Rethinking Flexible Connection Agreements

A contract to property-based multi value evaluation of flexible grid connection agreements for utility-scale battery storage in the Netherlands

M.C. Tulleken

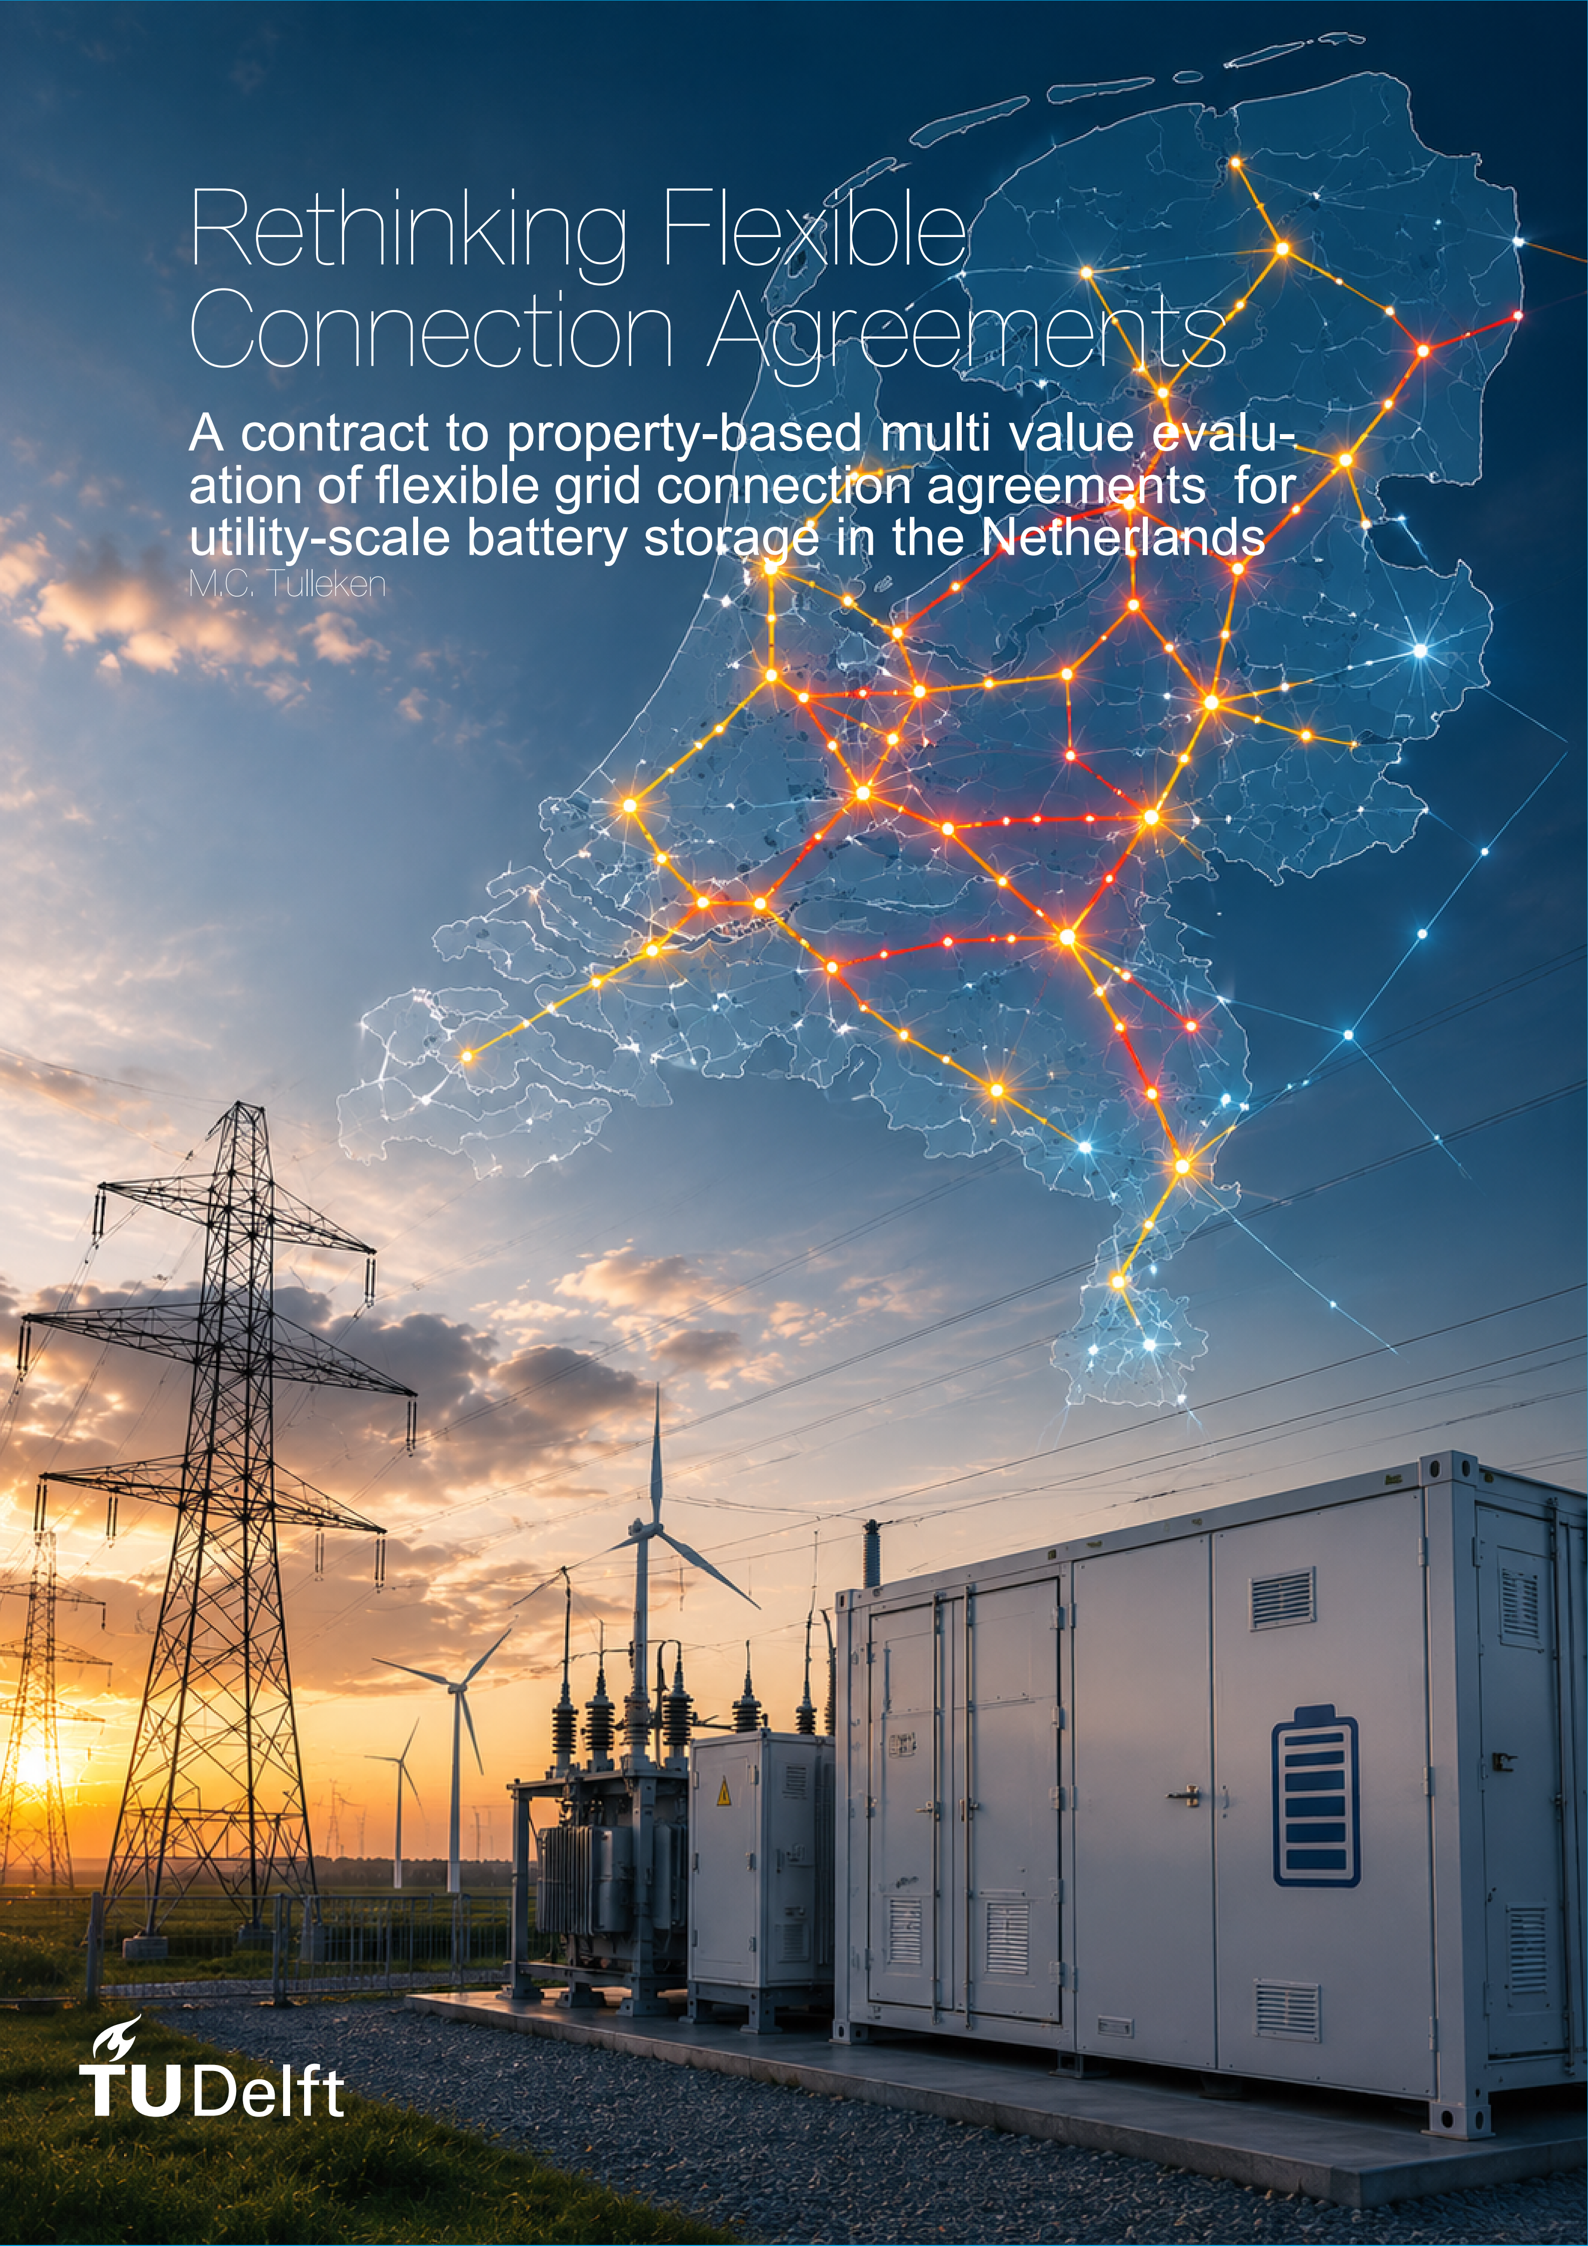

TUDelft

Rethinking Flexible Connection Agreements

A contract to property-based multi value evaluation of flexible grid connection agreements for utility-scale battery storage in the Netherlands

M.C. Tulleken

Master Thesis submitted in partial fulfillment of the requirements
for the degree of Master of Science

MSc Complex Systems Engineering & Management
Faculty of Technology, Policy and Management
Delft University of Technology

Student number:	5048648
Project duration:	February 2026 – July 2026
Chair & Second supervisor:	Dr. A. Correljé
First supervisor:	Dr.ir. E.J.L Chappin
Advisor:	Dr. A. Abdelshafy
External supervisor:	Dr.-Ing. D. Schmid (RWE Supply and Trading GmbH)

Delft, July 2026

TU Delft

Preface

This thesis marks the completion of my Master's programme in Complex Systems Engineering and Management (CoSEM) at Delft University of Technology. It explores the design and evaluation of flexible grid connection agreements for utility-scale battery energy storage in the Netherlands, a topic situated at the intersection of the energy transition, electricity markets, and institutional design.

What initially attracted me to this subject was its socio-technical complexity. Flexible connection agreements are often presented as technical instruments for managing grid congestion. Throughout this research, however, I became increasingly fascinated by the institutional questions underlying these arrangements: how contracts allocate risks, responsibilities, and value between market participants, and how those institutional choices ultimately shape technical and financial outcomes. Exploring these interactions, and the bankability challenges they create for battery storage, proved both intellectually stimulating and highly rewarding.

Conducting this research in collaboration with RWE provided a unique opportunity to work at the interface between academia and industry. Developing a research approach that continuously moved between institutional analysis, quantitative modelling, and policy evaluation confirmed my enthusiasm for addressing complex societal challenges through a combination of engineering, economics, and public policy.

I would like to express my sincere gratitude to everyone who supported me throughout this project. I am particularly grateful to Dr. Aad Correljé for his valuable guidance, critical questions, and continual encouragement to develop a clear and coherent storyline throughout the thesis. I would also like to thank Dr. Ali Abdelshafy for his constructive feedback, his confidence in my independence as a researcher, and his continuous encouragement to strengthen both the scientific rigour and broader perspective of this work. Special thanks also go to Dr.-Ing. Dennis Schmid for his enthusiasm from the very beginning, his continuous involvement throughout the project, and for constantly challenging me to connect the academic work with industrial practice. I would also like to thank Dr.ir. Emile Chappin for serving on my graduation committee.

I would especially like to thank the entire graduation committee for creating an open, constructive, and intellectually stimulating atmosphere during our meetings. The combination of discussion, diverse perspectives, and concentrated feedback provided clear direction while giving me the confidence and freedom to continue developing the research independently between meetings. That balance between guidance and independence has been invaluable throughout this project.

I am equally grateful to RWE Supply & Trading for providing the opportunity to conduct this research within such an inspiring professional environment. I would like to thank all colleagues at RWE for the many interesting discussions and for making the past months both professionally rewarding and highly enjoyable.

Finally, I would like to thank my family and friends for their support, encouragement, and patience throughout my studies. Knowing that they always believed in me has been an invaluable source of motivation, not only during this thesis but throughout my entire time at university.

Enjoy reading!

M.C. Tulleken
Delft, 2026

AI Statement

In accordance with the guidelines of Delft University of Technology and in the interest of transparency in academic research, I declare that generative artificial intelligence (AI) tools, including ChatGPT, Grammarly, and Claude, were used during the preparation of this thesis.

Writing support. AI tools were used to improve the clarity, readability, grammar, and overall writing style of the report. They also assisted in restructuring text, reducing unnecessary verbosity, and refining the presentation of complex institutional and technical concepts. All scientific content, interpretations, and final wording were done by the author.

Visualisation support. AI-assisted tools were used during the conceptual design of several figures and visual layouts. The final figures included in this thesis were manually recreated and adapted by the author to ensure consistency with the methodology, results, and overall presentation of the research.

Programming support. AI tools were used as programming assistants for debugging, code optimisation, discussing implementation choices, and validating parts of the Python code used for the MILP optimisation model, Monte Carlo simulations, and post-processing analyses. All code modifications, model implementation, validation, and verification were performed by the author.

Author responsibility. The research questions, methodological design, institutional analysis, synthesis of stakeholder frictions into structural tensions, development of the property-based evaluation framework, implementation of the simulation models, interpretation of the results, and all conclusions presented in this thesis are the original work of the author. The author accepts full responsibility for the accuracy, integrity, and academic quality of this research.

Executive Summary

The Dutch electricity grid is running out of capacity long before it runs out of electricity. Legally binding climate targets, large-scale electrification, and rapid renewable deployment have pushed demand for transport capacity past the pace of physical network expansion: by the end of 2025, more than 14,000 business connection requests were awaiting capacity, with access timelines extending to several years in parts of the country. Grid congestion has consequently shifted from a temporary engineering hurdle to a structural boundary condition on the Dutch energy transition. In response, Dutch regulators and network operators have introduced a portfolio of flexible connection agreements, the Alternative Transport Right (ATR85), the Time-Block Transport Right (TBTR), the Capacity Limiting Contract (CBC), and the Capacity Steering Contract (CSC), that make transport availability conditional on network conditions rather than guaranteed outright, exchanging transport certainty for earlier or continued grid access.

Battery Energy Storage Systems (BESS) sit at the centre of this shift. As bidirectional, highly responsive assets, they can absorb surplus renewable generation and release energy when the system needs flexibility. TenneT estimates that up to 9 GW of battery capacity may be required by 2030 to support the Dutch system. Yet this creates an institutional paradox: the technologies best positioned to relieve congestion are themselves increasingly exposed to congestion-related access restrictions. The existing connection agreements were designed around passive, predictable network users, not assets whose value depends on continuously adapting dispatch to changing system conditions. Whether these flexible agreements offer a sufficiently attractive and predictable investment environment for storage developers, while genuinely serving system-level congestion objectives, has remained poorly understood.

Existing research predominantly addresses this problem through three largely separate literature streams. Technical and economic studies of BESS operation evaluate revenue stacking and degradation but treat grid access rules as fixed external constraints. Congestion-management studies focus on system-level flexibility and redispatch but pay limited attention to the operational and financial consequences for individual assets. Institutional and regulatory studies examine non-firm access design and fairness of allocation but do not connect these regulatory choices back to concrete investment, degradation, and system-alignment outcomes. Few studies evaluate how contractual access agreements simultaneously shape battery operation, investment viability, asset degradation, and system-level outcomes within a single integrated framework. This thesis addresses that gap by asking to what extent existing and prospective flexible connection agreements influence the financial viability, asset degradation, and system alignment of utility-scale BESS, and which contractual design dimensions most effectively address the institutional frictions in Dutch non-firm transport arrangements.

The thesis answers this question through an exploratory sequential research design. A qualitative institutional analysis is conducted first, and its findings actively shape the structure of the quantitative model that follows, rather than the two being run in parallel, or the qualitative work being added afterward as interpretive colour. Concretely, 33 regulatory and stakeholder documents from the Authority for Consumers and Markets (ACM) were systematically coded, yielding 112 friction fragments. These were synthesised into three structural tensions: temporal misalignment between market commitments and real-time congestion needs, asymmetric allocation of risk and value between network operators and asset owners, and a disconnect between operational control and financial responsibility. These three tensions directly define what the subsequent simulation is built to test, rather than serving as background motivation.

The thesis's central contribution is the translation of these three tensions into four contractual design properties: compensation structure, firmness, volume orientation, and directive control. Contracts are then evaluated at the level of these underlying properties rather than as fixed products. This reframing shifts the unit of analysis from the contract to the property. The question is no longer which of these four contracts performs best, a question with only local Dutch relevance. Instead, the focus shifts to which underlying design properties drive which outcomes. The answer transfers naturally to congestion-management instruments not yet designed, in the Netherlands or elsewhere. Compensation structure and directive control are already represented by existing Dutch agreements; firmness and volume orientation were not represented by any existing agreement, which is why two prospective contract designs, the Firmness Gradient and the Volume Bound arrangement, were developed specifically to isolate them experimentally.

These four properties were tested in a Mixed-Integer Linear Programming Policy Testbed. A utility-scale 10 MW / 20 MWh lithium iron phosphate battery, co-optimising across the Day-Ahead, aFRR, and imbalance markets, was evaluated under six contractual arrangements, ATR85, TBTR, CBC, CSC, Volume Bound, and

Firmness Gradient, against an unconstrained baseline. Every arrangement was tested across 200 Monte Carlo market realisations spanning two stylised Dutch congestion archetypes: an urban, demand-driven regime and a rural, generation-driven regime. By evaluating every contractual arrangement under identical market realisations, the framework isolates the causal effect of contract design from stochastic market variation, ensuring that observed performance differences reflect contractual design rather than chance.

Two original metrics connect this quantitative evidence back to the institutional tensions above: Uncompensated Operational Risk (UOR), which isolates the market value displaced by a contract before compensation is considered, and the timing premium, which isolates the financial value of congestion-aligned versus randomly timed activation.

The unconstrained baseline reveals the scale of the underlying problem: even before any contractual restriction is applied, the Dutch high-voltage transport tariff alone absorbs about 70% of gross market revenues. As a result, the baseline is structurally non-bankable across all scenario families, reaching a best-case IRR of just 3.8% against the 8.5% hurdle rate. This establishes that the primary barrier to battery bankability in the Netherlands precedes contract design altogether.

The results indicate that, under the current Dutch tariff architecture, changes to operational contract design alone are insufficient to restore investment bankability. Properties such as firmness, volume orientation, and directive control primarily reduce operational risk or alter battery operation, but do not generate sufficient additional cash flow to overcome the structural tariff burden. Within the evaluated contract set, only the compensation-based design implemented through CBC restores bankability. CBC is the sole contract that achieves it, reaching IRRs of 11.4 to 20.6%. The underlying mechanism merits closer attention. The contract's financial transfer corresponds to a coverage ratio of up to 1,596% of demonstrated operational risk. CBC does not restore bankability because its compensation is closely calibrated to the operational harm it imposes. Rather, the financial transfer is sufficiently large to offset the underlying tariff shortfall, making the size of the transfer, rather than its proportionality to operational harm, the more consequential finding.

Directive control, tested through CSC, delivers by far the largest improvement in congestion alignment ($\Delta\Omega$ up to +0.192, more than six times the next-best contract), demonstrating that active dispatch instructions are uniquely effective at aligning battery behaviour with system needs. This benefit, however, is strongly conditional on accurate DSO activation timing: under random activation, the same contract adds financial cost without adding any system-alignment value, revealing that the case for directive control rests as much on DSO forecasting capability as on the contractual mechanism itself.

Firmness, tested through the prospective Firmness Gradient, reduces Uncompensated Operational Risk to near zero through a guaranteed 3 MW transport floor alone, with no financial compensation attached. This demonstrates that bounding downside exposure structurally can be more effective per euro than compensating for unbounded exposure after the fact, yet this risk reduction does not, by itself, translate into investment bankability.

Volume orientation, tested through the prospective Volume Bound design, extends implied battery lifetime by roughly two months but reverses sign on congestion alignment depending on network archetype. It improves alignment under demand-driven urban congestion but actively worsens it under generation-driven rural congestion. The reason is that the same daily energy budget that preserves flexibility reallocates dispatch toward high-value evening discharge, at the expense of the midday charging that rural congestion actually needs. This regime dependence shows that a contractual property's value is not portable across network topologies, and that congestion-management design may need to be locally, not nationally, calibrated.

Collectively, the results indicate that no single contract design property simultaneously delivers financial viability, congestion alignment, and technical performance. Instead, each property addresses a distinct aspect of the problem: compensation structure restores financial viability, directive control improves congestion alignment, firmness reduces operational risk, and volume orientation enhances technical performance through lower degradation, albeit with regime-dependent alignment effects. This differentiation of functional roles, rather than the performance of any individual contract, constitutes the central empirical finding of this thesis.

For policy and regulatory design, the results demonstrate that contractual design alone is insufficient. Compensation should be calibrated to demonstrated operational burden (as measured through UOR) rather than to historical redispatch prices or rated capacity, while directive-control instruments such as CSC should be complemented by investment in DSO forecasting and activation scheduling capabilities. Without accurate activation, the alignment benefits of directive control remain contingent on operational execution rather than inherent to the contract itself.

For the scientific literature, this thesis demonstrates that treating contract design as a set of underlying properties, rather than as fixed products, produces design insight that generalises beyond the specific Dutch instruments studied. It also reveals, empirically rather than only theoretically, that the institutional and technical dimensions of flexible grid access are not separable. Financial, technical, and system-alignment outcomes are structurally distinct performance axes that must be evaluated jointly. Optimising for one in isolation offers no guarantee of performance on the others and, in several cases, actively trades off against them.

Most importantly, the finding that even the best-performing contract achieves bankability only through compensation far exceeding demonstrated burden indicates that flexible connection agreements are currently being asked to perform two functions at once: managing congestion, and quietly compensating for a structural shortfall in the underlying transport tariff. This reframes the policy problem: the more consequential reform may lie upstream of contract design, in the tariff architecture itself, rather than in the contracts.

The principal contribution of this thesis is the property-based evaluation framework, which shifts the unit of analysis from contracts to contractual properties, compensation structure, firmness, volume orientation, and directive control, offering a transferable analytical lens for congestion-management instruments not yet designed. This is supported by three further contributions. First, an integrated exploratory sequential research design that systematically links qualitative institutional analysis to quantitative policy evaluation. Second, two original prospective contracts, Firmness Gradient and Volume Bound, each built to isolate a single design property from its usual remuneration; together, they provide direct evidence on how structural contract features, independent of how generously they are paid, shape investor risk and system outcomes. Third, two novel evaluation metrics, Uncompensated Operational Risk and the timing premium, alongside a structural Congestion Management Indicator, that give previously qualitative institutional concerns a quantitative, testable form.

Contents

1 Introduction

The availability of transport capacity has become one of the defining constraints on the Dutch electricity system. A combination of legally binding climate targets, large-scale electrification, and the rapid rollout of renewable generation has pushed demand for grid capacity well beyond the pace at which physical network infrastructure can expand [47, 54]. The consequences are already visible in the connection queue: by the end of 2025, more than 14,000 business connection requests remained unprocessed, and in several regions the wait for a high-voltage connection stretched to several years [49]. What was once treated as a temporary engineering bottleneck has hardened into a structural feature of the Dutch energy transition, one that now shapes which projects can proceed and on what timescale.

The resulting policy debate about this issue increasingly centres on two complementary pathways. The first focuses on accelerating physical network expansion through large-scale investments in transmission and distribution infrastructure. The second focuses on improving utilisation of the existing network through more flexible and efficient use of available capacity. Dutch network operators have increasingly emphasised that congestion cannot be resolved through network expansion alone and advocate a transition towards a more flexible electricity system in which network capacity is actively shared and managed [45]. While this transition is widely supported in principle, its implementation depends on the willingness of market participants to accept more flexible forms of grid access in exchange for earlier or continued network access. Whether they are willing to do so depends on the extent to which these agreements provide a predictable and economically attractive investment environment for flexibility providers.

To enable more efficient utilisation of existing network capacity, Dutch policymakers and regulators have introduced a series of flexible connection agreements intended to establish transparent rules for allocating transport rights, operational flexibility, and associated risks under congested network conditions. Examples include the Alternative Transport Right (ATR85), Time-Block-Based Transport Rights (TBTR), Capacity Limiting Contracts (CBC), and Capacity Steering Contracts (CSC). These agreements represent an institutional attempt to facilitate earlier grid access while enabling more efficient utilisation of existing network capacity by exchanging transport certainty for flexibility commitments [5, 19].

Although these agreements make the allocation of rights and responsibilities more explicit, they fundamentally alter how operational risks, financial value, and dispatch authority are distributed between network operators and connected assets. Across regulatory consultations and stakeholder discussions, concerns have repeatedly been raised regarding curtailment exposure, tariff structures, compensation mechanisms, and the allocation of operational control. These issues directly affect the predictability of future cash flows and, consequently, the ability of projects to attract long-term investment [24, 46, 7].

Among the technologies expected to operate under these new connection agreements, Battery Energy Storage Systems (BESS) occupy a particularly important position. As bidirectional and highly responsive assets, batteries can absorb surplus renewable generation during periods of oversupply and release energy when system conditions require additional flexibility, thereby supporting both system balancing and congestion management [16, 8]. TenneT estimates that up to 9 GW of battery capacity may be required by 2030 to support the Dutch energy system [59]. At the same time, storage assets derive much of their economic value from the operational flexibility that these agreements constrain, making BESS an ideal case through which the consequences of alternative contractual designs can be evaluated.

This creates an institutional paradox: the technologies that could contribute most to relieving congestion are themselves increasingly exposed to congestion-related access restrictions and contractual uncertainty. The underlying reason is that existing connection agreements were largely designed for passive and relatively predictable network users, whereas BESS derive value from continuously adapting their operation to changing system conditions [64]. As a result, the institutional assumptions embedded in traditional access agreements are poorly aligned with the operational logic of flexible storage assets.

The implications of these flexible arrangements for battery operation, degradation, investment viability, and congestion-management performance nonetheless remain poorly understood. At the same time, stakeholders continue to question whether the proposed allocation of risks, compensation mechanisms, and operational control provides a sufficiently bankable investment environment for utility-scale storage projects [25, 3]. It also remains unclear whether the currently proposed contractual agreements represent the most effective institutional designs. Alternative allocations of contractual properties might better reconcile investor incentives with public congestion-management objectives, but this possibility remains untested. This knowledge gap motivates the present research.

The academic literature offers only a partial perspective for evaluating the challenges associated with flexible connection agreements. A first stream of research explores the technical and economic operation of BESS. It mainly investigates optimisation strategies, bidding behaviour, revenue stacking, degradation modelling, and participation across electricity markets such as Day-Ahead, balancing, and ancillary-service markets [14, 43, 9, 64]. The studies provide a detailed understanding of battery operation and value creation. They typically evaluate batteries under predefined regulatory conditions, however, treating flexible connection agreements and congestion-management restrictions as exogenous constraints rather than objects of analysis [64, 39].

A second stream of research examines congestion management and flexibility procurement from the perspective of network operation. This literature primarily focuses on system-level objectives such as congestion relief, grid security, flexibility activation, queue management, redispatch, and market design [62, 19, 50]. While these studies increasingly acknowledge regulatory and governance considerations, their principal concern remains the effectiveness and efficiency of congestion-management mechanisms at the system level. Consequently, comparatively limited attention is paid to the operational and financial implications of these mechanisms for individual assets and investors.

A third stream of research addresses the institutional and regulatory dimensions of network access. Recent studies have begun to examine non-firm access agreements, flexible connection agreements, transport-right design, queue management, and questions of fairness in the allocation of scarce network capacity [19, 65, 64, 21]. These contributions provide valuable insights into regulatory design, access allocation, and governance structures, without evaluating how alternative contractual arrangements influence battery operation, degradation, system alignment, market participation, and investment viability across multiple electricity markets.

Taken together, these research streams provide important but fragmented insights. Technical studies explain how batteries operate, congestion-management studies explain how flexibility can support network operation, and institutional studies explain how transport rights are allocated. Limited attention has been paid to how these domains interact. Existing studies rarely evaluate how contractual access agreements simultaneously affect battery operation, investment viability, asset degradation, and system-level outcomes within a single integrated analytical framework. As a result, there remains limited understanding of how alternative connection agreements perform as socio-technical arrangements that jointly shape investor incentives, operational flexibility, and congestion-management objectives.

This fragmentation is particularly relevant because flexible connection agreements shape both technical system operation and the institutional conditions under which flexibility assets participate in electricity markets. As argued by Geels [30], socio-technical transitions are shaped not only by technological innovation but also by the institutional rules governing its deployment and operation. Flexible connection agreements can therefore be understood as institutional configurations that allocate transport rights, operational authority, compensation mechanisms, and risk among actors. Consequently, the same battery asset may experience substantially different operational and financial outcomes depending on how these contractual design choices are specified.

The challenge is therefore fundamentally institutional as well as technical. Emerging policy debates reveal recurring concerns regarding the allocation of curtailment risk and the predictability of future revenues [19, 62, 64]. Additional concerns relate to the interaction between market commitments and network constraints, and to the distribution of operational control between network operators and asset owners. Together, these concerns indicate that flexible connection agreements simultaneously function as market arrangements, operational restrictions, and governance instruments. Understanding their implications therefore requires an analytical perspective that integrates technical performance, economic viability, and institutional design.

Against this background, this thesis investigates how the institutional design of flexible connection agreements influences the financial viability, technical performance, and system alignment of utility-scale battery storage. Flexible connection agreements are treated not merely as operational instruments, but as institutional configurations characterised by different allocations of firmness, compensation, control authority, and operational risk. These contractual properties provide the analytical basis for the property-based evaluation framework developed in this thesis. Rather than evaluating individual contracts solely on technical or financial performance, the research investigates how underlying contractual design properties influence the interaction between investment viability, battery operation, asset degradation, and system alignment.

This leads to the following central research question:

> *To what extent do existing and prospective flexible grid connection agreements influence the financial viability, asset degradation, and system alignment of Battery Energy Storage Systems, and which contractual design dimensions most effectively address the institutional frictions identified in Dutch non-firm transport arrangements?*

To answer this research question, the thesis combines qualitative institutional discovery with a simulation-based Policy Testbed and a property-based evaluation framework. Together, these components enable existing and prospective connection agreements to be evaluated across financial, technical, and system-level dimensions. They also enable a systematic comparison of alternative contractual design properties. Crucially, the analysis does not end with these quantitative outcomes. The resulting evidence is mapped back onto the three structural tensions identified during the institutional discovery, thereby closing the analytical loop between stakeholder-identified governance frictions and simulation-based policy evaluation. In doing so, the thesis provides a structured framework for analysing flexible connection agreements as socio-technical design problems. It offers insight into how these designs may simultaneously support investment in storage assets and the broader objectives of congestion management within an increasingly constrained electricity system.

Research Objectives

This research pursues three interconnected objectives that together form a continuous analytical pipeline from institutional problem identification to quantitative policy evaluation.

First, it systematically identifies and structures the institutional frictions within the Dutch regulatory framework for flexible grid access. Second, it develops a reusable simulation-based Policy Testbed for the controlled comparison of alternative contractual arrangements. Third, it develops and applies a property-based evaluation framework that maps quantitative simulation outcomes back to the institutional frictions identified during the institutional discovery.

These objectives are formalised in the four sub-research questions below, each addressing a successive stage of the analytical pipeline.

Sub-Research Questions

- **SQ1: Institutional Discovery**
  How can institutional frictions in Dutch non-firm transport and flexible connection agreements be systematically identified from regulatory documents and stakeholder consultations, and translated into a structured policy design space?

- **SQ2: Baseline Model**
  What is the theoretical performance potential, in terms of financial value, physical degradation, and natural system alignment, of a BESS co-optimising across Day-Ahead, aFRR, and imbalance markets under unconstrained grid access conditions?

- **SQ3: Policy Testbed**
  How can a baseline optimisation model be transformed into a reusable policy testbed for evaluating non-firm access agreements for battery energy storage systems, and what operational, financial, and system-level outcomes emerge when these agreements are implemented as interchangeable constraint layers under identical market conditions?

- **SQ4: System Alignment Synthesis**
  To what extent do alternative connection agreement designs resolve the institutional frictions identified in the Dutch policy debate, and what trade-offs emerge when contracts are evaluated through a property-based framework decomposing contractual arrangements into compensation structure, firmness, volume orientation, and directive control?

Report Structure

The remainder of this thesis is organised into six chapters, following an analytical pipeline that links qualitative institutional discovery with quantitative policy evaluation.

Chapter 2 establishes the institutional and operational context of the study. It describes the transition from traditional firm grid access towards flexible connection agreements in the Netherlands and introduces the sequential electricity market environment in which utility-scale BESS operate. The chapter further develops the conceptual foundations of flexible connection agreement design by introducing the principal contractual dimensions of firmness, remuneration, volume orientation, and directive control.

Chapter 3 presents the research methodology. It describes the integrated methodological framework that combines qualitative institutional discovery with a quantitative Policy Testbed and property-based evaluation framework. The chapter is complemented by three methodological appendices. Appendix A documents the coding procedure and validation of the institutional discovery, Appendix B presents the complete mathematical formulation and technical implementation of the optimisation model, and Appendix C details the implementation of each contractual arrangement evaluated within the Policy Testbed.

Chapter 4 presents the research results. It identifies the principal institutional frictions within the Dutch policy debate, establishes an unconstrained operational baseline for utility-scale BESS, and evaluates the financial, technical, and system-alignment performance of six existing and prospective flexible connection agreement designs. The chapter concludes with an Integrated Evaluation that reinterprets these contract-level findings at the level of the four underlying design properties, compensation structure, firmness, volume orientation, and directive control, and maps the resulting evidence back onto the three structural tensions identified during the institutional discovery.

Chapter 5 discusses the findings in relation to the institutional tensions identified through the institutional discovery. It interprets the observed mechanisms underlying contract performance, evaluates the central methodological claim that property-level comparison yields more generalisable design insight than contract-level comparison, positions the findings within the existing literature, answers the research questions, and synthesises the principal scientific and practical contributions of the study.

Chapter 6 concludes the thesis with a broader reflection on the research process. It discusses the methodological choices made, the limitations of the study, and the researcher's personal development throughout the project.

Relevance to the Complex Systems Engineering and Management programme

This research closely aligns with the objectives of the MSc Complex Systems Engineering and Management (CoSEM) programme by addressing a complex socio-technical challenge at the intersection of engineering, economics, regulation, and public policy. Rather than treating BESS deployment as a purely technical optimisation problem, this thesis considers them as socio-technical assets whose successful deployment depends not only on engineering performance but also on the institutional arrangements governing their operation — reflecting the CoSEM philosophy that successful engineering solutions for complex societal challenges require the simultaneous consideration of technical, economic, organisational, and institutional dimensions.

The research combines qualitative institutional discovery of regulatory documents and stakeholder consultations with quantitative modelling through Mixed-Integer Linear Programming (MILP) optimisation and Monte Carlo simulation. By integrating these complementary perspectives into a property-based evaluation framework, the study demonstrates how contractual design, market incentives, regulatory arrangements, and physical system constraints jointly determine the feasibility, bankability, and system-level value of battery energy storage.

2
Institutional and Operational Context

2.1. The Transition from Firm to Flexible Grid Access

Historically, access to the Dutch electricity network was organised around the assumption of firm transport rights. Grid users could reasonably expect that a physical connection would be accompanied by continuous and unrestricted access to transport capacity. This arrangement reflected a centrally organised electricity system characterised by predictable demand patterns, dispatchable generation assets, and relatively limited electrification pressures. In this regime, the system operator assumed full responsibility for availability, effectively treating transport capacity as a guaranteed right that came automatically with the physical connection [30, 54].

The ongoing energy transition has increasingly challenged these assumptions. Rapid growth in renewable electricity generation, large-scale electrification of industry, transport, and heating, and the decentralisation of energy production have significantly increased pressure on the electricity network [32, 48]. As a result, transport capacity has become a scarce resource rather than an automatically available service.

The consequences of this scarcity have become increasingly visible. By the end of 2025, 14,044 unique requests for transport capacity were awaiting allocation, representing approximately 9.1 GW of requested capacity [49]. For many industrial projects, electrification initiatives, renewable generation assets, and storage developments, the availability of transport capacity has become a critical constraint on deployment and expansion. Network scarcity therefore increasingly influences which projects are able to participate in the energy transition and which remain on waiting lists [49].

In response, policymakers, regulators, and network operators have increasingly shifted attention towards the more efficient utilisation of existing infrastructure alongside continued network expansion [49, 58]. Within this debate, flexibility has emerged as a key system resource capable of creating additional transport capacity without immediate physical reinforcement of the network.

The growing emphasis on flexibility reflects a broader institutional transition. Rather than guaranteeing unrestricted transport rights under all conditions, the Dutch electricity system is gradually moving towards a regime in which access becomes conditional upon network conditions and system requirements. From this perspective, congestion is no longer solely an engineering problem concerning physical network capacity, but also an allocation problem concerning the distribution of rights, obligations, and risks among system actors.

Understanding this transition requires examining how the Dutch regulatory framework increasingly separates the physical connection from the right to transport electricity, thereby redefining how access, responsibility, and risk are allocated among system actors.

2.2. Grid Access as an Institutional Arrangement

The emergence of flexible transport rights is often presented as a technical response to physical congestion. However, from an institutional perspective, these developments represent a more fundamental transformation in the allocation of rights, responsibilities, and risks within the electricity system.

The legal foundation for this transformation can be found in the Dutch Electricity Act of 1998. A central feature of the legislation is the distinction between the obligation to provide a physical connection and the obligation to provide transport capacity. Article 23 establishes the connection obligation, requiring network operators to provide applicants with a physical connection to the electricity network. Article 24 establishes the transport obligation, but simultaneously permits network operators to refuse transport requests when sufficient capacity is not reasonably available [23, 7].

Under uncongested conditions, this distinction remains largely invisible because both rights are effectively delivered simultaneously. Under structural congestion, however, the distinction becomes highly consequential. A physical connection no longer guarantees unrestricted access to transport capacity. Instead, access to the network increasingly becomes contingent upon operational conditions, congestion forecasts, and contractual arrangements.

This development can be interpreted as a transformation in the nature of transport rights themselves. Historically, transport capacity functioned as a largely unconditional entitlement associated with a grid connection. Under emerging forms of non-firm access, transport rights become conditional rights whose availability varies across time and operating conditions [65]. In practical terms, this implies a redistribution of congestion-related risk from the network operator towards connected users.

The institutional significance of this transition extends beyond operational considerations. Flexible connection agreements effectively redefine the distribution of rights and obligations between network operators and market participants. Questions that were previously resolved through infrastructure expansion increasingly become questions of contractual design: who bears the risk of congestion, who retains operational control, and how should flexibility be compensated?

As the regulator responsible for maintaining the Dutch Electricity Code, the ACM plays a central role in operationalising the legal framework within which flexible connection agreements can be implemented [6, 44]. Through successive code amendments, the regulator has created the legal foundation for instruments such as Alternative Transport Rights (ATRs), Capacity Limiting Contracts (CBCs), and Capacity Steering Contracts (CSCs). In doing so, the ACM faces the challenge of balancing multiple objectives simultaneously: maintaining non-discriminatory access, safeguarding security of supply, facilitating the energy transition, and enabling more efficient utilisation of existing infrastructure.

The result is a regulatory environment in which transport capacity is no longer viewed solely as an engineering constraint, but increasingly as a scarce resource that must be actively allocated. Flexible connection agreements represent one response to this challenge. The following section introduces these arrangements in detail, before a later section examines why BESS occupy a particularly important position within this evolving framework.

2.3. Flexible Connection Agreements

To facilitate additional grid access under conditions of structural congestion, Dutch regulators and network operators have introduced a portfolio of flexible connection agreements. These arrangements differ from traditional firm access by making transport availability conditional upon network conditions, operational constraints, or dispatch instructions. Rather than representing a single policy instrument, they form a spectrum of contractual designs that differ in their allocation of transport certainty, operational control, and financial compensation [5].

ATR85 The Alternative Transport Right (ATR85) provides transport availability during at least 85% of annual hours; the contracted-capacity tariff component is waived in exchange, yielding tariff reductions of up to approximately 65% [64, 56, 5]. Flexibility is compensated ex ante through this standardised discount rather than through activation-specific payments.

TBTR The Time-Block Transport Right (TBTR) withdraws transport rights during predefined time windows aligned with periods of higher network utilisation, while remaining fully available at all other times [5]. Compensation is provided indirectly through a tariff reduction proportional to the share of time for which transport rights are unavailable, following the methodology specified in the code decisions [44]

CBC The Capacity Limiting Contract (CBC) is a congestion-management arrangement under which participants commit to a temporary operating limit below their contracted transport capacity when activated by the network operator, rather than losing access entirely. Participants receive direct financial compensation for this flexibility [16, 15].

CSC The Capacity Steering Contract (CSC) is the most interventionist arrangement currently available: the network operator actively directs asset behaviour through dispatch instructions during congestion events, which for BESS may include mandatory charging or discharging actions [31, 56, 53]. Rather than merely limiting network usage, the objective is to transform the battery from a congestion-neutral participant into an active provider of congestion relief.

Together, these agreements represent different responses to the challenge of allocating scarce transport capacity. While they pursue the common objective of improving network utilisation, they do so through different combinations of transport certainty, operational authority, and compensation mechanisms.

Although these agreements differ in their operational implementation, they can be understood as alternative combinations of a smaller set of underlying design characteristics. Differences between contracts primarily arise through variations in transport certainty, compensation structure, and the allocation of operational control. Their significance for Battery Energy Storage Systems is examined in the following section, before these dimensions are synthesised into the conceptual framework introduced in Section 2.5.

