.4% of its mean across the twelve years, and its correlation with g_k is +0.38. The decline of the magnitude term therefore sets both the trend and most of the year-to-year variation.

Figure C.5 repeats the bars of Figure C.4 at reduced opacity and draws the two decomposition terms on top. The magnitude term tracks the bar heights, falling from year one to year twelve with the effective capacity. The alignment term oscillates about its mean with little correspondence to the bar heights: year 11 carries the lowest bar of the twelve and the third-highest alignment value.




Figure C.4: Fused per-year sensitivity g_k , DK1 2022 prices, years 1 to 12, normalized to its annual mean. The dotted line marks the annual mean.




Figure C.5: The magnitude and alignment terms overlaid on the g_k bars of Figure C.4 at reduced opacity, DK1 2022 prices. The magnitude term follows the bar heights; the alignment term does not.

Interpretation

Two features of this decomposition argue against g_k as a capacity-update signal. It is dominated by the magnitude term, which is the product of the effective capacity and the mean sizes of the two vectors, so it reports the capacity fade rather than anything about how degradation and revenue coincide in time. And the alignment term, which is the only part that carries that information, contributes little and varies with the temporal placement of the sub-gradient rather than with any change in the battery.

An earlier version of this analysis, computed with the half-cycle attribution map that the spanning-set construction of Section 2.8.2 replaces, gave a different picture on the same run: g_k swung by a factor of 2.9 across the generation and the alignment term correlated with it at +0.97 against +0.17 for the magnitude term. The two sub-gradients agree at 88% of time steps and differ almost entirely at the steps where cycles nest, which are the steps a single-owner attribution cannot represent. Most of the year-to-year variation that first motivated examining this signal was therefore an artifact of that attribution rather than a property of the dispatch.

Either way, a single year of g_k is not a reliable signal for capacity updating. Under the attribution map it varied too much to be attributed to any change in the battery; under the exact construction it varies too little to carry information beyond the capacity fade, which the capital cost term already supplies. The capacity gradient used in this thesis avoids the question: its degradation term sums the per-cycle contributions analytically rather than contracting the sub-gradient against the dual, so neither the temporal alignment nor its dependence on the attribution rule enters. That gradient is the quantity checked in Test 1, Section C.1.

D

Verification on Synthetic Input

This appendix reports the verification of the rainflow counter and of the dispatch on synthetic input, where the correct answer is available before the code runs. It complements Section 4.2, which verifies the same implementation on one week of a solved year of real data, and Appendix C, which verifies the two gradients. No measured cell or pack data is used anywhere in this thesis, so nothing here or elsewhere is a validation against a physical battery.

Only the results that other parts of the thesis rely on are reported. Comparisons of a closed form against the same closed form are omitted, because they test the arithmetic of the reference rather than the implementation.

D.1. Rainflow Counter on Synthetic Traces

Two traces are used, both at the 300 MWh reference capacity and the 10 to 90% window. The constant-amplitude trace is ten identical charge and discharge cycles between the window bounds, closing at the lower bound, so every cycle has amplitude $\delta = 0.80$ and mean state of charge $\bar{\sigma} = 0.50$ by construction. The nested trace is one excursion from the lower to the upper bound and back, carrying four ripples of amplitude $\delta = 0.10$ on its charging leg at mean states of charge of 0.25, 0.35, 0.45 and 0.55. The second trace is the case that separates rainflow counting from a peak count, because the ripples must be extracted without destroying the excursion that contains them. A third trace holds the state of charge constant. Figure D.1 shows the first two with the counted cycles drawn on them.

On the constant-amplitude trace the counter returns 20 records, each a half cycle of count 0.5, for a total count of 10. The three-point rule of Section 2.4 never lets the residue exceed three points on a sequence of constant amplitude, so it reports half cycles throughout rather than full cycles. The number of records is therefore twice the number of physical cycles while the total count equals it. This is why record counts and cycle counts are reported separately throughout Chapter 4 and Appendix C, where the solved Year 1 trajectory gives 916 records against a total count of 717.0.

