From Firm to Flexible

A Sector-Comparative Assessment of TDTR Participation for Large Electricity Consumers in the Netherlands

Peebe ten Hoor

Master Thesis
Complex Systems Engineering & Management (TU Delft)

TU Delft

From Firm to Flexible

A Sector-Comparative Assessment of TDTR Participation for Large Electricity Consumers in the Netherlands

by

Peebe ten Hoor

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

Student number:	6291759	
Project duration:	1 February 2026 – 7 July 2026	
Thesis committee:	Dr. ir. K. Bruninx,	TU Delft, supervisor
	Prof. dr. M. Warnier,	TU Delft, second supervisor
	M. Hoeksma,	Alliander, company supervisor
	M. Bongaerts,	Alliander, company supervisor

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

TU Delft

Preface

This thesis is submitted in partial fulfilment of the requirements for the degree of Master of Science in Complex Systems Engineering and Management (CoSEM) at Delft University of Technology. The work was carried out in collaboration with Alliander, whose work on grid capacity and congestion management provided both the subject and the access to operational and regulatory experience that made it possible. What interested me in time-duration-based transport rights for DSOs is that the instrument is still so new. At the time of writing, it sits in a formal code-change trajectory and is barely described in the academic literature. As a result, almost every part of it had to be pieced together from regulatory documents and from practice before it could be represented in a model. What made it compelling is how directly a decision taken on paper, the terms of a conditional grid-access contract, shapes, or fails to shape, what actually happens on the physical network. Studying where those contractual terms and the realities of a congested distribution grid pull in the same direction, and where customer savings and grid relief quietly come apart, is the kind of socio-technical question the CoSEM programme is built around, and it is part of what made the project worth the effort. None of it would have come together without the people around me, whom I would like to thank here.

First of all I would like to thank my committee at TU Delft. My first supervisor, Kenneth Bruninx, whose sharp eye for the techno-economic side of the problem shaped how the optimisation and its cost assumptions were framed, and whose questions repeatedly pushed the analysis to be more precise than I would have left it on my own. I am also thankful to my second supervisor, Martijn Warnier, whose systems perspective kept the customer-level model connected to the wider question of what TDTR means for the grid and the operator, and helped me see where the results actually mattered. I am grateful to both of them for making time in their busy schedules to guide me throughout the process.

At Alliander I am very grateful to Martijn Hoeksma and Martijn Bongaerts for their support throughout the project and for giving me access to the operational knowledge, the regulatory context, and the network data the model rests on. I also appreciated their willingness to explain how grid access and congestion management work in practice, and to talk through the assumptions behind the model. Many such conversations were needed: with little prior research to build on, the model had to be established largely from regulatory documents, operational experience, and iterative discussion. Together with their colleagues, they also showed me how broad, urgent, and practically relevant the energy sector is, and helped me realise that this is the world in which I want to continue developing myself. I am sincerely thankful for that.

Finally I would like to thank my friends and family for their support and for the welcome distractions over the past year, with particular thanks to my parents for their encouragement throughout my studies.

Peebe ten Hoor
Delft, July 2026

Executive summary

Electricity grid congestion increasingly limits the ability of large electricity consumers to expand, electrify, or connect new assets. While grid reinforcement remains necessary, it is often slow due to spatial, permitting, and implementation constraints. This creates a need for complementary instruments that use existing grid capacity more efficiently. This thesis evaluates the time-duration-based transport right (TDTR) as such an instrument. TDTR is a form of Alternative Transport Right in which a customer receives firm capacity for part of its demand and conditional capacity for the remaining share. Under the Dutch design, customers are guaranteed firm access for 85% of the time in a year, while up to 15% may be restricted during predefined congestion-related call-hours. For the distribution system operator (DSO), TDTR is therefore both a customer contract and a potential congestion-management tool. The central question of this thesis is whether it can deliver value for the customer and the grid at the same time.

The question is examined with a customer-level optimisation model applied to three flexibility archetypes: logistics with an electric-vehicle fleet, cold storage, and dairy processing with an industrial heat pump. The logistics case represents charging that can be shifted within operational windows, the cold-storage case cooling that can be shifted within the thermal mass of the stored product, and the dairy case a largely continuous process with a shiftable heat-pump share. Each case is first evaluated with flexibility only under a fully firm connection, and then with flexibility combined with an optimised TDTR contract. The model selects the firm-to-TDTR ratio that minimises the customer's total annual cost of electricity supply, including energy costs, taxes and network tariffs, while respecting the technical constraints of the process.

Grid impact is evaluated through the reduction in the customer's annual peak, the change in system-wide grid pressure across the year, and, most importantly, the change in consumption during local call-hours. These call-hours are defined as the 15% busiest hours of the substation and therefore represent moments of local grid pressure. Each archetype is tested across four substation call-hour profiles whose timing differs relative to electricity prices. In well-aligned profiles, call-hours coincide with relatively high electricity prices, so the customer's cost incentive and the DSO's congestion objective point in the same direction. In misaligned profiles, call-hours coincide with relatively low electricity prices, so the customer has an incentive to consume exactly when the DSO would prefer load to fall.

On the customer side, TDTR creates a positive business case across all three archetypes, but the saving is driven less by connection size than by the share of load that can be shifted in time. Logistics achieves the largest saving, about 23%, because EV charging can be moved to cheaper hours within the charging window. Cold storage saves about 5%, because cooling can be shifted within the thermal flexibility of the stored product. Dairy processing saves about 2%, because the underlying process is relatively flat and continuous, leaving little room to reduce the monthly peak or shift substantial energy volumes. The largest electricity consumer is therefore not automatically the most suitable TDTR participant, the controllable share of the load is the more important determinant.

The grid-side result is more conditional. A positive customer business case does not automatically imply a lower load during the call-hours. TDTR reduces the customer's firm capacity right during predefined local congestion hours, but it does not by itself determine how the customer operates its flexible demand. That operation remains driven by the customer's own cost minimisation, in which electricity prices play an important role. This explains why the same product can deliver local relief on one substation and little or none on another.

On well-aligned substations, customer and grid incentives reinforce each other. The customer already has an incentive to shift flexible demand away from the expensive call-hours, even without TDTR. As a result, much of the observed call-hour relief is produced by price-driven flexibility rather than by the

TDTR contract itself. In these cases, TDTR still has value because it formalises the reduced capacity right during call-hours and prices conditional access, but its incremental effect on dispatch is limited. On misaligned substations, the mechanism is weaker: the customer's cost-minimising dispatch wants to use more electricity in the same hours in which the DSO wants load to fall. TDTR must then work against the market-price signal. The model shows that the customer accepts only a smaller TDTR share, or complies with the formal contract while delivering little additional reduction in total call-hour energy. The cost-optimal firm fraction thus becomes a practical bound on voluntary participation.

This distinction matters for how TDTR should be interpreted as a grid instrument. Its incremental grid value cannot be read directly from the difference between a reference profile and a TDTR profile, because part of the apparent relief may reflect price-driven flexibility that would have occurred even without the contract. TDTR can support reinforcement deferral when it reduces realised load during the hours that drive the local network constraint, or when it lowers the firm capacity claim that must be guaranteed during those same hours. However, neither effect should be assumed automatically. Whether TDTR translates into avoided or deferred investment depends on whether the reduced capacity right is actually binding during the call. Only where the customer's demand would otherwise reach the firm level does the restriction force load down, where realised demand sits below it, the customer keeps headroom to raise consumption up to its firm right, so call-hour load can still rise even though the formal capacity claim is lower.

The operator implication is therefore not that TDTR should be rolled out passively wherever a customer is willing to participate. Because voluntary participation is strongest where the customer's private incentives already align with the local congestion profile, TDTR may concentrate on the substations where it is easiest to sell rather than where it adds incremental grid value. On well-aligned substations, customers self-select into the product because the discount fits behaviour they would partly adopt anyway. On misaligned substations, where the contract would need to counteract the price signal, customers require higher compensation or a stricter firm-fraction requirement to participate. Without active targeting, TDTR therefore risks spreading fastest where it adds the least additional congestion relief.

This points to a more selective role for the DSO. Before offering TDTR, the operator should assess whether the customer's load profile, flexibility characteristics and price response fit the local call-hour profile. Customers with flexible demand that can be shifted away from congested hours are more likely to provide both a private business case and grid value. Customers whose flexibility is mainly used to follow low electricity prices may still find TDTR financially attractive, but their contribution to local or system-wide relief is more uncertain.

The results also qualify the role of adding more technical flexibility. More flexibility does not automatically mean more useful grid relief. In the dairy case, doubling the shiftable heat-pump share from 10% to 20% increases the saving only from 2.32% to 2.65%, while requiring twice the thermal-buffer capacity, because the underlying profile remains too flat to offer much steering potential. In the logistics case, adding stationary batteries improves the grid indicators but reduces the customer's saving as the battery grows: the smallest battery still leaves a positive saving of roughly 18%, below the roughly 23% of the battery-free logistics case, while the largest battery brings the business case close to zero or negative. The two cases differ. For dairy the saving still rises, but only marginally, while the buffer it needs doubles, adding capital cost and risk for little extra return. For the battery the value it earns from price arbitrage and peak reduction shrinks as it grows while its cost keeps rising, so beyond a point the saving drops.

In conclusion, TDTR can lower customer costs and can support local congestion relief, but only under specific conditions. Its value is strongest when the DSO's congestion signal, the customer's operational flexibility and the customer's economic incentives point in the same direction. Where these conditions hold, TDTR can align customer savings with local grid value. Where they do not, TDTR may still reduce the customer's formal firm capacity claim during call-hours or grant access within existing capacity that an unconditional firm connection could not, but realised congestion relief cannot be assumed in advance. The main lesson for the DSO is therefore to screen, monitor and steer: screen whether the customer's profile fits the local call-hours, monitor the load actually reduced rather than the TDTR share contracted, and steer contract terms or complementary instruments toward the substations where market incentives and grid needs diverge.

Contents

1 Introduction

Across Europe, electrification of heat, mobility, and industry is progressing faster than distribution grids can be reinforced. This increases the frequency and severity of distribution-grid congestion, creates long connection queues, and places additional operational pressure on Distribution System Operators (DSOs) (Skok et al., 2022; Thomaßen et al., 2024). In the Netherlands, recent policy and sector publications emphasise that congestion has become structural and that accelerated investment is necessary but insufficient to resolve constraints in the short term (Ministerie van Klimaat en Groene Groei, 2025; Netbeheer Nederland, 2025). Under persistent congestion, the challenge is therefore not only to expand network capacity, but also to use scarce capacity more efficiently and allocate it more effectively across users with different load profiles and flexibility potentials (Council of European Energy Regulators, 2023; Monterde et al., 2025; TNO, 2023).

In this context, Dutch regulators introduced the **Alternative Transport Right (ATR)** regime as a form of conditional grid access. ATRs allow users to obtain additional or earlier access in exchange for reduced firmness during predefined congestion periods (Autoriteit Consument & Markt, 2024; Netbeheer Nederland, 2024). Conceptually, ATRs align with a wider European move toward flexible or non-firm connection arrangements and congestion-management instruments. The academic literature suggests that such arrangements can increase hosting capacity and improve utilisation of existing networks, while raising questions about predictability, fairness, enforceability, and user acceptance (Bjarghov et al., 2024; Hennig et al., 2024; Monterde et al., 2025).

Despite growing attention, DSOs and policymakers still lack a consistent approach to assess the technical feasibility and economic attractiveness of conditional grid access for heterogeneous large electricity consumers. Large energy consumers differ strongly in operational flexibility, process constraints, and the economic consequences of curtailment (Golmohamadi, 2022; Leinauer et al., 2022; Santecchia et al., 2022; Zhang et al., 2025). The practical relevance of this thesis lies in providing a structured and transparent assessment logic that enables DSOs and large electricity consumers to evaluate how representative sector archetypes respond to TDTR-style restrictions in a Dutch distribution-grid context, thereby informing both participation decisions and the targeting of conditional grid-access arrangements to suitable customers.

1.1. Research objective

The objective of this thesis is to develop a structured and transparent techno-economic assessment framework for evaluating time-duration-based transport rights (TDTR) for large electricity consumers under substation-derived TDTR call-hour patterns. More specifically, the thesis aims to analyse how flexibility characteristics, curtailment costs, and TDTR contract conditions jointly shape the technical feasibility and economic attractiveness of TDTR participation.

To achieve this objective, the thesis develops an optimisation-based modelling approach from the perspective of a large electricity consumers. Within this approach, TDTR is not assessed only as a standalone non-firm access product, but as a conditional transport right that can be combined with a firm transport capacity component. The analysis takes representative large-consumer sectors as its starting point and models the characteristic flexibility asset of each at asset level. The economic attractiveness of TDTR participation is then evaluated at the case level, where the modelled flexibility is combined with the remaining demand, and the resulting findings are translated back to the sector archetype each case represents. This enables a comparative evaluation of alternative firm-plus-TDTR configurations and supports systematic reasoning about trade-offs between flexibility use, operational burden, and economic outcomes under different contract settings.

The broader aim is not to predict the exact real-world behaviour of every individual firm, but to provide a decision-support framework for DSOs assessing both the customer-side suitability of TDTR participation and its grid-side relevance under local congestion. The framework supports decisions on which types of large electricity consumers are promising candidates for TDTR, under which firm-to-TDTR ratios participation is technically feasible and financially beneficial for the customer, and whether the resulting flexibility is delivered during the moments that matter for the grid.

1.2. Research question and sub-questions

Based on the identified knowledge gap, this thesis addresses the following main research question:

> **Under which conditions does conditional grid access through TDTR deliver value, and to whom: when does a cost-optimal firm-to-TDTR ratio a large electricity consumer also relieve the local grid, and when do these two outcomes diverge?**

This main research question is operationalised through three sub-questions that follow the sequential logic of the study.

SQ1 How have prior studies modelled the interaction between demand-side flexibility and network-imposed restrictions, and which modelling choices follow from that body of work for a TDTR assessment at the level of large electricity consumers?

This sub-question reviews the existing literature on demand-side flexibility under network restrictions, distinguishing the network-side and user-side modelling streams, in order to derive the concrete modelling choices on which the framework of this thesis builds.

SQ2 For each representative sector and corresponding flexibility asset, what is the cost-optimal firm-to-TDTR ratio, and how much does participation at that ratio save the customer and relieve the grid relative to a firm-only and a flex-only scenario?

This sub-question applies the optimisation framework to each representative case. It uses the compliance-cost curve to identify the cost-optimal firm-to-TDTR ratio, and reports the resulting customer saving and grid-impact indicators against both a firm-only reference and a flex-only scenario, so that the value specific to TDTR is separated from the value of operational flexibility alone.

SQ3 How do these outcomes depend on the substation call profile and on the assumed flexibility of the assets, and under which conditions do the customer-side and grid-side outcomes align or diverge?

This sub-question repeats the optimisation across four substation call profiles and several asset-parameter variations. It distinguishes outcomes that remain stable from those that depend on local congestion timing, and identifies the conditions, the overlap between the local call-hours and the price-driven flexible window, under which an attractive customer contract does or does not translate into (local) grid relief.

1.3. Societal contribution

The societal relevance of this thesis follows directly from the congestion problem it addresses. Grid congestion now delays the connection of new industrial load, charging infrastructure, and electrified heat across the Netherlands, slowing the energy transition precisely where additional electricity demand is most needed (Ministerie van Klimaat en Groene Groei, 2025; Netbeheer Nederland, 2025). Because reinforcement cannot keep pace, instruments that use existing capacity more efficiently are not a convenience but a precondition for continued electrification under congestion.

This thesis contributes to that effort in three ways. First, it gives DSOs a transparent basis for deciding which large consumers to steer toward conditional access, so that scarce grid capacity is allocated to the connections where flexibility can relieve local congestion, rather than on a first-come-first-served basis. Second, it gives large electricity consumers a clearer view of whether and on what terms TDTR participation is worthwhile, lowering the information barrier that currently keeps suitable customers from engaging with conditional access at all. Third, by showing where a customer's private cost savings do and do not coincide with grid relief, it helps regulators and policymakers judge when a voluntary contract suffices and when additional steering or compensation is needed for the instrument to serve the public interest of a faster, more efficient energy transition. In each case the broader societal value is the same: helping connect more electrified demand within the grid that already exists, while the slower work of reinforcement catches up.

1.4. Scientific relevance

The scientific relevance of this thesis lies in connecting two literature streams that are often studied separately: non-firm or conditional grid access, and demand-side flexibility of large electricity consumers. Existing work on non-firm access mainly focuses on network, regulatory, or market-design perspectives, while end-user flexibility studies often analyse technical flexibility potential without linking it to explicit transport-right restrictions (Council of European Energy Regulators, 2023; Golmohamadi, 2022; Leinauer et al., 2022; Monterde et al., 2025; Santecchia et al., 2022; Zhang et al., 2025). As a result, limited insight exists into how TDTR-style restrictions affect individual flexibility assets from the customer perspective.

This thesis contributes by developing a techno-economic optimisation framework that evaluates TDTR participation at asset level. Against a firm-only reference, the framework compares a flexibility scenario without TDTR and a flexibility scenario with TDTR, while explicitly modelling shifting, buffering, curtailment costs, electricity-cost components, and substation-based call-hour patterns. This provides a structured way to assess under which firm-to-TDTR ratios participation is technically feasible and economically attractive for different large-consumer archetypes.

In doing so, the thesis translates the institutional concept of conditional access into optimisation-ready scenarios. It thereby contributes to the methodological development of customer-oriented TDTR assessment using substation-derived call-hour patterns.

1.5. Link with the CoSEM master programme

This thesis is closely connected to the Complex Systems Engineering and Management (CoSEM) master programme, as it studies a socio-technical problem in which technical infrastructure, regulation, economic incentives, and operational decisions of large electricity consumers come together. TDTR cannot be assessed from a purely technical perspective, since its feasibility as a congestion-management instrument also depends on contractual design, customer flexibility, tariff structures, and institutional implementation.

The thesis therefore applies a systems perspective by combining optimisation modelling with sector-specific operational constraints and policy-relevant contract conditions. In doing so, it reflects the CoSEM focus on analysing complex infrastructure systems and designing decision-support tools for problems where engineering, economics, and governance are strongly interdependent.

1.6. Research approach

The research process is summarised in Figure 1.1. The study starts from the problem of distribution-grid congestion and first reviews two relevant literature streams: non-firm or conditional grid access, and demand-side flexibility of large electricity consumers. This literature review is used to derive the main modelling choices, including the representation of TDTR as a stacked firm-plus-conditional transport right and the need to evaluate both customer-side and grid-operator-side outcomes.

The next step is the selection of representative large-electricity-consumer archetypes. Three cases are used as an evidence base: logistics with an electric vehicle fleet, cold storage, and dairy processing with an industrial heat pump. These cases are not intended to reproduce individual firms in full operational detail, but to represent different flexibility mechanisms: shiftable EV charging, thermal flexibility in cooling processes, and process-constrained heat-pump flexibility.

The selected cases are then translated into input data and modelling assumptions. For each case, a reference demand profile, flexibility envelope, cost assumptions, and TDTR call-hour signal are defined. These inputs feed a techno-economic optimisation model that compares a firm-only reference, a flexibility scenario without TDTR, and a flexibility-plus-TDTR scenario. The model is run across different firm-to-TDTR ratios, resulting in a compliance-cost curve and an optimal firm-to-TDTR configuration for each case.

The outputs are evaluated through two complementary lenses. First, customer-side performance is assessed through annual cost components and the corresponding economic attractiveness of TDTR participation. Second, grid-operator-side performance is assessed through indicators that measure whether load is actually reduced during relevant congestion-related call-hours. Sensitivity analyses are then used to test whether the findings are robust or depend on local congestion timing and available flexibility. The first of these varies the substation call profile: the temporal pattern of when call-hours occur at a given substation across the year, in terms of months and hours of the day. Because different substations become congested at different times, this pattern differs by location and determines when a customer's flexibility must respond.

This approach results in a comparative assessment of TDTR suitability across customer archetypes and local call profiles. The outcome is not a universal ranking of sectors, but a structured explanation of which flexibility characteristics make TDTR suitable for both customers and DSOs.

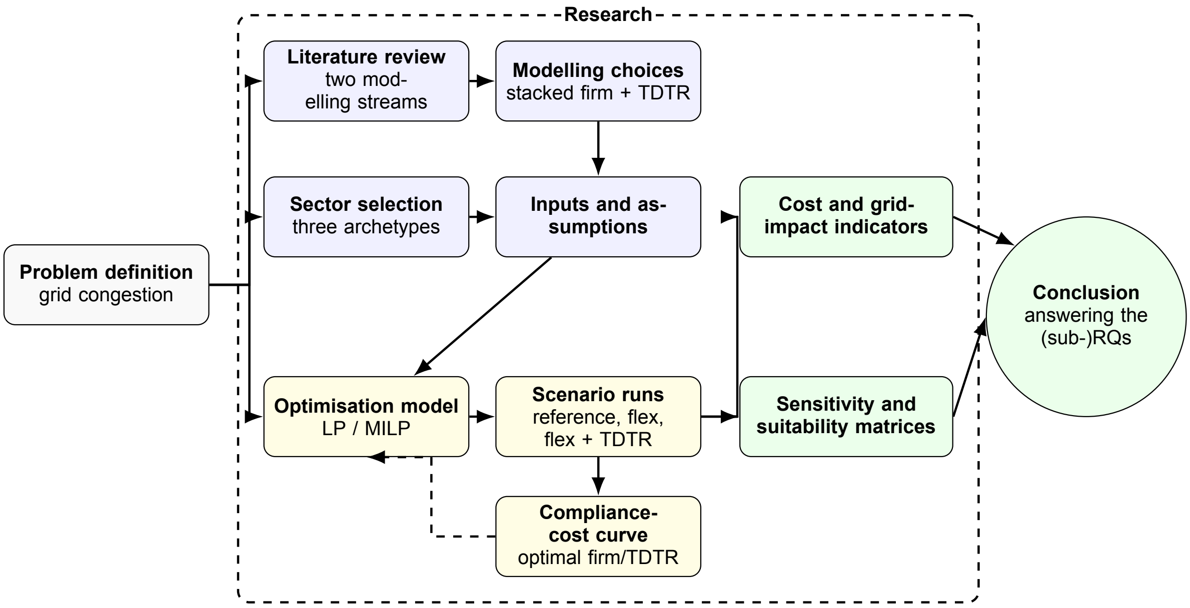

Figure 1.1: Schematic overview of the research process.

1.7. Thesis structure

The remainder of this thesis is structured as follows.

Chapter 2 – Background & Conceptual Framing Provides the conceptual and institutional background on Dutch distribution-grid congestion and congestion management, and positions ATRs, and TDTR in particular, as a stacked firm-plus-conditional access arrangement.

Chapter 3 – Modelling Flexibility under Network Restrictions Reviews how demand-side flexibility under network restrictions has been modelled, distinguishes the network-side and user-side streams, and derives the modelling choices on which this thesis builds.

Chapter 4 – Methodology Presents the techno-economic optimisation framework, including the system boundary, mathematical formalisation, TDTR call-hour signal, cost formulation, and grid-impact indicators.

Chapter 5 – Case Construction and Modelling Assumptions Translates the selected sector archetypes into three case models and develops their reference profiles, flexibility assumptions, cost assumptions, and substation call profiles.

Chapter 6 – Results Presents the optimisation results, including the compliance-cost curves, annual costs, grid-impact indicators, and sensitivity analyses across call profiles, asset assumptions, and the day-ahead price assumption.

Chapter 7 – Discussion Interprets what TDTR adds beyond price-driven flexibility, how it fits alongside other congestion instruments, when it may contribute to reinforcement deferral or conditional access, and where a network operator should steer it. It then sets out the limitations, recommendations, and directions for further work.

Chapter 8 – Conclusion Answers the research questions and states when TDTR delivers customer value, grid value, or both.

The appendices provide supporting material, including the economic framework, literature review tables, sector selection, call patterns, model verification, additional sensitivity results, and extra outputs.

2

Background & Conceptual Framing

This chapter provides the conceptual and institutional background for the thesis. It first explains why congestion in Dutch distribution grids has become a structural challenge and why this puts the traditional logic of firm transport access under pressure. It then positions Alternative Transport Rights (ATRs) within the broader landscape of congestion management and non-firm access arrangements. The chapter also describes the current Dutch ATR landscape and the implementation status of each variant, and concludes by defining the analytical perspective adopted in the remainder of the thesis.

2.1. Dutch distribution-grid congestion as a structural problem

Dutch DSOs operate under a security-of-supply mandate. In a traditional *firm-access* regime, contracted transport capacity is expected to be available whenever the user requires it. Responsibility for managing congestion thus rests with the system rather than with the individual customer. This logic holds as long as congestion is incidental. Once it becomes structural, DSOs face a growing tension between secure network operation and connecting additional demand, and the assumption that every customer can receive fully firm access at all moments becomes difficult to maintain. This is the context in which alternative access arrangements such as ATR have become relevant in the Netherlands.

2.1.1. Structural drivers of congestion

Congestion in Dutch distribution grids increasingly results from structural developments rather than short-lived operational problems. Three reinforcing dynamics stand out.

First, electrification raises peak demand across many end uses. Industrial electrification, electric mobility, and electrified heating add substantial loads that often coincide during already critical hours (Damianakis et al., 2023), pushing local network components closer to their limits more frequently and for longer. As Figure 2.1 illustrates, demand can approach the network limit during a few peak hours while substantial capacity remains unused for most of the day.

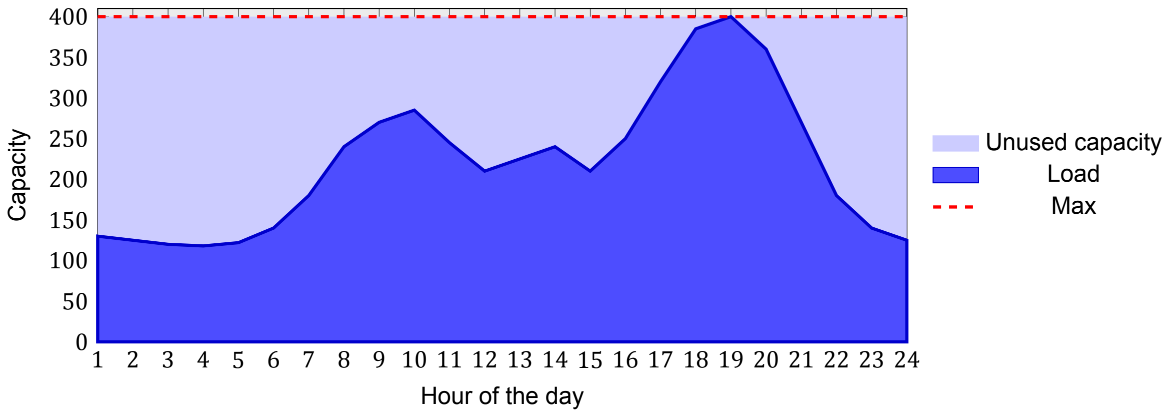

Figure 2.1: Illustrative demand profile: peak hours approach the network limit while substantial capacity remains unused for much of the day.

Second, the rapid growth of variable renewable generation shifts both the timing and the spatial distribution of network stress, creating new local congestion patterns on both the feed-in and withdrawal side (CE Delft & Merosch, 2024).

Third, reinforcement remains necessary but slow, delayed by lengthy planning and permitting procedures, supply-chain bottlenecks, and shortages of specialised labour (Kaya et al., 2026), so physical expansion cannot keep pace with demand growth.

Taken together, these developments make congestion a structural access problem rather than a purely operational one: in many congested areas capacity is not unavailable at all times, but scarce during specific hours and locations. The central issue therefore becomes how to allocate limited transport capacity more effectively under persistent scarcity.

2.2. From congestion management to conditional access

To understand where ATR fits, it is useful to distinguish congestion management from other electricity-system functions, and to position the various congestion-management instruments along a single conceptual spectrum. This section first defines the boundaries of the congestion-management domain and then introduces the three instrument families that organise both the international literature and Dutch practice.

2.2.1. Congestion management as a distinct system function

Congestion management addresses a different problem, it resolves location-specific network constraints that arise when the physical limits of particular network components are at risk of being violated (Trago & Kroondijk, 2024). Congestion management is therefore not primarily about energy balance, but about the spatially specific feasibility of transport through the network.

2.2.2. Three families of congestion-management instruments

The congestion-management literature and policy practice can be understood in terms of three broad instrument families, which differ in time horizon, certainty, and the way congestion risk is distributed between the system and the user.

A first family consists of *activation-based instruments*. These instruments contract and activate flexibility when congestion is expected or occurs. Examples include redispatch-type arrangements and local flexibility mechanisms in which market parties submit bids that can be called upon to relieve a specific bottleneck. Their strength is operational relevance: they target congestion moments directly. At the same time, their effectiveness depends on reliable activation, sufficient liquidity, and workable settlement and verification arrangements (Attar et al., 2025; Michaelis et al., 2025; Repo et al., 2020).

A second family consists of *incentive-based instruments*. These rely on price signals, such as time-of-use tariffs or peak-based network tariffs, to encourage users to shift demand away from system-stressing periods. Such instruments can improve average network utilisation, but they provide less direct control over highly localised congestion and do not guarantee a response when it is most needed (Hennig et al., 2023; Hofmann et al., 2025; Schittekatte et al., 2022).

A third family consists of *access arrangements*. Here, congestion is not managed only by activating flexibility on top of an existing firm right, but by modifying the quality of the transport right itself. Access can become conditional, limited, or differentiated, meaning that part of the congestion risk is contractually shifted from the system to the user. This family includes bilateral curtailment agreements as well as broader non-firm access regimes (Council of European Energy Regulators, 2023; Monterde et al., 2025).

This three-way distinction is central to the present thesis. Activation-based and incentive-based instruments contract or encourage flexibility on top of an existing firm access right, whereas ATR belongs to the access-arrangement family and changes the nature of access itself by making transport conditional under predefined rules.

2.3. ATR in the Dutch contractual landscape

Within the Dutch contractual landscape, two types of instruments are particularly relevant for unlocking demand-side flexibility under network scarcity: congestion-management contracts and alternative transport rights. Both instruments affect the customer's ability to use network capacity, but they differ in their contractual logic and in the direction of the capacity adjustment.

1. **Congestion-management contracts**: temporary arrangements applied to customers with an existing firm transport right. Under these contracts, the customer accepts a reduction of its available transport capacity during specific congestion periods. The instrument is therefore linked to the temporary management of local congestion until network reinforcement or another structural solution becomes available.
2. **Alternative transport rights**: structural conditional-access arrangements that provide additional or new transport rights when fully firm capacity is not available. Instead of reducing an existing firm right, ATR defines a non-firm or partly firm access product from the outset. The customer receives access under predefined conditions, for example by accepting restrictions during TDTR call-hours.

The distinction is therefore not simply that congestion-management contracts are activation-based whereas ATRs are access-based. Rather, congestion-management contracts can be understood as temporary reductions of existing firm access, while ATRs are closer to access-based instruments because they define the conditions under which additional or alternative network access is granted.

2.3.1. Alternative transport rights in the Dutch context

ATRs, by contrast, represent *new or additional* transport rights with a structural character. Table 2.1 summarises the ATR variants that are currently relevant in the Dutch context.

Table 2.1: Overview of Dutch ATR variants and their implementation status as of April 2026.

Variant	Core logic	Status
Time-duration-based transport right (TDTR)	Transport guaranteed for at least 85% of the hours per year. Restrictions during the remaining (up to) 15% of hours, communicated at least one day ahead. Dynamic in character.	Operational at TenneT (TSO) level
Time-block-based transport right (TBTR)	Transport access only within predefined time windows. Provides upfront predictability and targets off-peak capacity. Structural in character.	Operational at DSO level since April 2025, in limited form (still only the 00:00–06:00 block)
Fully variable transport right	No guaranteed minimum capacity, day-ahead allocation based on residual network space. Purely dynamic in character.	Intended to be operational at DSO level, but not yet implemented in practice
Energy-volume-based transport right	Customer contracts a daily energy volume rather than a guaranteed number of transport hours, available capacity may vary throughout the day while the contracted daily volume remains fixed. Dynamic in character.	Considered but not prioritised, requires a complex translation from contracted daily energy to time-varying power rights, and is no longer part of the active implementation pipeline (Netbeheer Nederland, 2024)

The *time-duration-based transport right* (TDTR) is dynamic in character and forms the central focus of this thesis. For now, TDTR is only available at the TSO level but movements are made to introduce it to the DSO level. However, extending this transport right to the DSO level is more complex, as it requires additional research to determine whether a fixed generic availability percentage can be defined for regional networks. According to the ACM, this depends on factors such as local load profiles, the impact on higher network levels, and the size of the flexible installation relative to local grid capacity (Autoriteit Consument & Markt, 2024).

Taken together, these variants share a common principle, they unlock residual grid capacity by differentiating the firmness and conditions of transport access in a transparent and contractually defined way. They differ, however, in the degree of predictability offered to the customer, in the voltage level at which they are available, and in the complexity of practical implementation.

2.3.2. Stacking firm and ATR capacity

An important practical feature of ATR is that it does not necessarily replace firm access altogether. Customers may instead combine a baseline of firm capacity for critical operations with an ATR-based top-up for more flexible demand. In such a stacked arrangement, the firm portion preserves reliability for the core process, while the ATR portion provides additional transport capacity subject to bounded exposure to restrictions. Figure 2.2 illustrates how this stacking logic translates the customer's contracted capacity into differentiated economic value.

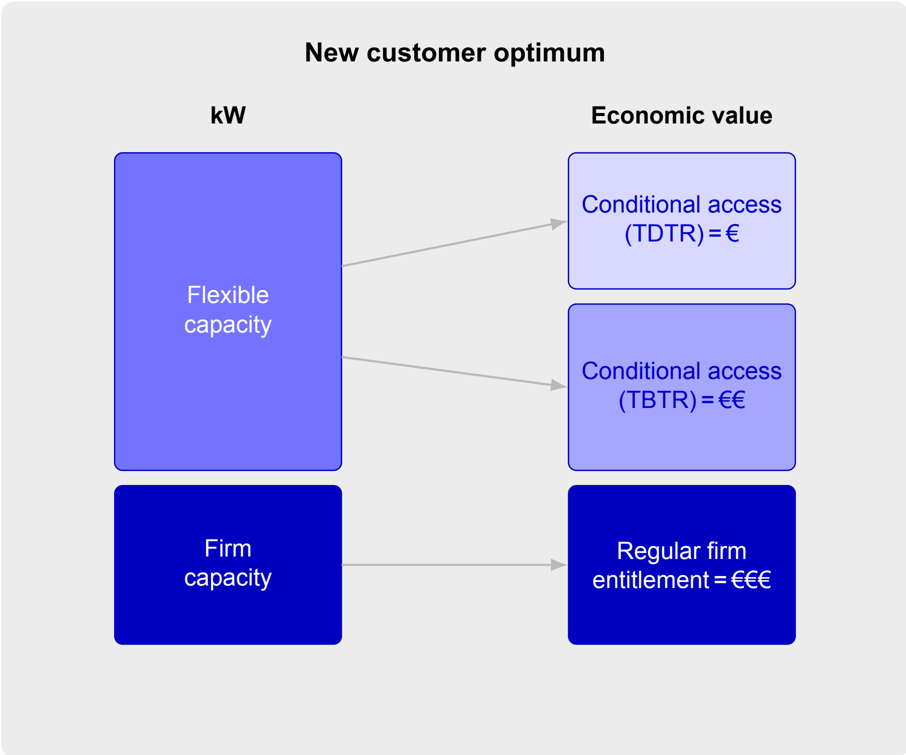

Figure 2.2: Stacking logic of firm and conditional transport capacity. The flexible portion can be contracted as TDTR or TBTR for instance, each with a different tariff treatment, while the firm portion carries the full tariff.

This logic is analytically important because it reflects how ATR may be used in practice. For many customers, the relevant decision is not whether to choose between fully firm and fully non-firm access, but how much of their required capacity can be placed under conditional access without creating unacceptable operational or economic risks, and which ATR variant best fits the temporal structure of their flexible demand.

2.3.3. Economic logic of TDTR participation

The stacking decision described above is ultimately an economic one, driven by how TDTR changes the customer's electricity and contract costs. This subsection summarises that logic at a high level, the full derivation of cost components and tariff treatment is provided in Appendix A.

For a large electricity consumer, the total electricity bill consists of three components: energy procurement costs paid to the supplier, network tariffs paid to the network operator, and government taxes. Of these, network tariffs are the component most directly affected by the choice of transport right. They are built from three drivers: a contracted-capacity charge (*KW-contract*), a realised-peak charge (*KW-max*), and a volumetric charge (*kWh*) (if applicable under type of connection, 2MVA and lower).

TDTR changes this structure in a specific way: for the capacity placed under conditional access, the *KW-contract* charge is removed, while the *KW-max* and *kWh* (if applicable) components remain unchanged. The economic appeal of TDTR therefore stems primarily from avoiding the contracted-capacity charge on the flexible portion of demand, while the retained *KW-max* charge preserves an incentive to limit realised peaks.

The value of TDTR participation is thus determined by weighing the network-tariff savings against the cost of adapting demand during restricted hours, which depends on the customer's load profile, flexibility options, and curtailment cost. Appendix A develops these components in full, including the day-ahead energy-cost structure, the tariff treatment of firm, TDTR, and TBTR capacity, and the proposed time-dependent network tariff which will be operationalised by the beginning of 2028.

2.4. Conclusion

The discussion above defines the analytical perspective adopted in this thesis. ATR is treated not as a generic flexibility product, but as a contract-based access arrangement designed to allocate residual network capacity by differentiating the firmness of access. Its core function is therefore not to guarantee fully firm transport rights for all users at all times, but to make conditional access possible for customers that can adapt part of their demand during restricted hours.

Among the Dutch ATR variants, this thesis focuses on TDTR for two reasons. First, TDTR is the dynamic ATR variant with a completed regulatory trajectory at the TSO level, providing a concrete contractual structure with a guaranteed availability share and a bounded number of restriction hours. Second, the extension of TDTR to the DSO level remains uncertain and requires further assessment, because local load profiles, interactions with higher network levels, and the size of flexible installations may strongly affect whether a generic availability percentage is appropriate for regional networks. This makes TDTR particularly relevant for this thesis: by testing TDTR-style restrictions against representative customer flexibility assets and local substation call profiles, the analysis contributes to understanding under which conditions TDTR could be technically feasible and economically attractive for large electricity consumers in a distribution-grid context.

