Modelling Flow-based Market Coupling in the CORE region

A Data-driven and Interpretable Approach

Balázs Riskutia

Delft University of Technology

TU Delft

CROSS OPTIONS

Modelling Flow-based Market Coupling in the CORE region

A Data-driven and Interpretable Approach

by

Balázs Riskutia

in partial fulfilment of the requirements for the degree of

Master of Science
in Complex Systems Engineering and Management

at the Delft University of Technology,
to be defended publicly on Friday July 14, 2023 at 15:00 PM.

Student number: 5487099
Project duration: February, 2023 – July, 2023
Thesis committee:

Dr. ir. K. Bruninx,	TU Delft, first supervisor
Dr. A. F. Correljé,	TU Delft, second supervisor
R. Timmer,	Cross Options, external supervisor
Dr. ir. R. A. Hakvoort,	TU Delft, committee chair

Cover: High voltage towers; Photo: kalafoto - Fotolia (Modified)

CROSS OPTIONS

TU Delft

Abstract

Integrating renewables in the electricity system in a cost-efficient way requires massive transmission system investments and the efficient use of available transmission capacity. Markets are pivotal in the latter, especially in coordinating flows between countries. In the European Union, flow-based market coupling (FBMC) arose as the preferred market-based cross-border capacity allocation method, which has recently been extended to the CORE region, involving 12 countries of the EU. While the expansion of the flow-based methodology brings the EU closer to a single internal electricity market, its complexity and scale hinders analytical efforts of market participants, system operators and regulators. They conduct analysis to obtain insights into price formation, to enhance coordination as well as a more efficient use of assets. The limited and fragmented data available on flow-based domain strategies of TSOs and cost structures of participants obstructs the analysis.

This thesis attempts to bridge the gap between stylized academic models on FBMC and real-world market outcomes, to be able to reason about operational day-ahead markets via these simplified models. To this end, a multi-step modelling process is carried out to forecast day-ahead zonal market prices and cross-border flows. Inverse optimisation is utilised to recover cost functions on a bidding zone and technology level. This is followed by a spatial reconstruction of the static grid, which is subsequently used to infer the flow-based domain based on historical observations. The model-based approach has the added benefit of being interpretable, and can be adjusted for structural and regulatory market changes.

Inverse optimisation has proven to be able to recover aggregated technology cost functions with which real-world market outcomes can be forecasted in a tractable way. The model developed in this work is shown to outperform a commercially available, machine-learning-based algorithm in forecasting day-ahead prices. The limitations of reconstructing flow-based domains using publicly available flow-based market data are identified. Analysis concludes that while the ability to recover cross-border flows is sensitive to the shape and size of the inferred domains, the performance of price forecasts is robust against the quality of domain inference. The delivered work is argued to yield valuable insights to both market participants optimising their assets, and regulators structurally assessing the effects of flow-based domain configurations on welfare outcomes of real-world day-ahead markets.

Contents

List of Figures

List of Tables

1 Introduction

1.1. Context

As mitigating the repercussions of climate change is more urgent than ever (IPCC, 2022), a strong focus needs to be kept on reforming emission-intensive industries. In the EU, energy production and consumption account for more than 75% of its greenhouse gas emissions (European Commission, 2018), making the energy sector a prominent player in reducing our carbon footprint. Upscaling renewable energy production has one of the highest prospects of tackling the challenge mentioned above. Another prospect is to better utilise non-carbon emitting base load power plants by enabling more cross-border electricity trade (Finck, 2021). Both approaches pose additional challenges to how the transmission network is managed.

A mechanism is necessary to effectively support cross-border electricity trades that enable transporting from overproducing regions (due to excess renewable generation) to regions with scarcity, while respecting transmission grid constraints. Such a mechanism should also provide market incentives for cheap base load power plants to produce more for exports if there is excess available capacity.

Network congestion management is an important concept to maintain the security of supply on electricity transmission networks. Network congestion occurs when flows scheduled for a given lines exceed their physical line capacities , requiring corrective action (De Vries et al., 2020). Cross-border congestion is an especially important issue, since it requires coordination of system operators of different countries. Recently, these cross-border transfer capacities received increasing attention from regulators and investors, as they play a central role in creating a liberalised internal European electricity market. Besides allowing for an higher overall social welfare, it also supports the integration of renewable energy that benefits from geographical smoothing.

In order to facilitate these cross-border exchanges, the market operators need to have knowledge about the available transfer capacities on the interconnectors. Previously, the system of Available Transfer Capacities (ATC) was employed in the Central Western European (CWE) region. The ATC approach entails Transmission System Operators (TSOs) nominating fixed capacities on their inter-zonal network connections available for commercial trade. They do so by establishing bilateral contracts, above already expected utilisation (Mohammed et al., 2019). The ATC system was found to still result in rather moderate amounts of cross-zonal exchanges, however (Felten et al., 2021). Its fixed and simplified nature (only having a border-level resolution) failed to respond to short-term spatial variability of generation, which is due to a high penetration of renewables. Additionally, it does not account for potential injections in other market zones that further have effect on the considered elements across a given border. When first introduced in 2015 (Tennet, 2015), the flow-based market coupling (FBMC) paradigm aimed to address these shortcomings, replacing ATC. FBMC accounts for the physical flows in parallel network elements, with a sufficient spatial resolution to capture load variability. It is found to allocate considerably more transfer capacities compared to its predecessor (Ovaere et al., 2023). This is due to the experience that a trading domain that is constrained by multiple (but single) line elements is found to be larger under the same operation limits, than using fixed, border-level constraints (Ovaere et al., 2023). This latter aspect is illustrated on Figure (1.1). Another reason that the FBMC approach performs better on the European grid is the highly meshed nature of the network. Accounting for par-

allel physical flows becomes even more important when the network is composed of several parallel elements, rather than more scarce tree-like branching structures. Despite enabling a higher overall social welfare with better resource allocation in the region, FBMC brings about additional complexities and challenges. By employing base cases and Generation Shift Keys, it makes TSOs rely on more forecasts, introducing significant parametric uncertainty (Kristiansen, 2020). By not having the exact algorithm published, it gives market participants considerable model uncertainty, as the complex underlying algorithm makes it difficult to forecast positions and prices in the day-ahead markets. With a lack of forecast on future transmission capacity, trading long-term contracts become inherently more difficult as well (Felten et al., 2021).

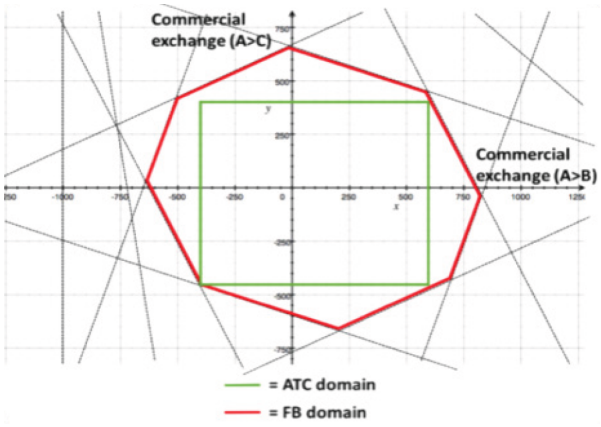

Figure 1.1: An illustration on the difference between an ATC and flow-based domain between bidding zones A and B (sourced from KU Leuven (2015))

The FBMC approach was first implemented in 2015 in the Central Western Europe (CWE) region involving the day-ahead markets of Belgium, the Netherlands, France and Germany/Luxembourg/Austria) (Tennet, 2015). When it was extended to the CORE region on the summer of 2022 (JAO, 2022c), additional uncertainties have been introduced, due to new bidding zones entering the system. The composition of the CORE region now involves 12 bidding zones: Austria, Belgium, Croatia, the Czech Republic, France, Germany, Hungary, Luxembourg, the Netherlands, Poland, Romania, Slovakia and Slovenia (JAO, 2022a). While this extension brings the EU closer to a unified electricity market, an in-depth knowledge on the transmission system and TSO capacity choices is more necessary than ever. Having better informed market participants is important for efficient resource allocation, while a better knowledge on the system might also help avoiding costly re-dispatch after coupling.

1.2. Role of the transmission network in the energy transition

In the previous section, the transmission network was introduced to play an essential role in the energy transition. This relation is further elaborated in this section, as it further frames the relevance of this thesis project in a broader societal context. A high penetration of renewables in the energy system introduce not only temporal intermittency, but also spatial variability. The latter aspect is especially relevant for wind, as while northwestern countries have abundant offshore and onshore wind power to utilise, the Eastern European countries are less rich in wind resources. To address these fluctuations, new flexibility measures could be deployed, ranging from demand-side management to storage technologies. However, these technologies have still low technological and institutional maturity (Chen et al., 2020).

Expanding transmission networks by TSOs to better interconnect regions can also provide increased spatial flexibility. Multiple large-scale transmission projects are already under construction, and several others are planned to be delivered between 2030 and 2050. The development plans have covered 153 billion EUR up until 2020, and projects commissioned for 2030 already represent 123 billion EUR (ENTSO-e, 2021). Chen et al. (2020) and Ovaere et al. (2022) shows the potential of optimal expansion of the grid to let more renewables substitute fossil fuel generation. It can smoothen out the geographical availability of wind to lend countries further inland access to cheap renewable power. An optimal expansion also helps to tackle grid congestion brought about the distributed power injection of rooftop

solar PV. Higher transmission capacities also enhance market coupling, which in the end increases economic efficiency. It does so by allowing more trade from low-cost regions to high-cost regions, leading to price convergence (which results in a higher overall social welfare across the continent) (Ovaere et al., 2022).