On the nested trace the counter returns four full cycles at $\delta = 0.10$ with the four expected mean states of charge, plus two half cycles at $\delta = 0.80$ and $\bar{\sigma} = 0.50$, for a total count of 5. The result does not change when the ramps between turning points are sampled more finely, which confirms that only the turning points enter. Nesting of this kind is what allows more than one cycle to span a single time step, which is the case the spanning-set sub-gradient of Section C.2 has to represent and a single-owner attribution cannot.

The counter itself is checked against the worked example given in ASTM E1049 [2], and the two copies of it in the code base return identical records on both traces. The second point is a precondition for the Xu and Shi branches of Section 4.4.1 being comparable.

The constant-state-of-charge trace returns a cycle term of 3.5×10^{-9} rather than zero. A trace with no reversals still closes its residue as one half cycle of zero amplitude, and the amplitude stress function S_δ clips its argument at 10^{-6} instead of returning zero, so the record contributes $0.5 S_\delta(10^{-6})$. The value

is 3.3 parts per million of the calendar term on the same trace and is neglected. It is reported because it is a property of the implementation rather than rounding, and because the same mechanism, at a different threshold, accounts for the residual in the gradient comparison of Appendix C.

The polynomial surrogate is refitted at run time from the 10 to 90% window and returns $k_3 = 3.2417795 \times 10^{-5}$ and $k_4 = 1.1785286$, reproducing Equation 3.10 to the precision at which it is quoted. The fitted exponent satisfies the convexity condition $k_4 > 1$ of Section 2.7.1, and at the center of the window $S_\sigma(\bar{\sigma}) = 1$ exactly. Against the Xu reference curve the surrogate returns 83.6% of S_δ at $\delta = 0.80$ and 68.9% at $\delta = 0.10$. The second point lies below the lower fitting bound of 0.15 (Section 3.5.2), so the surrogate is extrapolating there and understates the damage of shallow cycles by about 31%. The error is one-signed and optimistic, and it grows as the cycle gets shallower. These two values show the divergence for a single cycle. Figure 3.8 shows it across the full amplitude range, and Table 4.9 over a year.

The synthetic traces are exact by construction, so this section does not exercise the shallow phantom cycles that slack state-of-charge constraints inject into a solved dispatch trace (Section 3.4.3). Those are a property of the optimization and not of the counter.




Figure D.1: Rainflow read-back on the two synthetic state-of-charge traces. Top: three periods of the constant-amplitude trace, with the 10% and 90% window bounds shown as dashed lines. Bottom: the nested trace, with each counted cycle drawn as a vertical bar spanning its amplitude, placed at the midpoint of the interval it was extracted from. The four ripples at $\delta = 0.10$ are resolved separately from the excursion at $\delta = 0.80$.

D.2. Efficiency Threshold in the Dispatch

The dispatch is tested on a price signal whose optimal schedule is known before the solver runs. The wind series is held constant at 160 MW. The battery is the 300 MWh reference unit at the 10 to 90% window, rated at 150 MW. The grid limit is 325 MW, which leaves headroom for the battery to discharge at rated power while the wind is producing.

The test is a swept two-level price, twelve cheap hours followed by twelve expensive hours, whose ratio

is moved across the point at which arbitrage becomes profitable. That point is a closed form. Charging stores η_{in} of the energy diverted and discharging draws $1/\eta_{\text{out}}$ from the store, so moving one unit of energy from the cheap block to the expensive block pays only if the price ratio exceeds $1/(\eta_{\text{in}}\eta_{\text{out}})$. With the symmetric efficiencies of Section 3.4.3 this is 1.0989010. The prediction is falsifiable in a way a comparison against a closed form is not: a misplaced efficiency anywhere in the kernel moves the threshold, and the sweep locates it.