3

Modelling flexibility under network restrictions

This chapter is structured around one overarching literature question that bridges the conceptual framing of Chapter 2 with the methodology developed in Chapter 4.

This part will answer **SQ1**:
How have prior studies modelled the interaction between demand-side flexibility and network-imposed restrictions, and which modelling choices follow from that body of work for a TDTR assessment at the level of large electricity consumers?

The purpose of this chapter is to position the modelling approach of this thesis within an existing body of work and to derive from that work a small number of concrete modelling choices that carry through into Chapter 4. Because Alternative Transport Rights as operationalised in the Netherlands, and TDTR in particular, are still in a formal code-change trajectory for TSO at the time of writing, dedicated peer-reviewed studies on Dutch ATR are absent. The Dutch ATR-specific evidence consulted here therefore consists largely of regulatory documents, position papers, and sector publications (Council of European Energy Regulators, 2023; Netbeheer Nederland, 2024). International studies on non-firm access, capacity-limitation services, and industrial flexibility are used by analogy to derive insights that are subsequently interpreted in the Dutch TDTR context.

The chapter is organised around the way prior studies bridge, or fail to bridge, two analytical perspectives that are central to TDTR and ATRs in general. Section 3.1 introduces a network-side stream, in which restriction regimes are designed and evaluated from the operator's point of view, and a user-side stream, in which flexibility is represented in terms of process constraints and operational cost. Section 3.2 then turns to the small set of bridge studies that combine both perspectives, with Verhagen et al. (2025) as the closest methodological neighbour, identifies the knowledge gap that remains, translates it into the two concrete modelling choices that structure the remainder of this thesis, and states where this thesis positions itself relative to the reviewed work.

The detailed search protocol underpinning this review, the resulting core literature set, and the synthesis matrix linking each study to ATR-relevant design dimensions are reported in Appendix B. Table B.1 summarises the Scopus search strings (`TITLE-ABS-KEY`), the number of hits, and the selected core studies, complemented by backward and forward citation tracking and the inclusion of relevant regulatory and guidance documents. The inclusion and exclusion criteria (E1–E6) applied during the full-text eligibility assessment are listed in Figure B.1. Table B.2 presents the resulting core literature set, while Table B.3 provides a synthesis matrix linking each study to ATR-relevant contract and modelling dimensions. Together, these materials ensure transparency of the review process and clarify how the literature informs the analytical framework developed in the remainder of this thesis.

3.1. Two modelling streams

Existing studies that combine flexibility and network restrictions can be grouped into two broad streams that differ in where they place analytical detail. Each stream is internally coherent and well developed, yet each leaves a complementary blind spot that this thesis seeks to address.

3.1.1. Network-side stream

The first stream takes the network operator as its analytical starting point. Studies in this stream model congestion management through capacity-limitation services, local flexibility markets, or queue-allocation mechanisms, and ask under which conditions these instruments deliver hosting capacity, welfare, or efficiency gains. Heinrich et al. (2021) show that a capacity-limitation service can prevent transformer overloads and substantially raise hosting capacity, provided the DSO can reliably enforce the agreed limit. Chen et al. (2023) extend this by treating the implementation pattern of the limitation as a design variable, showing that predictable but more restrictive schemes can outperform flexible but uncertain ones in terms of practical workability. Alavijeh et al. (2024) emphasise that the resulting market design depends critically on monitoring, baselining, and operational coordination, without which the theoretical benefits cannot be realised in practice. De Santi et al. (2025) move the focus from a single user to the queue-allocation problem, demonstrating that the way restrictions are distributed among users affects both efficiency and welfare outcomes.

These studies provide detailed network-side assessments of conditional access and congestion management mechanisms. They show that non-firm or capacity-limited arrangements can increase the use of existing network capacity, but that their effectiveness depends on enforceability, predictability, and operational coordination. This thesis does not model the network side in comparable technical detail. Instead, it takes substation-derived TDTR call-hours as an exogenous input and analyses how different large electricity consumers respond to these restrictions from the customer side. The network-side stream therefore motivates the use of conditional access, while the contribution of this thesis lies in linking local TDTR call profiles to the operational and economic response of heterogeneous large energy consumer flexibility assets. This enables a structured assessment of which firm-plus-TDTR configurations are technically feasible, financially attractive for the customer, and relevant from a grid-utilisation perspective.

At a broader conceptual level, Monterde et al. (2025) review non-firm connection types across Europe and distinguish between capacity-limited, time-limited, and dynamic access models. They identify substantial cross-country variation and, crucially for this thesis, a lack of standardised evaluation frameworks. Regulatory guidance reinforces this point: the CEER framework on Alternative Connection Agreements stresses transparency, non-discrimination, proportionality, and consumer protection as core requirements for flexible access schemes (Council of European Energy Regulators, 2023). Together, these sources confirm that non-firm access is technically viable and increasingly institutionalised, but they evaluate it primarily from a network, market-design, or regulatory perspective.

As a consequence, the representation of users in this stream is often simplified. Users typically appear as homogeneous capacity reductions, simple price-takers, or fixed shifting profiles. This is appropriate for the network-side questions these studies ask, but it leaves limited insight into the heterogeneity of process constraints, buffering options, and curtailment costs across large energy consumer sectors. As a result, this stream cannot determine at what firm-to-TDTR ratio participation becomes technically feasible and economically attractive for a specific customer type. That question requires a complementary customer-side modelling perspective.

3.1.2. User-side stream

The second stream takes the industrial user as its analytical starting point. Studies in this stream characterise demand-side flexibility through process constraints, recovery requirements, and curtailment cost. Golmohamadi (2022) provides a comprehensive sector survey showing that flexibility potential, response speed, and recovery behaviour vary strongly across heavy industries. Zhang et al. (2025) reach a similar conclusion in a more applied review. Santecchia et al. (2022) propose an explicit modelling abstraction in which heterogeneous processes are represented as “equivalent batteries”, enabling quantification of flexibility cost and required compensation levels.

These three studies broadly agree that industrial flexibility is sector-specific, but they diverge in how they translate that heterogeneity into modelling practice. Golmohamadi and Zhang et al. remain largely descriptive, mapping flexibility potential without providing a unified constraint representation. Santecchia et al. propose precisely such a unified representation, but at the cost of abstracting away from sector-specific operational logic. This thesis adopts an intermediate position: generic modelling elements (shiftable load, thermally buffered load, optional storage) are used across sectors for comparability, while sector-specific operational windows and curtailment bounds preserve heterogeneity where it materially affects feasibility.

The literature further shows that industrial flexibility should not be treated as a single aggregate quantity. Different flexibility mechanisms vary in response speed, duration, reversibility, and recovery requirements, which makes their suitability highly dependent on the timing, frequency, and notice characteristics of the restriction (Golmohamadi, 2022; Golmohamadi & Asadi, 2020; Serpe et al., 2023). These constraints are particularly relevant for ATR assessment, because the feasibility of responding to conditional transport limitations depends not only on the total flexible volume available, but also on whether that volume can be delivered at the moments the restriction is called and sustained for the full duration of each call.

Empirical case studies illustrate this heterogeneity and, importantly, also reveal its limits. Serpe et al. (2023) analyse demand-side flexibility in the German steel industry and show that significant flexibility exists, but only within strict operational limits; the aggregate potential reported in review studies may therefore systematically overestimate the flexibility actually deployable under binding restrictions. Complementing these technical insights, Leinauer et al. (2022) show that, despite technical potential, industrial participation in demand response is often limited by organisational risk aversion, regulatory uncertainty, and insufficient financial incentives. This is an important counterweight to the more optimistic sector reviews: technical flexibility does not automatically translate into practically available flexibility, and firms may adopt conservative operating strategies or decline to participate altogether. This thesis does not model the participation decision itself: it assumes throughout that the customer takes up a TDTR contract and responds to it. The adoption barriers identified by Leinauer et al. therefore fall outside the model and are treated as a limitation on how its results should be read, namely as the technical and economic suitability of TDTR conditional on participation, rather than a prediction of whether firms will participate.

Across these studies, the network restriction itself is treated as exogenous and typically simple, a fixed price signal, a generic curtailment event, or an annual flexibility budget. The temporal structure of the restriction, when calls occur, how often, with how much notice, and at what intensity, is rarely modelled in a way that maps onto a contractual instrument such as TDTR. The result is the mirror image of the network-side stream, rich on the user side, sparse on the contract side. This thesis addresses that underdeveloped side directly: it represents the restriction as a stacked firm-plus-TDTR contract and drives the user-side optimisation with substation-derived call-hours that fix when calls occur, how often, and at what intensity, so that the temporal structure the user-side literature leaves generic becomes an explicit input to the model.

3.2. Knowledge gap and resulting modelling choices

Taken together, the reviewed work supplies each building block needed to assess TDTR participation at the customer level, but no study combines them. The network-side stream models restriction regimes in detail while representing users generically (Alavijeh et al., 2024; Chen et al., 2023; De Santi et al., 2025; Heinrich et al., 2021; Monterde et al., 2025), the user-side stream models flexibility in detail while representing restrictions generically (Golmohamadi, 2022; Leinauer et al., 2022; Santecchia et al., 2022; Serpe et al., 2023; Zhang et al., 2025), and the studies that join the two consider either a single inherently flexible asset or a binary firm/non-firm contract (Bjarghov et al., 2024; Verhagen et al., 2025). What remains unknown is therefore not whether flexibility or non-firm access works in isolation, but whether, and under which firm-to-TDTR ratio, a stacked firm-plus-TDTR contract can relieve substation-level congestion while remaining cost-effective for a connected customer, and how this depends on the match between a customer's flexibility type and the local call profile. Concretely, no prior study combines (i) a stacked firm-plus-TDTR contract, (ii) sector-specific flexibility representations, and (iii) substation-derived call profiles within a single comparative optimisation framework.

Two modelling choices follow directly from this gap and are carried into Chapter 4. First, because the Dutch TDTR is offered on top of a firm contract (Chapter 2), the contract is represented as a stacked firm-plus-TDTR arrangement: the firm share is swept as a parameter to produce a compliance-cost curve, rather than evaluating a single fixed configuration as in Verhagen et al. (2025). Second, because the user-side literature shows that the burden of a restriction reflects sector-specific consequences such as lost production, restart cost, service-quality loss, or buffer cycling (Golmohamadi, 2022; Santecchia et al., 2022; Serpe et al., 2023), curtailment cost is parameterised per sector archetype in Chapter 5 rather than as a single universal penalty.

One assumption is inherited from this body of work rather than chosen. Like the studies that bridge both streams (Bjarghov et al., 2024; Santecchia et al., 2022; Verhagen et al., 2025), the optimisation assumes perfect foresight over call schedules and electricity prices. This is standard practice because it isolates the cost-minimal response to TDTR restrictions, but it makes the results a best-case benchmark: real customers face price and operational uncertainty and limited response time after a call notice, so the model may overstate the value of flexibility and understate compliance cost (Leinauer et al., 2022). The assumption is therefore retained for tractability and revisited as a limitation in Chapter 7.

Chapter 4 translates these choices into a concrete optimisation formulation. Appendix C and Chapter 5 develop the sector archetypes and case parameters. Chapter 6 applies the framework, and Chapter 7 returns to the literature to interpret the results in light of the assumptions on which they rest.

4
Methodology

This chapter sets out how the techno-economic optimisation framework that supports this thesis is constructed. The starting point is the literature-derived requirement of Chapter 3: an optimisation framework that integrates a stacked firm-plus-TDTR contract, sector-specific flexibility representations, and substation-derived call-hour profiles within a single comparative analysis. The framework developed here is built to meet that requirement.

The framework is applied to three sector archetypes that were selected for investigation in this thesis: **cold storage** (thermal cooling flexibility), **distribution centres with an EV fleet** (charging flexibility), and **dairy processing** (thermal buffer flexibility). These three are the outcome of a sector-selection procedure, documented in full in Appendix C, that screens thirteen candidate large-consumer sectors on demand-profile structure and a weighted multi-criteria assessment. The detailed case parameters follow in Chapter 5.

4.1. Modelling approach and system boundary

The research approach introduced in Chapter 1 is operationalised through a techno-economic optimisation model. This section specifies the modelling approach, the system boundary, and the temporal scope of the optimisation. The purpose is to make explicit how TDTR participation is translated into a customer-side optimisation problem.

The framework is implemented as a linear programme (LP) in Python solved using the Gurobi solver. Where binary variables are required, for example to model on/off behaviour of flexible cooling components or mutually exclusive buffer charging and discharging, the formulation extends to a mixed-integer linear programme (MILP). The EV-charging case, by contrast, represents charger availability through fixed time windows and capacity limits rather than binary variables, so it remains a linear programme.

The system boundary is drawn around a single large energy consumer connection. Wholesale electricity prices, grid tariffs, the energy tax, the substation-derived TDTR call-hour signal, and the reference demand profile are exogenous inputs, while the internal flexibility of the connection, including load shifting, thermal buffering, scheduled curtailment, and charging dispatch, is endogenous. Network reinforcement, capacity-allocation decisions by the DSO, behavioural adoption dynamics across firms, and strategic interaction between users fall outside the boundary.

The model operates on a 15-minute time resolution over one representative calendar year, giving 35,040 time steps. This resolution matches the metering interval of large consumers in the Dutch system and preserves the temporal structure of TDTR call-hours, peak loads, and shifting windows. A full-year horizon is required because the kW-max tariff is determined per calendar month and because the TDTR call-hour signal varies seasonally.

4.2. Modelling framework

The modelling framework in Figure 4.1 translates each sector case into exogenous inputs, runs the operational optimisation, and evaluates the resulting profiles on cost and grid impact. The *exogenous inputs* are fixed before each run and split into sector inputs (the reference demand profile and the asset-specific flexibility envelope) and market-and-contract inputs (day-ahead prices, network tariff components, the energy tax, the TDTR call-hour signal, and the contract configuration). The *optimisation* block then determines, the operational decisions that minimise total annual cost subject to the demand balance, contract, and asset-specific flexibility constraints. Its output is an optimised load profile, which the *cost* and *grid-impact evaluation* blocks translate into the annual cost components of Appendix A and the grid-impact indicators of Section 4.4.

The central operation in this framework is the firm-to-TDTR sweep. Rather than optimise the contract split, the firm share is fixed at a series of values across its full range and the optimisation is solved independently at each one. A firm share of 100% reproduces the firm-only reference (no conditional capacity, the TDTR cap never binds), lower firm shares enlarge the conditional portion, so the firm cap that applies during call-hours tightens and the customer must shift, buffer, or curtail more load to comply. Solving at each firm share traces out how the customer's cost and grid impact change as the contract moves from fully firm to predominantly conditional, which is what produces the compliance-cost curve rather than a single profitability verdict. The cost-minimising firm-to-TDTR ratio is then read off this curve. The feedback arrow in Figure 4.1 represents exactly this loop: each sweep value feeds a new contract configuration back into the optimisation.

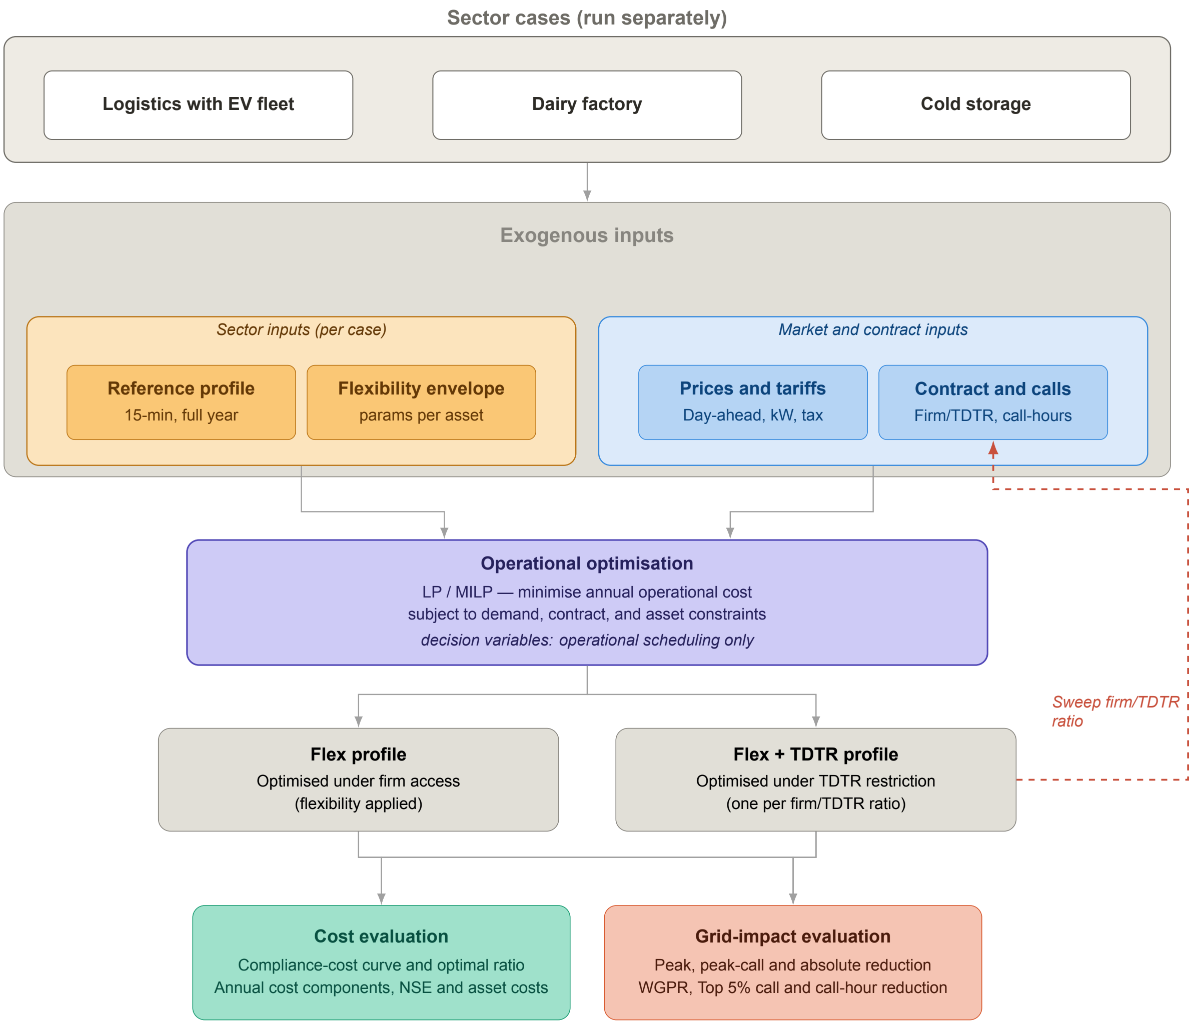

Figure 4.1: Modelling workflow for the techno-economic TDTR assessment.

4.3. Mathematical formalisation

This section develops the formal LP/MILP underlying the framework. All sets, indices, parameters, and decision variables used below are defined in the nomenclature section at the start of the thesis and are not reintroduced here.

4.3.1. Objective function

The optimisation minimises the total variable annual cost:

$$\begin{aligned}\min Z = {} & \underbrace{\sum_{t\in\mathcal{T}} \lambda_t\, p_t^{\text{tot}}\,\Delta t}_{\text{commodity cost}} + \underbrace{c^{\text{tax}} \sum_{t\in\mathcal{T}} p_t^{\text{tot}}\,\Delta t}_{\text{energy tax}} + \underbrace{c^{\text{kWh}} \sum_{t\in\mathcal{T}} p_t^{\text{tot}}\,\Delta t}_{\text{volumetric network cost}} \\ & + \underbrace{c^{\max} \sum_{m\in\mathcal{M}} p_m^{\text{peak}}}_{\text{monthly kW-max charge}} + \underbrace{c^{\text{NSE}}\, E^{\text{NSE}}}_{\text{non-served energy penalty}} .\end{aligned} \tag{4.1}$$

The first three terms are volumetric, all applied to total demand p_t^{tot}: the commodity cost values it at the hourly day-ahead price λ_t, while the energy tax and (where applicable) the volumetric network tariff apply the fixed rates c^{tax} and c^{kWh}. The fourth term applies the monthly kW-max tariff $c^{\max}$ to the highest demand p_m^{peak} in each month. The fifth term penalises non-served energy E^{NSE}.

The contracted-capacity cost is intentionally not part of the objective. For each optimization run, the contract configuration $(P^{\text{firm}}, P^{\text{TDTR}})$ is fixed before the optimisation and therefore does not affect the dispatch decision. Annual contracted-capacity costs are computed outside the optimisation as

$$C^{\text{contract}} = \underbrace{12\, c^{\text{contract}} P^{\text{firm}}}_{\text{firm capacity charge}} + \underbrace{12\, c^{\text{TDTR}} P^{\text{TDTR}}}_{\substack{\text{TDTR capacity charge} \\ \text{(discounted tariff)}}} , \tag{4.2}$$

where c^{TDTR} is the discounted contracted-capacity tariff for the TDTR portion. Total annual cost in any scenario is then

$$C^{\text{total}} = \underbrace{Z}_{\substack{\text{operational} \\ \text{cost}}} + \underbrace{C^{\text{contract}}}_{\substack{\text{firm + TDTR} \\ \text{capacity cost}}} + \underbrace{C^{\text{asset}}}_{\substack{\text{annualised buffer,} \\ \text{EMS, etc.}}} , \tag{4.3}$$

with C^{asset} capturing annualised asset-specific investments such as a thermal buffer or an energy management system.

Asset-specific investments are converted into equivalent annual costs so that one-off capital expenditures can be compared with recurring annual cost components in (4.3). For an investment I with technical lifetime L and discount rate r, the annualised asset cost is calculated as:

$$C^{\text{asset}} = \underbrace{I \cdot \text{CRF}(r, L)}_{\text{annualised CAPEX}} + \underbrace{C^{\text{O\&M}}}_{\text{fixed annual O\&M}} + \underbrace{C^{\text{deg}}}_{\text{degradation cost}} , \tag{4.4}$$

where the capital recovery factor is defined as:

$$\text{CRF}(r, L) = \frac{r(1+r)^L}{(1+r)^L - 1}. \tag{4.5}$$

4.3.2. Common constraints

The constraints below apply to all cases. Throughout, p_t^{tot} denotes the total demand at the connection in time step t, on which the network charges, the monthly peaks, and the TDTR cap are all evaluated. Each case defines how p_t^{tot} is composed from its own fixed and flexible components, including any shifting or curtailment, the case-specific power balances (Eqs. (4.21), (4.30), and (4.46)) are the concrete instances of this composition.

Monthly peak definition. The kW-max tariff is charged on the highest demand within each calendar month. The monthly peak variable p_m^{peak} is therefore bounded from below by the total demand at every time step in that month:

$$p_m^{\text{peak}} \geq p_t^{\text{tot}} \quad \forall m \in \mathcal{M},\ \forall t \in \mathcal{T}_m. \tag{4.6}$$

Because p_m^{peak} enters the objective with a positive coefficient, the optimiser drives it down to the smallest feasible value, which is exactly the maximum demand realised in that month.

Annual peak definition. The annual peak is defined in the same way over the full horizon:

$$p^{\text{peak,ann}} \geq p_t^{\text{tot}} \quad \forall t \in \mathcal{T}. \tag{4.7}$$

The annual peak is not part of the objective, it is tracked for reporting and for calculating the contracted-capacity cost outside the optimisation, and is used consistently across all cases.

TDTR call-hour constraint. During TDTR call-hours the total demand may not exceed the firm capacity, outside call-hours it may not exceed the full reference contract:

$$p_t^{\text{tot}} \leq P^{\text{firm}} \quad \forall t \in \mathcal{T}^{\text{call}}, \tag{4.8}$$

$$p_t^{\text{tot}} \leq P^{\text{contract}} \quad \forall t \in \mathcal{T} \setminus \mathcal{T}^{\text{call}}. \tag{4.9}$$

Flexible demand that cannot be served under the cap is absorbed by the demand-side NSE terms in the case-specific balances. Inflexible base load in excess of the cap is forcibly curtailed through a non-negative curtailment term n_t^{curt} that enters the case-specific power balance with a negative sign and is bounded by the excess of the inflexible load p_t^{inflex} over the applicable cap P_t^{cap}, with $P_t^{\text{cap}} = P^{\text{firm}}$ during call-hours and $P_t^{\text{cap}} = P^{\text{contract}}$ otherwise:

$$0 \ \leq \ n_t^{\text{curt}} \ \leq \ \max\!\left(0,\ p_t^{\text{inflex}} - P_t^{\text{cap}}\right) \qquad \forall t \in \mathcal{T}. \tag{4.10}$$

That the hard cap is implemented correctly, and that the resulting structural minimum firm fraction $f_{\min}$ matches an independently computed value, is verified in Appendix E.3.

Non-served energy. Two NSE components are tracked separately and penalised at the same rate c^{NSE} in the objective:

$$E^{\text{NSE}} = E^{\text{NSE,dem}} + E^{\text{NSE,curt}}, \tag{4.11}$$

where

$$E^{\text{NSE,dem}} = \sum_{t \in \mathcal{T}} e_t^{\text{NSE}} \tag{4.12}$$

is the case-specific demand-side NSE and

$$E^{\text{NSE,curt}} = \sum_{t \in \mathcal{T}} n_t^{\text{curt}}\, \Delta t \tag{4.13}$$

is the curtailed inflexible base load introduced in Eqs. (4.8)–(4.10). Splitting the two preserves the breakdown in the result tables while ensuring both shortfall types are priced consistently.

Non-negativity. All power and energy variables are non-negative. Case-specific upper bounds are added per case.

4.3.3. Cold storage flexibility optimisation

The cold-storage case represents temperature-constrained cooling demand, in which the flexible share of the cooling load can be shifted through the thermal inertia of the refrigerated space. Flexibility is modelled through an explicit thermal buffer with charge, discharge, and state-of-charge (SOC) variables: cooling delivered ahead of time is stored as cold and used to cover later cooling demand. Demand can therefore only be served earlier than originally required (pre-cooling), not later. The optimisation determines when the flexible cooling demand is served, subject to the installed capacity, the buffer capacity and daily runtime bounds, ramping constraints, and TDTR call-hour restrictions. The case is formulated as a MILP, because the compressor on/off logic and the mutually exclusive buffer charge/discharge operation are represented by binary variables. The case-specific parameter values, including the flexible share, buffer sizing, and runtime limits, are given in §5.1.

Flexible cooling-demand balance Let $p_t^{\text{flex,ref}}$ denote the flexible part of the reference cold-storage electricity demand at time step t. This demand is covered by direct cooling p_t^{direct}, by discharging previously stored cooling p_t^{dis}, or recorded as non-served cooling demand:

$$p_t^{\text{direct}} + p_t^{\text{dis}} + p_t^{\text{NSE}} = p_t^{\text{flex,ref}} \qquad \forall t \in \mathcal{T}. \tag{4.14}$$

All terms are expressed in electrical-equivalent power.

Flexible cooling power The flexible compressor power at time step t equals the direct cooling plus any cooling produced ahead of demand and charged into the buffer:

$$p_t^{\text{flex}} = p_t^{\text{direct}} + p_t^{\text{ch}} \qquad \forall t \in \mathcal{T}. \tag{4.15}$$

Charging the buffer represents pre-cooling: the compressor runs earlier than the original demand moment and the thermal mass of the cold store retains the delivered cooling until it is needed.

Buffer state-of-charge balance The buffer SOC evolves through charging and discharging. Because the storage medium is the thermal mass of the cold store itself, no standing-loss term is included:

$$SOC_t = SOC_{t-1} + p_t^{\text{ch}}\Delta t - p_t^{\text{dis}}\Delta t \qquad \forall t \in \mathcal{T}. \tag{4.16}$$

The first time step is linked cyclically to the final time step to avoid a free initial buffer level:

$$SOC_1 = SOC_{|\mathcal{T}|} + p_1^{\text{ch}}\Delta t - p_1^{\text{dis}}\Delta t. \tag{4.17}$$

Buffer capacity limit The SOC is bounded by the buffer capacity, set to two hours of flexible cooling capacity:

$$0 \le SOC_t \le E^{\text{CS,buf}}, \qquad E^{\text{CS,buf}} = 2\,P^{\text{max,flex}} \qquad \forall t \in \mathcal{T}. \tag{4.18}$$

Mutually exclusive buffer operation Because the buffer is lossless, simultaneous charging and discharging is not automatically suboptimal and must be excluded explicitly through a binary buffer-mode variable z_t^{buf}:

$$p_t^{\text{ch}} \le M^{\text{CS,buf}} z_t^{\text{buf}} \qquad \forall t \in \mathcal{T}, \tag{4.19}$$

$$p_t^{\text{dis}} \le M^{\text{CS,buf}}(1 - z_t^{\text{buf}}) \qquad \forall t \in \mathcal{T}, \tag{4.20}$$

where $z_t^{\text{buf}} \in \{0,1\}$ and $M^{\text{CS,buf}} = P^{\text{max,flex}}$ is an upper bound on the buffer charge and discharge power.

Total power balance Total cold-storage demand equals the inflexible demand, net of any forced curtailment under the hard TDTR cap (Eq. (4.10)), plus the flexible compressor power:

$$p_t^{\text{tot}} = p_t^{\text{inflex}} - n_t^{\text{curt}} + p_t^{\text{flex}} \qquad \forall t \in \mathcal{T}. \tag{4.21}$$

Installed-capacity limit The total electricity demand is bounded by the installed electrical capacity of the cold-storage facility:

$$0 \leq p_t^{\mathrm{tot}} \leq P^{\mathrm{max}} \qquad \forall t \in \mathcal{T}. \tag{4.22}$$

Flexible-capacity limit The flexible compressor power and the buffer charge and discharge power are limited by the flexible part of the installed capacity:

$$0 \leq p_t^{\mathrm{flex}} \leq P^{\mathrm{max,flex}}, \qquad 0 \leq p_t^{\mathrm{ch}}, p_t^{\mathrm{dis}} \leq P^{\mathrm{max,flex}} \qquad \forall t \in \mathcal{T}. \tag{4.23}$$

Flex-on logic A binary variable y_t indicates whether the compressor-driven flexible cooling component is active in a given time step. It is linked to the flexible electricity use through a big-M formulation:

$$p_t^{\mathrm{flex}} \leq P^{\mathrm{max,flex}} y_t \qquad \forall t \in \mathcal{T}, \tag{4.24}$$

$$p_t^{\mathrm{flex}} \geq \epsilon^{\mathrm{flex}} y_t \qquad \forall t \in \mathcal{T}^{\mathrm{feas}}. \tag{4.25}$$

where $y_t \in \{0, 1\}$ and ϵ^{flex} is a small positive threshold. If $y_t = 0$, the flexible cooling component is off and no flexible cooling electricity is used in that time step. If $y_t = 1$, the flexible component operates at least at the minimum threshold.

Daily maximum flex-off duration The cold-storage compressor cannot be left idle indefinitely, because the refrigerated space must remain within its temperature band. The model therefore limits the number of time steps per day in which the flexible component may be off:

$$\sum_{t \in \mathcal{T}_d^{\mathrm{feas}}} (1 - y_t) \leq N^{\mathrm{off,max}} \qquad \forall d \in \mathcal{D}. \tag{4.26}$$

Ramping constraint To avoid abrupt changes in cooling power, a ramp-rate limit is imposed between consecutive time steps:

$$p_t^{\mathrm{tot}} - p_{t-1}^{\mathrm{tot}} \leq R^{\mathrm{max}} \qquad \forall t = 2, \dots, |\mathcal{T}|, \tag{4.27}$$

$$p_{t-1}^{\mathrm{tot}} - p_t^{\mathrm{tot}} \leq R^{\mathrm{max}} \qquad \forall t = 2, \dots, |\mathcal{T}|. \tag{4.28}$$

Non-served energy The non-served flexible cooling demand is included in the generic non-served energy expression as:

$$E^{\mathrm{NSE}} = \sum_{t \in \mathcal{T}} p_t^{\mathrm{NSE}} \Delta t. \tag{4.29}$$

This non-served energy is penalised in the objective function using the case-specific NSE penalty.

4.3.4. Shiftable electric-vehicle charging

Electric-vehicle charging is modelled as shiftable electricity demand with bounded availability windows. The model does not change the underlying charging requirement, but determines when charging is scheduled within the feasible windows, allowing it to reduce electricity costs, avoid monthly peaks, and comply with TDTR call-hour restrictions where feasible.

Two charging types are distinguished. Fast charging, represented by the charging plaza, is modelled with explicit shift variables within a bounded time window, analogous to the shift-window formulation used for cold storage. Depot charging is modelled differently: vehicles are connected for a longer overnight period and must receive a required amount of energy before departure. Depot charging is therefore represented through an energy-balance constraint over each overnight availability window, subject to the number of simultaneously available depot chargers.

The total connection demand combines a fixed base load, realised depot charging, realised flexible fast charging, and any inflexible fast-charging demand outside the flexible operating window, net of any forced curtailment of the inflexible components under the hard TDTR cap (Eq. (4.10)):

$$p_t^{\mathrm{tot}} = p_t^{\mathrm{base}} + p_t^{\mathrm{fast,inflex}} - n_t^{\mathrm{curt}} + p_t^{\mathrm{depot}} + p_t^{\mathrm{fast}} \qquad \forall t \in \mathcal{T}, \tag{4.30}$$

where the curtailment bound of Eq. (4.10) applies to the joint inflexible component $p_t^{\mathrm{inflex}} = p_t^{\mathrm{base}} + p_t^{\mathrm{fast,inflex}}$.

Fast-charging flexibility

Let $p_t^{\mathrm{fast,ref}}$ denote the flexible part of the reference fast-charging demand at time step t, defined over the daytime fast-charging window $\mathcal{T}^{\mathrm{fast}}$. Fast-charging demand outside this window remains inflexible and is included in $p_t^{\mathrm{fast,inflex}}$.

The flexible fast-charging demand may be shifted to another time step τ within the allowed fast-charging window Ω_t^{fast}:

$$\Omega_t^{\mathrm{fast}} = \left\{\tau \in \mathcal{T} : |\tau - t| \le K^{\mathrm{fast}},\ \tau \in \mathcal{T}^{\mathrm{fast}}\right\} \qquad \forall t \in \mathcal{T}^{\mathrm{fast}}, \tag{4.31}$$

where $\mathcal{T}^{\mathrm{fast}}$ denotes the daytime fast-charging window and K^{fast} the permitted shift horizon.

The decision variable $x_{t,\tau}^{\mathrm{fast}}$ is the fast-charging power originally required at time step t but delivered at time step τ. The fast-charging balance is:

$$\sum_{\tau \in \Omega_t^{\mathrm{fast}}} x_{t,\tau}^{\mathrm{fast}} + n_t^{\mathrm{fast}} = p_t^{\mathrm{fast,ref}} \qquad \forall t \in \mathcal{T}^{\mathrm{fast}}, \tag{4.32}$$

where n_t^{fast} is non-served fast-charging demand. This variable is included to avoid infeasibility under restrictive charging and TDTR conditions and is penalised in the objective.

The realised flexible fast-charging load at time step τ is:

$$p_\tau^{\mathrm{fast}} = \sum_{t : \tau \in \Omega_t^{\mathrm{fast}}} x_{t,\tau}^{\mathrm{fast}} \qquad \forall \tau \in \mathcal{T}^{\mathrm{fast}}. \tag{4.33}$$

Fast charging is limited by the installed charging-plaza capacity:

$$0 \le p_t^{\mathrm{fast}} \le P^{\mathrm{fast,max}} \qquad \forall t \in \mathcal{T}^{\mathrm{fast}}. \tag{4.34}$$

The electrical non-served energy from fast charging is:

$$E^{\mathrm{fast,NSE}} = \sum_{t \in \mathcal{T}^{\mathrm{fast}}} n_t^{\mathrm{fast}} \Delta t. \tag{4.35}$$

Depot-charging flexibility

Depot charging represents vehicles connected for a longer overnight period. In contrast to fast charging, depot charging is not modelled with explicit time-step-to-time-step shift variables. Instead, the required depot energy is aggregated per overnight charging period and must be delivered within the corresponding availability window.

Let $\Omega_d^{\mathrm{depot}}$ denote the depot availability window for night d. Depot charging is allowed only inside the overnight window:

$$p_t^{\mathrm{depot}} = 0 \qquad \forall t \notin \bigcup_{d \in \mathcal{D}} \Omega_d^{\mathrm{depot}}. \tag{4.36}$$

Within each depot window, charging is bounded by the available depot-charging capacity:

$$0 \le p_t^{\mathrm{depot}} \le P^{\mathrm{depot,max}} \qquad \forall t \in \Omega_d^{\mathrm{depot}},\ \forall d \in \mathcal{D}. \tag{4.37}$$

The depot capacity is determined by the number of simultaneously available depot chargers:

$$P^{\text{depot,max}} = N^{\text{depot,ch}} P^{\text{ch}}, \tag{4.38}$$

where $N^{\text{depot,ch}}$ is the number of depot chargers and P^{ch} is the charging power per charger.

For each overnight period, the required depot energy must be delivered within the corresponding window:

$$\sum_{t \in \Omega_d^{\text{depot}}} p_t^{\text{depot}} \Delta t + E_d^{\text{depot,NSE}} = E_d^{\text{depot,req}} \qquad \forall d \in \mathcal{D}, \tag{4.39}$$

where $E_d^{\text{depot,NSE}}$ is non-served depot energy. The required depot energy follows from the unmanaged reference profile over the associated operational day:

$$E_d^{\text{depot,req}} = \sum_{t \in \mathcal{T}_d^{\text{op}}} p_t^{\text{depot,ref}} \Delta t \qquad \forall d \in \mathcal{D}. \tag{4.40}$$

Here, $\mathcal{T}_d^{\text{op}}$ denotes the operational day whose depot charging demand is assigned to night d.

To avoid unrealistically concentrating all depot charging in a short peak, the model also links depot charging power to the number of active depot chargers:

$$p_t^{\text{depot}} \le P^{\text{ch}} a_t^{\text{depot}} \qquad \forall t \in \Omega_d^{\text{depot}},\ \forall d \in \mathcal{D}, \tag{4.41}$$

where a_t^{depot} is the number of active depot chargers at time step t. This number cannot exceed the number of trucks that require charging in the corresponding night:

$$0 \le a_t^{\text{depot}} \le N_d^{\text{truck,req}} \qquad \forall t \in \Omega_d^{\text{depot}},\ \forall d \in \mathcal{D}. \tag{4.42}$$

The number of trucks requiring charging in night d follows from the required depot energy and the energy of one full charging session, capped at the size of the electrified fleet:

$$N_d^{\text{truck,req}} = \min\left(\left\lceil \frac{E_d^{\text{depot,req}}}{E^{\text{charge}}} \right\rceil,\ N^{\text{EV}} \right) \qquad \forall d \in \mathcal{D}, \tag{4.43}$$

where E^{charge} is the energy per full charging session and N^{EV} the number of electric vehicles. This number bounds the active depot chargers through Eq. (4.42).