Table 2.1: Comparison of Dutch Flexible Grid Agreements for BESS

Dimension	ATR85 (TDTR)	TBTR	CBC	CSC
Institutional logic	Modified transport right	Modified transport right	Congestion-management service	Congestion-management service
Firmness	85% availability guarantee	Deterministic time windows	Full access except when activated	Activation dependent
Directive control	User self-dispatch	User self-dispatch	User self-dispatch within imposed limit	Network operator directed
Directionality	Bidirectional	Contract dependent	Typically one-sided	Congestion dependent
Activation logic	Curtailment by operator	Fixed schedule	Activated during congestion	Active dispatch instruction
Compensation	Tariff discount	Tariff discount	Direct remuneration	Direct remuneration

2.4. BESS in a Sequential Market Environment

BESS differ fundamentally from most other electricity system assets because their value is created through temporal flexibility. Rather than generating electricity from a primary energy source or consuming electricity for a fixed productive purpose, batteries create value by shifting electricity across time. Their commercial viability therefore depends on the ability to respond continuously to changing market conditions.

This value creation increasingly relies on participation across multiple electricity markets. Because no single market typically provides sufficient revenues to justify investment, battery operators combine income streams from day-ahead arbitrage, balancing services, and ancillary-service markets such as aFRR and FCR, a practice commonly referred to as revenue stacking [39]. The economic rationale for large-scale battery investments in the Netherlands, therefore, depends not only on access to individual markets, but also on the ability to move flexibly between multiple revenue opportunities over time.

The operation of a battery is further shaped by the sequential structure of electricity markets. Decisions made in the day-ahead market influence opportunities available in balancing markets, while real-time dispatch affects future market participation through changes in the battery's state of charge. Battery operation is therefore inherently intertemporal: a charging or discharging decision at time t affects the set of feasible decisions at time $t + 1$ and beyond [39, 16].

This intertemporal logic distinguishes batteries from many traditional grid users. For a conventional industrial consumer, a temporary restriction primarily affects consumption during the restricted period itself. For a battery, a restriction may have cascading effects across multiple subsequent periods because it alters the available energy inventory. A battery that is prevented from charging may be unable to discharge later; a battery that is prevented from discharging may lose access to future charging opportunities. Consequently, contractual restrictions influence not only immediate revenues but also future operational flexibility.

These characteristics make battery investments particularly sensitive to the design of transport rights. Uncertainty regarding the timing, frequency, and magnitude of access restrictions directly affects expected cash flows and project risk profiles. This issue is amplified by the relatively high network tariffs faced by Dutch battery projects. Industry stakeholders have repeatedly argued that the combination of substantial fixed grid charges and uncertain transport availability creates a bankability challenge for new storage investments [65, 25]. Flexible connection agreements such as ATR85 attempt to address this challenge by offering significant tariff reductions in exchange for reduced transport certainty [5].

At the same time, batteries occupy a unique position within congestion policy. Unlike many other grid users, they are not merely affected by congestion but can potentially contribute to its mitigation. Through strategic charging and discharging behaviour, batteries can absorb surplus generation, reduce peak loading, and provide local flexibility services. However, this creates a fundamental governance challenge. The market incentives that maximise private value creation do not necessarily coincide with the operational objectives of network operators. Under certain circumstances, economically optimal charging behaviour may reinforce rather than alleviate local congestion.

This tension lies at the heart of the debate on flexible connection agreements. Batteries are simultaneously commercial assets, providers of system flexibility, and potential congestion mitigators. Consequently, the design of transport contracts is not merely an operational issue but a question of how financial incentives, operational control, and system needs can be aligned within a congested electricity system.

2.5. Conceptual Framework: Dimensions of Contract Design

Although flexible connection agreements differ in their operational details, they can be understood as combinations of a smaller set of underlying design dimensions. From an institutional perspective, contract design is fundamentally concerned with how rights, risks, and responsibilities are allocated between network operators and grid users. Across the academic literature, regulatory documents, and industry discussions, three dimensions consistently emerge as the primary lenses through which flexible connection agreements are described: firmness, compensation structure, and directive control. Together, these dimensions capture three fundamental institutional questions underlying flexible grid access: how much transport access is guaranteed, how value and risk are distributed, and who retains authority over operational decisions.

The first dimension is *firmness*, referring to the degree of certainty associated with transport access. Traditional grid connections provide effectively unconditional access to transport capacity, whereas flexible agreements introduce varying degrees of conditionality. Some arrangements, such as TBTR, provide deterministic access within predefined periods, while others, such as ATR85, provide stochastic access in which users know the probability of curtailment but not its exact timing [64, 19]. Because transport availability determines when and to what extent market opportunities can be accessed, firmness is closely linked to revenue uncertainty, cash-flow predictability, and ultimately investment bankability.

The second dimension is the *compensation structure*. Flexible contracts alter the distribution of financial value between network operators and users by compensating participants for accepting operational restrictions or providing flexibility services. Compensation can take the form of indirect tariff reductions, as in ATR85 and TBTR, or direct payments linked to flexibility provision, as in CBC and CSC arrangements [13, 64]. The design of remuneration mechanisms determines how operational risks are valued and how incentives are distributed between public and private actors.

The third dimension concerns *directive control*, referring to the extent to which operational dispatch authority is transferred from the asset owner to the network operator. Under conventional self-dispatch arrangements, asset operators determine charging and discharging behaviour based on market signals. Under more interventionist arrangements, operational authority is partially transferred through restrictions, activation requests, or direct dispatch instructions [62]. The degree of directive control is particularly relevant for BESS because battery value depends strongly on optimising dispatch decisions across multiple markets and time horizons.

Together, these three dimensions provide an initial conceptual lens through which flexible connection agreements can be compared in the existing literature. Rather than treating this conceptualization as fixed, Chapter 3 revisits it through an institutional analysis of Dutch flexible connection agreements. That analysis extends the literature-based perspective into a four-property evaluation framework by identifying an additional design property, volume orientation, and by refining the literature concept of directive control to reflect the practical allocation of dispatch authority observed in stakeholder evidence. This empirically grounded framework forms the basis for the simulation-based evaluation developed in the remainder of this thesis.

2.6. Emerging Governance Challenges

The transition from firm to flexible access has generated an active debate among regulators, network operators, market participants, and industry organisations. Although stakeholders often differ in their preferred solutions, the literature consistently suggests that the introduction of flexible transport rights raises a range of governance challenges that extend beyond purely technical congestion management.

Several recurring concerns can be identified throughout policy documents, consultation responses, industry reports, and academic literature. These concerns relate to a wide range of issues, including the interaction between market processes and network operations, the allocation of risks and benefits, questions of fairness and compensation, the distribution of operational authority, and the practical implementation of flexible connection agreements [52, 64, 41]. While these issues are frequently discussed, they are typically addressed from the perspective of individual stakeholders, specific contractual arrangements, or isolated regulatory interventions.

As a result, the broader structure of these concerns remains unclear. It is not yet evident which issues are perceived as most important, how different stakeholder perspectives relate to one another, or whether recurring themes can be identified across the wider Dutch policy debate. Consequently, despite extensive discussion surrounding flexible transport rights, there remains limited understanding of the underlying institutional frictions that shape stakeholder behaviour and influence the perceived attractiveness of flexible connection agreements.

This institutional character of flexible connection agreements, established in Chapter 1, motivates the systematic Institutional Discovery presented next: understanding how these concerns are articulated, and how they relate to one another, requires a systematic examination of the institutional debate surrounding flexible transport rights.

The following chapter addresses this challenge through an Institutional Discovery. Using regulatory documents, consultation responses, industry publications, and stakeholder reports, the analysis systematically identifies and structures the institutional frictions associated with flexible connection agreements. The objective is not to evaluate contractual performance at this stage, but to identify, structure, and interpret the institutional frictions that characterise the Dutch debate on flexible transport rights. The resulting problem structure subsequently provides the foundation for the simulation-based policy evaluation developed in later chapters.

3 Methodology

3.1. Institutional Discovery and Optimisation Framework Overview

Evaluating whether a contractual arrangement resolves an institutional friction requires a methodology that can both identify the friction and measure its operational consequences. Neither approach is sufficient on its own. Qualitative analysis identifies stakeholder concerns but cannot verify if they have a significant impact on performance. Conversely, a quantitative model can rank contracts without explaining the underlying institutional problem it aims to solve. The framework developed here therefore combines both perspectives within a single analytical structure, reflecting the mixed institutional and technical nature of the research question.

This logic results in an integrated methodological framework, illustrated in Figure 3.1 and summarised in Table 3.1. The framework begins with an institutional discovery of 33 regulatory documents, identifying recurring frictions that are synthesised into three structural tensions concerning time, risk, and control. These findings subsequently inform the development of an optimisation framework consisting of an unconstrained Baseline model and a Policy Testbed in which alternative flexible connection agreements are evaluated under identical Monte Carlo market scenarios. Reusing the same market conditions, battery characteristics, and congestion profiles across all contract variants isolates the causal impact of contractual design from stochastic variation. Finally, the simulation outcomes are interpreted in light of the institutional tensions identified during the institutional discovery, allowing qualitative stakeholder concerns and quantitative performance to be evaluated within a single analytical framework.

The analytical bridge between the Institutional Discovery and the Policy Testbed is provided by four contract design property dimensions that emerge directly from the institutional findings: compensation structure, firmness, volume orientation, and directive control. Rather than evaluating ATR85, TBTR, CBC, and CSC solely as fixed contractual products, the framework decomposes them into these underlying design properties, which can be compared consistently across both existing and prospective arrangements. Note that the evidential basis for these properties differs: compensation structure and firmness are evaluated across multiple contractual arrangements, whereas volume orientation and directive control are each represented by a single contract design within the current Policy Testbed. This asymmetry is addressed explicitly in Section 3.4.2. By evaluating contract designs at the level of underlying design properties rather than as fixed products, the framework moves beyond the comparison of individual contracts and enables the identification of which contractual characteristics are associated with financial viability, technical performance, and congestion alignment.

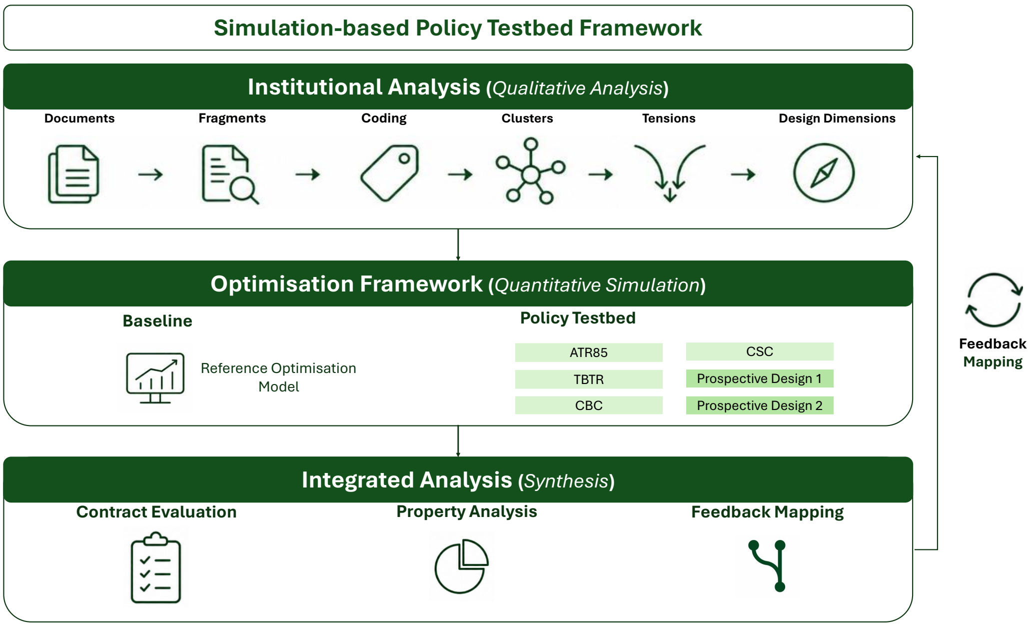

Figure 3.1: Overview of the integrated research framework. A qualitative institutional discovery of Dutch regulatory and stakeholder documents identifies recurring frictions, which are synthesised into structural tensions and translated into contract design properties. These properties inform the Policy Testbed, in which alternative flexible connection agreements are evaluated against an unconstrained Baseline under identical Monte Carlo market scenarios. Simulation outcomes are subsequently interpreted in light of the institutional tensions identified during the institutional discovery, integrating qualitative and quantitative findings within a single analytical structure.

3.2. Institutional Discovery

This component addresses the following sub-research question:

> *How can institutional frictions in Dutch non-firm transport and flexible connection agreements be systematically identified from regulatory documents and stakeholder consultations, and translated into a structured policy design space?*

The objective is not to evaluate performance, but to identify and structure the institutional frictions that emerge from existing agreements as reported in regulatory documents and stakeholder consultations. The outcome is a structured representation of recurring friction patterns and their underlying mechanisms, which are synthesised into a set of structural tensions and used to define a conceptual policy design space for subsequent modelling and evaluation.

The institutional discovery adopts a qualitative, document-based research design. The approach is exploratory, deriving analytical structure from stakeholder discourse, and is formalised through a structured coding and clustering procedure to ensure reproducibility.

The methodology follows a sequential process, illustrated in Figure 3.2. Starting from the construction of a stakeholder-informed document corpus, the analysis proceeds through successive stages of friction extraction, coding, clustering, synthesis, and translation into a policy design space. The resulting workflow provides the analytical bridge between the qualitative institutional discovery and the quantitative Policy Testbed developed in the subsequent research component.

As shown in Figure 3.2, each step builds upon the previous one, progressively reducing the complexity of the stakeholder corpus while preserving the underlying institutional mechanisms. This structure ensures that qualitative input is systematically organised into recurring friction patterns based on shared mechanisms and operational impacts, while preserving differences in stakeholder perspective [26].

Table 3.1: Overview of the integrated research framework

Methodological component	Role in research design
Institutional Discovery	Identification and structuring of institutional frictions from regulatory and stakeholder data through qualitative institutional discovery.
Optimisation Framework	Construction of an unconstrained Baseline and evaluation of alternative flexible connection agreements using a MILP-based Policy Testbed under identical Monte Carlo market scenarios.
Integrated Evaluation	Integration of qualitative and quantitative findings by evaluating financial viability, technical performance, and congestion alignment, and mapping these outcomes back to the institutional frictions identified during the institutional discovery.

3.2.1. Document Corpus Construction: ACM Regulatory Archive

The empirical basis for this institutional discovery consists of a systematically constructed corpus of 33 documents retrieved from the public archive of the Authority for Consumers and Markets (ACM). As the national regulator responsible for electricity market design and grid access, the ACM provides the authoritative source for both formal institutional rules and stakeholder consultation processes [2].

The dataset includes draft and final regulatory decisions, consultation documents related to congestion management and non-firm access, and stakeholder submissions.

Document selection followed three criteria. First, documents must explicitly address non-firm transport or congestion management mechanisms, including instruments such as ATR85, TBTR, CBC, and CSC. Second, documents must be part of formal consultation procedures, ensuring institutional relevance and comparability. Third, documents must contain substantive stakeholder input, allowing the inclusion of operational and market perspectives.

This selection ensures that the dataset captures both formal regulatory design and its interpretation in practice.

3.2.2. Stakeholder Classification

To facilitate a structured comparison of institutional perspectives, the analysed documents are classified according to the functional role of the author within the Dutch electricity system. This classification recognises that flexible connection agreements are evaluated from different institutional positions, each characterised by distinct responsibilities, incentives, and operational objectives. As a result, the same contractual arrangement may be perceived differently by regulators, system operators, market associations, and commercial actors. Table 3.2 summarises the stakeholder groups included in the institutional discovery together with the rationale for their classification.

Table 3.2: Stakeholder groups used in the institutional discovery.

Stakeholder group	Rationale for inclusion
Regulators (ACM)	Represent the regulatory perspective and are responsible for the design, approval, and oversight of flexible grid access arrangements.
Grid operators	Represent the system-operation perspective and are responsible for maintaining grid reliability, managing congestion, and implementing access arrangements.
Market associations	Represent collective industry interests and provide aggregated perspectives on market functioning, investment conditions, and practical implementation challenges.
Individual market actors	Represent commercial stakeholders directly affected by flexible connection agreements, including battery developers, renewable energy developers, and other market participants.

By distinguishing between these stakeholder groups, the analysis can identify whether reported frictions emerge consistently across the system or are concentrated within particular institutional perspectives. This dis-

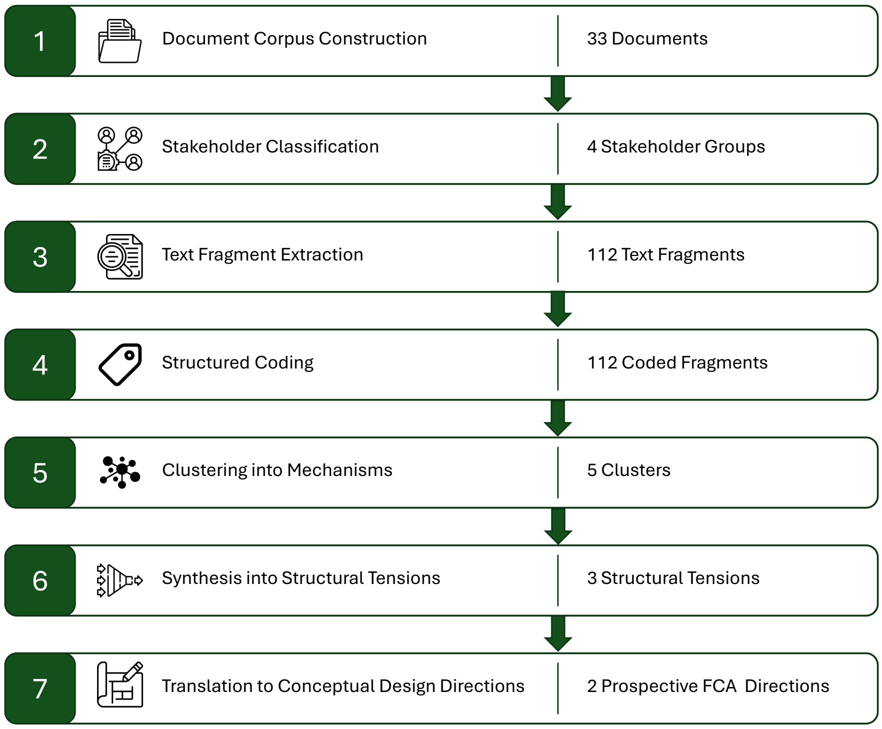

Figure 3.2: Institutional Discovery framework. Qualitative evidence from regulatory documents and stakeholder consultations is transformed through coding, clustering, and synthesis into an analytical design space that forms the conceptual foundation for the subsequent Policy Testbed.

tinction is important because broadly shared concerns may indicate systemic governance challenges, whereas concerns concentrated within a specific stakeholder group may reflect distributional conflicts related to the allocation of risks, responsibilities, or benefits.

3.2.3. Analytical Procedure: From Text Fragments to Codes

The unit of analysis is defined at the level of the text fragment rather than the document or stakeholder position for a specific reason: the same document frequently contains friction observations that point in different institutional directions. A single ACM consultation response may simultaneously identify tariff design problems, timing coordination failures, and control authority ambiguities. Aggregating to document level would obscure these distinctions; disaggregating to sentence level would strip the context necessary to interpret them. The fragment, a statement in which a stakeholder connects a concrete limitation to an identifiable institutional mechanism, represents the minimum unit that preserves both specificity and interpretive integrity for this research.

A fragment is included when it simultaneously satisfies two criteria. First, it must describe a concrete limitation affecting the operation, control, or economic viability of a BESS. Second, that limitation must be attributable to an identifiable institutional mechanism, a tariff design feature, dispatch rule, or contractual provision. Statements that are purely descriptive, procedural, or normative without reference to operational consequences are excluded.

This definition determines the extraction method. Institutional frictions in regulatory documents are rarely expressed as explicit complaints: they surface as qualifications, conditional statements, or technically-framed objections whose institutional significance depends on understanding the interaction between regulatory provisions and system operation. Keyword-based or pattern-based identification, the logic underlying tools such as Atlas.ti, would systematically misclassify fragments whose friction content is implicit rather than lexically marked [18, 10]. A comparative trial using Atlas.ti is reported in Appendix A and demonstrated that automated coding generated a large volume of procedural and context-free citations while failing to distinguish between operational symptoms and underlying institutional mechanisms. Fragment extraction is therefore performed through close reading of the full corpus, prioritising operational meaning over surface-level textual features. Given the moderate corpus size of 33 documents, this approach trades coding speed for the conceptual consistency that the subsequent clustering and synthesis steps require.

Structured Coding Procedure

Each identified fragment is transformed into a structured coding record. The objective is to ensure that observations remain traceable from the original stakeholder statement to the analytical interpretation used in the subsequent steps.

Table 3.3 presents an illustrative example of the coding structure applied throughout the analysis.

Table 3.3: Example of the coding procedure used to translate stakeholder statements into structured analytical observations.

Stakeholder	Agreement	Friction	Underlying mechanism	Operational impact	Original fragment	Source
Individual market actor	ATR85	Revenue uncertainty caused by unpredictable curtailment	Curtailment timing is uncertain and not known ex ante	Reduced revenue predictability and investment certainty	“The timing of curtailment remains uncertain, ... increasing investment risk.”	Consultation response

This coding structure ensures that each observation remains traceable from the original source statement to the corresponding analytical interpretation while maintaining a consistent level of abstraction across the dataset.

The focus of the coding procedure is on identifying recurring mechanisms rather than counting isolated statements. Similar fragments are therefore reformulated into consistent descriptions to enable comparison across documents and stakeholder groups.

Clustering of Frictions into Mechanisms

Following the coding step, the coded fragments, each comprising an identified friction together with its associated institutional mechanism and operational impact, are grouped into higher-level clusters. Cluster assignment is based on similarity in the underlying mechanism and operational impact rather than on similarity between the friction statements themselves.

Clustering is performed inductively, based on similarities in causal structure rather than document source or agreement type. A friction is assigned to a cluster if the underlying mechanism is consistent with the defining logic of the cluster and the operational impact affects the same dimension of system performance.

The objective of the clustering stage was not to identify a statistically unique classification, but to develop an analytically useful structure capable of capturing recurring friction patterns across the corpus. The resulting five-cluster solution was refined iteratively throughout the coding process and proved sufficiently distinct to support the subsequent synthesis into higher-level structural tensions. The clustering is therefore defended on substantive rather than statistical grounds: each cluster captures a distinct causal mechanism, and the later stages of coding primarily enriched and refined existing categories rather than necessitating the creation of additional ones. The resulting cluster structure is presented and discussed in Chapter 4.

3.2.4. Synthesis of Structural Tensions

The synthesis of friction clusters into structural tensions draws on two methodological foundations. Following Proka et al. [51], the identified frictions are interpreted not as isolated operational complaints but as symptoms of misalignment between an emerging technology regime, utility-scale BESS, and the established institutional regime governing congestion management and grid access. Following Enserink et al. [26], the objective is not to aggregate stakeholder positions or identify majority viewpoints, but to uncover the institutional mechanisms that repeatedly generate similar friction patterns across actors and contexts.

The synthesis proceeds through a second-order comparison of the causal mechanisms underlying each friction cluster. Rather than treating each cluster as an independent problem category, the analysis examines whether multiple clusters originate from the same underlying institutional misalignment. Clusters are synthesised when their recurring friction patterns consistently reinforce one another or when one category primarily amplifies the consequences of another. The objective is therefore explanatory rather than descriptive: to identify the smallest set of higher-order tensions capable of accounting for the observed friction patterns across the dataset.

This synthesis logic is illustrated by the relationship between information frictions and risk misalignment. Although initially identified as separate friction clusters, both ultimately concern the valuation of operational exposure under uncertain congestion conditions. Limited transparency regarding congestion frequency or activation behaviour does not merely create an informational burden in isolation; it increases financial uncertainty by preventing market participants from accurately assessing and pricing the risks associated with flexible grid access.

This shared causal mechanism justifies the synthesis of both clusters into the broader structural tension of risk and value allocation. The resulting structural tensions should therefore not be interpreted as additional coding categories, but as second-order analytical constructs that explain why multiple friction categories repeatedly co-occur throughout the dataset. In doing so, the synthesis reduces the complexity of the empirical observations while preserving their explanatory depth. The complete application of this synthesis process is presented in Chapter 4.

3.2.5. Translation to Contract Design Properties

The institutional discovery provides more than an explanation of existing contractual challenges. It also establishes the conceptual foundation for the quantitative evaluation developed in the remainder of this study. To enable simulation-based analysis, the identified structural tensions are translated into contract design properties that define the design space evaluated within the Policy Testbed.

While the structural tensions explain recurring institutional frictions, they cannot be represented directly within an optimisation model. A translation layer is therefore required that links higher-level institutional mechanisms to observable contractual characteristics. This translation proceeds from structural tensions, to contract design requirements, and ultimately to measurable contract design properties.

This translation is subject to an explicit scope assumption. Several institutional frictions identified during the analysis could, in principle, be mitigated through broader reforms of operational procedures, forecasting practices, information exchange, or market coordination. However, the objective of this research is specifically to evaluate the extent to which these frictions can be addressed through the design of flexible connection agreements themselves. Consequently, only design characteristics that can be embedded within contractual arrangements are considered.

Rather than treating existing connection agreements as fixed policy instruments, the analysis abstracts from individual contractual implementations and instead focuses on the underlying design dimensions that determine contractual behaviour. This abstraction enables both existing agreements and prospective contract concepts to be analysed within a common analytical framework.

The translation is guided by the premise that institutional tensions become operationally relevant through specific contractual design choices. Consequently, the identified structural tensions are reformulated as contract design requirements, which are subsequently represented through a limited number of observable property dimensions.

The first structural tension, **temporal misalignment**, concerns the interaction between contractual restrictions and the inherently intertemporal operating logic of utility-scale BESS. Across the analysed documents, stakeholders repeatedly expressed concerns regarding the timing, duration, and operational consequences of transport restrictions. These concerns point towards a fundamental design question that extends beyond the presence or absence of curtailment itself: whether contractual restrictions are imposed through instantaneous power limits or cumulative energy budgets. This distinction gives rise to the property dimension of volume orientation.

The second structural tension, **risk and value allocation**, emerged through recurring concerns regarding uncertain access rights, uncompensated operational losses, and the absence of mechanisms capable of limiting downside exposure. These concerns decompose into two institutionally distinct design questions. The first concerns proportionality: whether financial compensation reflects the operational harm imposed on the asset owner. The second concerns boundedness: whether contractual access rights limit worst-case operational exposure regardless of the compensation mechanism. These represent distinct institutional problems. A contract may compensate operational losses proportionally while still exposing the operator to unbounded downside risk if no minimum access level is guaranteed. Conversely, guaranteed access rights may bound downside exposure without necessarily providing adequate compensation. Accordingly, this structural tension gives rise to two separate property dimensions: compensation structure and firmness.

The third structural tension, **control versus responsibility**, concerns the allocation of operational authority and accountability. Stakeholders repeatedly highlighted situations in which network operators influenced dispatch decisions while the resulting financial consequences remained with the asset owner. This recurring concern is represented by the property dimension of directive control, describing the extent to which operational decision-making authority is transferred from the BESS operator to the DSO.

Together, these four property dimensions represent the smallest set of observable contractual characteristics through which the three structural tensions become operationally measurable within the Policy Testbed. The framework is therefore not imposed a priori, but emerges directly from the institutional discovery. It provides the conceptual bridge between stakeholder-reported institutional frictions and the subsequent quantitative evaluation of alternative flexible connection agreement designs.

Operationalisation into Contract Properties

The contract design properties identified above cannot be evaluated directly. To enable simulation-based analysis, they are operationalised as observable contractual features that can be represented within the optimisation model and evaluated using quantitative performance metrics.

The operationalisation follows three sequential steps. First, each structural tension is translated into a design requirement describing the institutional condition necessary to mitigate the underlying friction. Second, these design requirements are expressed as observable contract property dimensions. Finally, each property dimension is linked to quantitative performance metrics that provide evidence of its effectiveness under different operating conditions. The resulting operationalisation framework is summarised in Table 3.4.

Table 3.4: Operationalisation of structural tensions into contract design requirements, property dimensions, and evaluation metrics.

Structural tension	Design requirement	Property dimension	Evaluation metric
Temporal misalignment	Contractual restrictions should preserve operational flexibility across time while minimising interference with periods of highest operational value	Volume orientation	Congestion alignment reversal ($\Delta\Omega$) across urban/rural regimes; degradation; implied battery lifetime
Risk and value allocation	Financial compensation and access rights should be proportionate to the operational harm imposed and provide mechanisms to bound downside exposure	Compensation structure; Firmness	Compensation mismatch (UOR versus remuneration); IRR; bankability
Control rights versus responsibility	Operational control authority should be aligned with responsibility for the resulting operational and financial consequences	Directive control	$\Delta\Omega$; UOR under CSC; SoC displacement effects

The metrics named in this table are defined later in this chapter, not here: UOR in Section 3.3.8, the alignment score Ω and its Baseline-referenced version $\Delta\Omega$ in Section 3.3.7, and State-of-Charge (SoC) in Section 3.3.1. They are named here only to show which metric belongs to which property.

The timing premium (Section 4.2.2) is not listed as a row in this table, because it cuts across properties rather than belonging to one: it is computed for ATR85, CBC, CSC, and Firmness Gradient, but not for Volume Bound. For Volume Bound, the equivalent temporal effect shows up instead as the regime-dependent $\Delta\Omega$ reversal reported above. This is why the timing premium is discussed under Temporal Misalignment in Chapter 5, even though it has no row of its own here.

Two aspects of this operationalisation require clarification.

First, the identified property dimensions should not be interpreted as independent experimental factors. Individual contract designs typically combine multiple properties simultaneously. Consequently, the property-based analysis does not attempt to estimate isolated causal effects for individual properties. Instead, it identifies recurring performance patterns associated with different combinations of contract design properties across the evaluated agreements.

Second, the proposed property dimensions do not constitute an exhaustive taxonomy of flexible connection agreement design. They represent the dimensions that emerged most consistently from the institutional discovery and are therefore the most relevant to the research question. Other potentially relevant characteristics, such as contract duration, locational specificity, governance arrangements, or ownership structures, lie outside the scope of the present study.

This operationalisation establishes analytical continuity between the institutional discovery and the quantitative optimisation framework. The institutional discovery identifies the design challenges that flexible connection agreements must address; the property framework then translates these challenges into observable contractual characteristics that can be represented within the Policy Testbed. The optimisation framework subsequently evaluates how alternative manifestations of these properties influence financial performance, technical operation, and congestion alignment under identical market conditions.

3.3. Optimisation Framework

Having translated the institutional structural tensions into four observable contract design properties, the remainder of this chapter develops the quantitative Policy Testbed required to evaluate these properties under market and congestion uncertainty.

3.3.1. Design Philosophy

The central methodological challenge in evaluating flexible connection agreements is that their performance depends strongly on uncertain future market conditions. A contract that appears financially attractive under one price environment may become unattractive under another, while an agreement that performs well during infrequent congestion events may behave very differently when transport restrictions coincide with periods of high market value. Evaluating contractual designs under a single assumed future therefore provides only limited insight into their robustness or broader policy relevance.

The objective of the Policy Testbed is consequently not to forecast the revenues of a specific battery project, but to evaluate how alternative contract designs perform across a structured range of plausible operating conditions. Rather than predicting a single future outcome, the framework seeks to identify recurring behavioural patterns, financial outcomes, and system-level effects that remain consistent under varying market conditions.

This objective shapes the overall modelling philosophy. A Monte Carlo framework captures uncertainty by generating diverse market and congestion conditions, a rolling-horizon optimisation preserves the sequential and path-dependent nature of battery operation, and identical scenario realisations are reused across all contract variants to ensure that observed performance differences are attributable solely to contractual design.

Evaluating contractual performance additionally requires a counterfactual reference that is independent of any specific agreement. The optimisation framework therefore adopts an unconstrained Baseline representing a utility-scale BESS operating under identical market conditions, battery characteristics, and optimisation logic, but without DSO-imposed transport restrictions. Although such unconstrained operation does not represent a realistic institutional alternative under the Dutch regulatory framework, it provides a controlled experimental baseline against which the causal contribution of individual contract designs can be isolated.

3.3.2. Overall Architecture

Figure 3.3 summarises the architecture of the integrated optimisation Framework. To ensure methodological rigor, the framework is designed as a modular policy testbed where historical market data, technical battery specifications, and congestion archetypes form a single, consistent input base for Monte Carlo generation.

For every scenario, the model first executes an unconstrained Baseline to establish a counterfactual reference of maximum market value. Alternative flexible connection agreements are then implemented as interchangeable contractual constraint sets. Because every contract is evaluated under the identical scenario set described in Section 3.3.1, the framework isolates the causal impact of contractual design from market stochasticity.

A key design principle underlying this controlled experimental design is that the same congestion profile is used both to generate contractual restrictions and to evaluate ex-post system alignment.

This ensures methodological consistency across all variants: measured performance always reflects a response to the same underlying grid stress. As a result, differences in financial, technical, and alignment outcomes can be attributed to specific contract properties rather than to random variation in the market environment.

Asset and Market Representation

The framework represents a utility-scale lithium iron phosphate (LFP) battery system with a rated power of 10MW and an energy capacity of 20MWh. This two-hour duration is a deliberate design choice: it satisfies the minimum endurance requirements for participation in the Dutch automatic Frequency Restoration Reserve (aFRR) market [64, 55]. To preserve operational realism and avoid accelerated degradation, the battery operates with a one-way efficiency of 90% and a State-of-Charge (SoC) range between 5% and 95% [67, 42].

Multi-Market Participation To evaluate flexible connection agreements realistically, the optimisation framework must capture the revenue stacking logic of utility-scale battery operation. Following recent studies on Dutch battery valuation, the framework considers simultaneous participation in three principal value streams: Day-Ahead (DA) energy arbitrage, aFRR capacity reservation, and passive imbalance market participation [64, 20].

The economic value of battery flexibility depends on how capacity is allocated between these competing opportunities. For example, capacity reserved for balancing cannot be deployed for arbitrage, while real-time dispatch decisions directly affect future reserve availability through their impact on the State-of-Charge. This

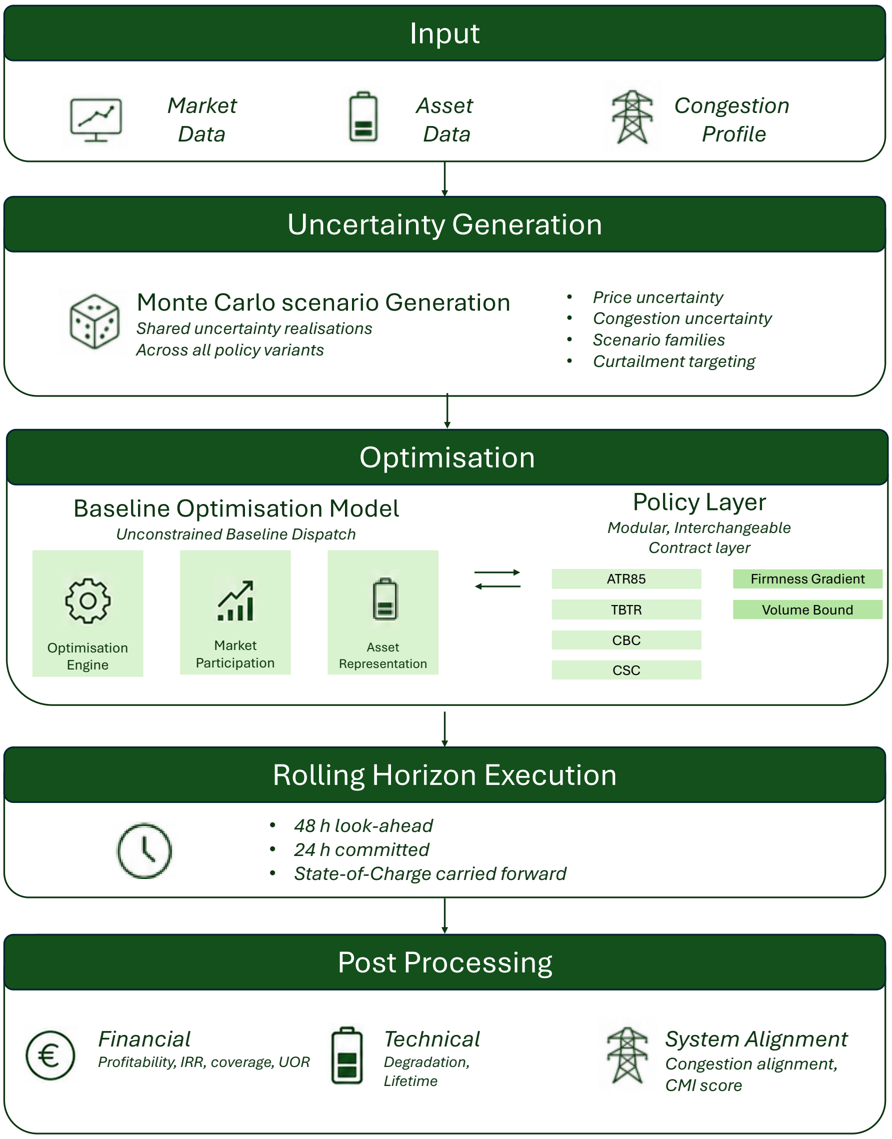

Figure 3.3: Overview of the policy testbed workflow. Historical market data, battery specifications, and congestion profiles are transformed into Monte Carlo scenario realisations, evaluated through a rolling-horizon optimisation model, and subsequently assessed using financial, technical, and system-alignment metrics.

study therefore adopts a co-optimised market representation, where all three revenue streams compete simultaneously for the battery's limited power and energy capacity [22, 63]. By internalising these market interactions, the framework ensures that any reduction in value observed under a policy variant represents a genuine operational burden rather than an artefact of inefficient market participation.

Market Participation Assumptions Several additional assumptions are introduced to ensure consistency with Dutch market design and realistic battery operation.

Reserve bids are constrained to remain constant within each four-hour aFRR procurement block, reflecting the institutional structure of TenneT's balancing market [57]. Furthermore, although the battery is physically capable of reserving up to 5 MW of balancing capacity, reserve participation is conservatively limited through an operational availability factor, ensuring sufficient flexibility remains available to deliver contracted balancing services under uncertain activation patterns and changing SoC conditions.

To preserve the price-taker nature of the asset, imbalance participation is limited to a maximum share of observed market volume within each Program Time Unit (PTU). This prevents the optimisation from implicitly assuming market power and ensures that simulated revenues remain representative of realistic Dutch market conditions [64]. In addition, imbalance participation is restricted to the system-supportive direction indicated by the national regulation-state signal, consistent with the underlying objective of maintaining system balance

[50].

MILP Formulation
Battery dispatch decisions are determined through a Mixed-Integer Linear Programming (MILP) formulation that maximises net operational value within each optimisation horizon.

The objective function is defined as:

$$\max\left(R^{\mathrm{DA}} + R^{\mathrm{aFRR}} + R^{\mathrm{imb}} + R^{\mathrm{contract}} - C^{DAFO} - \kappa^{\mathrm{win}} - \kappa^{\mathrm{cal}} - \kappa^{\mathrm{grid}}\right) \tag{3.1}$$

where R^{DA}, R^{aFRR}, and R^{imb} represent revenues from Day-Ahead trading, reserve capacity provision, and imbalance market participation respectively. The term R^{contract} captures remuneration associated with a specific grid access regime. Operational revenues are offset by Day-Ahead imbalance settlement costs (C^{DAFO}), cycling degradation costs (κ^{win}), calendar degradation costs (κ^{cal}), and annual grid access charges (κ^{grid}).

A key modelling choice is the explicit inclusion of degradation costs within the optimisation objective. Rather than treating degradation as an ex-post performance indicator, battery wear is monetised and internalised directly during dispatch optimisation. As a result, cycling decisions are only economically attractive when expected revenues exceed the associated degradation costs. This introduces an endogenous trade-off between short-term revenue generation and long-term asset preservation [12, 64].

This internalisation defines the behavioural logic of the framework: the battery operates as an economically rational agent that evaluates market opportunities against the opportunity cost of consuming battery lifetime. From an institutional perspective, this is critical because any reduction in value observed under a policy variant can be interpreted as a genuine operational burden rather than an artefact of inefficient dispatch. This logic provides the foundation for the UOR metric used to evaluate the proportionality of contract designs in later chapters.

The full mathematical formulation of the optimisation model, including physical operating constraints, market participation constraints, State-of-Charge dynamics, reserve deliverability constraints, variable definitions, solver settings and notation tables, is provided in Appendix B.

3.3.3. Battery Degradation Sub-Model

Building on the central behavioural assumption established in Section 3.3.1, this sub-model formalises the technical and economic representation of battery ageing. To implement the internalisation of degradation costs within the MILP framework, the model distinguishes between cycling-induced damage and time-dependent calendar ageing. These technical loss factors are subsequently monetised to determine the marginal opportunity cost of each dispatch decision [67, 40]. This internalisation is essential for the later evaluation of contractual impacts, as it ensures that any observed UOR reflects a genuine burden beyond the battery's own rational aging strategy

Cycling Degradation
Cycling degradation is represented using a convex power-law relationship between Depth of Discharge (DoD) and incremental battery damage. This formulation reflects the well-established observation that deeper discharge cycles impose disproportionately greater degradation than shallow cycles in lithium-ion battery systems [67].

The relationship is expressed as:

$$D(d) = D_{100} \cdot d^{\beta} \tag{3.2}$$

where d denotes the cycle Depth of Discharge, β is the degradation exponent, and D_{100} represents the damage associated with a full-cycle discharge. Following Xu et al. [67], a value of $\beta = 1.79$ is adopted for LFP battery systems. This introduces convexity into the degradation function, such that the marginal damage associated with cycling increases with discharge depth.

The degradation function is calibrated such that a full cycle at 100% DoD corresponds to approximately 8,000 equivalent full cycles before reaching End-of-Life (EoL), defined as a 20% reduction in usable capacity. This calibration ensures that the monetised degradation costs used during dispatch optimisation are representative of contemporary utility-scale LFP battery systems [52].

The adopted formulation deliberately prioritises computational tractability over electrochemical detail. The objective is not to predict cell-level ageing with high fidelity, but to capture the relative economic penalty associated with different cycling patterns within a large-scale policy evaluation framework. Because the resulting degradation function is non-linear, it cannot be incorporated directly into the MILP formulation. The function is therefore approximated using a piecewise-linear Special Ordered Set Type 2 (SOS2) representation.

Calendar Degradation
In addition to cycling-related degradation, the model accounts for calendar ageing. Calendar degradation represents the irreversible loss of battery capacity that occurs over time irrespective of operational intensity, primarily due to electrochemical ageing mechanisms such as Solid Electrolyte Interphase (SEI) growth [67].

Following the approach adopted by Xu et al. [67], a constant calendar degradation rate of 1% per year is applied throughout the simulation horizon.

Damage Monetisation
The final link between technical asset degradation and economic decision-making is established through damage monetisation. The framework adopts the damage-based levelised degradation methodology proposed by Massat et al. [40], which translates fractional capacity loss directly into an equivalent replacement cost per unit of State-of-Health consumed.

By incorporating degradation costs directly within the optimisation objective, the model evaluates operational decisions from an economic rather than purely technical perspective. This enables short-term market revenues and long-term asset preservation to be assessed within a unified decision framework [66, 40].
The detailed mathematical implementation of the degradation framework, including the SOS2 piecewise-linear approximation, breakpoint definitions, interpolation weights, and monetisation equations, is provided in Appendix B.7

3.3.4. Temporal Structure and Information Realism

In addition to physical and market constraints, real-world battery operators face significant informational limitations when making dispatch decisions. A common limitation in battery valuation studies is the assumption of perfect foresight, whereby the optimisation model has complete knowledge of future market conditions over the entire simulation horizon. While this assumption simplifies the optimisation problem, it systematically overestimates asset performance and does not reflect practical operational decision-making [32]. To improve operational realism, the optimisation framework combines a rolling-horizon architecture with a limited foresight framework that approximates the information available to market participants in practice.

Rather than optimising the entire year simultaneously, the annual simulation is decomposed into a sequence of overlapping optimisation windows. Each optimisation window considers a 48-hour look-ahead horizon, while only the first 24 hours of dispatch decisions are committed before the optimisation is repeated for the subsequent window [29, 32]. This architecture serves two complementary purposes. First, it reduces end-of-horizon effects, whereby an optimisation model may artificially deplete the battery because future operational opportunities are not represented. Second, it enables intertemporal decision-making beyond a single day, allowing the battery to balance immediate market opportunities against future operational flexibility.