Figure D.2 shows the result. The battery is idle at every ratio at or below 1.0975 and saturates the window at every ratio at or above 1.0995, discharging 686.83625 MWh over three days against a closed form of $3\bar{E}(\sigma_{\text{max}} - \sigma_{\text{min}})\eta_{\text{out}}$. The transition therefore brackets the predicted value within 0.18%. Six of the thirteen sampled ratios lie between the idle and the saturated band. They are plotted but not asserted. Close to the threshold the linear program is nearly degenerate, so the outcome would test the solver tolerance rather than the model.

The same constant is recovered independently in Appendix C. There the ratio of the discharge sensitivity to the charge sensitivity of the per-timestep sub-gradient equals $1/(\eta_{\text{in}}\eta_{\text{out}})$ at every time step of a solved year, to 4.4×10^{-16} . One result comes from the behavior of the linear program and the other from the algebra of the derivative, and they agree.

A sinusoidal price over seven days is used as a second case, and rainflow counting on the resulting trace returns exactly seven deep cycles, one per price period, at $\delta = 0.80$ and $\bar{\sigma} = 0.50$. This closes the loop between the two sections: the dispatch produces the cycle structure that the counter was verified to read correctly.

Two scope statements follow from the setup. On these dispatches the window saturates once per day, so the battery rests at the lower or the upper bound and never at an interior level. The diagnostic counting idle steps strictly inside the window returns zero, and no energy is discarded, so the phantom cycle mechanism is not exercised here either. And the wind series is constant and never exceeds the grid limit, so no case curtails. The lifetime net present value reported in Chapter 4 uses the marginal revenue basis, which includes curtailment recovery. At this site that term is worth about 52 kEUR, or 0.02% of net present value, because the grid connection equals the installed wind rating.




Figure D.2: Energy discharged over three days against the price ratio between the expensive and the cheap block. Left: the full sweep. Right: detail around the transition, markers only, since no ratio between the two was solved. The battery is idle at 1.0975 and saturates at 1.0995, bracketing the predicted threshold $1/(\eta_{\text{in}}\eta_{\text{out}}) = 1.0989$ within 0.18%. Six of the thirteen sampled ratios lie between the idle and the saturated band and are shown but not asserted, because the linear program is close to degenerate in a narrow interval around the threshold.

E

Monolithic NLP Supporting Material

E.1. Per-Day Exemplar Dispatch Panels

Figures E.1 to E.4 show one representative day for each of the four behavior types of Section 4.5.2: Type A a converged gain, Type B a converged loss, Type C a stalled solve, and Type D a diverged solve. In each panel the upper plot is the day-ahead price and the lower plot is the state-of-charge trajectory for the LP (solid) and the NLP (dashed), as a fraction of installed energy capacity, with the 0.10 and 0.90 operating limits marked by dotted lines. Each figure is preceded by a short description of the day it shows. The caption gives the change in objective value against the LP, the revenue difference, the degradation cost difference, the final optimality, and the solver termination reason. The termination reason is one of three: `gtol`, the first-order optimality condition was met; `xtol`, the step length fell below tolerance before an optimum was certified; and `maxiter`, the solver reached its iteration limit.

The representative day for Type A, a converged gain, is D049, a low-variability day on which the solver reached a first-order optimum that improves on the LP. The two trajectories coincide through the morning and midday, when both run a deep cycle. They differ in the afternoon. Between hours 14 and 22 the LP charges to a peak near 0.71 and the NLP holds a shallower peak near 0.63, then returns to the lower limit sooner. Removing this one shallow afternoon cycle lowers cycle degradation by more than the revenue it gives up, so the objective change is positive. This is one of seventeen converged days that improve on the LP, and the improvement comes from dropping a single low-value cycle rather than from a different operating strategy.




Figure E.1: Type A, converged gain: day D049 (Low regime). Converged solve, final optimality 9.5×10^{-7} (gto1). Revenue lower by 28 EUR, degradation cost lower by 115 EUR, objective change +87 EUR.