4.3.5. Dairy heat-pump optimisation

The dairy case represents process-constrained electrified heat, in which only a limited share of the heat-pump demand can be shifted through a thermal buffer. The underlying dairy-process electricity demand is fixed, while a predefined fraction of the heat-pump demand is treated as flexible. Because the COP is assumed constant, heat flows are converted to electrical-equivalent values before entering the electricity-based cost formulation. The optimisation therefore determines when the flexible heat demand is produced, subject to the heat-pump capacity, buffer dynamics, ramping constraints, and TDTR call-hour restrictions.

Analogous to the cold-storage case, flexibility is represented through an explicit thermal buffer with charge, discharge, and state-of-charge (SOC) variables. Flexible heat can be produced ahead of its original demand time and stored in the buffer until it is needed, demand at each time step is covered by direct production or by discharging previously stored heat. Unlike the cold-storage buffer, the dairy buffer carries a thermal standing loss: stored heat that is lost must be compensated by additional heat-pump production, so standing losses translate into additional electricity consumption and cost. A binary buffer-mode variable prevents simultaneous charging and discharging, which makes the dairy heat-pump formulation a MILP. The buffer capacity is sized exogenously rather than optimised, the case-specific sizing rule and parameter values are given in §5.3.

Flexible heat balance Let $q_t^{\text{flex,ref}}$ denote the flexible part of the heat-pump heat demand at time step t, and $q_t^{\text{HP,flex}}$ the flexible heat produced by the heat pump. For every time step, production plus buffer discharge plus non-served heat must equal the flexible heat demand plus buffer charge:

$$q_t^{\text{HP,flex}} + q_t^{\text{dis}} + n_t^{\text{HP}} = q_t^{\text{flex,ref}} + q_t^{\text{ch}} \qquad \forall t \in \mathcal{T}. \tag{4.44}$$

Here, q_t^{ch} and q_t^{dis} are the buffer charge and discharge flows, and n_t^{HP} is non-served flexible heat demand. All heat-flow variables are expressed as thermal power.

Flexible heat-pump output The corresponding electrical heat-pump power follows from the constant COP:

$$p_t^{\text{HP,flex}} = \frac{q_t^{\text{HP,flex}}}{COP} \qquad \forall t \in \mathcal{T}. \tag{4.45}$$

Total electrical power balance The total connection demand equals the fixed dairy-process demand and the inflexible heat-pump demand, net of any forced curtailment under the hard TDTR cap (Eq. (4.10)), plus the optimised flexible heat-pump demand:

$$p_t^{\text{tot}} = p_t^{\text{dairy}} + p_t^{\text{HP,inflex}} - n_t^{\text{curt}} + p_t^{\text{HP,flex}} \qquad \forall t \in \mathcal{T}, \tag{4.46}$$

where the curtailment bound of Eq. (4.10) applies to the joint inflexible component $p_t^{\text{inflex}} = p_t^{\text{dairy}} + p_t^{\text{HP,inflex}}$.

Heat-pump capacity limit The flexible heat-pump power is limited by the remaining heat-pump capacity after the inflexible heat-pump demand has been supplied:

$$p_t^{\text{HP,flex}} \leq P^{\text{HP,max}} - p_t^{\text{HP,inflex}} \qquad \forall t \in \mathcal{T}. \tag{4.47}$$

Ramping constraint To avoid unrealistically abrupt changes in heat-pump operation, the flexible heat-pump power is subject to a ramp-rate limit:

$$p_t^{\text{HP,flex}} - p_{t-1}^{\text{HP,flex}} \leq R^{\text{HP}} \qquad \forall t = 2, \dots, |\mathcal{T}|, \tag{4.48}$$

$$p_{t-1}^{\text{HP,flex}} - p_t^{\text{HP,flex}} \leq R^{\text{HP}} \qquad \forall t = 2, \dots, |\mathcal{T}|, \tag{4.49}$$

Buffer state-of-charge balance The thermal-buffer state of charge evolves through charging, discharging, and thermal standing losses:

$$SOC_t = SOC_{t-1}(1 - \ell) + q_t^{\text{ch}}\Delta t - q_t^{\text{dis}}\Delta t \qquad \forall t \in \mathcal{T}. \tag{4.50}$$

Here, ℓ denotes the fractional standing-loss rate per time step.

The first time step is linked cyclically to the final time step to avoid a free initial buffer level:

$$SOC_1 = SOC_{|\mathcal{T}|}(1 - \ell) + q_1^{\text{ch}}\Delta t - q_1^{\text{dis}}\Delta t. \tag{4.51}$$

Over a full cycle, total charged heat therefore equals total discharged heat plus the accumulated standing losses, this closed-cycle energy balance is verified in Appendix E.1.

Buffer capacity limit The buffer state of charge is bounded by the installed thermal-buffer capacity:

$$0 \leq SOC_t \leq E^{\text{buffer}} \qquad \forall t \in \mathcal{T}. \tag{4.52}$$

Mutually exclusive buffer operation To prevent simultaneous charging and discharging of the buffer, a binary buffer-mode variable z_t is introduced:

$$q_t^{\mathrm{ch}} \leq M^{\mathrm{buffer}} z_t \qquad \forall t \in \mathcal{T}, \tag{4.53}$$

$$q_t^{\mathrm{dis}} \leq M^{\mathrm{buffer}} (1 - z_t) \qquad \forall t \in \mathcal{T}, \tag{4.54}$$

where $z_t \in \{0, 1\}$ and M^{buffer} is an upper bound on the buffer charge and discharge power, set to the maximum of the buffer energy-to-power ratio $E^{\mathrm{buffer}}/\Delta t$ and the maximum available flexible heat flow.

Non-served energy Non-served flexible heat demand is converted into an electrical-equivalent shortfall before entering the generic NSE expression:

$$e_t^{\mathrm{NSE}} = \frac{n_t^{\mathrm{HP}}}{COP} \Delta t \qquad \forall t \in \mathcal{T}. \tag{4.55}$$

4.3.6. Construction of the TDTR call-hour signal

The TDTR call-hour signal used in this thesis is derived from internal substation load data provided by Liander. These data consist of hourly net load measurements at the 110 kV and 150 kV substations operated by Liander, and are not publicly available. A substation operates as a bidirectional node in the electricity network. During periods of high demand, electricity flows from the higher-voltage level to the lower-voltage level to supply downstream consumers. Conversely, when local generation, such as solar PV, exceeds local demand, electricity flows upward from the lower-voltage level to the higher-voltage level. As a result, the hourly substation profile contains both positive values, representing net downstream consumption, and negative values, representing net feed-in to the upstream grid.

In the Dutch context, congestion at the substation level predominantly occurs during periods of high downstream demand, especially in winter, when power must be transported from the higher-voltage level to the lower-voltage level. Since TDTR restrictions are intended to manage this type of network loading, only the positive load values are used to derive ATR call hours. Negative values are excluded from the analysis, as they represent hours with net feed-in rather than downstream transport congestion, and therefore do not lead to TDTR call events.

Defining call hours: top 15% of positive-load hours The positive hourly substation load values are sorted in descending order to form a load duration curve. The highest 15% of these hours, corresponding to 1314 hours per year, are identified as ATR call hours during which the full conditional capacity is called.

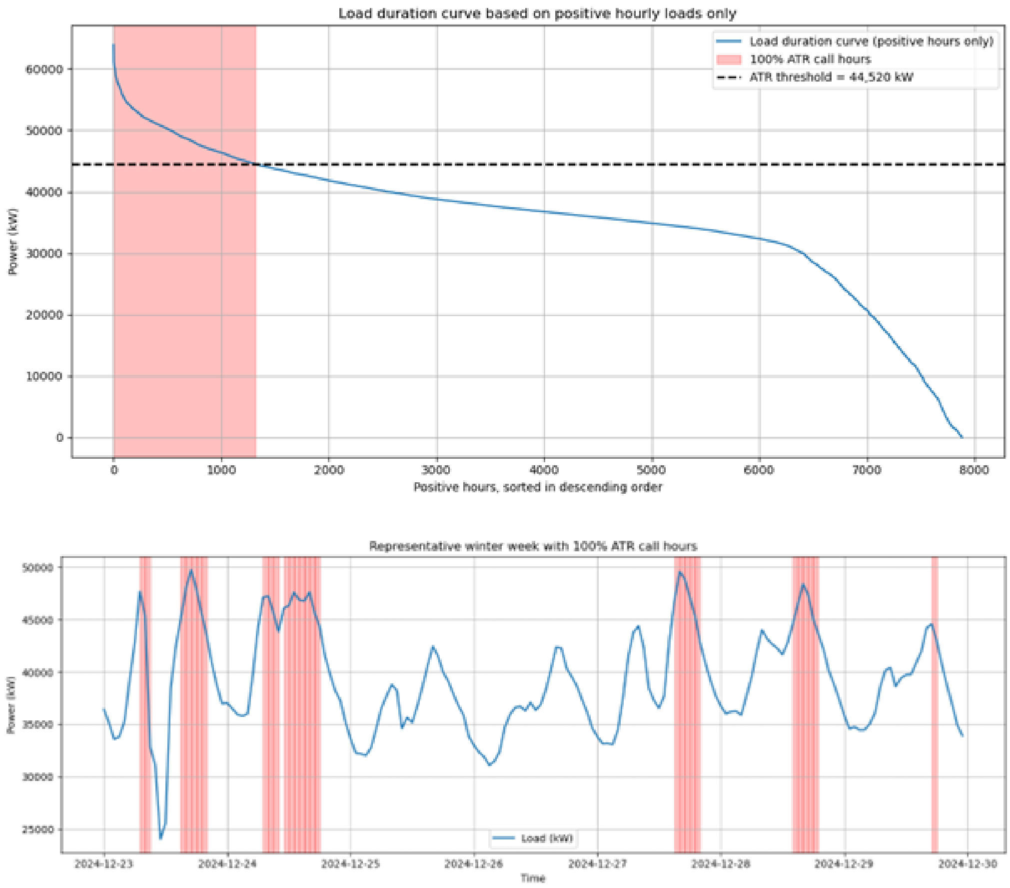

Figure 4.2: TDTR call-hour derivation for the Oudehaske 110 kV substation: the top 15% of positive-load hours (top) are called, resulting in the call blocks shown in a representative winter week (bottom). Monthly and hourly call-hour distributions across all four substation profiles are reported in Appendix D.

4.4. Evaluation indicators

The assessment rests on two families of key performance indicators (KPIs), reflecting the two perspectives this thesis keeps separate. *Customer-cost KPIs* capture the business case for the connected customer, and *grid-impact KPIs* capture whether TDTR participation actually relieves the network.

The customer-cost KPIs are the annual cost components of the electricity bill: commodity cost, energy tax, volumetric network cost, the monthly kW-max charge, the contracted-capacity cost, non-served energy cost, and any annualised asset-specific cost. Each is computed from the optimised load profile and the contract configuration exactly as defined in the economic framework of Appendix A, where the bill structure and tariff drivers are developed in full.

The grid-impact KPIs are the focus of the remainder of this section. Because the framework abstracts away from network simulation, grid impact is evaluated directly on the optimised load profile. The evaluation starts from two absolute load levels, the overall peak and the peak during call-hours, and from these derives four indicators: the absolute peak reduction, the weighted grid-pressure reduction, the share of high-load moments falling in call-hours, and the call-hour load reduction. The two peak levels describe how heavily the customer loads the grid in absolute terms, while the four indicators express, in relative terms, how far the optimised profile reduces that loading and how well it aligns with the congested hours. Of these, the call-hour load reduction is the primary measure, since it captures load reduction in exactly the hours the network is congested. Each is described below.

4.4.1. Peak load and peak load during call-hours

The evaluation reports two absolute peak load levels, both read directly from the optimised profile. The first is the overall peak load, the highest load reached at any time step:

$$P_s^{\text{peak}} = \max_t p_t^s. \tag{4.56}$$

The second is the peak load during the call-hours, the highest load reached within the local call-window:

$$P_s^{\text{peak,call}} = \max_{t \in \mathcal{T}^{\text{call}}} p_t^s, \tag{4.57}$$

where $\mathcal{T}^{\text{call}}$ is the set of call-hours. The overall peak determines the customer's monthly kW-max charge, while the call-hour peak indicates how heavily the customer loads the network during exactly the hours it is congested. A flexibility response that lowers the overall peak does not necessarily lower the call-hour peak, the gap between the two shows whether the customer's highest demand has been moved out of the congested window or merely rescheduled within it.

4.4.2. Absolute peak reduction

The absolute peak reduction compares the highest peak in the optimised scenario with the highest peak in the reference scenario, irrespective of timing:

$$P_{\text{red},s}^{\text{abs}} = 100 \cdot \left(1 - \frac{\max_t p_t^s}{\max_t p_t^{\text{ref}}}\right). \tag{4.58}$$

A higher value means that the customer's maximum peak demand is lower. The indicator measures peak reduction at the customer level, but does not account for the timing of the peak or its coincidence with congested hours.

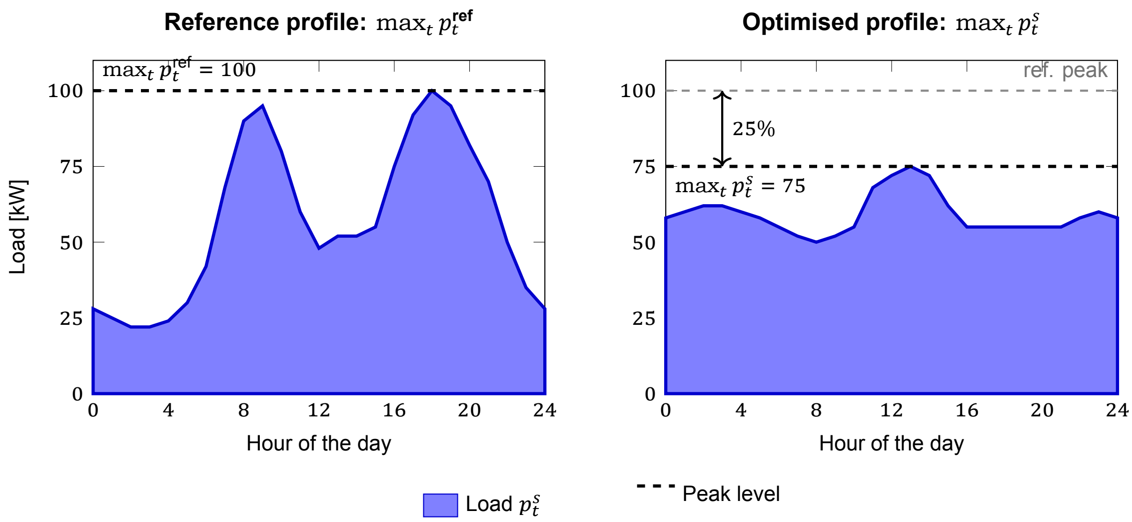

Figure 4.3: Illustrative effect of peak shaving on the absolute peak reduction. The indicator compares the highest load in each profile, irrespective of timing. For this example $\max_t p_t^{\text{ref}} = 100$ and $\max_t p_t^s = 75$, giving an absolute peak reduction of $P_{\text{red},s}^{\text{abs}} = 25\%$.

4.4.3. Weighted grid-pressure reduction

The weighted grid-pressure reduction measures how much of the customer's demand falls in hours the proposed time-dependent network tariff treats as critical. Each time step carries a weight w_t reflecting expected network loading, high for peak periods (up to 1.0), low otherwise. These are not set here but taken from the published weighting factors of the proposed tariff reform, which vary by month, hour, and day type (Appendix A.2). A higher reduction means more demand shifted out of high-weight hours. Because the weights are set at system level rather than per substation, this indicator captures alignment with the system-wide tariff signal rather than local congestion relief, which the call-hour-based indicators capture instead.

The weighted grid-pressure score of scenario s is defined as:

$$G_s^{\mathrm{w}} = \sum_{t=1}^{T} p_t^s w_t \Delta t, \tag{4.59}$$

and the reduction relative to the reference profile as:

$$G_{\mathrm{red},s}^{\mathrm{w}} = 100 \cdot \left(1 - \frac{G_s^{\mathrm{w}}}{G_{\mathrm{ref}}^{\mathrm{w}}}\right). \tag{4.60}$$

A positive value means demand has been shifted toward lower-weight hours, a negative value means the opposite. Figure 4.4 illustrates this for a representative working day in January, using the actual proposed tariff weights (see Appendix A.2). Both profiles serve the same total daily energy, but the optimised profile shifts demand away from the $w = 1.00$ peak hours, lowering G_s^{w} relative to $G_{\mathrm{ref}}^{\mathrm{w}}$.

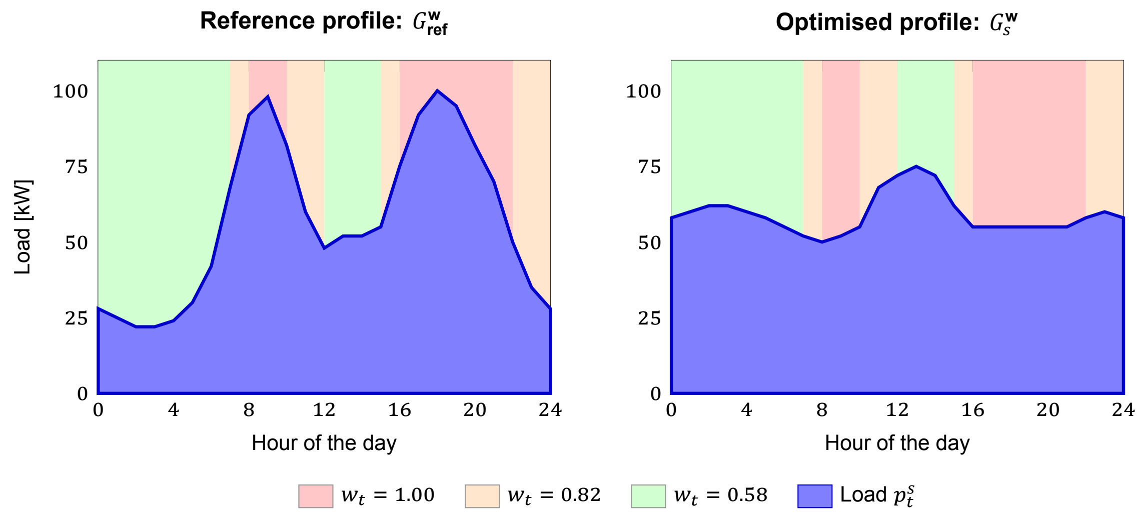

Figure 4.4: Illustrative effect of load shifting on the weighted grid-pressure score for a representative working day. For this example $G_{\mathrm{ref}}^{\mathrm{w}} \approx 1{,}067$ and $G_s^{\mathrm{w}} \approx 868$, giving a weighted grid-pressure reduction of $G_{\mathrm{red},s}^{\mathrm{w}} \approx 19\%$.

4.4.4. Share of high-load moments in call-hours

This indicator checks whether the customer's highest-load moments fall inside or outside the call-hours. Even when total energy use during call-hours is low, a few remaining peaks at the wrong moment can still stress the grid.

The top 5% of loads in scenario s are taken as the high-load moments:

$$H_s = \left\{t : p_t^s \geq Q_{95}(\{p_t^s\}_{t=1}^{T})\right\}. \tag{4.61}$$

The share of these moments falling in call-hours is:

$$S_s^{\mathrm{call}} = 100 \cdot \frac{|H_s \cap \mathcal{T}^{\mathrm{call}}|}{|H_s|}. \tag{4.62}$$

A low value means peaks occur outside call-hours (desired), a high value means peaks coincide with the moments the network operator considers critical. This is the only indicator that directly links the load profile to the timing of network constraints.

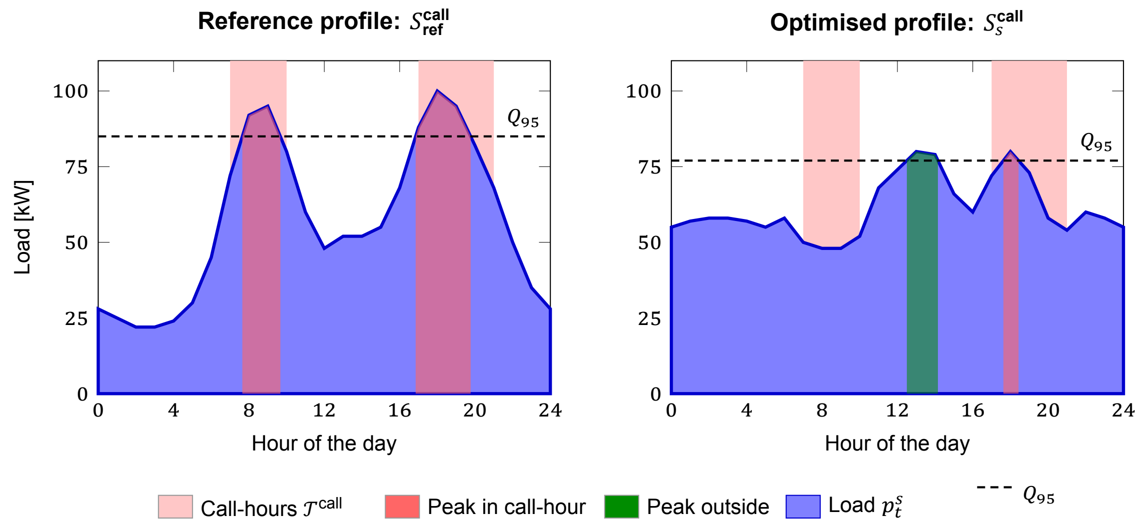

Figure 4.5: Illustrative effect of load shifting on the share of high-load moments (top 5%, above Q_{95}) falling in call-hours. In the reference profile all peaks fall inside the call-hours ($S^{\text{call}}_{\text{ref}} = 100\%$); the optimised profile shifts most to midday, leaving one in the evening call-hour, giving $S^{\text{call}}_{s} \approx 33\%$.

4.4.5. Call-hour load reduction

The call-hour load reduction quantifies to what extent the optimisation lowers demand specifically during the contractually constrained periods. Whereas the previous indicator looks at *where* the highest-load moments fall, this indicator looks at *how much* load is actually drawn during the call-hours themselves. Together, the two indicators capture both the redistribution and the absolute reduction of load in the constrained periods.

Let $\mathcal{T}^{\text{call}}$ denote the set of timesteps that fall within the call-hours. The total energy drawn during call-hours in scenario s is:

$$E^{\text{call}}_{s} = \sum_{t \in \mathcal{T}^{\text{call}}} p^s_t \, \Delta t, \tag{4.63}$$

and the call-hour load reduction relative to the reference profile:

$$E^{\text{call}}_{\text{red},s} = 100 \cdot \left(1 - \frac{E^{\text{call}}_{s}}{E^{\text{call}}_{\text{ref}}}\right). \tag{4.64}$$

A positive value indicates that energy demand during call-hours is reduced because flexibility is shifted from call-hours to non-call-hours. A negative value indicates the opposite. Figure 4.6 illustrates this for a representative winter day, with the shaded bands marking the TDTR call-hours. The reference profile delivers substantial energy during call-hours (high $E^{\text{call}}_{\text{ref}}$), while the optimised profile shifts this demand to non-call-hours (lower E^{call}_{s}).

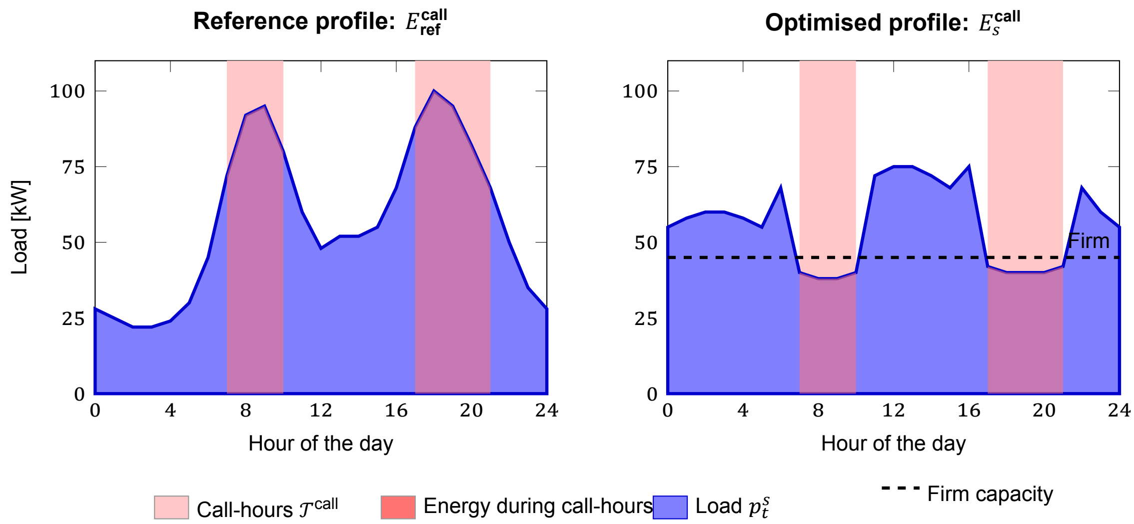

Figure 4.6: Illustrative effect of load shifting on the call-hour energy demand for a representative working day.For this example $E_{\text{ref}}^{\text{call}} \approx 704$ and $E_s^{\text{call}} \approx 318$, giving a call-hour load reduction of $E_{\text{red},s}^{\text{call}} \approx 55\%$.

4.5. Conclusion

This chapter has translated the literature-derived requirements of Chapter 3 into a concrete techno-economic optimisation framework. The correctness of this framework's implementation, together with checks that the buffer and battery dynamics behave as physically expected, is documented in Appendix E. Four methodological choices structure the remainder of the thesis.

First, the framework is formulated as an LP/MILP that takes the customer perspective on a single large energy consumer connection, with wholesale prices, network tariffs, the TDTR call-hour signal, and the reference demand profile as exogenous inputs, and operational flexibility as the endogenous decision space. This scope places the analysis on the user side of the bridge identified in Chapter 3 and ensures that the network restriction enters the model as a contractual condition rather than as a physical network simulation.

Second, the contract representation is reformulated as a stacked firm-plus-TDTR arrangement in which the firm-to-TDTR ratio is swept as an exogenous parameter rather than optimised. This produces a compliance-cost curve over the contract space rather than a single profitability verdict, which is the principal extension relative to Verhagen et al. (2025).

Third, the TDTR call-hour signal is derived from internal Liander substation load data, using the Oudehaske 110 kV profile as the default specification and three additional substations (Hemweg, Eerbeek, Zaltbommel) for sensitivity analysis. This grounds the analysis in empirically observed congestion patterns rather than a generic restriction regime.

Fourth, customer cost and grid impact are evaluated as complementary dimensions through five indicators. Annual cost components capture the customer business case, while peak reduction, weighted grid-pressure, share of high-load moments in call-hours, and call-hour load reduction together capture the network-side effects. The separation between the system-wide tariff signal (weighted grid-pressure) and the local congestion signal (call-hour-based indicators) makes explicit that alignment with one does not imply alignment with the other.

Together, these choices form a general framework that is parameterised separately for each sector archetype.

5

Case construction and modelling assumptions

The case construction specifies which flexible assets are modelled and how their reference demand is composed. For each sector case, a baseline electricity profile is combined with the flexibility available in that sector, whether from a separate electrified asset or from flexibility intrinsic to the core process itself.

Each case is built up in three modelling layers: a baseline electricity demand from an empirical or normalised sector profile, the sector's flexibility, and operational constraints that determine how that flexibility can respond to TDTR call-hours. For each case, the reference profile, taken without any optimisation and under a fully firm contract, serves as the baseline against which the flexibility and TDTR scenarios are compared in Chapter 6. For each case, the contract reference capacity P^{contract} is derived from the annual peak of the reference profile and adds a 10% contracting margin to the case.

5.1. Case 1: Cold storage

The cold-storage case represents a refrigerated warehouse in which electricity demand is mainly driven by the cooling installation. No separate add-on asset is added to the baseline profile, the cold-storage load itself is the flexibility asset. Flexibility is embedded in the refrigeration system, because the thermal inertia of the refrigerated space, stored goods, and cooling system allows compressor operation to be shifted within operational temperature limits. Stored products thereby act as passive thermal mass that delays temperature changes during compressor curtailment or load shifting (Akerma et al., 2020).

5.1.1. Reference profile construction

The construction of the cold-storage reference profile, including its baseline profile, scaling basis, and connection category, is summarised in Table 5.1.

Table 5.1: Construction of the cold-storage case

Element	Implementation in this study	Source / basis
Baseline profile	Normalised cold-storage electricity profile	SBI 52102
Scaling basis	Representative peak demand of 290 kW	EAN analysis and case Kahmann and Kleiburg (2023)
Connection category	MS / AC5 / <2 MVA	Based on the tariff classification in Table A.3

The resulting reference profile is shown in Figure 5.1. It displays a continuous cooling demand pattern with recurring short-term fluctuations around the inflexible base load, and an annual peak of approximately 290 kW.

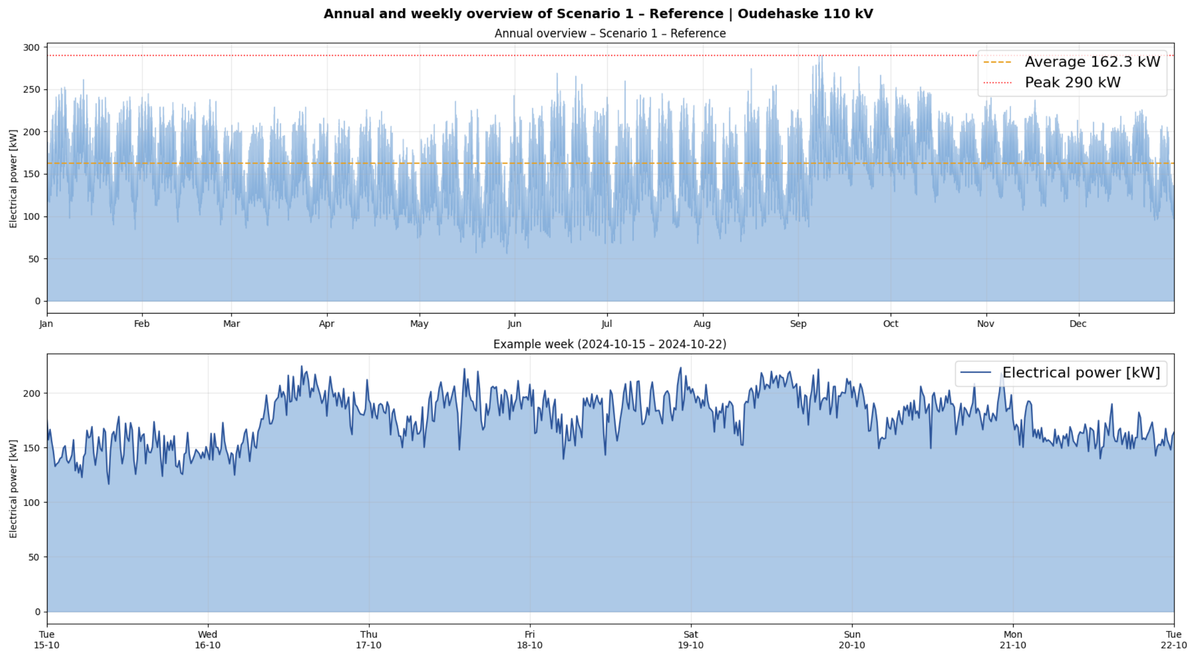

Figure 5.1: Annual and weekly overview of the cold-storage reference profile (Scenario 1).

5.1.2. Cooling and flexibility assumptions

The flexibility mechanism is based on *pre-cooling*: the refrigeration system can cool ahead of demand, storing “cold” in the thermal mass of the refrigerated space and the stored goods, and drawing on that stored cold later. Flexibility acts in one direction only: cooling can be brought forward and stored, but the cooling demand itself must be served at its original moment, either directly, from the buffer, or, against a penalty, as non-served energy. In addition, the flexible part of the cooling system may be switched off for at most four hours per day, representing the maximum temperature drift permitted within operational limits. The optimisation follows the thermal-buffer formulation introduced in §4.3.3, with the case-specific parameter values listed in Table 5.2.

Table 5.2: Scaling and flexibility assumptions for the cold-storage case

Parameter	Value	Source / basis
Maximum electrical demand P^{max}	290 kW	Kahmann and Kleiburg (2023)
Flexible capacity $P^{\mathrm{max,flex}}$	145 kW (50%)	CE Delft (2025)
Inflexible load share	50%	CE Delft (2025)
Pre-cooling horizon t^{shift}	2 hours	CE Delft (2025)
Buffer capacity $E^{\mathrm{CS,buf}}$	290 $\mathrm{kWh_e}$ (= $P^{\mathrm{max,flex}} \times t^{\mathrm{shift}}$)	Derived sizing rule, identical to the dairy case
Maximum flex-off duration $N^{\mathrm{off,max}}$	16 time steps (4 h/day)	CE Delft (2025)
Ramp limit R^{max}	10 kW per 15 min	Technical modelling assumption
Minimum flex-on threshold ϵ^{flex}	0.5 kW	Technical modelling assumption

Because the buffer represents the thermal mass of the refrigerated space itself rather than a separate storage asset, no buffer CAPEX is added, the cost of activating the flexibility is captured by the EMS cost (Table 5.3).

5.1.3. Cost assumptions

Table 5.3 summarises the cost assumptions used for the cold-storage case. These EMS cost assumptions are also used in the Logistics and Dairy cases.

Table 5.3: Cost assumptions for the cold-storage case

Cost component	Value / method	Source / basis
EMS CAPEX	5,000 €	Cost assumption
EMS lifetime	10 years	Cost assumption
EMS fixed OPEX	500 €/year	Cost assumption
Discount rate	5%	Cost assumption
Capital recovery factor (CRF)	0.1295	Computed with $r = 0.05$, $n = 10$
NSE penalty	10 €/$\mathrm{kWh_e}$	Anchored to sectoral value of lost load (de Nooij et al., 2007; Praktiknjo, 2016)

The NSE penalty represents the cost of cooling demand that cannot be served. Temporal shifting of the cooling load is already captured endogenously by the scheduling decision, so the NSE variable represents only the residual energy that cannot be shifted at all, that is, *shed* energy. Its cost is therefore not the compensation for shifting consumption (0.11–0.50 €/$\mathrm{kWh_e}$ (Common Futures et al., 2026)), which would understate it by an order of magnitude, but the value of lost load. A value of 10 €/$\mathrm{kWh_e}$ is used, at the lower end of the Dutch sectoral VOLL range (de Nooij et al., 2007; Praktiknjo, 2016). It acts as a feasibility safeguard: curtailment draws first on the thermal mass of the stored product, which tolerates several hours without immediate consequence (Ndraha et al., 2018), so residual NSE appears only under physical infeasibility.

No asset-parameter sensitivity is run for cold storage. Its flexibility is intrinsic to the cooling process and is already represented through the shifting constraints, installed cooling capacity, ramping limits, and maximum flex-off duration. Adding a separate flexibility asset would change the nature of the case rather than test the robustness of the existing mechanism.

5.2. Case 2: Logistics with EV fleet

The second case represents a large distribution-centre operator with an electrified vehicle fleet. The case combines the existing electricity demand of a logistics site with additional electricity demand from electric vehicle charging. The existing site demand represents non-EV activities such as warehouse operation and office use. The EV charging demand represents the electrification of part of the vehicle fleet and consists of two charging types: overnight depot charging and daytime fast charging.

This case is selected because it captures a key electrification pathway in the logistics sector. Unlike cold storage, where flexibility is embedded in the existing cooling load, the flexibility in this case is introduced through a new electrified asset. The relevant operational question is therefore whether charging demand can be shifted within vehicle-availability windows without violating the charging requirements of the fleet.

5.2.1. Reference profile construction

The construction of the logistics reference profile for this case is summarised in Table 5.4.

Table 5.4: Construction of the distribution-centre EV case

Element	Implementation in this study	Source / basis
Baseline profile	Normalised land-transport electricity profile representing non-EV site demand	SBI 52109
Scaling basis	Building electricity demand of a large logistics site	CE Delft (2025)
Depot charging	Overnight charging profile for vehicles returning to the site	ElaadNL (2026)
Daytime fast charging	High-power opportunity charging during operational windows	ElaadNL (2026)
Connection category	AC6 / >2 MVA	Tariff classification, see Table A.3

The electrified fleet consists of 71 vehicles: 4 vans, 41 rigid trucks and 26 tractor-trailer combinations. This represents partial electrification of a larger logistics fleet. Depot charging is used for the full electrified fleet, while daytime fast charging represents additional high-power charging during operational activity.

Table 5.5: Assumed vehicle fleet and electrification in the distribution-centre EV case

Vehicle type	Total fleet	Electric vehicles	Electric share	Role in charging profile
Vans	10	4	40%	Depot and daytime charging
Rigid trucks	150	41	27%	Mainly depot charging
Tractor-trailers	150	26	17%	Depot and fast charging
Total	310	71	23%	Combined EV charging demand

The resulting reference profile is shown in Figure 5.2.