These capital-intensive investments have to be guided by well-informed decisions to actually arrive at optimal transmission expansion configurations for the above-mentioned purposes. While estimating the CAPEX of expansion projects is relatively straightforward (accounting for the costs of raw materials, engineering expertise and construction/maintenance labor), assessing their OPEX is getting increasingly challenging (Konstantelos et al., 2017). Assessing operational costs also involve quantifying the benefits that the new transmission capacities induce in the market, influencing the cost recovery of the initial infrastructural investments. The fact that the pan-European day-ahead market clearing algorithm is becoming more complex also increases the complexity of evaluating how the continent-wide social welfare changes with different network arrangements. The cost structures of market participants are not known, therefore the realistic quantification of social welfare as a market outcome is lacking. Cost formation is further determined by the mechanism of flow-based market coupling, as it determines cross-border exchanges and hence price convergence. In conclusion, there is a need for models that quantify the effects of different network configurations in an increasingly complex market setting, informing the optimal allocation of the billions of euros to develop a both environmentally and economically efficient energy infrastructure. This is one of the needs that this thesis subsequently addresses.

1.3. A socio-technical perspective

In order to describe the energy infrastructure at hand, and to qualitatively formulate the impact of this thesis project, the report takes a holistic approach. This holistic approach considers the energy infrastructure as a complex socio-technical system, based on the formulation of Scholten and Künneke (2016). Liberalised energy markets have been progressively considered as socio-technical systems, giving rise to interactions between technical components, energy market dynamics and institutional arrangements (Scholten and Künneke, 2016). This perspective departs from perceiving the electricity network only by its engineered characteristics (generation, network flows, topology and robustness), or only by its economic characteristics (optimising monetary flows to obtain social welfare). Indeed, flow-based market coupling considers both the physical limits of the transmission network and the underlying market structure of power exchanges.

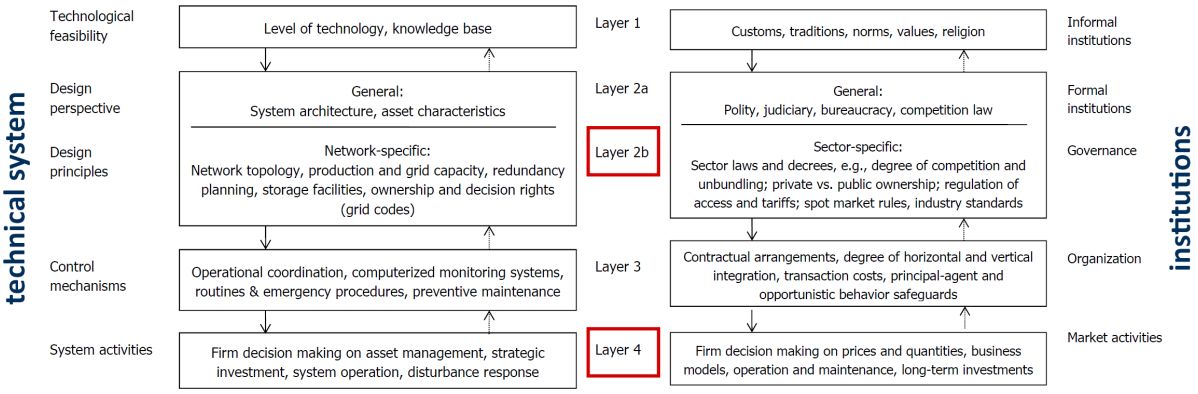

Figure 1.2: Framing the energy infrastructure with the Williamson scheme (Williamson, 2000), adapted from Scholten and Künneke (2016)

.

Describing socio-technical systems often starts with positioning the relevant activities and transactions within the well-known 4-layer Williamson scheme (Williamson, 2000). Transmission capacity management is a critical technical function that lies in the focus of this thesis. Allocating transmission network capacities poses as a critical transaction (Künneke et al., 2010) carried out by system and market operators. As this critical transaction concerns technical design principles and governance, regulation of access and market rules, it resides on the 2nd level of the Williamson scheme (Scholten and Künneke, 2016). This 2nd level comprises the institutional environment: setting the formal rules of the game, with special emphasis on property rights (Williamson, 2000). This thesis is also concerned

with delivering market forecasts, which affects market participants and their decision making on prices and quantities. This latter aspects resides in the 4th level of the Williamson scheme, which involves resource allocation: getting the marginal conditions right (Williamson, 2000). The two identified layers that are relevant in this thesis are visualised on Figure 1.2.

The coordination of activities in both dimensions is crucial for the energy system to perform in accordance with the 'triple A' goals: availability, affordability and acceptability (De Vries et al., 2020). The technical performance is described in terms of reliability and robustness, while the socio-economic performance lies in the efficient and effective resource allocation (Scholten and Künneke, 2016). One will see in the upcoming chapters, that a performance trade-off exist between each of the two dimensions.

The regulatory responsibilities in layer 2 evolve around the efficient allocation of property rights. This is especially challenging in network industries, such as the energy industry. In a networked context, property rights are overshadowed by negative network externalities. Congestion is such a negative externality, which can be resolved by allocating a network access as a scarce resource between alternative usage scenarios (Glachant, 2012). This latter process is called congestion management. Regulators play an important part in facilitating a transparent and predictable congestion management. Having sufficient information at hand enables provisional simulations and analyses to eliminate future structural congestion on the long term, and to also coordinate reactions to congestion in the short-term (Glachant, 2012). At this point, one can identify another important performance criteria, that spans both dimensions: transparency. In an increasingly complex, unbundled and networked socio-technical environment, sufficiently keeping information and knowledge as public good is essential for an efficient coordination.

This section concludes with a remark on what a typical research problem within the *Complex Systems Engineering and Management* (CoSEM) curriculum entails. As stipulated in Correljé and Künneke (2021), CoSEM is concerned with designing an intervention that is aimed to modify or create some processes within the socio-technical system. Rather than designing an intervention, this thesis explores a quantitative approach, to attempt restoring transparency and tractability in an increasingly complex, networked system. It does so to aid both market participants to optimise their resource allocation, and regulators to explore new interventions within the institutional sphere of energy infrastructures. The rationale behind these holistic objectives are identified in the next section, taking the context of flow-based market coupling.

1.4. Knowledge gap identification

In the flow-based approach, the transmission network is represented by a DC load flow model (Van den Bergh and Delarue, 2014). In the FBMC, one is interested in how power injection at a given node affects the flow between all nodes via the established lines. The first important concept that needs to be established here, is the nodal Power Transfer Distribution Factor (PTDF) matrix. Given a network $\mathcal{G}$ with a set of nodes $\mathcal{N}$ and a set of lines $\mathcal{L}$, $PTDF^{N}_{l,n}$ describes the fraction of injected power at node n that flows through line l. This matrix can be determined purely based on the physical properties of the network (by knowing the transmission lines between the nodes, their physical capacity, susceptance, etc.).

However, during market coupling, only zonal information is available, therefore these nodal properties have no direct use. To capture how zonal injection (and hence zonal net positions) affect flows on the lines, zonal PTDFs are introduced. Upon importing power into a zone, $PTDF^{Z}_{l,z}$ describes the fraction of this power that flows through line l (Schonheit et al., 2021). In order to represent this zonal knowledge, a mapping needs to be found between nodal and zonal PTDFs. This mapping is established by Generation Shift Keys (GSKs). It is also represented by a matrix, that maps $\mathcal{N}$ to the set of zones $\mathcal{Z}$. $GSK_{n,z}$ shows how an injection at node n contributes to the change in the net position of zone z. GSKs pose a major uncertainty factor in the FBMC model, as they naturally depend on the dispatch of the system, which in turn depends on the market outcome. Several strategies exists for calculating these values (Finck et al., 2018). Only a handful of TSOs publish their approach (JAO et al., 2021), but this calculation method is generally unknown for most bidding zones.

To circumvent the circular dependency of the flow-based domains and the market clearing described above, the TSOs have introduced the following approach (Schonheit et al., 2021). Initially, TSOs model base cases, where they attempt to forecast how the market will clear D-2 (meaning a two-day-ahead forecast). They use it to provide reference net positions and reference flows to the market operators,

but it can also be used in certain strategies for determining GSKs. TSOs also need to make available zonal PTDFs for the lines managed by them, for which they need to determine the GSKs (assuming they have a knowledge over the physical network). Upon having the data described above, the flow-based constraints can be constructed for the market clearing algorithm, which generally has the form shown in (1.1).

$$\sum_{z} \left[PTDF_{l,z}^{Z} \cdot NP_z \right] \leq RAM_l \qquad \forall l \in \mathcal{L} \tag{1.1}$$

On the left side, net position changes are translated into changes in line-specific flows via the zonal PTDFs. These flow changes are restricted by the Remaining Available Margin (RAM), which can be calculated for each considered line from its physical capacity, the reference flow assumed to be present there, and from the so-called Flow Reliability Margins (FRM). The latter essentially represents the uncertainty faced by the TSOs. Constraint (1.1) is then fed into the market clearing algorithm, which in essence, can be represented by an optimisation of social welfare across the bidding zones with regards to the acceptance of bids and offers, subject to supply-demand and capacity constraints. After the market has cleared, additional congestion relieving measures can take place to mitigate any remaining congestion, but that is not the focus of this thesis.

To aid market participants, the market operator requires TSOs to publish key properties that are being used for the FBMC algorithm (such as the zonal PTDFs for CNEs, RAMs, reference flows) (ACER, 2019) following the market clearing, which helps market and transmission network analysis in hindsight. However, since these are published *ex post*, it still hinders the attempts of market participants to forecast prices, transfer capacities and cross-border flows. Having better informed market participants is a key for more efficient resource allocation, and hence, an increased social welfare (Finck, 2021).

It is worth mentioning that several price forecasting approaches have been developed, that are used by market participants to improve their risk management and strategic choices. These mainly involve econometric/statistical (e.g. ARIMA in Zhao et al. (2017)), and machine learning (ML) (e.g. Support Vector Machine in Shiri et al. (2015)) techniques. These models are preferred due to their ability to capture short-term price volatility. Machine learning models are also dubbed to be able to reveal strategic behavioural aspects of market participants. Even though these techniques have proven to have good predictive performance, they have several drawbacks. Both econometric and ML approaches are black-box models, lacking interpretability, or transparency in other words. Black-box models tend to have difficulties to capture market dynamics, structural market changes and regulatory interventions (de Marcos et al., 2019). This aspect cannot be neglected in today's continuously evolving energy markets that try to accommodate the net-zero goals of the EU. Furthermore, black-box models tend to explicitly depend on historical data. Even though the market structure was not changed when expanding the flow-based domain to the CORE region, past data on market outcomes turned to be inappropriate, as the new zones brought in new participants and additional infrastructure, changing the landscape of market optimisation. Due to the limitations of black-box models mentioned above, this thesis focuses on a model-based forecasting method, explicitly employing an FBMC market-clearing algorithm. Model-based forecasting approaches are referred to as fundamental methods in the literature (de Marcos et al., 2019).