The representative day for Type B, a converged loss, is D179, a high-variability day. The solver reached a first-order optimum that is worse than the LP on both terms. The NLP trajectory remains above the LP trajectory for the whole day and does not return as far toward the lower operating limit. A higher mean state of charge raises the stress factor, so annual degradation increases rather than decreases, and the day also earns less revenue. The reported degradation difference is therefore negative and the objective change is -45 EUR. This day shows that a certified first-order point is not guaranteed to improve on the LP: the embedded gradient can lead the solver to a point that is worse in both revenue and degradation. The direction the solver took is consistent with the gradient it was given. Raising the trough of a cycle lowers its amplitude and raises its mean, and the supplied gradient carries only the first effect (Section 3.8.1). At the trough of this day's deep cycle it overstates the reduction in degradation by about a third, so a move the gradient reports as an improvement increases degradation instead. A second converged pattern, in which the two trajectories coincide and revenue falls for a near-zero reduction in degradation cost, is common among converged losses; day D217, for example, gives up 71 EUR of revenue for a 0.8 EUR reduction.




Figure E.2: Type B, converged loss: day D179 (High regime). Converged solve, final optimality 2.2×10^{-7} (gto1). Revenue lower by 22 EUR, degradation cost higher by 24 EUR, objective change -45 EUR.

The representative day for Type C, a stalled solve, is D084, a mid-variability day on which the solver stopped at a final optimality of 3.9×10^{-3} , three orders of magnitude above the convergence threshold, without certifying a local optimum. The LP and NLP coincide until hour 14. They then separate: the NLP holds a flat shelf near 0.46 through the afternoon while the LP drops to about 0.29, and the two rejoin near hour 22. Holding the higher afternoon level reduces the cycle depth and is reported as a large reduction in degradation cost, giving an objective change of +107 EUR. Because the solver did not converge, this reallocation is not certified, and the reported saving is attached to a trajectory the optimality conditions do not support. This type carries the greatest risk of misreading, because a stalled optimality can sit close to the threshold and be taken for a near-optimal gain until the residual is inspected.




Figure E.3: Type C, stalled: day D084 (Mid regime). Final optimality 3.9×10^{-3} (xtol, step collapse), above the convergence threshold. Revenue lower by 31 EUR, degradation cost lower by 138 EUR, objective change +107 EUR.

The representative day for Type D, a diverged solve, is D363, a low-variability day on which the solver stopped at its iteration limit with a final optimality of 8.6×10^9 . The NLP state of charge is a flat line near 0.52 for the whole day: the battery neither charges nor discharges. The LP, by contrast, runs a full deep cycle, charging to the upper limit and discharging into the evening. The reported 986 EUR reduction in degradation cost is an accounting result of a trajectory that does not change over the day, and only about 0.5 EUR of revenue separates the two dispatches on this low-price day. This day is the clearest example of the effect that produces the largest points in the positive tail of Figure 4.15: a saving computed on a solution the solver never made feasible.




Figure E.4: Type D, diverged: day D363 (Low regime). Final optimality 8.6×10^9 (maxiter). Reported degradation cost lower by 986 EUR on a trajectory that does not change over the day, with about 0.5 EUR of revenue separating the two dispatches; objective change +986 EUR.

E.2. Distribution of Outcomes

The complete per-day dataset is summarized here in two distribution tables rather than as a full 356-row listing, since the individual day values carry no information beyond the distributions they form. Table E.1 bins all active days by the change in objective value and splits each bin by solver convergence. Table E.2 gives the sign of that change and the positive mass for each convergence band, which is the tabular form of the scatter in Figure 4.15. Together they are the source for the aggregate statistics quoted in Section 4.5.

Table E.1: Distribution of the change in objective value across the 356 active days, split by solver convergence. Most days cluster just below zero, and the positive bins are dominated by non-converged solves.