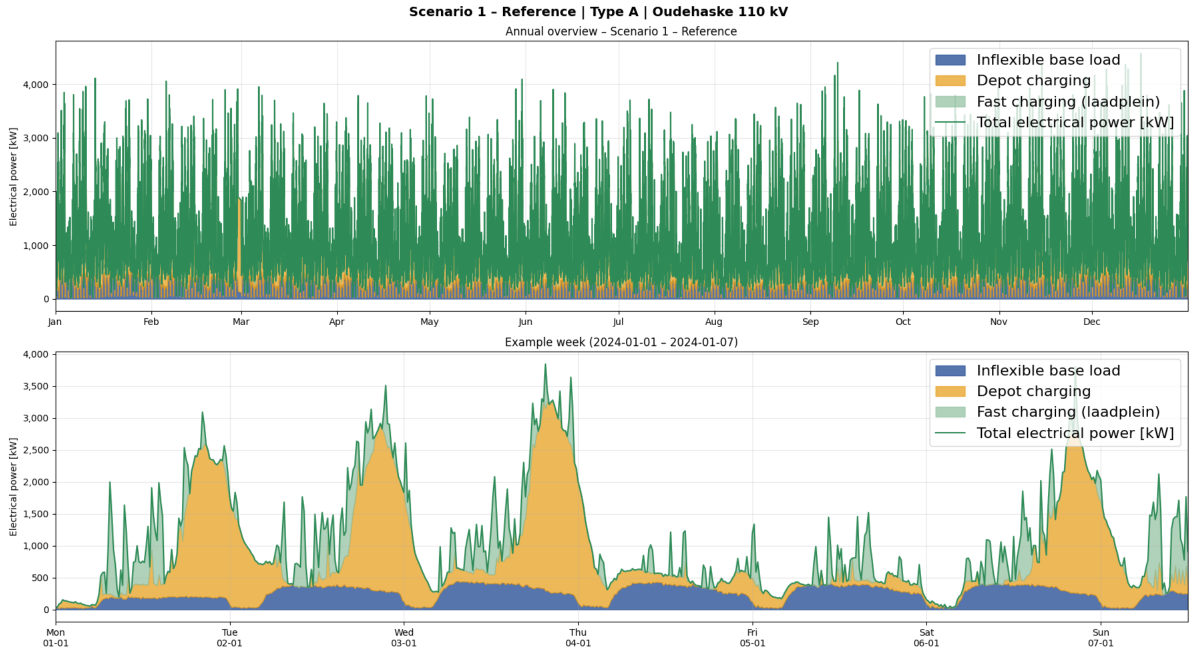

Figure 5.2: Annual and weekly overview of the logistics reference profile (Scenario 1).

5.2.2. Flexibility and cost assumptions

Flexibility is applied only to the EV charging components. The non-EV logistics baseload remains fixed. Depot charging and daytime fast charging are treated as separate flexible components because they have different operational constraints. Depot charging can be shifted within the overnight vehicle-availability window. Daytime fast charging is treated as flexible only within the operational window from 06:00 to 18:00, where it can be shifted within a short window around the original charging moment. Fast-charging demand outside this operational window is treated as inflexible and is added to the fixed load, consistent with the formulation in Section 4.3.4.

Table 5.6: Flexibility and cost assumptions for the distribution-centre EV case

Parameter	Value	Source / basis
Daytime fast-charging shift window	± 1 hour	CE Delft (2025)
Fast-charging points	10	Operational assumption
Fast-charger power	350 kW	Operational assumption
Maximum fast-charging power	3.50 MW	10×350 kW
Flexible fast-charging window	06:00–18:00	Operational daytime window assumption
Depot-charging window	22:00–06:00	CE Delft (2025)
Depot chargers / electric vehicles	71	CE Delft (2025)
Depot charger power	70 kW	Operational assumption
Maximum depot-charging power	4.97 MW	71×70 kW
Battery capacity per truck	400 kWh	Fleet-average assumption, truck-dominated fleet
Required charging session	10% to 100% SOC	Modelling assumption
Energy per full charging session	360 kWh	$(1.00 - 0.10) \times 400$ kWh
NSE penalty	5 €/kWh$_\mathrm{e}$	Anchored to transport-sector value of lost load (de Nooij et al., 2007; Praktiknjo, 2016)

The NSE penalty prevents infeasibility when charging cannot be fully scheduled within the available windows or under TDTR restrictions. As shifting charging to off-peak hours is captured endogenously, the NSE variable represents only non-shiftable, *shed* energy: a charging session that cannot be completed and a missed or curtailed trip. It is therefore anchored to the value of lost load rather than to shifting compensation (de Nooij et al., 2007; Praktiknjo, 2016). A value of 5 €/kWh$_\mathrm{e}$ is used, below the cold-storage and dairy values because a partially served charging session is an operational disruption rather than the loss of a continuous, interruption-sensitive process.

The number of trucks requiring charging per night is bounded through Eq. (4.43), using $E^{\mathrm{charge}} = 360$ kWh per full charging session ($(1.00 - 0.10) \times 400$ kWh) and $N^{\mathrm{EV}} = 71$ electric vehicles.

5.2.3. Battery sensitivity

The logistics case already contains substantial flexibility through smart charging, but that flexibility is tied to vehicle availability: charging can only be shifted within the windows in which vehicles are connected. During a call-hour in which few or no vehicles are plugged in, smart charging cannot respond. A stationary battery is tested here precisely because it breaks this dependency: it can discharge during call-hours regardless of vehicle presence, providing compliance capability that charging flexibility alone cannot. This is why the battery is examined for logistics specifically and not as a generic add-on, it relieves the binding constraint of this case (the availability windows), whereas the dairy and cold-storage cases are limited by process-bound flexibility rather than by availability timing.

To avoid basing the conclusion on a single arbitrary size, three batteries are tested that share the same power rating but differ in storage duration: Battery X (two-hour), Battery Y (four-hour), and Battery Z (eight-hour). Spanning these durations tests whether TDTR performance is constrained more by short-duration peak shaving or by longer-duration energy shifting during extended call blocks. The cost and performance assumptions are listed in Table 5.7 and the operational formulation in Eq. (5.1), the results across the three durations are reported in Chapter 6.

5.2.4. Battery cost assumptions

Table 5.7 lists the cost and performance assumptions for the battery energy storage system (BESS). Three configurations are evaluated, sharing the same power rating but differing in storage duration: Battery X (two-hour), Battery Y (four-hour), and Battery Z (eight-hour).

Table 5.7: Battery (BESS) cost and performance assumptions

Parameter	Value	Source / basis
Investment (CAPEX)	400 €/kWh nominal	Assumption
Economic lifetime	10 years	Assumption
Discount rate	5.0%	Assumption
Annuity factor (CRF)	0.1295	Derived from rate and lifetime
Fixed O&M	25 €/kWh nominal/year	Assumption
Degradation cost	0.05 €/kWh throughput	Assumption
Charge/discharge efficiency	0.922 (each)	Assumption
Round-trip efficiency	0.850	Derived ($\eta_{ch} \cdot \eta_{dis}$)
Battery scenarios (same power, increasing duration)		
Battery X	1,000 kWh / 500 kW	Two-hour system
Battery Y	2,000 kWh / 500 kW	Four-hour system
Battery Z	4,000 kWh / 500 kW	Eight-hour system

Operational battery formulation The battery operates under a standard storage formulation. The state of charge evolves with the charge and discharge power and the one-way efficiency $\eta = 0.922$ applied on both sides:

$$soc_t = soc_{t-1} + \eta\, p_t^{\mathrm{ch}} \Delta t - \frac{p_t^{\mathrm{dis}}}{\eta} \Delta t \qquad \forall t \in \mathcal{T}. \tag{5.1}$$

Charge and discharge power are bounded by the rated power of the selected configuration, and the state of charge by its energy capacity. The initial state of charge is set to 50% of the energy capacity, and a cyclic end condition $soc_{|\mathcal{T}|} \geq soc_{\mathrm{init}}$ rules out end-of-horizon depletion. No binary mode variable is used: simultaneous charging and discharging is rendered economically unattractive by the round-trip loss, and its absence in the optimised schedules is verified ex post in Appendix E.2.

5.3. Case 3: Dairy processing with industrial heat pump

The dairy case represents a large food-processing customer that electrifies part of its process heat demand with an industrial heat pump. The case combines a fixed dairy-process electricity profile with an additional heat-pump load. The fixed profile captures existing process electricity demand, while the heat-pump add-on captures the extra electricity demand caused by process-heat electrification.

5.3.1. Reference profile construction

The construction of the dairy reference profile, including its non-flexible dairy-process demand, the heat-pump add-on and its sizing, and the connection category, is summarised in Table 5.8.

Table 5.8: Construction of the dairy heat-pump case

Element	Implementation in this study	Source / basis
Baseline profile	Normalised dairy-sector electricity profile	SBI 105
Scaling basis	Representative dairy producer connected to the regional grid	Liander case data
Add-on asset	Industrial heat pump	Liander dairy heat-pump profile
Add-on sizing	Based on process electricity, process heat, and electrified heat demand	CE Delft (2025)
Connection category	AC6 / >2 MVA	Based on the tariff classification in Table A.3

The resulting reference profile is shown in Figure 5.3. It combines the inflexible dairy-process demand with the reference heat-pump demand, producing a relatively stable profile, as the dairy process requires a continuous supply of heat and only part of the heat-pump demand can vary over time.

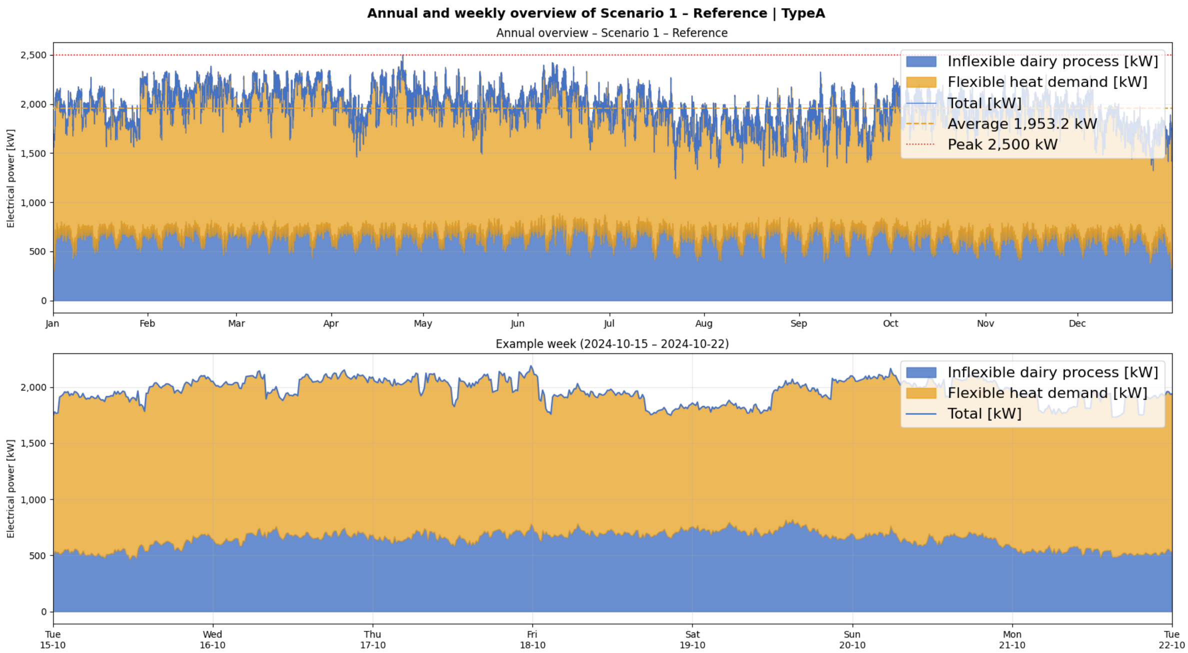

Figure 5.3: Annual and weekly overview of the dairy heat-pump reference profile (Scenario 1).

5.3.2. Heat-pump and flexibility assumptions

The useful process heat demand is converted into electricity demand using a COP of 4.02, based on the dairy heat-pump application analysed by Cox et al. (2022). This value is consistent with the temperature levels reported in the RVO industrial heat-pump case for dairy production (Rijksdienst voor Ondernemend Nederland, 2024). The resulting heat-pump electricity profile is used as the add-on load in the optimisation.

Table 5.9: Scaling and flexibility assumptions for the dairy heat-pump case

Parameter	Value	Source / basis
Dairy-process electricity	5,850 MWh/year	Liander case data
Heat-pump electricity	11,260 MWh/year	Model calculation using COP = 4.02
Combined annual demand	17,110 MWh/year	Model calculation
COP	4.02	Cox et al. (2022)
Heat-pump capacity $P^{\mathrm{HP,max}}$	1,725 $\mathrm{kW_e}$	Peak of the reference heat-pump profile
Flexible heat demand	10%	CE Delft (2025)
Flexibility horizon t^{shift}	2 hours (pre-production, buffer sizing basis)	CE Delft (2025)
Buffer capacity E^{buffer}	1,387 $\mathrm{kWh_{th}}$ ($= 0.10 \times 1{,}725 \times 4.02 \times 2$)	Exogenous sizing rule, see §5.3.3
Ramp limit R^{HP}	34.5 $\mathrm{kW_e}$ per 15 min	20% of flexible HP capacity per time step ($0.20 \times 0.10 \times P^{\mathrm{HP,max}}$)
Buffer loss	2.5%/hour	CE Delft (2025)

5.3.3. Cost assumptions

Table 5.10: Cost assumptions for the dairy heat-pump case

Cost component	Value / method	Source / basis
Flexible heat share	10% of heat-pump demand	CE Delft (2025)
Thermal-buffer capacity	$f_{\mathrm{flex}} \times P^{\mathrm{HP}}_{e,\mathrm{max}} \times COP \times t_{\mathrm{shift}}$	Exogenous sizing rule (see below)
Buffer CAPEX	50 €/$\mathrm{kWh_{th}}$	Conservative upper-bound estimate based on TTES literature (Energy Innovation NL, n.d.)
Buffer lifetime	15 years	Cost assumption
Discount rate	5%	Cost assumption
Fixed O&M	1.5% of CAPEX/year	Cost assumption
Annualised buffer cost factor	11.13% of CAPEX/year	CRF plus fixed O&M
NSE penalty	15 €/$\mathrm{kWh_e}$	Upper end of sectoral value of lost load, reflecting continuous food-process demand (de Nooij et al., 2007; Praktiknjo, 2016)

The thermal buffer is sized exogenously from the flexible heat share rather than optimised, with a capacity holding two hours of flexible heat production. Flexibility acts in one direction only: flexible heat can be produced ahead of demand and stored until needed, production cannot be deferred, because the process-heat demand must be served continuously.

The NSE penalty represents shed process heat, the residual demand that buffering cannot cover, and is anchored to the value of lost load (de Nooij et al., 2007; Praktiknjo, 2016). A value of 15 €/$\mathrm{kWh_e}$ is used, at the upper end of the sectoral VOLL range and above the other two cases because dairy is a continuous food process: hygiene constraints and low interruption tolerance raise the consequence of non-served energy.

5.3.4. Flexible-share sensitivity

The binding limitation in the dairy case is the small share of heat-pump demand that can be shifted through the thermal buffer. The base case assumes 10% of the heat-pump load is flexible, in line with the share that CE Delft reports as achievable for a continuous industrial process (CE Delft, 2025). To test whether the weak TDTR performance of this archetype is caused mainly by limited flexibility volume rather than by the timing of the calls, the flexible share is raised from 10% to 20%, a more favourable assumption for a continuous process. The buffer capacity scales accordingly through the sizing rule of §5.3.3, the result is reported in Chapter 6.

5.4. The four substation call profiles

Four substation profiles were selected for this thesis, each representing a distinct temporal congestion pattern. The Oudehaske 110 kV substation is used as the default case for the main specification, because it represents the most recognisable and recurring demand-driven congestion pattern, with call-hours concentrated in the winter months and around the morning and evening peak. This makes it a relevant reference profile for the asset-level analysis. The other three profiles are applied identically to every case as a shared robustness dimension: each case is solved under all four profiles, varying the call-hour timing while holding everything else fixed, so that the comparison isolates the effect of the local congestion shape. The case-by-case results across the four profiles are reported in Chapter 6.

Although all four profiles contain 1,314 call-hours, corresponding to 15% of the year, their temporal distribution differs substantially. The full call-hour distributions by month, hour of day, and day of week are given in Appendix D.1–D.3, the four can be characterised as follows:

- **Type A – Oudehaske 110 kV: seasonal morning/evening congestion.** A clear winter-oriented pattern, with call-hours concentrated mainly in January–March and October–December and, within the day, around the morning and evening peak. This is the most recognisable and recurring demand-driven pattern, and is used as the default specification.
- **Type B – Hemweg 150 kV: concentrated industrial call blocks.** The lowest number of call-days, but with call-hours concentrated within those days. The pattern is strongest in winter and mostly falls on working-day daytime and evening hours, representing a dense industrial load area where congestion appears in intense blocks rather than spread across the year.
- **Type C – Eerbeek 150 kV: strongly seasonal autumn/winter congestion.** The strongest seasonal concentration, with many call-hours in autumn and early winter, especially around October and November, mainly in daytime and evening (some night) hours. Compared with Type A, it is less evenly spread across winter and more concentrated in a narrower part of the year.
- **Type D – Zaltbommel 150 kV: dispersed year-round exposure.** The highest number of call-days and the most dispersed pattern, with call-hours in winter but also in summer and autumn. In contrast to the other patterns, a lot of call-hours are in the night hours and during cheaper electricity moments (12:00-14:00).

These differences matter because the operational burden of TDTR is not determined only by the total number of annual call-hours. A profile with fewer but longer call blocks may be manageable for assets with sufficient short-term storage or shifting capacity, while a profile with many dispersed call-days may interfere more often with normal operation. Solving every case under all four profiles therefore tests whether the economic and technical attractiveness of TDTR is robust across congestion shapes, or whether it depends on the specific timing of the local profile.

6 Results

This chapter presents the optimisation results for each sector and its corresponding flexibility asset. The analysis compares **three scenarios**, which together isolate the effect of operational flexibility and of TDTR participation:

Firm-only (reference). The asset meets its demand under a conventional firm contracted-capacity connection, sized to cover its peak load, without exploiting temporal flexibility.

Flex. The asset's operation is cost-optimised: its flexible load is scheduled over time to minimise total annual cost, subject to the asset's physical and operational limits, but *without* any TDTR restrictions.

Flex+TDTR. Identical to Flex, but the schedule must additionally respect the TDTR call-hour restrictions at the optimal case-specific firm-to-TDTR ratio determined with the compliance-cost curve.

All case results in this chapter use the Oudehaske (Type A) call profile as the default and the sensitivity analysis (§6.4) repeats the optimisation across the other three substation profiles.

For each case, four outputs are reported for the default Oudehaske (Type A) profile:

1. **Compliance-cost curve** showing how the total annual cost changes across different firm-to-TDTR ratios. This curve is used to identify the cost-minimising contract structure for each case.
2. **Scenario dispatch overviews** for the two optimised scenarios (Flex and Flex+TDTR), illustrating how the asset schedules its load with and without TDTR call-hour restrictions. The Firm-only reference is not shown here, as it involves no scheduling decision.
3. **Annual cost summary** comparing all three scenarios, decomposed into commodity cost, energy tax, contracted-capacity cost, monthly kW-max charge, non-served energy (where relevant), montly kWh charge (only relevant at cold storage), and any annualised asset-specific costs such as buffer or EMS costs.
4. **Grid-impact indicators** reporting the six metrics defined in §4.4.

This part answers **SQ2**:
For each representative sector and corresponding flexibility asset, what is the cost-optimal firm-to-TDTR ratio, and how much does participation at that ratio save the customer and relieve the grid relative to a firm-only and a flex-only scenario?

6.1. Case 1: Cold storage

Key findings. The cold-storage customer reduces annual electricity costs from €214,203 to €202,548 with Flex + 45% TDTR, a 5.4% reduction relative to the firm-only reference (Table 6.1). Flexibility without TDTR reduces annual costs by 4.4%, the operational cooling flexibility accounts for the larger share of the reduction (€9,366 of the €11,655 total). The 45% TDTR contract adds the remaining saving by lowering the kW-contract on the TDTR share.

Why this works. Cold storage has intrinsic short-term flexibility because cooling demand can be shifted within the thermal inertia of the refrigerated space. This allows the optimiser to reduce compressor use during TDTR call-hours without a separate thermal-buffer investment: the cooling service requirement is fixed, while the timing of electricity use is adjustable.

Where the savings come from. The Flex scenario reduces costs through lower energy cost and a lower monthly kW-max charge. The 45% TDTR contract reduces the kW-contract from €8,766 to €4,821, a saving of €3,945. This is partly offset by the NSE penalty of €1,326 and a €98 increase in the kW-max charge, giving a net additional saving of €2,290 relative to Flex. The total improvement is 5.4%, against 4.4% for Flex alone.

Residual non-served energy. With an NSE penalty of 10 €/kWh, the NSE cost of €1,326 in the Flex + 45% TDTR scenario corresponds to 0.13 MWh of non-served cooling energy per year, or 0.009% of the total annual demand of 1,422 MWh. This is cooling that the optimiser chooses not to serve during the most constrained call-hours, after the shiftable flexibility within the thermal inertia has already been exhausted. At 0.009% of annual demand it is confined to a small number of call-hour intervals, so the result is a cost-optimal trade-off under the chosen penalty value rather than a systematic shortfall in the cooling service.

Grid impact. The annual peak decreases from 290.00 kW in the reference to 255.45 kW under Flex + TDTR, an absolute reduction of 11.91% (Table 6.2). The peak during call-hours falls from 268.38 kW to 175.45 kW. The Top 5% call indicator decreases from 6.22% to 0.00%. The TDTR constraint relocates the highest-load moments out of constrained hours rather than further reducing the annual peak.

6.1.1. Compliance cost curve

The compliance-cost curve identifies the contract split that minimises total annual cost under the TDTR restriction (Figure 6.1). The minimum is reached at a firm fraction of 55%, corresponding to a 45% TDTR contract.

At lower firm fractions the customer receives a larger contract discount, but the available firm capacity becomes too restrictive during call-hours. This raises compliance costs and, at the lowest firm fractions, non-served-energy costs. At higher firm fractions compliance is easier, but the TDTR discount is smaller and the kW-contract cost increases. The optimum lies between these two effects: enough firm capacity to avoid operational disruption, but enough TDTR capacity to reduce annual contract costs. Its position is consistent with the asset's physical flexibility: 50% of the cold-storage load is shiftable, and the compressor ramping constraints limit how quickly that share can be moved out of the call-hours. The firm fraction therefore cannot fall much further before the non-shiftable load and the ramping limits force non-served energy, which is why the cost-minimising point settles at a 55% firm share rather than at a lower one.

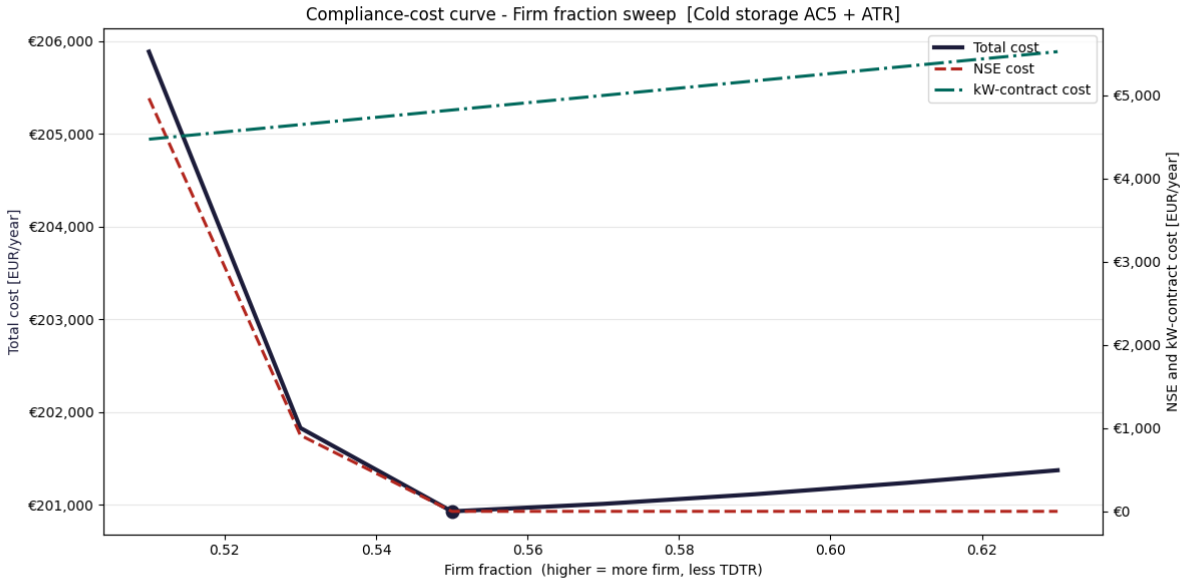

Figure 6.1: Compliance-cost curve for the cold-storage case (firm-fraction sweep).

6.1.2. Scenario-specific overviews

In the Flex scenario the flexible cooling load is redistributed within the operational constraints of the cold store, concentrating consumption in lower-price periods and lowering the annual peak to 255.89 kW (Figure 6.2, top plot). In the Flex + TDTR scenario (middle plot) the shaded bands mark the active TDTR call-hours, the optimiser shifts flexible cooling out of these hours and holds the load below the 175 kW firm limit where feasible. The day-ahead price (bottom plot) shows that, without TDTR, the flexible load already tracks low-price hours, under TDTR the schedule additionally avoids call-hours, including call-hours that coincide with low prices. As a result, the call-hour peak falls further than the annual peak.

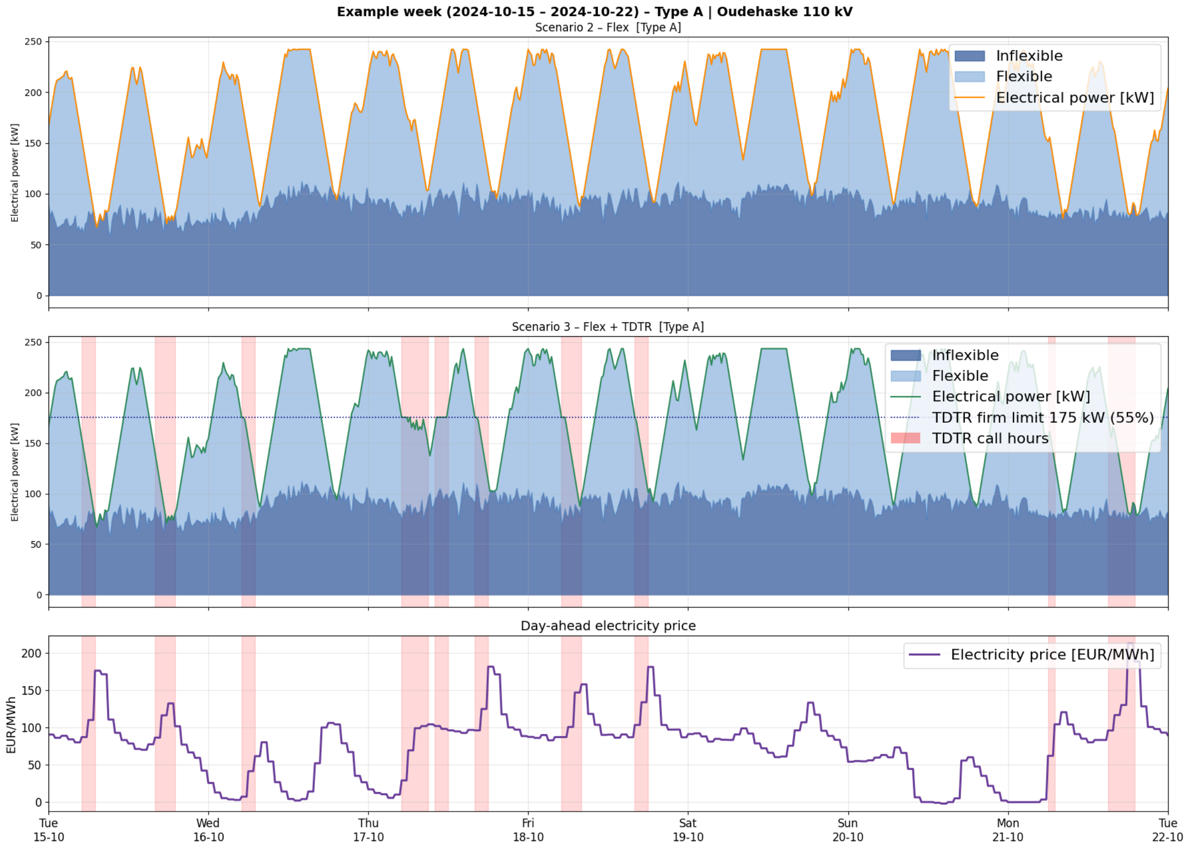

Figure 6.2: Weekly dispatch overview of Scenario 2 (Flex) and Scenario 3 (Flex + 45% TDTR) for the Oudehaske (Type A) profile, with the day-ahead electricity price (bottom plot).

6.1.3. Annual costs

In the Flex scenario, total annual cost falls from €214,203 to €204,838, a 4.4% reduction relative to the reference (Table 6.1). This is driven by lower energy cost and a lower monthly kW-max charge, while the kW-contract is unchanged because the contract basis is still fully firm.

The 45% TDTR contract reduces total annual cost to €202,548, 5.4% below the reference. The additional saving comes from the lower kW-contract, which falls from €8,766 to €4,821 because the TDTR share receives a discount on the contracted-capacity component. This €3,945 reduction is partly offset by the NSE penalty of €1,326 and a €98 increase in the kW-max charge, leaving a net additional saving of €2,290 over Flex. The operational cooling flexibility accounts for the larger share of the total reduction, TDTR contributes the remaining 1.0 percentage point.

Table 6.1: Annual cost summary of the cold-storage case with a 45% TDTR contract.

Metric	Reference	Flex	Flex + 45% TDTR
Energy cost [€]	109,403.51	100,117.59	100,357.09
Energy tax [€]	53,117.00	53,117.00	53,112.05
kWh transport [€]	32,140.41	32,140.41	32,137.41
kW-contract [€]	8,766.12	8,766.12	4,821.37
kW-max [€]	10,776.21	9,548.93	9,646.58
NSE cost [€]	0.00	0.00	1,326.00
EMS cost [€]	0.00	1,147.52	1,147.52
Total annual cost [€]	**214,203.24**	**204,837.57**	**202,548.02**
Cost improvement vs. reference [%]	**–**	**4.4%**	**5.4%**

6.1.4. Grid-impact indicators

The annual peak decreases from 290.00 kW in the reference to 255.89 kW under Flex and 255.45 kW under Flex + TDTR, absolute reductions of 11.76% and 11.91% respectively (Table 6.2). The difference between the two flexible scenarios in annual-peak reduction is 0.15 percentage points; the general peak-shaving is achieved by the operational flexibility of the cooling load.

The peak during call-hours falls from 268.38 kW in the reference to 255.89 kW under Flex and to 175.45 kW under Flex + TDTR. The TDTR-constrained optimisation shifts cooling demand out of locally constrained hours.

The Top 5% call indicator decreases from 6.22% in the reference to 1.94% under Flex and 0.00% under Flex + TDTR: in the TDTR case the highest-load moments fall entirely outside the call-hour set. The call-hour load reduction increases from 16.11% under Flex to 18.09% under Flex + TDTR.

The weighted grid-pressure reduction is 4.86% under Flex and 4.60% under Flex + TDTR. The lower WGPR in the TDTR case indicates that the TDTR constraint reduces load most during call hours, while relocating part of the load to moments of congestion or stress elsewhere in the electricity system (local versus system-wide).

Table 6.2: Grid-impact indicators for the cold-storage scenarios with a 45% TDTR contract.

Scenario	Peak [kW]	Peak call [kW]	Abs. red. [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Reference	290.00	268.38	0.00	6.22	0.00	0.00
Flex	255.89	255.89	11.76	1.94	4.86	16.11
Flex + TDTR	255.45	175.45	11.91	0.00	4.60	18.09

6.2. Case 2: Logistics with EV fleet

Key findings. Smart charging combined with a 74% TDTR contract reduces the annual electricity bill of the logistics customer by 23.4% relative to the firm-only reference, from €1,395.7k to €1,069.9k per year (Table 6.3). Smart charging without TDTR reduces annual costs by 13.2%, through lower energy costs and lower monthly kW-max charges. The 74% TDTR contract adds a further saving by reducing the kW-contract cost on the TDTR share.

Why this works. Most electricity demand in the logistics case relates to EV charging, which can be shifted within operational charging windows. Smart charging reduces the annual peak from 4,573.15 kW to 3,937.89 kW. Adding TDTR shifts charging demand out of call-hours: the peak during call-hours falls from 4,573.15 kW in the reference to 1,307.92 kW under smart charging + TDTR (Table 6.4), held just below the 1,308 kW firm limit. The Top 5% call indicator decreases from 33.05% in the reference to 0.00% under smart charging + TDTR, so the highest-load moments fall outside the call-hour set.

Where the savings come from. The 23.4% reduction in the smart charging + 74% TDTR scenario is driven by three components. Energy cost decreases from €822.7k to €670.9k because charging is shifted to lower-price hours. The kW-max charge decreases from €302.8k to €258.9k because smart charging reduces the monthly peak. The kW-contract cost decreases from €240.9k to €62.6k because 74% of the contract is assigned to TDTR. These reductions outweigh the NSE cost of €47.0k and the EMS cost of €1.1k.

Residual non-served energy. With an NSE penalty of 5 €/kWh, the NSE costs of €14.0k (smart charging) and €47.0k (smart + 74% TDTR) correspond to 2.8 MWh and 9.4 MWh of non-served charging energy per year. Relative to the annual electricity demand of 9,479 MWh, this is 0.03% and 0.10% of total demand respectively.

In production terms, the non-served energy is charging that the optimiser chooses not to deliver because fully serving it would require additional firm capacity or a higher monthly peak. At an energy of 360 kWh per full charging session, 9.4 MWh corresponds to about 26 full-charge equivalents per year, distributed as partial shortfalls across multiple nights rather than as complete missed charges. In operational terms this means a number of vehicles leave with a lower state of charge or require a rescheduled top-up. The result is a cost-optimal trade-off under the chosen penalty value, not evidence that unserved charging is acceptable in practice.

6.2.1. Compliance cost curve

The compliance-cost curve reaches its minimum at a firm fraction of 26%, corresponding to a 74% TDTR contract (Figure 6.3). The logistics case operates with a large TDTR share because EV charging can shift peaks away from call-hours, and a large part of the charging demand occurs at night, when TDTR restrictions are typically not active. The customer can therefore reduce firm contracted capacity while keeping compliance costs limited. At lower firm fractions the contract discount is larger, but firm capacity becomes too restrictive during call-hours, raising NSE costs, at higher firm fractions compliance is easier but the TDTR discount is smaller.

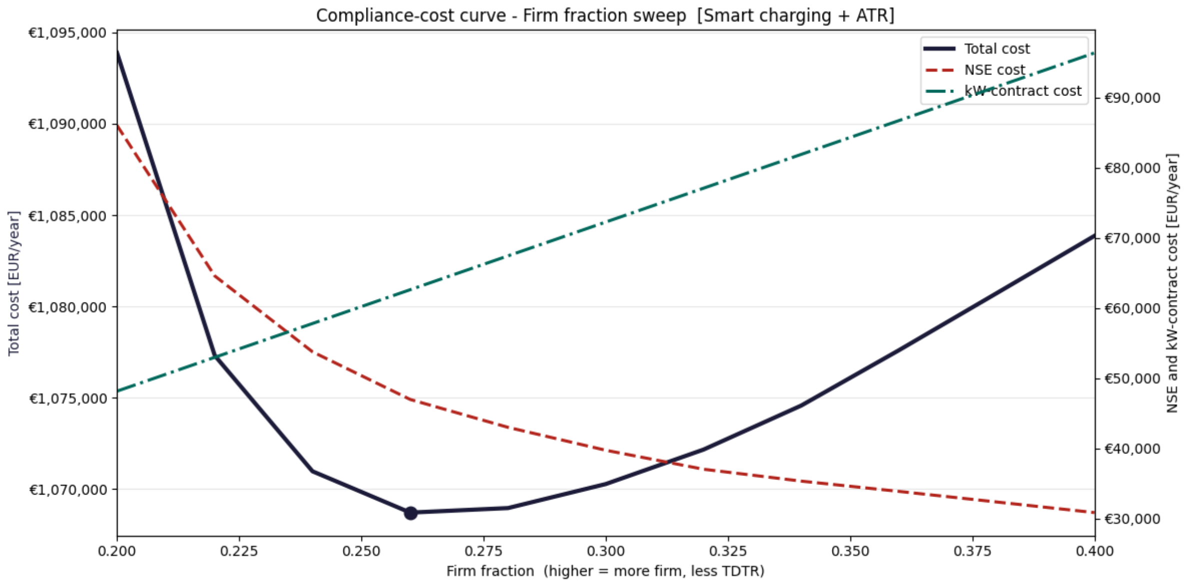

Figure 6.3: Compliance-cost curve for the logistics case (firm-fraction sweep).

6.2.2. Scenario-specific overviews

In the Flex (smart-charging) scenario, charging demand is shifted within the available charging windows, reducing the high and irregular peaks of the reference (Figure 6.4, top plot). In the Flex + 74% TDTR scenario (middle plot) the shaded bands mark the active call-hours; the optimiser shifts depot and fast-charging load out of these hours and holds the total load below the 1,308 kW firm limit where feasible. The day-ahead price (bottom plot) shows that smart charging already concentrates depot charging in low-price periods, under TDTR the schedule additionally avoids call-hours, including call-hours that coincide with low prices.