In order to construct a fundamental forecast model, one might attempt to approximate how the market will clear based on well-established forecasts of exogenous factors, such as the weather, and electricity demand of consumers. Such a venture would generally require to characterize the transmission network, the way how TSOs calculate zonal PTDFs through GSKs, and one would also need an appropriate approximation of the clearing algorithm in place. The lines of the transmission networks and their physical properties are published with a joint effort of TSOs, with network nodes representing the substations. However, the physical location of power plants are generally unknown, so as the information on which network node they actually inject to. Some TSOs publish their method of calculating the Generation Shift Keys (JAO et al., 2021), but for most of them it is still unknown. Some of the known methods require running base cases, for which the methodology is never published. The manifestation of the large-scale market algorithm that also includes the flow-based approach is called EUPHEMIA (Nemo, 2020). Although a large part of the algorithmic choices is documented, there is no public implementation of the algorithm.

Given the multi-layered uncertainty around the FBMC, several research articles have been dedicated to study parts of the approach from different perspectives. Attempts have been made to fully

characterize the EUPHEMIA algorithm (Lam et al., 2018). However, this model does not include how the TSOs determine the network constraints, reference net positions and flows with respect to a base case. The model presented in (Schonheit et al., 2021) serves as a good academic starting point for assessing the whole FBMC pipeline starting from the base case until the intra-day congestion management. However, it presents the FBMC pipeline only on stylized cases, and while enumerating a variety of GSK approaches, it remains uncertain about which ones are actually being used in real-life. A sophisticated data completion framework has been introduced by (Puiu and Hauser, 2021), which approximates zonal PTDF and RAM values from data published on market outcomes. It also uses the actual CWE grid reconstructed by Matke et al. (2016). Although the paper shows results of a good reconstruction power, their formulation is lacking the ability to forecast future constraint values based on the available exogenous factors.

In conclusion, in order to support a better informed decision-making of market participants, there is a need for characterizing the flow-based domain in the newly formed CORE region. An interpretable forecast model is necessary that can continue to be useful in a continuously changing market landscape. Such a forecast model will not only aid market participants, but also regulators to analyse scenarios with changing infrastructural and market landscapes, promoting interventions. To achieve this objective, one faces a knowledge gap for two reasons. On the one hand, a combined formulation is lacking to both realistically approximate the flow-based constraints with regards to exogenous factors and reference cases of the domain, and to fit these constraints to a clearing algorithm to yield a forecast of the market results. On the other hand, all the previous studies that were dealing with the real-world have been implemented for the CWE region (or for other parts of the EU, mainly Italy), the newly formed CORE region has not been examined and characterised yet. Based on this knowledge gap, the main research question of this thesis is formulated as follows:

> **MQ**: *How can the flow-based electricity market domain be replicated in the newly formed CORE region to forecast day-ahead market clearing results?*

In order to replicate the flow-based domain, a proxy model of the actual system needs to be considered, that adequately captures the way the electricity market responds to exogenous factors. Subsequently, the parameters of such a model need to be fine-tuned so that its produced results best resemble the real-world observable outcomes. To this end, a data-driven *modelling approach* is proposed. The approach inherently has a modelling nature, as multiple real-world socio-technical artefacts need to be represented by their simplified models: the transmission grid, the TSOs decision-making process on available cross-border transfer capacities, and the day-ahead market clearing algorithm. This thesis aims to explore available models and approaches in the literature for the artefacts mentioned above and subsequently adjust and combine them to formulate an answer for the main research question. The approach is also heavily data-driven, as historical records of exogenous parameters and observed market outcomes will be gathered and used to fit the parameters of the otherwise stylized models, to produce outcomes close to the observed ones. While market outcome forecasting brings a strong *quantitative* nature to the thesis, the model-based formulation will also bring about a *descriptive* and *explanatory* research perspective. It does so by attempting to infer how system operators account for their zonal injection patterns, and how prices form within an increasingly interconnected electricity market.

After the identification of the research approach to be followed, sub-questions are formulated accordingly to aid answering the main research question. These sub-questions structure the modelling process into steps, and also provide further clarification on the nature of forecasts the project aims to achieve:

> **SQ1**: *What is a static electricity grid representation of the CORE region that can be used to model the flow-based domain?*
>
> **SQ2**: *How can the generation bid curves of market participants be approximated to have the resulting generation levels resemble the real-world positions after day-ahead market clearing?*
>
> **SQ3**: *How can the potential strategies of system operators behind choosing their flow-based feasibility domain be reconstructed?*
>
> **SQ4**: *To what extent is the established model able to predict day-ahead zonal prices and cross-border exchanges within the region, compared to alternative data-driven techniques?*

1.5. An unprecedented period

Flow-based market coupling in the CORE region went live in the summer of 2022. Since the CORE region lies in the centre of this thesis, observed market data is only available from this point, up until the modelling stage of the thesis. The resulting time period examined is hence between September 2022 and February 2023. This section aims to emphasize that this chosen period is an unprecedented one. A post-COVID recession looms over the economy, which affects the demand side of energy markets (Ingham, 2023). The Ukrainian war has further brought about an energy crisis, where European countries are increasingly trying to reduce their dependence on Russian supply of natural gas (Galetto, 2022). The peculiarity of this period, as relevant for this thesis, is shown in Figure 1.3, displaying day-ahead electricity prices for the Netherlands, between September 2022 and March 2023. As one can see, the price levels are being consolidated as approaching February, albeit still high when compared to historical values. Between September and December, prices not only go above 500 €/MWh, but also subject to extreme volatility.

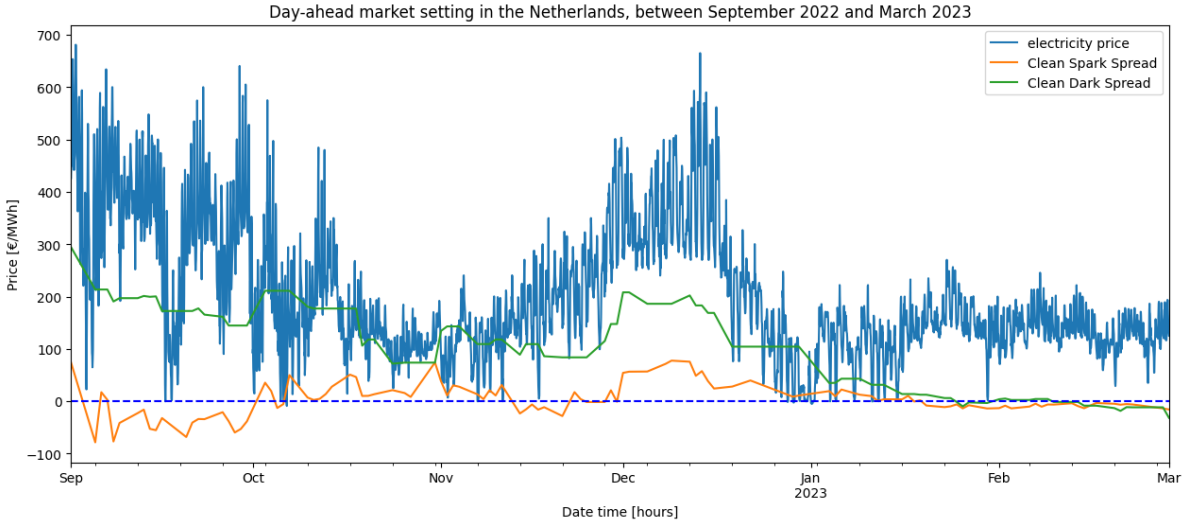

Figure 1.3: The day-ahead market setting in the Netherlands between September 2022 and February 2023, displaying the clearing price (ENTSO-E, 2023), Clean Spark and Dark Spreads (retrieved from Refinitv Eikon)

In Figure 1.3, the *Clean Spark Spread* (CSS) and *Clean Dark Spread* (CDS) are also presented. The Clean Spark Spread represents the theoretical gross margin of a gas-fired power plant after selling 1 MWh of electricity, preceded by buying the natural gas to produce this 1 MWh of electricity (ICE, 2023). The Clean Dark Spread means the same respectively, but defined for coal-fired plants, having bought coal as fuel. The ”clean” part comes from the inclusion of carbon credits in the margin, to cover emission quotas for power generation (Dechezleprêtre et al., 2023). The first thing one could notice is how the spread between the CSS and CDS correlates with the volatility of prices. A larger spread implies more volatile clearing prices. When the spread between CSS and CDS is large, the price gap between gas-fired and coal-fired generation is also wider. In the Netherlands, both gas-fired and coal-fired plants are dominant price setters. When the prices setter technology changes over time mainly due to fluctuations in electricity demand, the clearing price jumps will be wider as well.

Figure 1.4 illustrate that these margins were historically high in this period, implying that the extremely high electricity prices could not have been explained only with high fuel prices, but other exogenous or behavioural market mechanisms must have played a role as well, induced by the crisis. (If high electricity prices had been solely caused by high fuel prices, the profit margins would have stayed closer to zero.) As these further aspects go beyond the scope of the fundamental modelling exercise of this thesis, the model is expected to perform more modestly in the first 4 months. A better performance is expected from January-February onwards in forecasting day-ahead electricity prices.