To preserve continuity across the annual simulation, the terminal SoC of each committed optimisation window is propagated as the initial condition for the subsequent window. This carry-over mechanism prevents artificial resets of the battery state and ensures that both energy trajectories and degradation pathways remain physically consistent throughout the simulation period.

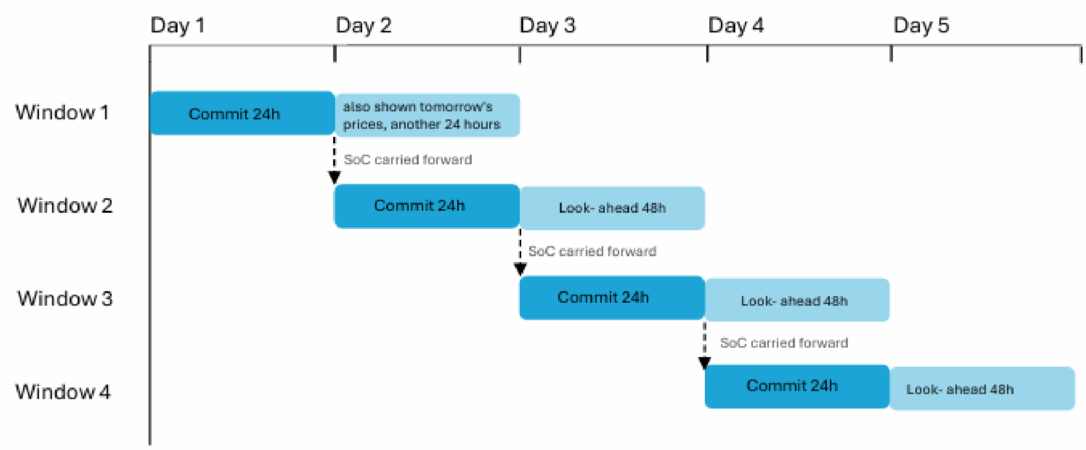

Figure 3.4: Illustration of the rolling-horizon architecture. The optimiser evaluates a 48-hour look-ahead horizon but commits only the first 24 hours of dispatch decisions, after which the optimisation is repeated using the realised terminal State-of-Charge as the initial condition for the subsequent window.

Operational realism is further improved by restricting the optimiser's access to future market information.

Rather than assuming perfect foresight, market-specific forecasting assumptions are introduced that reflect the institutional organisation of the Dutch electricity markets and prevent the optimisation from exploiting information that would not be available to market participants in practice.

For the imbalance market, realised prices are assumed to be known only over the immediate one-hour forecasting horizon. Beyond this horizon, future imbalance prices are replaced by rolling-average values, preventing the optimiser from exploiting short-lived price spikes that cannot be predicted reliably in practice. Similarly, aFRR reserve opportunities are evaluated at the level of the four-hour procurement block. Consistent with the institutional design of the Dutch balancing market, reserve bids remain constant throughout each procurement block and are therefore evaluated using block-level price expectations rather than perfect knowledge of future within-block price developments [55].

To account for forecast uncertainty and execution risk, conservative derating factors are applied to each revenue stream. Day-Ahead price spreads are compressed around the annual mean price by a factor of 0.90, following the DRF_{DAM} assumption proposed by Verhagen [63] (the full spread-compression formulation is given in Appendix B.4). A derating factor of 0.70 is applied multiplicatively to aFRR revenues to account for bid-stack competition and uncertainty regarding reserve-market capture rates. Finally, imbalance revenues are scaled by a factor of 0.50 to reflect ex-post settlement uncertainty and to avoid unrealistically optimistic valuation of imbalance market opportunities.

Together, the rolling-horizon architecture, limited foresight assumptions, and revenue derating factors provide a realistic representation of battery operation under uncertainty. Consequently, the baseline reflects both the physical constraints of utility-scale BESS operation and the informational limitations faced by real-world market participants.

The detailed implementation of the temporal execution and limited foresight assumptions is provided in Appendix B.3.

3.3.5. Uncertainty and Scenario Design

Evaluating flexible connection agreements under a single market trajectory provides only limited policy insight, as battery revenues are inherently sensitive to uncertain developments in electricity markets, renewable generation patterns, and the future saturation of flexibility assets. Rather than predicting a single future market trajectory, this study evaluates contract performance across a structured range of plausible operating conditions. To this end, the optimisation framework employs a Monte Carlo architecture that generates $N = 200$ independent market realisations.

The objective of the Monte Carlo framework is not to estimate future market probabilities with high statistical precision, but to expose each contractual arrangement to a broad, internally consistent uncertainty space. As noted by Koehler, Brown, and Haneuse [38], no universally applicable number of Monte Carlo replications exists; the required sample size depends on the variance of the model outputs and the desired estimation accuracy. The selected value of $N = 200$ therefore represents a pragmatic modelling choice that provides broad coverage of plausible operating conditions while remaining computationally tractable within the rolling-horizon optimisation framework.

Uncertainty is represented through two structural market drivers: volatility and market saturation. Volatility determines the magnitude of price spreads available for energy arbitrage and balancing services and is expected to increase with the growing penetration of weather-dependent renewable generation, which amplifies short-term supply fluctuations and contributes to phenomena such as the deepening of the duck curve. Conversely, increasing market saturation compresses these spreads as larger numbers of battery systems compete for the same flexibility revenues, reducing market value through price cannibalisation [39, 13, 20].

For each Monte Carlo realisation, both drivers are sampled independently from uniform distributions, $vol \sim \mathcal{U}(0.70, 1.30)$ and $sat \sim \mathcal{U}(0.70, 1.30)$, generating a broad uncertainty space that spans market environments ranging from favourable high-volatility conditions to highly saturated futures characterised by compressed revenue opportunities.

Collectively, these modelling choices transform the optimisation framework from a conventional revenue forecasting model into a policy testbed for evaluating institutional design choices under uncertainty. Rather than asking which future is most likely to occur, the framework evaluates which contractual arrangements remain robust across a broad range of plausible futures and how individual contract design choices influence financial performance, technical operation, and congestion alignment under uncertainty.

3.3.6. Congestion Environment and Timing

In addition to market uncertainty, the optimisation framework incorporates uncertainty regarding future congestion conditions. This is essential because the operational consequences of flexible connection agreements depend not only on electricity market dynamics, but also on the frequency, direction, and timing of transport restrictions imposed by the distribution system operator. To capture this uncertainty, the framework systematically varies both the congestion environment and the temporal distribution of transport restrictions [32, 52].

Congestion intensity is represented through six representative curtailment levels corresponding to restrictions during 5%, 10%, 15%, 20%, 25%, and 30% of all Program Time Units (PTUs). These levels span operating environments ranging from relatively unconstrained networks to heavily congested conditions, allowing contract performance to be evaluated across a broad spectrum of congestion severity.

To reflect the heterogeneous nature of congestion within the Dutch electricity system, the framework distinguishes between two stylised congestion archetypes: urban and rural. Urban congestion is primarily demand-driven and is associated with periods of high local electricity consumption, whereas rural congestion is predominantly generation-driven and associated with periods of high renewable electricity production, particularly solar generation [52]. These different congestion mechanisms are represented through asymmetric transport restrictions. In urban environments, 80% of congestion events restrict charging capacity and 20% restrict discharging capacity, reflecting the predominance of demand-driven congestion. In rural environments, this relationship is reversed, with 80% of events restricting discharging capacity and 20% restricting charging capacity, reflecting congestion associated with periods of high renewable generation. These proportions constitute modelling assumptions that capture the dominant directional characteristics of Dutch congestion while allowing for a limited proportion of atypical congestion events.

The economic impact of transport restrictions depends on their frequency, and on their timing. Battery revenues are concentrated within a relatively small number of high-value trading opportunities, meaning that restrictions imposed during these periods may substantially reduce project value, whereas restrictions of identical duration occurring during lower-value periods may have only limited financial consequences [20, 52, 64]. To isolate this effect, the framework varies the temporal allocation of congestion events independently of congestion intensity. In some scenarios, transport restrictions are distributed randomly throughout the year, whereas in others they are concentrated during periods of elevated network stress identified using the Congestion Management Indicator (CMI), introduced in Section 3.3.7. This experimental design isolates the economic consequences of congestion timing independently from congestion frequency.

The resulting difference in economic performance is referred to as the timing premium: the change in annual profit associated with fully congestion-aligned relative to randomly timed activation [64]. A positive timing premium indicates that congestion-aligned activation coincides with, and thereby reinforces, the battery's commercially preferred dispatch; a negative timing premium indicates that aligned activation displaces market value the battery would otherwise have captured under random timing. Together with the market uncertainty framework presented in the previous section, these congestion scenarios define the operating environments used throughout the Policy Testbed. By systematically varying market conditions, congestion intensity, and congestion timing while holding all remaining model components constant, the framework isolates the causal contribution of contractual design and identifies behavioural patterns and performance drivers that remain robust across a broad range of plausible futures. This provides the experimental foundation for the subsequent property-based evaluation of flexible connection agreement designs.

The mathematical implementation of volatility scaling and market saturation is provided in Appendix B.4; the probabilistic congestion-event generation procedure is provided in Appendix B.5.

3.3.7. Congestion Management Indicator Construction

Evaluating the financial performance of flexible connection agreements alone does not indicate whether battery operation supports the underlying needs of the electricity system. The optimisation framework therefore constructs a *Congestion Management Indicator* (CMI): a structured, time-varying signal that represents the relative desirability of charging or discharging from a congestion-management perspective, based on stylised congestion patterns observed within the Dutch electricity system [59].

A fundamental design principle is that the CMI remains entirely external to the optimisation problem. It is not available to the optimiser when dispatch decisions are determined, but is applied exclusively as an ex-post evaluation metric. Consequently, dispatch decisions remain driven solely by private market incentives, while the CMI provides an independent reference against which the resulting system-level alignment can be assessed. This distinction is central to the research objective. It allows the framework to quantify how much profit-maximising battery behaviour naturally coincides with congestion-supportive behaviour, and to assess how alternative flexible connection agreements change that relationship.

Rather than deriving congestion signals directly from operational redispatch data, the CMI is constructed using deterministic temporal weighting profiles. An initial approach considered using GOPACS redispatch transactions as the underlying congestion signal.

This approach was rejected on methodological grounds. GOPACS transactions lack the nodal specificity needed for asset-level evaluation, they provide only episodic rather than continuous congestion information, and their underlying spatial logic is inconsistent with the non-spatial MILP framework used in this study [37]. Instead, deterministic weighting profiles grounded in Dutch regulatory literature and stylised congestion behaviour provide a transparent, reproducible, and internally consistent representation of system stress across all evaluated scenarios.

The framework distinguishes between two stylised congestion archetypes representing the dominant forms of congestion currently observed within the Dutch electricity system. The first archetype represents urban, demand-driven congestion, where local electricity demand exceeds available transport capacity. Consistent with observed Dutch network loading patterns, this archetype is characterised by a morning demand ramp and a more pronounced evening demand peak, during which battery discharge is considered congestion-supportive because it reduces local network loading and mitigates thermal transport constraints [21, 32]. The second archetype represents rural, generation-driven congestion, where renewable electricity production, particularly solar photovoltaic generation, exceeds local transport capacity. Under these conditions, congestion primarily occurs during midday hours (10:00–16:00), during which battery charging is considered congestion-supportive because it absorbs excess local generation and reduces network export requirements [52].

For each congestion archetype, a deterministic weighting profile is assigned to every Program Time Unit (PTU) over the annual simulation horizon. Each interval receives a congestion weight $w_t \in [-1, 1]$, where the sign indicates the congestion-supportive dispatch direction and the magnitude represents the relative importance of that interval for congestion relief. Positive weights indicate periods during which charging is system-supportive, whereas negative weights indicate periods during which discharging is system-supportive. These weighting profiles serve two complementary purposes within the Policy Testbed. First, they guide the stochastic allocation of congestion events during scenario generation. Second, they provide the independent reference signal against which realised battery dispatch is evaluated.

The resulting congestion alignment score is defined as

$$\Omega = \frac{\sum_{t \in T_{\text{com}}} \left(p_{\text{dis},t} - p_{\text{ch},t}\right) w_t \Delta t}{\sum_{t \in T_{\text{com}}} |w_t|}, \tag{3.3}$$

where $p_{\text{dis},t}$ and $p_{\text{ch},t}$ denote the realised discharge and charge power (MW), respectively, w_t is the congestion weight assigned to interval t, $\Delta t = 0.25\,\text{h}$ denotes the Program Time Unit duration, and T_{com} represents the set of committed dispatch intervals. By construction, more negative values of Ω indicate stronger congestion-supportive alignment, whereas values close to zero indicate no systematic relationship between battery dispatch and the underlying congestion profile.

Because the CMI remains independent of the optimisation objective, the resulting alignment score provides an unbiased measure of system alignment. Throughout the remainder of this thesis, changes in congestion alignment are evaluated relative to the unconstrained baseline using the metric $\Delta\Omega$, allowing the system-level implications of alternative flexible connection agreements to be assessed alongside their financial and technical performance.

The complete seasonal weighting profiles are provided in Appendix B.6.5; the sign convention and mathematical implementation of the CMI alignment score are provided in Appendix B.6.2.

3.3.8. Financial Evaluation Framework

The Policy Testbed evaluates whether contractual arrangements remain financially viable under a broad range of future market and congestion conditions. While the preceding submodels focus on operational behaviour, this submodel translates these effects into investor-relevant performance metrics. The objective is not only to determine which contracts generate the highest financial returns, but also to assess whether those returns remain proportionate to the operational risks imposed by the agreement. Financial performance is assessed using Net Present Value and Internal Rate of Return over a 15-year investment horizon.

Scenario Affinity Weighting

Future battery revenues depend strongly on the long-term evolution of electricity markets. Because each Monte Carlo realisation represents only one plausible future, financial performance cannot be evaluated meaningfully at the level of individual runs. Instead, simulation outcomes are aggregated into four scenario families representing combinations of market volatility and market saturation.

Rather than assigning each Monte Carlo realisation to a single discrete scenario, the framework adopts a continuous affinity-weighting approach adapted from Kalluri and Gervás [36]. This avoids artificial scenario boundaries by allowing each realisation to contribute to multiple scenario families according to its position within the volatility-saturation space. Consequently, the scenario families should be interpreted as representative market regimes rather than mutually exclusive categories.

For each scenario family, a scenario-conditional annual operating profit is calculated as the affinity-weighted average across all $N = 200$ Monte Carlo realisations:

$$\bar{\pi}_m^{\mathrm{sc}} = \frac{\sum_{r=1}^{N} (A_r^{\mathrm{sc}})^{\alpha_m} \cdot \pi_r}{\sum_{r=1}^{N} (A_r^{\mathrm{sc}})^{\alpha_m}} \tag{3.4}$$

where A_r^{sc} denotes the affinity of Monte Carlo realisation r to scenario family sc, π_r denotes the annual operating profit of realisation r, and α_m is a sharpening exponent associated with milestone year m.

The purpose of the sharpening exponent is to represent the gradual reduction of uncertainty over the project lifetime. At project inception ($\alpha_m = 0$), all Monte Carlo realisations contribute equally, reflecting the absence of information regarding future market conditions. As the project progresses through the representative milestone years ($m \in \{1, 3, 6, 9\}$), the weighting gradually concentrates on those realisations that are increasingly representative of the evaluated scenario family. Early project years therefore reflect the full range of plausible futures, whereas later years increasingly capture the structural characteristics of the underlying market regime.

The same affinity-weighting procedure is applied to both annual operating profits and annual degradation rates. Together, these scenario-conditional milestone estimates form the input to the lifetime cashflow model described in the following section, allowing both financial performance and physical asset degradation to evolve consistently with the assumed market regime.

The complete mathematical formulation of the affinity scores and scenario-family construction is provided in Appendix B.8.

Cashflow Construction

A utility-scale Battery Energy Storage System (BESS) is not a static asset. As the battery ages, its usable capacity gradually declines, reducing its ability to participate in electricity markets and thereby limiting future revenue generation. A financial model that ignores this physical degradation would systematically overestimate long-term profitability and produce misleading assessments of project bankability [40]. The cashflow model therefore explicitly links the operational outputs of the optimisation framework to the projected physical condition of the battery throughout the 15-year investment horizon.

The scenario-weighted annual operating profit trajectory, $\bar{\pi}_y^{\mathrm{sc}}$, is adjusted for battery degradation through the projected State-of-Health (SoH) trajectory, which represents the remaining usable capacity of the asset. The projected State-of-Health in project year y is estimated as

$$\mathrm{SoH}_y = \max\left(0,\ 1 - \bar{D}_y^{\mathrm{sc}} \cdot y\right), \tag{3.5}$$

where $\bar{D}_y^{\mathrm{sc}}$ denotes the interpolated annual degradation rate in year y, obtained from the degradation model described in Section 3.3.3. A battery with $\mathrm{SoH}_y = 0.85$ is therefore assumed to retain 85% of its initial revenue-generating capability across all market activities, implying a proportional reduction in annual operating profits.

Annual project cashflows are subsequently constructed as

$$\mathrm{CF}_y = \bar{\pi}_y^{\mathrm{sc}} \cdot \mathrm{SoH}_y - \mathrm{OPEX}, \tag{3.6}$$

where $\mathrm{CF}_0 = -\mathrm{CAPEX} = -€6.0\,\mathrm{M}$ represents the initial investment and $\mathrm{OPEX} = 0.03 \cdot \mathrm{CAPEX} = €180\,\mathrm{k}$ denotes the annual fixed operating expenditure. Operating expenditures are assumed to remain constant throughout the project lifetime, reflecting the assumption that maintenance contracts, insurance, land lease, and operational overheads do not decline proportionally with battery degradation [20]. Consequently, degradation progressively erodes project profitability through the combination of declining revenues and largely fixed operating costs.

Investment Performance Metrics

The lifetime cashflow series is subsequently translated into investment performance metrics that quantify the financial attractiveness of each contractual arrangement from the perspective of a project developer. Two complementary metrics are employed throughout this study.

Net Present Value (NPV) measures the absolute value created by the investment at the required rate of return, which is assumed to be $r = 6\%$:

$$NPV = \sum_{y=0}^{15} \frac{\mathrm{CF}_y}{(1+r)^y}, \tag{3.7}$$

where a positive NPV indicates that the project generates value in excess of the required return.

The Internal Rate of Return (IRR) complements the NPV by identifying the project-specific discount rate r^* at which the investment exactly breaks even:

$$\sum_{y=0}^{15} \frac{\mathrm{CF}_y}{(1+r^*)^y} = 0, \tag{3.8}$$

To evaluate investment attractiveness against a market-consistent baseline, a hurdle rate of 8.5% is adopted. Within the context of this study, projects with an IRR below 8.5% are classified as non-bankable, as such returns are generally considered insufficient to attract utility-scale BESS project finance under current Dutch market conditions [39]. Both the discount rate and the hurdle rate were validated through expert consultation with a Dutch utility-scale BESS developer and are consistent with contemporary market practice [11].

Throughout the results chapter, the IRR serves as the primary indicator of project bankability because it facilitates direct comparison with investor return requirements. The NPV is reported alongside the IRR to provide additional insight into the absolute magnitude of value creation or destruction associated with each contractual design.

Uncompensated Operational Risk and Compensation Proportionality

Conventional financial performance indicators, such as Net Present Value (NPV) and the Internal Rate of Return (IRR), quantify the financial attractiveness of an investment but provide limited insight into whether contractual compensation remains proportionate to the operational burden imposed on a battery. A contractual arrangement may appear financially attractive because remuneration offsets the economic consequences of transport restrictions, while simultaneously transferring substantial operational risk to the asset owner. Financial viability alone therefore does not reveal whether a contract resolves the asymmetric allocation of risk and value identified during the institutional discovery.

Evaluating this institutional tension requires a metric that isolates the operational consequences of contractual restrictions independently from the compensation subsequently provided. The framework therefore introduces the concept of *Uncompensated Operational Risk* (UOR), which quantifies the gross market value displaced by a contractual arrangement relative to the unconstrained baseline before contractual remuneration is considered.

UOR is defined as

$$UOR = R^{\mathrm{gross}}_{\mathrm{baseline}} - R^{\mathrm{gross}}_{\mathrm{policy}} \tag{3.9}$$

where $R^{\mathrm{gross}}_{\mathrm{baseline}}$ denotes the gross market revenues obtained under unconstrained firm grid access and $R^{\mathrm{gross}}_{\mathrm{policy}}$ denotes the corresponding gross market revenues under the evaluated flexible connection agreement. Because both cases are evaluated under identical market conditions, battery characteristics, and congestion profiles, the resulting difference isolates the operational value displaced solely by the contractual design.

By excluding contractual remuneration from the calculation, UOR measures the intrinsic economic burden imposed by transport restrictions rather than the net financial outcome experienced by the asset owner. This distinction enables the framework to differentiate between contractual arrangements that reduce operational value displacement through their design and arrangements that primarily compensate for that displacement through financial transfers.

To evaluate whether contractual remuneration remains proportionate to the operational burden imposed, the framework introduces the *Compensation Coverage Ratio* (CR):

$$CR = \frac{R^{\mathrm{contract}}}{UOR}, \tag{3.10}$$

where R^{contract} denotes the contractual remuneration received under the evaluated agreement. A value of $CR = 1$ indicates that contractual compensation exactly offsets the displaced market value. Values below unity indicate under-compensation, whereas values above unity indicate compensation exceeding the operational burden imposed.

Together, UOR and the Compensation Coverage Ratio translate the qualitative concern regarding asymmetric risk allocation into measurable quantitative indicators. Rather than evaluating flexible connection agreements solely on financial profitability, the framework assesses whether contractual remuneration remains proportionate to the operational burden created by transport restrictions. These metrics therefore establish the analytical bridge between the institutional discovery and the quantitative Policy Testbed, enabling alternative contract designs to be evaluated in terms of both economic performance and the proportional allocation of risks and rewards.

3.3.9. Policy Testbed

The baseline Model establishes how a utility-scale Battery Energy Storage System (BESS) behaves under unconstrained grid access. Building upon this reference case, the Policy Testbed evaluates how alternative flexible connection agreements modify battery behaviour, financial performance, degradation, and congestion alignment under identical operating conditions. The objective is not to redesign the underlying optimisation framework for each contractual arrangement, but to isolate the causal effects of contractual design by introducing policy-specific restrictions and remuneration mechanisms while leaving all other model components unchanged.

To achieve this, each flexible connection agreement is implemented as a modular policy layer on top of the baseline Model. Operational provisions, such as transport restrictions, capacity limits, or dispatch instructions, modify the feasible operating region of the optimisation problem, whereas financial provisions enter the objective function through contract-specific remuneration and penalty components. The underlying optimisation structure, including the decision variables, battery dynamics, degradation modelling, market participation logic, rolling-horizon architecture, and Monte Carlo scenario generation, remains identical across all policy variants.

This modular architecture enables controlled comparison between contractual arrangements. As established in Section 3.3.1, every policy variant reuses the same Monte Carlo realisations, battery characteristics, and congestion profiles as the corresponding baseline case, so that observed differences in dispatch behaviour, financial performance, degradation, or congestion alignment can be attributed directly to contractual design.

Throughout both the optimisation and post-processing stages, contract-specific remuneration is tracked as a separate accounting stream. This enables financial performance to be decomposed into intrinsic operating value and contractual compensation, allowing investment performance, compensation proportionality, and UOR to be evaluated independently.

The Policy Testbed therefore provides a controlled experimental environment in which alternative flexible connection agreements can be evaluated in terms of their financial performance, and also in terms of their ability to mitigate the institutional frictions identified during the institutional discovery. In doing so, it provides the quantitative counterpart to the qualitative analysis presented in the first part of the methodology, enabling the institutional mechanisms identified earlier to be evaluated through observable operational, financial, and system-level outcomes.

Contract Implementations

Having established the Baseline Model and the common optimisation framework, the flexible connection agreements identified during the institutional discovery are implemented by translating their contractual provisions into mathematical constraints and remuneration mechanisms. Each implementation reflects the institutional logic of the respective agreement while preserving the common optimisation architecture. Consequently, all policy variants share identical battery dynamics, market participation logic, degradation modelling, and uncertainty assumptions, ensuring that observed differences in operational behaviour, financial performance, and congestion alignment can be attributed directly to contractual design rather than changes in the underlying model.

ATR85: Alternative Transport Right 85% The Alternative Transport Right 85% (ATR85) represents an availability-based congestion management arrangement that exchanges temporary transport interruptions for a fixed reduction in annual grid tariffs. During TSO-flagged congestion periods, three constraints simultaneously force charging power, discharging power, and Day-Ahead commitments to zero, resulting in a complete bidirectional operational shutdown irrespective of whether the triggering event originates from charging or discharging restrictions.

Although ATR85 formally permits curtailment during up to 15% of the year, this allowance is implemented as a stochastic per-window activation rate of 14 PTUs within each 96-PTU optimisation window rather than as a cumulative annual counter. This implementation remains consistent with the Monte Carlo uncertainty framework adopted throughout this study and reflects the uncertain occurrence of congestion events rather than deterministic quota exhaustion. The contractual tariff reduction of 65% is incorporated as a fixed annual revenue component that is independent of realised activation frequency, reflecting the institutional principle that remuneration is granted in exchange for transport availability rather than realised curtailment.

CBC: Capacity Limiting Contract The Capacity Limiting Contract (CBC) represents a passive congestion management instrument. Rather than interrupting transport access entirely, the agreement limits the maximum import and export capacity available to the battery during congestion periods while preserving operational flexibility within those limits. Consequently, the battery remains free to determine its economically optimal dispatch within the residual capacity envelope.

Accordingly, CBC imposes a symmetric power limit of 3 MW on both charging and discharging during congestion-flagged PTUs. This implementation reflects the institutional objective of reducing network loading without prescribing a specific operating strategy.

Contractual remuneration is linked directly to the curtailed transport capacity and is calculated as

$$R_{\mathrm{CBC}} = (P_{\mathrm{rated}} - P_{\mathrm{limit}}) \cdot r_{\mathrm{comp}} \cdot \Delta t \cdot n_{\mathrm{flagged}}, \tag{3.11}$$

where P_{rated} denotes the rated battery power, P_{limit} the contractual power limit, r_{comp} the compensation rate, and n_{flagged} the number of congestion-flagged PTUs. The compensation rate is set at €100/MWh, calibrated to the historical unweighted GOPACS average of €93.5/MWh [28]. As the ESNL working group proposes €120/MWh as a potential industry target, additional sensitivity analyses evaluate compensation rates between €40/MWh and €140/MWh. Unlike ATR85, CBC does not provide a fixed tariff reduction; contractual remuneration is therefore event-triggered and scales directly with the realised number of congestion activations.

CSC: Capacity Steering Contract The Capacity Steering Contract (CSC) represents the only evaluated arrangement in which the Distribution System Operator (DSO) actively influences battery operation during congestion events. Whereas ATR85 and CBC reduce network loading through predefined transport restrictions, CSC transfers part of the operational decision authority to the DSO through explicit dispatch instructions. The contract therefore provides a direct evaluation of the directive-control design dimension and enables the analysis to assess how increased operator control influences battery behaviour, financial performance, and congestion alignment.

During activation periods, the DSO requests a fixed dispatch setpoint of ± 5 MW, corresponding to 50% of the battery's rated power. The use of a fixed setpoint is a deliberate modelling choice intended to represent a substantial congestion-management intervention while maintaining experimental comparability across all Monte Carlo realisations. Although real-world dispatch requests would depend on the instantaneous magnitude and location of congestion, adopting a constant setpoint isolates the effect of directive control itself from the complexities of real-time congestion forecasting and system-specific operating conditions.

Compliance is modelled as a soft obligation rather than a physically enforced dispatch command. Instead of forcing the battery to follow the requested setpoint, the contract creates economic incentives through activation payments while imposing financial penalties for undelivered instructed power. This implementation reflects the institutional design of CSC, where compliance is achieved through contractual incentives rather than direct operational control. The complete mathematical formulation of the setpoint-tracking mechanism is provided in Appendix C.3.

An activation payment of €150/MWh is provided for both charging and discharging instructions, reflecting the greater operational burden associated with active steering compared with passive capacity limitation. This compensation level exceeds the CBC benchmark of €100/MWh, consistent with the recommendation of the ESNL working group that steering services should receive higher remuneration than limitation services because they impose greater operational constraints on flexibility providers [28]. Undelivered instructed energy is penalised at €400/MWh, corresponding to the reference level proposed during the ACM regulatory consultation. Additional sensitivity analyses evaluate lower penalty levels of €250/MWh and €150/MWh to assess the robustness of investment outcomes under alternative compliance regimes.

The activation structure reflects the dominant congestion characteristics of the two stylised network archetypes. Urban congestion is primarily demand-driven and therefore results predominantly in discharge instructions, whereas rural congestion is primarily generation-driven and therefore results predominantly in charging instructions, as summarised in Table 3.5.

Table 3.5: Location-specific CSC activation direction split reflecting dominant congestion mechanisms.

Location regime	Forced discharge	Forced charge
Urban (demand-driven congestion)	70%	30%
Rural (generation-driven congestion)	10%	90%

Unlike the congestion timing used elsewhere in the scenario framework, CSC activations are always scheduled according to the CMI-weighted timing profile. As CSC is explicitly designed as a congestion-response mechanism, dispatch instructions are assumed to occur during periods of elevated network stress rather than at random moments. This implementation ensures that the evaluation isolates the effectiveness of directive control under realistic congestion conditions rather than confounding it with random activation timing.

TBTR: Time-Block Transport Right The Time-Block Transport Right (TBTR) represents a deterministic congestion-management arrangement in which transport restrictions are imposed during predefined daily time windows. Unlike ATR85, where transport interruptions occur stochastically in response to congestion events, TBTR applies fixed calendar-based restrictions irrespective of prevailing network conditions. The agreement therefore provides a direct evaluation of how temporally rigid transport rights influence battery operation under dynamic market conditions.

The contract imposes a complete bidirectional shutdown during the fixed periods of 07:00–09:00 and 17:00–19:00, corresponding to 16 PTUs per day. Charging, discharging, and Day-Ahead commitments are simultaneously constrained to zero during these windows, independent of the congestion flags generated by the scenario framework.

Consistent with the institutional $t/24$ principle, the connection holder receives a tariff reduction proportional to the fraction of daily transport rights surrendered [5]. Consequently, giving up on 4 hours of transport availability results in a fixed tariff discount of approximately 16.7%. Unlike ATR85, remuneration is therefore entirely deterministic and independent of realised congestion conditions. This implementation allows the framework to isolate the financial and operational consequences of temporally rigid access rights and evaluate the extent to which fixed restriction windows coincide with high-value battery operating periods.

Prospective Contract Implementations

In addition to the existing contractual arrangements, the Policy Testbed evaluates prospective contract designs that emerged from the institutional discovery as potential solutions to the structural frictions identified in the current regulatory framework. Rather than replicating existing market arrangements, these contracts isolate individual design properties to evaluate their potential contribution to more effective and proportionate congestion management.

Energy-Volume-Bound Right The Energy-Volume-Bound Right replaces instantaneous transport restrictions with a cumulative daily energy budget. Whereas existing contracts regulate battery behaviour through power-based constraints applied at individual points in time, the proposed design limits the total energy exchanged over a predefined period while allowing the battery to determine the economically optimal timing of that exchange. This directly addresses the temporal misalignment identified during the institutional discovery, where stakeholders highlighted the conflict between instantaneous restrictions and the inherently intertemporal operating logic of battery storage.

The contract imposes a daily energy throughput limit of

$$E^{\mathrm{VB}} = 10\,\mathrm{MWh}$$

per direction over each 24-hour optimisation period. This corresponds to approximately 50% of the battery's rated energy capacity, allowing roughly half an equivalent cycle per day while preserving substantial intraday scheduling flexibility.

To prevent short-duration operation at full rated power, an additional instantaneous power envelope of

$$P_{\mathrm{env}}^{\mathrm{VB}} = 6\,\mathrm{MW}$$

is imposed alongside the cumulative energy budget. Together, these constraints provide both long-term energy control and realistic short-term operating limits.

No contractual remuneration or tariff discount is associated with this prospective design. This is a deliberate experimental choice that isolates the operational consequences of volume-oriented restrictions from financial incentives, allowing the contribution of the volume-orientation property to congestion alignment, battery degradation, and investment performance to be evaluated independently. The complete mathematical formulation is provided in Appendix C.5.

Firmness Gradient (Prospective) The Firmness Gradient contract was developed to evaluate whether guaranteed minimum transport rights can reduce investor risk more effectively than ex-post compensation mechanisms. Whereas existing arrangements primarily compensate battery operators after operational restrictions have occurred, the Firmness Gradient introduces a structural guarantee of network access by maintaining a permanent firm transport floor while exposing only the remaining connection capacity to congestion management. The contract therefore isolates the design property of *firmness* and evaluates whether guaranteed access alone improves both operational and financial performance.

A guaranteed transport floor of $P_{\mathrm{firm}} = 3\,\mathrm{MW}$ is maintained continuously in both charging and discharging directions, irrespective of congestion conditions. During congestion events, only the non-firm portion of the connection may be curtailed, while the guaranteed transport floor remains available at all times. The discharge-side implementation is represented by

$$p_{\mathrm{dis},t} \leq \begin{cases} P_{\mathrm{rated}}, & \text{if no discharge curtailment is active,} \\ P_{\mathrm{firm}}, & \text{if discharge curtailment is active,} \end{cases} \tag{3.12}$$

where P_{rated} denotes the rated battery power and $P_{\mathrm{firm}} = 3\,\mathrm{MW}$ represents the guaranteed firm transport capacity. The same implementation is applied symmetrically to the charging direction. Consequently, the battery retains access to at least 30% of its rated transport capacity under all operating conditions, while the remaining capacity remains subject to congestion management interventions.

The guaranteed transport floor is calibrated to preserve aFRR prequalification throughout the project lifetime. Maintaining uninterrupted access to balancing markets reflects the practical observation that investors seek to maximise expected revenues, and to preserve access to strategically important revenue streams under adverse operating conditions. The selected threshold therefore represents a stylised, yet institutionally grounded, minimum level of firm transport rights.

No contractual remuneration or tariff discount is associated with this arrangement. This is a deliberate experimental design choice intended to isolate the contribution of structural firmness from financial compensation. By excluding contractual transfers, any observed changes in profitability, degradation, congestion alignment, or UOR can be attributed directly to the guaranteed access right itself rather than to compensatory payments.

The Firmness Gradient therefore tests the hypothesis that structurally limiting downside exposure through guaranteed transport rights provides a more effective mechanism for improving project bankability and reducing investor risk than compensating unrestricted operational losses after they occur.

The complete mathematical implementation of all contract designs is provided in Appendix C

Compensation and Tariff Treatment

To ensure financial comparability across all evaluated arrangements, the full annual grid connection tariff of €143,570 per MW (2024 Dutch high-voltage transport tariff) is applied to every policy variant [63]. Consequently, all reported financial outcomes incorporate the complete cost of grid access, ensuring that differences in profitability reflect contractual design rather than differences in the underlying tariff structure.

Only the existing contractual arrangements, ATR85 and TBTR, receive an explicit tariff reduction. Consistent with their institutional design, these reductions are implemented as fixed positive revenue terms that are independent of realised battery dispatch. All remaining arrangements, including CBC, CSC, and the prospective contract designs, pay the full transport tariff. Where applicable, compensation is instead provided through contract-specific remuneration mechanisms linked to congestion events or dispatch instructions.

The prospective contracts, Volume Bound and Firmness Gradient, deliberately exclude both contractual remuneration and tariff reductions. This is an experimental design choice rather than a policy recommendation. By removing financial transfers from these prospective designs, the Policy Testbed isolates the operational consequences of individual contractual design properties, allowing the effects of firmness and volume orientation to be evaluated independently of compensation incentives.

This treatment establishes the counterfactual required for the subsequent property-based analysis. Any differences in financial performance, degradation, congestion alignment, or UOR can therefore be attributed either to the operational design of the contract itself or to the presence of financial compensation. Consequently, the Policy Testbed distinguishes between contractual arrangements that create value through more efficient operational design and those that achieve comparable outcomes primarily through financial transfers.

Sensitivity Analysis

The objective of the sensitivity analysis is to evaluate whether the principal contract-level and property-level findings remain robust under plausible alternative modelling assumptions and contractual parameterisations. Several policy designs considered within the Policy Testbed contain implementation parameters that are not yet standardised in regulatory practice, while the financial attractiveness of flexible connection agreements depends on compensation levels, tariff structures, penalty mechanisms, and investment assumptions that remain subject to market and regulatory uncertainty [6, 28, 25]. Consequently, the sensitivity analysis is not intended to optimise individual contract designs, but to assess whether the conclusions of the study remain stable across a realistic range of alternative assumptions.

A distinction is made between sensitivities that affect only the financial evaluation and those that alter the optimisation problem itself. Parameters such as tariff discounts, compensation levels, capital expenditure (CAPEX), discount rates, or hurdle rates modify only the post-processing stage. These analytical sweeps leave the battery dispatch trajectories unchanged, allowing the financial consequences of alternative assumptions to be evaluated directly from the existing Monte Carlo simulations without additional optimisation.

By contrast, parameters that modify either the objective function or the feasible operating region alter the battery's optimal operating strategy. These re-optimisation sweeps include, for example, the Capacity Steering Contract (CSC) penalty rate, activation timing quality (b_{dis}), or other parameters that influence the operational incentives faced by the battery. Because economically rational dispatch responds endogenously to these changes, the optimisation model must be solved again under the modified contractual conditions to capture the resulting behavioural adaptation.

This distinction reflects the underlying philosophy of the Policy Testbed. Financial assumptions influence the valuation of a dispatch trajectory, whereas operational assumptions influence the dispatch trajectory itself. Separating these two forms of uncertainty ensures that behavioural effects are not conflated with purely financial effects and enables a transparent interpretation of the resulting sensitivity analyses. The evaluated parameter ranges, default values, and methodological motivations are summarised in Table 3.6.

Table 3.6: Overview of the sensitivity analyses conducted. Analytical sweeps modify post-processing assumptions without altering battery dispatch and therefore produce exact financial metrics for previously simulated trajectories. Re-optimisation sweeps modify operational incentives or constraints and consequently require the MILP to be solved again. The default value denotes the parameterisation adopted throughout the primary Policy Testbed evaluation.

Parameter	**Contract**	**Parameter values evaluated**	**Model default**	**Purpose**
Analytical sweeps (post-processing only)				
Tariff discount	ATR85	40%–75%	65%	Assess the sensitivity of investment performance to tariff-based remuneration.
Tariff discount	TBTR	10%–30%	16.7% ($t/24$)	Assess the sensitivity of investment performance to tariff-based remuneration under deterministic transport restrictions.
Capacity payment rate	CBC	€40–140/MWh	€100/MWh	Assess the sensitivity of financial performance to alternative capacity-based remuneration levels.
Project CAPEX	CBC	€4M–€9M	€6.0M	Assess the sensitivity of investment performance to alternative capital cost assumptions.
Discount rate	CBC	6%–12%	6%	Assess the sensitivity of project valuation to alternative investor return requirements.
Re-optimisation sweeps (MILP re-solved)				
Deviation penalty rate	CSC	€150, €250, €400/MWh	€400/MWh	Assess the behavioural and financial effects of alternative penalty levels.
Activation payment rate	CSC	€100–180/MWh	€150/MWh	Assess the sensitivity of CSC performance to alternative activation remuneration levels.
Timing quality (b_{dis})	CSC	0.00, 0.25, 0.50, 0.75, 1.00	0.00 and 1.00	Isolate the contribution of activation timing quality to operational and financial performance.

Analytical Parameter Sweeps Several contractual remuneration parameters and financial assumptions enter the cashflow model as fixed scalars and therefore do not influence battery dispatch. These parameters can consequently be varied analytically, producing exact financial metrics for previously simulated trajectories without requiring additional optimisation runs. Table 3.6 summarises the evaluated parameters, parameter values, and default settings for the ATR85, TBTR, and CBC analytical sweeps.

For each remuneration-related sensitivity, the analysis additionally determines the *bankability compensation rate*, defined as the minimum remuneration level required for the project Internal Rate of Return (IRR) to reach the adopted 8.5% hurdle rate. This provides a direct measure of the remuneration required for each contractual design to satisfy the bankability criterion under different market regimes and thereby quantifies the financial implications of alternative contractual remuneration structures.

CSC Penalty Rate Sensitivity The deviation penalty of the Capacity Steering Contract (CSC) directly influences the optimisation objective and therefore affects the battery's operational decision-making during activation periods. Because the battery responds endogenously to economic incentives, varying the penalty requires the optimisation model to be solved again under the modified contractual conditions. Three penalty levels are evaluated: €400/MWh (the default value, corresponding to the level proposed in the ACM regulatory consultation), €250/MWh, and €150/MWh [4]. This sensitivity analysis assesses the effect of alternative compliance incentives on operational behaviour, congestion alignment, and financial performance.

CSC Timing Quality Sensitivity A dedicated sensitivity analysis evaluates the influence of DSO activation timing on the performance of the Capacity Steering Contract. The timing quality parameter b_{dis} is varied across five levels ranging from 0.00 (uniformly random activation timing) to 1.00 (fully congestion-aligned activation timing). Because all other contractual characteristics remain unchanged, this within-contract comparison isolates the contribution of activation timing quality to operational behaviour, congestion alignment, and financial performance. The analysis therefore distinguishes the contribution of the directive control mechanism from that of DSO targeting quality.

3.4. Integrated Evaluation

Having established the institutional discovery, the optimisation framework, and the Policy Testbed, the final methodological component defines how the resulting evidence is synthesised into conclusions regarding flexible connection agreement design. Rather than evaluating individual simulation outputs in isolation, this framework translates the operational and financial outcomes generated by the Policy Testbed into structured evidence concerning contractual performance, underlying design properties, and their institutional implications.

The objective is not simply to identify which contractual arrangement performs best, but to determine whether recurring performance patterns can be associated with identifiable design properties and whether these properties address the structural tensions identified during the institutional discovery. To preserve methodological transparency and rigour, the evaluation is organised into three sequential stages, each building explicitly upon the evidence established in the preceding stage:

1. **Contractual Assessment** The first stage establishes what the Policy Testbed directly demonstrates. Each contractual arrangement is evaluated relative to the unconstrained Baseline Model across three performance dimensions: financial viability, technical performance, and system alignment. The resulting performance indicators constitute direct model outputs and therefore require no interpretation beyond the simulated evidence itself.

2. **Property-Based Interpretation** The second stage examines whether recurring performance patterns across different contractual arrangements can be associated with the underlying design properties identified in the institutional discovery. These properties comprise compensation structure, firmness, volume orientation, and directive control. Because no two evaluated contracts differ exclusively in a single design dimension, this stage is explicitly observational. The analysis therefore identifies recurring associations and trade-offs rather than isolated causal effects.

3. **Institutional Feedback Mapping** The final stage interprets the property-level findings in relation to the three structural tensions identified in the institutional discovery: temporal misalignment, asymmetric risk allocation, and the disconnect between operational control and financial responsibility. This stage is explicitly interpretive, translating quantitative evidence into institutional insights regarding the extent to which alternative contractual designs address the governance challenges underlying flexible grid access.

This sequential framework ensures that every conclusion remains traceable to the evidence on which it is based. Direct simulation outputs are first established, recurring patterns are subsequently interpreted at the design-property level, and only then are institutional implications derived. By separating empirical observation from interpretation, the framework ensures that all conclusions remain within the evidential scope of the simulation-based methodology.

3.4.1. Contractual Assessment

The three performance dimensions, financial viability, technical performance, and system alignment, are not treated as independent analytical tracks but as complementary perspectives applied to the same simulation outputs. Collectively, they characterise the trade-offs associated with each flexible connection agreement. All contract-level findings presented in Chapter 4 are derived within this evaluation framework, ensuring that every conclusion remains directly traceable to the underlying simulation outputs.

Financial Performance

The central financial question is whether a contractual arrangement provides a viable basis for private investment in utility-scale battery storage under current Dutch market conditions.