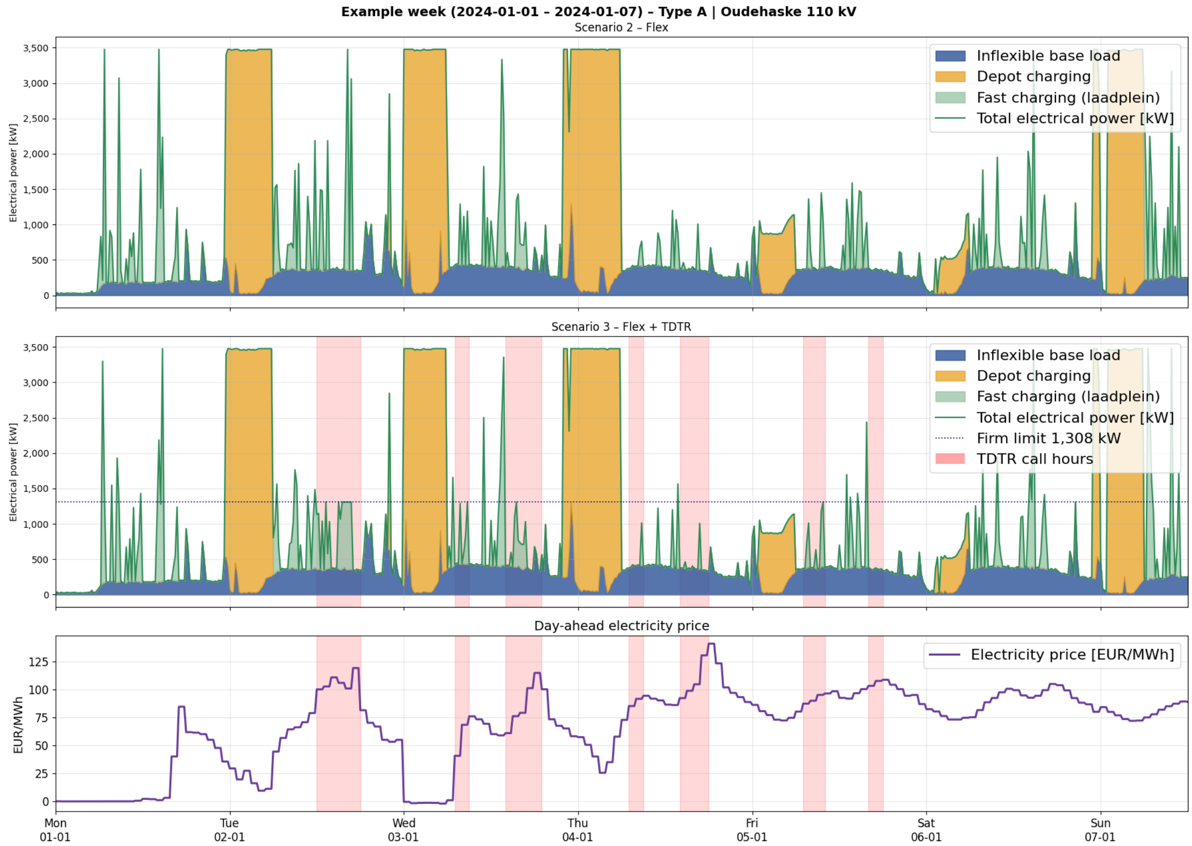

Figure 6.4: Weekly dispatch overview of Scenario 2 (Flex / smart charging) and Scenario 3 (Flex + 74% TDTR) for the Oudehaske (Type A) profile, with the day-ahead electricity price (bottom plot).

6.2.3. Annual costs

In the smart-charging scenario, total annual cost falls from €1,395.7k to €1,210.9k per year, a 13.2% reduction relative to the reference (Table 6.3), through lower energy costs and lower monthly kW-max charges. The 74% TDTR contract reduces total annual cost to €1,069.9k, 23.4% below the reference. This is the lowest-cost configuration in the logistics case. The driver is the reduction in kW-contract cost, from €240.9k to €62.6k, because the TDTR share receives a discount on the contracted-capacity component.

Table 6.3: Annual cost summary of the logistics smart-charging case for the Oudehaske 110 kV substation (Type A call profile).

Metric	Reference	Smart charging	Smart + 74% TDTR
Energy cost [k€]	822.7	670.6	670.9
Energy tax [k€]	29.4	29.3	29.3
kW-contract [k€]	240.9	240.9	62.6
kW-max [k€]	302.8	255.0	258.9
NSE cost [k€]	0.0	14.0	47.0
EMS cost [k€]	0.0	1.1	1.1
Total annual cost [k€]	**1,395.7**	**1,210.9**	**1,069.9**
Cost improvement vs. reference [%]	**–**	**13.2**	**23.4**

6.2.4. Grid-impact indicators

The annual peak decreases from 4,573.15 kW in the reference to 3,937.89 kW under smart charging, an absolute reduction of 13.89% (Table 6.4). Under smart charging + TDTR the annual peak is 4,019.19 kW, an absolute reduction of 12.11%, so the TDTR-constrained solution does not reduce the annual peak as far as smart charging alone.

The peak during call-hours decreases from 4,573.15 kW in the reference to 3,937.89 kW under smart charging and to 1,307.92 kW under smart charging + TDTR. The call-hour load reduction increases from 52.00% under smart charging to 56.08% under smart charging + TDTR. The Top 5% call indicator decreases from 33.05% in the reference to 5.82% under smart charging and to 0.00% under smart charging + TDTR, so the TDTR-constrained optimisation removes the highest-load moments from the call-hour set.

The weighted grid-pressure reduction is 6.93% under smart charging and 6.83% under smart charging + TDTR. The lower WGPR in the TDTR case indicates that the TDTR constraint reduces load most during call hours, while relocating part of the load to moments of congestion or stress elsewhere in the electricity system (local versus system-wide).

Table 6.4: Grid-impact indicators for the logistics scenarios with a 74% TDTR contract.

Scenario	Peak [kW]	Peak call [kW]	Abs. red. [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Reference	4,573.15	4,573.15	0.00	33.05	0.00	0.00
Smart charging	3,937.89	3,937.89	13.89	5.82	6.93	52.00
Smart charging + TDTR	4,019.19	1,307.92	12.11	0.00	6.83	56.08

6.3. Case 3: Dairy with industrial heat pump

Key findings. A 19% TDTR contract combined with 10% process-heat flexibility reduces the annual electricity bill of the dairy customer by 2.28% relative to the firm-only reference, from €1,673,267 to €1,635,110 per year (Table 6.5). Flexibility without TDTR reduces annual costs by 0.70%, the standalone value of thermal-buffer flexibility is limited. The TDTR contract adds value mainly by reducing the kW-contract cost on the TDTR share. The TDTR constraint shifts all Top 5% load moments out of call-hours, from 17.35% in the reference to 0.00% under Flex + TDTR + buffer (Table 6.6).

Why the cost effect is constrained. Only 10% of the heat-pump load is shiftable through the thermal buffer, while the underlying dairy-process electricity demand is fixed. The annual peak reduction is therefore 4.77%, and most of the TDTR value comes from the reduction in kW-contract cost rather than from operational peak reduction.

Where the savings come from. The Flex + 19% TDTR + buffer scenario reduces energy cost from €1,313,908 to €1,301,141 and kW-contract cost from €131,670 to €106,653. The monthly kW-max charge decreases from €174,649 to €165,724. These reductions outweigh the annual buffer cost of €7,720 and the NSE cost of €751, giving a total annual saving of €38,157.

Residual non-served energy. With an NSE penalty of 15 €/kWh, the NSE cost of €751 in the Flex + 19% TDTR + buffer scenario corresponds to 0.05 MWh of non-served process-heat energy per year, or 0.0003% of the total annual demand of 17,110 MWh. In production terms, this is heat-pump electricity that the optimiser does not deliver during the most constrained call-hours, met from the thermal buffer rather than served at the moment of demand. At this magnitude the non-served energy affects only a few call-hour intervals and does not require the dairy process itself to be reduced. The result is a cost-optimal trade-off under the chosen penalty value.

Grid impact. The annual peak decreases from 2,500.00 kW in the reference to 2,380.73 kW under Flex + TDTR + buffer. The peak during call-hours falls from 2,386.01 kW to 2,227.50 kW. The weighted grid-pressure reduction is 0.41% and the call-hour load reduction is 1.92%. The dairy case responds to TDTR restrictions, but the grid benefit is bounded by the small flexible share and the continuous process load.

6.3.1. Compliance cost curve

The compliance-cost curve reaches its minimum at a firm fraction of 81%, corresponding to a 19% TDTR contract (Figure 6.5). The dairy case requires a higher firm share than the cold-storage and logistics cases. At lower firm fractions the TDTR discount is larger, but the firm capacity becomes too restrictive for heat-pump operation, raising NSE costs, because the 10% thermal-buffer flexibility cannot shift enough process-heat demand away from call-hours.

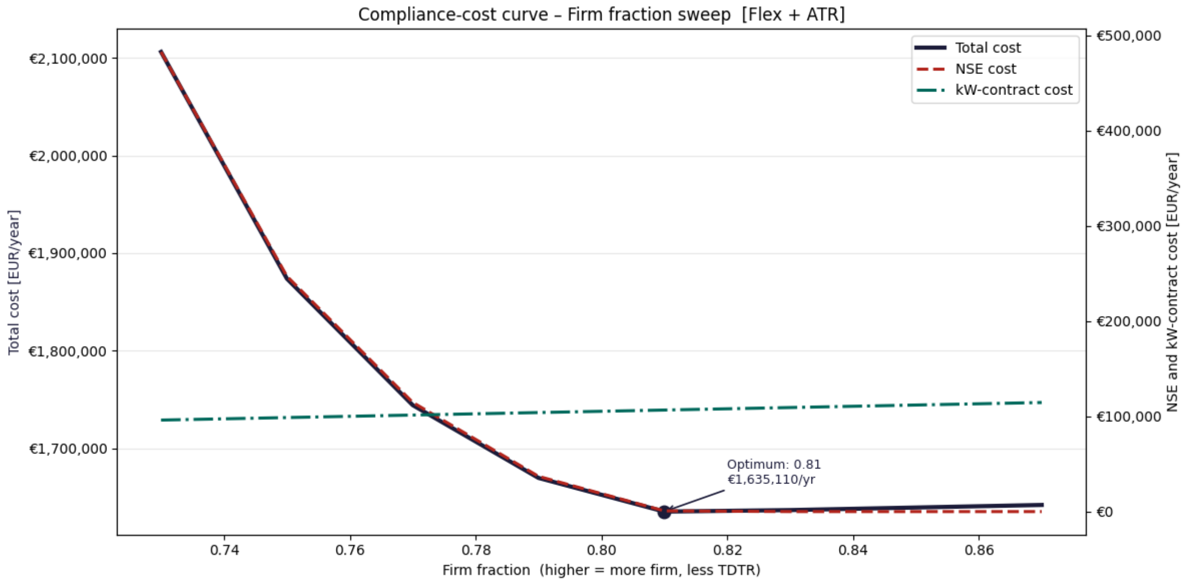

Figure 6.5: Compliance-cost curve for the dairy case with 10% flexibility (firm-fraction sweep).

6.3.2. Scenario-specific overviews

In the Flex scenario the flexible heat-pump demand is shifted within the thermal-buffer constraints, reducing the annual peak from 2,500 kW to 2,380.73 kW and smoothing short-term peaks in the heat-pump component, while the dairy-process load is fixed (Figure 6.6, top plot). In the Flex + TDTR scenario (middle plot) the shaded bands mark the active call-hours; the optimiser reduces heat-pump operation in these hours and uses the buffer to hold the load below the 2,200 kW firm limit. The change in the total profile is small, because only a small share of the heat demand is shiftable while most of the electricity demand is tied to continuous production. The day-ahead price (bottom plot) shows that the limited flexible share is dispatched towards lower-price hours, with the TDTR schedule additionally avoiding call-hours.

Figure 6.6: Weekly dispatch overview of Scenario 2 (Flex) and Scenario 3 (Flex + 19% TDTR + buffer) for the Oudehaske (Type A) profile, with the day-ahead electricity price (bottom plot).

6.3.3. Annual costs

The Flex + buffer scenario reduces annual costs from €1,673,267 to €1,661,557, a 0.70% improvement relative to the reference (Table 6.5), driven by lower energy cost and a lower monthly kW-max charge, partly offset by the annual buffer cost. The 19% TDTR contract reduces total annual cost to €1,635,110, a 2.28% improvement. The additional saving is driven by the kW-contract cost, which decreases from €131,670 to €106,653 because the TDTR share receives a discount on the contracted-capacity component, outweighing the buffer cost of €7,720 and the NSE cost of €751.

Relative to the logistics and cold-storage cases, the total cost improvement is bounded by the limited flexibility: 10% of the heat-pump load is shiftable through the thermal buffer, while most of the process electricity demand is fixed.

Table 6.5: Annual cost summary of the heat-pump case with 10% flexibility and a 19% TDTR contract for the Oudehaske 110 kV substation.

Metric	Reference	Flex + buffer	Flex + 19% TDTR + buffer
Annual peak [kW]	2,500.00	2,381.00	2,381.00
Total annual demand [MWh]	17,109.90	17,137.40	17,137.30
HP annual demand [MWh]	11,260.30	11,287.70	11,287.70
Energy cost [€]	1,313,907.86	1,302,190.29	1,301,140.91
Energy tax [€]	53,040.70	53,125.81	53,121.11
kW-contract [€]	131,670.00	131,670.00	106,652.70
kW-max [€]	174,648.90	165,703.83	165,724.49
NSE cost [€]	0.00	0.00	751.17
Buffer cost [€]	0.00	7,720.01	7,720.01
Total annual cost [€]	**1,673,267.46**	**1,661,557.46**	**1,635,110.39**
Cost improvement vs. reference [%]	**–**	**0.70%**	**2.28%**

6.3.4. Grid-impact indicators

The annual peak decreases from 2,500.00 kW in the reference to 2,380.73 kW under both Flex + buffer and Flex + TDTR + buffer, an absolute reduction of 4.77% (Table 6.6). The annual peak reduction is achieved by the thermal-buffer flexibility; adding the TDTR constraint does not reduce the annual peak further but shifts the highest-load moments out of call-hours, visible in the Top 5% call indicator, which falls from 17.35% in the reference to 9.49% under Flex + buffer and to 0.00% under Flex + TDTR + buffer.

The peak during call-hours falls from 2,386.01 kW in the reference to 2,303.81 kW under Flex + buffer and to 2,227.50 kW under Flex + TDTR + buffer. The weighted grid-pressure reduction is 0.40% under Flex + buffer and 0.41% under Flex + TDTR + buffer, the call-hour load reduction is 1.80% and 1.92% respectively. The dairy case responds to TDTR call-hours, but the effect is bounded by the shiftable share of the heat-pump load and the continuous nature of the dairy process.

Table 6.6: Grid-impact indicators for the heat-pump scenarios with 10% flexibility and a 19% TDTR contract for the Oudehaske 110 kV substation.

Scenario	Peak [kW]	Peak call [kW]	Abs. red. [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Reference	2,500.00	2,386.01	0.00	17.35	0.00	0.00
Flex + buffer	2,380.73	2,303.81	4.77	9.49	0.40	1.80
Flex + TDTR + buffer	2,380.73	2,227.50	4.77	0.00	0.41	1.92

This part will answer **SQ3**:
How do these outcomes depend on the substation call profile and on the assumed flexibility of the assets, and under which conditions do the customer-side and grid-side outcomes align or diverge?

6.4. Sensitivity analysis

The case results presented for cold storage, logistics, and dairy processing in Sections 6.1–6.3 are based on two main sets of assumptions: the temporal call-hour profile of the substation to which the customer is connected, and the flexibility characteristics of the modelled asset. The different substation call-hour profiles used in this analysis are defined in Section 4.3.6. In short, these profiles represent four distinct local congestion patterns: seasonal morning and evening congestion (Type A), concentrated industrial call-blocks (Type B), strongly seasonal autumn and winter congestion with some night calls (Type C), and dispersed, year-round exposure whose call-hours fall largely in cheap-price and night hours (Type D).

This section tests how robust the TDTR outcomes are to these assumptions. It first evaluates how the results change when the same customer archetype is connected to different local call-hour profiles. It then assesses selected asset-parameter sensitivities, such as changes in the available flexibility capacity or operational constraints. Finally, the combined results are interpreted from both the customer and network-operator perspectives, because a positive customer business case does not necessarily imply meaningful grid relief.

6.4.1. Customer-side and grid-operator-side outcomes per call profile

This section reads the scenario outcomes profile by profile. The four profiles are four distinct substation call patterns, Type A (Oudehaske), Type B (Hemweg), Type C (Eerbeek), and Type D (Zaltbommel), that differ in *when* the call hours fall and, decisively, in how those hours overlap with the low day-ahead-price hours and with each asset's flexible window, as seen more in detail in . For each profile the analysis follows the same two-step logic: first what flexibility alone achieves relative to the firm-only reference, and then what the TDTR constraint adds on top of that flexible scenario.

Because the NSE penalty is anchored to the value of lost load (5, 10, and 15 €/kWh for logistics, cold storage, and dairy respectively), non-served energy is expensive, so the optimiser will only accept a large TDTR share where compliance can be met by *shifting* load within the asset's flexible window, not by *shedding* demand. The cost-optimal firm fraction is therefore itself an alignment indicator: a low firm fraction (large TDTR share) signals that the call hours overlap the flexible, low-price window, while a high firm fraction (small TDTR share) signals that compliance would require shedding, which the optimiser avoids.

Table 6.7: Scenario outcomes for call profile Type A – Oudehaske (flex and flex + TDTR, relative to the firm-only reference).

Case	Scenario	Firm ratio [%]	Annual saving (+) [kEUR/year, %]	Peak reduction [%]	Weighted grid-pressure reduction [%]	Call-hour load reduction [%]
Cold storage	Flex	100	+9.4 (4.4%)	11.76	4.86	16.11
Cold storage	Flex + TDTR	55	+11.7 (5.4%)	11.91	4.60	18.09
Logistics	Smart charging	100	+184.8 (13.2%)	13.89	6.93	52.00
Logistics	Smart charging + TDTR	26	+325.9 (23.4%)	12.11	6.83	56.08
Dairy	Flex + buffer	100	+11.7 (0.7%)	4.77	0.40	1.80
Dairy	Flex + TDTR + buffer	81	+38.2 (2.3%)	4.77	0.41	1.92

Under the Oudehaske profile (Table 6.7), flexibility alone moves all three loads in the right direction: every flex case posts a positive saving and a positive call-hour reduction, so price-driven scheduling and the local call pattern are aligned. The cost-optimal TDTR share confirms this: logistics retains a 74% TDTR share (firm 26), which lifts its saving to 23.4% and its call-hour reduction from 52.00% to 56.08%. Cold storage takes a 45% TDTR share (call-hour reduction 16.11% to 18.09%) and dairy a 19% share. The weighted grid-pressure reduction is essentially unchanged across all three (logistics 6.93% to 6.83%), so on this profile TDTR works by sharpening behaviour during the call hours rather than by reshaping the load across the rest of the year.

Table 6.8: Scenario outcomes for call profile Type B – Hemweg (flex and flex + TDTR, relative to the firm-only reference).

Case	Scenario	Firm ratio [%]	Annual saving (+) [kEUR/year, %]	Peak reduction [%]	Weighted grid-pressure reduction [%]	Call-hour load reduction [%]
Cold storage	Flex	100	+9.4 (4.4%)	11.76	4.86	10.08
Cold storage	Flex + TDTR	57	+11.6 (5.4%)	9.66	4.42	11.79
Logistics	Smart charging	100	+184.8 (13.2%)	13.89	6.93	44.19
Logistics	Smart charging + TDTR	46	+293.1 (21.0%)	13.89	7.02	45.65
Dairy	Flex + buffer	100	+11.7 (0.7%)	4.77	0.40	0.95
Dairy	Flex + TDTR + buffer	85	+31.4 (1.9%)	4.26	0.40	0.96

The Hemweg profile (Table 6.8) differs from Oudehaske in that its call hours overlap the logistics flexible window less. The cost-optimal logistics TDTR share drops accordingly, from 74% to 54% (firm 46), and the call-hour reduction TDTR adds falls from +4.1 to +1.5 points (44.19% to 45.65%). The customer saving is still largely set by load type and connection size (21.0%, close to Type A), but the lower TDTR share is the quantitative readout of weaker alignment: the optimiser holds more firm capacity because more of the call hours fall outside the window where charging can be shifted without shedding. Cold storage (43% share) and dairy (15% share) show the same direction. Across A and B, the call-hours still overlap the flexible window and the high-price hours enough that compliance is met by shifting. Across A and B, the call-hours still overlap the flexible window and the high-price hours enough that compliance is met by shifting.

Table 6.9: Scenario outcomes for call profile Type C – Eerbeek (flex and flex + TDTR, relative to the firm-only reference).

Case	Scenario	Firm ratio [%]	Annual saving (+) [kEUR/year, %]	Peak reduction [%]	Weighted grid-pressure reduction [%]	Call-hour load reduction [%]
Cold storage	Flex	100	+9.4 (4.4%)	11.76	4.86	11.90
Cold storage	Flex + TDTR	61	+8.8 (4.1%)	5.43	4.69	13.14
Logistics	Smart charging	100	+184.8 (13.2%)	13.89	6.93	37.66
Logistics	Smart charging + TDTR	70	+251.6 (18.0%)	18.37	6.92	37.80
Dairy	Flex + buffer	100	+11.7 (0.7%)	4.77	0.40	1.28
Dairy	Flex + TDTR + buffer	81	+34.2 (2.0%)	4.77	0.41	1.34

Table 6.9 reports the Eerbeek (Type C) outcomes, the profile on which the realistic NSE penalty changes the conclusion most. The logistics TDTR share collapses to 30% (firm 70), and the call-hour reduction TDTR adds is +0.14 points (37.66% to 37.80%) with the weighted grid-pressure reduction essentially unchanged (6.93% to 6.92%). The customer saving still rises to 18.0%, because the 30% TDTR share carries a kW-contract discount, but the load profile is essentially unchanged: the customer collects the discount while delivering no additional grid relief. Cold storage makes the point even sharper: the best available TDTR split yields a 4.1% saving, *below* the 4.4% of flexibility alone, so a cost-minimising cold-storage customer would decline TDTR on this profile entirely. Under the earlier, low NSE penalty this profile appeared to be where TDTR did the most structural work (a 78% share driving call-hour reduction to 57%). That outcome was the model shedding energy to satisfy the constraint, once shedding is priced realistically, the optimiser declines, the TDTR share retreats, and the apparent structural correction disappears.

Table 6.10: Scenario outcomes for call profile Type D – Zaltbommel (flex and flex + TDTR, relative to the firm-only reference).

Case	Scenario	Firm ratio [%]	Annual saving (+) [kEUR/year, %]	Peak reduction [%]	Weighted grid-pressure reduction [%]	Call-hour load reduction [%]
Cold storage	Flex	100	+9.4 (4.4%)	11.76	4.86	-4.68
Cold storage	Flex + TDTR	61	+12.4 (5.8%)	11.59	4.48	-0.96
Logistics	Smart charging	100	+184.8 (13.2%)	13.89	6.93	-41.85
Logistics	Smart charging + TDTR	70	+253.0 (18.1%)	16.40	6.90	-41.50
Dairy	Flex + buffer	100	+11.7 (0.7%)	4.77	0.40	-0.68
Dairy	Flex + TDTR + buffer	83	+33.7 (2.0%)	4.77	0.40	-0.68

Zaltbommel (Table 6.10) is the adverse profile and must be read through the flex scenario. Here flexibility alone is not neutral but raises load during the call hours: the call-hour reduction is negative under flex for all three loads, reaching -41.85% for logistics, -4.68% for cold storage and -0.68% for dairy. The cost-minimising schedule, driven by day-ahead prices, concentrates load into the low-price hours that on this profile coincide with the substation's call hours, low price signals and network needs point in opposite directions.

The two flexible assets then diverge, and the divergence is the central evidence for the mechanism. Cold storage corrects most of the misalignment (-4.68% to -0.96%) because it complies through real thermal-buffer shifting at a 39% TDTR share. Logistics does not: its TDTR share is held at 30% (firm 70), the call-hour reduction moves from -41.85% to -41.50%, and the weighted grid-pressure reduction is essentially unchanged (6.93% to 6.90%). The logistics customer still records a positive saving (18.1%) from the contract discount, but the load during the call hours is unchanged. The difference is that compliance by cold storage is met by shifting, whereas compliance by logistics on this profile would require shedding charging sessions, which the NSE penalty makes uneconomic.

Synthesis across profiles

The four profiles together make the central point of the chapter. TDTR value is not an intrinsic property of a sector but emerges from the interaction between the customer load profile and the local call pattern. Under a realistic shedding cost this interaction is read directly from the cost-optimal TDTR share, which tracks alignment: a high share (logistics Type A, 74%) signals that compliance is met by shifting and is accompanied by real call-hour relief, while a low share (logistics Types C and D, 30%) signals that compliance would require shedding and is accompanied by near-zero added relief. Grid benefit follows the share. The exception that proves the rule is cold storage on Type D, where thermal-buffer flexibility lets a small TDTR share correct an adverse profile that logistics cannot.

This creates a mismatch between where TDTR is most attractive for the customer and where it is most valuable for the grid. On Types A–B, flexibility is already aligned with the call hours, so TDTR mainly fine-tunes an already favourable load response. On Types C–D the grid needs stronger correction, but the customer chooses only a small TDTR share because deeper compliance would require costly load shedding. TDTR therefore adds the least grid relief where it is most needed: a positive customer business case can coexist with no grid benefit at all, the logistics customer on Types C and D saving 18% while leaving the call-hour load unchanged.

For the network operator, the implication is not that this thesis determines how TDTR should be differentiated across substations. That would require a separate assessment of tariff design, regulatory constraints, and equal-treatment requirements. What the results do show is that a uniform roll-out cannot be interpreted as uniform grid value. Because cost-optimisation self-selects high TDTR shares onto well-aligned substations, where flexibility already delivers relief, the operator should actively screen where the product is likely to add incremental grid value and where additional steering is needed. On substations where congestion and market signals diverge, real relief requires a response beyond the customer's voluntary cost optimum, either through stricter firm-fraction requirements, payments tied to delivered call-hour reduction, or compensation for the non-served- energy cost. The firm fraction is therefore a possible operational lever, but using it shifts part of the misalignment cost from the grid to the customer. The limited relief on these profiles is thus not a technical impossibility, but a cost that the voluntary, discount-based contract does not by itself price.

6.4.2. Asset-parameter sensitivity

The final sensitivity dimension concerns the assumed characteristics of the flexibility assets themselves.

Dairy: flexibility share. The dairy sensitivity tests whether increasing the shiftable share of the heat-pump load changes the TDTR outcome. The annual-cost and grid-impact results for the 20% flexibility case are reported in Tables 6.11 and 6.12; the corresponding compliance-cost curve and weekly dispatch profile are shown in Appendix Figures F.1 and F.2.

Increasing the flexible share from 10% to 20% raises the cost-optimal TDTR share from 19% to 20% and the cost improvement relative to the firm-only reference from 2.28% to 2.46%. In absolute terms, the annual saving increases from €38,157 to €41,214. Additional thermal flexibility creates value, but the marginal improvement is bounded.

The grid-impact indicators follow the same pattern. The annual peak reduction increases from 4.77% to 5.92%, and the peak during call-hours decreases from 2,227 kW in the 10% case to 2,200 kW in the 20% case. The Top 5% of customer load moments are removed from call-hours in both TDTR cases.

The weighted grid-pressure reduction increases from 0.41% to 0.74%, and the call-hour load reduction from 1.92% to 3.55%.

The additional flexibility requires a larger thermal buffer: buffer capacity doubles from 1,386.7 to 2,773.4 $\mathrm{kWh_{th}}$, raising the annual buffer cost from €7,720 to €15,440. The 20% case records no NSE cost. The sensitivity confirms that more thermal flexibility improves TDTR performance with diminishing returns. For process-constrained sectors such as dairy, the TDTR business case is bounded by the available flexibility of the heat-pump system, rather than by the contract structure alone.

Table 6.11: Annual cost summary of the heat-pump case with 20% flexibility and a 20% TDTR contract.

Metric	Reference	Flex + buffer	Flex + 20% TDTR + buffer
Annual peak [kW]	2,500.00	2,352.00	2,352.00
Total annual demand [MWh]	17,109.90	17,163.86	17,164.62
HP annual demand [MWh]	11,260.25	11,314.20	11,314.96
Energy cost [€]	1,313,907.86	1,292,190.42	1,291,997.44
Energy tax [€]	53,040.70	53,207.96	53,210.32
kW-contract [€]	131,670.00	131,670.00	105,336.00
kW-max [€]	174,648.90	164,105.44	164,921.92
NSE cost [€]	0.00	0.00	0.00
Buffer cost [€]	0.00	15,440.02	15,440.02
EMS cost [€]	0.00	1,147.52	1,147.52
Total annual cost [€]	**1,673,267.46**	**1,657,761.36**	**1,632,053.21**
Cost improvement vs. reference [%]	**–**	**0.93%**	**2.46%**

Table 6.12: Grid-impact indicators for the heat-pump scenarios with 20% flexibility and a 20% TDTR contract.

Scenario	Peak [kW]	Peak call [kW]	Abs. red. [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Reference	2,500.00	2,386.01	0.00	17.35	0.00	0.00
Flex + buffer	2,351.89	2,351.89	5.92	5.31	0.72	3.26
Flex + TDTR + buffer	2,351.89	2,200.00	5.92	0.00	0.74	3.55

Logistics: battery addition. The battery sensitivity tests whether stationary storage changes the conclusion of the logistics smart-charging + TDTR case. Three BESS configurations of increasing size are compared, from the smallest (X) to the largest (Z). The results are summarised in Table 6.13. Across every profile, adding a battery lowers the customer's annual saving relative to the no-battery smart-charging + TDTR case. For Type A the saving falls from 23.4% (no battery) to 18.0% with Battery X, 12.4% with Battery Y, and 0.89% with Battery Z; the energy-oriented battery drives the business case to near zero or negative on all four profiles (0.89% to -6.28%). The annualised CAPEX, fixed O&M and degradation costs exceed the additional bill reduction the battery enables, so no battery configuration is part of the cost-optimal customer strategy. That strategy remains smart charging + TDTR without stationary storage. The reason to report the battery is therefore not the customer business case but the grid outcome, which the battery does improve, and the distinction between the two is the point. On the well-aligned profiles the battery raises the call-hour load reduction (Type A from 56.08% without a battery to 63.24–68.19%) and the weighted grid-pressure reduction (to 8.22–10.63%). At the same time the cost-optimal TDTR share *falls* (Type A from 74% to 78–80% TDTR), because the battery can cover the firm limit during call hours, so the customer no longer needs a large TDTR share to comply.

The battery thus substitutes for the TDTR contract as the compliance mechanism, delivering grid relief while reducing the contractual TDTR exposure. Type D again behaves differently and confirms the alignment mechanism. The call-hour load reduction is negative for all three battery sizes (-44.64%

to -40.25%), more negative than the no-battery case (-41.50%), because the battery is dispatched on price rather than on the local call signal: it charges in the low-price hours that coincide with the call hours, reinforcing rather than relieving the local peak. The weighted grid-pressure reduction stays positive (8.37–10.63%), so the battery still flattens the overall annual profile, but it cannot turn the call-hour effect positive on a profile where price and network needs oppose each other. In summary, the battery is a grid-side enhancement, not a customer-side strategy. Under the assumed costs it always weakens the private business case, so a cost-minimising logistics customer would not install it voluntarily. Its relevance is for the network operator: on well-aligned profiles (Types A–C) it deepens call-hour relief and can take over the compliance role from the TDTR share, while on a misaligned profile (Type D) it neither relieves the call hours nor pays for itself. If the operator values the additional relief, the implication is that it would have to be procured or compensated separately, because the TDTR contract alone does not make storage privately worthwhile.

Table 6.13: Key economic and grid-impact indicators for the logistics battery sensitivity analysis.

Type	Substation	Battery case	Firm [%]	Saving [%]	WGPR [%]	Call-hour red. [%]
A	Oudehaske 110 kV	Battery X	22	18.04	8.22	63.24
A	Oudehaske 110 kV	Battery Y	20	12.39	9.37	66.86
A	Oudehaske 110 kV	Battery Z	20	0.89	10.63	68.19
B	Hemweg 150 kV	Battery X	40	15.44	8.50	49.93
B	Hemweg 150 kV	Battery Y	38	10.05	9.71	51.44
B	Hemweg 150 kV	Battery Z	36	-1.22	10.93	52.15
C	Eerbeek 150 kV	Battery X	70	11.70	8.39	42.49
C	Eerbeek 150 kV	Battery Y	68	6.30	9.55	44.60
C	Eerbeek 150 kV	Battery Z	62	-4.62	10.65	44.77
D	Zaltbommel 150 kV	Battery X	68	11.72	8.37	-44.64
D	Zaltbommel 150 kV	Battery Y	66	6.36	9.52	-43.96
D	Zaltbommel 150 kV	Battery Z	62	-4.55	10.63	-40.25

6.4.3. Robustness to the day-ahead price assumption

The optimisation uses the 2024 Dutch day-ahead series. This section asks, without re-running the model, whether the alignment ordering of the profiles is a property of the chosen year or of the model mechanism. Two descriptive comparisons are used: the price level and intraday shape across 2023–2025 (Table 6.14), and the overlap between each profile's call-hours and the year's low- and negative-price hours (Table 6.15). On price level, 2024 is the middle of the three years (mean 77.3 against 95.8 in 2023 and 87.2 in 2025), and the intraday shape is stable, with the evening average above the midday average in every year (120.1 against 66.0 EUR/MWh in 2025). The only monotone trend is the number of negative-price hours, rising from 314 in 2023 to 458 in 2024 and 584 in 2025 (3.6% to 6.7% of all hours). The 2024 series is therefore not an extreme year, so the relevant test is whether the *ordering* of the profiles, rather than the price level of any single year, survives across 2023–2025.

Table 6.14: Dutch day-ahead price sensitivity for 2023–2025. Prices are in EUR/MWh; the 2025 quarter-hourly data are averaged to hourly values for comparability. Midday is defined as 10:00–16:59 and evening as 17:00–21:59.

Year	Mean	Median	P20	P80	Negative hours [h]	Negative share [%]	Midday avg	Evening avg
2023	95.8	99.2	69.1	129.2	314	3.6	76.7	124.1
2024	77.3	80.0	47.5	107.0	458	5.2	55.1	106.7
2025	87.2	90.8	60.0	119.0	584	6.7	66.0	120.1

The overlap comparison carries the structural result. The average call-hour price exceeds the average non-call-hour price in every profile and year, the only exception being Zaltbommel in 2023, where the two are effectively level (95.6 against 95.9 EUR/MWh). Cost-minimisation therefore avoids the call-hours on price grounds alone, independently of the year, the condition under which voluntary, price-driven behaviour relieves the call hours rather than working against them. The endpoints of the alignment ordering are preserved throughout: the gap between the average call-hour and non-call-hour price is widest for Oudehaske (Type A; 121.4 against 81.2 EUR/MWh in 2025) and narrowest for Zaltbommel (Type D), with Hemweg and Eerbeek (Types B and C) close together in between, matching the scenario results. The ordering is thus a property of the price–call overlap, not of the single year used to generate it.

Table 6.15: Overlap between TDTR call-hours and low or negative day-ahead prices for 2023–2025. Low-price overlap is defined relative to the lowest 20% of prices within each year.

Call profile	Year	Avg call price [EUR/MWh]	Avg non-call price [EUR/MWh]	Low-price overlap [%]	Negative-price overlap [%]
Type A – Oudehaske	2023	118.6	91.8	11.4	0.9
	2024	116.5	70.4	1.9	0.0
	2025	121.4	81.2	3.4	0.9
Type B – Hemweg	2023	110.8	93.2	17.9	3.4
	2024	111.6	71.3	3.8	0.3
	2025	107.0	83.7	12.9	4.3
Type C – Eerbeek	2023	112.8	92.8	13.7	0.8
	2024	110.7	71.4	7.3	0.9
	2025	112.6	82.7	5.9	1.1
Type D – Zaltbommel	2023	95.6	95.9	23.3	3.2
	2024	96.0	74.0	5.9	0.4
	2025	98.7	85.2	9.5	0.7

6.5. Conclusion

The three cases differ clearly in how well they suit TDTR. Logistics is the best fit, cold storage is moderate, and dairy is the weakest. What matters is not the sector itself, but how much of the load can be shifted and how well that demand lines up with the local call-hours, all under the current cost assumptions for non-served energy, contracts, and flexibility assets.

Logistics is the strongest case. Smart charging plus a 74% TDTR contract cuts annual costs by 23.4% on the default Oudehaske profile. Most of its demand is EV charging, which can be moved to other hours. Smart charging already lowers the yearly peak, and TDTR moves the highest-load moments out of the call-hours. Across all four call profiles the customer keeps saving (roughly 18–23%).

Cold storage is a moderate fit. Flex plus a 45% TDTR contract cuts costs by 5.4%. Cooling can be shifted within the cold room's thermal inertia, so no separate buffer is needed.

Dairy is the weakest case. A 19% TDTR contract with 10% flexibility cuts costs by 2.28%. The process runs almost continuously and only a small share of the heat-pump load can be shifted, so the grid benefit stays small.

The sensitivity analysis adds four points, all conditional on the current cost assumptions:

1. **The customer's savings are stable across call profiles.** They depend mainly on the load type and connection size, not on the local call timing.

2. **The grid benefit depends on whether market signals and local congestion align, and the cost-optimal TDTR share is the readout of that alignment.** What drives the load schedule is the response to day-ahead electricity-market signals, not flexibility as such: the optimiser shifts load toward low-price hours. Where those low-price hours coincide with the call hours (Types A–B), a large TDTR share is taken and TDTR adds real call-hour relief, where they do not (Types C–D), the TDTR share retreats to a small level and TDTR adds almost nothing (on Eerbeek, Type C, the logistics call-hour reduction moves only from 37.66% to 37.80%, a negligible 0.14 percentage-point change). The weighted grid-pressure reduction shows that price-driven scheduling does

flatten the overall profile and so eases system-wide congestion, even on profiles where it fails to relieve the specific local call hours, local and system-wide relief are not the same thing.

3. **Market-driven scheduling can harm the local grid.** Under the Zaltbommel profile (Type D), the cheapest hours fall exactly in the call hours, so cost-minimising charging concentrates load there and the call-hour effect turns negative (down to -41.9% for logistics under flex alone). A realistic non-served-energy penalty means TDTR cannot force this back to zero voluntarily, because compliance would require shedding charging sessions the customer will not pay for, the firm fraction is a lever the operator can pull, but only by transferring the misalignment cost onto the customer. A good deal for the customer therefore does not guarantee relief for the grid, and can run against it.

4. **Extra flexibility gives diminishing returns under current costs.** Doubling dairy flexibility barely improves the result, and adding a battery to the logistics case improves the grid indicators but is too expensive: customer savings shrink to near zero or turn negative as the battery grows. Under the assumed asset costs, smart charging plus TDTR, without a battery, remains the cost-optimal customer strategy.

5. **The alignment reading does not hinge on the price year.** Across 2023–2025, the profile ordering is preserved: Oudehaske remains the most clearly aligned profile and Zaltbommel the weakest, with Hemweg and Eerbeek in between. The rising number of negative-price hours sharpens this mechanism, because it increases the pull toward cheap hours.

In short, TDTR works best for large, market-responsive loads like EV charging, less well for shiftable but smaller or continuous loads like cold storage, and least well for process-bound cases like dairy. But the grid benefit always depends on whether the price-driven schedule lines up with the local call-hours, a question of the customer–substation pairing rather than the sector.