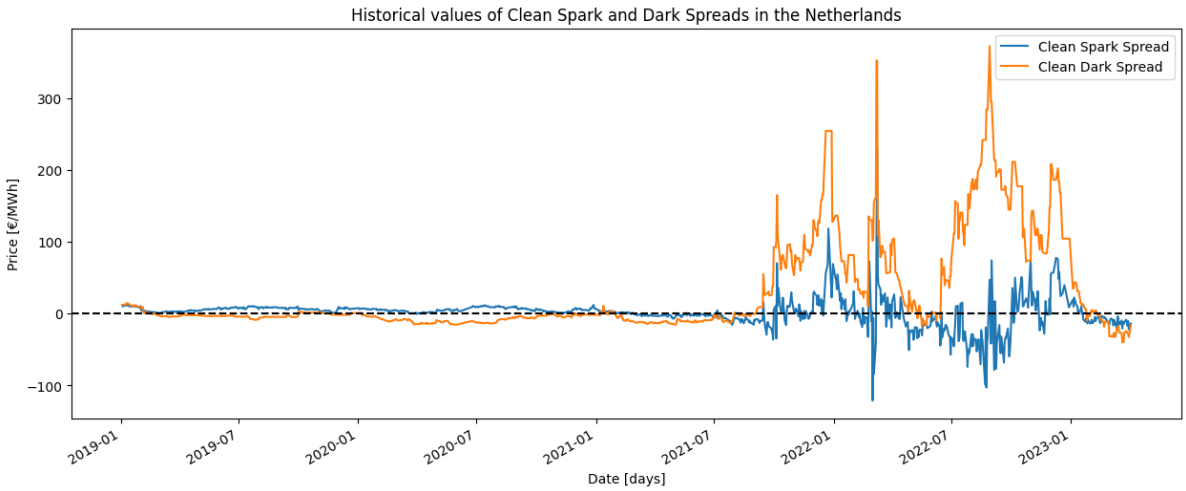

Figure 1.4: Historical values of Clean Spark and Dark Spreads in the Netherlands, between January 2019 and February 2023 (retrieved from Refinitv Eikon)

1.6. Thesis outline

Governed by the research questions, the structure of the thesis is outlined in Figure 1.5. First, literature and theory are reviewed for mathematical optimisation, as well as for elaborating on the concepts of flow-based market coupling. Subsequently, the methodology is discussed for obtaining a static grid model, reconstructing cost curves with inverse optimisation, and reconstructing the flow-based domain by modelling a base case and inferring the GSK strategies used within each zone. The thesis proceeds with an *Analysis* chapter. It first quantitatively compares the forecast results delivered by the cost curve reconstruction and the flow-based domain inference separately, to available observations. Then it follows with the model assembly, where both the inferred flow-based domain calculations and the reconstructed cost curves are used in the forward clearing model. The resulting day-ahead price and cross-border flow forecasts from this assembly then are compared to commercial forecasts via error metrics. To further understand how the reconstructed and inferred pieces make one interpretable forecast model, the information flow of the assembly is depicted on Figure 1.6. The analysis is followed by a *Discussion* chapter on the applicability, strengths and caveats of the developed model. The thesis concludes with a *Conclusions* chapter, summarizing the work and formally answering the research questions.

1.7. Contributions

This chapter introduced the need for having a structurally interpretable day-ahead market clearing model that captures the behaviour of real-world markets. After discussing the main research question and its sub-questions, the chapter concludes with highlighting the main contributions of this thesis. First of all, it conceives an inverse formulation of the flow-based day-ahead market clearing model that can be used to infer aggregated technology cost curves that approximate offers of market participants. This inverse optimisation problem is implemented for the newly formed CORE region. The thesis further constructs an optimisation that attempts to reveal which GSK strategies are most likely being used by (TSOs), from a given set of academically identified strategies. An attempt is made to model the base case calculations of TSOs with a nodal pricing approach. Electricity demand is only reported on a zonal level, therefore deducing nodal demand values for the nodal approach is not straight-forward. The thesis explores whether distributing the zonal demand over the nodes based on the GDP and population size reported for the node brings the results closer to the actual reference calculations of the TSOs. Since public market data is available on reference flows of a subset of (critical) network elements, the performance (and hence the relevance) of the proposed model can be verified. During the modelling process, the quality and the availability of public market data is assessed, as the obstacles occurred during the reconstruction process, caused by the (un-)availability of data are subsequently identified.

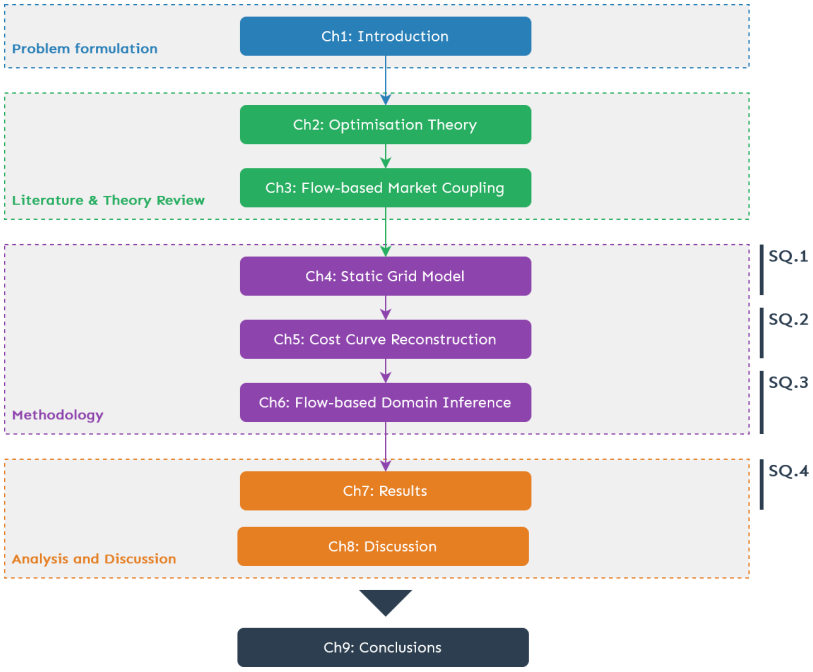

Figure 1.5: Outline of the thesis

As far as the academic relevance is concerned, this thesis contributes to the understanding of the complex flow-based market clearing model, when it is viewed through the lens of interpretable mathematical programs. A missing connection is attempted to be made between the real-world energy markets and otherwise stylized academic models. By directly comparing the established model to real-world observations, the benefits and limitations of the approach are identified.

Finally, concluding with the holistic view, this thesis attempts to restore transparency and tractability in an increasingly complex, networked system. This work aids both market participants to optimise their resource allocation, and regulators to explore new interventions within the institutional sphere of energy infrastructures.

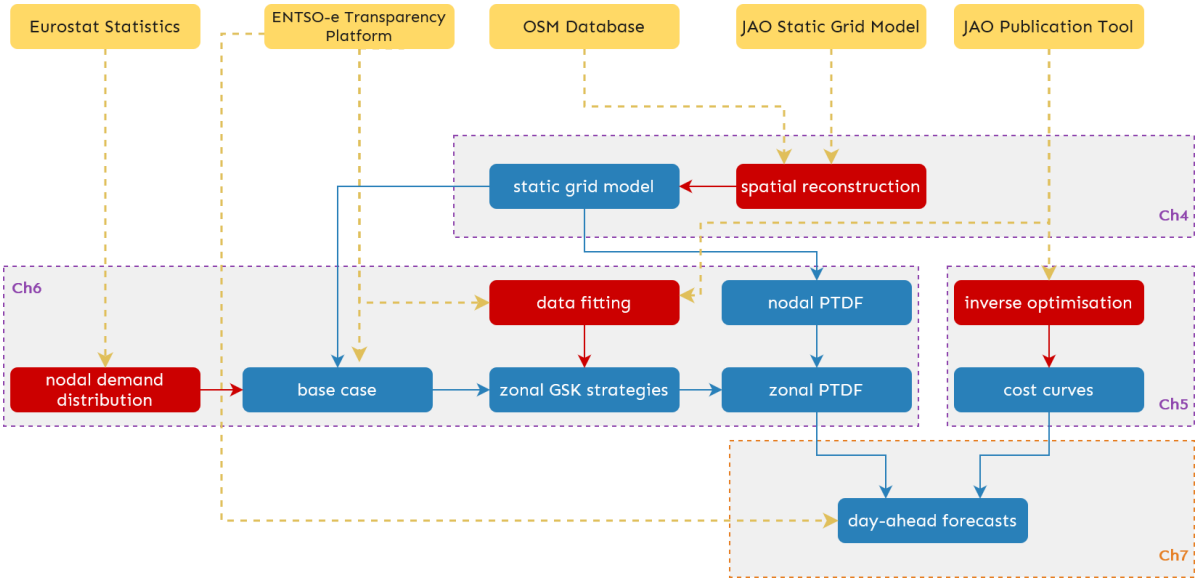

Figure 1.6: The complete model assembly. The red boxes represent the modelling processes carried out in this thesis, delivering the artefacts indicated by the blue boxes. The dependencies between these artefacts are further indicated, leading up to the day-ahead forecasts. The containers show the chapters in which these components are discussed. The yellow boxes represent the data sources that are used by the modelling processes and the artefacts themselves

2

Optimisation Theory

The broad field of mathematical optimisation is concerned with *optimising* an objective conceived to quantify the measure of fitness of a decision among a set of alternative solutions. More specifically, the term optimising can mean *maximising* or *minimizing* an objective, where the latter can be a function in a mathematical sense. In order to obtain a solution, a set of *decision variables* are considered, which affect the quality of the decision (quantified by the *objective function*). The set of alternative solutions can be contained within a *feasibility domain*, which constrain the values that decision variables can take in order to keep the decision feasible with regards to some criteria. Such a domain is defined with *equality* and *inequality constraints* expressed for the decision variables. The aspects described above can be summarized in the following mathematical terms (assuming minimisation):

$$\min f(x) \tag{2.1a}$$

subject to

$$g_j(x) \leq 0 \tag{2.1b}$$
$$h_l(x) = 0 \tag{2.1c}$$
$$\underline{x_i} \leq x_i \leq \overline{x_i} \quad \forall i \tag{2.1d}$$

One can suspect from the highly generalised description above that the field of optimisation is quite broad, encompassing several practical fields where developing optimisation models lies at the centre: decision-making in operations research, design optimisation in mechanics, infrastructure planning in civil engineering, resource allocation in economics and finance, model predictive control in control engineering, among many others. In all of these fields, the practice of optimisation starts with building a model that adequately captures the real-world system, and which can in turn be optimised with regards to a set of objectives.