Revenue composition is decomposed into contributions from Day-Ahead energy arbitrage, automatic Frequency Restoration Reserve (aFRR) capacity provision, imbalance market participation, and contract-specific remuneration. The cost structure comprises cycling and calendar degradation costs, annual grid access fees, and fixed operating expenditure (OPEX). This decomposition identifies which revenue streams are most affected by contractual restrictions and quantifies their contribution to overall project cashflows.

Long-term investment viability is assessed over a fifteen-year project horizon using State-of-Health (SoH)-adjusted cashflow projections, as described in Appendix B.8. The primary investment metrics are Net Present Value (NPV), evaluated at a discount rate of 6%, and Internal Rate of Return (IRR). Bankability is defined as the condition under which the project IRR exceeds the adopted hurdle rate of 8.5%, representing the minimum required return for utility-scale battery storage projects in the Netherlands, as established through expert consultation [11].

The concept of UOR, introduced in Section 3.3.8, isolates the gross operational market value displaced by contractual restrictions before any compensation is considered. Comparing UOR with contractual remuneration enables the proportionality of compensation to be assessed and allows the bankability compensation rate, defined as the minimum remuneration required for the project to satisfy the bankability criterion, to be determined for each contractual arrangement and scenario family.

Technical Performance

The central technical question is whether contractual restrictions alter the battery's physical ageing trajectory relative to unconstrained operation.

Degradation is evaluated through the annualised State-of-Health reduction and the implied asset lifetime, defined as the projected year in which the State-of-Health reaches the End-of-Life threshold of 0.80. Because degradation costs are internalised within the optimisation objective, the battery endogenously moderates its cycling intensity in response to operational constraints. Technical outcomes therefore reflect the combined effect of the contractual restriction and the asset's endogenous operational response rather than an externally imposed cycling profile.

Market participation is further evaluated through revenue contributions by market segment and daily utilisation rates. Together, these indicators identify which operational value streams are most affected by each contractual arrangement and provide the technical evidence supporting the subsequent property-based interpretation.

System Alignment

The central alignment question is whether contractual arrangements produce battery dispatch that is more consistent with underlying grid congestion needs than profit-maximising unconstrained operation.

Alignment is evaluated using the ex-post Congestion Management Indicator (CMI), described in Section 3.3.7. The CMI is applied exclusively as a post-hoc evaluation metric; it is not available to the optimisation model during dispatch decisions and does not form part of the optimisation objective. Consequently, alignment is evaluated independently of the optimisation process, allowing the framework to assess whether market-driven behaviour coincides with system-supportive operation.

Contract performance is expressed relative to the unconstrained Baseline Model through the improvement metric

$$\Delta\Omega_{\pi,r} = \Omega_{\mathrm{Baseline},r} - \Omega_{\pi,r} \tag{3.13}$$

where r denotes the Monte Carlo realisation and π the evaluated contractual arrangement. Positive values of $\Delta\Omega$ indicate that the contractual arrangement produces more congestion-supportive dispatch than the unconstrained Baseline Model under identical market conditions.

The unconstrained Baseline Model nevertheless exhibits a non-zero degree of natural alignment because commercial market incentives partially coincide with congestion patterns. Consequently, $\Delta\Omega$ isolates the incremental contribution of each contractual arrangement beyond this naturally occurring alignment.

Integrated Contractual Assessment

The individual performance dimensions are synthesised into an integrated contractual assessment to identify trade-offs that are not apparent from single-metric comparisons. Rather than evaluating financial viability, technical performance, and system alignment independently, the assessment considers their combined performance across all contractual arrangements.

Two complementary analyses structure this synthesis. First, a multi-criteria scorecard summarises the performance of each contractual arrangement across four performance dimensions: annual profit, battery lifetime, congestion alignment, and bankability. For each dimension, the scorecard reports the percentage of Monte Carlo realisations in which a contractual arrangement outperforms the unconstrained Baseline Model, providing a measure of the consistency with which each arrangement delivers improvements under uncertainty. Second, the relative performance of the contractual arrangements is compared across performance dimensions and scenario families to identify recurring strengths, weaknesses, and trade-offs. Together, these analyses distinguish arrangements that perform consistently well across multiple objectives from those whose advantages are confined to specific performance dimensions or market conditions.

The integrated contractual assessment establishes the evidential basis for the subsequent property-based interpretation. Contract-level performance patterns are first established directly from the simulation outputs before broader conclusions regarding contractual design properties are drawn. This separation ensures that property-level interpretations remain grounded in the evidence generated by the Policy Testbed.

3.4.2. Property-Based Interpretation

The property-based interpretation examines whether the performance patterns established during the contractual assessment can be associated with the four design properties identified during the institutional discovery. The objective is not to determine which specific contractual arrangement performs best, but to identify recurring relationships between design properties and observed operational, financial, and system-level outcomes across different contractual configurations and scenario families.

Because the evaluated contractual arrangements simultaneously embody multiple design properties, the analysis is explicitly observational rather than experimental. No two arrangements differ exclusively in a single design property, and causal attribution is therefore beyond the evidential scope of the Policy Testbed. Instead, the analysis identifies recurring associations and trade-offs, which subsequently form the basis for the institutional feedback mapping.

Compensation structure Compensation structure describes how financial remuneration is linked to operational restrictions. The evaluated arrangements range from the absence of compensation (Volume Bound and Firmness Gradient), through fixed tariff reductions (ATR85 and TBTR), to event-triggered capacity payments (CBC) and activation-based energy payments (CSC). The analysis examines whether the relationship between contractual remuneration and UOR is systematically associated with investment viability, using comparisons between UOR and contractual remuneration together with the bankability compensation rates derived from the analytical parameter sweeps.

Firmness Firmness describes the degree of guaranteed transport availability provided by a contractual arrangement. The evaluated designs range from probabilistic bidirectional shutdown (ATR85), through partial transport availability (CBC), to a structurally guaranteed transport floor (Firmness Gradient). The analysis assesses whether increasing firmness is associated with lower UOR magnitude, reduced variability in operational risk across Monte Carlo realisations, and improved investment performance, thereby providing evidence on the potential role of structural transport guarantees in reducing investor exposure.

Volume orientation Volume orientation describes whether contractual restrictions are expressed as instantaneous power limits or cumulative energy budgets. Within the evaluated arrangements, only the Volume Bound contract adopts an energy-based restriction. Consequently, the available evidence is limited to a single-contract observation and the resulting findings are interpreted as exploratory. The analysis evaluates whether energy-based restrictions produce systematically different operational behaviour, degradation outcomes, or congestion alignment compared with conventional power-based restrictions.

Directive control Directive control describes the allocation of dispatch authority between the asset owner and the network operator. Only the Capacity Steering Contract (CSC) incorporates active directive control. The CSC timing-quality sensitivity provides an informative within-contract comparison because the activation timing parameter is varied while the contractual structure remains unchanged. This analysis evaluates the contribution of activation timing quality to operational behaviour, congestion alignment, and financial performance, thereby distinguishing the influence of activation timing quality from the directive control mechanism itself.

3.4.3. Institutional Feedback Mapping

The final stage of the Integrated Evaluation reconnects the simulation outcomes to the structural tensions identified during the institutional discovery. Rather than treating the quantitative results as an endpoint, this stage provides the interpretive framework through which contractual performance is evaluated in relation to the governance challenges that motivated the study. The objective is not to establish causal relationships between contractual design properties and institutional outcomes, but to assess how the evidence generated by the Policy Testbed informs the stakeholder concerns identified during the institutional discovery.

The mapping evaluates whether the observed simulation outcomes are consistent with the structural tensions identified during the institutional discovery and whether particular contractual characteristics are systematically associated with improved or worsened performance with respect to those tensions. Consistent with the preceding property-based interpretation, the resulting findings remain explicitly interpretive and should not be understood as evidence of causal relationships.

Risk and Value Allocation This tension concerns the relationship between the operational risks borne by battery investors and the financial compensation provided under flexible connection agreements. Quantitative evidence is derived from the relationship between UOR, contractual remuneration, profitability, and investment viability. Particular attention is given to compensation coverage ratios, Internal Rate of Return (IRR), and the bankability compensation rate. The mapping assesses whether observed compensation structures appear consistent with mitigating the mismatch between operational risk exposure and financial value allocation identified during the institutional discovery.

Temporal Misalignment This tension concerns the mismatch between the timing of network congestion and the timing of battery market commitments. Quantitative evidence is derived from timing-sensitive performance indicators, including programme imbalance costs, congestion alignment outcomes (Ω), and the comparative performance of energy-based and power-based contractual restrictions. The mapping assesses whether alternative contractual designs appear associated with reduced temporal frictions and whether particular restriction types impose disproportionate losses during economically valuable operating periods.

Control Rights and Responsibility This tension concerns the allocation of operational authority to network operators without a corresponding allocation of financial responsibility. Quantitative evidence is derived from congestion alignment performance ($\Delta\Omega$), programme imbalance costs, compliance penalties, and the financial consequences associated with directive interventions. The mapping assesses whether arrangements granting greater operational control to the network operator are associated with additional financial burdens for asset owners and whether these burdens are accompanied by commensurate system-level benefits.

The findings from this mapping are presented in Section 4.3.3. Each institutional tension is evaluated using the quantitative evidence generated by the Policy Testbed. The objective is not to validate or reject the structural tensions identified during the institutional discovery, but to assess the extent to which the simulation evidence is consistent with, contradicts, or remains inconclusive regarding the institutional concerns raised by stakeholders.

4

Results

This chapter presents the findings of the research following the analytical framework established in Chapter 3. The results are organised according to the successive stages of the evaluation framework introduced in the methodology. The chapter begins with the institutional discovery, identifying the stakeholder frictions and structural tensions that motivate the study. It subsequently presents the unconstrained baseline Model, establishing the financial and operational reference against which all contractual arrangements are evaluated. Building on this reference, the contractual assessment evaluates the performance of the six flexible connection arrangements. Finally, the Integrated Evaluation synthesises the contract-level findings at the design-property level and interprets them in relation to the structural tensions identified during the institutional discovery.

4.1. Institutional Discovery

The institutional discovery establishes the qualitative foundation of the research results by identifying the core stakeholder frictions and structural tensions inherent in the Dutch flexible grid access framework. By defining this problem space, the analysis provides the interpretive framework through which the subsequent quantitative evidence generated by the Policy Testbed is evaluated

4.1.1. Identification of Friction Clusters

Following the systematic retrieval and selection procedure described in Section 3.2.1, a corpus of 33 regulatory and stakeholder documents was analysed. The analysis identified 112 substantive friction fragments, each coded together with its underlying institutional mechanism and operational impact (Section 3.2.1). These coded fragments were classified into five institutional friction clusters on the basis of shared mechanism and impact, summarised in Table 4.1.

The distribution of identified friction fragments is highly uneven across the five clusters. Nearly half of all fragments ($n = 47$) concern institutional allocation and design gaps, whereas temporal misalignment represents the smallest category ($n = 11$). This concentration suggests that stakeholders primarily locate congestion-related challenges in the institutional architecture underpinning flexible connection agreements rather than in the operational implementation of individual contractual arrangements.

Lower fragment counts should not be interpreted as evidence that a friction is less important. Operational control and temporal coordination arise only in relatively specific operational contexts, affecting a narrower subset of stakeholders, yet they are raised consistently wherever those situations occur.

Table 4.1: Overview of identified institutional friction clusters. Clusters were formed on the basis of shared underlying institutional mechanisms and operational impact, not similarity between friction statements alone.

Institutional friction cluster	Fragments (n)
Institutional allocation and design gaps	47
Incentive and risk misalignment	25
Information and uncertainty frictions	22
Temporal misalignment	11
Control and dispatch authority conflict	7
Total	**112**

This concentration reflects two closely related stakeholder concerns: the separation between the legal right to obtain a grid connection and the right to receive meaningful transport capacity, and the compatibility of first-come-first-served allocation logic with the policy objective of activating flexibility. These concerns do not question how existing contractual arrangements perform, but whether the underlying allocation framework is capable of producing the intended incentives at all, suggesting that incremental contractual redesign alone may be insufficient if the institutional architecture generating these frictions remains unchanged. The complete coding spreadsheet, listing all 112 fragments together with their assigned stakeholder, agreement, friction, mechanism, and cluster, is made publicly available via a GitHub repository to support transparency and reproducibility of the institutional discovery [1].

Incentive and risk misalignment ($n = 25$) and information and uncertainty frictions ($n = 22$) form a closely related pair. The former captures a perceived imbalance between operational obligations and financial compensation, including uncompensated curtailment, uncertain activation patterns, and penalty accumulation, all of which stakeholders consistently identify as barriers to investment. The latter captures the structural precondition underlying these concerns: without transparency regarding congestion conditions and activation expectations, operational risk cannot be assessed or priced appropriately. Although analytically distinct, the two clusters are institutionally intertwined, as information opacity prevents market participants from evaluating the financial consequences of flexible operation.

Temporal misalignment ($n = 11$) and control and dispatch authority conflict ($n = 7$) occur less frequently but are highly specific in nature. Temporal frictions arise at the interface between electricity market gate closures and congestion management timelines, where battery operators commit to Day-Ahead market positions before network operators identify curtailment requirements. Control and dispatch authority conflicts arise when operational authority is transferred to the network operator while imbalance settlement responsibility remains with the asset owner. Despite their lower frequency, both clusters motivate the quantitative evaluation undertaken in the Policy Testbed, where these operational interactions are examined explicitly.

Stakeholder Distribution of Frictions

Whereas the overall cluster totals identify the dominant institutional frictions, the stakeholder distribution reveals important differences in how these frictions are perceived across stakeholder groups. These differences are as informative as the areas of agreement because they illustrate how the congestion problem is interpreted from different institutional perspectives.

Table 4.2: Distribution of identified institutional frictions across stakeholder groups. Percentages represent the share of each friction category within the total number of fragments (n) attributed to the corresponding stakeholder group.

Stakeholder	Timing	Risk	Control	Information	Design
Regulators (ACM)	0.0%	22.2%	5.5%	22.2%	50.0%
Grid operators	13.3%	26.7%	13.3%	6.7%	40.0%
Market associations	13.8%	10.3%	10.3%	20.7%	44.9%
Market actors	12.0%	28.0%	2.0%	18.0%	44.0%

Three observations emerge from the stakeholder distribution.

First, institutional allocation and design gaps constitute the largest friction cluster for every stakeholder group. This consistency is notable because it spans actors with fundamentally different institutional responsibilities, suggesting broad agreement that the principal challenges lie in the institutional architecture of flexible connection agreements rather than in the operational performance of individual contractual arrangements.

Second, incentive and risk misalignment is particularly prominent among market actors (28.0%) and grid operators (26.7%). Although these stakeholder groups occupy different institutional roles, both identify operational risk allocation and the financial consequences of flexibility obligations as recurring concerns. This pattern suggests that balancing contractual obligations with appropriate economic incentives is perceived as a central challenge throughout the flexible access framework.

Third, regulators report no temporal misalignment fragments, whereas all other stakeholder groups consistently identify temporal coordination as a recurring concern. This asymmetry points to a fundamental difference in institutional perspective. Regulatory documents primarily frame flexible connection agreements as allocation mechanisms, whereas network operators and market participants place greater emphasis on the operational coordination between congestion management actions and electricity market commitments.

[1] GitHub repository

These observations provide the empirical basis for the synthesis presented in the following section. Rather than aggregating the five friction clusters mechanically, the synthesis identifies the broader institutional tensions that recur across stakeholder groups despite their different institutional perspectives and responsibilities.

4.1.2. Synthesis of Institutional Design Tensions

While the five identified friction clusters provide a structured description of recurring stakeholder concerns, they do not yet explain the broader institutional mechanisms from which those concerns arise. To move from descriptive classification towards policy evaluation, the qualitative findings are synthesised into three institutional design tensions. Rather than representing additional empirical findings, these tensions capture recurring institutional misalignments that underlie multiple observed friction clusters and recur across stakeholder groups.

Table 4.3 summarises how the empirical friction clusters are synthesised into the three institutional tensions that provide the conceptual foundation for the quantitative evaluation.

Table 4.3: Synthesis of empirical friction clusters into institutional design tensions.

Empirical friction clusters	Institutional design tension
Temporal misalignment ($n = 11$)	Temporal misalignment
Incentive & risk misalignment ($n = 25$) Information & uncertainty frictions ($n = 22$)	Asymmetric allocation of risk and value
Control & dispatch authority conflict ($n = 7$) Institutional allocation & design gaps ($n = 47$)	Control rights and operational responsibility

Temporal misalignment between system operation and market coordination Temporal misalignment is retained as a standalone institutional tension because it originates from a fundamentally different mechanism than the remaining tensions. Whereas the others concern the allocation of rights and responsibilities, temporal misalignment arises from the interaction between electricity market gate closures and real-time congestion management. Battery operators commit to market positions hours before network operators determine whether congestion intervention is required, leaving no contractual mechanism to reconcile these different decision horizons.

The tension manifests itself through late congestion notifications, restrictions that conflict with pre-committed reserve positions, and contract durations that do not align with long-term investment horizons. Stakeholder perspectives on this tension differ systematically. Grid operators primarily frame it as an operational forecasting challenge, whereas market participants experience it as a commercial exposure because post-gate-closure interventions create imbalance positions that cannot be anticipated or hedged ex ante. Although this cluster contains relatively few fragments, its operational specificity makes it disproportionately important for the subsequent quantitative evaluation.

Asymmetric allocation of risk and value This tension combines the incentive and risk misalignment cluster with the information and uncertainty cluster because both originate from the same underlying institutional mechanism. Information deficits do not create operational risk independently; rather, they prevent that risk from being assessed, priced, and managed appropriately. A curtailment whose timing, frequency, and financial consequences cannot be estimated in advance is structurally unhedgeable.

The resulting tension reflects a structural mismatch between those who bear the financial consequences of congestion management and those who capture its system-level benefits. Downside risks, including curtailment losses, uncertain activation patterns, and penalty exposure, remain concentrated on asset owners, while the value of avoided congestion accrues primarily to network operators and, more broadly, electricity consumers. Information asymmetry therefore acts as the mechanism through which operational uncertainty is transformed into financial risk that cannot be reliably assessed or priced, making proportionate risk allocation inherently difficult.

Control rights and operational responsibility This tension synthesises the control and dispatch authority conflict cluster with the broader institutional allocation and design gaps. The dominance of institutional design fragments across stakeholder groups indicates that observed operational conflicts are manifestations of a more fundamental institutional issue rather than isolated contractual problems.

The resulting tension reflects a structural mismatch between the authority to direct asset behaviour and the obligation to bear the financial consequences of that direction. Under contractual arrangements that transfer dispatch authority to the network operator, operational decisions may be taken by the DSO while imbalance settlement costs and secondary market consequences remain the responsibility of the asset owner. Missing prioritisation rules, inconsistencies between TSO and DSO congestion regimes, and incomplete definitions of non-firm transport rights all point towards the absence of a coherent allocation principle that aligns operational authority with financial responsibility.

Synthesis Together, the three institutional design tensions demonstrate that the challenges surrounding flexible grid access extend beyond the operational characteristics of individual contractual arrangements. Instead, they point towards more fundamental questions concerning the allocation of rights, responsibilities, and financial risks within the Dutch congestion-management framework.

4.1.3. Emergent Design Requirements

The three institutional design tensions identified in the preceding section define the functional requirements that contractual arrangements intended to improve flexible grid access should satisfy. Rather than prescribing optimal contractual solutions, this study translates these requirements into contractual design properties that can be instantiated through alternative connection agreements and evaluated systematically within the Policy Testbed.

Two of these design properties, firmness and volume orientation, were not represented by existing Dutch flexible connection agreements at the time of this research. Two prospective contractual arrangements were therefore developed to operationalise these properties: the Firmness Gradient and the Volume Bound arrangement. The remaining design properties, compensation structure and directive control, are represented by existing contractual arrangements included within the Policy Testbed.

Firmness Gradient: bounding downside exposure Across the institutional discovery, stakeholders consistently expressed concerns regarding the unbounded nature of operational exposure. The absence of guaranteed minimum transport rights, combined with uncertain activation patterns, leaves downside risk structurally open-ended. The resulting design requirement is therefore not the elimination of congestion management interventions, but the introduction of a mechanism that establishes a predictable upper bound on operational risk.

Within this study, this requirement is operationalised through the proposed Firmness Gradient arrangement. The arrangement instantiates the firmness design property by guaranteeing a minimum level of transport capacity under all congestion conditions while allowing the remaining capacity to remain subject to congestion management. Rather than removing operational restrictions, it transforms unbounded exposure into bounded exposure, providing an experimental configuration through which the influence of structural transport guarantees can be evaluated.

Volume Bound: aligning contractual constraints with storage operation The temporal misalignment tension originates not primarily from the frequency of operational restrictions, but from the incompatibility between instantaneous power constraints and the inherently intertemporal operating logic of battery storage. Batteries optimise charging and discharging decisions across time, whereas conventional transport restrictions constrain operation at individual moments.

Within this study, this requirement is operationalised through the proposed Volume Bound arrangement. The arrangement instantiates the volume orientation design property by expressing transport obligations in terms of cumulative energy exchange rather than instantaneous power. Consequently, the contractual constraint shifts from limiting dispatch at individual moments to constraining aggregate energy exchange over a predefined period. This preserves operational flexibility while remaining compatible with congestion management objectives.

Neither proposed arrangement is presented as a preferred contractual solution. Instead, both serve as experimental configurations through which the influence of an individual design property can be isolated and evaluated relative to existing Dutch flexible connection agreements.

4.2. Optimisation Framework

Having identified the structural tensions and the prospective contract designs intended to address them, the analysis now turns to testing these ideas quantitatively. The unconstrained Baseline model establishes the reference performance of the battery system against which all contractual arrangements are subsequently evaluated.

The baseline establishes the unconstrained reference performance of the battery system, providing the baseline against which all contractual arrangements are evaluated. By representing profit-maximising market dispatch in the absence of grid-imposed restrictions, it provides the counterfactual required to isolate the operational and financial consequences of alternative contractual designs.

Within the Policy Testbed, the baseline and all contractual arrangements are evaluated under identical market realisations. Consequently, observed performance differences can be attributed to differences in the evaluated contractual arrangements rather than to stochastic variation in the market environment.

The results presented in this section characterise the financial performance, technical behaviour, and natural congestion alignment of a utility-scale LFP battery operating under unconstrained conditions.

4.2.1. Baseline Model

The unconstrained baseline is evaluated across financial, technical, and alignment dimensions to establish the reference performance against which all subsequent policy variants are compared. This characterisation begins with an assessment of the battery system's financial viability and its sensitivity to structural market drivers.

Financial Evaluation and Scenario Sensitivity

Unconstrained market operation generates a mean annual profit of €437k (urban) and €439k (rural), a difference of less than 0.5% that reflects the revenue-side equivalence of the two congestion archetypes before any contractual layer is applied. Gross revenues are broadly distributed across DA arbitrage (17.7%), aFRR capacity reservation (48.7%), and imbalance settlement (33.5%). The annual grid fee of €1,436k absorbs approximately 70% of gross revenues. OPEX and operating costs account for the remainder. The grid fee therefore represents the dominant cost component before any contractual restriction is imposed, and the dominance of aFRR in the revenue stack means that any contract disrupting reserve participation attacks the largest value stream of the asset.

The baseline does not provide a bankable investment case on its own. Even in its most favourable scenario family, the Baseline reaches an IRR of only 3.0% (urban) and 3.8% (rural), substantially below the 8.5% hurdle rate. Performance deteriorates further as volatility and saturation move away from this best case. This shortfall follows directly from the cost structure established above: with the grid fee absorbing roughly 70% of gross revenues before any contractual restriction is applied, no scenario realisation generates sufficient residual margin to clear the hurdle rate unconstrained. Any contract that fails to improve on this baseline therefore adds operational risk to an already unviable investment.

Table 4.4 reports affinity-weighted annual profit and IRR at planning year 9 for each scenario family.

Table 4.4: Scenario-conditional Baseline performance at planning year 9. Affinity-weighted mean annual profit (k€/yr) and project IRR per scenario family and congestion regime. Dashes indicate that the cashflow trajectory does not converge to a valid IRR.

Scenario family	Profit yr 9 (k€/yr)		IRR	
	Urban	Rural	Urban	Rural
High Vol + Low Sat	767	812	3.0%	3.8%
High Vol + High Sat	554	542	−1.8%	−2.2%
Low Vol + Low Sat	164	152	-	-
Low Vol + High Sat	−9	54	-	-
Monte Carlo mean	437	439	-	-

The scenario spread reaches €776k/yr (urban) and €758k/yr (rural) at planning year 9, approximately 1.8 times the mean annual profit, reflecting the sensitivity of battery revenues to arbitrage spread compression as saturation increases. The Low Vol / High Sat scenario produces near-zero annual profit (€−9k urban; €54k rural), with the two regimes converging under compressed spread conditions. Minor differences between archetypes across scenario families reflect stochastic variation across independently sampled Monte Carlo ensembles rather than structural differences, as both regimes share identical market prices, dispatch logic, and cost assumptions under unconstrained operation.

Baseline CMI Alignment

Even in the absence of contractual restrictions, profit-maximising dispatch exhibits a degree of natural congestion alignment. The Baseline mean Ω score is -0.045 in the urban regime and -0.150 in the rural regime, indicating stronger natural congestion alignment under the rural congestion archetype.

Across all Monte Carlo runs, the battery remains idle during 84.9% of congestion-weighted intervals in the urban regime and 87.1% in the rural regime. This suggests that much of the observed baseline alignment is associated with commercial optimisation frequently resulting in inactivity during congestion-critical periods, rather than with explicit responses to network conditions.

The 5th and 95th percentile values confirm that this natural alignment is consistent across Monte Carlo runs rather than being driven by a small number of favourable realisations. The interpercentile range spans 0.021 units (−0.056 to −0.035) in the urban regime and 0.125 units (−0.218 to −0.093) in the rural regime.

Because the rural Baseline Ω is more than three times larger in magnitude than the urban value, the congestion-alignment improvement of each contractual arrangement ($\Delta\Omega$) is evaluated relative to its regime-specific baseline rather than a common zero reference. Consequently, identical $\Delta\Omega$ values should be interpreted relative to the natural alignment already present within the corresponding congestion archetype.

Baseline Degradation

Baseline unconstrained operation produces an implied battery lifetime of approximately 12.9 years before reaching the 80% State-of-Health End-of-Life threshold, slightly below the 15-year investment horizon. SoH-adjusted revenue projections in the cashflow model therefore reflect a modest capacity fade over the project lifetime. The implied lifetime is effectively identical across the urban and rural congestion archetypes, indicating no meaningful difference in baseline battery utilisation under unconstrained operation. This value serves as the reference against which contractual degradation impacts are evaluated in Section 4.2.2.

4.2.2. Contract Evaluation

Building upon the unconstrained reference case, the six contractual arrangements are evaluated across financial, system-alignment, and technical dimensions to identify the trade-offs associated with alternative contractual designs. The evaluation begins with financial performance, examining how different combinations of contractual restrictions and remuneration mechanisms influence annual profitability, long-term investment viability, investment sensitivity, and the proportionality between contractual remuneration and operational burden.

Financial Performance

The evaluated contractual arrangements can be divided into two groups: remunerated (ATR85, CBC, CSC, and TBTR) and unremunerated (Volume Bound and Firmness Gradient). Within the remunerated group, profitability differs substantially across contracts, while the financial performance of the unremunerated arrangements depends entirely on the extent to which their operational restrictions affect market participation. The resulting financial outcomes are presented in four stages. First, annual profitability is compared across contracts. Second, long-term investment viability is assessed using project IRR. Third, the robustness of investment performance is examined through sensitivity analysis. Finally, contractual remuneration is evaluated relative to the uncompensated operational risk imposed by each arrangement.

Annual profitability CBC yields the highest mean annual profit across both regimes, representing a factor of approximately 3.1 (urban) to 3.2 (rural) improvement over the Baseline. The profit ranking is identical across both regimes: CBC > CSC > ATR85 > Baseline ≈ Firmness Gradient > Volume Bound > TBTR. Figure 4.1 decomposes the revenue and cost structure per contract.

Table 4.5: Mean annual profit per contract and congestion regime (Monte Carlo mean). Δ denotes deviation from the unconstrained Baseline. Scenario spread is the range of affinity-weighted annual profit across the four scenario families at planning year 9.

Contract	Mean profit (k€/yr)		Δ vs Baseline (k€/yr)		Spread yr 9 (k€/yr)	
	Urban	Rural	Urban	Rural	Urban	Rural
Baseline	437	439	-	-	776	758
ATR85	670	731	+233	+292	422	472
CBC	1,369	1,388	+932	+949	651	625
CSC	823	886	+386	+447	360	586
TBTR	180	181	−257	−258	520	509
Volume Bound	206	208	−231	−231	693	679
Firmness Gradient	434	437	−3	−2	775	758

Diagnostic analysis confirms that CBC's improvement originates from remuneration rather than improved market access: DA revenue under CBC exceeds the Baseline by only €7k/yr (urban) and €12k/yr (rural), while DAFO costs remain close to the Baseline. The battery continues to participate fully in all three market revenue

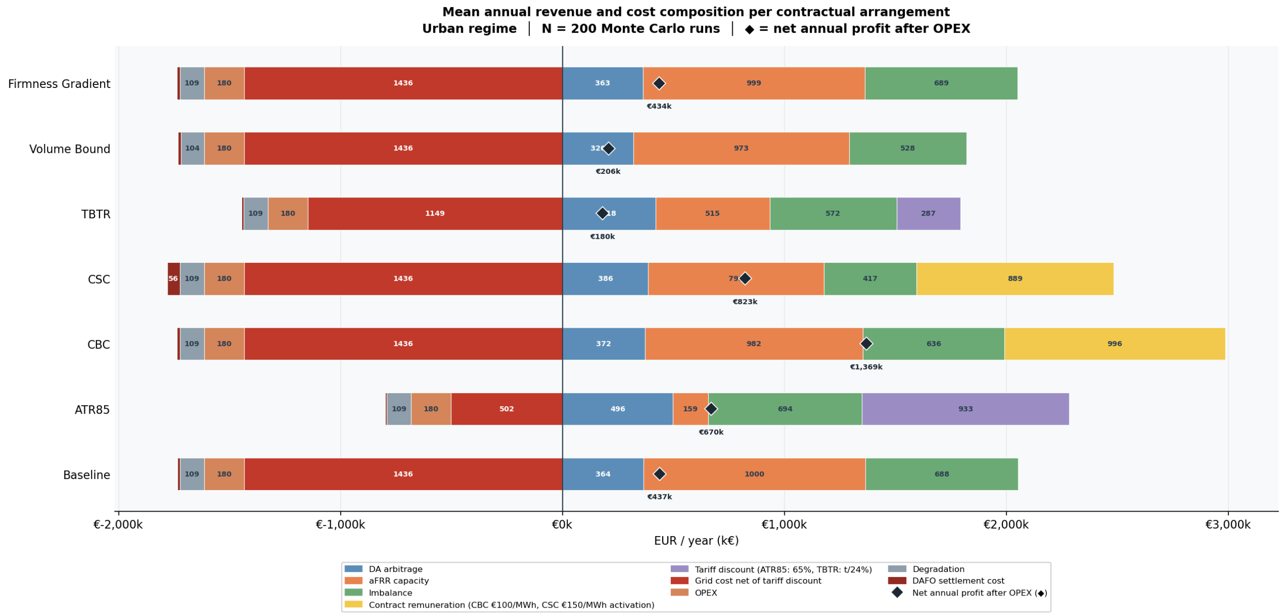

Figure 4.1: Revenue and cost decomposition per contract, urban congestion regime. Bars to the right show market revenues (DA arbitrage, aFRR capacity, imbalance) and contract remuneration; bars to the left show costs (degradation, OPEX, grid fee net of any tariff discount). The diamond marker (◆) indicates mean net profit. Light-purple segments mark the tariff discount under ATR85 (65%) and TBTR ($t/24 \approx 16.7\%$). The rural counterpart is shown in Figure D.1 (Appendix D).

streams while collecting the compensation payment. CSC and CBC are the two arrangements yielding the highest financial returns, although through structurally different mechanisms. CBC's compensation is added on top of largely unrestricted market participation, whereas CSC's remuneration only partially offsets a genuine operational cost. CSC's DAFO charges, arising from forced dispatch conflicting with pre-committed positions, reach €55,805/yr (urban) and €69,632/yr (rural), the highest in the contract set. This difference in mechanism, not merely magnitude, is why CBC's compensation structure sustains bankability where CSC's does not (see Investment Viability, below). The remaining four contracts reduce profit through restriction alone, without compensation large enough to recover Baseline-level returns: ATR85's fixed tariff discount and TBTR's tighter restriction window both fall short despite full sensitivity testing of the discount range, while Volume Bound and Firmness Gradient forgo remuneration entirely, leaving their financial outcome to be determined solely by how binding each constraint proves in practice.

Investment Viability and Bankability The final dimension of the financial evaluation assesses long-term investment viability using the project Internal Rate of Return (IRR) over a 15-year investment horizon. Tables 4.6 and 4.7 summarise the IRR achieved under each contractual arrangement across the four scenario families. Following expert validation, an IRR of 8.5% is adopted as the bankability threshold; projects exceeding this hurdle rate are considered capable of attracting project finance under current Dutch market conditions.

Table 4.6: Project IRR by contract and scenario family for the urban congestion regime. Entries exceeding the 8.5% hurdle rate are considered bankable. Dashes indicate that the cashflow trajectory does not converge to a valid IRR.

Contract	High Vol Low Sat	High Vol High Sat	Low Vol Low Sat	Low Vol High Sat
Baseline	2.99%	−1.82%	-	-
ATR85	4.80%	2.74%	−3.57%	−6.89%
CBC	20.31%	18.69%	11.92%	11.37%
CSC	8.16%	5.66%	2.64%	0.74%
TBTR	−7.85%	−15.14%	-	-
Volume Bound	−3.95%	−11.92%	-	-
Firmness Gradient	2.92%	−1.88%	-	-

The unconstrained Baseline remains non-bankable across all scenario families in both congestion regimes, reaching a maximum IRR of only 3.8%, consistent with the structural revenue shortfall established in Section 4.2.1.

Table 4.7: Project IRR by contract and scenario family for the rural congestion regime. Entries exceeding the 8.5% hurdle rate are considered bankable. Dashes indicate that the cashflow trajectory does not converge to a valid IRR.

Contract	High Vol Low Sat	High Vol High Sat	Low Vol Low Sat	Low Vol High Sat
Baseline	3.77%	−2.22%	-	-
ATR85	7.18%	4.62%	−2.08%	−3.83%
CBC	20.56%	18.05%	13.18%	12.29%
CSC	11.06%	8.83%	−0.12%	1.52%
TBTR	−6.80%	−15.83%	-	-
Volume Bound	−2.89%	−12.76%	-	-
Firmness Gradient	3.71%	−2.29%	-	-

ATR85 improves project returns relative to the Baseline through its 65% transport tariff reduction, but remains below the 8.5% hurdle rate in every scenario family. Although the tariff reduction consistently improves investment performance, it is insufficient to restore bankability.

CSC achieves bankability only in the two rural High Volatility scenario families, reaching IRRs of 11.06% and 8.83%. In all urban scenario families and both Low Volatility rural families, project returns remain below the hurdle rate.

CBC is the only contractual arrangement achieving bankability across all four scenario families in both congestion regimes. Project IRRs range from 11.37% to 20.31% in the urban regime and from 12.29% to 20.56% in the rural regime, consistently exceeding the bankability threshold under both favourable and unfavourable market conditions.

TBTR, Volume Bound, and Firmness Gradient remain non-bankable across all evaluated scenario families, frequently producing negative project returns. Despite substantially reducing operational risk, Firmness Gradient does not translate this improvement into higher long-term investment returns and therefore remains broadly comparable to the Baseline in terms of project bankability.

Investment sensitivity CBC's bankability is robust to capital cost uncertainty: the CAPEX ceiling for bankability reaches €9.0M in both High Vol families and €6.2M in both Low Vol families (urban), substantially above the €6.0M model assumption (Figure 4.2). Bankability is also preserved across discount rates up to 11–12% in the Low Vol urban families and 12–13% in the Low Vol rural families.

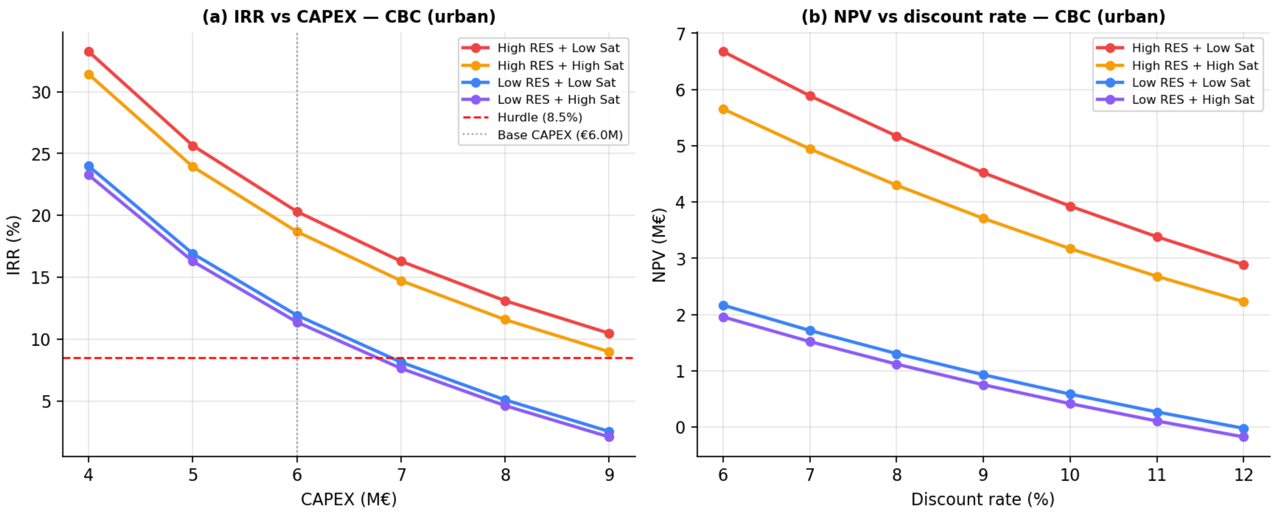

Figure 4.2: CBC investment sensitivity in the urban regime. Panel (a): IRR as a function of CAPEX by scenario family. Panel (b): NPV as a function of discount rate. Red dashed line: 8.5% hurdle rate. Dotted vertical line: base assumptions (€6.0M CAPEX; 6% discount rate). Rural counterpart in Figure D.4 (Appendix D).

The CBC compensation rate is the most consequential calibration parameter, and the only one among the three tested contracts capable of moving the arrangement across the bankability threshold: mean urban IRR crosses the 8.5% hurdle rate at approximately €60–80/MWh depending on scenario family (Figure 4.3), well within the plausible range of DSO compensation schemes. ATR85 and TBTR, by contrast, cannot reach bankability at any tested discount level within the evaluated range. Their tariff-discount mechanism scales financial return too weakly relative to the operational restriction it imposes, leaving too large a gap to the hurdle rate to close. The UOR-matching rate, the level at which remuneration exactly covers mean operational risk displacement, is €6.3/MWh (urban) and €6.5/MWh (rural), a factor of approximately 19× below the ESNL benchmark used in this study. This gap is calibration-conditional: it reflects both the specific €100/MWh compensation rate and the low UOR that results from the 3 MW ceiling chosen here. The structural finding, that CBC remuneration is substantially decoupled from operational burden, holds across the evaluated calibration range, but the precise factor depends on modelling choices.

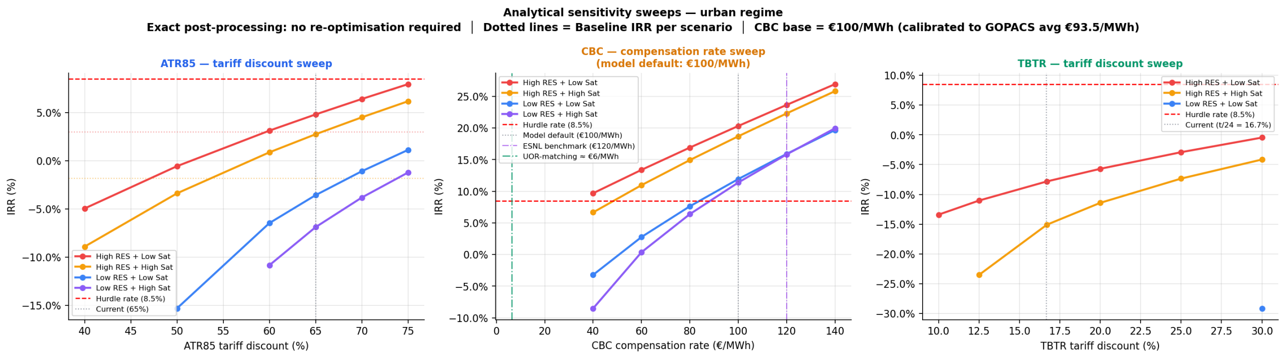

Figure 4.3: Analytical sensitivity sweeps, urban regime. Three contracts are shown: ATR85 (tariff discount), CBC (compensation rate), TBTR (tariff discount). Dotted lines mark scenario-conditional Baseline IRRs. Red dashed line: 8.5% hurdle rate. The green dash-dot line on the CBC panel marks the UOR-matching compensation rate (€6.3/MWh). Rural counterpart in Figure D.2 (Appendix D).

Compensation Structure and Uncompensated Operational Risk The relationship between contractual remuneration and UOR provides an indication of whether financial compensation is proportionate to the operational burden imposed by each contractual arrangement. Figure 4.4 plots mean annual remuneration against mean UOR, while Table 4.8 summarises the corresponding UOR, remuneration, and coverage ratios for each contract.

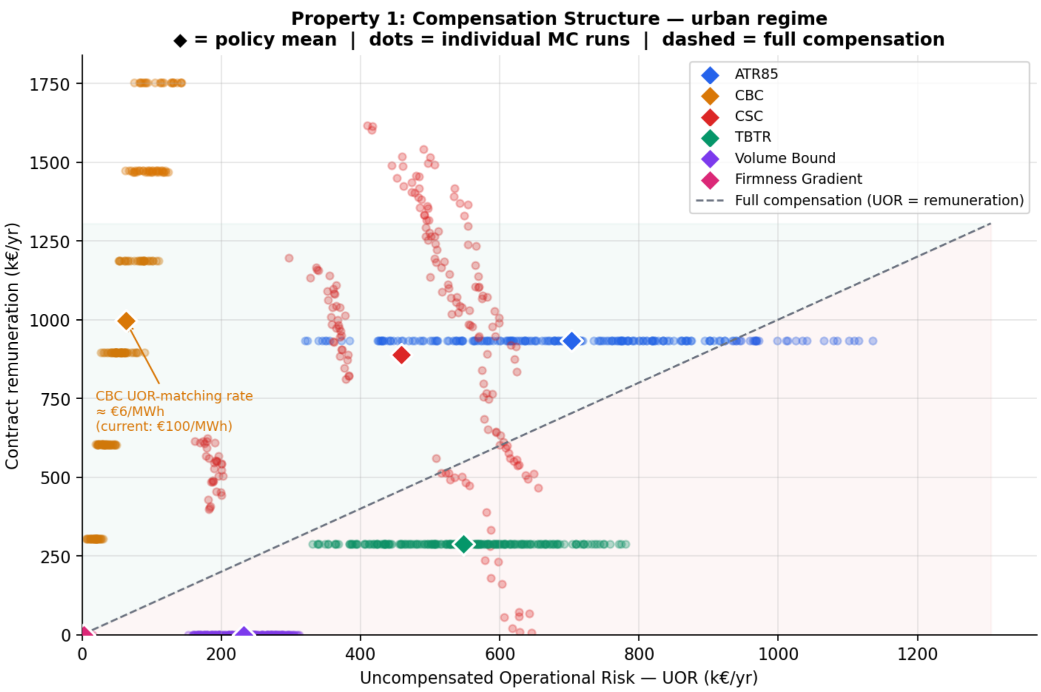

Figure 4.4: Compensation structure, urban regime. Mean UOR versus mean annual remuneration per contract (♦ = policy mean; semi-transparent dots = individual Monte Carlo runs). Dashed diagonal: full compensation line (UOR = remuneration). Green shading: overcompensated region; red shading: undercompensated. Rural counterpart in Figure D.3 (Appendix D).