7 Discussion

The results show that TDTR can lower customer costs, but that a positive business case does not by itself produce local grid relief: whether load actually falls during the call-hours depends on the fit between the customer's flexibility, the local call-hour profile, and the electricity-price signal. This chapter interprets that finding, asking first what the contract adds on top of price-driven flexibility, then positioning TDTR within the operator's wider instrument portfolio and against its potential to defer reinforcement, considering where voluntary participation tends to concentrate, setting out the limitations that bound these results, and finally translating the analysis into recommendations for the network operator and directions for further work.

7.1. What the contract adds on top of price-driven flexibility

A central finding is that a positive customer business case does not guarantee a lower load during the call hours. TDTR is a contract: it lowers the customer's firm capacity right during predefined local congestion hours, but it does not by itself determine how the customer dispatches its flexible demand. That dispatch is still driven by the customer's cost minimisation, in which electricity prices play an important role.

This distinction explains why the same product has different grid effects across the substation profiles. When call-hours coincide with high electricity prices, the customer's private incentive and the DSO's congestion objective point in the same direction. The customer already has an incentive to shift flexible demand away from these hours, even without TDTR. In such cases, much of the observed call-hour relief comes from price-driven flexibility rather than from the contract itself. TDTR then mainly formalises and prices a response that the customer would partly have provided anyway. Its added value is therefore less a marginal change in dispatch and more the creation of an explicit conditional capacity right.

When call-hours coincide with low electricity prices, the mechanism changes. The customer's cost-minimising dispatch wants to use more electricity in the same hours in which the DSO wants load to fall. TDTR then has to work against the price signal. The optimiser will only accept a lower firm fraction as long as the discount outweighs the cost of changing operation. Beyond that point, additional call-hour reduction may be technically possible but is no longer privately cost-optimal. This explains why the cost-optimal firm fraction forms a practical bound on voluntary participation.

The implication is that TDTR's incremental grid value cannot be read directly from the difference between a reference profile and a TDTR profile. If the reference profile is price-inelastic, part of the relief may simply be the effect of price-driven flexibility that would occur even without the contract. A cleaner interpretation is therefore to separate the two effects. The restriction imposed by TDTR is most valuable where the customer's natural price response does not already relieve the local grid.

7.2. TDTR among the operator's other instruments

TDTR also has to be positioned alongside the operator's other congestion management instruments. Congestion-management (CM) contracts are typically temporary, on-call measures linked to an existing firm right. TDTR differs because it defines an ex-ante conditional capacity right: the customer accepts a bounded restriction during predefined call-hours in exchange for a lower tariff.

The two instruments are therefore complementary, but their value is not additive. CM, TDTR, wholesale-market optimisation and balancing-market participation can all draw on the same flexible capacity. Flexibility used for price arbitrage or imbalance response cannot automatically be counted again as congestion relief. This matters for how TDTR is credited: the DSO should avoid treating all observed call-hour reduction as a contract effect when part of that response may already be driven or rewarded elsewhere.

For the operator, the relevant question is therefore not only whether TDTR works in isolation, but how it fits within the broader portfolio of flexibility instruments. TDTR is most useful where ex-ante certainty about reduced capacity rights during local congestion is more valuable than relying solely on short-term activation through CM.

7.3. TDTR and avoided reinforcement

The results show that TDTR should not be interpreted as an automatic substitute for grid reinforcement. Its contribution to avoided reinforcement depends on whether the selected call-hours correspond to the moments that drive the local network constraint, and whether the customer's load can actually be reduced in those hours.

In this thesis, call-hours are defined as the 15% busiest hours of the substation. They therefore represent moments of local grid pressure. However, a lower firm right during these hours only creates physical grid relief if the customer's realised demand is also lower when the constraint is binding. This is most likely when the customer's flexibility, operational limits and price response fit the local congestion profile.

TDTR can affect reinforcement in two related but distinct ways. First, it can reduce realised load during call-hours. This is the strongest case for avoided reinforcement: if the call-hour peak falls at the constrained component, less capacity is needed at the moments that matter most for the local grid. Second, TDTR reduces the customer's firm capacity claim during those same hours. Under a firm-only contract the customer retains a firm right up to the original contracted capacity, whereas under TDTR this right is reduced to the firm level during call-hours. This gives the DSO more certainty about the maximum capacity that must be guaranteed locally.

The second effect is important, but it should not be overstated. A lower firm right does not automatically prove that reinforcement can be deferred. That depends on the actual binding constraint, the simultaneity of other users and the realised behaviour of the customer. If the customer profile does not fit the call-hour profile, TDTR may limit the formal capacity claim without producing a large reduction in total call-hour energy. In that case, the reinforcement value is more uncertain and has to be assessed case by case.

Therefore, TDTR can support reinforcement deferral when the product is targeted at customers whose flexibility can reduce demand during the hours in which the local grid is actually constrained. Where that fit is absent, TDTR may still be useful as a conditional access product, but its contribution to physical grid relief cannot be assumed in advance.

7.4. Recognising where TDTR concentrates

Because voluntary participation is strongest where the customer's private incentives already align with the local congestion profile, TDTR may concentrate on the substations where it is easiest to sell rather than where it adds most incremental grid value. This creates a steering problem for the DSO. On well-aligned substations, customers self-select into the product because the discount fits behaviour they would partly adopt anyway. On misaligned substations, where the contract would need to counteract the price signal, customers require a higher compensation or a lower firm-fraction requirement to participate.

This does not mean that the TDTR contract should simply be differentiated by location. Any differentiation in discounts, call percentages or eligibility has to respect the ACM equal-treatment principles: similar cases must be treated similarly, and only materially different cases can justify different treatment. A first requirement is therefore to express those differences in objective terms, for example through the local congestion profile, the timing of call-hours and the expected fit with customer flexibility. How durable such criteria can be is itself debatable, because the alignment between local congestion and electricity prices shifts over time as renewable generation, market behaviour and connected assets evolve: a classification that is objective today may no longer hold a few years later. Low or negative prices are increasingly linked to weather-driven renewable output, and customers that optimise against those prices may shift demand into the same hours that already stress parts of the grid. Type D illustrates this: local call-hours coincide with low-price periods, so customer optimisation and network relief move in opposite directions. As market conditions and asset portfolios change, similar misalignment may deepen on some substations and ease on others.

The implication is that TDTR requires active targeting rather than passive roll-out. The DSO should screen not only where congestion occurs, but also which customers are likely to reduce demand during those specific hours. Without such screening, TDTR risks spreading fastest where it adds the least incremental relief, while remaining less attractive where the network need is strongest.

7.5. Methodological limitations

Several limitations qualify the results, and most point in the same direction: the model is a clean, optimistic benchmark by design, rather than a prediction of practice.

Representativeness. The cases are archetypes, not full sector populations. Firms within a sector differ in process design, electrification level, connection size, and risk tolerance, so the results give structured insight into the underlying mechanisms rather than sector-wide classifications.

Sector selection. Some high-potential sectors are excluded because they combine multiple interacting flexibility sources behind one connection. Greenhouse horticulture is the clearest example, where lighting control, CHP, heat buffering, and electric heat production jointly determine the net profile, modelling it would require a multi-asset framework beyond the three transparent archetypes compared here.

Perfect foresight. Day-ahead prices, demand, call-hours, and flexibility availability are treated as known in advance, which makes the results an upper bound. The assumption is not evenly spread: the energy cost saving rests on predictable price-following against day-ahead prices known the evening before, so it stays close to reality, while the call-hour relief is more exposed, since a real customer facing limited notice or a forecast error cannot schedule as perfectly around the call-hours. This error is asymmetric, it overstates the relief on the aligned profiles, where the model already delivers the maximum, but barely moves the misaligned profiles, where relief is near zero regardless.

Single-customer optimisation. Each case optimises one customer against one call signal in isolation. This cannot capture the composition effect, where many customers shift toward the same cheap, uncall hours and TDTR simply moves the system peak to a new common hour rather than removing it. Realised, system-level relief could therefore be smaller than the single-customer results suggest, and this multi-customer rebound is arguably the most important open question for the DSO.

Rational-customer assumption. The model assumes a perfectly cost-minimising customer. Real customers deviate in two directions, both weakening the grid outcome: taking the discount without responding to the call signal beyond what cost-minimisation already implies, or declining a TDTR contract altogether out of risk aversion even where the expected saving is positive. Modelled participation and relief are thus an optimistic bound.

NSE penalty. Non-served energy enters as a single, sector-level VOLL-anchored penalty, a device that lets the optimiser trade restriction against curtailment rather than a statement that unmet demand is acceptable. Because this value drives the trade-off, the results are sensitive to it, and because it does not represent each firm's own margin per MWh, the model overstates the retreat to a small share for high-margin or interruption-tolerant customers and overstates uptake where the true NSE cost exceeds the anchor.

Implementation costs. The model captures EMS and thermal-buffer costs but not the full implementation burden, organisational change, monitoring, contract management, and coordination, which weigh especially on smaller customers and make the model optimistic about TDTR's practical attractiveness.

Grid-side indicators. The weighted grid-pressure reduction and the call-hour load reduction are not fully independent: both measure a time-weighted load reduction and differ only in the time axis of the weighting, so they do not give orthogonal evidence of grid value. Their divergence on Type D is informative, but its cause is the customer's price response, not a system-wide congestion signal. A more robust indicator set would pair the call-hour measure with a different dimension, such as the reliability of the delivered reduction.

Unidirectional charging. The logistics case models unidirectional charging only. Bidirectional charging, V2G, and V2B are excluded to keep the case focused on compliance through demand shifting, including them would require further assumptions on availability, degradation, discharge limits, and market participation, and is better treated as a separate extension.

Offtake-driven congestion only. The model addresses substations whose congestion arises from demand drawing on the grid. Feed-in-driven congestion, where local generation injects into the grid, reverses the alignment logic and is outside the scope of the instrument modelled here.

Call-hour clustering. The model constrains TDTR through an annual share of call hours (15%), not through their sequencing, so it does not capture the consecutive-call episodes that real contracts bound with maximum-duration or maximum-succession terms. A prolonged run of call hours, two weeks back to back, cannot in practice be met from within the process. The modelled participation is therefore optimistic, and the extra non-served-energy risk this imposes could scare customers.

7.6. Recommendations for the network operator

Taken together, the results support offering TDTR, but not as a passive or generic congestion solution. The contract reduces the customer's firm capacity right during predefined call-hours, but it does not by itself steer the dispatch of flexible demand. Whether TDTR produces local grid relief therefore depends on the fit between the customer's load profile, its flexibility characteristics, its price response, and the local call-hour profile. The operator should therefore treat TDTR as a targeted access and congestion-management product: useful where it aligns customer savings with local grid value, but uncertain where the customer's cost-minimising behaviour works against the local congestion signal.

The recommendation is therefore not a blanket yes or no. TDTR is most defensible where the selected call-hours represent the actual hours that drive the local network constraint and where the customer has flexibility that can reduce demand in those hours. In such cases, the product can lower customer costs while also reducing the capacity used or guaranteed during local congestion. Where this fit is absent, TDTR may still grant access within existing capacity or reduce the formal firm capacity claim during call-hours, but its realised congestion relief cannot be assumed in advance. For existing firm customers, in particular, the operator should avoid offering TDTR purely because the customer is willing to accept a discount, the relevant question is whether the contract is likely to reduce load at the constrained component when it matters.

Because each customer sits at a fixed connection, the operator's main lever is not where to place participants, but which customers to steer toward TDTR, under what terms, and with what monitoring. The recommendations below set out where TDTR is the right instrument, what it reliably delivers, how to make customer suitability visible and steer it toward the substations that need relief, and how to avoid crediting TDTR for relief that price-driven flexibility or other instruments already provide.

Value TDTR for the access it enables, not only for call-hour relief. On a congested substation, TDTR lets demand be served within existing capacity that an unconditional firm connection could not be granted, and this holds regardless of whether the call-hour peak falls. Where price and call-hours align, the price-driven dispatch additionally relieves the call-hours, where they do not, that relief is small or absent. The operator should treat access as the reliable benefit of the product and call-hour relief as a profile-dependent bonus, rather than judging every connection on realised reduction alone.

Make the benefits of TDTR visible to customers. A customer rarely sees on its own how its load profile lines up with the local call pattern, and so underestimates what participation could yield. The operator could close this gap by building a decision support dashboard by extending the optimisation framework used here: a customer enters its size, profile, and flexible assets, selects its substation, and is shown an approximate firm fraction and the associated saving. The same tool serves the operator in reverse, flagging which connections are well or poorly aligned. This keeps participation voluntary while surfacing the value and bringing well-aligned customers to the table.

Actively target the substations that need relief, rather than waiting for uptake. Because voluntary participation is strongest where the customer's incentives already align with the local congestion profile, TDTR tends to spread fastest where it adds the least, and a suitability dashboard, by bringing well-aligned customers forward, can reinforce this. The operator should therefore use the same assessment in reverse: identify the substations where relief is needed but uptake will not come on its own, and concentrate steering and stronger terms there. Because alignment itself shifts over time as renewable output and connected assets evolve, this screening should be repeated periodically rather than fixed once, since a connection that is misaligned today may not be in a few years.

Where price alignment is absent, add a positive reason to reduce load. On these substations the operator should look beyond the access-only contract for an instrument that makes call-hour reduction privately worthwhile: a payment tied to the delivered reduction, compensation for the resulting non-served energy, or a targeted congestion-management measure that bridges the gap between TDTR roll-out and short-term CM activation.

Let customers optionally declare demand-side flexibility. Storage and on-site generation must already be reported, but the demand-side mechanisms that determine TDTR suitability, shiftable charging and thermal buffering, do not. A voluntary declaration route lets a customer that documents its flexible share credibly receive a firm fraction matched to that share, instead of a generic profile. This rewards transparency without penalising its absence and keeps the initiative with the customer.

Monitor the realised reduction in the call-hours, not the contracted share. A customer can keep to its contracted firm-to-TDTR ratio on paper while the load it actually draws during the call-hours barely changes. Two things make the contracted share an unreliable proxy for relief. First, a lower capacity right only relieves the grid where it actually binds: if the customer's demand would otherwise have reached the firm level, the restriction forces load down, but where realised demand already sits below it, the customer keeps headroom up to its firm right and call-hour load need not fall at all. Second, part of any observed reduction may be flexibility that price signals, congestion-management contracts, or balancing participation already drive or reward, and that cannot be credited to TDTR again. The operator should therefore track the delivered reduction against a baseline that reflects what the customer would have scheduled on price alone, net of flexibility rewarded through other instruments. As a customer applies its own flexibility, this baseline drifts from the registered reference profile, which should be updated accordingly.

7.7. Directions for further work

The reflections above point to two kinds of further research: directions that deepen the model itself, and directions that move beyond the single-customer frame to the system and the operator.

Deepening the model

A first direction is to **relax the perfect-foresight assumption** and test how robust the results are under forecast errors, limited notice times, and imperfect operational response. Re-solving under a rolling horizon (for example 24–36 hours) with day-ahead call notice and stochastic demand and availability would turn the optimistic benchmark into a realistic operating range, and would show how far the aligned-profile relief retreats once perfect scheduling is no longer possible, while the central decoupling result is expected to hold.

A second is to **measure the realised flexible share empirically**, rather than assuming a given technical potential, since what a firm is willing and able to offer depends on organisational and operational factors and is usually smaller than its theoretical potential.

A third is to **estimate firm-specific NSE costs empirically**, in consultation with sector associations and individual large consumers, replacing the sector-level anchor whose sensitivity is noted among the limitations and sharpening the restriction-versus-curtailment trade-off.

A fourth is to **extend the modelling scope**: relax the unidirectional-charging assumption in the logistics case to include V2G or V2B operation, which the present model leaves out.

A fifth avenue is to **model the symmetric instrument for feed-in-driven congestion**, and to test how each archetype's suitability changes when the binding congestion direction reverses. Extending the model with behind-the-meter PV, battery storage, and local renewable generation such as wind parks would make it possible to assess substations exposed to congestion in both directions.

From the connection to the system

A first system-level direction is to **weigh the system cost of offering TDTR against the operator's revenue and deferral benefit**. This thesis shows that the operator earns connection revenue and may defer reinforcement even where delivered relief is small, while the administration and monitoring the product requires are real costs. Whether TDTR is net positive for the operator, and for the system, is a portfolio-level cost-benefit question the single-customer optimisation here cannot answer.

A second is to **quantify the multi-customer composition effect**, where many customers shift into the same cheap, uncall hours and TDTR moves the system peak rather than removing it. A multi-customer or market model would estimate realised system-wide relief against the sum of the individual contracts.

A third is to **trace the distributional effects of the discount**. The saving a non-responding customer collects on a misaligned substation is recovered elsewhere, through the tariffs other network users pay or through the operator, so the question is whether TDTR creates net social value or mainly redistributes it toward self-selecting customers on the substations where it delivers least.

A fourth is to **test the interaction with the congestion-management and balancing markets at scale**, since capacity locked into TDTR is no longer available for arbitrage or imbalance, and wide deployment could thin those markets and shift their price formation.

A fith is to **optimise the contract terms themselves**, the discount and the call-percentage, rather than taking them as given. These are the operator's own levers, yet a flat discount overpays for relief on aligned substations and underbuys it on misaligned ones, setting the terms against the relief actually delivered would turn a flat price for access into a targeted one.

Taken together, these directions suggest a layered research agenda in which the case-specific optimisation developed here is one part of a broader approach: rapid archetype-based screening at the top, detailed optimisation for the promising cases, and an ex-post, system-level check of the realised grid impact to close the loop back to the DSO's planning process. Two issues are especially important. The first is information: how well the framework works depends in the end on flexibility data that is not yet collected in a systematic way. The second is level of analysis: the value of TDTR to a single customer is well captured by the model developed here, but its value to the grid and to the operator only becomes visible once many connections, the wider sector portfolio, and the cost of the product itself are brought into view.

8
Conclusion

This thesis investigated under which conditions conditional grid access through time-duration-based transport rights (TDTR) delivers value, and to whom. The central research question was:

> *Under which conditions does conditional grid access through TDTR deliver value, and to whom: when does a cost-optimal firm-to-TDTR ratio for a large electricity consumer also relieve the local grid, and when do these two outcomes diverge?*

To answer this, a customer-level techno-economic optimisation model was developed and applied to three flexibility archetypes across four substation call profiles. This conclusion synthesises the findings through the three sub-questions before returning to the main research question.

The literature review (SQ1, Chapter 3) showed that existing work divides into two streams that rarely meet. A network-side stream models restriction regimes, capacity limitation, and queue allocation in detail, but represents users generically as homogeneous capacity reductions or simple price-takers. A user-side stream models industrial flexibility in detail through process constraints, recovery needs, and curtailment cost, but treats the network restriction itself as exogenous and simple. The small set of bridge studies, with Verhagen et al. (2025) as the closest neighbour, shows that the two can be combined, but each covers only part of the problem: one asset class under a binary contract, or heterogeneous users without a stacked TDTR-type structure. Two modelling choices followed and structured the rest of the thesis: a stacked firm-plus-TDTR contract representation rather than a binary firm/non-firm choice, and curtailment cost parameterised as a sector-specific, value-of-lost-load-anchored input rather than a single universal penalty. Together these define the gap addressed here: the absence of any prior study combining a stacked contract, sector-specific flexibility, and substation-derived call profiles in one comparative framework.

The case results (SQ2) show that every case can technically comply with TDTR call-hour restrictions, but the cost-optimal firm-to-TDTR ratio, and the value it delivers, differs by case. Both the customer saving and the grid-impact indicators are measured against the firm-only reference, but reported separately for the flex and the flex+TDTR scenario. Comparing the two columns is what isolates the value specific to TDTR from the value of operational flexibility alone: flexibility's own contribution is already visible in the flex scenario, and the step to flex+TDTR shows what the call-hour restriction adds on top of it. On the default Oudehaske profile, logistics reaches its optimum at a 74% TDTR contract (a 26% firm share) and reduces annual costs by 23.4% relative to the firm-only reference, because most of its demand is EV charging that can be shifted within operational charging windows. Cold storage is optimal at a 45% TDTR contract and saves 5.4%, with most of the reduction already coming from flexible cooling operation and TDTR adding a lower contracted-capacity cost. Dairy is optimal at a 19% TDTR contract and saves 2.3%, because continuous process demand and a small shiftable heat-pump share leave little firm capacity to give up. Attractiveness therefore scales with the share of demand that can be moved away from call-hours at low cost: a larger TDTR share only pays off when the customer can release firm capacity without incurring substantial non-served energy, which is why the cost-minimising firm fraction settles at 26% for logistics, 55% for cold storage, and 81% for dairy.

The sensitivity analysis (SQ3) reveals a clear split: the customer-side economics are robust, but the grid value is not, and the two can diverge sharply. Across the four call profiles the cost-optimal ratios and annual savings vary only modestly: dairy stays near 2%, cold storage near 5%, and logistics within roughly an 18–24% band. The realised grid value behaves differently. Because non-served energy is expensive, the optimiser accepts a large TDTR share only where compliance can be met by shifting load rather than by shedding it. The cost-optimal firm fraction therefore indicates how much load the customer can move at low cost, not whether moving it relieves the local grid, that depends on a separate property, the alignment between the customer's price-driven schedule and the local call-hours. Where call-hours coincide with high electricity prices (Types A–B), the customer's price response and the DSO's congestion signal point in the same direction: the customer already has an incentive to shift demand away from the called hours, and TDTR can reinforce this. Where call-hours coincide with low electricity prices (Types C–D), TDTR adds little realised reduction, but for two different reasons. On Eerbeek (Type C), price-driven scheduling already reduces call-hour load substantially in the flex baseline, so little is left for TDTR to correct: the reduction moves only from 37.7% to 37.8%, and the customer simply collects the contracted-capacity discount while the realised load profile is essentially unchanged. Under the adverse Zaltbommel profile (Type D), the price signal instead pulls load into the call-hours (-41.9% before TDTR); because forcing deeper compliance would require uneconomic shedding, the cost-optimal TDTR contract leaves this largely uncorrected (-41.5%). The exception that proves the mechanism is cold storage on Type D, where thermal-buffer shifting corrects most of the adverse effect (-4.7% to -1.0%): TDTR relieves the call hours only where the asset can respond by shifting rather than by shedding. The asset-parameter analysis confirms diminishing returns under the current cost assumptions: doubling dairy heat-pump flexibility improves the result only modestly while requiring a much larger thermal buffer, and adding a battery to the logistics case improves the grid indicators but erodes the customer saving as battery size grows, falling from roughly 18% for the smallest battery, below the roughly 23% of the battery-free logistics case, to near zero or negative for the largest, without reversing the Type D effect, because the battery too is dispatched on the customer's cost drivers (mainly the day-ahead price and the monthly peak charge), not on the local call signal.

Returning to the main research question, TDTR can be both technically feasible and economically attractive, but only delivers local grid value under additional conditions. Three conditions govern the outcome. The first is a large, controllable flexible load share: the decisive factor is how much demand can be moved without disrupting the underlying service, which is why logistics scores highest and dairy lowest, and why connection size matters less than controllability. The second is a cost-optimal firm-to-TDTR ratio at which the contracted-capacity discount outweighs the cost of compliance while substantial non-served energy is avoided, for TDTR to add value over firm capacity alone, this optimum must fall at a meaningful TDTR share rather than a marginal one. The third condition separates customer value from grid value and is a property of the customer–substation pair rather than of the sector: alignment between the customer's price-driven flexible schedule and the local call-hours. Under the assumed shedding cost, voluntary participation does not correct such misalignment on its own, and it can self-select the largest TDTR shares onto substations where price-driven flexibility already helps the grid. This sharpens the need for active deployment: the DSO should screen whether the customer profile fits the local call-hours, monitor the realised load reduction in the call hours rather than the contracted share, and add steering or compensation where market incentives and grid needs diverge.

The result is an archetype-based assessment showing that customer value and grid value do not always move together. Suitability is not a sector property but a property of the flexibility behind the connection and its overlap with the local call pattern. TDTR is most valuable where contractual cost savings, operational flexibility, and local congestion relief point in the same direction. Where they do not, TDTR may still reduce the customer's formal firm capacity claim during call-hours or help structure conditional access, but its congestion-relieving value must be judged on realised call-hour load reduction rather than assumed from the contracted TDTR share.

9
AI Statement

AI tools were used to support the preparation of this thesis. Specifically, they were employed to assist with language refinement, formatting of BibTEX references, and the creation and structuring of figures and tables in LATEX (Overleaf). The conceptualisation of the research, formulation of research questions and methodological design were carried out entirely by the author.

A

Economic Framework of TDTR

Chapter 2 positioned TDTR as a contract-based form of conditional access alongside firm transport capacity. This chapter translates that institutional concept into the economic inputs required for the optimisation model: how TDTR affects electricity costs, and when restrictions apply during the year.

Section A.1 describes the three main electricity-bill components for large electricity consumers: energy procurement costs, network tariffs, and taxes. The focus is on how TDTR changes the network-tariff structure by reducing the *KW-contract* charge for the conditional capacity component. Section A.2 introduces the proposed time-dependent network tariff reform, which adds a timing-related network-cost incentive relevant for flexible operation.

A.1. Structure of the electricity bill

For the purposes of this study, the total electricity bill of a LEC is divided into three components: electricity procurement costs paid to the supplier, network tariffs paid to the network operator, and government taxes (Figure A.1). Together they form the financial basis for assessing whether flexibility and alternative transport contracts can create a positive business case. Each component is developed in turn below.

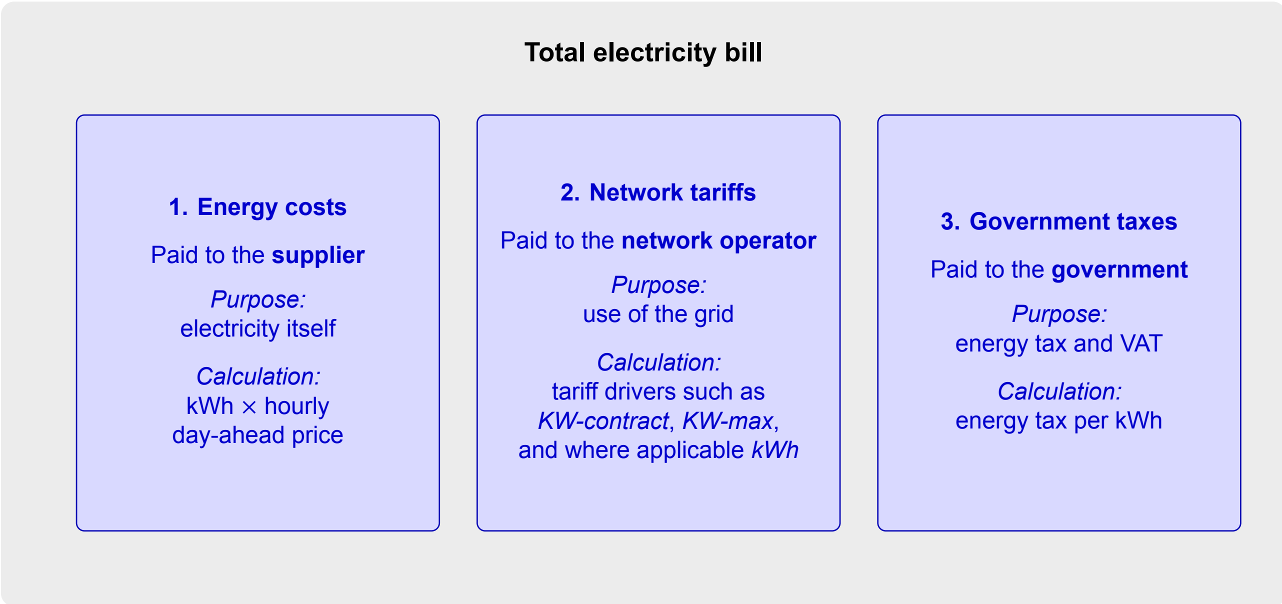

Figure A.1: Main components of the total electricity bill.

A.1.1. Energy costs

Electricity procurement costs are represented using hourly day-ahead prices, which reflect the wholesale market value of electricity for next-day delivery. As a result, the timing of electricity consumption directly affects total procurement costs. A flexible asset such as a heat pump, e-boiler, or EV charger that shifts consumption from expensive to cheaper hours can reduce procurement costs even when total electricity consumption remains unchanged.

In the model, procurement cost is calculated as the product of hourly demand and the corresponding day-ahead price:

$$C^{\text{energy}} = \sum_{t \in \mathcal{T}} P_t \cdot \pi_t \cdot \Delta t$$

where P_t is the electricity demand in time step t, π_t is the day-ahead electricity price in EUR/kWh (ENTSO-E, 2026), and Δt is the duration of the time step in hours. Day-ahead prices therefore serve two purposes in the business-case assessment: they determine the direct cost of consumption, and they create an additional source of flexibility value because load shifting may reduce exposure to high-price hours.

A.1.2. Network tariffs

Three tariff drivers are used in the network-tariff structure. The *KW-contract* component is the contracted transport capacity reserved for the customer and charged in EUR/kW/month. The *KW-max* component is based on the realised monthly peak demand and is also charged in EUR/kW/month. The *kWh* component is a volumetric transport charge based on total electricity consumption and charged in EUR/kWh.

Network tariffs are the cost component most directly affected by flexibility and conditional transport rights, because they depend not only on total consumption but also on contracted capacity and realised peak demand. They therefore drive the financial value of peak shaving, load shifting, and TDTR participation.

Large-consumer connections from 3×80 A onward (Liander, 2026b) are grouped by grid level and tariff structure in Table A.1. Smaller and medium-sized large consumers face a tariff structure that includes *KW-contract*, *kWh*, and *KW-max*; larger high-capacity consumers face a more capacity-driven structure with *KW-contract* and *KW-max* only.

Table A.1: Grouping of large consumers and tariff structure used in this study.

Group	Grid level	AC-category	Indicative connection size	Tariff split
Smaller large consumers	LV / lower MV	AC4a–AC4b	>3×80 A to 160 kVA	25% KW-contract, 50% kWh, 25% KW-max
Medium large consumers	MV	AC5a–AC5b–AC5	>160 kVA to 2 MVA	25% KW-contract, 50% kWh, 25% KW-max
Larger high-capacity consumers	HV	AC6a–AC6b	>2 MVA	50% KW-contract, 50% KW-max

The contract form determines how the tariff drivers are charged. Under firm transport, all tariff components apply. Under TDTR, the *KW-contract* component is removed for the conditional capacity, while *KW-max* remains unchanged. This preserves an incentive to reduce realised peak demand.

Table A.2: Financial effect of firm transport

Contract form	AC5 and below	AC6 and above
Firm transport	25% *KW-contract*, 50% *kWh*, 25% *KW-max*	50% *KW-contract*, 50% *KW-max*
TDTR	0% *KW-contract*, 50% *kWh*, 25% *KW-max*	0% *KW-contract*, 50% *KW-max*

The tariff values used in this study are taken from Liander's 2026 tariff sheet for large-consumption electricity connections (Liander, 2026a) and are listed in Table A.3. The applicable values depend on the connection category of the case under consideration.

Table A.3: Relevant Liander 2026 transport tariff components used in this study.

Connection category	kWh tariff	KW-contract	KW-max
MS > 136 to $2{,}000$ kW	0.0226 EUR/kWh	2.29 EUR/kW/month	3.57 EUR/kW/month
TS/MS of HS/MS $> 2{,}000$ kW	–	3.9900 EUR/kW/month	6.24 EUR/kW/month
TS $> 2{,}000$ kW	–	3.9900 EUR/kW/month	5.39 EUR/kW/month

A.1.3. Energy tax

Energy tax is a government levy on electricity consumption and is treated as a separate cost component, distinct from network tariffs. The applicable rate depends on the annual consumption bracket of the case, for business consumers in the Netherlands, the brackets relevant for this study are 50,001 kWh to 10 million kWh and more than 10 million kWh (Table A.4). Energy tax is calculated as

$$C^{\text{tax}} = E \cdot \tau^{\text{tax}},$$

where E is annual electricity consumption in kWh and τ^{tax} is the applicable rate. Unlike the network-tariff drivers, energy tax depends only on total electricity volume, not on its temporal distribution: load shifting without a change in total consumption does not reduce energy tax payments, whereas a reduction in total electricity use would.

Table A.4: Dutch electricity tax rates for business users in 2026 (Belastingdienst, 2026).

Year	50,001 kWh to 10 million kWh	More than 10 million kWh (business)
2026	€ 0.03735	€ 0.00310

A.2. Proposed time-dependent network tariffs for large consumers

In addition to the current tariff drivers, a proposal has been submitted to introduce time-dependent transport tariffs for large consumers connected to distribution networks (Netbeheer Nederland, 2026). The core idea is that network charges should depend not only on contracted capacity, peak demand, and volumetric use, but also on *when* the grid is used, creating a time-based incentive to shift consumption away from periods of high expected system loading.

Under the proposal, the *KW-max* driver is replaced by a weighted variant in which each quarter-hourly demand value is multiplied by a predefined weighting factor, and the highest weighted value in the month determines the chargeable peak. For MS, Trafo MS/LS, and LS categories the *kWh* driver is weighted in the same way. The *KW-contract* component remains unchanged. The weighting factors reflect expected relative network loading rather than wholesale prices and differ by month, hour, and day type. High-load periods receive a higher relative weight (up to 1.0), low-load periods a lower one. Figures A.2 and A.3 show the proposed schemes for working days and non-working days.

Maand \ Uur	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	1,00	1,00	0,82	0,82	0,58	0,58	0,58	0,82	1,00	1,00	1,00	1,00	1,00	1,00	0,82	0,82
2	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	1,00	1,00	0,82	0,82	0,58	0,58	0,58	0,82	1,00	1,00	1,00	1,00	1,00	1,00	0,82	0,82
3	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,25	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,82
4	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,25	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,82
5	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,10	0,10	0,25	0,25	0,25	0,25	0,58	0,58	0,82	0,82	0,82
6	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,10	0,10	0,25	0,25	0,25	0,25	0,58	0,58	0,82	0,82	0,82
7	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,10	0,10	0,25	0,25	0,25	0,25	0,58	0,58	0,82	0,82	0,82
8	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,10	0,10	0,25	0,25	0,25	0,25	0,58	0,58	0,82	0,82	0,82
9	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,25	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,82
10	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,25	0,25	0,25	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,82
11	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	1,00	1,00	0,82	0,82	0,58	0,58	0,58	0,82	1,00	1,00	1,00	1,00	1,00	1,00	0,82	0,82
12	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	1,00	1,00	0,82	0,82	0,58	0,58	0,58	0,82	1,00	1,00	1,00	1,00	1,00	1,00	0,82	0,82

Figure A.2: Proposed scheme of weighting factors for working days in the time-dependent network tariff proposal for large consumers.

Uur / Maand	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
1	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,58	0,58	0,58
2	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,58	0,58	0,58
3	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58	0,58	0,58
4	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58	0,58	0,58
5	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58
6	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58
7	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58
8	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58
9	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58	0,58	0,58
10	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,25	0,10	0,10	0,10	0,10	0,10	0,10	0,25	0,25	0,58	0,58	0,58	0,58	0,58	0,58
11	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,58	0,58	0,58
12	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,58	0,82	0,82	0,82	0,82	0,58	0,58	0,58

Figure A.3: Proposed scheme of weighting factors for non-working days in the time-dependent network tariff proposal for large consumers.

The weighting factors are not a direct discount on current tariffs. Since the allowed revenues of the network operator must still be recovered, the tariff per weighted unit is set in the annual tariff decision: the weighting factors therefore redistribute network charges across time periods and users rather than reducing them in aggregate.

For this thesis, the proposal is relevant as a forward-looking extension of the current tariff logic. It adds a timing-related network-tariff incentive on top of the existing flexibility value from lower procurement costs, lower peak demand, and lower contracted capacity. Flexible assets can reduce their weighted network-tariff exposure by shifting operation from high-weight to low-weight periods, and the model's sensitivity to this proposal is examined in the weighted-tariff specification used in Chapter 6.

B Literature Review Tables

Table B.1: Initial literature review: search terms and results

Search terms (Scopus, TITLE-ABS-KEY)	Hits	Selected
Search string A (capacity-based congestion and access)		
`("capacity limitation" OR "capacity-limiting" OR "capacity constraint*")` `AND ("distribution network*" OR "distribution grid*")` `AND (flexibilit* OR "non-firm connection*" OR "flexible connection*")`	17	[1], [3], [5], [10]
Search string B (non-firm and flexible grid connections)		
`("non-firm connection*" OR "flexible connection*" OR "conditional access")` `AND ("distribution network*" OR "distribution grid*" OR DSO*)`	18	[7], [8]
Search string C (industrial demand-side flexibility)		
`("industrial flexibilit*" OR "industrial load flexibilit*" OR "industrial demand response" OR` `("demand-side flexibilit*" AND industrial))` `AND (electricity OR electric* OR power)`	125	[2], [4], [6], [11]
Snowballing (backward/forward citation tracking and regulatory documents)	–	[9], [12], [13], [14]

Numbers in brackets refer to the corresponding studies listed in Table 2.