The problem at hand in this thesis is confined to a deterministic, continuous, convex single objective formulation. Constraints will always be rendered as linear, whereas some objective functions will further contain quadratic elements. The convexity criterion brings an important feature to the formulation, which also enables the inverse formulation of the original problem in the first place. Convex problems always have a single global optimum, therefore one does not need to be concerned about convergence, local optima or noisy functions. Solving a convex problem also does not require a starting point or derivatives of the objective function.

The rest of the chapter proceeds with discussing linear and quadratic programs, touches upon the nature of duality, then introduces the notion of inverse optimisation and its different approaches. The theoretical discussion is followed by an intermezzo on comparing the field of machine learning and regression models to mathematical optimisation. The chapter concludes with discussing available numerical solvers and presenting the solver choice made in this thesis.

2.1. Linear programming

Linear programming (LP) involves a linear objective function with a set of linear constraints, that can be generally written in a matrix formulation as:

$$\min \mathbf{c}^\mathsf{T}\mathbf{x} \tag{2.2a}$$

subject to

$$\mathbf{A}\mathbf{x} = \mathbf{b_1} \tag{2.2b}$$
$$\mathbf{B}\mathbf{x} \leq \mathbf{b_2} \tag{2.2c}$$
$$\tag{2.2d}$$

where $\mathbf{x}$ is the vector of decision variables. The objective function simplifies to a weighted sum of the decision variables, weighted by their respective 'cost'. The naming convention comes from the common application of minimising the cost of decisions, hence $\mathbf{c}$ is called the *cost vector*. $\mathbf{A}$ and $\mathbf{B}$ are coefficient matrices for the equality and inequality constraints depicting which variables participate in a given constraint, $\mathbf{b_1}$ and $\mathbf{b_2}$ vectors are the constant components in the constraints. All LP problems are convex optimisation problems. LP formulations mainly involve efficient resource allocation problems, finding out which tasks should be assigned to which employee, which products should be produced, which stocks should be kept in a portfolio or in which order carry out shipments, to mention a few.

2.2. Quadratic programming

Quadratic programs (QP) also have linear constraints, but involve a quadratic objective function, taking the following general form:

$$\min \frac{1}{2}\mathbf{x}^\mathsf{T}\mathbf{Q}\mathbf{x} + \mathbf{f}^\mathsf{T}\mathbf{x} \tag{2.3a}$$

subject to

$$\mathbf{A}\mathbf{x} = \mathbf{b_1} \tag{2.3b}$$
$$\mathbf{B}\mathbf{x} \leq \mathbf{b_2} \tag{2.3c}$$
$$\tag{2.3d}$$

A QP problem is only convex if the $\mathbf{Q}$ matrix is positive semidefinite. A convex quadratic problem is a special case of nonlinear programs that is more convenient to solve. This is due to the fact that they still have only one feasible region, surrounded by flat surfaces (due to the linear constraints). But while the solutions of linear programs are found on the intersections of the bounding surfaces, solutions of QP problems may be found anywhere inside the feasible region or on its bounding surfaces. The most reliable and fast way to solve QP problems is by extended version of LP solvers, which give better results than general NLP solvers. QP problems are often formulated for data fitting purposes.

2.3. Duality

Before proceeding with the concept of inverse optimisation, it is important to introduce the notion of *duality*. For every linear program (hereon referred to as *primal* problem), one can find a dual formulation in such a way that if one is of a minimisation kind, then the other one is a maximisation (and vice versa), and the optimal values of the corresponding objective functions are equal. The dual pair of (2.2) is formulated below:

$$\max \mathbf{b}_1^\mathsf{T}\boldsymbol{\lambda} + \mathbf{b}_2^\mathsf{T}\boldsymbol{\mu} \tag{2.4a}$$

subject to

$$\mathbf{A}^\mathsf{T}\boldsymbol{\lambda} + \mathbf{B}^\mathsf{T}\boldsymbol{\mu} \geq \mathbf{c} \tag{2.4b}$$

where $\boldsymbol{\lambda}$ is said to be the dual variable vector of the equality constraints, and $\boldsymbol{\mu}$ is the dual variable vector of the inequality constraints. These dual variables can be interpreted as prices associated with the constraints in the primal problem, hence they are often referred to as *shadow prices*. They represent the unit change in the optimal value of the objective function if the right-hand side of the constraint that the dual variable refers to is changed marginally.

This economic meaning of shadow prices directly correspond to the outcome of auctions on the day-ahead electricity market. As one will see in the next chapter, the market clearing price is the marginal cost of the most expensive generator that still needs to produce to cover the demand. In other words, if an additional unit of power needs to be produced (upon increasing demand, the right hand side of the equality constraint), it will need to be produced by this marginal power plant, having its marginal cost as the unit increase in the objective function value. With that said, the market clearing price directly corresponds to the $\boldsymbol{\lambda}$ dual variable of the equality (supply-demand) constraint.

The primal-dual pair has another important trait that is captured by the *Strong Duality Theorem*, and that will be important for formulating an inverse problem. The Strong Duality Theorem states the following:

> *If either of the problems (2.2) or (2.4) has a finite optimal solution, so does the other, and the corresponding values of the objective functions are equal. If either problem has an unbounded objective, the other problem has no feasible solution.* (Bradley et al., 1977)

2.4. Inverse optimisation

The problems having discussed up until now will be further referred to as *forward problems*, producing an optimal solution with regards to the decision variables and model parameters (cost function and constraints). In inverse optimisation (IO), one aims to impute the parameters of the cost function and/or constraints of the problem, in order to render a solution (of decision variables) optimal. In order words, if one has an observed solution of the optimisation problem, one seeks to infer the otherwise unknown parameters of the optimisation model itself.

The formulation will be discussed here on linear programs, but one will see in Chapter 5 that the same approach can be applied on convex quadratic programs as well. Furthermore, as mentioned above, it is possible to aim for reconstructing the cost function as well as the constraint parameters. Given the focus of this thesis, only cost function inference will be considered subsequently.

Let us consider the general LP problem (2.2) again. Given that $\mathbf{c}$ is also assumed to be unknown here, one needs to find a way to include the cost vector in the optimisation problem as a bounded decision variable. This can be achieved by considering the dual counterpart of the original problem (2.4), and enforcing the Strong Duality Theorem. The latter means that the two objective functions are enforced to be equal.

With this primal-dual formulation, two new set of decision variables enter the picture, which are the $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ vectors, the dual variables of the equality and inequality constraints respectively. The $\mathbf{A}$ and $\mathbf{B}$ matrices as well as $\mathbf{b_1}$ and $\mathbf{b_2}$ vectors remain as exogenous parameters. The original decision variable vector $\mathbf{x}$ turns to be a parameter as well, as it is exogenously observed in this case, subsequently denoted as $\mathbf{x_0}$. Given the above described setting, the classic inverse optimisation problem can be formulated as presented in (2.5). It will subsequently abbreviated as $\mathbf{IO}_0$.

$$\min_{\mathbf{c},\boldsymbol{\lambda},\boldsymbol{\mu}} \; 0 \tag{2.5a}$$

subject to

$$\mathbf{A^T}\boldsymbol{\lambda} + \mathbf{B^T}\boldsymbol{\mu} \geq \mathbf{c} \tag{2.5b}$$

$$\mathbf{c^T x_0} = \mathbf{b_1^T}\boldsymbol{\lambda} + \mathbf{b_2^T}\boldsymbol{\mu} \tag{2.5c}$$

$$\tag{2.5d}$$

One can recognise the dual feasibility constraints in (2.5b). The reason the primal constraints are omitted here is because $\mathbf{x_0}$ is already assumed to be optimal, satisfying its primal constraints. The strong duality is enforced in (2.5c). What changes here is that on the left-hand side now $\mathbf{c}$ is the variable and $\mathbf{x_0}$ is the parameter, keeping the problem a linear one. The classic formulation only aims

to find a cost vector that makes the observed decision vector an optimal one, assuming that there exists such a cost vector. This leaves us without any particular objective function to minimise.

2.4.1. L-norm formulation

The classic formulation makes the strong assumption that $\mathbf{x_0}$ belongs the optimal set of the original problem to be reconstructed. In realistic applications, it is rarely the case, however. For example, the observation might come with a measurement noise, which could render a solution infeasible to a problem which is highly sensitive to perturbation. Alternatively, the model whose parameters are being inferred may be a simpler and lower dimensional proxy of the real-world model being investigated. This could also render the inverse problem infeasible, as the observation coming from the real-world model might not posit as an optimal solution of the structurally simplified model.

The above-mentioned situation naturally brings forward the need to find approximately optimal solutions to the inverse problem, thus finding a cost vector that renders the observed solution as optimal as possible with regards to the forward problem. An intuitive way to explore near optimal solutions is to introduce an error term in the optimality conditions that otherwise need to be met in order to arrive at an optimal solution.

An error term can be introduced with regards to the observed values directly, by bringing back $\mathbf{x}$ as an internal decision variable next to the observed $\mathbf{x_0}$ values. The objective is then to minimise the distance between the observed and internally resulting $\mathbf{x}$ values. This approach is hereby referred to as the L-norm formulation (abbreviated as $\mathbf{IO}_L$), as the formerly mentioned distance can be measured with an arbitrary choice of norm. The L1-norm corresponds to the absolute-value norm, while the L2-norm is the Euclidean distance (with quadratically measured differences). The L-norm formulation of the inverse problem can be written as follows.

$$\min_{\mathbf{c},\mathbf{x},\boldsymbol{\lambda},\boldsymbol{\mu}} ||\mathbf{x} - \mathbf{x_0}||_L \tag{2.6a}$$

subject to

$$\mathbf{A^T}\boldsymbol{\lambda} + \mathbf{B^T}\boldsymbol{\mu} \geq \mathbf{c} \tag{2.6b}$$

$$\mathbf{c^T x} = \mathbf{b_1^T}\boldsymbol{\lambda} + \mathbf{b_2^T}\boldsymbol{\mu} \tag{2.6c}$$

$$\tag{2.6d}$$

If the resulting optimal value of the objective function turns out to be zero, one can know that the observed solution was actually part of the optimal set. It is also visible that the number of decision variables grew by reintroducing $\mathbf{x}$. What further makes this problem harder to solve is the multiplication between $\mathbf{c}$ and $\mathbf{x}$, creating a *bilinear term*. The inverse problem stops being a linear one, although it can still be solved as a convex quadratic problem.