No evaluated contractual arrangement lies close to the full compensation line. Instead, the results separate into two distinct groups. ATR85, CBC, and CSC provide remuneration substantially exceeding their measured

Table 4.8: UOR and remuneration summary by contract and congestion regime. Coverage ratio = remuneration / UOR. Positive compensation gaps indicate remuneration exceeding measured operational burden, whereas negative gaps indicate undercompensation. Volume Bound and Firmness Gradient provide no remuneration by design.

Contract	Urban			Rural		
	UOR (k€)	Rem. (k€)	Coverage	UOR (k€)	Rem. (k€)	Coverage
ATR85	703	933	133%	644	933	145%
CBC	62	996	1,606%	66	1,016	1,539%
CSC	458	889	194%	296	800	270%
TBTR	547	287	52%	549	287	52%
Volume Bound	232	0	0%	234	0	0%
Firmness Gradient	2	0	0%	3	0	0%

operational burden, whereas TBTR, Volume Bound, and Firmness Gradient remain below full compensation, with the latter two providing no remuneration by design. The observed distribution therefore indicates that the evaluated contracts tend towards either substantial overcompensation or substantial undercompensation, rather than remuneration proportional to operational burden.

CBC represents the most pronounced case of overcompensation, achieving coverage ratios exceeding 1,500%. The resulting coverage ratios indicate that contractual remuneration substantially exceeds the measured operational burden associated with congestion management under the assumptions of this study. At the opposite end of the spectrum, Firmness Gradient reduces UOR to near-zero levels through its guaranteed transport rights but provides no direct remuneration, whereas Volume Bound leaves a substantially larger operational burden uncompensated. TBTR occupies an intermediate position, compensating approximately half of the measured UOR in both congestion regimes.

CSC illustrates that nominal coverage ratios do not fully capture the effective operational burden. The reported coverage ratios exclude DAFO charges generated by compliant forced dispatch, amounting to €55,805/yr and €69,632/yr, respectively. Incorporating these charges gives a combined operational burden of UOR plus DAFO. Because these two cost terms come from disjoint parts of the objective function, they can be added without double-counting. On this burden-adjusted basis, the coverage ratio falls to approximately 173% in the urban regime and 219% in the rural regime, still above the full-compensation threshold.

This adjustment is unique to CSC because only its directive-control mechanism generates this additional cost category. Consequently, UOR-based coverage ratios should be interpreted alongside the underlying contractual cost structure rather than as directly comparable measures across arrangements with structurally different operational cost profiles.

System Alignment Outcomes

Beyond financial performance, the contractual arrangements are evaluated on the extent to which battery dispatch aligns with congestion management needs. Throughout this section, more negative Ω values indicate stronger congestion-supportive alignment, while positive $\Delta\Omega$ values indicate improvement relative to the regime-specific unconstrained Baseline. Table 4.9 summarises the mean alignment score Ω and the corresponding Baseline-referenced improvement $\Delta\Omega$ across both congestion regimes.

Table 4.9: CMI alignment outcomes per contract and congestion regime. Ω is the mean ex-post alignment score (more negative indicates stronger congestion-supportive alignment). $\Delta\Omega$ is the improvement relative to the unconstrained Baseline (positive values indicate improved alignment). Results represent Monte Carlo means across $N = 200$ runs.

Contract	Mean Ω		$\Delta\Omega$ vs Baseline	
	Urban	Rural	Urban	Rural
Baseline	−0.0448	−0.1504	–	–
ATR85	−0.0726	−0.1870	+0.028	+0.037
CBC	−0.0489	−0.1392	+0.004	−0.011
CSC	−0.2366	−0.3184	+0.192	+0.168
TBTR	−0.0772	−0.1802	+0.032	+0.030
Volume Bound	−0.0755	−0.0835	+0.031	−0.067
Firmness Gradient	−0.0466	−0.1504	+0.002	≈0

Alignment performance Three observations emerge from Table 4.9.

First, CSC delivers by far the largest improvement in congestion-supportive alignment under both congestion regimes. Mean alignment improves by $\Delta\Omega = +0.192$ in the urban regime and $\Delta\Omega = +0.168$ in the rural regime, representing an improvement approximately six times larger than the next-best contract in the urban regime. No other contractual arrangement approaches this level of alignment improvement.

Second, ATR85, TBTR, and Volume Bound all produce modest improvements under the urban congestion archetype. In contrast, CBC and Firmness Gradient remain close to the unconstrained Baseline, indicating that these arrangements alter battery operation only marginally from a congestion alignment perspective despite producing substantially different financial outcomes.

Third, alignment performance is not equally robust across congestion regimes. Whereas CSC maintains a substantial improvement under both urban and rural conditions, CBC and Volume Bound reduce alignment relative to the unconstrained Baseline in the rural regime. This indicates that alignment improvements observed under one congestion archetype do not necessarily generalise to another.

Across the three observations above, the results demonstrate substantial variation in the ability of contractual arrangements to improve congestion-supportive dispatch. While several contracts produce small improvements relative to the unconstrained Baseline, only CSC consistently achieves a large alignment benefit across both congestion regimes.

The magnitude of CSC's alignment improvement is partly determined by the fixed 5 MW activation setpoint assumed in the simulations. Consequently, the reported $\Delta\Omega$ values should be interpreted as representing a high-intervention configuration. Alternative setpoint strategies would be expected to alter the magnitude of the alignment improvement, although not the relative comparison between contracts presented here.

Volume orientation: rural alignment reversal The Volume Bound rural reversal is directionally consistent across all 200 Monte Carlo runs. In the urban regime the throughput budget produces a modest positive alignment effect ($\Delta\Omega = +0.031$); in the rural regime the same constraint produces a negative effect ($\Delta\Omega = -0.067$). The mechanism underlying this reversal is examined in Section 4.3.1.

Technical Outcomes

Table 4.10 summarises annual degradation outcomes across all evaluated contracts.

Table 4.10: Degradation outcomes per contract and congestion regime. ΔD is the median run-matched annual State-of-Health reduction relative to the unconstrained Baseline (negative = less wear than Baseline; positive = more wear). Implied lifetime is the projected year at which SoH reaches the 0.80 End-of-Life threshold. Values within solver tolerance ($|\Delta D| < 0.0001$ SoH/yr) are not meaningfully distinguishable from zero.

Contract	ΔD (SoH/yr)		Implied lifetime (yr)	
	Urban	Rural	Urban	Rural
Baseline	-	-	12.90	12.90
ATR85	$+0.000010$	$+0.000010$	12.90	12.90
CBC	-0.000005	-0.000006	12.91	12.91
CSC	$+0.000008$	$+0.000009$	12.90	12.90
TBTR	$+0.000010$	$+0.000010$	12.90	12.90
Volume Bound	-0.000227	-0.000229	13.10	13.10
Firmness Gradient	≈ 0	≈ 0	12.91	12.91

Five contracts, ATR85, CBC, CSC, TBTR, and Firmness Gradient, produce ΔD values within the predefined solver tolerance of 0.0001 SoH/yr and are therefore degradation-equivalent to the unconstrained Baseline. Volume Bound is the sole exception: its ΔD of approximately -0.000227 to -0.000229 SoH/yr is 20 to 25 times larger in magnitude than the next-largest effect in the contract set, directionally consistent across both congestion regimes, and associated with an implied lifetime extension from approximately 12.90 to 13.10 years.

Figure 4.5 visualises the distribution of run-matched degradation outcomes across the full Monte Carlo ensemble.

The Value of Activation Timing (Timing Premium)

The timing premium quantifies the financial value of accurate congestion targeting by comparing identical optimisation runs under fully random and fully CMI-aligned activation timing. Each sampled scenario was solved twice while holding all market conditions, congestion events, and technical parameters constant, thereby isolating the contribution of activation timing to annual profitability.

The timing premium differs markedly across contractual arrangements, and its ranking is not uniform across congestion archetypes. In the urban regime, CSC exhibits by far the largest financial benefit from accurate

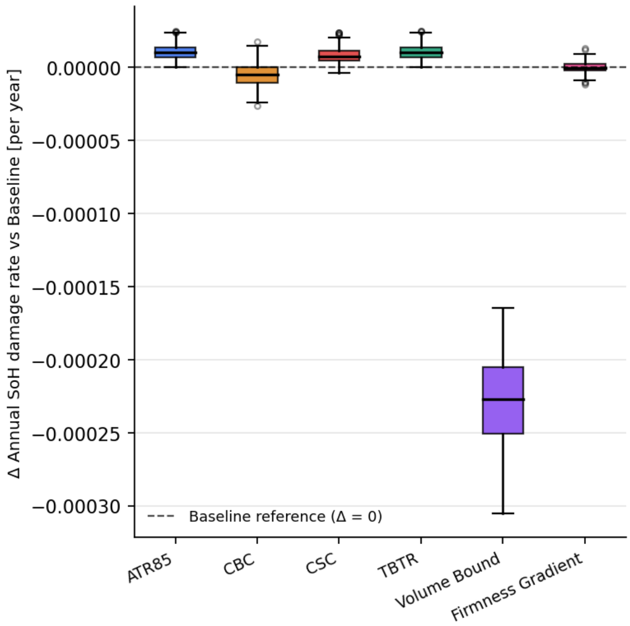

Figure 4.5: Distribution of annual State-of-Health damage rate relative to the unconstrained Baseline per contractual arrangement, urban regime ($N = 200$ Monte Carlo runs per contract). Negative values indicate lower annual degradation than the Baseline; positive values indicate higher degradation. Box boundaries denote the interquartile range, the horizontal line the median, whiskers 1.5 times the interquartile range, and circles outliers. The dashed horizontal line marks the Baseline reference ($\Delta = 0$).

Table 4.11: Timing premium per contract: difference in mean annual profit between fully CMI-aligned and fully random activation timing (k€/yr), obtained via paired re-solve ($N = 16$ stratified runs, 4 per scenario family). Positive values indicate that aligned timing improves profitability relative to random timing; negative values indicate that aligned timing displaces more market value than random timing. The CSC rural figure is derived from the endpoints ($b_{dis} = 0.0$ and $b_{dis} = 1.0$) of the CSC timing-quality sensitivity sweep (Section 4.2.2); the underlying computation is equivalent. TBTR is excluded: its restriction window is fixed and does not respond to the dis_bias parameter used to construct the random/aligned counterfactual for the other contracts, so a comparable timing premium cannot be defined under this methodology.

	Urban (k€/yr)		Rural (k€/yr)	
Contract	Mean	Median	Mean	Median
ATR85	+25.7	+27.1	+126.8	+118.3
CBC	+7.2	+6.1	+0.5	−0.5
CSC	+115.0	+97.7	+87.2	+99.7
Firmness Gradient	−0.9	−0.4	+0.5	+1.4

activation timing, more than four times the premium observed for ATR85, indicating that its performance depends strongly on whether dispatch instructions are aligned with congestion-critical periods. In the rural regime, however, ATR85 exhibits the largest timing premium of all evaluated contracts, exceeding that of CSC. This indicates that timing quality is not a benefit confined to directive control: for ATR85's full bidirectional shutdown, whether curtailment coincides with the midday solar-peak hours that drive rural congestion is at least as financially consequential as the precision of CSC's steering instructions. By contrast, CBC and Firmness Gradient remain largely insensitive to activation timing, with timing premia close to zero under both regimes. ATR85 therefore occupies an intermediate position in the urban regime, benefiting from improved timing despite not explicitly incorporating dispatch targeting within its contractual design, while in the rural regime its timing sensitivity exceeds that of CSC. TBTR is not evaluated in this comparison, since its restriction schedule is fixed by calendar rather than by market or congestion conditions, and therefore has no activation-timing decision for the paired re-solve methodology to test.

The contrast between CSC and CBC/Firmness Gradient reflects a fundamental difference in how timing enters the contractual design. Under CSC, activation timing directly determines when forced dispatch occurs, allowing well-targeted interventions to coincide more closely with commercially favourable battery operation. ATR85's fixed-probability, bidirectional shutdown similarly derives substantial value from timing, particularly in the rural regime, even though its activation is not directed by an explicit control signal in the way CSC's is. Under CBC and Firmness Gradient, by contrast, profitability is driven primarily by capacity rights and remuneration rather than the precise timing of congestion events, leaving little opportunity for additional value creation through improved targeting.

Although the financial benefit of accurate timing is evident under both congestion archetypes, the underlying operational behaviour differs. In the urban regime, the additional value arises primarily from discharge actions that coincide with evening demand peaks, whereas in the rural regime it originates predominantly from charging during periods of midday solar generation. Despite these different dispatch patterns, improved targeting consistently increases the economic value of CSC while leaving the performance of CBC and Firmness Gradient largely unchanged.

Sensitivity Analysis

Robustness and Scenario Performance The preceding sections compared contracts using mean performance. This section examines whether those rankings remain robust across the full Monte Carlo uncertainty space. Table 4.12 reports the proportion of run-matched simulations in which each contract outperforms the unconstrained Baseline.

Table 4.12: Multi-criteria scorecard, urban congestion regime ($N = 200$ Monte Carlo runs). Each cell reports the proportion of runs in which the contract outperforms the run-matched unconstrained Baseline on the stated dimension. Annual profit: net annual cashflow exceeds Baseline. Lifetime: implied SoH trajectory reaches End-of-Life later than Baseline. CMI: realised $\Delta\Omega > 0$. Bankable: scenario-family IRR exceeds the 8.5% hurdle rate. Rural counterpart in Table D.1 (Appendix D).

Contract	Annual profit	Battery lifetime	CMI alignment	Bankable scenarios
ATR85	87.5%	0%	100%	0%
CBC	100%	77.0%	82%	100%
CSC	79.5%	4.0%	100%	0%
TBTR	0%	0%	100%	0%
Volume Bound	0%	100%	100%	0%
Firmness Gradient	24.5%	53.5%	73%	0%

Four observations emerge from the robustness analysis.

First, the financial ranking established in the mean-value analysis remains largely unchanged across the full uncertainty space. CBC is the only contractual arrangement that consistently combines superior annual profitability with robust investment viability, whereas the remaining contracts fail to achieve bankability despite improving annual profitability under a substantial share of scenarios.

Second, profitability, technical performance, and congestion alignment prove to be largely independent dimensions of contract performance. ATR85 and CSC improve annual profitability in most realisations without consistently extending battery lifetime, whereas Volume Bound exhibits the opposite behaviour by consistently improving battery lifetime while sacrificing financial performance. TBTR, in turn, consistently improves congestion alignment but does not translate this improvement into financial or technical gains.

Third, Firmness Gradient exhibits the most balanced performance profile. Rather than dominating any single evaluation dimension, it delivers moderate improvements across several criteria without consistently outperforming the Baseline on either financial or technical outcomes. This pattern is consistent with the earlier finding that reducing operational risk alone does not necessarily translate into higher investment returns.

The rural scorecard (Appendix D) preserves the overall ordering of contractual performance. The principal difference is a reduction in the robustness of CBC's alignment performance, consistent with the rural alignment reversal identified in Section 4.2.2, while CSC achieves bankability in a larger share of scenario families. The remaining comparative relationships are largely unchanged across the two congestion archetypes.

Scenario spread Table 4.13 reports the scenario spread at planning year 9, defined as the range of affinity-weighted mean annual profit across the four scenario families.

The scenario spread quantifies the sensitivity of annual profitability to differences in market volatility and renewable saturation. Smaller spreads indicate more stable financial performance across scenario families.

In the urban regime, CSC exhibits the smallest scenario spread, followed by ATR85, indicating that these arrangements produce the most consistent financial outcomes despite their different contractual designs. TBTR and CBC occupy an intermediate position, while Volume Bound retains a relatively wide spread. Firmness Gradient is virtually indistinguishable from the unconstrained Baseline, indicating that guaranteeing a minimum transport right does not materially reduce financial sensitivity to changing market conditions.

The ordering changes in the rural regime. ATR85 exhibits the smallest spread, whereas CSC becomes substantially more sensitive to scenario conditions than in the urban case. Despite these shifts, Firmness

Table 4.13: Scenario spread at planning year 9: range of affinity-weighted mean annual profit across the four scenario families (k€/yr).

Contract	Urban (k€/yr)	Rural (k€/yr)
Baseline	776	758
ATR85	422	472
CBC	651	625
CSC	360	586
TBTR	520	509
Volume Bound	693	679
Firmness Gradient	775	758

Gradient and Volume Bound again remain among the widest-spread arrangements, while CBC and TBTR occupy intermediate positions.

Overall, the scenario spread complements the mean-value analysis by showing that contractual arrangements differ in both expected profitability, and in the consistency with which that profitability is maintained across different market environments.

CSC Sensitivity Analysis Three parameters governing CSC's design are examined in turn, activation rate, deviation penalty, and DSO targeting quality, to isolate which aspect of the contract each performance dimension depends on. The CSC activation payment sensitivity evaluates whether varying the remuneration rate between €100/MWh and €180/MWh, holding the deviation penalty fixed at €400/MWh, alters battery dispatch behaviour or congestion alignment. Re-optimisation is conducted over $N = 40$ stratified runs (10 per scenario family) using the delta method relative to the €150/MWh default, with the limited-foresight assumption applied consistently with the main simulation. The sanity check at the default rate confirms a mean delta of €0k, validating the approach. Table 4.14 summarises the corrected urban and rural results.

Table 4.14: CSC activation rate sensitivity summary. Mean annual profit and mean CMI alignment score Ω are reported for each evaluated activation rate. Penalty fixed at €400/MWh throughout. $N = 40$ stratified runs (10 per scenario family), foresight-corrected. Ω is measured only over PTUs during which the contract is actively steering dispatch (delta-method scope), and is therefore not directly comparable in absolute magnitude to the full-year Ω reported in Table 4.9; the relative flatness across activation rates remains the relevant finding regardless of scope.

Activation rate	Mean profit (k€/yr)		Mean Ω (active windows)	
(€/MWh)	Urban	Rural	Urban	Rural
100	149	220	−0.1425	−0.2155
120	447	583	−0.1423	−0.2155
140	745	786	−0.1426	−0.2156
150 (default)	**893**	**927**	**−0.1427**	**−0.2156**
160	1,042	1,068	−0.1427	−0.2156
180	1,339	1,349	−0.1426	−0.2154
Range	+1,190	+1,129	<0.0004	<0.0002

Financial returns increase approximately linearly with the activation rate in both congestion regimes; the rural IRR-by-scenario-family trend is shown in Figure D.5. The two High Volatility scenario families reach the 8.5% bankability threshold at or below the default activation payment, whereas the two Low Volatility families require higher remuneration before becoming bankable. In contrast, congestion alignment remains effectively unchanged throughout the evaluated range, with total variation in Ω below 0.0004 in the urban regime and below 0.0002 in the rural regime. Within the tested parameter range, the activation payment therefore has a much stronger influence on financial performance than on realised dispatch behaviour.

Deviation penalty sensitivity The deviation penalty sensitivity examines the effect of reducing the default penalty from €400/MWh to €250/MWh and €150/MWh while maintaining the activation payment at its default value of €150/MWh. Table 4.15 summarises the resulting changes in financial performance and congestion alignment for both congestion regimes.

Reducing the deviation penalty consistently improves financial performance in both congestion regimes, with the vast majority of re-optimised runs outperforming the reference case. The financial benefit increases progressively as the penalty is reduced, indicating that lower penalties allow the optimiser to recover commercially attractive dispatch opportunities that would otherwise be sacrificed to contractual compliance.

Table 4.15: CSC penalty rate sensitivity summary. Mean CMI alignment score Ω (active-window scope) and profit change relative to the €400/MWh reference, per penalty level. Activation rate fixed at the €150/MWh default. $N = 40$ stratified runs (10 per scenario family), foresight-corrected. Δ figures denote paired, run-matched deltas against the reference penalty.

Penalty rate	Mean Ω		Median / Mean Δ profit (k€/yr)		Runs improving
(€/MWh)	Urban	Rural	Urban	Rural	Urban / Rural
150	−0.1202	−0.2308	+423 / +710	+403 / +687	39/40 / 39/40
250	−0.1312	−0.2126	+159 / +351	+166 / +364	35/40 / 34/40
400 (ref.)	**−0.1427**	**−0.2156**	**0 / 0**	**0 / 0**	—

The alignment response differs between the two congestion archetypes. In the urban regime, congestion alignment weakens consistently as the penalty is reduced, indicating that lower penalties make deviations from DSO dispatch instructions progressively more attractive from a market perspective. The rural results exhibit the same tendency between the two highest penalty levels but do not continue this pattern at the lowest evaluated penalty. As this behaviour could not be independently verified within the scope of this study, it is reported without further interpretation and is not used in the principal conclusions of this thesis.

Overall, the sensitivity analysis indicates that the deviation penalty governs a trade-off between financial performance and congestion alignment. Lower penalties consistently improve financial outcomes, whereas their effect on alignment is robust only in the urban regime. The corresponding relationship in the rural regime requires further verification before broader conclusions can be drawn.

CSC timing quality sensitivity The main simulation identified CSC as the strongest-performing contract in terms of congestion alignment. This sensitivity analysis isolates the contribution of DSO targeting quality by varying the dispatch-bias parameter (b_{dis}) while holding all other model assumptions constant. Unlike the activation-rate and penalty-rate sensitivities, alignment is evaluated over the full annual simulation horizon, making the reported Ω values directly comparable with the main results in Table 4.9.

Table 4.16: CSC timing quality sensitivity summary, urban regime. Mean CMI alignment score Ω (full-year scope) and scenario-weighted IRR by scenario family for five DSO timing quality levels. $N = 16$ freshly re-solved stratified runs per level (4 per scenario family), foresight corrected. IRR is calculated using the scenario-affinity-weighted methodology of Section 3.3.8 (Eq. 3.4).

b_{dis}	Ω (urban)	IRR Sc1	IRR Sc2	IRR Sc3	IRR Sc4	Bankable families
0.00 (random)	−0.005	9.71%	6.82%	−0.21%	5.35%	1/4
0.25	−0.071	9.03%	6.25%	−1.20%	5.05%	1/4
0.50 (50/50 mix)	−0.132	9.77%	6.78%	−0.82%	5.05%	1/4
0.75	−0.189	11.13%	7.80%	0.61%	5.77%	1/4
1.00 (fully aligned)	−0.244	12.95%	9.71%	2.23%	7.16%	2/4

Improving DSO targeting quality produces a consistent improvement in congestion alignment in the urban regime. Random activation generates little additional alignment beyond the unconstrained Baseline, whereas progressively more accurate targeting steadily strengthens CSC's congestion-supportive behaviour. Under fully aligned activation, the resulting alignment closely matches the main simulation, indicating that CSC's performance depends primarily on activation timing rather than on directive control alone. The small difference relative to the main simulation is consistent with the smaller stratified sample used here ($N = 16$ versus $N = 200$).

The financial response is more conditional. Better targeting improves returns across all scenario families, but only the scenarios already close to the investment hurdle become bankable as timing quality improves. Favourable market conditions therefore remain the primary determinant of bankability, while accurate DSO targeting can determine whether marginal projects cross the investment threshold.

These scenario-level results should not be compared directly with the main scorecard (Table 4.12), where timing quality is randomly sampled across the full Monte Carlo ensemble rather than held fixed at individual levels.

The rural regime exhibits the same monotonic improvement in alignment with increasing targeting quality, confirming that accurate dispatch timing is essential for realising the system benefits of directive control. Unlike the urban regime, however, investment performance is only weakly affected by timing quality. The same scenario families remain bankable throughout the sweep, indicating that favourable market conditions dominate the investment outcome, while improved DSO targeting primarily increases congestion alignment.

Table 4.17: CSC timing quality sensitivity summary, rural regime. Mean CMI alignment score Ω (full-year scope) and scenario-weighted IRR by scenario family for five DSO timing quality levels. $N = 16$ freshly re-solved stratified runs per level (4 per scenario family), foresight corrected, using the same methodology as Table 4.16.

b_{dis}	Ω (rural)	IRR Sc1	IRR Sc2	IRR Sc3	IRR Sc4	Bankable families
0.00 (random)	−0.130	13.89%	8.65%	−10.64%	2.58%	2/4
0.25	−0.178	14.38%	9.02%	−9.42%	2.86%	2/4
0.50 (50/50 mix)	−0.223	14.67%	9.41%	−9.24%	2.83%	2/4
0.75	−0.265	15.26%	9.50%	−8.14%	2.65%	2/4
1.00 (fully aligned)	−0.304	16.44%	10.68%	−6.22%	2.86%	2/4

Across both regimes, better DSO targeting consistently strengthens congestion alignment, but its financial value is regime-dependent: in the urban regime it can move near-bankable projects across the investment hurdle, whereas in the rural regime it improves alignment without materially affecting bankability.

The implications of these findings for the institutional design tensions identified in Section 4.1.2 are discussed in Section 4.3.3.

4.3. Integrated Evaluation

The Integrated Evaluation interprets the contract-level findings generated within the optimisation Framework at a higher level of abstraction by examining whether the observed performance patterns can be associated with specific design properties. The objective is to assess how these design properties relate to the institutional frictions and structural tensions identified during the institutional discovery. By shifting the unit of analysis from individual contractual arrangements to their underlying design properties, the evaluation identifies recurring patterns that extend beyond the performance of individual contracts.

The integrated evaluation consists of three complementary components. First, the property-based interpretation examines how the four design properties, compensation structure, firmness, volume orientation, and directive control, are associated with financial, technical, and system-level outcomes. Second, the integrated property synthesis evaluates the trade-offs that emerge across the different design properties. Finally, the institutional feedback mapping reconnects the quantitative findings to the structural tensions identified during the institutional discovery, thereby closing the analytical loop of the study.

Consistent with the research design, all property-level findings are explicitly observational. Because no two contractual arrangements differ exclusively in a single design property, the analysis identifies associations between design properties and performance outcomes rather than isolated causal effects.

4.3.1. Property-based interpretation

Compensation Structure

Compensation structure describes how financial transfers are linked to operational restrictions. Across the six evaluated contractual arrangements, compensation ranges from the absence of financial transfers (Volume Bound and Firmness Gradient), through fixed proportional tariff reductions decoupled from realised curtailment (ATR85 and TBTR), to event-triggered capacity payments (CBC) and activation-based energy payments linked to directive dispatch events (CSC).

At the property level, this spectrum maps onto a single clear association: only the event-triggered capacity payment (CBC), whose coverage ratio exceeds 1,500% of measured UOR (Table 4.8), is associated with investment viability. The fixed proportional tariff reductions (ATR85, TBTR) and the two uncompensated arrangements (Volume Bound, Firmness Gradient) all remain non-bankable regardless of their own coverage ratio. This indicates an association between the scale of compensation coverage and bankability, rather than between precise compensation calibration and bankability, and shows that restricting operational flexibility without a commensurate financial transfer is insufficient to restore investment viability.

The CBC sensitivity analysis (Section 4.2.2) further shows that simply compensating the measured operational burden is insufficient: bankability is only achieved at compensation levels of approximately €60–80/MWh in the urban regime, far above the €6.3/MWh UOR-matching rate. Among the evaluated arrangements, only CBC provides remuneration of this order of magnitude, explaining why it is the only contract that consistently satisfies the bankability criterion.

Firmness

Firmness captures the extent to which transport availability remains guaranteed under a contractual arrangement. The evaluated contracts span a broad spectrum of firmness levels, ranging from probabilistic shutdown without guaranteed transport rights (ATR85) to a structurally guaranteed 3 MW firm floor (Firmness Gradient). Between these extremes lie passive capacity ceilings with preserved residual access (CBC), deterministic time-window restrictions (TBTR), and daily energy budgets with a soft power envelope (Volume Bound).

The central question for the firmness property is whether greater guaranteed access is associated with lower and less variable UOR. The evaluated contract set supports this: Firmness Gradient, the only arrangement providing a structural firm floor, achieves the lowest and least variable UOR of all six contracts (Table 4.8), while CBC's preserved residual access, without an explicit firmness guarantee, achieves the second-lowest, indicating that operational access itself, not only formal firmness, limits downside exposure.

At the opposite end of the spectrum, ATR85 exhibits both the highest mean UOR and the widest distribution of operational risk. This observed association between lower firmness and higher, more variable UOR is directionally consistent with the institutional discovery, in which probabilistic access rights were repeatedly identified as a source of uncertain downside exposure. However, the relationship cannot be attributed to firmness alone, since ATR85 also differs from the remaining arrangements in its shutdown mechanism, operational activation logic, and compensation structure.

Firmness Gradient provides the most analytically isolated observation of the firmness property. By design, it contains no explicit remuneration, allowing the effect of guaranteed transport access to be evaluated independently of financial transfers. Despite reducing UOR to near-zero levels, it produces mean annual profits within €3k/yr of the unconstrained Baseline and fails to achieve bankability in any evaluated scenario family. Within the evaluated contract set, greater firmness therefore appears to reduce operational downside exposure without materially improving investment viability under the analysed conditions.

Volume Orientation

Volume orientation captures whether contractual restrictions are expressed through instantaneous power limits or cumulative energy budgets. Among the evaluated arrangements, only Volume Bound adopts an explicitly energy-oriented design by limiting daily throughput rather than instantaneous power. The evidence for this property is therefore based on a single contractual arrangement and should be interpreted as exploratory.

Relative to the power-oriented arrangements, the Volume Bound design is associated with three consistent outcomes: improved technical performance through reduced battery degradation, lower financial performance, and a congestion-alignment effect that depends on the underlying congestion regime. Together, these outcomes characterise the behaviour of the only energy-oriented contract evaluated in this study.

The technical benefit is consistent with the daily energy budget limiting cumulative throughput and thereby reducing high depth-of-discharge cycling. Because degradation costs increase non-linearly with cycle depth, constraining daily energy throughput naturally suppresses the most damaging operating patterns and extends battery lifetime. This benefit should not be interpreted as unique to energy-budget orientation itself, however, as any sufficiently restrictive throughput constraint would be expected to produce a similar effect.

The most analytically significant finding concerns congestion alignment. Under the urban congestion archetype, the daily energy budget produces a modest improvement relative to the unconstrained Baseline. Under the rural archetype, however, this relationship reverses and alignment deteriorates. This reversal follows directly from the optimisation logic. Because the battery operates under a fixed daily energy allowance, it allocates the available throughput towards the commercially most valuable actions. In the rural regime, these are evening discharge and aFRR reservation rather than the midday charging behaviour that most effectively relieves generation-driven congestion. The limited energy budget is therefore allocated away from the congestion-supportive action, reducing system alignment despite improving technical performance.

Financially, Volume Bound remains below the unconstrained Baseline and does not produce bankable outcomes under the evaluated assumptions. The throughput restriction reduces exposure to high-value market opportunities without providing compensating remuneration, indicating that the technical benefit alone is insufficient to offset the associated revenue loss.

As only a single uncompensated contract represents this property, no general conclusion regarding energy-oriented contractual design can be drawn. Within the evaluated scenarios, however, the results suggest that energy-budget orientation represents a clear technical-financial trade-off, while its effect on congestion alignment depends on the underlying congestion archetype.

Directive Control

Directive control describes the extent to which operational dispatch authority is transferred from the asset owner to the network operator. Among the evaluated contractual arrangements, only CSC embodies active directive control, whereby the DSO issues explicit dispatch instructions during congestion events, with compliance incentivised through activation payments and penalised through deviation charges.

Among the evaluated contractual arrangements, directive control is associated by far with the strongest congestion alignment, consistent with the margin already reported in Table 4.9 and Section 4.2.2. This improvement is observed consistently across the Monte Carlo ensemble and reflects the fact that dispatch instructions are weighted by the Congestion Management Indicator (CMI), concentrating activation during congestion-critical periods.

Two observations qualify this result. First, the alignment benefit is strongly dependent on targeting quality. The timing-quality sensitivity analysis (Table 4.16) shows that alignment improves progressively as dispatch instructions become more consistent with the CMI. Under random activation, CSC performs little better than the unconstrained Baseline, whereas fully aligned activation reproduces the strong congestion-supportive behaviour observed in the main simulation. The effectiveness of directive control therefore depends not only on the ability to issue dispatch instructions, but also on the ability to issue them at the right time.

Second, directive control is associated with the largest operational burden among the evaluated contractual arrangements. CSC produces the highest DAFO costs because instructed dispatch can conflict with previously committed market positions. Although activation payments offset part of this burden, they do not eliminate it. Consequently, improved targeting enhances financial performance in some scenario families, but does not make the contract uniformly bankable. Under the evaluated compensation parameters, improved targeting allows one additional scenario family to exceed the 8.5% hurdle rate, while the remaining scenarios remain below the investment threshold.

Taken together, these findings show that directive control delivers the strongest congestion-alignment performance among the evaluated design properties, while simultaneously imposing the greatest operational burden on the asset owner. Within the evaluated contract set, directive control therefore exhibits the clearest trade-off between system alignment and investment performance.

4.3.2. Integrated Property Synthesis

None of the evaluated design properties is associated with favourable outcomes across all three evaluation dimensions simultaneously. The preceding analyses reveal consistent patterns of association across both congestion regimes and throughout the Monte Carlo distribution.

Financial viability and congestion alignment remain separated across the evaluated contract set. CBC achieves universal bankability with negligible alignment improvement, whereas CSC achieves the strongest alignment improvement while remaining non-bankable in the urban regime under the evaluated calibration. No evaluated contractual arrangement simultaneously achieves both outcomes.

Firmness is associated with reduced operational risk but is not associated with improved financial viability in isolation. Firmness Gradient exhibits the lowest UOR in the evaluated contract set while remaining non-bankable under the evaluated calibration.

Volume Bound's energy-budget design is associated with reduced degradation and improved congestion alignment in the urban regime, but reduced alignment in the rural regime. The same property therefore exhibits different alignment outcomes across the two congestion archetypes, illustrating that the observed association is regime-dependent.

Table 4.18 summarises the observed associations between each property dimension and the three evaluation outcomes.

Table 4.18: Summary of observed property-outcome associations. Entries reflect the direction and consistency of associations observed across the evaluated contracts. All associations are observational: no two contracts in this study differ on a single property dimension in isolation. ↑ = positive association; ↓ = negative association; ∼ = no consistent association; ± = regime-dependent association.

Property dimension	Financial viability	Congestion alignment	Technical (degradation)
Compensation structure (high coverage ratio)	↑ assoc. with bankability	∼ no consistent association	∼ no consistent association
Firmness (guaranteed floor)	∼ assoc. with reduced operational risk	∼ no consistent association	∼ no consistent association
Volume orientation (energy budget)	↓ assoc. with lower financial performance	± positive urban; negative rural	↑ assoc. with reduced degradation rate
Directive control (DSO dispatch authority)	↓ assoc. with lower financial performance	↑ assoc. with largest alignment improvement	∼ no consistent association

Table 4.18 summarises this partitioning of the outcome space; its implications for institutional design are developed in Chapter 5.

4.3.3. Institutional Feedback Mapping

This section maps the contractual assessment and property-based evaluation findings onto the three structural tensions identified during the institutional discovery, reporting the relevant quantitative evidence for each tension. No assessment of whether this evidence confirms, contradicts, or resolves the underlying institutional concern is offered here; that interpretation is developed in Chapter 5. No causal claims are made: the Policy Testbed evaluates performance associations, not institutional mechanisms.

Tension 1: Asymmetric Risk and Value Allocation

The institutional discovery identified a recurring stakeholder concern that flexible connection agreements impose operational risks on battery investors without providing compensation commensurate with the resulting operational burden.

The UOR-remuneration scatter (Figure 4.4) shows that none of the evaluated arrangements lies close to the full UOR-matching line. Instead, the contracts span the full compensation spectrum, ranging from no compensation (Volume Bound), through partial compensation (TBTR), to remuneration substantially exceeding measured operational burden (ATR85, CSC and particularly CBC).

ATR85 illustrates a different aspect of the same tension. Although its mean remuneration exceeds mean UOR, the upper tail of the UOR distribution approaches, and in some cases exceeds, the fixed annual compensation. Fixed remuneration may therefore provide adequate average compensation while still exposing investors to substantial downside risk under adverse market conditions.

Two observations qualify these findings. First, the evaluated arrangements support the existence of a structural mismatch between operational burden and financial remuneration, but the direction of that mismatch is not uniform. Rather than consistently undercompensating investors, the results show that both the magnitude and direction of the mismatch depend strongly on contractual design.

Second, UOR captures only the direct market value displaced by contractual restrictions. Secondary costs, such as DAFO charges under directive control, and broader strategic benefits associated with access certainty fall outside this metric. The reported relationships should therefore be interpreted as describing proportionality between remuneration and measured operational burden, rather than the overall financial attractiveness of a contractual arrangement.

Tension 2: Temporal misalignment

The institutional discovery identified a recurring concern that contractual restrictions interfere with the intertemporal operating logic of battery storage by constraining dispatch during periods of highest market value rather than periods of low operational significance.

The timing-premium analysis (Table 4.11) indicates that this concern is not universal across the evaluated contractual arrangements. Instead, the financial consequences of activation timing vary considerably with contractual design. CSC and ATR85 both exhibit positive timing premia across the evaluated congestion regimes, indicating that congestion-aligned activation generally coincides with, rather than displaces, commercially valuable dispatch. CBC, by contrast, exhibits only a negligible timing premium, suggesting that its financial performance is largely insensitive to activation timing within the evaluated operating range.

The clearest evidence of temporal misalignment is observed for TBTR, though not through the timing-premium metric used for the other contracts, which cannot be applied to a fixed restriction schedule. TBTR's restriction windows (07:00–09:00 and 17:00–19:00) are determined purely by calendar time and do not adapt to congestion conditions, market prices, or scenario realisation, in contrast to CSC, whose activation timing responds directly to congestion signals. A direct comparison of realised market values during these fixed windows against the annual mean, computed across the same stratified scenario sample used for the timing-premium analysis, confirms this is systematically unfavourable rather than neutral: mean Day-Ahead prices during TBTR's restricted hours exceed the annual mean by 18.5% (all runs), and imbalance-up prices by 14.1%. Volume Bound illustrates a different form of temporal tension: rather than fixing restriction windows in time, its daily energy budget changes how throughput is allocated across the day, reducing congestion alignment under the rural archetype for the reasons set out in Section 4.3.1.

Taken together, these findings suggest that temporal misalignment is not an inherent characteristic of flexible connection agreements, but depends strongly on contractual design. Within the evaluated contract set, responsive arrangements generally preserve the battery's intertemporal operating logic more effectively than deterministic restriction schedules.

Tension 3: Control Rights and Responsibility

The institutional discovery identified a recurring concern that network operators may obtain greater influence over battery dispatch while financial responsibility for the resulting operational consequences remains with the asset owner.

The quantitative results support the presence of this tension under CSC, the only evaluated contractual arrangement that incorporates directive control. CSC delivers the strongest congestion alignment of all evaluated arrangements, but it is also the only contract that systematically generates substantial DAFO costs. These costs arise because DSO dispatch instructions can conflict with market positions that were committed earlier, leaving the asset owner responsible for the resulting settlement costs.

This imbalance between operational authority and financial responsibility is also reflected in the burden-adjusted compensation analysis. Once secondary DAFO costs are incorporated, the effective compensation coverage is reduced in both congestion regimes, indicating that the contractual remuneration only partially offsets the additional operational burden associated with directive control. The combination of strong congestion-alignment performance and limited investment viability therefore places operational benefits and financial consequences on different actors.

The timing-quality sensitivity further qualifies this relationship. The benefits of directive control depend strongly on the quality of DSO targeting. Under poorly targeted activation, CSC provides little improvement over the natural congestion alignment achieved by unconstrained market operation, whereas progressively more accurate targeting steadily strengthens alignment. The effectiveness of directive control therefore depends on transferring operational authority, and on the precision with which that authority is exercised.

Summary of Institutional Evidence

Table 4.19 summarises the simulation evidence associated with each institutional tension identified during the institutional discovery. The table reports the principal quantitative observations from the Policy Testbed without interpreting whether they confirm, contradict, or qualify the institutional concerns. That synthesis is developed in Chapter 5.

Table 4.19: Summary of simulation evidence associated with the institutional tensions identified in the institutional discovery. Entries report direct observations from the Policy Testbed without interpretive judgement.

Structural tension	Simulation evidence
Asymmetric risk and value allocation	Coverage ratios span the full spectrum, from no remuneration (Volume Bound and Firmness Gradient) to substantial over-compensation (CBC). UOR-remuneration mismatches are observed across all evaluated arrangements, with the CBC UOR-matching rate approximately 19× below the ESNL benchmark.
Temporal misalignment	Positive timing premia are observed for ATR85 and CSC across both regimes, with negligible timing effects for CBC and Firmness Gradient. Volume Bound exhibits a rural alignment reversal. TBTR could not be evaluated on this metric, since its fixed restriction window does not respond to the activation-timing parameter used to construct the comparison; a direct comparison instead shows its fixed hours coincide with Day-Ahead prices averaging 18.5% above the annual mean, consistent across all 16 tested scenarios.
Control rights and responsibility	CSC delivers the strongest congestion-alignment improvement while simultaneously generating the highest DAFO costs. The realised alignment benefit depends strongly on DSO targeting quality.

The implications of this evidence for the institutional design of flexible connection agreements are discussed in Chapter 5.

5

Discussion

This chapter does not re-rank the contractual arrangements evaluated in Chapter 4. The performance patterns established there require no revision. Instead, this chapter interprets what those findings mean, considers the mechanisms that may explain the observed patterns, and discusses their implications for the institutional design of flexible connection agreements in the Netherlands. Key figures are restated only where necessary to support new interpretation; readers are assumed to be familiar with the full evidence base presented in Chapter 4.

The discussion is organised into eight sections. Section 5.1 revisits the three structural tensions identified in the institutional analysis in light of the simulation results. Section 5.2 discusses the implications of the property-based evaluation. Section 5.3 positions the findings within the existing literature. Section 5.4 considers their implications for policy and contract design. Section 5.5 reflects on the study's limitations and methodological choices. Section 5.6 synthesises the preceding discussion into an overall conclusion. Section 5.7 summarises the scientific contribution of the study. Section 5.8 then outlines directions for future research and institutional development.

Throughout this chapter, the observed relationships are interpreted as associations rather than causal effects, and the distinction between model results and institutional inference is maintained explicitly.

5.1. Revisiting the Institutional Tensions

A structural cost floor that contract design cannot cross

The most fundamental finding of this study is not about any individual contractual arrangement. It concerns the institutional context within which all evaluated arrangements operate. Even under the most favourable market conditions considered, the unconstrained Baseline does not constitute a bankable investment case when evaluated against the 8.5% hurdle rate identified through expert consultation [39], and viability deteriorates further under less favourable conditions. This indicates that the principal barrier to bankability precedes contractual design itself.

The mechanism appears to be structural rather than operational. Within the evaluated setting, a substantial share of gross market revenue is absorbed by the Dutch high-voltage transport tariff before any contractual restriction is applied, and no evaluated arrangement is able to overcome a cost burden of this magnitude through operational optimisation alone. This should be understood as an institutional observation rather than a modelling critique: battery storage utilises transport capacity intermittently, only when market opportunity and state-of-charge conditions align, whereas the transport tariff is largely independent of actual network utilisation. This mismatch between infrastructure use and cost recovery is consistent with the largest friction cluster identified in the institutional analysis — the allocation of costs, risks, and responsibilities within the existing access architecture — and it is against this structural floor that the contractual results below should be interpreted: compensation sufficient to restore investment viability must offset not only the operational burden of restriction, but this underlying cost disadvantage as well.

Asymmetric risk and value allocation: confirmed in structure, reversed in direction

The institutional analysis identified a recurring stakeholder concern that flexible connection agreements transfer operational risk to asset owners without providing commensurate compensation. The simulation results are consistent with the existence of this structural mismatch, but not with the direction anticipated by stakeholders. Rather than systematic undercompensation, three of the four remunerated arrangements provide compensation exceeding demonstrated UOR, while only one fails to cover the burden it imposes in expectation.

The clearest example is CBC. Bankability across all scenario families is achieved not through operational performance, but through remuneration that substantially exceeds measured operational burden. As shown in Section 4.3.1, the contract's coverage ratio exceeds 1,500% in both congestion regimes, corresponding to compensation of approximately fifteen to sixteen times demonstrated UOR. What explains this scale of overcompensation is the mechanism examined next.