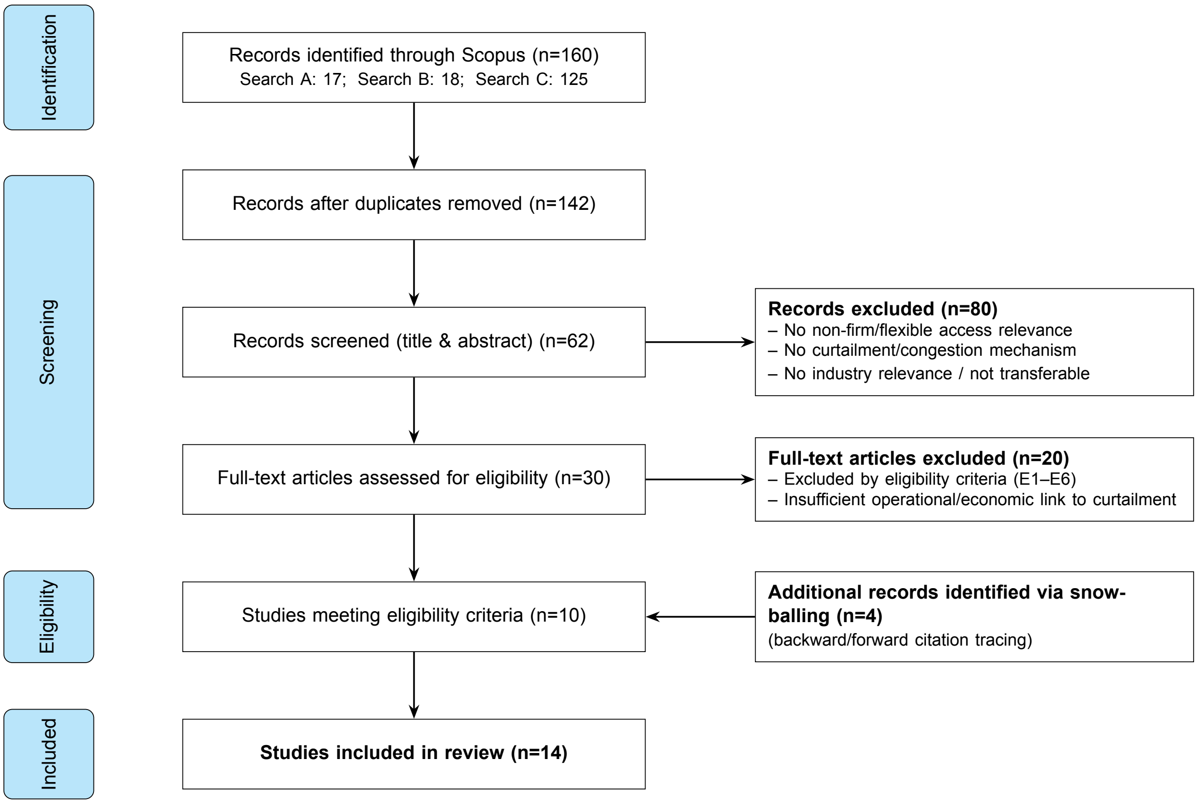

Figure B.1: PRISMA-style screening flowchart of the systematic literature search. Records were identified via Scopus (Search A–C) and supplemented through backward/forward citation snowballing, yielding 14 studies included in the review

Table B.2: Overview of core literature on non-firm grid access and industrial flexibility informing the analysis of Alternative Transport Rights (ATR).

No.	Authors	Year	Title	Cites	DOI
1	Heinrich et al.	2021	*A local flexibility market mechanism with capacity limitation services*	51	10.1016/j.enpol.2021.112335
2	Golmohamadi	2022	*Demand-Side Flexibility in Power Systems: A Survey of Residential, Industrial, Commercial, and Agricultural Sectors*	47	10.3390/su14137916
3	Chen et al.	2023	*Modeling and characterization of the impact of capacity limitation services on distribution networks*	6	10.1016/j.segan.2023.101081
4	Santecchia et al.	2022	*Industrial flexibility as demand side response for electrical grid stability*	23	10.3389/fenrg.2022.831462
5	Mirzaei Alavijeh et al.	2024	*Capacity-limitation based local flexibility market for congestion management in distribution networks: Design and challenges*	23	10.1016/j.ijepes.2023.109742
6	Serpe et al.	2023	*Demand-side flexibility of the German steel industry: A case study*	5	10.1109/EEM58374.2023.10161699
7	De Santi et al.	2025	*Managing connection queues in distribution networks with flexible connection agreements*	3	10.1016/j.apenergy.2025.126260
8	Bjarghov et al.	2024	*Enhancing grid hosting capacity with coordinated non-firm connections in industrial energy communities*	0	10.2139/ssrn.4713408
9	Leinauer et al.	2022	*Obstacles to demand response: Why industrial companies do not adapt their power consumption to volatile power generation*	50	10.1016/j.enpol.2022.112876
10	Hennig et al.	2024	*Risk vs. restriction—An investigation of capacity-limitation based congestion management mechanisms*	18	10.1016/j.enpol.2023.113976
11	Zhang et al.	2025	*A Review of Industrial Load Flexibility Enhancement for Demand-Response Interaction*	5	10.3390/su17114938
12	Monterde et al.	2025	*Non-firm grid connections: A review of access types, mechanisms, and regulatory frameworks*	1	10.1007/s40518-025-00268-7
13	Verhagen, Hu & Harmsen	2025	*Evaluating non-firm grid connections for battery energy storage systems: A co-optimization case study of the Netherlands*	0	10.1016/j.enpol.2025.114903
14	CEER	2023	*Paper on Alternative Connection Agreements*	–	–

Table B.3: Synthesis matrix of core literature mapped to ATR-relevant contract and modelling dimensions

Study	Non-firm access	Curtailment design	Sector / process explicit	Cost / flexibility quantified	Method
1	✓				Opt.
2			✓	✓	Review
3	✓	✓			Opt./Design
4			✓	✓	Opt.
5	✓	✓			Design
6		✓	✓	✓	Case/Opt.
7	✓	✓			Opt.
8	✓	✓	✓		Opt.
9			✓		Empirical
10	✓	✓			Risk/Opt.
11			✓	✓	Quant/Emp.
12	✓	✓			Review
13	✓	✓		✓	Opt.
14	✓	✓			Guidance

C
Sector Selection

C.1. Purpose and approach

This appendix documents the sector-selection procedure used to identify the sector archetypes that are carried forward into the case construction in Chapter 5. The purpose is not to model the sectors in detail here, but to justify why the selected cases are relevant, comparable, and suitable for a TDTR assessment.

The selection starts from thirteen candidate large-electricity-consumer sectors. These sectors were included because they are relevant within the large-consumer segment, can be linked to available KO standard profiles, and are discussed in relation to electrification, flexibility, or distribution-grid congestion. For each candidate sector, a sector-specific SBI profile is used where available. Where no suitable SBI profile is available, the closest KO profile is used as a proxy.

The assessment consists of two parts. First, the available sector profiles are screened to characterise the temporal structure of demand. This provides a quantitative basis for judging whether a sector is likely to be relevant under conditional access. Second, the candidate sectors are evaluated through a multi-criteria scoring matrix that combines the screening results with information on annual electricity use, operational similarity, flexibility potential, ATR product fit, and data availability.

The translation of the selected archetypes into optimisation-ready model inputs, including flexible assets, operational constraints, curtailment costs, and TDTR parameters, is developed in Chapter 5.

C.2. Candidate sectors

Table C.1 lists the thirteen candidate sectors considered in the sector-selection procedure. Each sector is linked to the KO standard profile that provides the closest generic load-profile representation. In addition, the table reports the SBI profile used for the sector-specific screening. For offices and supermarkets, no suitable SBI profile was available within the dataset, so the corresponding KO profile is used as a proxy. For data centres, the dedicated KO data-centre profile is used.

The set of candidate sectors is not intended as a complete classification of Dutch economic activity. Instead, it is a pragmatic selection of large-electricity-consumer sectors that are relevant for assessing TDTR participation.

Table C.1: Candidate sectors considered in the sector-selection procedure.

Sector	KO profile	SBI profile
Cut flowers under glass	KO_GLASTUINBOUW	01192
Annual crops	KO_AGRARIER	011
Dairy products	KO_INDUSTRIE	105
Bread, pastry and pasta	KO_INDUSTRIE	107
Paper and cardboard	KO_INDUSTRIE	17
Cold stores	KO_LOGISTIEK	52102
Distribution and storage	KO_LOGISTIEK	52109
Land transport	KO_LOGISTIEK	49
Hospitals and inpatient care	KO_KANTOOR_ONDERWIJS	8610
Offices	KO_KANTOOR_ONDERWIJS	KO profile proxy
Data centres	KO_DATACENTERS	KO profile
Supermarkets	KO_LOGISTIEK	KO profile proxy
Education	KO_KANTOOR_ONDERWIJS	8542

C.3. Load-profile screening

A load-profile screening is applied to the sector profiles listed in Table C.1. The purpose of this screening is not to select sectors mechanically, but to provide a quantitative indication of how relevant each sector may be for TDTR assessment. Sectors with peaky, concentrated, or variable demand profiles are more likely to benefit from conditional access arrangements than sectors with flat and continuous demand, because their demand can in principle be shifted, reduced, or managed during a limited number of constrained hours.

Four indicators are used: overall peakiness, rarity of peak events, temporal compactness of peak events, and year-round variability. These indicators are summarised in Tables C.2–C.3.

Table C.2: Indicator formulas and interpretation.

Indicator	Formula	Description	Interpretation of a high score
Peak-to-Average Ratio	$1 - \frac{\bar{r}}{r_{\max}}$	Relative distance between the average load and the maximum quarter-hour load.	A highly peaky profile with substantial peak-shaving potential.
Peak Infrequency	$1 - \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}(r_t \geq 0.8\, r_{\max})$	How rarely the profile operates above 80% of its maximum value.	Peaks occur only occasionally, suggesting they may be easier to shift or curtail.
Time Concentration	$1 - \min\left(1, \frac{\bar{L}_{\text{peak}}}{24}\right)$	How concentrated peak events are in time, based on the average duration of consecutive peak intervals.	Peak events occur in short, compact blocks rather than long sustained periods.
Variability	$\min\left(1, \frac{P_{90}(r) - P_{10}(r)}{2}\right)$	Spread of the normalised load profile over the year.	Greater variability, implying more potential value of flexible control.

Table C.3: Notation used for the screening indicators.

Symbol / term	Definition
r_t	Normalised quarter-hour load at time step t.
T	Total number of quarter-hour time steps in the full-year profile.
$r_{\max}$	Maximum load: $r_{\max} = \max_t r_t$.
$\bar{r}$	Mean load: $\bar{r} = \frac{1}{T}\sum_{t=1}^{T} r_t$.
$P_{10}(r)$, $P_{90}(r)$	10th and 90th percentile of the normalised load profile.
$\mathbf{1}(\cdot)$	Indicator function, equal to 1 if the condition holds and 0 otherwise.
$\overline{L}_{\text{peak}}$	Average duration in hours of consecutive peak episodes with $r_t \geq 0.8\, r_{\max}$.

C.3.1. Screening results

Table C.4 reports the screening results for the thirteen candidate sectors. The results show substantial differences in temporal demand structure. Cut flowers under glass, annual crops, cold stores, and distribution and storage all combine high Peak Infrequency with high Time Concentration, indicating that high-load moments occur in relatively concentrated periods. Distribution and storage also has the highest Peak-to-Average Ratio and the highest Variability score, suggesting a strongly non-flat profile. Data centres score low on most indicators, reflecting their flat and continuous demand pattern. Dairy products also score relatively low on peakiness, but remain relevant as an industrial electrification case because their flexibility is linked to process heat rather than to profile shape alone.

The screening results are used as input for the multi-criteria assessment, especially for the congestion-relevance criterion. They do not determine the final selection by themselves, because TDTR suitability also depends on operational flexibility, process constraints, data availability, and whether the sector adds a distinct flexibility mechanism to the comparative case set.

Table C.4: Sector screening results based on four load-profile indicators.

Sector	Peak-to-Avg.	Peak Infreq.	Time Conc.	Variability
Cut flowers under glass	0.589	0.982	0.961	0.153
Annual crops	0.516	0.977	0.943	0.168
Dairy products	0.242	0.648	0.877	0.108
Bread, pastry and pasta	0.379	0.830	0.861	0.231
Paper and cardboard	0.450	0.859	0.877	0.241
Cold stores	0.367	0.945	0.958	0.133
Distribution and storage	0.618	0.958	0.960	0.279
Land transport	0.454	0.898	0.872	0.275
Hospitals and inpatient care	0.417	0.907	0.910	0.185
Offices	0.379	0.800	0.806	0.224
Data centres	0.156	0.138	0.882	0.051
Supermarkets	0.425	0.873	0.849	0.235
Education	0.458	0.870	0.839	0.237

C.4. Multi-criteria assessment

The thirteen candidate sectors are evaluated against six criteria, defined in Table C.5. The criteria combine quantitative and qualitative evidence. Annual electricity use is informed by EAN-level analysis. Congestion relevance is informed by the load-profile screening. Similarity of business operations, estimated flexibility potential, and ATR product fit are assessed using sector literature, flexibility studies, and sector-lead interviews. Data availability is based on the availability and quality of sector profiles and additional modelling inputs.

The scores are therefore informed ordinal judgements rather than measured optimisation outcomes. They are used to select a set of representative archetypes. The detailed quantitative parameterisation of the selected cases follows in Chapter 5.

Table C.5: Evaluation criteria for sector archetype selection.

Criterion	Low (10)	Medium (55)	High (100)	Wt.
(1) Annual electricity use *(EAN analysis)*	< 1 GWh/yr, limited contribution to regional congestion.	1–3 GWh/yr, relevant for local transformer and feeder constraints.	> 3 GWh/yr, structurally relevant for regional grid congestion.	15%
(2) Congestion relevance *(profile screening)*	Flat load, weak coincidence with congestion hours.	Moderate peakiness, partial overlap with local congestion periods.	Strong, temporally concentrated peaks coinciding with congestion moments.	10%
(3) Similarity of business operations *(literature + interviews)*	Highly heterogeneous operations and load profiles.	Moderately comparable, dominant processes similar with site-level variation.	Largely homogeneous processes and load profiles.	15%
(4) Estimated flexibility potential *(literature + interviews)*	Little controllable load, flexibility requires major investments.	Moderate flexibility through one dominant mechanism.	High flexibility through multiple controllable assets with bounded constraints.	20%
(5) ATR product fit *(screening + operational assessment)*	Long notice times and limited activation tolerance.	Conditional compatibility, binding constraints on duration or frequency.	Short notice acceptable, repeated activations feasible.	20%
(6) Data availability for modelling *(data availability check)*	Fragmented data, no representative profiles.	Profiles available with gaps or limited representativeness.	High-quality, representative data suitable for quantitative modelling.	20%

Table C.6: Sector-by-sector scoring on criteria (1–6), with weighted final score.

Sector description	(1)	(2)	(3)	(4)	(5)	(6)	Final
Cut flowers and ornamental shrubs under glass	**100**	**100**	**55**	**100**	**100**	**55**	**84**
Cultivation of annual crops	**10**	**100**	**10**	**55**	**55**	**100**	**55**
Manufacture of dairy products	**100**	**10**	**100**	**55**	**55**	**100**	**71**
Manufacture of bread, pastry goods, and pasta products	**55**	**55**	**55**	**55**	**55**	**100**	**64**
Manufacture of paper, cardboard, and paper/-cardboard products	**100**	**100**	**55**	**55**	**55**	**100**	**75**
Storage in cold stores	**55**	**100**	**100**	**100**	**100**	**100**	**93**
Storage in distribution centres and other storage facilities with EV fleet	**55**	**100**	**55**	**100**	**100**	**100**	**87**
Land transport	**55**	**55**	**55**	**55**	**55**	**100**	**64**
Hospitals and inpatient mental health/addiction care	**100**	**55**	**55**	**55**	**10**	**100**	**62**
Offices	**55**	**55**	**55**	**55**	**55**	**10**	**46**
Data centres	**100**	**10**	**100**	**55**	**10**	**100**	**64**
Supermarkets	**10**	**55**	**55**	**100**	**100**	**10**	**57**
Education	**10**	**55**	**100**	**55**	**55**	**100**	**64**

C.5. Selected archetypes

Based on the multi-criteria assessment, three archetypes are carried forward into the comparative TDTR assessment. The final selection is not based on the ranking alone, but also on the need to construct a set of cases that is analytically distinct, practically modelable, and representative of different flexibility mechanisms.

1. **Cold storage** receives the highest final score and is selected as the cooling-flexibility case. The sector combines a strong ATR fit with a relatively homogeneous operational logic: electricity demand is closely linked to maintaining temperature within acceptable bounds. This makes the flexibility mechanism technically bounded and suitable for optimisation, because demand can be shifted within thermal constraints without directly changing the core service provided.
2. **Distribution centres with EV fleet** are selected as the logistics case. This archetype captures electrification-driven demand growth through depot or workplace EV charging. The flexibility mechanism is operationally intuitive: charging can be shifted within vehicle-availability and energy-demand constraints. This makes the case directly relevant for future distribution-grid congestion, where transport electrification is expected to increase local peak demand.
3. **Dairy processing** is selected as the industrial heat case. Although it does not receive the highest overall score, it adds an important industrial archetype that is not covered by the two logistics-related cases. Its flexibility is represented through an industrial heat pump and thermal buffering, making it suitable for analysing how electrified process heat can respond to TDTR restrictions. Dairy is preferred over paper manufacturing because its heat-pump electrification pathway can be represented more transparently within the modelling framework.

Several high-scoring candidates are not selected. Cut flowers and ornamental shrubs under glass score highly because of their peaky load profile and apparent flexibility potential, but greenhouse horticulture combines several interacting flexibility mechanisms, such as lighting, CHP operation, heat buffering, and electric heat production. Modelling this would require a more complex multi-asset framework than used in this thesis. Paper and cardboard manufacturing is also not selected because the final case set already includes an industrial process-heat archetype through dairy processing. Hospitals, data centres, and supermarkets are considered relevant but less suitable for the core case set because of service-continuity constraints or weaker data availability.

D

Call patterns for different substations

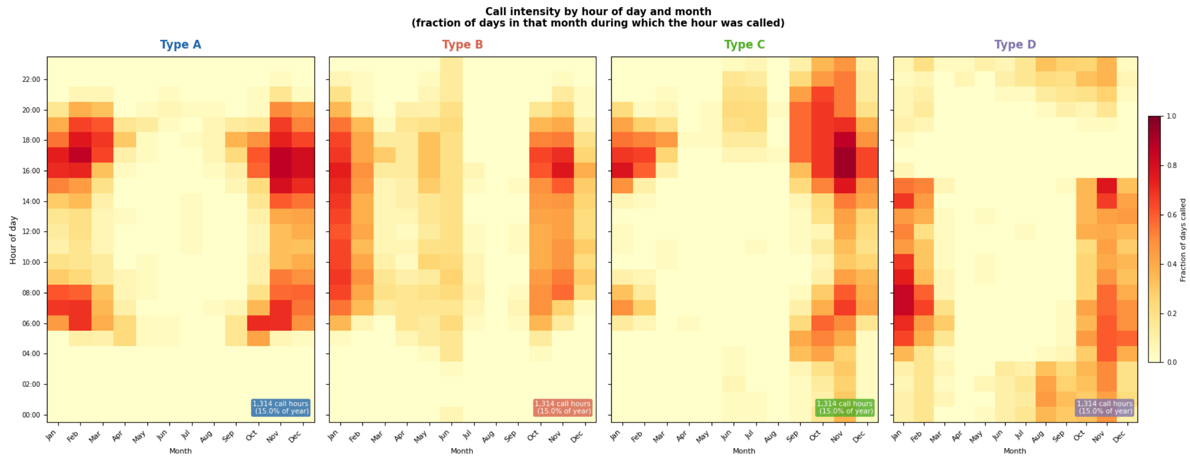

Figure D.1: Call intensity by hour of day and month for the analysed substation profiles.

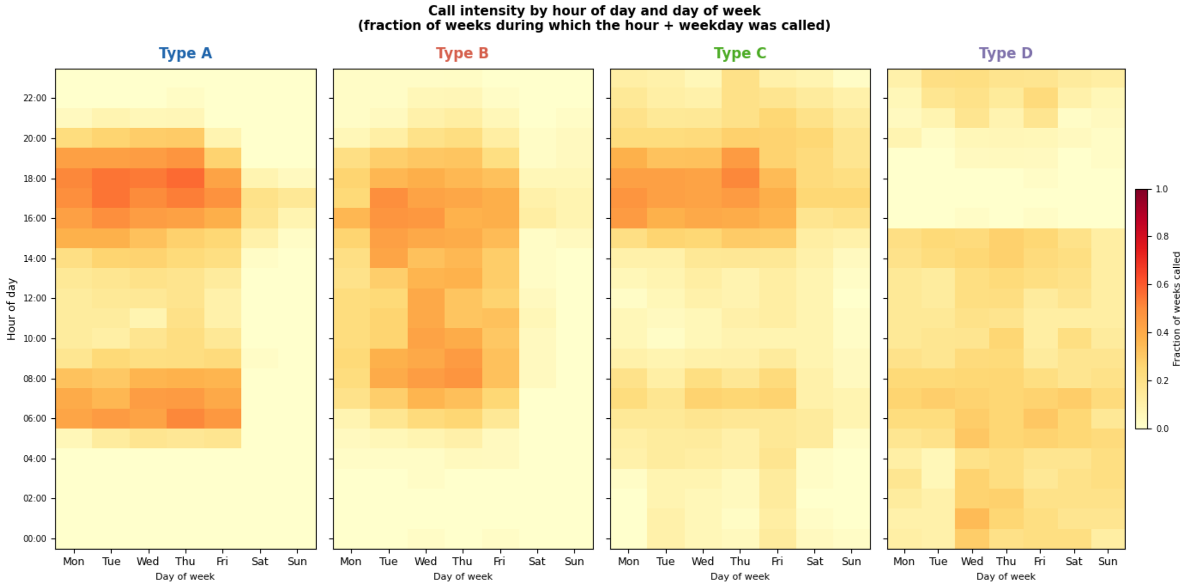

Figure D.2: Call intensity by hour of day and day of week for the analysed substation profiles.

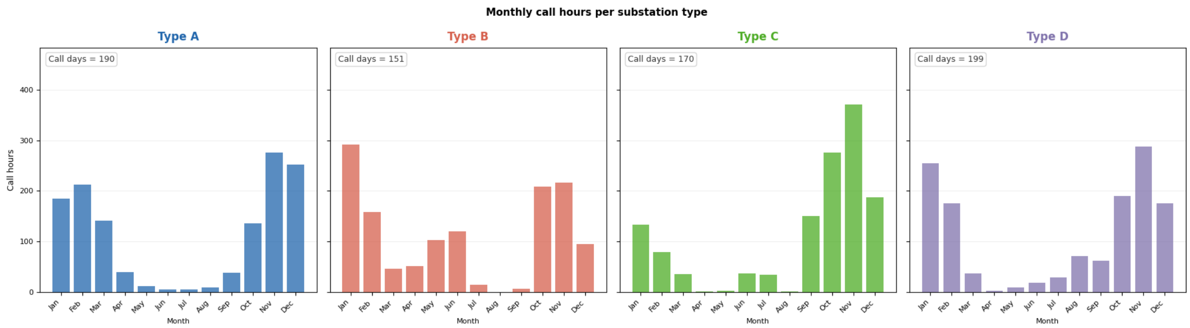

Figure D.3: Monthly call hours per substation profile.

E

Verification and validation of the model

Following Sargent (2010), this chapter distinguishes *verification*, whether the conceptual MILP formulation is implemented correctly in code, from *validation*, whether the model's behaviour is a plausible representation of the intended system. The bound, mutual-exclusivity and cyclic-closure checks (Sections E.1 and E.2), together with the analytical cross-check of the structural minimum firm fraction $f_{\min}$ (Section E.3), constitute the *verification* evidence: they confirm that the intended constraints are active and that the model output matches independently derived values. Each section additionally reports *validation* evidence through face validity, by showing that the buffer use, battery cycling and compliance-cost curve shapes respond to the call profile and differ across assets in the way the intended flexibility mechanisms predict.

E.1. Dairy thermal-buffer dynamics

The thermal-buffer dynamics in the dairy heat-pump model were checked for the Flex + TDTR scenario across all four call-hour profiles (Figures E.1–E.4). Each figure shows the buffer state of charge (SOC) in the upper panel and the charge and discharge power in the lower panel. Green shading indicates charging periods and orange shading indicates discharging periods.

The plots verify three implementation properties. First, the SOC remains within its capacity bounds: it never drops below zero and does not exceed the installed buffer capacity. Second, charging and discharging are separated in time, confirming that the binary buffer-mode logic introduced in the MILP formulation is active. Third, the SOC evolves consistently with the buffer recursion: it increases when the buffer is charged, decreases when it is discharged, and remains close to zero when the buffer is not used.

For face validity, the buffer-use pattern differs across call profiles, as expected when the optimisation responds to the timing of local TDTR calls. Types A, B and D show a relatively regular daily cycle, while Type C produces a less regular SOC trajectory because its call-hour pattern is more concentrated during the selected period. Together, these checks confirm that the binary logic prevents artificial simultaneous charging and discharging and that the buffer responds to the call profile in a physically plausible way.

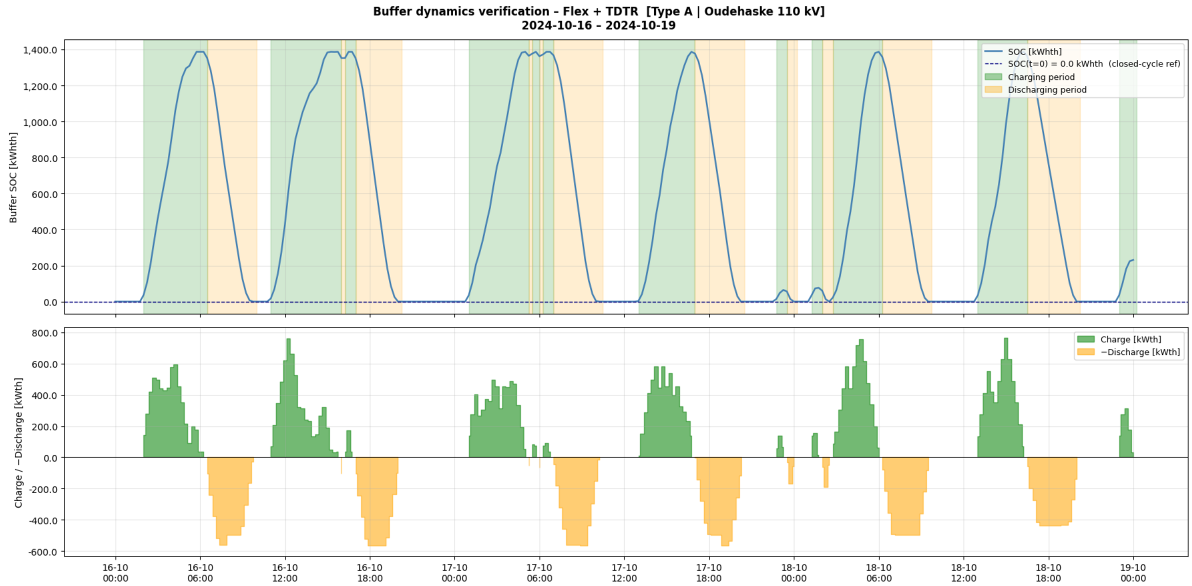

Figure E.1: Dairy thermal-buffer dynamics under call profile Type A (Oudehaske 110 kV), Flex + TDTR scenario.

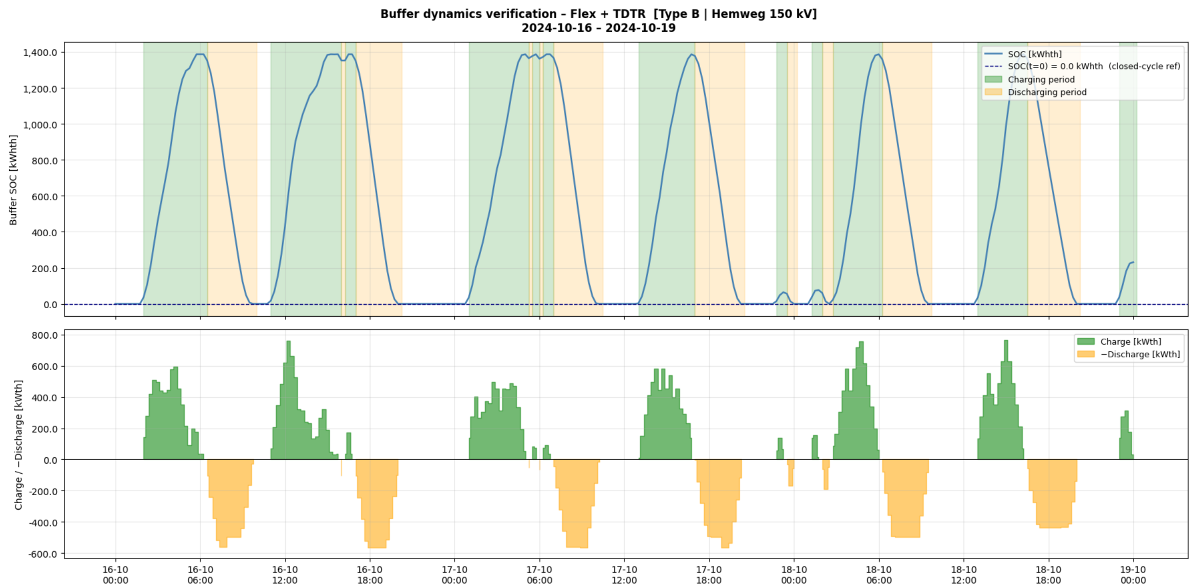

Figure E.2: Dairy thermal-buffer dynamics under call profile Type B (Hemweg 150 kV), Flex + TDTR scenario.

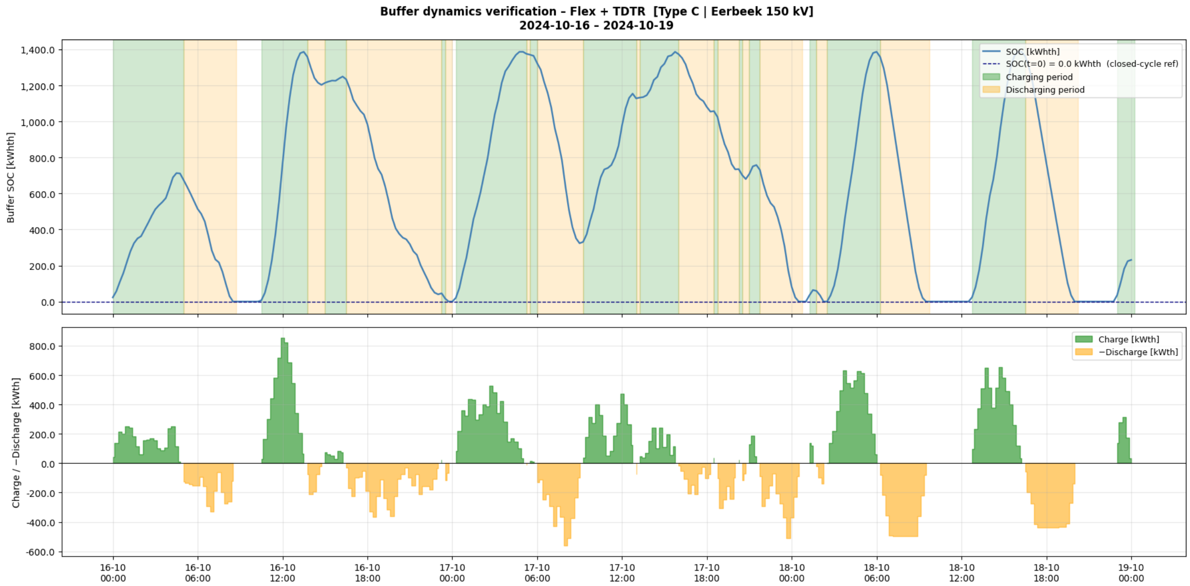

Figure E.3: Dairy thermal-buffer dynamics under call profile Type C (Eerbeek 150 kV), Flex + TDTR scenario.

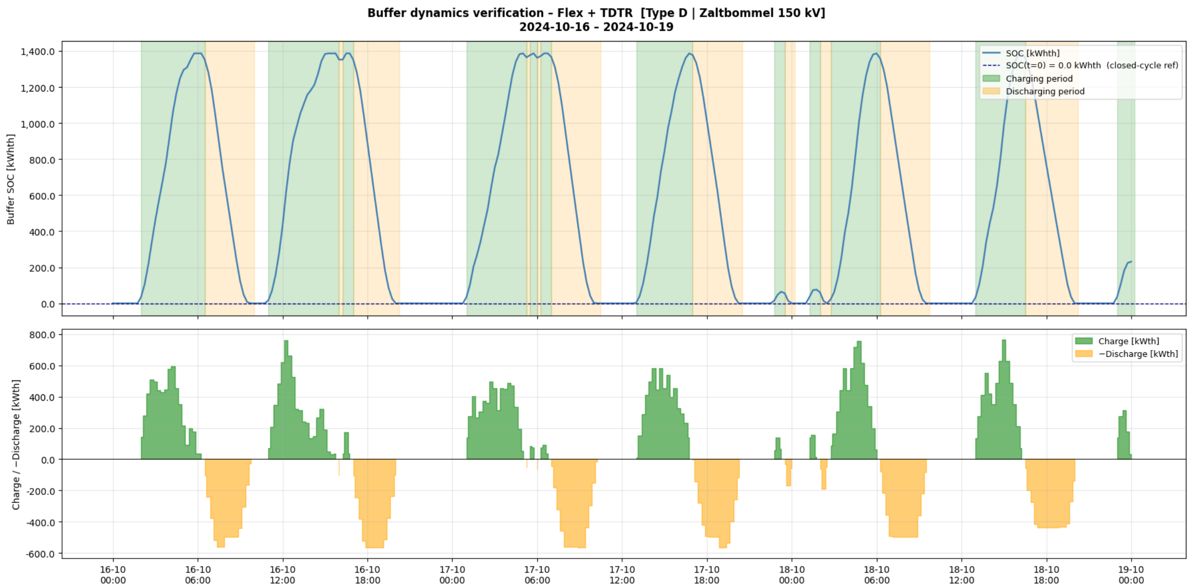

Figure E.4: Dairy thermal-buffer dynamics under call profile Type D (Zaltbommel 150 kV), Flex + TDTR scenario.

E.2. Battery charging

The battery SOC dynamics were checked for the Oudehaske case (Type A, scenario Smart charging + TDTR + Battery) across all three battery sizes (Figures E.5–E.7). Each figure shows the SOC trajectory (top) and the charge/discharge power (bottom) over a representative period, with the ATR call-hours shaded.

Three properties verify that the SOC dynamics are implemented correctly. First, the SOC stays within its bounds at all times: it never exceeds 100% nor drops below 0%, so the energy-capacity constraint is respected. Second, the charge and discharge power never exceed the rated power of each battery (±1000 kW for battery A, ±2000 kW for battery B, ±4000 kW for battery C), and charging and discharging never occur simultaneously, confirming the power limits and the charge/discharge separation. Third, the SOC closes over the cycle: it returns to (or above) the start level marked by the dashed line, which confirms the cyclic end constraint soc$[T-1] \geq$ soc$_{\text{init}}$ and rules out artificial end-of-horizon depletion. Together with the closed-cycle energy balance (charge into SOC equals discharge from SOC plus the small end-of-cycle surplus, within numerical tolerance), this confirms the implementation is correct.

The behaviour is also physically consistent with the intended use of the battery under TDTR, which provides face validity. Discharging (orange) is concentrated within the shaded ATR call-hours while charging (blue) takes place outside them: the battery is emptied to keep the total load under the firm limit during a call and recharged once the call ends. This pattern is visible across all three sizes and becomes more pronounced as the battery grows. Battery A (Fig. E.5) cycles fully and frequently because its limited energy content forces rapid swings, whereas battery C (Fig. E.7) shows deeper, smoother cycles and sustains discharge over longer call windows thanks to its larger capacity and higher power rating.

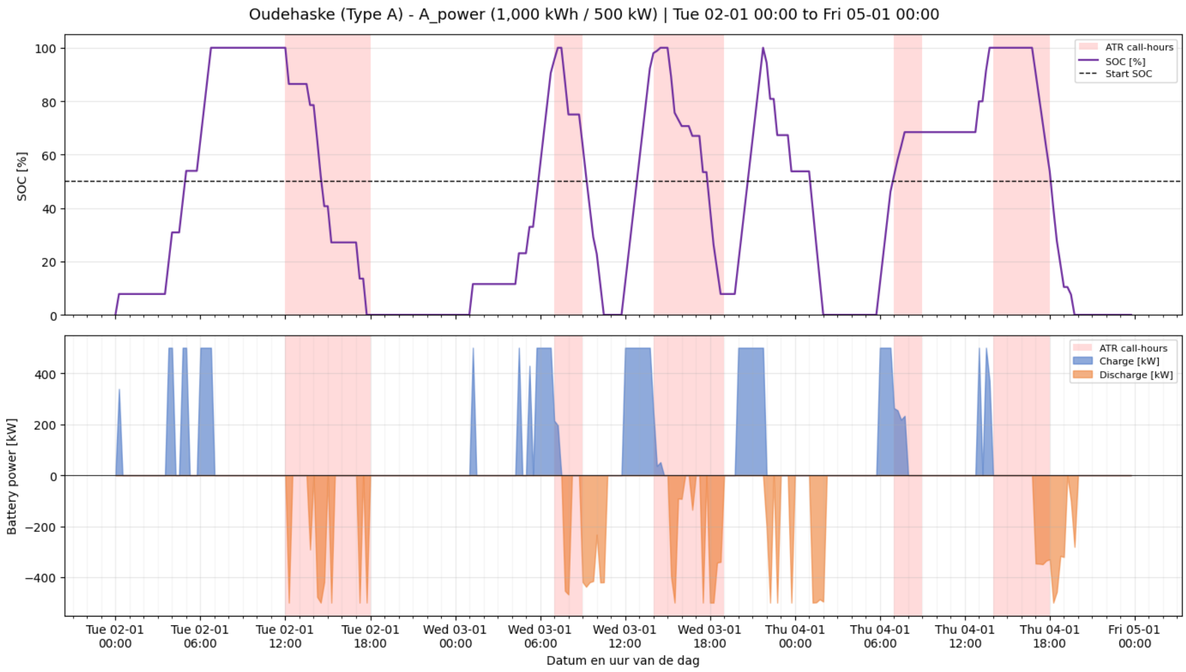

Figure E.5: Battery dynamics for battery A (Oudehaske, Type A, Smart charging + TDTR + Battery). Top: SOC trajectory closing over the cycle; bottom: charge (positive) and discharge (negative).