2.4.2. Duality gap minimization

An error term can also be introduced in terms of the objective function value. The cost vector $\mathbf{c}$ can be found such that it minimises the mismatch between the objective values produced by $\mathbf{x_0}$ and an optimal solution found for the forward model. It relaxes the Strong Duality Theorem and it focuses on minimising the duality gap between the primal and dual formulations. Hence this approach will subsequently be called *duality gap minimisation* abbreviated as $\mathbf{IO}_d$ and will take the mathematical form presented in (2.7). One can already notice that the attractive feature of duality gap minimisation compared to $\mathbf{IO}_L$ is that it remains linear.

$$\min_{\mathbf{c},\varepsilon,\boldsymbol{\lambda},\boldsymbol{\mu}} \varepsilon_d \tag{2.7a}$$

subject to

$$\mathbf{A^T}\boldsymbol{\lambda} + \mathbf{B^T}\boldsymbol{\mu} \geq \mathbf{c} \tag{2.7b}$$

$$\mathbf{c^T x_0} = \mathbf{b_1^T}\boldsymbol{\lambda} + \mathbf{b_2^T}\boldsymbol{\mu} + \varepsilon_d \tag{2.7c}$$

(2.7d)

Besides computational considerations, the choice of IO approach can also be dictated by the application itself. If the focus is on producing a solution of $\mathbf{x}$ that is optimal for the recovered $\mathbf{c}$ and close to the observed $\mathbf{x_0}$, a norm-based model seems to fit the purpose. If the goal is rather to recover the objective function value produced by x_0, duality gap minimisation might come as a more natural choice.

2.5. Perspective with machine learning and regression models

Since the optimisation approach employed in this thesis for forecast is often compared to machine learning and regression forecast methodologies, it is important to spend some time on explaining how these terms relate to each other. This delineation with regards to the energy system itself is left for the subsequent chapters, here the mathematical connection is explained.

A popular form of regression analysis is the least-square regression used for data fitting. It attempts to fit the coefficients of a model function to an available set of observations from the real-world system. It does so by minimising the sum of the squares of the residuals, where the residual is the difference between the observation and the value provided by the model function. One can perceive this as a special case of a quadratic problem, being an unconstrained one in the general sense (Martins and Ning, 2021).

Optimisation also lies at the core of machine learning (ML) problems. Solving problems with machine learning starts with transforming the input data to an adequate format for the input layer of a neural network (Sun et al., 2019). Then one needs to choose an appropriate family of neural network models. The model is trained by optimising the parameters of the model (weights of the network) with regards to a loss function (the objective in this case). Research in the field of neural networks subsequently moved from the traditional back-propagation method towards experimenting with unconstrained nonlinear programming techniques (Piccialli and Sciandrone, 2018).

In conclusion, both machine learning and regression techniques involve optimisation processes, therefore they can be discussed under the umbrella of the broader field of mathematical optimisation. While in general linear or nonlinear programs produce the most optimal decision within a space of feasible decisions (solutions) of a model, the output of ML and regression processes is a fitted model to be recurrently used for inference, classification or forecast purposes (Martins and Ning, 2021). In the latter case, optimisation is the tool to train the parameters of these models.

3

Flow-based Market Coupling

In order to position the concept of flow-based market coupling (FBMC), let us take a step back, and consider the electricity system within Europe (Figure 3.1). The following modular structure is the result of market liberalisation efforts carried out in the 90s, and is a perfect example of a highly intertwined socio-technical system. The physical layer comprises power plants, either owned by private parties of backed by governments, the high-voltage transmission network that transmits electricity across regions, the low-voltage distribution network that distributes electricity within a region/neighbourhood, and that connects consumer devices to the grid. On the institutional level one can find the electricity producer companies, the wholesale markets where electricity is maintained as a liquid good, and the retail market which connects small consumers to the system on a contractual basis (De Vries et al., 2020). Transmission System Operators (TSOs) provide security of electricity supply by maintaining the electricity network and by collaborating with wholesale markets, all while being an independent party.

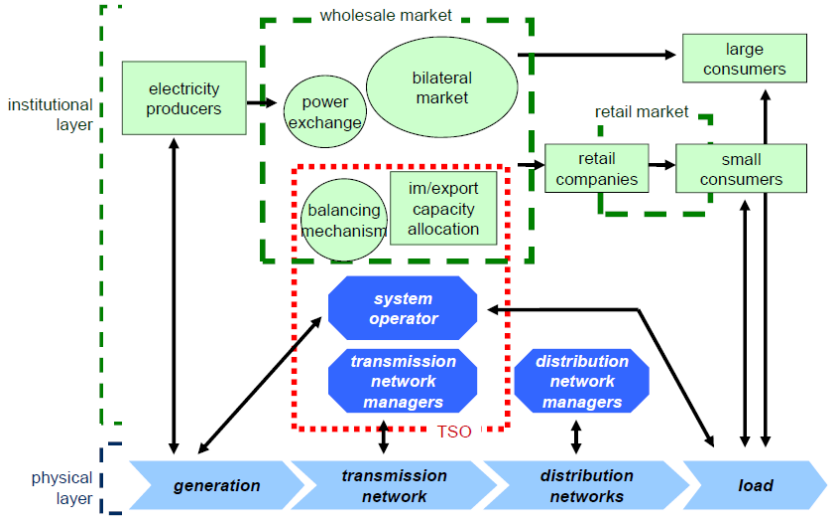

Figure 3.1: The composition of the liberalised European electricity system (taken from De Vries et al. (2020)).

In order to understand FBMC, one needs to focus on the role of the TSOs and power exchanges. Power exchanges establish short-term markets, where electricity is traded on a daily basis, always the day before delivery (hence the term day-ahead). Participants of the exchange are required to submit orders for buying and/or selling electricity, for every hour of a given day of delivery. All orders are collected within an order book, based on which the market operator determines a single market price (called the clearing price) for each hour of the day. The so-called generation offers are then put into an ascending order based on the bid prices, while the consumption bids are put into a descending order. The intersection of the the bid-offer curves will result in the market clearing price for the given hour (Chatzigiannis et al., 2016). The process is illustrated on Figure 3.2.

The choice of market clearing price is based on the tenets of welfare economics. Welfare economics stipulates that the total social welfare is maximised, and hence allocative efficiency is achieved, when the market price both equals the marginal cost (of generation), and marginal benefit (of consumption) (Cannan and Pigou, 1921). Ordering the generation offers and consumption bids might be interpreted as such marginal curves. Moving up on the monotonously increasing offer curve can tell us the cost of generating one additional MWh of electricity. Based on this parallel, social welfare is maximised at the intersection of the bid-offer curves. At the intersection, there is always one power plant that represent that price level, this power plant is hereby referred to as the marginal plant. For the remainder of the thesis, the technology of that plant will be relevant, therefore this price setter is referred to as the marginal (technology) type.

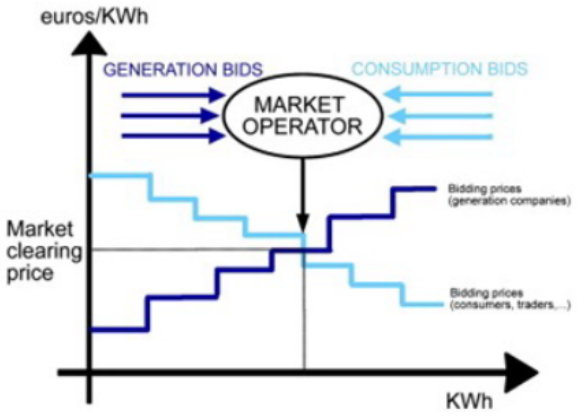

Figure 3.2: An illustration of the market clearing procedure (taken from lecture slides of Dr. Stefan Pfenninger at TU Delft).

Based on its properties mentioned above, the market clearing can be formulated as a social welfare maximisation algorithm, or in other terms, an algorithm that minimises system costs (Chatzigiannis et al., 2016). If one assumes hourly consumption bids (hereby referred to as demand) highly inelastic, one only needs to focus on minimising the costs of generation. To end up with a more stylized algorithm, the market orders can be abstracted away by only considering the (marginal) costs of generation, and assuming a continuously controllable generation level. Based on Section 2.1, this problem can be formulated as a linear problem, resembling (2.2). In this case, $\mathbf{c}$ represents the cost of different power plants, $\mathbf{x}$ is the vector of their respective generation level. The values of $\mathbf{b_2}$ represent the generation limit (installed capacity) of power plants, while $\mathbf{b_1}$ represent the inelastic demand.

To further increase social welfare, European market coupling allows trading to occur across borders, distributing electricity more efficiently over the continent (SPOT, 2022). This means that so-called bidding zones are established, and market orders are collected separately within each zone. When running the continental clearing algorithm (the previously introduced EUPHEMIA), the zonal orders are freely allocated within the zones, whereas inter-zonal allocations are constrained (Nemo, 2020). These constraints are there to represent the physical transmission limits across borders. This also means that in a perfect infrastructural setting, the market could clear with a single market price for the whole continent. As long as the constraints are there, zonal clearing prices differ, however. As introduced in Chapter 1, these constraints can be formulated either via the traditional ATC approach, or via flow-based market coupling, a methodology introduced in the last decade.

The concepts of FBMC have been gradually introduced in Chapter 1. Since accounting for all the network elements in the clearing algorithm would be computationally too intensive, only a subset of these elements are included in the constraint sets of the welfare optimisation, that are deemed the most critical. These elements hence are referred to as Critical Network Elements (CNEs). To determine CNEs, TSOs first compute the zonal PTDFs for all their network elements, based on the nodal PTDF matrix and a GSK calculation strategy appointed by the TSO internally. Zonal PTDFs represent the sensitivity of a line to power injection in the respective zone. To select CNEs, one is interested in the sensitivity of a line to zone-to-zone exchanges. To obtain these zone-to-zone sensitivities, the absolute difference is taken between two zonal PTDF values of the same line. In the selection process, if at least one of these zone-to-zone sensitivities is above a threshold for a line, the line is appointed as critical.