This finding is consistent with the theoretical concern raised by Hirth [33] that capacity-based remuneration can become decoupled from the market value of the flexibility service provided. This interpretation is also consistent with the design principles proposed by ACER [1], who identify proportionality and constraint cost recovery as central principles for flexible connection agreement design. The CBC results indicate that the current Dutch compensation calibration, derived from historical voluntary redispatch prices, is not anchored to the demonstrated operational cost of capacity limitation in a co-optimised multi-market battery context. Within the evaluated framework, investment viability is restored primarily through financial transfer rather than through compensation proportional to the demonstrated operational burden imposed by the contractual restriction.

Across the evaluated scenarios, the 3 MW power ceiling imposed by CBC rarely binds on profitable dispatch decisions. This is reflected in CBC's consistently low UOR across both congestion regimes, indicating that the contractual ceiling displaces only a limited amount of economically preferred dispatch. Because the battery co-optimises across Day-Ahead arbitrage, aFRR reservation, and imbalance participation simultaneously, the optimal dispatch solution already allocates capacity across multiple market segments. The ceiling therefore rarely alters the battery's preferred dispatch strategy under the evaluated conditions.

Remuneration, however, accrues for every congestion-flagged PTU regardless of whether the ceiling altered dispatch. The coverage ratio, which exceeds 1,500% in both regimes, is therefore not the result of generous calibration applied to a genuinely constrained asset. Rather, it reflects compensation calculated against rated capacity curtailment — capacity that, under the evaluated conditions, the battery would frequently not have used anyway under unconstrained operation. CBC's bankability should therefore not be interpreted as evidence that the compensation mechanism is proportionate to the operational burden imposed. Instead, it demonstrates that a sufficiently large unconditional transfer can bridge the underlying tariff gap identified above, rather than compensating the operational restriction itself.

This is precisely what makes the policy significance double-edged: this overcompensation would, in practice, be funded through grid costs socialised across network users, while bearing little relationship to the congestion relief delivered. The sensitivity analysis identifies a compensation range that could retain investor attractiveness while reducing, though not eliminating, the structural overcompensation.

Fixed availability payments, as used by ATR85, nominally exceed full UOR coverage in expectation yet expose investors to undercompensation during adverse market realisations. As established in Section 4.3.1, ATR85 exhibits the widest distribution of operational risk of all evaluated contracts, meaning the fixed transfer is insensitive to when restrictions coincide with market value, so mean coverage conceals tail exposure that erodes the surplus in adverse realisations. This mechanism explains why ATR85 achieves no bankable scenarios despite a coverage ratio exceeding unity.

Under CSC, the UOR metric overstates effective coverage because it does not capture DAFO costs generated by compliant forced dispatch. These costs are material, and the activation payment structure provides no direct compensation mechanism for them. Even after incorporating these charges, coverage remains above the compensation line (Section 4.2.2). This demonstrates that even nominally remunerated contracts can materially understate the actual financial friction imposed by directive grid management. The asymmetric risk and value allocation tension is therefore confirmed at the structural level: compensation remains decoupled from demonstrated operational burden, although its manifestation is more nuanced than the direction originally described by stakeholders.

Temporal misalignment: contract-design conditional, not universal

The temporal misalignment tension receives mixed support across the evaluated contractual arrangements. It is most evident under fixed-window restrictions, is not supported for directive control in the evaluated urban regime, and is reversed under capacity-limiting and energy-budget designs through mechanisms that differ by contract.

Under CSC, directive instructions are CMI-weighted by design, so activation timing tracks the congestion profile rather than market value. In the urban regime, where forced dispatch is predominantly discharge (70% of activations, Table 3.5) during the evening demand peak, this timing premium is strongly positive (Table 4.11). Under the evaluated urban conditions, CMI-aligned activation therefore coincides with the battery's preferred discharge behaviour rather than displacing it. The urban results consequently do not support the temporal misalignment concern for directive control.

In the rural regime, where CSC is predominantly a forced-charging instrument (90% of activations, Table 3.5) targeting the midday solar-driven congestion window, the paired timing analysis again yields a strongly positive premium (Table 4.11). This indicates that, as in the urban regime, CMI-aligned activation reinforces rather than displaces the battery's commercially preferred dispatch behaviour under rural operating conditions, further supporting the conclusion that the temporal misalignment concern does not hold uniformly for directive control across congestion archetypes.

Volume Bound reverses the temporal misalignment tension through a different mechanism. Whereas CBC constrains instantaneous power independently in each interval, meaning a restriction in one PTU has no bearing on any other, Volume Bound pools all charging and discharging within a single daily energy budget, so a restriction in one part of the day changes what is economically available in another. See Section 5.2 for the resulting mechanism in the rural regime. The resulting misalignment therefore arises not from restriction timing itself, but from the interaction between a pooled energy constraint and the directional relationship between congestion timing and market value.

TBTR exhibits the clearest form of temporal misalignment in the evaluated contract set. Its restriction window is predetermined and applied irrespective of realised congestion or market conditions, removing operational flexibility on every restricted day, in contrast to CSC's congestion-responsive timing or Volume Bound's daily allocation discretion. The paired timing-premium methodology cannot be applied to this fixed schedule; instead, a direct comparison of realised market values during TBTR's windows against the annual mean (Section 4.3.3) confirms this is systematically unfavourable rather than neutral. Since both windows fall within the peaks that structurally drive Dutch electricity prices, TBTR's fixed schedule does not merely fail to adapt to market conditions; it systematically coincides with the hours of highest commercial value.

ATR85 and CBC do not exhibit the same pattern, producing neutral or positive timing premia. ATR85's rural timing premium is in fact the largest of any evaluated contract in that regime, exceeding even CSC's directively-controlled premium; this indicates that avoiding curtailment during the wrong hours can matter as much for a probabilistic, non-directive arrangement as for one in which the DSO actively times dispatch. Neither ATR85 nor CBC pools intertemporal decisions within a shared energy budget nor imposes deterministic restriction windows. This contrast across contracts suggests that temporal misalignment is not an inherent consequence of operational restrictions themselves, but depends on the specific contractual mechanisms through which those restrictions are implemented.

Control rights versus responsibility: confirmed through the mechanism the framework predicted

The control rights versus responsibility tension is directly instantiated by CSC and is strongly supported by the quantitative results. The most important finding is not the magnitude of the alignment improvement, already the largest observed among the evaluated contractual arrangements, but the conditions under which that improvement materialises. The timing-quality sensitivity analysis shows that the benefits of directive control depend on congestion-aligned DSO activation timing. As established in Section 5.2, this benefit is conditional on targeting quality: absent it, CSC delivers no improvement over natural market alignment. High targeting precision therefore appears to be a necessary condition for directive control to deliver meaningful system value.

Poor DSO targeting also imposes an investor cost. In the urban regime, randomly timed activation yields both lower profit and lower IRR than precisely targeted activation. Poor targeting is therefore costly on two dimensions simultaneously, reducing congestion alignment while also reducing investor performance, rather than presenting a trade-off between system and commercial objectives. This finding is broadly consistent with the emphasis placed by Joskow [35] and Hirth and Glismann [34] on forecasting quality and effective real-time system operation. It also demonstrates that the effectiveness of directive battery dispatch depends critically on activation timing.

The financial burden associated with directive control originates from a structural mismatch between real-time system operation and the sequential organisation of Dutch electricity markets. Once the Day-Ahead gate closes, market positions are fixed. A subsequent DSO instruction therefore forces a deviation from the committed position, with the resulting imbalance settled through the regulation-state mechanism. The DAFO costs observed under CSC are consequently not an incidental modelling artefact but a predictable outcome of combining real-time dispatch authority with forward market commitments. This institutional mismatch originates in market architecture rather than in the contractual penalty itself.

Penalty calibration nevertheless influences how this institutional tension is expressed. The default penalty level (€400/MWh), corresponding to the value proposed in the ACM consultation, functions both as a financial parameter and as a compliance mechanism. As established in Section 4.2.2, this trade-off is not incidental to the calibration chosen but reflects the underlying mechanism directly: lower penalties reduce the financial consequences of directive control only by reducing the level of compliance on which its system benefits depend.

The evidence therefore suggests that the control rights versus responsibility tension cannot be resolved through penalty calibration alone. Although lower penalties improve investor outcomes, they simultaneously erode the congestion-alignment benefits that justify directive control. A durable institutional solution therefore requires mechanisms that address the secondary settlement costs created by real-time interventions, alongside DSO scheduling practices capable of aligning system needs with investor opportunity costs.

5.2. What Matters More: Contracts or Properties?

Having revisited each institutional tension in turn, it is worth stepping back to ask a question that cuts across all three: does this property-based way of looking at contract design actually add value over comparing contracts directly? The central methodological claim of this thesis is that moving from contract comparison to property comparison produces more generalisable design insights. The evidence supports this claim.

Section 4.3.2 established that each of the four design properties maps onto a distinct evaluation outcome, with limited overlap between them. No single property simultaneously improves all three evaluation dimensions. This partitioning of the outcome space is the most important finding of the property-based analysis: it implies that contract design choices are inherently trade-offs across dimensions rather than optimisation problems within a single one.

Compensation structure

The key lesson is not that higher compensation rates produce better financial outcomes. It is that the mechanism linking payment to burden determines whether the risk-allocation tension is addressed or merely concealed. Fixed proportional transfers are structurally unable to track operational harm: they provide the same financial amount whether restrictions fall on the most or the least valuable hours of the operating year. Event-triggered capacity payments couple payment to restriction frequency but not to harm magnitude. Activation-based energy payments are the most closely linked to realised events but systematically exclude secondary settlement consequences. A compensation mechanism that would address the risk-allocation tension identified in the institutional analysis would need to link transfers more closely to demonstrated operational burden, covering both direct revenue displacement and secondary consequences simultaneously. No evaluated contract meets this dual condition. The design requirement identified here is therefore not a calibration target but a mechanism requirement: the form of the compensation link determines what information is embedded in the payment.

Firmness

The Firmness Gradient result is the study's most analytically isolated property observation. By carrying no compensation and imposing only a structural transport floor, the contract separates the access guarantee from any financial transfer. Its near-zero UOR confirms that a firm floor at a moderate fraction of rated capacity eliminates nearly all measurable operational risk displacement at modelled congestion frequencies, not by compensating for adverse outcomes but by preventing the high-curtailment scenarios in which large displacement would materialise. This effect is structurally distinct from, and potentially more reliable than, compensating for unbounded exposure after the fact. The contrast with ATR85 further illustrates this point. Whereas ATR85 reduces investor risk primarily through ex post financial compensation while leaving operational exposure effectively unbounded, the Firmness Gradient limits downside risk structurally by guaranteeing a minimum level of transport access, consistent with bounding exposure through guaranteed access being more effective than compensating unrestricted exposure after the fact. As established in Section 4.3.1, however, the Firmness Gradient remains non-bankable: firmness alone reduces operational risk but does not generate return, and bounded risk is therefore insufficient for investment viability on its own.

Volume orientation

The Volume Bound alignment reversal in the rural archetype is the most institutionally significant finding for this property. The energy budget, motivated by the temporal misalignment tension and designed to preserve the battery's intertemporal operating logic, amplifies rather than mitigates misalignment in generation-driven networks. As established in Section 4.3.1, the daily energy budget is reallocated toward higher-value evening discharge at the expense of the midday charging that relieves rural congestion. The directional consistency of this reversal across the full Monte Carlo distribution reflects the interaction between the budget constraint and the stable structural features of the rural congestion archetype. Unlike market prices, which vary substantially across scenario families, the timing of solar-peak congestion is a physical rather than an economic phenomenon. The rural CMI profile concentrates congestion weights during midday hours throughout the year. The stability of this pattern means that the optimiser repeatedly faces the same fundamental trade-off between charging during midday congestion windows and preserving the daily energy budget for higher-value evening discharge. Market uncertainty changes the magnitude of that trade-off but not its direction, which is why the rural reversal persists regardless of the scenario family drawn. This echoes the socio-technical regime transition logic described by Geels [30]: a design intended for one congestion topology generates adverse effects in a structurally different environment. The evidence from the Volume Bound contract suggests that energy-budget orientation may not be a universally portable design property. Its value therefore appears strongly location-conditional, depending on whether the market value distribution and the dominant congestion direction are mutually compatible at the deployment location.

As established in Section 4.2.2, the degradation benefit for Volume Bound arises from the same throughput constraint that generates the profit penalty. Crucially, the degradation benefit is a contingent consequence of the experimental design rather than an inherent property of energy-budget orientation, since the same throughput constraint that suppresses cycling would be expected to dissolve under remuneration levels sufficient to restore the suppressed dispatch. A scope boundary developed further in Section 5.5. The near-universal degradation equivalence across the other five contracts reflects the internalization mechanism described by Massat et al. [40]: when wear costs are embedded in the dispatch objective, the battery absorbs most contractual restrictions as modified feasibility boundaries without measurable change in physical ageing

Directive control

Among the evaluated design dimensions, directive control is the only one that consistently produces alignment improvements of policy-relevant magnitude. The gap between CSC and every other evaluated arrangement is not marginal (Table 4.9) and is robust across the full Monte Carlo distribution and both congestion archetypes. These results indicate that if meaningful congestion-alignment improvement is the primary regulatory objective, passive mechanisms alone are insufficient.

The institutional cost of this performance level, however, is equally consistent with the control-rights/responsibility tension discussed above: directive control as currently structured transfers dispatch authority while leaving the associated DAFO exposure with the asset owner, with no mechanism to redistribute those secondary settlement consequences.

This timing-quality dependency, established above, carries a direct governance implication: the case for granting directive authority rests on the DSO's ability to exercise it with sufficiently accurate activation scheduling. Where that capability cannot be achieved, the institutional trade-off is not between alignment gain and financial cost, but between financial cost and little or no alignment gain. In this context, investments in DSO forecasting capability and activation scheduling quality may be a precondition for directive control to deliver its intended system benefit, rather than a marginal improvement to an already functioning instrument.

A further qualification concerns the institutional context in which directive control is applied. The rural Baseline already exhibits strong natural alignment ($\Omega = -0.150$, more than three times the urban value), because commercial price signals happen to coincide with congestion conditions in that regime. This contrasts with the urban regime, where unconstrained market dispatch is substantially less congestion-supportive. This result suggests that the baseline governance challenge differs across congestion regimes: generation-driven networks benefit from a greater degree of naturally occurring market support, whereas demand-driven networks exhibit a larger gap between commercial optimisation and system objectives. While CSC substantially improves alignment in both regimes, delivering a 112% increase in the rural archetype, it does so from structurally different starting points. The scale of intervention required for system alignment is therefore regime-conditional, with market incentives partially substituting for contractual intervention in rural generation-heavy topologies.

These property-level findings can now be situated within the wider literature on battery storage economics, congestion management, and institutional design.

5.3. Positioning within Existing Literature

The Baseline revenue composition of this study confirms the qualitative ordering established by Verhagen, Hu, and Harmsen [64] and Dedousi [20]: reserve capacity markets dominate over energy arbitrage in Dutch utility-scale battery economics, with imbalance participation occupying an intermediate position. The Baseline asset lifetime is consistent with Schouten [52]'s estimates for grid-scale LFP systems under realistic Dutch operating conditions and reflects the calibration grounded in Xu et al. [67]. That this lifetime is achieved without aggressive cycling, because reserve capacity provides income through availability rather than throughput, is consistent with Verhagen, Hu, and Harmsen [64]'s observation that aFRR reservation delivers high revenue density at relatively modest cycling intensity, and is reinforced by the internalisation mechanism of Massat et al. [40].

Where this study extends rather than confirms the existing literature is in three specific directions.

First, it identifies the transport tariff, not market revenue compression, as the binding institutional constraint on battery investment viability. Verhagen [63] and Dedousi [20] document the revenue-side sensitivity to saturation-driven spread compression; this study adds that the dominant cost constraint on viability, established in Section 5.1, is structural rather than revenue-side. This shifts the policy diagnosis: the required reform may lie upstream of any contractual instrument.

Second, it operationalises the proportionality principle of ACER [1] through the UOR metric and the coverage ratio, translating a governance principle into a quantifiable property evaluable consistently across alternative designs. The finding that no evaluated contract achieves actuarially fair compensation from either direction provides direct empirical support for Hirth and Glismann [34]'s theoretical prediction that calibration to rated capacity rather than realised operational cost produces systematic misalignment in both directions. The GOPACS price of approximately €93.5/MWh forms the basis for current CBC calibration. But this price reflects scarcity conditions in a voluntary redispatch market, a market with no institutional link to the operational cost of capacity limitation that this study demonstrates for a co-optimised, multi-market battery. The gap between the GOPACS-anchored calibration rate and the demonstrated UOR-matching rate is therefore substantial, exceeding a factor of 14 in both congestion regimes, a gap that widens further, to approximately 19, when benchmarked instead against the ESNL working group's proposed €120/MWh target (Section 4.2.2). Third, it provides empirical evidence that natural market alignment is substantial and regime-dependent, and that passive mechanisms contribute only marginally beyond this natural baseline. Heemsbergen [32] identifies natural coincidence between commercial battery dispatch and system needs; this study quantifies it across urban and rural archetypes and shows that the Baseline battery is idle during the large majority of congestion-weighted intervals, contributing to alignment incidentally through market-driven inactivity during misaligned periods. The finding that CSC's alignment improvement depends on congestion-aligned DSO activation timing provides empirical evidence that the effectiveness of directive battery dispatch is closely linked to the quality of DSO scheduling, reinforcing the importance of forecasting accuracy and effective real-time control emphasised by Hirth and Glismann [34], while remaining consistent with Heemsbergen [32]'s spatial analysis of coincidence patterns across Dutch network topologies.

More broadly, the exploratory sequential research design demonstrates a structured approach for evaluating technical instruments whose performance depends fundamentally on institutional context. By deriving quantitative evaluation criteria from qualitative institutional analysis, implementing them within a modular Policy Testbed, and interpreting the resulting outcomes through the institutional tensions from which those criteria originated, the research design establishes a systematic link between regulatory analysis and quantitative energy systems modelling. In doing so, it provides one possible operationalisation of the socio-technical perspective advanced by Geels [30] while addressing the institutional misalignment challenges highlighted by Proka et al. [51]. Rather than evaluating technical performance in isolation, it embeds quantitative contract evaluation within the governance structures that shape operational behaviour, illustrating how institutional analysis and energy systems modelling can be integrated into a coherent policy evaluation methodology.

5.4. Implications for Flexible Connection Agreement Design

Three governance implications emerge from the analysis. Each is grounded in the evidence and framed as a direction suggested by the findings rather than as a regulatory prescription.

Anchor compensation calibration to demonstrated operational burden. Current institutional practice calibrates compensation to institutional risk, the possibility of receiving a connection with reduced transport rights, rather than to the operational harm that restriction actually imposes. The result is compensation that is simultaneously excessive in expectation and potentially inadequate under adverse realisations. A regulatory framework that bases calibration on demonstrated operational burden, for example, through ex-post UOR reporting by licensed battery operators, gradually anchoring the rate to observed displacement, could reduce this misalignment without requiring immediate tariff reform. Whether this is achievable within the existing ACM mandate is a governance question the model motivates but cannot resolve. What the evidence indicates is the scale of the current gap and the direction of correction.

Reform the compensation structure of directive agreements, not only their penalty rate. The primary financial burden of CSC does not arise from the penalty applied to non-compliant deviations but from programme-imbalance charges generated by compliant forced dispatch conflicting with pre-committed market positions. Regulatory attention focused exclusively on the penalty rate addresses compliance incentives without resolving the cost structure that makes compliance financially damaging. A compensation framework that covers secondary settlement consequences alongside primary activation events would structurally change the trade-off. This reform is made more urgent by the timing-quality sensitivity findings presented in Section 5.1: because randomly timed DSO activations impose financial costs without delivering meaningful alignment improvements, the institutional case for directive control depends on pairing improved compensation with demonstrated scheduling quality rather than on penalty calibration alone.

Treat firmness and location specificity as design requirements, not contract features. The firmness findings suggest that guaranteed minimum access rights may provide more effective risk reduction per euro of remuneration than proportionally higher compensation payments calibrated to the same risk level. A contract combining a structural floor with evidence-anchored compensation could address both downside-exposure and viability concerns simultaneously, without the scale of overcompensation currently required for bankability. The Volume Bound rural alignment reversal simultaneously demonstrates that contracts designed for one congestion topology may be counterproductive in another. Both observations point toward a design approach in which location-specific parameters are embedded within standardised contractual structures, rather than a single national template applied uniformly across structurally differentiated network environments.

5.5. Limitations

Several limitations should be considered when interpreting the findings discussed in this chapter. These limitations relate to specific analyses and conclusions presented here. Broader methodological limitations of the research design are discussed separately in Chapter 6.

The quantitative magnitude of the reported CBC remuneration-to-UOR ratio is conditional on the evaluated compensation calibration and the assumed 3 MW capacity ceiling. While the precise ratio would change under alternative calibrations, the broader finding that remuneration substantially exceeds measured UOR under the evaluated calibration remains unchanged.

A related limitation concerns the comparability of the CSC timing-quality sensitivity analysis (Table 4.16) with the main simulation results. The timing-quality sweep, based on sixteen freshly re-solved stratified simulations, reports Scenario 1 IRR exceeding the hurdle rate even under random activation timing, whereas the main simulation reports a lower Scenario 1 IRR and no bankable urban CSC scenarios overall. The same pattern is present, and somewhat larger in magnitude, in the rural regime (Table 4.17, Table 4.7). This discrepancy reflects two compounding factors: the smaller stratified sample underlying the sensitivity analysis and CSC's activation timing in the original simulation dataset not yet being fully responsive to the assigned timing-bias parameter when that dataset was generated. Consequently, the timing-quality sensitivity analysis should be regarded as the more reliable representation of CSC behaviour under controlled timing conditions, while the main simulation's CSC bankability results should be interpreted with corresponding caution.

A specific limitation of the Policy Testbed is the use of a fixed ± 5 MW setpoint for CSC activations. In operational practice, a DSO would optimise the instruction level based on physical grid state, whereas the model treats every activation as a uniform 50% rated power event. This simplification likely overstates the alignment potential of directive control in scenarios where only marginal congestion relief is required, while potentially understating the financial DAFO burden if real-world instructions frequently required full-rated (10 MW) responses. Consequently, the reported CSC results should be interpreted as representative of a stylised activation policy rather than a fully adaptive operational strategy. Future research could treat the setpoint magnitude as a dynamic variable linked to physical nodal sensitivity.

Furthermore, the property-based evaluation contains a structural confound between design properties and compensation. While the existing Dutch contracts (ATR85, TBTR, CBC, CSC) all retain their institutional remuneration, the two prospective designs (Firmness Gradient and Volume Bound) were tested with zero remuneration by deliberate experimental design. Consequently, the finding that 'firmness and volume orientation are not associated with bankability' is, to an extent, true by construction: any uncompensated contract paying the full transport tariff identified in Section 5.1 faces a structural financial hurdle that cannot be overcome by operational design properties alone. The analysis identifies the operational potential of these properties, while acknowledging that financial viability would require them to be paired with a separate compensation mechanism.

The prospective contract calibrations, the energy budget level and the access floor, are modelling choices, not empirically grounded parameters; alternative values would change quantitative outcomes while leaving the property-level observations intact. The degradation benefit observed under the Volume Bound arrangement is similarly conditional on the adopted experimental design. It arises from the combination of a throughput restriction and the absence of compensatory remuneration, which together reduce economically attractive cycling. Consequently, this lifetime extension should not be interpreted as an inherent property of energy-budget contracts, but as an outcome of the evaluated contract calibration. Alternative remuneration structures that restore the economic incentive to utilise the available energy budget could reduce or eliminate the observed degradation benefit.

Taken together, these limitations affect the quantitative magnitude of the CBC ratio, the CSC timing-quality and setpoint results, and the Volume Bound degradation finding, rather than the qualitative relationships identified throughout this chapter. The principal conclusions concerning the associations between contractual design properties, institutional tensions, and investment performance remain robust within the stated modelling assumptions.

5.6. Conclusions

This section synthesises the preceding discussion into direct answers to the four sub-research questions, before addressing the central research question in Section 5.6.2.

5.6.1. Answers to the Research Questions

SQ1: Institutional Discovery. The institutional analysis identified 112 friction fragments across 33 regulatory and stakeholder documents, distributed unevenly across five thematic clusters. These clusters were synthesised into three structural tensions: temporal misalignment, asymmetric allocation of risk and value, and a disconnect between control rights and operational responsibility. The dominance of institutional allocation and design gaps across every stakeholder group, together with the complete absence of temporal concerns among regulators, suggests that these are not isolated implementation issues but symptoms of a deeper mismatch between the intertemporal operating logic of battery storage and a governance architecture that was not designed with that logic in mind. These three tensions subsequently defined the property dimensions, compensation structure, firmness, volume orientation, and directive control, evaluated throughout the Policy Testbed.

SQ2: Baseline Model. An unconstrained battery co-optimising across the Day-Ahead, aFRR, and imbalance markets is structurally non-bankable under current Dutch conditions. Best-case IRR reaches only 3.0–3.8%, well below the 8.5% hurdle rate. The principal cause is institutional rather than technical: the Dutch high-voltage transport tariff absorbs approximately 70% of gross market revenue, independent of how intermittently the asset actually exercises its connection.

Physically, the unconstrained battery reaches an implied lifetime of approximately 12.9 years before crossing the 80% State-of-Health End-of-Life threshold, modestly short of the 15-year investment horizon and essentially identical across the urban and rural congestion archetypes, reflecting a degradation pathway shaped primarily by aFRR's availability-based revenue model rather than by intensive cycling.

In terms of system alignment, profit-maximising unconstrained dispatch already exhibits meaningful congestion-supportive behaviour without any contractual intervention, with a mean alignment score of $\Omega = -0.045$ in the urban regime and $\Omega = -0.150$ in the rural regime. This natural alignment arises largely from strategic inactivity rather than directed dispatch: the battery remains idle during 84.9% (urban) and 87.1% (rural) of congestion-weighted intervals. The alignment is therefore incidental rather than structural: it does not generalise, scale, or persist under different market or network conditions. The institutional design task is not to create alignment where none exists, but to convert this fragile overlap into something regulators can rely on.

Together, these findings establish the financial, technical, and system-alignment reference case against which all evaluated contractual arrangements are assessed. They further demonstrate that every contractual arrangement is evaluated not against a neutral baseline, but against a baseline that is already financially fragile while simultaneously providing a meaningful degree of congestion-supportive behaviour under unconstrained market operation.

SQ3: Policy Testbed. The six evaluated contractual arrangements produced markedly different operational, financial, and system-level outcomes. CBC is the only arrangement achieving bankability across all scenario families, but its coverage ratio exceeds 1,500% of demonstrated UOR, indicating that its investment viability is associated with the scale of the financial transfer rather than with proportionate compensation for operational burden. CSC achieves the largest improvement in congestion alignment within the evaluated contract set, but this improvement depends on congestion-aligned DSO activation timing. Under random activation timing, the contract provides little or no improvement in alignment relative to unconstrained market dispatch while maintaining comparable financial costs. Degradation outcomes are largely contract-invariant, with five of the six arrangements changing implied battery lifetime by less than four weeks relative to the Baseline. Volume Bound is the sole exception, extending implied battery lifetime by approximately two months. This benefit is accompanied by lower profitability and a regime-dependent alignment response, improving congestion alignment under the urban congestion archetype while reducing alignment under the rural archetype.

SQ4: System Alignment Synthesis. The integrated evaluation demonstrates that no evaluated contractual arrangement resolves more than one of the three structural tensions identified in the institutional analysis. Compensation structure is the property most consistently associated with financial viability, although only under compensation levels substantially exceeding measured operational burden. Directive control is the property most strongly associated with congestion alignment, but introduces additional DAFO costs that are not fully reflected in the current remuneration structure. Firmness, evaluated through the prospective Firmness Gradient arrangement, is associated with the largest reduction in operational risk, reducing UOR to near zero without independently improving investment viability. Volume orientation, evaluated through the prospective Volume Bound arrangement, is associated with reduced degradation and a regime-dependent alignment response, improving alignment under the urban congestion archetype while reducing it under the rural archetype.

The combined picture is that the four evaluated property dimensions are each associated with a distinct performance outcome, compensation structure with financial viability, directive control with congestion alignment, firmness with reduced operational risk, and volume orientation with degradation and regime-dependent alignment, even though compensation structure and firmness originate from the same underlying tension of risk and value allocation. No single property dimension, nor any evaluated contractual arrangement, simultaneously improves financial viability, congestion alignment, and technical performance while addressing all three institutional tensions.

5.6.2. Answer to the Central Research Question

Existing and prospective flexible connection agreements influence the performance of utility-scale Battery Energy Storage Systems primarily through financial transfers and operational restrictions rather than through changes in technical asset life. Across the evaluated contract set, financial viability, asset degradation, and system alignment remain structurally separate performance dimensions: no single contractual arrangement simultaneously optimises all three.

Financial viability is achieved only when contractual remuneration effectively functions as an implicit tariff relief mechanism. Because the Dutch high-voltage transport tariff absorbs approximately 70% of gross market revenue, the unconstrained Baseline remains structurally non-bankable across all evaluated market regimes. Consequently, bankability is achieved only under the CBC, whose remuneration exceeds demonstrated operational burden by more than an order of magnitude.

System alignment is most strongly improved by the CSC, which produces the largest congestion alignment improvement within the evaluated contract set. However, this benefit depends critically on congestion-aligned DSO activation timing. Under random activation timing, directive control provides little improvement over the Baseline while introducing substantial DAFO costs for the asset owner.

Technical performance is comparatively insensitive to contractual design. Five of the six evaluated arrangements alter implied battery lifetime by less than four weeks relative to the Baseline. The Volume Bound arrangement is the sole exception, extending implied lifetime by just over two months, although this improvement is accompanied by lower profitability and a regime-dependent alignment response.

Combined with the property-level pattern already established in Sections 4.3.2 and 5.2, namely that no single property or contractual arrangement resolves more than one institutional tension, these findings indicate that the institutional frictions identified in Dutch non-firm transport arrangements cannot be resolved through contractual design alone. Current flexible connection agreements are required to perform two fundamentally different functions: managing congestion and compensating for a structural shortfall in the underlying tariff architecture. The principal policy challenge therefore lies not only in the design of flexible connection agreements themselves, but in the broader institutional framework of grid-access pricing within which those agreements operate.

5.7. Scientific Contribution

This research makes four principal contributions to the literature on energy storage, congestion management, and electricity market design.

First, it introduces an integrated exploratory sequential research design that systematically links qualitative institutional analysis with quantitative policy evaluation. Whereas existing research has largely treated the technical performance of battery storage and the institutional design of access arrangements as separate domains, the proposed research design provides a reproducible bridge between stakeholder-identified governance frictions and simulation-based policy evaluation. The resulting modular Policy Testbed provides a reusable platform for evaluating flexible connection arrangements beyond those considered in this study.

Second, it develops a property-based evaluation framework that decomposes flexible connection agreements into four underlying design dimensions: compensation structure, firmness, volume orientation, and directive control. These dimensions are derived directly from the institutional analysis rather than predefined from technical characteristics. The finding that each property is associated with a distinct and largely non-overlapping aspect of contract performance demonstrates that flexible connection agreement design is inherently a multi-objective problem rather than a single optimisation task. The framework therefore provides a transferable analytical lens for evaluating future congestion management instruments beyond the specific Dutch contractual arrangements examined here.

Third, the study operationalises this framework through two original prospective contract designs: the Firmness Gradient and the Volume Bound arrangement. Each was constructed to isolate a single design property from financial remuneration, allowing its institutional and operational effects to be evaluated more directly. Together, these prospective designs provide experimental evidence on how individual contractual design properties influence investor risk and system performance independently of compensation. The Firmness Gradient provides evidence that guaranteed access can substantially reduce operational risk independently of financial transfers, whereas the Volume Bound arrangement provides evidence that contractual solutions to temporal misalignment are strongly dependent on the underlying congestion regime.

Fourth, the research introduces three complementary evaluation metrics: UOR, the timing premium, and the CMI. Together, these metrics translate institutional concerns into measurable quantities that can be evaluated quantitatively. UOR operationalises the institutional concern regarding risk allocation by measuring the operational burden remaining after contractual remuneration. The timing premium quantifies the economic value of restriction timing, while the CMI measures the extent to which commercial battery dispatch aligns with congestion management objectives across different congestion archetypes. Collectively, these metrics provide a quantitative basis for evaluating institutional design questions that have previously been addressed primarily qualitatively and establish a common analytical framework through which future contractual designs can be evaluated as market and regulatory conditions evolve.

5.8. Future outlook

The findings of this study extend beyond the evaluation of individual flexible connection agreements and raise broader questions regarding the future design of Dutch electricity markets. While the Policy Testbed evaluates contractual arrangements under current institutional conditions, several ongoing policy developments may substantially alter the economic and operational environment in which BESS operate. The results therefore provide a foundation for exploring how future market reforms could interact with flexible connection design.

The interaction between congestion contracts and an emerging Dutch capacity market One of the most prominent policy discussions in the Netherlands concerns the potential introduction of a capacity market to support long-term security of supply, addressing the broader ”missing money” problem in which energy-only markets may not provide sufficient revenue certainty to incentivise investment in dispatchable and flexible capacity. Although such a mechanism remains subject to political and regulatory debate, the findings of this study suggest it could materially alter the role of congestion-management contracts.

The bankability analysis shows that financially viable investment outcomes currently emerge only under compensation levels, such as those embedded in CBC, that substantially exceed the operational risk the contract actually displaces. In the absence of a capacity market, congestion contracts are implicitly being asked to perform two functions simultaneously: providing operational congestion relief to the network operator while also supplying the investment certainty required for project financing, a tension made more acute by a transport tariff that already absorbs roughly seventy percent of gross battery revenues.

If a future capacity market provided a stable revenue stream linked to availability rather than congestion participation, the burden currently placed on congestion contracts to secure investment viability could be reduced. Under such conditions, contractual designs could in principle be calibrated more closely around system-alignment objectives, with remuneration levels reflecting the value of the congestion-management service delivered rather than the scale of financial support required to reach bankability. While the present study does not model capacity market revenues and cannot establish whether such a mechanism would in fact resolve the bankability paradox identified here, the finding suggests that future research should treat congestion contracts and capacity markets as interacting policy instruments rather than independent market designs.

The implications of future tariff reform A second major uncertainty concerns the evolution of the Dutch tariff framework. The analysis demonstrates that grid-related charges constitute one of the dominant cost components affecting battery project viability. Under the tested regulatory framework, transport tariffs absorb a substantial share of gross revenues across all scenario families, contributing directly to the observed tension between financial viability and proportionate compensation.

Several ongoing policy discussions, including proposals regarding producer-related network charges and alternative tariff structures for storage assets, have the potential to alter this relationship. The findings suggest that contractual redesign alone may be insufficient to resolve the viability challenge if the underlying tariff architecture remains unchanged. Future research should therefore investigate tariff reform and contract design jointly, treating access pricing and congestion-management arrangements as complementary elements of a single institutional system rather than separate policy domains.

Inertia procurement and the diversification of flexibility revenue The future revenue stack of utility-scale battery storage is likely to differ substantially from the revenue environment represented in the present model. Emerging developments in electricity market design, including discussions on inertia procurement, voltage support, and other stability services, suggest that future battery business cases may increasingly rely on a broader portfolio of system-service revenues in addition to energy arbitrage and balancing markets. The growing importance of these prospective revenue streams was also consistently highlighted during the expert validation, where participants identified the expansion of ancillary service markets as an important factor shaping the long-term attractiveness of flexible connection agreements.

In particular, the anticipated introduction of an inertia market represents a significant shift towards remunerating frequency stability. As conventional synchronous generation is progressively replaced by inverter-based renewable generation, the resulting decline in natural system inertia creates a growing need for synthetic or virtual inertia, a service that utility-scale battery energy storage systems are technically well positioned to provide through advanced power electronic converter control. Recent market design proposals suggest that such services may be procured through competitive auction mechanisms, providing relatively stable remuneration that could reduce reliance on volatile energy-market revenues.

The emergence of additional system-service revenues could consequently reduce the relative importance of congestion-contract remuneration within the overall battery business case and thereby alter the attractiveness of alternative contract designs. Contractual arrangements that preserve operational flexibility while providing partial access guarantees may become increasingly attractive as a larger share of project revenues is derived from system-service provision rather than exclusively from energy-market optimisation. Although these prospective markets are outside the scope of the present study, the property-based Policy Testbed developed in this thesis provides a suitable framework for evaluating how future revenue streams may interact with flexible connection agreements and influence their financial viability, operational performance, and system value.

Towards location-specific congestion management One of the findings of the Policy Testbed is the divergence between urban and rural network archetypes. Several contract designs exhibit materially different performance depending on whether congestion is primarily demand-driven or generation-driven. Most notably, the Volume Bound contract improves alignment in one context while reducing it in another.

This finding suggests that nationally uniform congestion-management arrangements may not always be appropriate for geographically heterogeneous congestion conditions. Future research should therefore explore location-specific contract calibration, allowing contractual parameters to adapt to the dominant congestion mechanism within a particular network area. Such an approach could improve alignment outcomes without requiring fundamentally different contractual architectures, thereby preserving regulatory simplicity while increasing operational effectiveness.

Directive congestion management at deployment scale The Capacity Steering Contract demonstrates that direct operational control can substantially improve alignment with network objectives, albeit at the cost of reduced financial performance. However, the present study evaluates this mechanism only at the level of a single utility-scale battery.

At deployment scale, where distribution system operators may coordinate large portfolios of storage assets simultaneously, additional questions emerge regarding activation efficiency, portfolio interactions, operational coordination, and aggregate market impacts. The extent to which the observed trade-off between alignment and profitability persists under portfolio-level optimisation remains an open question. Extending the framework developed in this study to multi-asset environments therefore represents an important next step in evaluating the practical feasibility of directive congestion-management approaches.

Linking alignment to societal value Finally, the Policy Testbed intentionally evaluates congestion-management contracts from the perspective of battery operators while measuring system alignment through the CMI. Although this approach identifies which contracts improve alignment, it does not monetise the societal value of that improvement.

Future research could link alignment outcomes to avoided reinforcement investments, reduced renewable curtailment, lower congestion-management expenditures, or other measurable system benefits. Establishing such a connection would make it possible to compare investor costs and societal benefits within a common economic framework. This would represent an important step towards assessing whether the societal benefits of improved congestion management justify the financial burden imposed on participating market actors.

Taken together, these research directions suggest that the future effectiveness of flexible connection agreements will depend on how contractual design interacts with broader developments in tariff regulation, capacity remuneration, emerging flexibility markets, and the governance of increasingly congested electricity systems. Rather than representing isolated policy instruments, flexible connection agreements should therefore be understood as one component of a wider institutional architecture governing the integration of flexibility resources.

6 Reflection

Having established the study's findings, scientific contribution, and implications in the previous chapter, this final chapter steps back from the results themselves. Rather than asking what the study found, it reflects on how the research was conducted, the methodological choices that shaped those findings, their broader limitations, and the lessons learned throughout the project.

6.1. Methodological Reflection

Perhaps the most important methodological insight gained during this research is that flexible connection agreements cannot be evaluated meaningfully without first understanding the institutional problem they are intended to address. Many energy systems studies begin with an optimisation framework and subsequently evaluate policy options within that predefined technical setting. This study deliberately followed the opposite sequence. The institutional analysis defined the research problem, identified the structural tensions requiring investigation, and only then informed the development of the quantitative evaluation framework.

Looking back, this sequence fundamentally shaped both the research process and its contributions. Beginning with stakeholder concerns prevented the optimisation model from becoming an exercise in technical performance evaluation detached from the institutional context in which flexible connection agreements operate. Instead, the modelling framework was developed to evaluate questions that had already emerged from the regulatory debate. This created a stronger connection between the qualitative and quantitative components of the study than would have been achieved through a conventional optimisation-first approach.

A second methodological lesson concerns the value of abstraction. Existing contractual arrangements differ simultaneously in multiple design dimensions, making direct contract-to-contract comparisons difficult to interpret. Decomposing these arrangements into four underlying design properties, compensation structure, firmness, volume orientation, and directive control, provided a more robust analytical framework. Rather than evaluating contracts as isolated cases, this property-based perspective made it possible to identify broader design relationships that extend beyond the specific Dutch arrangements evaluated here. In hindsight, this abstraction step was what made the difference between a simulation study and a transferable analytical lens: a contract-by-contract framing would have produced findings specific to six Dutch instruments rather than a framework applicable to instruments not yet designed. This was an important lesson in balancing realism with generalisability, and one I expect to carry into future engineering research: modelling every detail of a specific case is often less valuable than identifying the smaller set of underlying properties that make a finding transferable.

6.2. Research Limitations

The findings of this study should be interpreted within the scope of the adopted research design and modelling assumptions. Limitations affecting the interpretation of specific results are discussed at the relevant point in Chapter 5; the limitations below concern the research design as a whole.

First, the institutional analysis is based on publicly available consultation documents and stakeholder reports. Although these sources provide a comprehensive representation of formally articulated positions within the Dutch regulatory debate, they do not capture informal negotiations, commercial considerations, or internal decision-making processes that may also influence contractual design. An independent analyst applying the same procedure could plausibly derive different, but equally defensible, categorisations. The scope of this study, evaluating what contractual design can achieve within the existing institutional environment, intentionally excludes tariff reform scenarios and allocation architecture changes. This boundary is coherent with the research question but means that the most consequential source of institutional friction identified in this thesis is not modelled as a design variable.

Second, the Policy Testbed necessarily abstracts from operational reality. The rolling-horizon optimisation with limited foresight improves on perfect-foresight models but remains a simplification. Intraday trading, reactive power provision, inertia services, and dynamic network constraints are not modelled. The battery is represented as a price taker, an assumption that becomes increasingly restrictive under high storage deployment futures, where the aggregate behaviour of storage assets would itself compress price spreads. The degradation model captures the dominant engineering mechanisms affecting lifetime but does not represent temperature-dependent ageing, electrolyte decomposition, or state-of-health feedback on power capability; degradation results should therefore be read as comparative indicators rather than lifetime forecasts. These simplifications primarily affect the quantitative magnitude of the reported outcomes rather than the comparative relationships between contractual arrangements.

Third, market uncertainty is represented through independent Monte Carlo variation in volatility, market saturation, congestion intensity, and congestion timing. In practice these variables are correlated: high renewable penetration simultaneously generates price volatility, storage competition, and congestion stress. This independence assumption may underrepresent the most adversarial scenario combinations, meaning tail-risk findings should be interpreted as lower bounds rather than comprehensive estimates of adverse exposure. Future research could extend the framework by modelling correlated uncertainty processes.

Finally, the property-based evaluation remains observational rather than experimental. Because existing contractual arrangements differ across multiple properties simultaneously, observed performance differences cannot be attributed exclusively to individual design dimensions. The prospective contract designs partially mitigate this limitation by isolating selected properties, but a more controlled architecture, holding three dimensions fixed while varying only a fourth, would provide stronger mechanistic evidence at the cost of evaluating hypothetical rather than institutionally observed arrangements. Accordingly, the relationships identified throughout this thesis should be interpreted as robust associations rather than definitive evidence of causal mechanisms.

6.3. Personal Reflection

This research marked a transition from applying established methods towards developing and defending an original research approach. Perhaps the most significant personal development was recognising that optimisation models are not only computational tools, but also instruments for testing institutional hypotheses. Rather than viewing the optimisation model as the objective of the research, I increasingly came to regard it as a means of evaluating how governance structures shape technical and economic outcomes.

Looking back, one of the biggest challenges was resisting the temptation to begin with a sophisticated optimisation model. Early in the project, I naturally approached the problem from a technical perspective. Through discussions with my supervisors and the institutional analysis, I gradually realised that the more fundamental question was not how to optimise a battery, but which institutional assumptions should actually be tested.

Working across disciplinary boundaries required continuous movement between qualitative institutional reasoning and quantitative optimisation modelling. Decisions made during the institutional analysis frequently required adjustments to the simulation framework, while unexpected model behaviour often prompted a reconsideration of the original institutional assumptions. This iterative dialogue between qualitative and quantitative analysis became one of the defining characteristics of the project and fundamentally changed how I approach complex engineering problems.