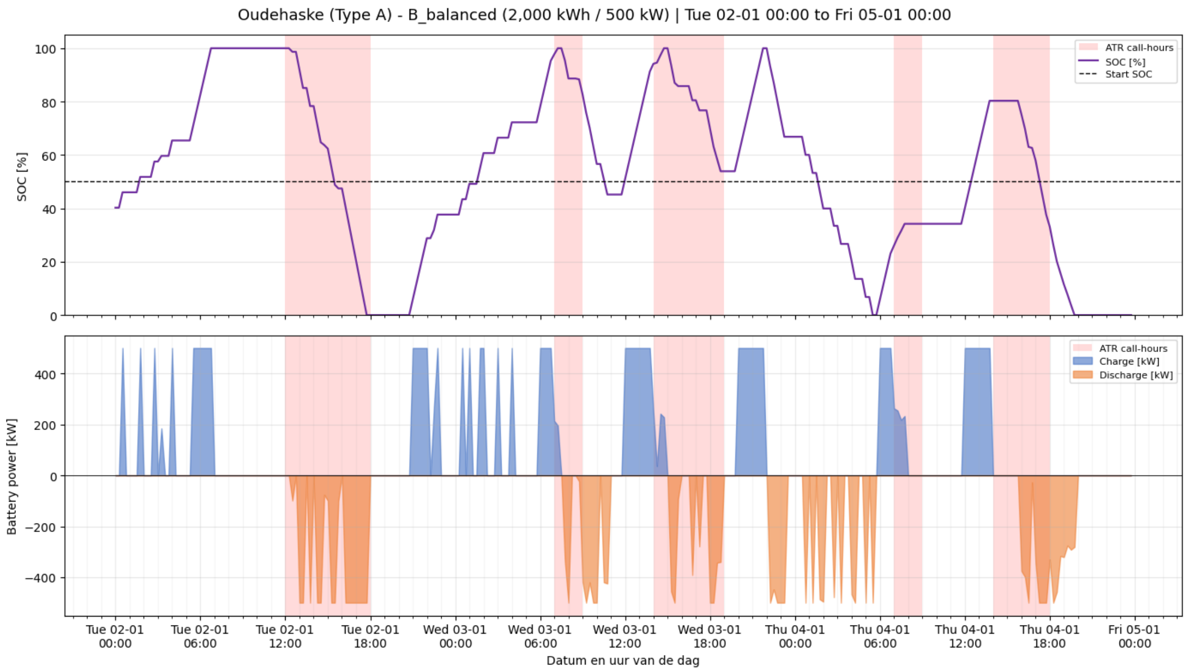

Figure E.6: Battery dynamics for battery B (Oudehaske, Type A, Smart charging + TDTR + Battery). Top: SOC trajectory closing over the cycle; bottom: charge (positive) and discharge (negative).

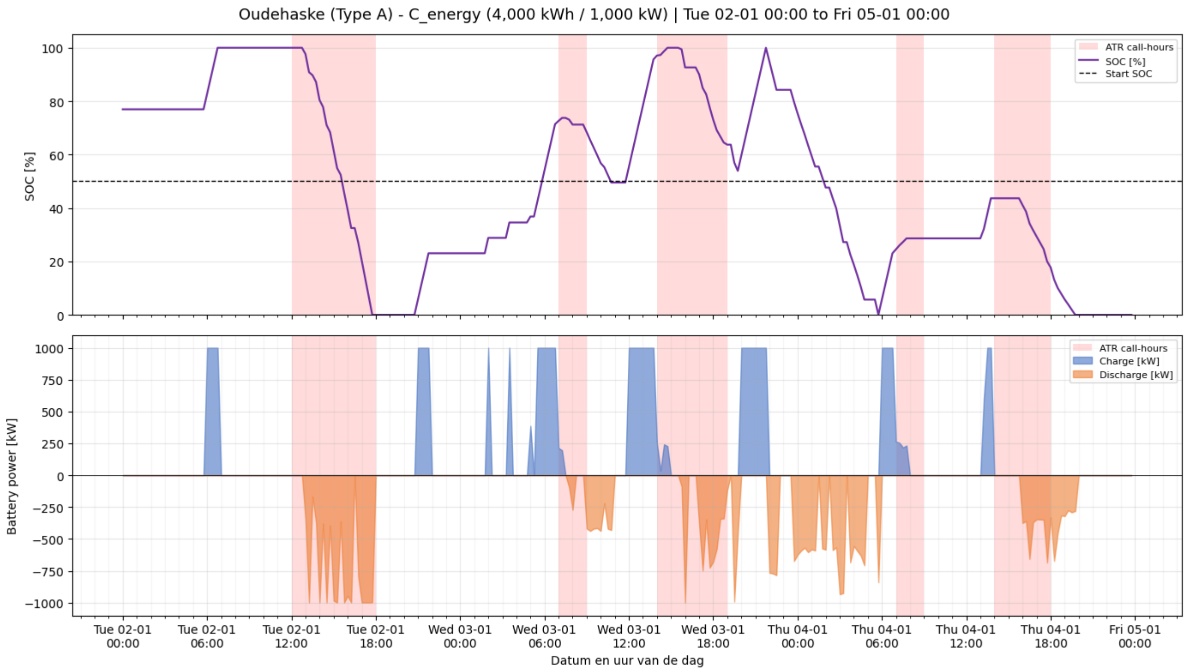

Figure E.7: Battery dynamics for battery C (Oudehaske, Type A, Smart charging + TDTR + Battery). Top: SOC trajectory closing over the cycle; bottom: charge (positive) and discharge (negative).

E.3. Analytical cross-check using a high NSE penalty

To verify that the energy-balance and TDTR-cap constraints are implemented correctly, the NSE penalty was set to €10,000/kWh, so that feasibility dominates cost entirely. The optimiser then accepts no avoidable non-served energy; under the hard cap the only remaining NSE is the structurally unavoidable curtailment of inflexible load (Eq. (4.10)), which vanishes exactly from the structural minimum firm fraction onward. In the compliance-cost curves (Figures E.9–E.8) this is the point where the total-cost and NSE curves kink to zero while the kW-contract cost stays flat. Verification succeeds if this $f_{\min}$ matches the value computed independently from the load profile, namely the maximum unavoidable load during call hours divided by the contract reference capacity.

The three curve shapes differ per asset, which is itself consistent with the intended flexibility representations (face validity): logistics (Fig. E.9) is convex with a kink at $f_{\min} = 0.80$, reflecting exploited EV-charging flexibility; dairy (Fig. E.10) is linear with a kink at $f_{\min} = 0.85$, reflecting an almost fully inflexible continuous process (a convex curve here would signal a modelling error); and cold storage (Fig. E.8) is convex with the lowest kink at $f_{\min} = 0.54$, reflecting thermal-mass buffering.

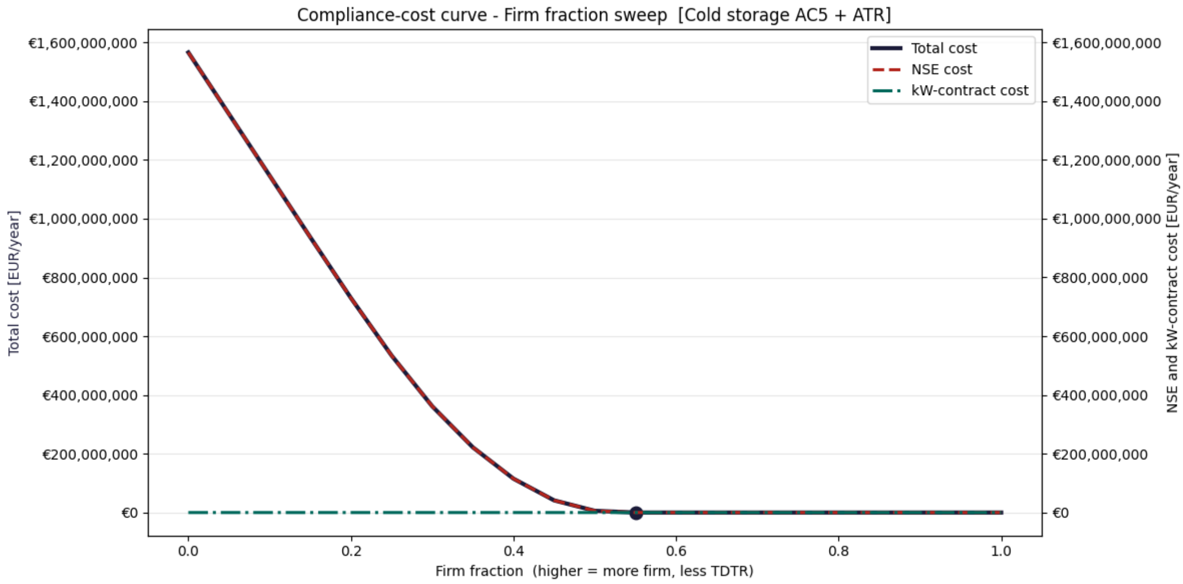

Figure E.8: Compliance-cost curve for the cold storage case (firm fraction sweep, AC5 + ATR) under an extremely high NSE penalty ($\lambda_{\text{NSE}} \to \infty$). The curve kinks to zero at $f_{\min} = 0.54$.

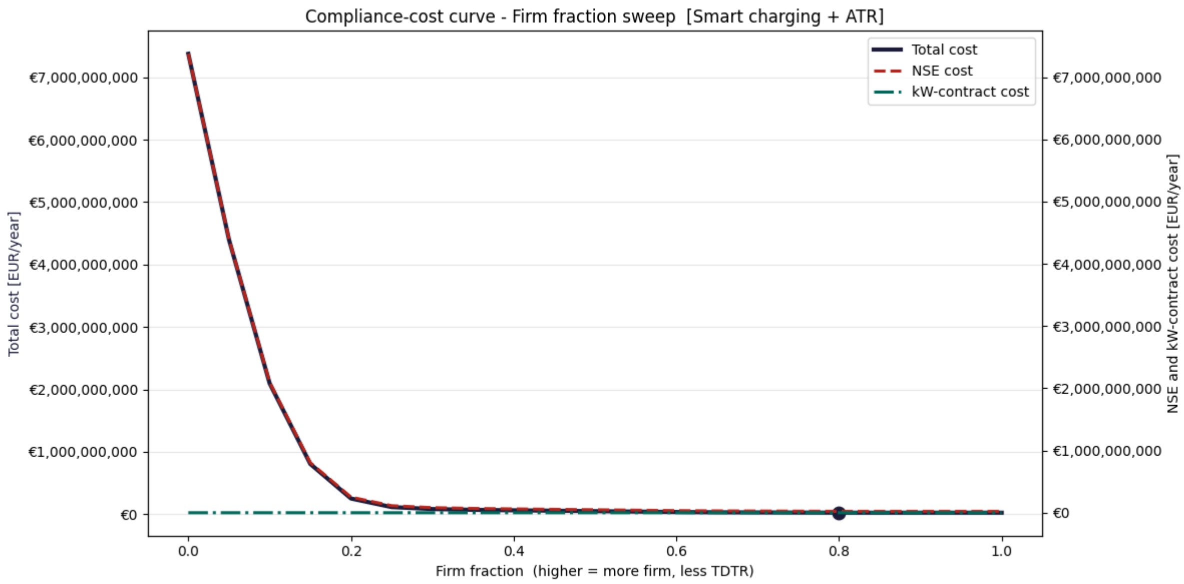

Figure E.9: Compliance-cost curve for the logistics case (firm fraction sweep, smart charging + ATR) under an extremely high NSE penalty ($\lambda_{\text{NSE}} \to \infty$). The curve kinks to zero at $f_{\min} = 0.80$.

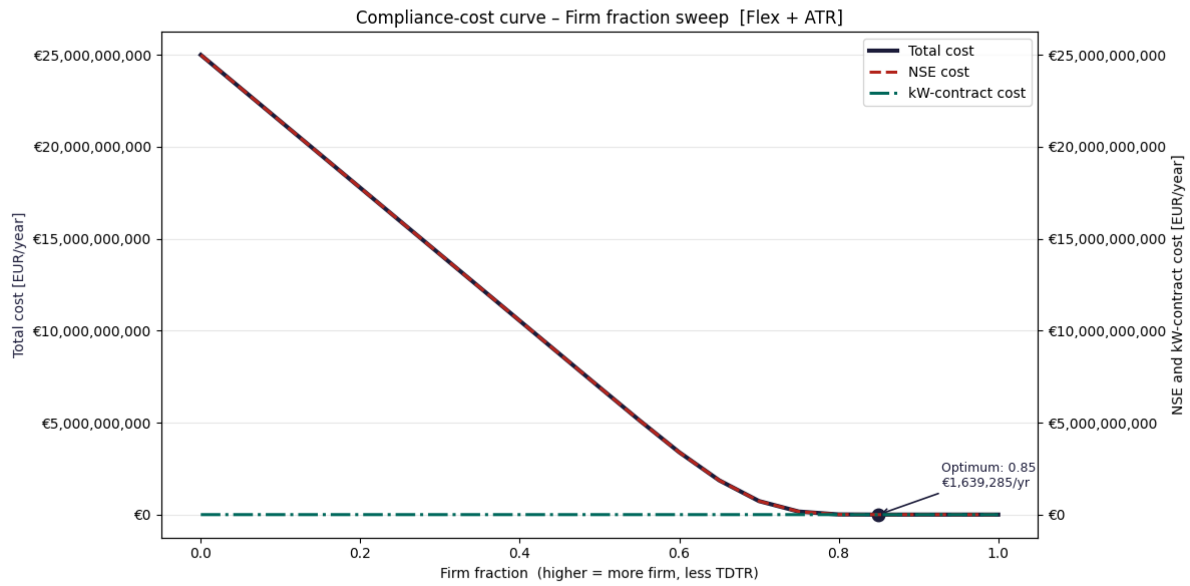

Figure E.10: Compliance-cost curve for the dairy case (firm fraction sweep, flex + ATR) under an extremely high NSE penalty ($\lambda_{\text{NSE}} \to \infty$). The linear decline kinks to zero at $f_{\min} = 0.85$.

F
Sensitivity analysis

F.1. Asset parameters sensitivity

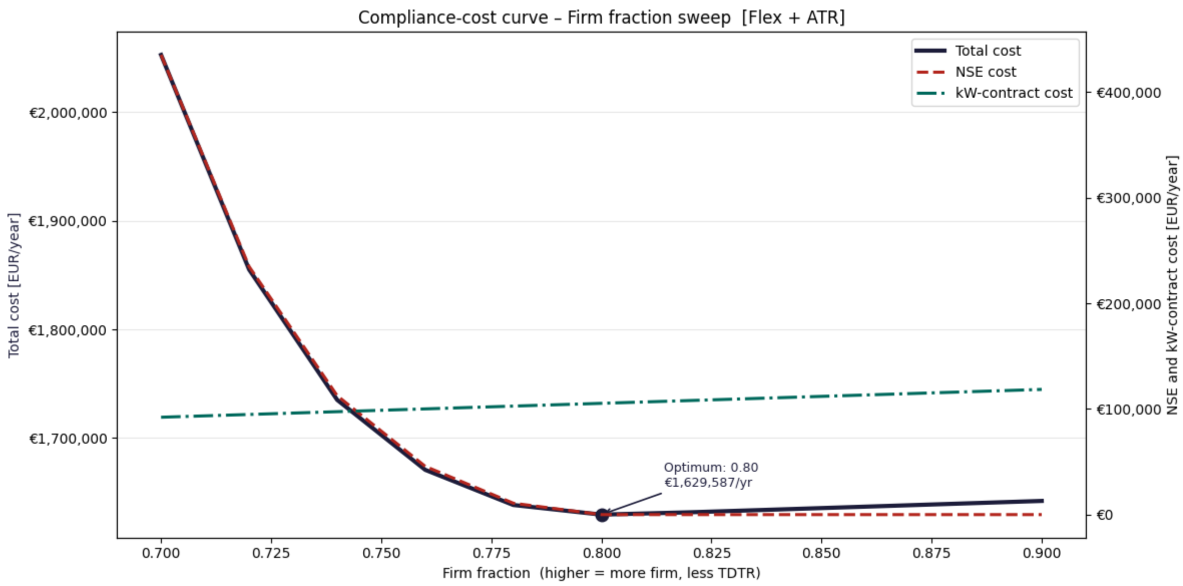

Figure F.1: Compliance-cost curve for the dairy case with 20% flexibility (firm-fraction sweep).

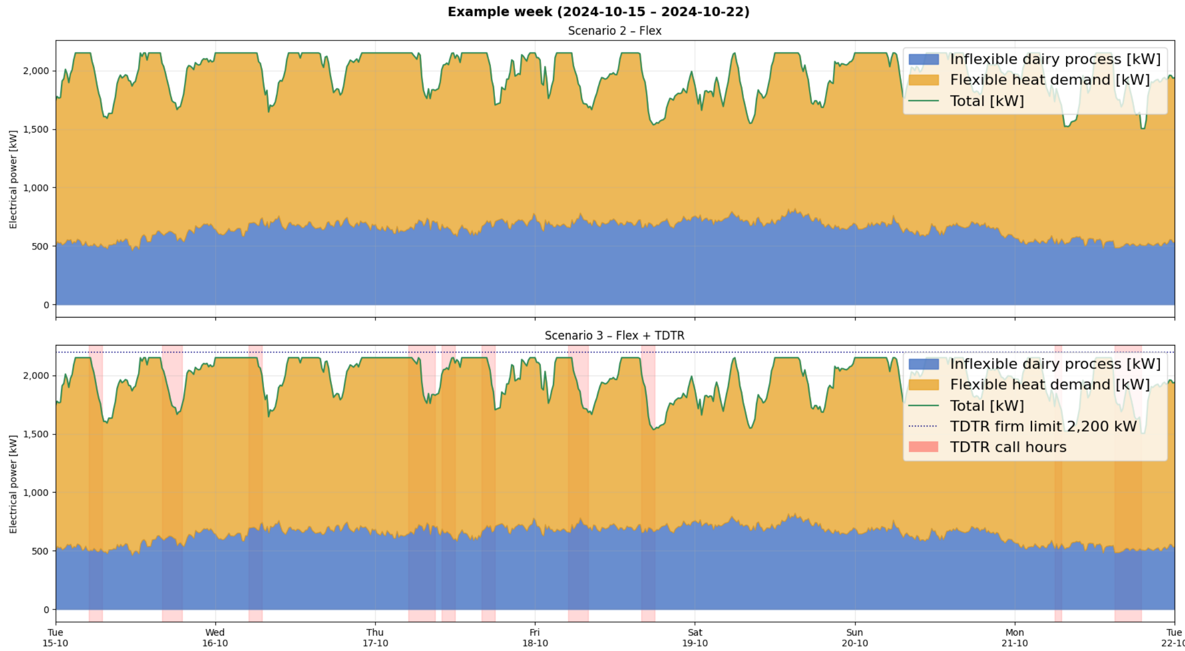

Figure F.2: Weekly dispatch overview of Scenario 2 (Flex) and Scenario 3 (Flex + 20% TDTR) for the dairy case with 20% flexibility.

G

Extra outputs

Table G.1: Top 5% in call-hours for customer load profiles and local call profiles.

Call profile	Reference: top 5% in call-hours			Flex: top 5% in call-hours			Flex + TDTR: top 5% in call-hours		
	Cold storage [%]	Logistics [%]	Dairy [%]	Cold storage [%]	Logistics [%]	Dairy [%]	Cold storage [%]	Logistics [%]	Dairy [%]
Type A	6.22	33.05	17.35	1.94	5.82	5.31	0.00	0.00	0.00
Type B	10.33	20.21	17.69	5.93	4.14	13.58	0.00	0.00	0.00
Type C	24.14	33.39	6.62	15.69	19.63	1.48	0.00	0.00	0.00
Type D	8.11	7.31	7.36	15.06	30.13	2.36	0.00	0.00	0.00

Table G.2: Grid-impact indicators for the cold-storage scenarios across local call profiles.

Call profile	Scenario	Peak [kW]	Peak call [kW]	Peak non-call [kW]	Abs. red. [%]	Rel. red. [%]	Adj. LF [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Type A	Reference	290.00	268.38	290.00	0.00	0.00	74.56	6.22	0.00	0.00
Type A	Flex	255.89	255.89	255.89	11.76	11.76	69.92	1.94	4.86	16.11
Type A	Flex + TDTR	255.45	175.45	255.45	11.91	11.91	68.55	0.00	4.60	18.09
Type B	Reference	290.00	269.24	290.00	0.00	0.00	74.56	10.33	0.00	0.00
Type B	Flex	255.89	242.18	255.89	11.76	11.76	69.92	5.93	4.86	10.08
Type B	Flex + TDTR	261.98	181.83	261.98	9.66	11.76	68.82	0.00	4.42	11.79
Type C	Reference	290.00	278.94	290.00	0.00	0.00	74.56	24.14	0.00	0.00
Type C	Flex	255.89	255.89	255.89	11.76	11.76	69.92	15.69	4.86	11.90
Type C	Flex + TDTR	274.25	194.59	274.25	5.43	22.56	70.82	0.00	4.69	13.14
Type D	Reference	290.00	266.86	290.00	0.00	0.00	74.56	8.11	0.00	0.00
Type D	Flex	255.89	255.89	255.89	11.76	11.76	69.92	15.06	4.86	-4.68
Type D	Flex + TDTR	256.39	194.59	256.39	11.59	11.59	70.91	0.00	4.48	-0.96

Table G.3: Grid-impact indicators for the logistics scenarios across local call profiles.

Call profile	Scenario	Peak [kW]	Peak call [kW]	Peak non-call [kW]	Abs. red. [%]	Rel. red. [%]	Adj. LF [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Type A	Reference	4573.15	4573.15	4403.20	0.00	0.00	39.68	33.05	0.00	0.00
Type A	Smart charging	3937.89	3937.89	3937.89	13.89	68.38	31.04	5.82	6.93	52.00
Type A	Smart charging + TDTR	4019.19	1307.92	4019.19	12.11	71.40	31.02	0.00	6.83	56.08
Type B	Reference	4573.15	4363.69	4573.15	0.00	0.00	39.68	20.21	0.00	0.00
Type B	Smart charging	3937.89	3911.29	3937.89	13.89	68.38	31.04	4.14	6.93	44.19
Type B	Smart charging + TDTR	3937.89	2314.01	3937.89	13.89	68.38	31.01	0.00	7.02	45.65
Type C	Reference	4573.15	4573.15	4164.85	0.00	0.00	39.68	33.39	0.00	0.00
Type C	Smart charging	3937.89	3937.89	3937.89	13.89	68.38	31.04	19.63	6.93	37.66
Type C	Smart charging + TDTR	3733.13	3521.32	3733.13	18.37	68.38	31.03	19.71	6.92	37.80
Type D	Reference	4573.15	3818.27	4573.15	0.00	0.00	39.68	7.31	0.00	0.00
Type D	Smart charging	3937.89	3937.89	3937.89	13.89	68.38	31.04	30.13	6.93	-41.85
Type D	Smart charging + TDTR	3823.25	3521.32	3823.25	16.40	68.38	31.03	30.41	6.90	-41.50

Table G.4: Grid-impact indicators for the dairy heat-pump 10% flex scenarios across local call profiles.

Call profile	Scenario	Peak [kW]	Peak call [kW]	Peak non-call [kW]	Abs. red. [%]	Rel. red. [%]	Adj. LF [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Type A	Reference	2500.00	2386.01	2500.00	0.00	0.00	86.62	17.35	0.00	0.00
Type A	Flex with buffer	2364.56	2364.56	2364.56	5.42	5.42	85.75	10.02	0.40	1.84
Type A	Flex + TDTR with buffer	2364.56	2229.41	2364.56	5.42	5.42	85.21	0.00	0.41	1.96
Type B	Reference	2500.00	2430.89	2500.00	0.00	0.00	86.62	17.69	0.00	0.00
Type B	Flex with buffer	2364.56	2364.56	2364.56	5.42	5.42	85.75	19.94	0.40	0.98
Type B	Flex + TDTR with buffer	2410.12	2282.50	2410.12	3.60	3.60	85.74	0.00	0.40	1.04
Type C	Reference	2500.00	2386.01	2500.00	0.00	0.00	86.62	6.62	0.00	0.00
Type C	Flex with buffer	2364.56	2364.56	2364.56	5.42	5.42	85.75	3.85	0.40	1.27
Type C	Flex + TDTR with buffer	2364.56	2229.41	2364.56	5.42	5.42	84.88	0.00	0.43	1.48
Type D	Reference	2500.00	2335.71	2500.00	0.00	0.00	86.62	7.36	0.00	0.00
Type D	Flex with buffer	2364.56	2361.87	2364.56	5.42	5.42	85.75	13.11	0.40	-0.66
Type D	Flex + TDTR with buffer	2364.56	2227.50	2364.56	5.42	5.42	85.57	0.00	0.40	-0.54

Table G.5: Grid-impact indicators for the dairy heat-pump 20% flex scenarios across local call profiles.

Call profile	Scenario	Peak [kW]	Peak call [kW]	Peak non-call [kW]	Abs. red. [%]	Rel. red. [%]	Adj. LF [%]	Top 5% call [%]	WGPR [%]	Call-hour red. [%]
Type A	Reference	2500.00	2386.01	2500.00	0.00	0.00	86.62	17.35	0.00	0.00
Type A	Flex with buffer	2347.64	2347.64	2347.64	6.09	6.09	85.06	2.85	0.73	3.29
Type A	Flex + TDTR with buffer	2347.64	2172.50	2347.64	6.09	6.09	84.39	0.00	0.75	3.75
Type B	Reference	2500.00	2430.89	2500.00	0.00	0.00	86.62	17.69	0.00	0.00
Type B	Flex with buffer	2347.64	2347.64	2347.64	6.09	6.09	85.06	20.39	0.73	1.87
Type B	Flex + TDTR with buffer	2383.67	2282.50	2383.67	4.65	4.65	85.80	0.00	0.73	1.91
Type C	Reference	2500.00	2386.01	2500.00	0.00	0.00	86.62	6.62	0.00	0.00
Type C	Flex with buffer	2347.64	2347.64	2347.64	6.09	6.09	85.06	1.75	0.73	2.38
Type C	Flex + TDTR with buffer	2347.64	2072.82	2347.64	6.09	6.09	83.87	0.00	0.76	3.09
Type D	Reference	2500.00	2335.71	2500.00	0.00	0.00	86.62	7.36	0.00	0.00
Type D	Flex with buffer	2347.64	2347.64	2347.64	6.09	6.09	85.06	1.21	0.73	-1.17
Type D	Flex + TDTR with buffer	2347.64	2172.50	2347.64	6.09	6.09	85.05	0.00	0.74	-0.69

Bibliography

Akerma, M., Hoang, H. M., Leducq, D., & Delahaye, A. (2020). Experimental characterization of demand response in a refrigerated cold room. *International Journal of Refrigeration*, *113*, 131–141. https://doi.org/10.1016/j.ijrefrig.2020.02.012

Alavijeh, N. M., Steen, D., Tuan, L. A., & Nyström, S. (2024). Capacity limitation based local flexibility market for congestion management in distribution networks: Design and challenges. *International Journal of Electrical Power and Energy Systems*, *156*, 109742. https://doi.org/10.1016/j.ijepes.2023.109742

Attar, M., Repo, S., Heikkilä, V., & Suominen, O. (2025). Congestion management in distribution grids using local flexibility markets: Investigating influential factors. *International Journal of Electrical Power & Energy Systems*. https://doi.org/10.1016/j.ijepes.2025.110601

Autoriteit Consument & Markt. (2024). *Codewijzigingsvoorstel alternatieve transportrechten* (tech. rep.). Autoriteit Consument & Markt. https://www.acm.nl/system/files/documents/codewijzigingsvoorstel-alternatieve-transportrechten.pdf

Belastingdienst. (2026). *Energiebelasting*. https://www.belastingdienst.nl/wps/wcm/connect/bldcontentnl/belastingdienst/zakelijk/overige_belastingen/belastingen_op_milieugrondslag/energiebelasting/energiebelasting

Bjarghov, S., Foslie, S. S., Askeland, M., Rana, R., & Taxt, H. (2024). Enhancing grid hosting capacity with coordinated non-firm connections in industrial energy communities. *Smart Energy*, *15*, 100154. https://doi.org/10.1016/j.segy.2024.100154

CE Delft. (2025). *Achtergrondrapport tijdsafhankelijke nettarieven grootverbruikers* (Report No. 240472). CE Delft. https://www.netbeheernederland.nl/sites/default/files/2025-05/ce_delft_240472_achtergrondrapport_tijdsafhankelijke_nettarieven_grootverbruikers_def_002.pdf

CE Delft & Merosch. (2024). *Oplossingen voor netcongestie bij bedrijven: Praktijkvoorbeelden met toekomstig potentieel en kosten per oplossing* (Onderzoeksrapport No. Projectnummer 3493). CE Delft and Merosch. https://ce.nl/wp-content/uploads/2024/06/230235_Onderzoeksrapport_Oplossingen_voor_netcongestie_bij_bedrijven.pdf

Chen, Z., Ziras, C., & Bindner, H. W. (2023). Modeling and characterization of the impact of capacity limitation services on distribution networks. *Sustainable Energy, Grids and Networks*, *35*, 101081. https://doi.org/10.1016/j.segan.2023.101081

Common Futures, KWA Bedrijfsadviseurs, & EV Consult. (2026). *Het ontsluiten van afnameflexibiliteit voor het elektriciteitsnet: Verdiepende analyse van het flexibiliteitspotentieel in vier sectoren* (Rapport in opdracht van Rijksdienst voor Ondernemend Nederland (RVO)). Topsector Energie Systeemintegratie. https://energy-innovation.nl/documents/1764/CF_KWA_EVC_Het_ontsluiten_van_afnameflexibiliteit_feb_2026.pdf

Council of European Energy Regulators. (2023). *Ceer paper on alternative connection agreements* (tech. rep. No. C23-DS-83-06). CEER. https://www.ceer.eu

Cox, J., Belding, S., & Lowder, T. (2022). Application of a novel heat pump model for estimating economic viability and barriers of heat pumps in dairy applications in the united states. *Applied Energy*, *310*, 118499. https://doi.org/10.1016/j.apenergy.2021.118499

Damianakis, N., Chandra-Mouli, G. R., Bauer, P., & Yu, Y. (2023). Assessing the grid impact of electric vehicles, heat pumps & pv generation in dutch lv distribution grids. *Applied Energy*. https://doi.org/10.1016/j.apenergy.2023.121878

De Santi, F., Meeus, L., Beckstedde, E., Delarue, E., & Vitiello, S. (2025). Managing connection queues in distribution networks with flexible connection agreements. *Applied Energy*, *396*, 126260. https://doi.org/10.1016/j.apenergy.2025.126260

de Nooij, M., Koopmans, C., & Bijvoet, C. (2007). The value of supply security: The costs of power interruptions: Economic input for damage reduction and investment in networks. *Energy Economics*, *29*(2), 277–295. https://doi.org/10.1016/j.eneco.2006.05.022

ElaadNL. (2026). ElaadNL Charging Profiles Generator [Online tool for generating electric-vehicle charging profiles.]. https://charging.elaad.nl/

Energy Innovation NL. (n.d.). High Temperature Tank Thermal Energy Storage (HT-TTES). https://energy-innovation.nl/documents/990/Factsheet-HT-TTES.pdf

ENTSO-E. (2026). *Transparency platform: Day-ahead energy prices for the netherlands* [Dutch bidding zone (10YNL———-L).]. https://transparency.entsoe.eu/market/energyPrices

Golmohamadi, H. (2022). Demand-side management in industrial sector: A review of heavy industries. *Renewable and Sustainable Energy Reviews*, *156*, 111963. https://doi.org/10.1016/j.rser.2021.111963

Golmohamadi, H., & Asadi, A. (2020). Integration of joint power-heat flexibility of oil refinery industries to uncertain energy markets. *Energies*, *13*(18), 4874. https://doi.org/10.3390/en13184874

Heinrich, C., Ziras, C., Jensen, T. V., Bindner, H. W., & Kazempour, J. (2021). A local flexibility market mechanism with capacity limitation services. *Energy Policy*, *156*, 112335. https://doi.org/10.1016/j.enpol.2021.112335

Hennig, R. J., de Vries, L. J., & Tindemans, S. H. (2023). Congestion management in electricity distribution networks: Smart tariffs, local markets and direct control. *Utilities Policy*, *85*, 101660. https://doi.org/10.1016/j.jup.2023.101660

Hennig, R. J., de Vries, L. J., & Tindemans, S. H. (2024). Risk vs. restriction: An investigation of capacity-limitation-based congestion management in electric distribution grids. *Energy Policy*, *186*, 113976. https://doi.org/10.1016/j.enpol.2023.113976

Hofmann, M., et al. (2025). Grid tariff design and peak demand shaving: A comparative assessment of peak pricing schemes. *Energy Economics*. https://doi.org/10.1016/j.eneco.2024.107559

Kahmann, R., & Kleiburg, R. (2023). *Onderzoek ter bewustwording van de waarde van flexibiliteit*. Nederlandse Vereniging Duurzame Energie. https://www.nvde.nl/wp-content/uploads/2023/02/23-01-24-NVDE-Flex-bewustzijn_final.pdf

Kaya, H. D., Leijten, M., Schraven, D., & Chan, P. W. (2026). Navigating lock-ins for adaptation: A case study of grid capacity planning in the dutch energy transition. *Sustainability Science*. https://doi.org/10.1007/s11625-025-01793-6

Leinauer, C., Schott, P., Fridgen, G., Keller, R., Ollig, P., & Weibelzahl, M. (2022). Obstacles to demand response: Why industrial companies do not adapt their power consumption to volatile power generation. *Energy Policy*, *165*, 112876. https://doi.org/10.1016/j.enpol.2022.112876

Liander. (2026a). *Tarieven voor aansluiting en transport elektriciteit 2026* (tech. rep.). Liander. https://www.liander.nl/-/media/files/tarieven/grootzakelijk/tarieven-2026/tarieven-voor-aansluiting-en-transport-elektriciteit-2026_2.pdf

Liander. (2026b). *Welke elektriciteitsaansluiting heeft u nodig?* https://www.liander.nl/grootzakelijk/aansluiting/soorten-elektriciteitsaansluitingen

Michaelis, A., Schneider, M., & Weibelzahl, M. (2025). Designing local flexibility markets: A toolbox for policymakers and market operators. *Energy*. https://doi.org/10.1016/j.energy.2025.136051

Ministerie van Klimaat en Groene Groei. (2025). *Schakelen naar de toekomst: Over bekostiging elektriciteitsinfrastructuur* (Report). Rijksoverheid. https://www.rijksoverheid.nl/documenten/rapporten/2025/03/07/schakelen-naar-de-toekomst-over-bekostiging-elektriciteitsinfrastructuur

Monterde, M. R., Alvarez, E. F., & Valarezo, O. (2025). Non-firm grid connections: A review of access types, mechanisms, and regulatory frameworks. *Current Sustainable/Renewable Energy Reports*, *12*, 23. https://doi.org/10.1007/s40518-025-00268-7

Ndraha, N., Hsiao, H.-I., Vlajic, J., Yang, M.-F., & Lin, H.-T. V. (2018). Time-temperature abuse in the food cold chain: Review of issues, challenges, and recommendations. *Food Control*, *89*, 12–21. https://doi.org/10.1016/j.foodcont.2018.01.027

Netbeheer Nederland. (2024). *Position paper: Alternatieve transportrechten* (Position paper). Netbeheer Nederland. https://www.netbeheernederland.nl/sites/default/files/2024-03/position_paper_alternatieve_transportrechten_v1.1_-_jan_2024.pdf

Netbeheer Nederland. (2025). *Ecorys: 'Versnelde netinvesteringen zouden miljarden aan maatschappelijke waarde opleveren'*. https://www.netbeheernederland.nl/artikelen/nieuws/ecorys-versnelde-netinvesteringen-zouden-miljarden-aan-maatschappelijke-waarde

Netbeheer Nederland. (2026). *Codewijzigingsvoorstel tijdsafhankelijke transporttarieven voor grootverbruikers bij de distributiesysteembeheerders* (tech. rep.) (Submitted to the Authority for Consumers and Markets (ACM)). Netbeheer Nederland. https://www.acm.nl/system/files/documents/

codewijzigingsvoorstel - tijdsafhankelijke - transporttarieven - voor - grootverbruikers - bij - de - distributiesysteembeheerders.pdf

Praktiknjo, A. (2016). The value of lost load for sectoral load shedding measures: The german case with 51 sectors. *Energies*, *9*(2), 116. https://doi.org/10.3390/en9020116

Repo, S., Kilkki, O., Annala, S., Terras, J.-M., & Almeida, B. (2020). Local flexibility market balance settlement. *CIRED 2020 Berlin Workshop (CIRED 2020)*, 642–645. https://doi.org/10.1049/oap-cired.2021.0180

Rijksdienst voor Ondernemend Nederland. (2024). *Industriële warmtepompen*. https://www.rvo.nl/onderwerpen/energie-besparen-de-industrie/industriele-warmtepompen

Santecchia, A., Kantor, I., Castro-Amoedo, R., & Maréchal, F. (2022). Industrial flexibility as demand side response for electrical grid stability. *Frontiers in Energy Research*, *10*, 831462. https://doi.org/10.3389/fenrg.2022.831462

Sargent, R. G. (2010). Verification and validation of simulation models. *Proceedings of the 2010 Winter Simulation Conference (WSC)*, 166–183. https://doi.org/10.1109/WSC.2010.5679166

Schittekatte, T., et al. (2022). *An analysis of time-of-use and critical peak pricing* (tech. rep. No. 2022-015). MIT CEEPR Working Paper.

Serpe, L., Fürmann, T., Qussous, R., & Weidlich, A. (2023). Demand-side flexibility of the german steel industry: A case study. *2023 19th International Conference on the European Energy Market (EEM)*, 1–6. https://doi.org/10.1109/EEM58374.2023.10161699

Skok, M., Wagmann, L., & Baričević, T. (2022). Use of flexibility in distribution networks: An overview of eu and croatian legal framework and trends. *Journal of Energy*, *71*(1), 3–17. https://doi.org/10.37798/2022711343

Thomaßen, G., Fuhrmanek, A., Cadenovic, R., Pozo Camara, D., & Vitiello, S. (2024). *Redispatch and congestion management: Future-proofing the european power market* (tech. rep. No. JRC137685). European Commission, Joint Research Centre. https://doi.org/10.2760/853898

TNO. (2023, September). *Flexibiliteit in het elektriciteitssysteem: Interim rapport deel 1* (tech. rep. No. TNO 2023 P11051). TNO.

Trago, B., & Kroondijk, J. (2024). Hoe zitten de elektriciteits-, balans- en congestiemarkten in elkaar? https://www.dep.nl/wp-content/uploads/DEP-Presentatie-Elektriciteitsmarkten-Externe-versie.pdf

Verhagen, P., Hu, J., & Harmsen, R. (2025). Evaluating non-firm grid connections for battery energy storage systems: A co-optimization case study of the netherlands. *Energy Policy*, *208*, 114903. https://doi.org/10.1016/j.enpol.2025.114903

Zhang, Y., Liu, H., Wang, J., & Li, X. (2025). A review of industrial load flexibility enhancement for demand-response interaction. *Sustainability*, *17*(11), 4938. https://doi.org/10.3390/su17114938