This selection criteria is mathematically formulated in (3.1), and the threshold is hereafter referred to as α.

$$|PTDF^{Z}_{l,z} - PTDF^{Z}_{l,z'}| \geq \alpha \qquad \forall z \in \mathcal{Z} \qquad \forall z' \in \mathcal{Z} \setminus \{z\} \qquad \forall l \in \mathcal{L} \tag{3.1}$$

These α values are also under the discretion of TSO choices. The threshold of 50Hertz is reported to be at 5%, while it is 8% for Amprion (two of the four the German TSOs, published in 50Hertz et al. (2017)).

While the PTDFs represent the participation of lines in the constraints, the actual limiting factor is the Remaining Available Margin (RAM), which is calculated in the following way:

$$RAM_{t,l} = cap_l - frm_l - F^{D-2}_{t,l} \tag{3.2}$$

where cap_l is the maximum enabled power flow on a line, frm_l is the Flow Reliability Margin, and $F^{D-2}_{t,l}$ is the reference flow calculated for the line, in a given market time unit (MTU) (Nemo, 2020). The Flow Reliability Margin is also subject to TSO choices, these are being publicly reported however (JAO, 2022b). They represent potential errors made in the reference calculations, defined for lines, but they tend to be consistent over a zone, mainly subject to seasonal changes. To calculate reference flows, TSOs run so-called base cases a day before the day-ahead clearing (so two days ahead), where they attempt to forecast flows and congestion that can assumed to be present over the lines, before making them available to flow-based market coupling. There is no commonly agreed academic approach on how to model base cases. They are most likely experience-based, potentially taking into account historical generation of conventional power plants, and they are subsequently adjusted for load forecasts and renewable feed-in.

FBMC defined over the CORE region does not operate in isolation, cross-border exports and imports are also incorporated from the neighbouring non-FBMC zones. Trades happen via ATC over these borders. Following this addition, the welfare optimisation problem for FBMC is formulated in (3.3), based on Schonheit et al. (2021).

$$\min \sum_{t,z} c_p \cdot G_{t,p} \tag{3.3a}$$

subject to

$$G_{t,p} \leq g^{\max}_{p} \qquad \forall t \in \mathcal{T} \quad \forall p \in \mathcal{P} \tag{3.3b}$$

$$d_{t,z} = \sum_{p \in mp(z)} G_{t,p} + ren_{t,z} - NP_{t,z} + \sum_{x \in mz(z)} [EX_{t,x,z} - EX_{t,z,x}] \qquad \forall t \in \mathcal{T} \quad \forall z \in \mathcal{Z} \tag{3.3c}$$

$$EX_{t,z,x} \leq atc_{t,z,x} \qquad \forall t \in \mathcal{T} \quad \forall x \in mz(z) \quad \forall z \in \mathcal{Z} \tag{3.3d}$$

$$NP_{t,z} = 0 \qquad \forall t \in \mathcal{T} \quad \forall z \notin \mathcal{Z}^{FB} \tag{3.3e}$$

$$\sum_{z} NP_{t,z} = 0 \qquad \forall t \in \mathcal{T} \tag{3.3f}$$

$$\sum_{z} ptdf^{Z}_{j,z} \cdot NP_{t,z} \leq ram_{j,t} \qquad \forall t \in \mathcal{T} \quad \forall j \in \mathcal{J} \tag{3.3g}$$

The total cost of generation is minimised in the objective function. The supply-demand equality is enforced in (3.3c) for each zone, where the supply side consists of the level of conventional generation, the aggregated renewable generation, the zonal net position, and the export-import balance with neighbouring non-FBMC zones. The non-FBMC exports and imports are capped with the Available Transfer Capacity in (3.3d), defined over each border. Constraint (3.3g) represents the flow-based domain within the FBMC region, constraining the zonal net positions via the PTDFs and RAMs of the identified set $\mathcal{J} \subset \mathcal{L}$ of CNEs. Finally, here generation is represented on a power plant level, but as the Reader will see, from Chapter 5, it will be aggregated on a technology and zonal level, while the marginal cost vector will be replaced by generation-dependent cost functions.

Subsequently, the same FBMC model is formulated with a simplified matrix notation in (3.4), which will be used to reformulate the program to its inverse counterpart in Chapter 5. (3.4b) represents the supply-demand equality in (3.3c), while (3.4c) represents the equality constraints of the exchange (net

position) variables, which corresponds to (3.3f) in this case. (3.4d) depicts the inequality constraints for generation and ATC limits, which is represented by (3.3b) and (3.3d). Finally, (3.4e) corresponds to the flow-based constraints as they are in (3.3g). The corresponding dual variables are displayed in parenthesis next to each constraint. Among these, $\boldsymbol{\lambda}$ represents the zonal clearing price, which is the shadow price for each zonal supply-demand equality constraint.

$$\min \mathbf{1}^T \mathbf{c} \tag{3.4a}$$

subject to:

$$\mathbf{A_g} \cdot \mathbf{g} = \mathbf{b_{1,g}} \quad (\boldsymbol{\lambda}) \tag{3.4b}$$
$$\mathbf{A_e} \cdot \mathbf{np} = \mathbf{b_{1,e}} \quad (\boldsymbol{\kappa}) \tag{3.4c}$$
$$\mathbf{B_g} \cdot \mathbf{g} \leq \mathbf{b_{2,g}} \quad (\boldsymbol{\mu}) \tag{3.4d}$$
$$\mathbf{B_e} \cdot \mathbf{np} \leq \mathbf{b_{2,e}} \quad (\boldsymbol{\nu}) \tag{3.4e}$$

4

Static Grid Model

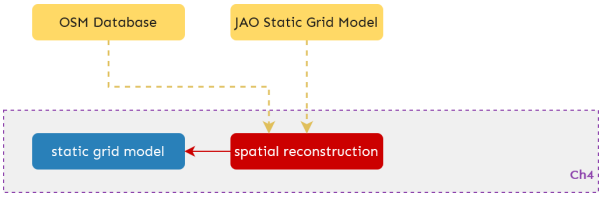

Figure 4.1: Modelling stage of spatial grid reconstruction

As the first piece in the methodological part of the thesis, this chapter discusses the spatial reconstruction of the static grid model of the CORE region, based on the available JAO publication and publicly available geographical data from OpenStreetMap. To visually follow the modelling process, one is referred to Figure 4.1, which highlights the first stage. First, the role of a static grid model in the spatial granularity of electricity markets is presented. Afterwards, the published JAO Static Grid Model is introduced. This is followed by obtaining the geographical locations of substations as the network nodes. As a next step, the geographical locations of power plants reported on the ENTSO-e Transparency Platform are also retrieved. Since these power plants will be modelled to inject power in one of the network nodes, the plants are subsequently mapped to their closest substation nodes based on their location. All data was retrieved via APIs provided by the respective platforms and was subsequently analysed in a Python Jupyter Notebook environment.

4.1. Spatial granularity of electricity markets

In the European markets, network nodes (substations) are aggregated into bidding zones, within which uniform prices are formed (day-ahead zonal clearing prices). However, in an age of increasing geographically distributed power (Bauknecht and Brunekreeft, 2008) and intermittent renewable generation, there is a trend towards smaller bidding zones to better reflect local scarcity (Bichler et al., 2022). This is reflected by the recent splitting of the previous Germany-Austria-Luxembourg zone into Germany-Luxembourg and Austria (Hurta et al., 2022), or the proposal of ACER to split the Dutch zone into two (ACER, 2022).

At the same time, FBMC also handles the grid with a higher spatial granularity. Even though prices are formed uniformly within a zone, inter-zonal transmission constraints are further accounted for in FBMC to quantify their affect on inter-zonal exchanges. Accounting for these intra-zonal transmission constraints is crucial both from an operational and long-term investment perspective. Since traditionally intra-zonal transmission lines were ignored during market clearing, TSOs needed to employ frequent expensive corrective measures to mitigate congestion (Bichler et al., 2022). In Germany, redispatch costs have been reported to be over 1 billion euros in 2019 (Bundesnetzagentur and Bundeskartellamt, 2021). Within the European energy system, this contributes to considerable welfare losses (Grimm et al., 2016). From a long-term perspective, displaying congestion signals for more transmission lines can lead to better investment signals for grid expansion.

Based on the developments mentioned above, there is a general need to model the grid on a transmission line, and hence on a nodal level. While a static grid model is already compiled by JAO (JAO, 2023), a spatial reconstruction can account for power injection patterns of spatially distributed power plants and further contribute to analysing different zonal reconfigurations. From a socio-technical perspective, while the nodal grid model purely lies on the technical side, the zonal configuration is already on the institutional side, concerning the allocation of property rights (in the form of access rights) (Bromley, 1989).

4.2. JAO Static Grid Model

The JAO Static Grid Model (JAO, 2023) is a data set compiled and maintained by JAO, collected from the publications of each TSO included in the CORE region. They contain high-voltage internal lines, tielines (cross-border lines) and transformers. For each transmission line, the data set includes its name, EIC (unique identifier maintained by ENTSO-e), maintainer TSO, the two substation names it connects, and its physical properties. The physical properties comprise the voltage level, maximum current rating, line resistance, reactance, susceptance and length. In total, it contains 2837 high-voltage lines.

4.3. Locating substations

In the Static Grid Model, only the substation names are disclosed, which is usually a name of a municipality in which the substation is situated. To spatially reconstruct the grid, the latitude and longitude coordinates are needed to be obtained for these substations. This was achieved by geocoding, using the OpenStreetMap Nominatim API (OpenStreetMap Contributors, 2017). Geocoding is when one searches for a geographical coordinates based on a name or address of the location. On the API, the name of the substation was searched for. As suggested by the documentation (OpenStreetMap Wiki, 2023), the following keywords were also specified in the API requests: *power*, *substation* and *electricity*. Since this was an automated process, it was possible that geocoding found a wrong match for some of the substations. This was possible to check by validating the resulting transmission line lengths. To obtain an estimation of the line length based on its two terminal locations, the geodesic distance between the two substations were calculated, then it was multiplied by 1.1 (based on Puiu and Hauser (2021)), to account for a longer length in reality. This estimated length then were compared to the ones published in the Static Grid Model. For the lines where the mismatch was larger than 50%, the returned substation coordinates were checked manually, and if needed, were looked up and corrected manually. Out of the 1477 substations, 67 manual corrections were necessary.