The project also strengthened my confidence in conducting independent scientific research. Throughout the study, methodological choices increasingly emerged from critical evaluation of the research problem itself rather than from reproducing approaches adopted in previous literature, and defending these choices through continuous discussion with academic supervisors and industry experts was an important part of that development. Perhaps the most valuable lesson was recognising that scientific contribution is not achieved by developing increasingly complex models, but by asking more meaningful questions and selecting methods appropriate for answering them: looking back, the research increasingly shifted from asking how a battery should be optimised towards asking which institutional assumptions should be tested in the first place, a shift in thinking likely to remain the most enduring outcome of this thesis, extending well beyond its specific findings.

Bibliography

[1] ACER. *ACER's Final Assessment of the EU Wholesale Electricity Market Design*. Tech. rep. Mar. 2022. URL: https://www.acer.europa.eu/sites/default/files/documents/Publications/Final_Assessment_EU_Wholesale_Electricity_Market_Design.pdf.

[2] ACM. *Besluit tot wijziging van de tariefstructuren en voorwaarden als bedoeld in artikelen 27 en 31 van de Elektriciteitswet 1998 betreffende de non-firm transportovereenkomst*. Tech. rep. Sept. 2023.

[3] ACM. *Besluit van de Autoriteit Consument en Markt van [datum ondertekening], kenmerk definitieve besluit tot wijziging van de tariefstructuren en voorwaarden als bedoeld in artikelen 27 en 31 van de Elektriciteitswet 1998 betreffende alternatieve transportrechten*. Tech. rep. Mar. 2024.

[4] ACM. *Besluit van de Autoriteit Consument en Markt van 11 december 2025, kenmerk ACM/UIT/662511 tot wijziging van de voorwaarden als bedoeld in artikel 31 juncto artikel 32, tweede lid van de Elektriciteitswet 1998 betreffende het capaciteitssturingscontract*. Tech. rep. Nov. 2025.

[5] ACM. *Overzicht en inzicht congestiemaatregelen ACM*. Tech. rep. ACM, July 2025. URL: https://www.acm.nl/system/files/documents/overzicht-congestiemaatregelen-acm-juli-2025.pdf.

[6] ACM. *Transparantie voor onbalanssystematiek*. Tech. rep. Authority for Consumers and Markets (ACM), 2024. URL: https://www.acm.nl/sites/default/files/old_publication/bijlagen/6935_12_19168.pdf.

[7] ACM. *Vragen en antwoorden – Congestiemanagement en het capaciteitsbeperkingscontract*. Tech. rep. May 2022. URL: https://www.acm.nl/system/files/documents/q-en-a-over-congestiemanagement-en-capaciteitsbeperkingsconctract.pdf.

[8] International Energy Agency et al. *World Energy Outlook Special Report Batteries and Secure Energy Transitions*. Tech. rep. IEA, 2023. URL: https://iea.blob.core.windows.net/assets/cb39c1bf-d2b3-446d-8c35-aae6b1f3a4a0/BatteriesandSecureEnergyTransitions.pdf.

[9] Luca Argiolas et al. "Optimal battery energy storage dispatch in energy and frequency regulation markets while peak shaving an EV fast charging station". In: *IEEE Open Access Journal of Power and Energy* 9 (Jan. 2022), pp. 374–385. DOI: 10.1109/oajpe.2022.3198553. URL: https://doi.org/10.1109/oajpe.2022.3198553.

[10] ATLAS.ti Scientific Software Development GmbH. *ATLAS.ti User Manual*. Tech. rep. Version x. 2026. URL: https://atlasti.com.

[11] Bart.O. "Expert validation of input values". Personal communication. 2026.

[12] Kjersti Berg, Rubi Rana, and Hossein Farahmand. "Quantifying the benefits of shared battery in a DSO-energy community cooperation". In: *Applied Energy* 343 (May 2023), p. 121105. DOI: 10.1016/j.apenergy.2023.121105. URL: https://doi.org/10.1016/j.apenergy.2023.121105.

[13] Sanne de Boer and RaboResearch. "The Dutch electricity sector faces major challenges". In: *In focus* (July 2024). URL: https://www.rabobank.com/knowledge/b011436116-the-dutch-electricity-sector-explained.

[14] Merve Bulut and Evrencan Ozcan. *How to build a state-of-the-art battery energy storage market? Challenges, opportunities, and future directions*. Tech. rep. 2024, p. 111174. URL: https://doi.org/10.1016/j.est.2024.111174.

[15] GOPACS BV. *Flexproduct: hoe werkt een CBC-T? - GOPACS*. Apr. 2026. URL: http%5Cbibliography%7Breferences%7Ds://www.gopacs.eu/producten/cbc-t/.

[16] Lucas van Cappellen et al. *Beleid voor grootschalige batterijsystemen en afnamenetcongestie*. Tech. rep. CE Delft, Mar. 2023. URL: https://ce.nl/wp-content/uploads/2023/04/CE_Delft_220376_Achtergrondrapport_Beleid_voor_grootschalige_batterijsystemen_en_afnamenetcongestie_DEF.pdf.

[17] A. Castillejo-Cuberos, J.M. Cardemil, and R. Escobar. *Techno-economic assessment of photovoltaic plants considering high temporal resolution and non-linear dynamics of battery storage*. Tech. rep. 2023, p. 120712. URL: https://doi.org/10.1016/j.apenergy.2023.120712.

[18] Lincoln Creswell John W & University of Nebraska. *Qualitative Inquiry & Research design: choosing among five approaches*. Second Edition. SAGE Publications, 2007. URL: https://eclass.uowm.gr/modules/document/file.php/MATHMAST103/Creswell%20Qualitative_Inquiry_and_Research.pdf.

[19] Federico De Santi et al. “Managing Connection Queues in Distribution Networks with Flexible Connection Agreements”. In: *Energy Policy* 198 (2025), p. 115127. ISSN: 0301-4215. DOI: 10.1016/j.enpol.2024.115127.

[20] Korina Dedousi. *Financial feasibility study of utility-scale batteries in the Dutch Electricity Market*. 2025.

[21] Chuma Ebere et al. “Electricity grid and societal vulnerability interconnection: stakeholder implications and integrated solutions in Europe”. In: *Environ. Res. Lett.* 20 (2025), p. 103007. URL: https://doi.org/10.1088/1748-9326/ae0492.

[22] Nora Ekström. *Maximizing the Value of Large-Scale Solar PV Parks through Battery Storage and Ancillary Services*. Tech. rep. Uppsala Universitet, 2023. URL: https://www.diva-portal.org/smash/get/diva2:1777292/FULLTEXT01.pdf.

[23] *Elektriciteitswet 1998*. Geraadpleegd op 8 maart 2026. Overheid.nl, 2025. URL: https://wetten.overheid.nl/BWBR0009755/2025-02-22.

[24] Corre Energy. *Zienswijze Corre Energy Netcodes/Consultatie ATR85*. Tech. rep. 2025.

[25] Energy Storage NL. *Zienswijze op het ontwerp-codebesluit Capaciteitssturingscontract (CSC)*. Consultation response submitted to the Autoriteit Consument & Markt (ACM). Public consultation document on the draft code decision for the Capacity Steering Contract (CSC). The Hague, Netherlands, Oct. 2025.

[26] Bert Enserink et al. *Policy analysis of Multi-Actor Systems*. 2022. URL: https://doi.org/10.5074/T.2022.004.

[27] ENTSO-E. *Transparency Platform*. 2026. URL: https://transparency.entsoe.eu/ (visited on 05/12/2026).

[28] Werkgroep Grootschalige Opslag ESNL. *Werkgroep Grootschalige opslag ESNL – 13-05-2026*. Tech. rep. Apr. 2026.

[29] Daniel Fernández-Muñoz and Juan I. Pérez-Díaz. “Optimisation models for the day-ahead energy and reserve self-scheduling of a hybrid wind–battery virtual power plant”. In: *Journal of Energy Storage* 57 (Dec. 2022), p. 106296. DOI: 10.1016/j.est.2022.106296. URL: https://doi.org/10.1016/j.est.2022.106296.

[30] Frank W. Geels. “From sectoral systems of innovation to socio-technical systems: Insights about dynamics and change from sociology and institutional theory”. In: *Research Policy* 33.6-7 (2004), pp. 897–920. DOI: 10.1016/j.respol.2004.01.015.

[31] GOPACS. *Capacity Steering Contract (CSC) - GOPACS*. Apr. 2026. URL: https://www.gopacs.eu/en/glossary/capacity-steering-contract-csc/.

[32] Tanne Heemsbergen. *Congestion-Neutral Operation of Residential Battery Storage in the Dutch Electricity Grid*. Delft University of Technology, 2025. URL: http://repository.tudelft.nl/.

[33] Lion Hirth. “The market value of variable renewables: The effect of solar wind power variability on their relative price”. In: *Energy Economics* 38 (2013), pp. 218–236. ISSN: 0140-9883. DOI: https://doi.org/10.1016/j.eneco.2013.02.004. URL: https://www.sciencedirect.com/science/article/pii/S0140988313000285.

[34] Lion Hirth and Samuel Glismann. *Congestion Management: From Physics to Regulatory Instruments. A Guide to the European Approach of Managing Congestion in Electricity Networks with Zonal Markets*. Working Paper. Kiel and Hamburg: ZBW – Leibniz Information Centre for Economics, 2018. URL: https://hdl.handle.net/10419/189641 (visited on 04/24/2026).

[35] Paul L. Joskow. “Incentive Regulation and Its Application to Electricity Networks”. In: *Review of Network Economics* 7.4 (2008), pp. 547–560. DOI: 10.2202/1446-9022.1161.

[36] Pratyusha Kalluri and Pablo Gervás. “Affinity-based interpretation of triangle social scenarios”. In: *Scitepress* (Jan. 2017), pp. 640–647. DOI: 10.5220/0006205506400647. URL: https://doi.org/10.5220/0006205506400647.

[37] Kris Kessels et al. *Specification and design criteria for local flexibility markets*. Tech. rep. European Commission, 2026. URL: https://op.europa.eu/en/publication-detail/-/publication/93fda27d-1453-11f1-8870-01aa75ed71a1/language-en.

[38] Elizabeth Koehler, Elizabeth Brown, and Sebastien J.-p. A. Haneuse. "On the Assessment of Monte Carlo Error in Simulation-Based Statistical Analyses". In: *The American Statistician* 63.2 (May 2009), pp. 155–162. DOI: 10.1198/tast.2009.0030. URL: https://doi.org/10.1198/tast.2009.0030.

[39] KYOS. *Battery Landscape Report*. Tech. rep. Sept. 2025, pp. 1–34.

[40] Simon Massat et al. *Development and analysis of a battery degradation model for use in mixed-integer linear programming*. Tech. rep. Technical University of Darmstadt, Jan. 2026. DOI: 10.1016/j.est.2026.120511. URL: www.elsevier.com/locate/est.

[41] Madeleine McPherson et al. *Impacts of storage dispatch on revenue in electricity markets*. Tech. rep. June 2020. URL: https://doi.org/10.1016/j.est.2020.101573.

[42] Michael Merten M. Sc. *Participation of Battery Storage Systems in the Automatic Frequency Restoration Reserve Market based on Machine Learning*. Tech. rep. 2020. URL: https://publications.rwth-aachen.de/record/811207/files/811207.pdf.

[43] Ahmed Mohamed, Rémy Rigo-Mariani, and Vincent Debusschere. *Stacked Revenues for Energy Storage Participating in Energy and Reserve Markets with an Optimal Frequency Regulation Modeling*. Tech. rep. Univ. Grenoble Alpes, CNRS, Grenoble INP, G2Elab, Atos Worldgrid Solutions for Energy, and Utilities, 2022.

[44] Netbeheer Nederland. *Codewijzigingsvoorstel alternatieve transportrechten*. Tech. rep. Jan. 2024.

[45] Netbeheer Nederland. *Kies voor beter benutten van het elektriciteitsnet*. Netbeheer Nederland. 2025. URL: https://www.netbeheernederland.nl/tijdomtekiezen/kies-voor-beter-benutten-van-het-elektriciteitsnet (visited on 06/14/2026).

[46] Energy Storage NL. *Reactie Energy Storage NL op ACM-consultatie 'Alternatieve transportrechten' en 'Use it or Lose it'*. Tech. rep. Nov. 2022.

[47] European Parliament and Council of the European Union. *REGULATION (EU) 2021/1119 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL*. Tech. rep. Publication Office, June 2021.

[48] Zsuzsanna Pató and Regulatory Assistance Project. *Gridlock in the Netherlands*. 2024, pp. 2–3. URL: https://www.raponline.org/wp-content/uploads/2024/01/RAP-Pato-Netherlands-gridlock-2024.pdf.

[49] Peter Wennink. *De route naar toekomstige welvaart*. Tech. rep. Commissie Wennink, Dec. 2025. URL: https://www.rapportwennink.nl/downloads/rapport_wennink_12december2025.pdf.

[50] Ksenia Poplavskaya et al. "Integration of day-ahead market and redispatch to increase cross-border exchanges in the European electricity market". In: *Applied Energy* 278 (Aug. 2020), p. 115669. DOI: 10.1016/j.apenergy.2020.115669. URL: https://doi.org/10.1016/j.apenergy.2020.115669.

[51] Antonia Proka a et al. *When top-down meets bottom-up: Is there a collaborative business model for local energy storage?* Tech. rep. June 2020. URL: https://doi.org/10.1016/j.erss.2020.101606.

[52] Rick Schouten. *Economic Value of BESS Participation in Market-Based Redispatch: A Market-Based Analysis of congestion Management revenue stacking in the Dutch Grid*. Nov. 2025.

[53] Stedin. *CONGESTIEMANAGEMENT-ONDERZOEK*. Tech. rep. Stedin, 2024.

[54] *Stedin Group Annual Report 2023*. Tech. rep. 2023.

[55] Laura Stok. *A new market design for balancing and redispatch*. Tech. rep. 2024.

[56] TenneT. *Annual Market Update 2025*. Tech. rep. TenneT, 2025.

[57] TenneT. *Balancing markets*. n.d. URL: https://www.tennet.eu/node/268 (visited on 05/13/2026).

[58] TenneT. *Integrated Annual Report 2024*. Tech. rep. 2024.

[59] TenneT. *TenneT's position on Battery Energy Storage Systems (BESS)*. Tech. rep. 2023, pp. 3–18. URL: https://tennet-drupal.s3.eu-central-1.amazonaws.com/default/2023-06/TenneT_s_position_large_BESS_-_Public_Info_-_update.pdf.

[60] TenneT. *Transparency Platform Download Page*. 2026. URL: https://www.tennet.eu/nl/nl-en/markets/transparency/download-page-transparency (visited on 05/12/2026).

[61] Musa Terkes et al. *Hybrid Renewable Energy Systems in Türkiye: A Multi-Scenario Assessment of demand, metering, and Carbon tax policies*. Tech. rep. University of New York, Yildiz Technical University,Department of Electrical Engineering, 2025, pp. 1–. DOI: 10.1109/ACCESS.2025.3587879.

[62] George Trienekens, Rick Scharrenberg, and Longjian Piao. “A Dutch Market Design on Congestion Management: The Role of (Fully) Independent Aggregators”. In: *Proceedings of the International Conference on the European Energy Market (EEM)*. TenneT TSO B.V. and TenneT TSO GmbH. Arnhem, The Netherlands, 2025.

[63] Paul Verhagen. *Impact of Dutch annual grid fees on large-scale battery energy storage systems*. June 2024.

[64] Paul Verhagen, Jing Hu, and Robert Harmsen. “Evaluating non-firm grid connections for battery energy storage systems: A co-optimization case study of the Netherlands”. In: *Energy Policy* 208 (Oct. 2025), p. 114903. DOI: 10.1016/j.enpol.2025.114903. URL: https://doi.org/10.1016/j.enpol.2025.114903.

[65] Arthur de Visser. *Institutional dynamics and systemic intermediaries in the Dutch electrification transition: Navigating flexibility challenges*. July 2025.

[66] Ying Wang et al. “Stochastic coordinated operation of wind and battery energy storage system considering battery degradation”. In: *Journal of Modern Power Systems and Clean Energy* 4.4 (Oct. 2016), pp. 581–592. DOI: 10.1007/s40565-016-0238-z. URL: https://doi.org/10.1007/s40565-016-0238-z.

[67] Bolun Xu et al. *Modeling of Lithium-Ion battery degradation for cell life assessment*. Tech. rep. 2016. DOI: 10.1109/TSG.2016.2578950.

A

Institutional Analysis Coding Validation: Manual versus Automated Analysis

This appendix documents a validation exercise conducted to assess whether the manual coding procedure adopted in the institutional analysis was methodologically justified. Rather than serving as the primary analytical method, Atlas.ti was used retrospectively to evaluate whether automated keyword-based coding would have produced an equivalent evidence base for the institutional analysis.

A.1. Atlas.ti Trial Procedure

The complete corpus of 33 stakeholder documents was imported into Atlas.ti 23. A keyword-based auto-coding procedure was applied using search terms related to congestion management, grid access, transport rights, flexibility, Battery Energy Storage Systems (BESS), and operational restrictions.

The automated procedure generated 2,541 codes across 1,573 quotations. These quotations were subsequently reviewed against the inclusion criteria defined in Section 3.2 to assess whether they represented analytically meaningful institutional frictions.

A.2. Assessment of Automated Output

A manual inspection of a stratified sample of automatically generated quotations revealed three recurring categories of non-analytical matches.

- **Procedural and legal references**, including consultation procedures, legal articles, formal definitions, and administrative descriptions that contained relevant terminology but did not describe operational or institutional limitations.

- **Metadata and document structure**, including document identifiers, addresses, headings, captions, and reference lists that were captured solely because they contained matching keywords.

- **Context-dependent passages**, in which terms such as “transport”, “capacity”, “connection”, or “congestion” were correctly identified but whose institutional meaning depended on the surrounding discussion. When extracted in isolation, these quotations often no longer represented identifiable institutional frictions.

These observations indicate that keyword-based coding was effective for identifying candidate passages, but could not reliably distinguish between descriptive text, procedural content, and analytically relevant stakeholder statements.

This limitation became particularly apparent in consultation documents, which frequently contained implementation questions, legal references, procedural explanations, and figure captions. Such sections generated large numbers of keyword matches despite often referring to the same underlying institutional issue or containing no operational consequence for BESS. Determining whether a quotation represented a genuine institutional friction therefore required interpretation of the surrounding document context rather than evaluation of the keyword occurrence in isolation.

A.3. Manual Coding Procedure

The manual coding procedure was conducted iteratively across the complete document corpus. Each document was first read in full to preserve its institutional context. Potentially relevant passages were subsequently highlighted, extracted, and reviewed against the predefined inclusion criteria.

A fragment was retained only when it simultaneously described:

1. a concrete operational, technical, or economic consequence for a BESS, and
2. an identifiable institutional mechanism, such as a tariff design, contractual provision, dispatch rule, or governance arrangement that caused or contributed to that consequence.

Fragments that were purely descriptive, procedural, administrative, or normative without an identifiable operational implication were excluded. Where multiple quotations referred to the same underlying institutional issue, they were interpreted collectively rather than counted as independent observations.

Applying these criteria resulted in a final analytical dataset of 112 coded fragments, which formed the basis for the institutional clustering presented in the institutional analysis.

Table A.1 summarises the comparison between both approaches.

Table A.1: Comparison between the automated Atlas.ti validation exercise and the manual coding procedure adopted.

Dimension	Atlas.ti auto-coding	Manual coding
Documents analysed	33	33
Primary purpose	Candidate identification	Institutional interpretation
Selection basis	Keyword occurrence	Explicit inclusion criteria
Context considered	Individual quotations	Complete document context
Output	2,541 codes (1,573 quotations)	112 analytical fragments
Role in this study	Validation exercise	Primary analytical method

The validation exercise demonstrates that automated keyword matching was useful for identifying potentially relevant passages but was not sufficient for constructing the institutional evidence base required in this study. Because the objective of the analysis was to identify causal relationships between operational limitations and institutional mechanisms rather than to count keyword occurrences, contextual interpretation formed an essential part of the analytical process. Manual coding was therefore not adopted as a matter of researcher preference, but as a methodological necessity for producing a coherent and internally consistent dataset of institutional frictions.

B
Technical Model Documentation

This appendix provides the detailed technical specification of the BESS Policy Testbed. The objective is to ensure transparency, reproducibility, and model traceability while keeping the main methodology chapter focused on conceptual design choices and modelling logic.

B.1. Notation and Parameters

B.1.1. Decision Variable Reference

Table B.1 summarises the principal decision variables used within the optimisation model.

Table B.1: Decision variable reference.

Symbol	Description	Unit	Bounds
p_t^{ch}	Physical charging power at PTU t	MW	$[0, P^{\mathrm{rated}}]$
p_t^{dis}	Physical discharging power at PTU t	MW	$[0, P^{\mathrm{rated}}]$
u_t	Binary operating mode variable	–	$\{0, 1\}$
p_t^{DA}	Day-Ahead scheduled position	MW	$[-P^{\mathrm{rated}}, P^{\mathrm{rated}}]$
$c_t^{\uparrow}$	Upward aFRR reserve capacity	MW	$[0, \bar{c}^{\mathrm{aFRR}}]$
$c_t^{\downarrow}$	Downward aFRR reserve capacity	MW	$[0, \bar{c}^{\mathrm{aFRR}}]$
$p_t^{\mathrm{imb,up}}$	Upward imbalance dispatch	MW	$[0, P^{\mathrm{rated}}]$
$p_t^{\mathrm{imb,dn}}$	Downward imbalance dispatch	MW	$[0, P^{\mathrm{rated}}]$
d_t^{+}	Positive DAFO deviation	MW	$[0, P^{\mathrm{rated}}]$
d_t^{-}	Negative DAFO deviation	MW	$[0, P^{\mathrm{rated}}]$
s_t	State of Charge	–	$[0.05, 0.95]$
κ^{win}	Cycling degradation cost	EUR	$[0, \infty)$
λ_i	SOS2 interpolation weights	–	$[0, 1]$

B.1.2. Model Parameter Reference

Table B.2 summarises the principal model parameters.

Table B.2: Key model parameters.

Symbol	Value	Description
P^{rated}	10 MW	Rated battery power
E^{rated}	20 MWh	Rated energy capacity
η	0.90	One-way efficiency
$\underline{s}$	0.05	Minimum State-of-Charge bound
$\bar{s}$	0.95	Maximum State-of-Charge bound
α^{AR}	0.08	Historical aFRR activation ratio
T^{end}	0.25 h	Reserve endurance requirement
ϕ	0.20	Imbalance participation cap
ρ^{RVP}	0.60	Relative reserve participation factor
β	1.79	Cycle degradation exponent
γ^{bidir}	0.50	Bidirectional aFRR activation cancellation factor

The bidirectional cancellation factor $\gamma^{bidir} = 0.50$ reflects the empirical observation that aFRR activations are approximately symmetric in direction over any extended period [64]. Expected net throughput from aFRR reservation is therefore approximately $\alpha^{AR} \cdot \gamma^{bidir}$ per unit of reserved capacity per PTU.

B.2. Benchmark MILP Formulation

The optimisation problem is solved over sequential rolling-horizon windows consisting of $H = 192$ Program Time Units (PTUs), corresponding to 48 hours. The objective is to maximise net operational value while respecting physical battery constraints, market participation rules, and contractual obligations.

B.2.1. Optimisation Objective

The benchmark model represents a profit-maximising battery operator participating simultaneously in the Day-Ahead, aFRR capacity, and imbalance markets. The objective of the optimisation is to maximise net operational value while internalising both technical degradation costs and contractual charges.

The optimisation problem is formulated as:

$$\max Z = R^{\mathrm{DA}} + R^{\mathrm{aFRR}} + R^{\mathrm{imb}} + R^{\mathrm{contract}} - C^{DAFO} - \kappa^{\mathrm{win}} - \kappa^{\mathrm{cal}} - \kappa^{\mathrm{grid}} \tag{B.1}$$

where:

- R^{DA} denotes revenues obtained through Day-Ahead energy arbitrage;
- R^{aFRR} denotes revenues from aFRR capacity reservation;
- R^{imb} denotes revenues from imbalance market participation;
- R^{contract} denotes contract-specific remuneration;
- C^{DAFO} represents imbalance settlement costs arising from deviations between scheduled and realised Day-Ahead positions;
- κ^{win} represents cycling degradation costs;
- κ^{cal} represents calendar degradation costs;
- κ^{grid} represents annual grid connection charges.

The objective function reflects the central modelling assumption that battery operation is driven by long-term economic value rather than short-term revenue maximisation alone. By incorporating degradation costs directly into the optimisation objective, the model internalises the opportunity cost of consuming battery lifetime. Consequently, the battery only elects to cycle when expected market revenues exceed the marginal value of the associated degradation. This allows the simulated asset to behave as a rational investor facing the same economic trade-offs as a real-world Battery Energy Storage System (BESS) operator.

B.2.2. Physical Capacity Constraints

Charging and discharging are mutually exclusive and jointly constrained by the rated power of the battery:

$$p_t^{\mathrm{dis}} + c_t^{\uparrow} + p_t^{\mathrm{imb,up}} \leq P^{\mathrm{rated}} u_t \tag{B.2}$$

$$p_t^{\mathrm{ch}} + c_t^{\downarrow} + p_t^{\mathrm{imb,dn}} \leq P^{\mathrm{rated}} (1 - u_t) \tag{B.3}$$

B.2.3. Day-Ahead Imbalance Settlement (DAFO) Cost

Deviations between the Day-Ahead scheduled position p_t^{DA} and the realised physical dispatch $(p_t^{\mathrm{dis}} - p_t^{\mathrm{ch}})$ are settled through the imbalance mechanism. The realised deviation is decomposed into non-negative slack variables:

$$C^{DAFO} = \sum_{t \in T^{\mathrm{com}}} \left(r_t^{\mathrm{imb,up}} \cdot d_t^{+} + r_t^{\mathrm{imb,down}} \cdot d_t^{-} \right) \cdot \Delta t \tag{B.4}$$

where $r_t^{\mathrm{imb,up}}$ and $r_t^{\mathrm{imb,down}}$ denote the directional imbalance settlement prices at PTU t (the same 15-minute-resolution series used in the imbalance revenue term R^{imb}, described in Table B.3), applied respectively to positive and negative DAFO deviations.

B.2.4. Imbalance Market Constraints

The optimisation assumes that the battery behaves as a price-taker and cannot materially influence market outcomes. To preserve this assumption, participation in the imbalance market is limited to a maximum share of observed system volume:

$$p_t^{\text{imb,up}} \leq \phi V_t^{\text{imb}} \tag{B.5}$$

$$p_t^{\text{imb,dn}} \leq \phi V_t^{\text{imb}} \tag{B.6}$$

where V_t^{imb} denotes the observed system imbalance volume and $\phi = 0.20$ represents the maximum participation share. This constraint prevents the 10 MW battery from capturing an unrealistically large fraction of the imbalance market, ensuring consistency with observed Dutch market liquidity and maintaining the price-taker assumption.

Directional participation is determined by the national regulation state signal.

B.2.5. State-of-Charge Dynamics

The battery state of charge (SoC) evolves as a function of net energy flows into and out of the battery. Energy inflows arise from charging, downward imbalance participation, and expected downward aFRR activation, while energy outflows arise from discharging, upward imbalance participation, and expected upward aFRR activation.

The initial state of charge is defined as

$$s_0 = s^{\text{init}} \tag{B.7}$$

The expected charging energy is defined as

$$P_t^{in} = p_t^{\text{ch}} + p_t^{\text{imb,dn}} + \alpha^{AR} c_t^{\downarrow} \tag{B.8}$$

and the expected discharging energy as

$$P_t^{out} = p_t^{\text{dis}} + p_t^{\text{imb,up}} + \alpha^{AR} c_t^{\uparrow}. \tag{B.9}$$

The state of charge then evolves according to

$$s_t = s_{t-1} + \frac{\Delta t}{E^{\text{rated}}} \left(\eta P_t^{in} - \frac{1}{\eta} P_t^{out} \right). \tag{B.10}$$

The first term inside the brackets represents energy entering the battery and is multiplied by the charging efficiency η. The second term represents energy leaving the battery and is divided by η to account for discharge losses. The parameter α^{AR} denotes the expected activation fraction of reserved aFRR capacity and converts reserve commitments into their expected energy impact on the battery state.

B.2.6. Reserve Deliverability Constraints

Reserve capacity can only be offered if the battery is able to sustain the requested activation throughout the required endurance period. Upward reserve therefore requires sufficient energy to remain available in the battery, while downward reserve requires sufficient empty capacity.

The resulting SoC constraints are:

$$s_t \geq \underline{s} + \frac{c_t^{\uparrow} T^{end}}{E^{\text{rated}} \eta} \tag{B.11}$$

$$s_t \leq \bar{s} - \frac{c_t^{\downarrow} T^{end}}{E^{\text{rated}}} \tag{B.12}$$

where T^{end} denotes the reserve endurance requirement, the duration for which reserved capacity must remain continuously deliverable following activation. These constraints ensure that committed reserve capacity remains physically deliverable throughout the endurance period.

B.2.7. Block Bidding Constraints

Consistent with TenneT procurement rules, reserve capacity bids are submitted per four-hour procurement block and must remain constant throughout the block. The following constraints therefore enforce identical upward and downward reserve bids for all PTUs within a given block [55]:

$$c_t^{\uparrow} = c_{t_0^b}^{\uparrow} \qquad \forall t \in b \tag{B.13}$$

$$c_t^{\downarrow} = c_{t_0^b}^{\downarrow} \qquad \forall t \in b \tag{B.14}$$

where t_0^b denotes the first PTU of block b. These constraints ensure that upward and downward reserve bids remain fixed within each four-hour procurement block, preventing the optimisation from adjusting reserve commitments at finer time resolutions than permitted by market rules.

B.2.8. Imbalance Residual Budget Constraints

Reserve commitments take precedence over imbalance-market positions. Consequently, any capacity reserved for aFRR is no longer available for imbalance dispatch.

The remaining imbalance capacity is therefore limited by the battery's rated power minus the committed reserve volume:

$$p_t^{\text{imb,up}} \le P^{\text{rated}} u_t - c_t^{\uparrow} \tag{B.15}$$

$$p_t^{\text{imb,dn}} \le P^{\text{rated}} (1 - u_t) - c_t^{\downarrow} \tag{B.16}$$

where u_t indicates the discharge mode of the battery. These constraints prevent the optimisation from simultaneously allocating the same physical capacity to both reserve provision and imbalance trading.

B.2.9. Daily Throughput Cap

The total directional energy throughput within the optimisation window, Θ^{win}, combines physical dispatch, imbalance dispatch, and the bidirectional-cancelled contribution of reserved aFRR capacity:

$$\Theta^{\text{win}} = \sum_t \left(p_t^{\text{dis}} + p_t^{\text{imb,up}} + \alpha^{AR} \gamma^{bidir} c_t^{\uparrow} + p_t^{\text{ch}} + p_t^{\text{imb,dn}} + \alpha^{AR} \gamma^{bidir} c_t^{\downarrow} \right) \Delta t \tag{B.17}$$

To prevent unrealistic cycling behaviour under extreme price conditions, throughput within each window is capped at a maximum number of equivalent daily cycles:

$$\Theta^{\text{win}} \le 2 \cdot E^{\text{rated}} \cdot 0.75 \cdot n_{\text{days}}^{\text{win}} \tag{B.18}$$

This corresponds to a maximum of 1.5 equivalent full cycles per day [39], preventing the optimisation from exploiting short-lived price opportunities through excessive charge-discharge cycling.

B.3. Temporal Execution Framework

B.3.1. Rolling Horizon Structure

The optimisation model operates on a 15-minute temporal resolution ($\Delta t = 0.25$ h). Each optimisation window is defined over a 48-hour horizon consisting of 192 Program Time Units (PTUs):

$$T = \{0, 1, \dots, 191\} \tag{B.19}$$

Only the first 96 PTUs (24 hours) are committed and recorded:

$$T^{\text{com}} = \{0, 1, \dots, 95\} \tag{B.20}$$

The remaining horizon functions as a look-ahead period used to reduce end-of-horizon distortions.

B.3.2. Limited Foresight Implementation

Prior to each optimisation window, two price series are modified to reflect realistic operator information constraints.

Imbalance price truncation Real-time imbalance prices are not perfectly predictable. To avoid granting the optimiser unrealistic foresight, imbalance prices are assumed to be known only one hour ahead ($T^{\text{foresight}} = 4$ PTUs).

Beyond the one-hour foresight horizon, future imbalance prices are approximated using the average imbalance price observed during the preceding foresight window.

This prevents the optimiser from exploiting short-lived imbalance price spikes that would not be foreseeable in real-time settlement.

aFRR block-average preprocessing Because reserve bids must remain constant within each four-hour procurement block (Section B.2.7), market participants cannot respond to price fluctuations occurring within the block after the bid has been submitted. To reflect this, all aFRR prices within a procurement block are replaced by the corresponding block-average price before optimisation, so that the optimiser evaluates reserve opportunities using the same information that would be available at the time of bid submission, rather than exploiting within-block price variation that could not be acted upon in practice.

B.3.3. State Carry-Over Logic

To ensure continuity between consecutive optimisation windows, the initial State-of-Charge (SoC) of each window is set equal to the final committed SoC of the preceding window:

$$s_{w+1}^{\text{init}} = s_{95,w} \tag{B.21}$$

where $s_{95,w}$ denotes the SoC at the final PTU of the committed 24-hour horizon. This prevents artificial resets of the battery state between optimisation windows.

B.3.4. Annualisation Procedure

The simulation evaluates a sequence of rolling optimisation windows rather than a complete calendar year. To obtain annual performance metrics, aggregated revenues, costs, and degradation outcomes are scaled according to the fraction of the year represented by the successfully solved windows:

$$f^{\text{ann}} = \frac{n_{windows} \cdot 96 \cdot 0.25}{8760} \tag{B.22}$$

where $n_{windows}$ denotes the number of successfully solved windows. Annual performance indicators are obtained by dividing aggregated simulation outputs by f^{ann}. This annualisation fraction is reused without modification in the financial post-processing stage (Section B.8).

B.4. Uncertainty Engine

This section provides the exact implementation of uncertainty modelling.

B.4.1. Data Acquisition and Preprocessing

This section documents the data acquisition and preprocessing procedures used to construct the market dataset underlying the policy testbed. All market data are based on the Dutch bidding zone for the calendar year 2024 and are converted to a common annual index of 35,040 Program Time Units (PTUs).

Data Sources

The benchmark model combines historical market information from multiple Dutch electricity market platforms. Table B.3 summarises the datasets used throughout the study.

Table B.3: Market datasets used in the benchmark model.

Dataset	Resolution	Purpose
Day-Ahead prices	1 hour	Energy arbitrage
aFRR capacity prices	4-hour blocks	Reserve capacity revenues
Imbalance prices	15 minutes	Imbalance settlement
Regulation state	15 minutes	Directional constraints
Imbalance volumes	15 minutes	Market participation limits

Day-Ahead price data were obtained from the ENTSO-E Transparency Platform, while balancing market information was obtained from publicly available TenneT transparency datasets [27, 60].

Resolution Harmonisation
The optimisation model operates on a 15-minute resolution. Consequently, all datasets were transformed to a common PTU-based time index.

Day-Ahead prices were originally available at hourly resolution and were disaggregated by assigning each hourly value to its four constituent PTUs. aFRR capacity prices were available at four-hour block resolution and were therefore assigned as constant values within each corresponding block. Imbalance prices, imbalance volumes, and regulation state signals were already available at PTU resolution and required no further temporal disaggregation.

Data Cleaning
All datasets were aligned to a common annual timestamp index. Missing observations were handled through forward filling, ensuring a continuous annual time series suitable for rolling-horizon optimisation. Where no previous observation was available, zero-filling was applied.

Price Processing
Prior to scenario generation, Day-Ahead prices were clipped at the 1st and 99th percentiles. This procedure follows Verhagen, Hu, and Harmsen [64] and reduces the influence of rare extreme observations that could otherwise dominate simulated revenue outcomes. The clipping procedure preserves structural market spreads while limiting the impact of exceptional price spikes.

B.4.2. Revenue Derating Factors

Three derating factors are applied to all revenue streams prior to Monte Carlo scenario scaling. These adjustments account for systematic gaps between observed market prices and the revenues a real-world battery operator can capture, reflecting forecast error, execution risk, and market-structural constraints.

Table B.4: Revenue derating factors applied during scenario generation.

Stream	Symbol	Value	Rationale
Day-Ahead	δ^{DA}	0.90	Forecast imprecision and execution risk
aFRR capacity	δ^{aFRR}	0.70	Bid-stack position
Imbalance	δ^{imb}	0.50	Ex-post settlement uncertainty

Day-Ahead market The DA spread is compressed symmetrically around the annual mean:

$$\tilde{\pi}_t^{\mathrm{DA}} = \bar{\pi}^{\mathrm{DA}} + \underbrace{vol \cdot sat^{-0.5} \cdot \delta^{\mathrm{DA}}}_{\text{effective spread factor}} \cdot \left(\pi_t^{\mathrm{DA}} - \bar{\pi}^{\mathrm{DA}}\right) \tag{B.23}$$

The $sat^{-0.5}$ term captures the inverse relationship between market saturation and price spread: as storage deployment increases, arbitrage spreads compress. The DA floor is bounded below at -50 EUR/MWh to prevent physically implausible prices.

aFRR capacity market aFRR revenues are adjusted using the same volatility and saturation factors. In addition, a derating factor of $\delta^{\mathrm{aFRR}} = 0.70$ is applied to reflect the fact that a battery may not consistently occupy the most favourable position in the capacity auction merit order.

$$\tilde{\pi}_t^{\mathrm{aFRR}} = \pi_t^{\mathrm{aFRR}} \cdot vol \cdot sat^{-0.5} \cdot \delta^{\mathrm{aFRR}} \tag{B.24}$$

Imbalance market Imbalance revenues are adjusted for volatility and forecast uncertainty only. Unlike the Day-Ahead and aFRR markets, no saturation factor is applied because imbalance prices are assumed to be driven primarily by system imbalances rather than by the level of battery deployment.

$$\tilde{\pi}_t^{\mathrm{imb}} = \pi_t^{\mathrm{imb}} \cdot vol \cdot \delta^{\mathrm{imb}} \tag{B.25}$$

Regulation State Processing
The national regulation state signal was processed prior to optimisation. Long (-1), short $(+1)$, and neutral (0) states were preserved. Dual-pricing intervals were conservatively mapped to a short system state $(+1)$.

The processed regulation state determines the direction in which imbalance participation is permitted. The battery is only allowed to take positions that support the prevailing system need, ensuring consistency with Dutch imbalance settlement principles.

Market Participation Limits
As established in the MILP formulation (Section B.2.4), imbalance participation is limited to a 20% share of observed market volume to preserve the price-taker assumption.

B.5. Curtailment Scenario Generation

This section documents the procedure by which congestion events are generated for each Monte Carlo run. Curtailment timing, directionality, and frequency are drawn stochastically from distributions calibrated to the stylised urban and rural congestion archetypes defined in Section B.6.

B.5.1. Directional Weight Decomposition

The CMI weight array w_t encodes, at each PTU, which dispatch direction is grid-supportive, using the sign convention introduced in Section B.6: $w_t > 0$ indicates that charging is grid-supportive (generation-driven congestion), while $w_t < 0$ indicates that discharging is grid-supportive (demand-driven congestion). To construct probability distributions for event sampling, the signed weight array is decomposed into two non-negative components:

$$w_t^{\text{dis}} = \max(w_t, 0) \tag{B.26}$$

$$w_t^{\text{ch}} = \max(-w_t, 0) \tag{B.27}$$

w_t^{dis} is the weight used to time discharge curtailment: it is largest where charging is grid-supportive. w_t^{ch} is the weight used to time charge curtailment, largest where discharging is grid-supportive. PTUs with larger weights are more likely to be selected as congestion events, concentrating curtailment in periods that are most relevant from a congestion-management perspective rather than distributing it uniformly across the year.

For standard curtailment contracts, restrictions are concentrated in periods where battery behaviour would worsen congestion. For CSC, the logic is reversed: activations are concentrated in periods where battery dispatch can actively relieve congestion. Consequently, forced-discharge instructions are weighted by w_t^{ch} and forced-charge instructions by w_t^{dis}.

B.5.2. Activation Frequency and PTU Selection

For each Monte Carlo run, a curtailment probability p^{curt} is drawn from discrete set of predefined probability levels, representing different levels of DSO curtailment frequency.

Within each optimisation window, the total number of curtailed PTUs is determined as

$$n^{\text{curt}} = \lfloor p^{\text{curt}} |T^{\text{com}}| \rfloor, \tag{B.28}$$

where $|T^{\text{com}}|$ denotes the number of PTUs in the committed horizon. The selected PTUs are then sampled without replacement from T^{com} using the directional weighting arrays defined earlier. When timing bias equals one, PTUs are sampled proportionally to the congestion weighting arrays; otherwise, they are sampled uniformly at random. This weighted sampling concentrates curtailment during congestion-critical periods while preserving stochastic variation across Monte Carlo runs.

B.5.3. Reproducibility and Scenario Reuse

To ensure fair comparison between contractual arrangements, each Monte Carlo scenario realisation, comprising both market conditions and curtailment timing, was generated once per run and reused identically across all evaluated contracts. A deterministic random seed guaranteed this reproducibility: identical scenario parameters always generated identical curtailment patterns, regardless of which contract was subsequently evaluated. Consequently, differences in financial, technical, and alignment outcomes between contracts reflect contractual design rather than incidental differences in the underlying scenario. Curtailment timing and intensity were drawn from a structured set of parameter combinations (curtailment probability and timing bias), while market volatility and saturation were sampled independently and continuously for every run, ensuring that the 200 Monte Carlo realisations per location still span a wide range of distinct market conditions.

B.6. Congestion Management Indicator (CMI) Formalisation

The Congestion Management Indicator is represented as a deterministic time-varying weighting signal $w_t \in [-1, 1]$ assigned to every Program Time Unit (PTU) throughout the simulation horizon. Each weight reflects the relative desirability of charging or discharging at that moment from a congestion-management perspective.

The CMI signal is computed once before optimisation begins and remains fixed; it is not visible to the dispatch optimiser and therefore does not influence operational decisions. It is used exclusively for ex-post evaluation.

B.6.1. Congestion Archetypes and Weight Profiles

Two stylised congestion environments are modelled:

- **Urban (demand-driven):** congestion arises from evening demand peaks. Weights are derived from the temporal structure of TenneT's time-of-use transmission tariff [59]. Discharge during peak-demand hours is grid-supportive.
- **Rural (generation-driven):** congestion arises from midday solar surplus. Weights reflect the ACM taxonomy and Dutch solar irradiance profiles [5]. Charging during solar-peak hours is grid-supportive.

The sign convention is as follows. Negative weights ($w_t < 0$) indicate that *discharging* is grid-supportive at time t; positive weights ($w_t > 0$) indicate that *charging* is grid-supportive. The magnitude of each weight reflects the relative importance of that interval for congestion relief. The full weight table is provided in Table B.5 at the end of this appendix.

B.6.2. Alignment Score

The alignment score Ω quantifies ex-post how closely battery dispatch coincides with congestion-supportive behaviour over the full simulation horizon. Net battery dispatch is:

$$p_t^{net} = p_t^{\mathrm{dis}} - p_t^{\mathrm{ch}} \tag{B.29}$$

and the alignment score is:

$$\Omega = \frac{\sum_{t\in T^{\mathrm{com}}} p_t^{net} \cdot w_t \cdot \Delta t}{\sum_{t\in T^{\mathrm{com}}} |w_t|} \tag{B.30}$$

Sign convention. The sign of Ω is counterintuitive but consistent. In the urban regime, $w_t < 0$ during discharge-helpful evening hours. Helpful discharge (large positive p_t^{dis}) therefore produces a negative numerator contribution, driving Ω toward more negative values. In the rural regime, $w_t > 0$ during charge-helpful midday hours. Helpful charging (large positive p_t^{ch}) produces $p_t^{net} = -p_t^{\mathrm{ch}}$, again a negative numerator contribution.

The universal interpretation holds for both archetypes: More negative Ω = stronger congestion-supportive alignment. Positive Ω = net misalignment with congestion needs.