Objective change (EUR)	Days	Converged	Non-converged
Below -50	8	7	1
-50 to -20	38	33	5
-20 to -5	114	74	40
-5 to 0	133	54	79
0 to 5	4	1	3
5 to 20	17	6	11
20 to 50	19	6	13
Above 50	23	4	19
Total	356	185	171

Table E.2: Sign of the objective change and positive mass by convergence band. The positive mass is the sum of all positive changes in the band. Converged days are the majority of the sample but carry 13% of the positive mass; the remainder comes from solves that did not reach an optimum.

Band	Days	Change ≤ 0	Change > 0	Positive mass (EUR)
Converged ($< 10^{-6}$)	185	168	17	666
Stalled (10^{-6} – 10^{-2})	95	71	24	1 088
Diverged ($> 10^{-2}$)	76	54	22	3 492
Total	356	293	63	5 246

For reference, the nine excluded zero-dispatch days are D024, D027, D045, D094, D098, D279, D315, D320, and D321. The within-year price-spread tertile cuts for 2022 are 127 and 210 EUR/MWh.

E.3. Representative Sample of Per-Day Results

Table E.3 lists individual per-day results for two groups of interest: the ten largest objective gains and a set of typical converged losses. Nine of the ten largest gains are non-converged, which is the per-day form of the result that the positive tail of Figure 4.15 is carried by non-convergence. The exception is D189, a converged gain of 136 EUR and the largest certified improvement over the LP in the year, which is the per-day form of the result that a small number of converged days do beat the LP. The typical converged losses show the other converged pattern, in which revenue falls for a near-zero reduction in degradation cost. A day is counted as converged when its final optimality is below 10^{-6} , the first-order optimality tolerance, so the convergence status of each day can be read directly from the optimality column. The four exemplar days of Section E.1 are shown there; two of them, D363 and D084, also fall among the ten largest gains and so appear in the upper block below.

Table E.3: Representative per-day results: the ten largest objective gains and six typical converged losses. Reg: regime (H/M/L). Spr: daily price spread (EUR/MWh). Δ Rev: revenue difference, positive when the reported dispatch earns less (EUR). Δ Deg: degradation cost difference, positive when it degrades less (EUR). Rat: ratio of the two. Δ Obj: change in objective value against the LP (EUR). It: iterations. Opt: final optimality; a day is converged when this is below 10^{-6} . The full dataset is summarized in Tables E.1 and E.2.

Day	Reg	Spr	Δ Rev	Δ Deg	Rat	Δ Obj	It	Opt
<i>Ten largest net gains (nine non-converged, one converged)</i>								
D363	L	54	0.5	986.2	2148.8	+985.8	500	8.6×10^9
D099	L	107	4.2	962.0	228.3	+957.8	500	4.4×10^{12}
D312	L	93	97.1	383.4	4.0	+286.4	394	9.8×10^{-4}
D147	M	142	16.0	289.4	18.1	+273.5	500	6.2×10^{-2}
D145	M	176	35.5	258.6	7.3	+223.1	421	3.1×10^{-2}
D073	H	262	92.0	289.9	3.1	+197.9	418	6.2×10^{-1}
D189	L	85	85.9	222.4	2.6	+136.5	223	9.5×10^{-7}
D097	M	136	131.2	243.4	1.9	+112.2	449	1.2×10^{-1}
D114	M	191	2.3	111.9	49.5	+109.6	457	4.0
D084	M	170	31.1	138.3	4.5	+107.2	363	3.9×10^{-3}
<i>Typical converged losses</i>								
D241	H	355	71.3	0.3	0.0	−71.0	118	1.8×10^{-7}
D175	M	194	34.5	0.1	0.0	−34.4	186	1.5×10^{-7}
D119	L	102	21.3	0.1	0.0	−21.2	164	9.7×10^{-7}
D172	H	247	16.0	0.2	0.0	−15.8	167	9.5×10^{-7}
D344	M	165	8.2	0.2	0.0	−7.9	72	3.3×10^{-7}
D101	L	120	4.8	0.1	0.0	−4.8	284	2.4×10^{-7}