B.6.3. Baseline-Referenced Improvement

Because the unconstrained Baseline already produces some natural alignment (see Section 4.2.1), contracts are evaluated on their incremental alignment contribution rather than on their raw Ω score. For each matched run pair:

$$\Delta\Omega_{\pi,r} = \Omega_r^{Baseline} - \Omega_r^{\pi} \tag{B.31}$$

where r is the run index and π is the policy. Positive $\Delta\Omega$ means the contract produces stronger congestion-supportive alignment than unconstrained market dispatch under the same market realisation. By construction, $\Delta\Omega^{Baseline} = 0$ for all runs. This metric is the primary alignment outcome reported in Chapter 4.

B.6.4. Idle Fraction

The alignment score Ω does not distinguish between a battery that actively dispatches in a congestion-helpful direction and one that simply remains inactive during congestion-critical periods. To separate these effects, an idle-fraction metric is computed:

$$\Omega_{idle} = \frac{\sum_t |w_t|\, idle_t}{\sum_t |w_t|} \tag{B.32}$$

where $idle_t$ equals 1 when the battery is inactive during PTU t and 0 otherwise. A high Ω_{idle} combined with a strong Ω score indicates that alignment arises primarily from strategic inactivity rather than directed dispatch. This distinction is discussed in Section 4.2.1 for the Baseline and in Section 4.2.2 for the contracted variants.

B.6.5. CMI Weight Profiles

The CMI signal is constructed by defining a representative daily weight profile for each calendar month and tiling it across the 96 PTUs of each day in that month. Both archetypes incorporate seasonal variation: the urban profile scales the evening demand peak by season, while the rural profile scales the midday solar congestion window according to seasonal PV generation. Tables B.5 and B.6 present the resulting weight profiles.

Table B.5: Urban CMI weight profile by time window and season. Constructed from the urban demand-congestion archetype described in Section B.6.5. Negative weights indicate discharge is grid-supportive; positive weights indicate charging is grid-supportive. Values after normalisation to $[-1, +1]$.

Time window	Winter	Shoulder	Summer	Congestion interpretation
00:00–06:00	$+0.20$	$+0.20$	$+0.20$	Overnight off-peak; mild charge-supportive signal reflecting periods of relatively low demand.
06:00–08:00	-0.35	-0.35	-0.35	Morning demand ramp; moderate discharge-supportive signal.
08:00–17:00	-0.10	-0.10	-0.10	Daytime period; weak discharge-supportive signal.
17:00–22:00	-1.00	-0.80	-0.55	Evening demand peak; strongest discharge-supportive signal. Seasonal variation reflects the stronger winter demand peak and reduced summer loading.
22:00–23:00	-0.45	-0.30	-0.20	Late-evening transition as demand declines.
23:00–24:00	$+0.10$	$+0.10$	$+0.10$	Overnight low-demand period.

Table B.6: Rural CMI weight profile by time window and season. Derived from the ACM *invoedingscongestie* taxonomy and Dutch solar PV irradiance profiles. Positive weights indicate charging is grid-supportive. Base weights (06:00–10:00 column) are not seasonally scaled. All other solar-window weights equal base $\times$ seasonal scale factor. Values after normalisation to $[-1, +1]$, anchored at summer midday peak $= +1.00$.

Time window	Summer	Shoulder	Winter	Congestion interpretation
00:00–06:00	0.00	0.00	0.00	Night; no solar generation, no congestion signal.
06:00–10:00	$+0.25$	$+0.25$	$+0.25$	Morning solar ramp-up; mild charge-helpful signal to create SoC headroom. Not seasonally scaled.
10:00–11:00	$+0.60$	$+0.42$	$+0.24$	Solar congestion flanks; base = 0.60.
11:00–14:00	$+0.80$	$+0.56$	$+0.32$	Solar congestion shoulder; base = 0.80.
12:00–13:00	$+1.00$	$+0.70$	$+0.40$	Peak solar irradiance; strongest. Base = 1.00.
14:00–17:00	$+0.60$	$+0.42$	$+0.24$	Solar congestion flanks; base = 0.60.
17:00–19:00	$+0.15$	$+0.11$	$+0.06$	Solar ramp-down; residual mild congestion signal. Base = 0.15.
19:00–24:00	0.00	0.00	0.00	Evening and night; no solar congestion.

B.7. Detailed Degradation Formulation

This section provides the detailed technical specification of the battery degradation sub-model used within the BESS Policy Testbed. While Section 3.3.6 presents the conceptual rationale and modelling assumptions, the present section documents the mathematical formulation, calibration procedure, and linearisation approach required for full computational reproducibility. The level of detail here is deliberate: the null-degradation finding across five of the six contracts, and the Volume Bound contract as its sole exception, rests directly on the formulation below, and is only interrogable if that formulation is fully specified.

B.7.1. Calibration of the Degradation Curve

The cycling degradation model is calibrated to reflect the performance characteristics of contemporary utility-scale LFP battery systems [52]. The calibration rests on three principal assumptions:

- A full cycle at 100% Depth of Discharge (DoD) corresponds to approximately 8,000 equivalent full cycles before reaching End-of-Life (EOL) [52].
- End-of-Life is defined as a 20% reduction in usable energy capacity, corresponding to a State-of-Health of 80% [52, 40].
- Total battery capital expenditure is assumed to equal 6 million euros, of which 66% is attributable to the replaceable battery stack component [17].

Together these assumptions fix the relationship between physical degradation and its corresponding economic value.

B.7.2. Degradation Function

Cycling degradation is represented through a convex power-law relationship between Depth of Discharge and incremental capacity loss:

$$Damage \propto DoD^{\beta}, \qquad \beta = 1.79 \tag{B.33}$$

where the degradation exponent β is taken from Terkes et al. [61] and Xu et al. [67]. The convexity of this relationship is the mechanism behind the degradation results: a single full-depth cycle inflicts roughly eighteen times the damage of a shallow 0.2-DoD cycle, so contracts that induce deep cycling are penalised disproportionately.

Anchoring the curve to the calibration above, the incremental fractional capacity loss at breakpoint i is:

$$\delta_i = \frac{0.20}{8000} \cdot \xi_i^{1.79} = 2.5 \times 10^{-5} \cdot \xi_i^{1.79} \tag{B.34}$$

B.7.3. MILP Implementation

Because the optimisation framework is formulated as a Mixed-Integer Linear Program (MILP), the convex degradation curve is approximated by piecewise-linear interpolation over the breakpoints, implemented with Special Ordered Set Type 2 (SOS2) variables [66]. The interpolation weights λ_i form a convex combination ($\sum_i \lambda_i = 1, \lambda_i \geq 0$) subject to the SOS2 condition that at most two adjacent weights are non-zero, which enforces interpolation along a single segment of the curve.

The effective DoD proxy ξ and the corresponding cyclic damage δ^{cyc} follow directly from the weights:

$$\xi = \sum_i \xi_i \lambda_i, \qquad \delta^{cyc} = \sum_i \delta_i \lambda_i \tag{B.35}$$

This representation captures the non-linear degradation response while preserving the linear structure required by the solver.

B.7.4. Damage Monetisation

Following the damage-based levelised degradation framework of Massat et al. [40], fractional capacity loss is converted into an equivalent monetary penalty. The cost attributed to breakpoint i is the fractional loss scaled by the replaceable stack value per unit of tolerable degradation:

$$\kappa_i = \delta_i \cdot \frac{0.66\, C^{\text{CAPEX}}}{0.20} \tag{B.36}$$

where C^{CAPEX} is total battery capital expenditure, $0.66\,C^{\mathrm{CAPEX}}$ the replaceable stack cost, and 0.20 the total fractional capacity loss tolerated before EOL. The degradation cost over an optimisation window is then:

$$\kappa^{\mathrm{win}} = n^{\mathrm{com}} \cdot \sum_i \kappa_i \lambda_i \tag{B.37}$$

where n^{com} is the number of days in the committed horizon. Because κ^{win} enters the objective function directly, degradation is internalised within every dispatch decision rather than evaluated post hoc.

B.7.5. Annual Degradation Accounting

Total battery ageing combines cycling and calendar degradation. Calendar ageing is represented through a constant annual degradation rate of 1% per year. Degradation accumulated within an optimisation window is therefore:

$$\delta^{\mathrm{win}} = n^{\mathrm{com}} \left(\sum_i \delta_i \lambda_i + \frac{0.01}{365} \right) \tag{B.38}$$

Window-level degradation is aggregated and rescaled to a full year by dividing by the fraction of the calendar year represented by the committed simulation horizon, f_{year}:

$$D^y = \frac{\sum \delta^{\mathrm{win}}}{f_{\mathrm{year}}} \tag{B.39}$$

where $f_{\mathrm{year}} \in (0, 1]$ is the committed simulation duration expressed as a fraction of 365 days. The projected State-of-Health trajectory is estimated by linear extrapolation of this constant annual rate:

$$\mathrm{SoH}_y = \max(0,\ 1 - D^y \cdot y) \tag{B.40}$$

The linear extrapolation is a deliberate simplification: it assumes a constant annual degradation rate rather than the mild deceleration typically observed empirically, and therefore yields a conservative (slightly pessimistic) long-run SoH estimate. This keeps long-term degradation accounting tractable within the annual simulation framework while remaining transparent about its assumptions.

B.8. Financial Post-Processing Documentation

This section documents the technical post-processing framework used to translate annual operational outcomes into long-term investment performance metrics. The objective of this section is to ensure full computational reproducibility of the financial evaluation framework presented in Section 3.4, while keeping the main methodology focused on the strategic rationale underlying the investment assessment.

The financial post-processing chain converts annual operational results into 15-year investment trajectories through four sequential steps:

1. annualisation of operational outcomes;
2. scenario-affinity weighting;
3. lifetime trajectory interpolation;
4. cashflow projection and investment evaluation.

B.8.1. Annualisation and Scaling

Operational revenues, costs, and degradation outcomes are accumulated across all optimisation windows within a simulation run and scaled to annual estimates using the annualisation fraction f^{ann} defined in Section B.3 (Equation B.22):

$$x^{annual} = \frac{\sum_w x_w^{window}}{f^{\mathrm{ann}}} \tag{B.41}$$

This scaling procedure ensures that annual performance metrics remain comparable across simulation runs, regardless of minor variations in solver convergence or window completion.

B.8.2. Scenario Affinity Weighting

Future market conditions are represented through four canonical scenario families defined by combinations of market volatility and market saturation.

Rather than assigning each Monte Carlo run to a single discrete scenario, the framework adopts a continuous affinity approach. Each simulation run receives a degree of membership for all scenario families based on the sampled volatility (vol) and saturation (sat) parameters.

The normalised coordinates are calculated as:

$$v = \frac{vol - 0.70}{0.60} \tag{B.42}$$

$$s = \frac{sat - 0.70}{0.60} \tag{B.43}$$

The resulting affinity scores are defined as:

$$A^{Sc1} = v(1 - s) \tag{B.44}$$

$$A^{Sc2} = vs \tag{B.45}$$

$$A^{Sc3} = (1 - v)(1 - s) \tag{B.46}$$

$$A^{Sc4} = (1 - v)s \tag{B.47}$$

where:

- Sc1 represents High Volatility / Low Saturation;
- Sc2 represents High Volatility / High Saturation;
- Sc3 represents Low Volatility / Low Saturation;
- Sc4 represents Low Volatility / High Saturation.

Scenario-conditional performance is subsequently calculated using affinity-weighted averages:

$$\bar{\pi}_y^{sc} = \frac{\sum_r (A_r^{sc})^{\alpha_y} \cdot \pi_r^{annual}}{\sum_r (A_r^{sc})^{\alpha_y}} \tag{B.48}$$

where α_y is a time-dependent sharpening parameter and π_r^{annual} denotes annual profit for run r.

B.8.3. Lifetime Trajectory Interpolation

The Monte Carlo framework generates annual performance estimates rather than complete 15-year trajectories. To represent long-term project evolution, performance is evaluated at four milestone years:

$$Y^{milestone} = \{1, 3, 6, 9\} \tag{B.49}$$

The scenario-affinity exponent is increased over time:

Table B.7: Scenario-affinity sharpening schedule.

Milestone Year	α_y
1	0
3	1
6	2
9	3

Year 1 assumes complete uncertainty, resulting in uniform weighting across all scenario families. As the project progresses, the weighting structure increasingly reflects the statistical characteristics of the sampled market regime.

Intermediate years and years beyond the final milestone are determined using linear interpolation between milestone values.

B.8.4. Cashflow Projection

Annual profit estimates are adjusted using the projected State-of-Health trajectory to reflect the capacity-dependent revenue reduction that accompanies battery ageing:

$$\mathrm{CF}_y = \bar{\pi}_y^{sc} \cdot \mathrm{SoH}_y - \mathrm{OPEX} \tag{B.50}$$

where $\bar{\pi}_y^{sc}$ denotes the scenario-affinity-weighted mean annual profit at milestone year y (net of degradation, grid, and DAFO costs as produced by the MILP), SoH_y denotes the projected State-of-Health at year y, and $\mathrm{OPEX} = 0.03 \cdot C^{\mathrm{CAPEX}}$ is fixed independently of battery health, reflecting that maintenance, insurance, and connection costs do not scale proportionally with capacity fade.

This formulation applies SoH derating at the profit level rather than the revenue level. A battery at 85% SoH can generate approximately 85% of year-1 operational value across all market streams simultaneously, while the fixed cost base remains unchanged.

These annual cashflows form the basis for Net Present Value and Internal Rate of Return calculations reported in the main analysis.

B.8.5. Statistical Interpretation

Each policy configuration is evaluated using $N = 200$ Monte Carlo runs. Financial outcomes are reported per scenario family using the affinity-weighted averaging procedure defined above (Section B.8.2). Because the continuous affinity approach assigns a non-zero weight to every run for every scenario family, all 200 runs contribute to each scenario-family estimate.

The sharpening exponent α_y increases across the projection horizon, assigning progressively greater weight to runs whose sampled volatility and saturation parameters are most consistent with a given scenario family. However, no run is excluded from the calculation. Scenario-family outcomes therefore remain weighted summaries of the full Monte Carlo distribution rather than estimates derived from separate scenario-specific subsamples.

The scenario families should consequently be interpreted as idealised regions within a continuous uncertainty space rather than as mutually exclusive future states. Differences between scenario-family outcomes reflect the influence of volatility and saturation on battery revenues under the affinity-weighting scheme, rather than the effects of discrete scenario assignment.

B.9. Solver settings

The optimisation problem is solved using Gurobi 11.0 with a relative optimality gap of 5% and a time limit of three seconds per optimisation window. These settings represent a deliberate trade-off between solution quality and computational tractability across the full simulation framework, which evaluates multiple policy variants over annual rolling-horizon simulations and Monte Carlo market realisations.

C

Contract Implementation Details

This section provides the exact implementation of each contractual arrangement within the Policy Testbed. All agreements are represented as additional constraint layers applied to the benchmark optimisation model described in Chapter 3. The objective is to ensure transparency and reproducibility of the policy simulations while keeping the main methodology focused on conceptual model design.

C.1. Alternative Transport Right 85% (ATR85)

ATR85 is implemented as a stochastic full-curtailment mechanism combined with a fixed tariff discount. The contractual 15% annual curtailment allowance is represented through a stochastic per-window activation rate rather than an annual cumulative counter.

C.1.1. PTU Selection Procedure

For each optimisation window, all PTUs flagged for either charge-side or discharge-side congestion are identified:

$$A = \{t : f_t^{\mathrm{dis}} = 1\} \cup \{t : f_t^{\mathrm{ch}} = 1\} \tag{C.1}$$

A maximum of 14 PTUs is selected within each committed horizon of 96 PTUs, corresponding to approximately 15% of the operational period. If more than 14 PTUs are flagged, a deterministic subset is selected using a reproducible random seed derived from the activation pattern.

C.1.2. Curtailment Constraints

During selected ATR85 activation periods, all battery operation is suspended. Charging, discharging, balancing participation, and market commitments are forced to zero:

$$p_t^{\mathrm{dis}} + c_t^{\uparrow} + p_t^{\mathrm{imb,up}} = 0 \tag{C.2}$$

$$p_t^{\mathrm{ch}} + c_t^{\downarrow} + p_t^{\mathrm{imb,dn}} = 0 \tag{C.3}$$

$$p_t^{\mathrm{DA}} = 0 \tag{C.4}$$

The shutdown is bidirectional and independent of the congestion direction that triggered the activation.

C.1.3. Tariff Discount

ATR85 receives a fixed tariff reduction equal to 65% of the annual grid fee. The remuneration is applied independently of realised activation frequency.

C.2. Capacity Limiting Contract (CBC)

CBC is implemented as a passive congestion-management mechanism. Rather than shutting down battery operation, the agreement limits the magnitude of grid usage during congestion periods.

C.2.1. Capacity Ceiling Constraints

During congestion events, both charging and discharging are restricted to a maximum transport capacity of 3 MW:

$$p_t^{\text{dis}} + c_t^{\uparrow} + p_t^{\text{imb,up}} \leq P^{\text{CBC}} \tag{C.5}$$

$$p_t^{\text{ch}} + c_t^{\downarrow} + p_t^{\text{imb,dn}} \leq P^{\text{CBC}} \tag{C.6}$$

where:

$$P^{\text{CBC}} = 3 \text{ MW} \tag{C.7}$$

The battery remains free to optimise within the remaining capacity envelope.

C.2.2. Compensation Formulation

Compensation is linked to curtailed transport capacity rather than realised energy losses:

$$R^{\text{CBC}} = (P^{\text{rated}} - P^{\text{CBC}}) \cdot r^{\text{comp}} \cdot \Delta t \cdot N^{\text{flagged}} \tag{C.8}$$

where N^{flagged} denotes the number of congestion-flagged PTUs and r^{comp} represents the contractual compensation rate.

C.3. Capacity Steering Contract (CSC)

The Capacity Steering Contract (CSC) represents an active congestion-management arrangement in which the Distribution System Operator (DSO) may issue dispatch instructions to participating Battery Energy Storage Systems (BESS). Unlike passive congestion-management contracts that restrict access rights, CSC grants the DSO limited operational control over battery dispatch during activated congestion periods.

C.3.1. Setpoint Tracking

Two mutually exclusive activation types are distinguished:

- Forced discharge activation ($f_t^{\text{CSC,dis}} = 1$), with target dispatch $\hat{p}_t = +5$ MW;
- Forced charge activation ($f_t^{\text{CSC,ch}} = 1$), with target dispatch $\hat{p}_t = -5$ MW.

During activated PTUs, battery dispatch is governed by the setpoint-tracking constraint

$$\left(p_t^{\text{dis}} - p_t^{\text{ch}}\right) + s_t^{+} - s_t^{-} = \hat{p}_t \tag{C.9}$$

where $s_t^{+}, s_t^{-} \geq 0$ are slack variables representing deviations from the requested setpoint.

The constraint is applied only when either $f_t^{\text{CSC,dis}} = 1$ or $f_t^{\text{CSC,ch}} = 1$. During non-activated PTUs, the battery operates according to the unconstrained market optimisation model.

C.3.2. Power Constraint During Forced Discharge

The general physical capacity constraint of Appendix B.2 already applies at every PTU regardless of contractual arrangement; during forced-discharge activations it therefore continues to cap the combined discharge commitment from physical dispatch, upward aFRR reservation, and upward imbalance participation at P^{rated}. No CSC-specific power constraint is introduced.

C.3.3. Activation Logic

CSC activations are generated using CMI-weighted sampling. Activation timing is therefore concentrated in periods associated with elevated congestion conditions.

The directional distribution of activations differs by network regime, as shown in Table C.1.

Table C.1: CSC activation direction split by network regime.

Regime	Forced discharge	Forced charge
Urban (demand congestion)	70%	30%
Rural (generation congestion)	10%	90%

C.3.4. Remuneration and Penalties

Activated dispatch instructions receive an activation payment of

$$r^{\mathrm{act}} = 150\ €/\mathrm{MWh} \tag{C.10}$$

applied to the requested dispatch volume for both charge and discharge activations.

Deviations from the requested setpoint are penalised according to

$$r^{\mathrm{pen}} = 400\ €/\mathrm{MWh} \tag{C.11}$$

in the reference configuration, corresponding to the value proposed in the ACM consultation.

The resulting CSC remuneration is incorporated directly into the optimisation objective as activation revenues minus deviation penalties. The activation payment rate and penalty rate are treated as policy parameters and are varied in sensitivity analysis. In particular, penalty levels of 150 and 250 €/MWh are evaluated.

C.4. Time-Block Transport Right (TBTR)

The Time-Block Transport Right (TBTR) represents an existing Dutch flexible connection arrangement in which access restrictions are imposed during predetermined time windows. Unlike congestion-responsive arrangements, the restrictions are independent of realised network conditions and are applied according to a fixed daily schedule.

C.4.1. Fixed Shutdown Windows

Battery operation is fully suspended during the following daily windows:

- 07:00 – 09:00
- 17:00 – 19:00

Together, these windows correspond to four restricted hours per day. During restricted periods, all charging, discharging, reserve capacity provision, imbalance participation, and Day-Ahead nominations are prohibited:

$$p_t^{\mathrm{dis}} + c_t^{\uparrow} + p_t^{\mathrm{imb,up}} \leq 0 \qquad \forall t \in \mathcal{T}_{TBTR} \tag{C.12}$$

$$p_t^{\mathrm{ch}} + c_t^{\downarrow} + p_t^{\mathrm{imb,dn}} \leq 0 \qquad \forall t \in \mathcal{T}_{TBTR} \tag{C.13}$$

$$p_t^{\mathrm{DA}} = 0 \qquad \forall t \in \mathcal{T}_{TBTR} \tag{C.14}$$

where $\mathcal{T}_{TBTR}$ denotes the set of PTUs falling within the predefined restriction windows.

Unlike ATR85, these restrictions are independent of realised congestion conditions and remain identical throughout the simulation year.

C.4.2. Tariff Discount

The financial compensation mechanism follows the institutional *t/24* principle, whereby the annual transport tariff discount is proportional to the fraction of daily access rights surrendered by the connection holder:

$$\delta^{TBTR} = \frac{t_{restricted}}{24} \tag{C.15}$$

where $t_{restricted}$ denotes the total number of restricted hours per day.

For the reference configuration,

$$\delta^{TBTR} = \frac{4}{24} = 16.7\% \tag{C.16}$$

corresponding to the four daily shutdown hours imposed by the agreement. The resulting annual tariff reduction is therefore calculated as:

$$R^{TBTR} = \delta^{TBTR} \cdot F_{grid} \tag{C.17}$$

where F_{grid} denotes the annual grid connection tariff. The tariff discount is implemented as a fixed financial benefit independent of actual contract utilisation, congestion occurrence, or market conditions.

C.5. Volume Bound

The Volume Bound agreement represents a prospective contract design derived from the institutional discovery, operationalising the volume orientation property dimension identified.

C.5.1. Daily Energy Budget

Instead of limiting instantaneous power, the contract restricts total daily energy throughput independently per direction. For each day d in the committed optimisation window:

$$\sum_{t \in D_d} p_t^{\mathrm{dis}} \, \Delta t \leq E_{dis}^{\mathrm{VB}} \qquad \forall d \tag{C.18}$$

$$\sum_{t \in D_d} p_t^{\mathrm{ch}} \, \Delta t \leq E_{ch}^{\mathrm{VB}} \qquad \forall d \tag{C.19}$$

where D_d denotes the set of PTUs belonging to day d, and $E_{dis}^{\mathrm{VB}} = E_{ch}^{\mathrm{VB}} = 10$ MWh in the reference configuration. This budget corresponds to 50% of rated energy capacity per direction and permits approximately half an equivalent cycle per day while preserving substantial intraday scheduling flexibility. The budget resets at each daily boundary; only the committed horizon is constrained, while the look-ahead horizon operates without a budget constraint to avoid artificial end-of-horizon effects.

C.5.2. Power Envelope

A fixed instantaneous power envelope of $P_{env}^{\mathrm{VB}} = 6$ MW is applied in addition to the daily energy budget, constraining the combined physical dispatch and reserve capacity in both directions:

$$p_t^{\mathrm{dis}} + c_t^{\uparrow} + p_t^{\mathrm{imb,up}} \leq P_{env}^{\mathrm{VB}} \qquad \forall t \tag{C.20}$$

$$p_t^{\mathrm{ch}} + c_t^{\downarrow} + p_t^{\mathrm{imb,dn}} \leq P_{env}^{\mathrm{VB}} \qquad \forall t \tag{C.21}$$

This envelope prevents unrealistic short-duration full-rated dispatch that the energy budget constraint alone would not exclude. The power envelope applies at all PTUs regardless of congestion status.

In addition to the physical dispatch envelope, the Day-Ahead nomination is constrained to the same maximum level:

$$p_t^{\mathrm{DA}} \leq P_{env}^{\mathrm{VB}} \qquad \forall t \tag{C.22}$$

This constraint prevents the optimisation from scheduling Day-Ahead market positions that would be infeasible under the contract's instantaneous power restriction.

No compensation or tariff discount is associated with this agreement. The absence of financial transfers is a deliberate experimental design choice: it isolates the operational effect of energy-based restriction from any remuneration incentive, enabling the property-based analysis in the integrated evaluation to assess volume orientation independently of compensation structure.

C.6. Firmness Gradient

The Firmness Gradient represents a prospective contract design derived from the institutional analysis. The arrangement separates transport rights into a guaranteed firm component and a congestion-sensitive non-firm component. The underlying design objective is to investigate whether investor risk can be reduced through guaranteed minimum access rights rather than through financial compensation mechanisms.

C.6.1. Firm Capacity Floor

A guaranteed firm transport right is maintained at all times:

$$P^{\mathrm{firm}} = 3 \text{ MW} \tag{C.23}$$

This capacity remains available irrespective of congestion conditions and therefore provides the connection holder with a minimum level of operational certainty.

C.6.2. Non-Firm Tranche

The remaining transport capacity,

$$P^{NF} = P^{\text{rated}} - P^{\text{firm}}, \tag{C.24}$$

is classified as non-firm and may be withdrawn during congestion periods. Consequently, the available discharge capacity is defined as:

$$p_t^{\text{dis}} \leq \begin{cases} P^{\text{rated}}, & \text{if no congestion event is active} \\ P^{\text{firm}}, & \text{if a congestion event is active} \end{cases} \tag{C.25}$$

An equivalent constraint is applied to charging.

Consequently, the battery retains guaranteed access to at least 3 MW of transport capacity under all operating conditions, while the remaining 7 MW remains subject to congestion management interventions.

The firm floor is calibrated to preserve balancing-market participation throughout the contract period while allowing the DSO to retain operational flexibility over the non-firm portion of the connection. This calibration reflects the institutional premise that investors may be more concerned with preserving access to critical value streams than with maximising unrestricted connection capacity.

No compensation or tariff discount is associated with this arrangement. This is a deliberate experimental design choice. By isolating the effect of guaranteed access rights from financial transfers, the contract enables the Policy Testbed to evaluate the contribution of firmness itself to investor risk reduction and financial viability.

D

Supplementary Results

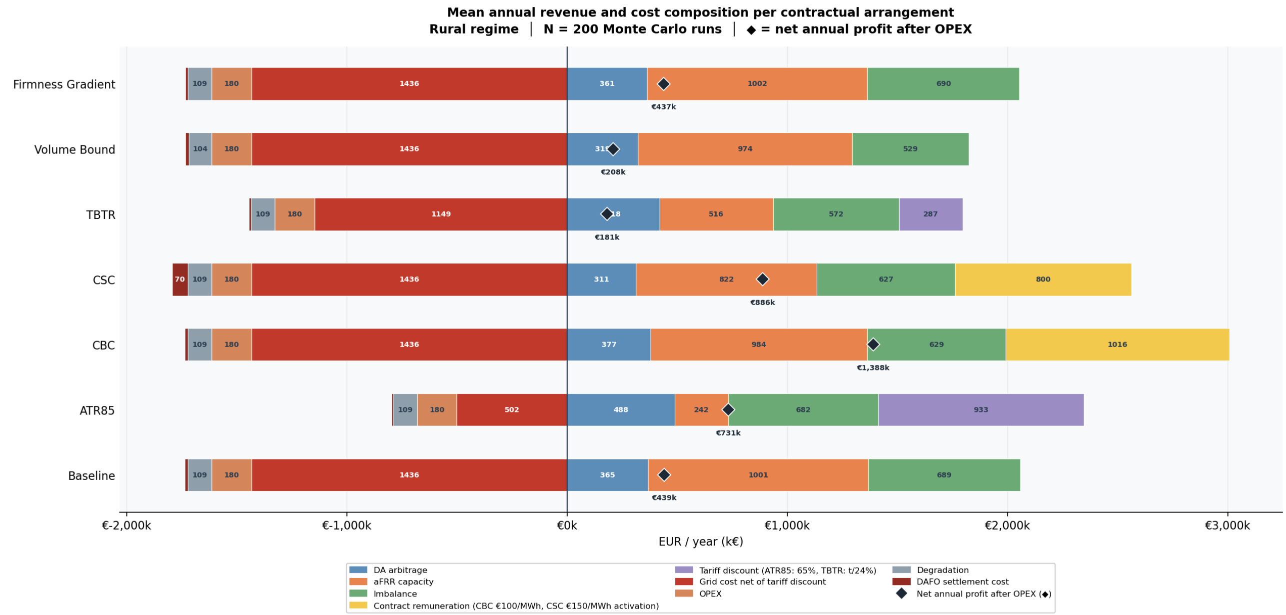

Figure D.1: Revenue and cost decomposition per contract, rural congestion regime. Bars to the right show market revenues (DA arbitrage, aFRR capacity, imbalance) and contract remuneration; bars to the left show costs (degradation, OPEX, grid fee net of any tariff discount). The diamond marker (◆) indicates mean net profit. Light-purple segments mark the tariff discount under ATR85 (65%) and TBTR ($t/24 \approx 16.7\%$). Urban counterpart in Figure 4.1.

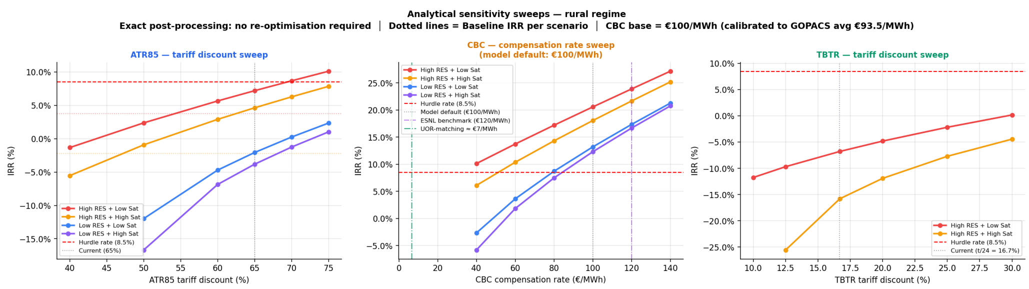

Figure D.2: Analytical sensitivity sweeps, rural congestion regime. Three contracts are shown: ATR85 (tariff discount), CBC (compensation rate), TBTR (tariff discount). Dotted lines mark scenario-conditional Baseline IRRs. Red dashed line: 8.5% hurdle rate. The green dash-dot line on the CBC panel marks the UOR-matching rate (€6.5/MWh). Urban counterpart in Figure 4.3.

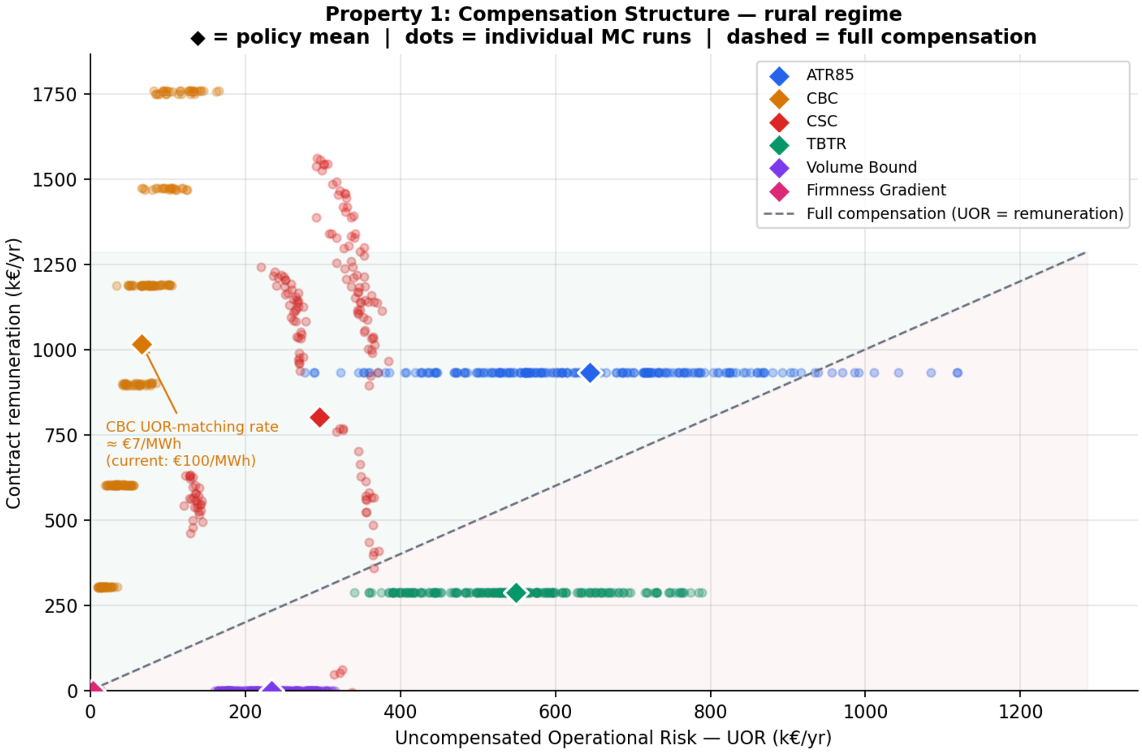

Figure D.3: Compensation structure, rural congestion regime. Mean UOR versus mean annual remuneration per contract (◆ = policy mean; semi-transparent dots = individual Monte Carlo runs). Dashed diagonal: full compensation line (UOR = remuneration). Green shading: overcompensated region; red shading: undercompensated region. Urban counterpart in Figure 4.4.

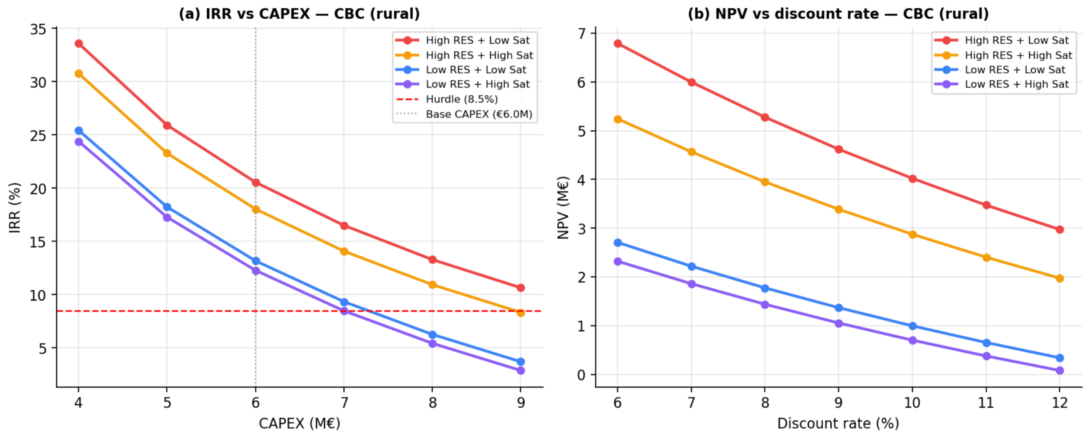

Figure D.4: CBC investment sensitivity, rural congestion regime. Panel (a): IRR as a function of CAPEX by scenario family. Panel (b): NPV as a function of discount rate. Red dashed line: 8.5% hurdle rate. Dotted vertical line: base assumptions (€6.0M CAPEX; 6% discount rate). Urban counterpart in Figure 4.2.

Table D.1: Multi-criteria scorecard, rural congestion regime ($N = 200$ Monte Carlo runs). Each cell reports the proportion of runs in which the contract outperforms the run-matched unconstrained Baseline on the stated dimension. Annual profit: net annual cashflow exceeds Baseline. Lifetime: implied SoH trajectory reaches End-of-Life later than Baseline. CMI: realised $\Delta\Omega > 0$. Bankable: scenario-family IRR exceeds 8.5% hurdle rate. Urban counterpart in Table 4.12.

Contract	Annual profit	Battery lifetime	CMI alignment	Bankable scenarios
ATR85	95%	0%	100%	0%
CBC	100%	78%	30%	100%
CSC	81%	2%	100%	50%
TBTR	0%	0%	100%	0%
Volume Bound	0%	100%	0%	0%
Firmness Gradient	22%	56%	48%	0%

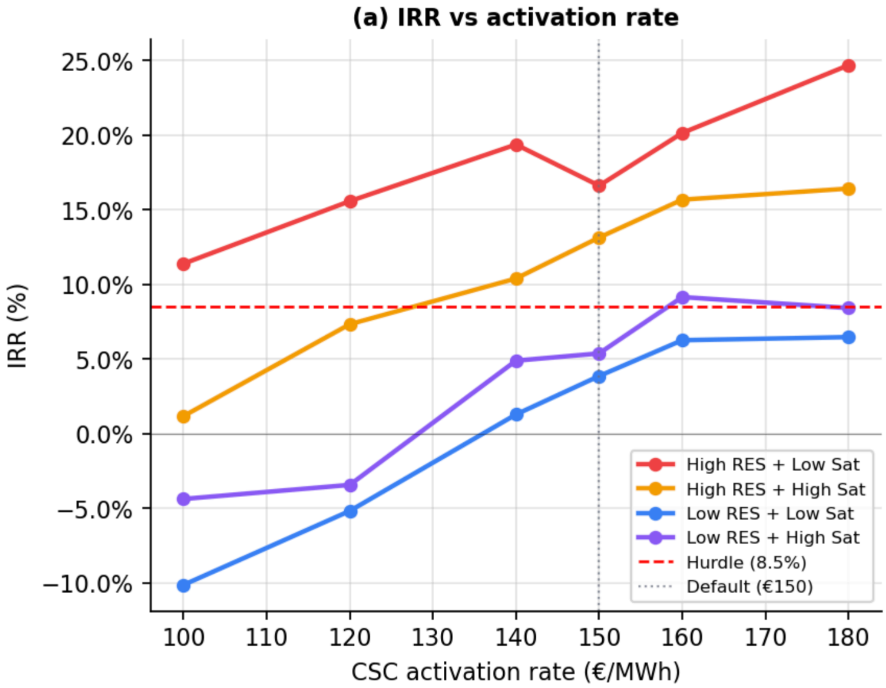

Figure D.5: CSC activation rate sensitivity, rural regime: project IRR by scenario family as a function of activation rate (€/MWh). Red dashed line: 8.5% hurdle rate. Dotted vertical line: €150/MWh default.

E
Pseudocode

This appendix summarises the computational framework at a methodological level: the Monte Carlo simulation and MILP dispatch model, and the post-processing pipeline that turns its output into the results reported in the thesis. Where this appendix abstracts from implementation-level detail, the full mathematical formulation, notation, and parameter values are provided in the Technical Model Documentation (Appendix B) and Contract Implementation Details (Appendix C); this chapter is intended as a self-contained algorithmic overview that can be read independently of those appendices. The framework is implemented as a modular Python simulation environment; this appendix describes the method, not the software architecture.

E.1. Overall pipeline

Algorithm 1 Monte Carlo framework

Data: Annual market dataset, location regime (urban/rural)
Result: Annualised operational and financial results per (run, policy)

Construct the location-specific congestion weight profile **for** *run* $= 1$ **to** N_{runs} **do**
  Sample a Monte Carlo scenario and generate the corresponding annual market and curtailment realisation (Algorithm 2) **foreach** *policy* ∈ *{Baseline, ATR85, CBC, CSC, TBTR, Volume Bound, Firmness Gradient}* **do**
    Solve the rolling-horizon dispatch optimisation (Algorithm 3) Store the annualised operational and financial outputs

Classify each run into a scenario family based on its sampled parameters Compute lifecycle NPV and IRR per (policy, scenario family) (Section E.4)

E.2. Scenario generation

Each Monte Carlo realisation samples a market scenario characterised by five parameters: a volatility factor and a market saturation factor (jointly scaling Day-Ahead, aFRR, and imbalance prices), a curtailment probability, and two independent congestion-targeting quality parameters, $\beta_{\text{dis}}, \beta_{\text{ch}} \in [0, 1]$, governing the timing of discharge- and charge-side curtailment respectively. Each parameter interpolates between random DSO activation timing ($\beta = 0$) and activation fully targeted at the most congestion-critical hours ($\beta = 1$), using a signed Congestion Mitigation Index (CMI) weight profile derived from literature. CSC activation timing reuses the discharge-side parameter β_{dis} rather than introducing a separate timing parameter, reflecting its role as an explicit congestion-response instrument. The battery dispatch itself never observes these CMI weights.

Algorithm 2 Scenario generation

Input: Location regime, congestion weight profile
Output: Annual market and curtailment scenario

Sample $(v, s, p_{\text{curt}}, \beta_{\text{dis}}, \beta_{\text{ch}})$ Scale Day-Ahead, aFRR, and imbalance prices according to (v, s) Generate the annual discharge curtailment schedule, targeted by β_{dis} Generate the annual charge curtailment schedule, targeted by β_{ch} Generate the annual CSC discharge-activation schedule, targeted by β_{dis} Generate the annual CSC charge-activation schedule, targeted by β_{dis} **return** *annual scenario data*

E.3. Rolling-horizon MILP dispatch

Algorithm 3 Rolling-horizon MILP dispatch, one run and one policy

```
Input: Annual scenario data, connection agreement (policy)
Output: Annualised profit, degradation, and ex-post CMI alignment score
Initialise State-of-Charge  for each optimisation window (48h look-ahead, 24h commitment) do
    Construct the MILP objective and constraints for the current window  Apply the selected connection agree-
      ment, including its policy-specific operational constraints and remuneration terms  Solve the optimisation
      if infeasible then
        Retain the previous battery state and continue
    else
        Store dispatch, revenues, and degradation for the committed 24h  Evaluate ex-post CMI alignment using
          the realised dispatch and the corresponding congestion weight profile  Carry the terminal State-of-
          Charge to the next optimisation window
Annualise all recorded quantities over the simulated period
```

The complete set of physical operating constraints, market participation constraints, and State-of-Charge dynamics underlying the MILP objective is provided in Appendix B. Each connection agreement modifies this common structure in one of two ways: by adding operational constraints to the feasible dispatch region (e.g. curtailment, capacity limits, dispatch instructions), by adding a remuneration or penalty term to the objective, or both. The complete set of contract-specific constraints and remuneration terms is given in Appendix C.

Battery degradation is represented by a convex, depth-of-discharge dependent cycling model, piecewise-linearised for MILP compatibility, and combined with a fixed calendar-ageing component. Degradation and State-of-Charge both carry over continuously between optimisation windows.

E.4. Financial post-processing

Each Monte Carlo realisation contributes to every scenario family through an affinity weight based on its sampled (v, s) location in parameter space, with the weighting sharpening at later milestone years. For each (policy, scenario family), profit and degradation at the milestone years are aggregated under this weighting, interpolated to a full lifecycle trajectory, and combined with capital and operating expenditure into a cashflow series to compute IRR and NPV.

E.5. Post-processing and analysis

The post-processing stage processes the stored simulation outputs and, where necessary, executes targeted re-optimisations to derive:

- CMI alignment improvement per policy, relative to the Baseline run under the same market realisation;
- institutional friction diagnostics: compensation coverage (contract remuneration relative to uncompensated operational risk) and timing premium (profit under CMI-targeted versus random curtailment timing);
- design-property indicators (firmness, degradation impact, DAFO cost impact, remuneration coverage, revenue retention), each computed relative to the matching Baseline run;
- sensitivity analyses over key contract parameters (tariff discounts, compensation and penalty rates, targeting quality), using either direct recomputation from existing results (for purely financial parameters) or targeted re-solves (for parameters that affect the dispatch itself);
- an integrated scorecard comparing every policy against Baseline across profit, battery lifetime, CMI alignment, and bankability.

Together, these outputs connect the simulation results back to the three institutional tensions introduced in Chapter 3: temporal misalignment, asymmetric risk allocation, and the disconnect between operational control and financial responsibility.