4.4. Locating power plants

The conventional power plant fleets of each zone have been retrieved from the ENTSO-e Transparency Platform (ENTSO-E, 2023), where both the installed capacities and technologies of units were disclosed. Based on this data set, the thesis operates with 554 conventional plants in total. The zone level compositions are visualised in Figure 4.2, grouped by the 10 identified technologies that are reported on ENTSO-e.

For flow-based domain inference, the locations of these power plants are also necessary to be known, which are not disclosed in public data sets as a whole. Therefore, a similar geocoding procedure was followed through the OpenStreetMap API. The names of power plants were searched for, and as suggested by the OpenStreetMap Wiki, the following keywords were added to the query: *power*, *plant*, *generator*. Most of the plants were successfully located via the automated script, 41 of them needed to be located manually. After retrieving their locations, power plants were mapped to the closest substation based on geographical proximity. This is a necessary simplification, as the data set of the high-voltage static grid does not necessary contain the lines and nodes that connect the power plants to the network. This way one can make sure to map all the generation capacity to one of the network junctions. The final results of the spatial grid reconstruction is presented in Figure 4.3. In the upcoming chapters, this grid will be used to facilitate calculations on the flow-based domain.

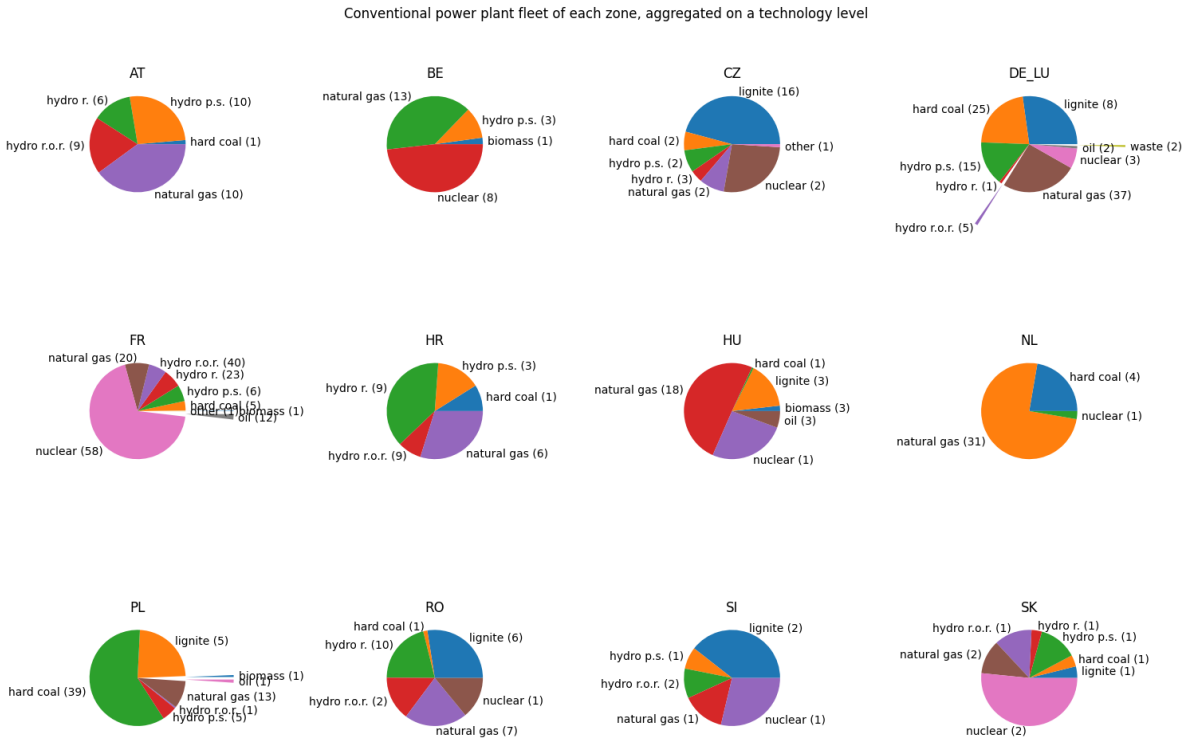

Figure 4.2: Conventional power plant for each zone, aggregated on a technology level, compiled from ENTSO-E (2023). The pie-chart shares represent the distribution of generation capacity across the technologies, while the numbers next to the labels represent the number or plants reported for that technology.

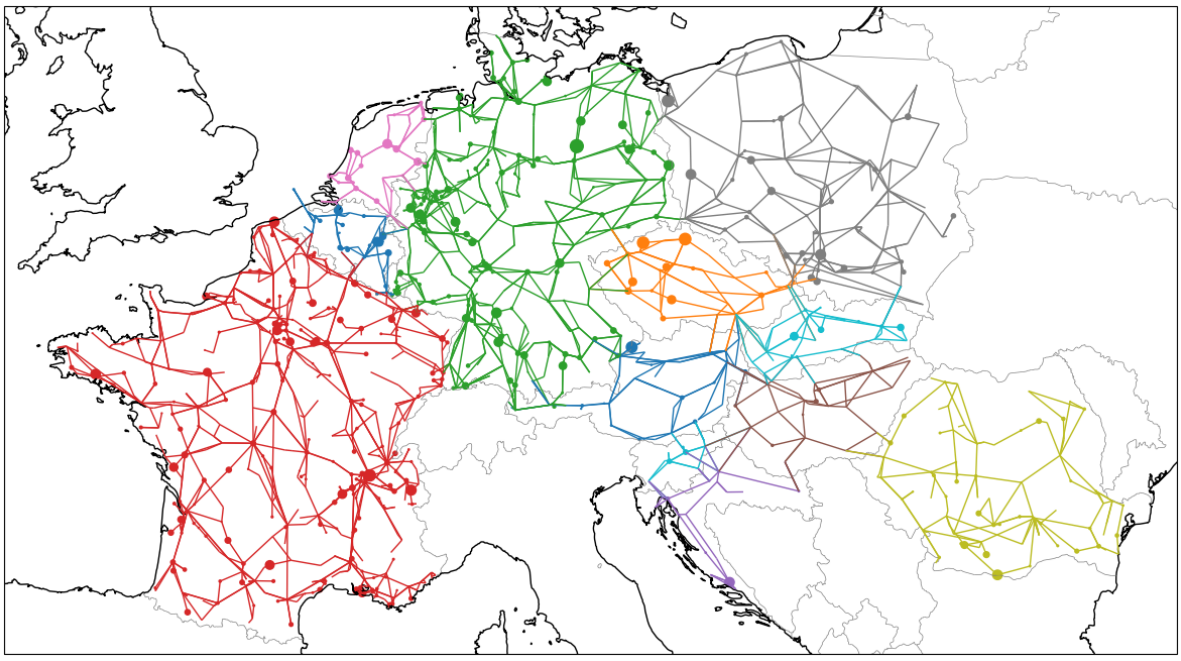

Figure 4.3: The overall result of reconstructing the static grid within the CORE region. The node sizes indicate the amount of generation capacity mapped to those nodes. The different colours correspond to different bidding zones.

5

Cost Curve Reconstruction

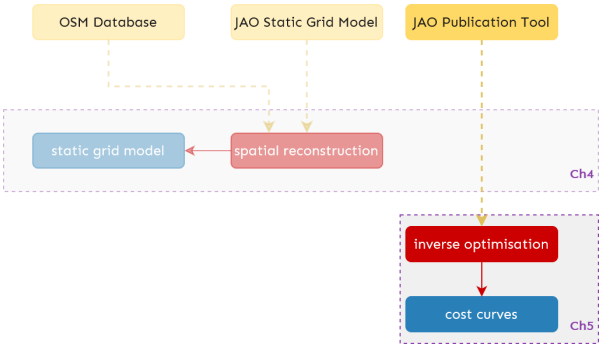

Figure 5.1: Modelling stage of cost curve reconstruction

Having the theoretical foundations of inverse optimisation laid down in Chapter 2, and the model of flow-based market coupling discussed in Chapter 3, this chapter aims to develop an inverse optimisation model to impute the offer curves of power plants. To visually follow the modelling process, one is referred to Figure 5.1, which highlights the current stage. The chapter starts with discussing the formation of electricity prices within the energy infrastructure and presents the relevance of cost curve reconstructing for forecasting market outcomes. Subsequently, it discusses how the cost characteristics of power plants will be captured, then proceeds with enumerating the types of observations available at hand to conduct inverse optimisation. Finally, the IO model is presented.

5.1. Electricity price formation in the energy system

The prices with which energy is being offered on the zonal spot markets directly affect the zonal clearing prices, which affect the zonal net positions and cross-border flows via the resulting inter-zonal price spreads. It is therefore crucial to realistically capture these price levels if one aims to simulate how the market operates within the CORE region. Prices on the spot market partly correspond to production assets. These prices can incorporate fuel prices of the corresponding fuel for a thermal plant, further accounting for its thermal efficiency. Prices can further include government-imposed components, such as the carbon credits (Lovcha et al., 2022). While market participants pricing their assets affect the zonal spot prices, these spot prices also affect the assets, constituting a feedback loop. Production units offering under the resulting market price will produce electricity, while the ones offering above will not. In more techno-economic terms, the price duration curve (showing the frequency for certain market price levels for a given period) will determine how much a power plant produces over the period (De Vries et al., 2020). This production amount can be related to the generated income, and when considered over the lifespan of the unit, it can be related to its profitability. Hence, long-term market price signals affect the investment climate on given power plant technologies. While the focus of this

thesis is short-term optimisation, this latter aspect is utilised within long-term capacity optimisation models (De Vries et al., 2020).

The main uncertainty within the above-mentioned price composition is the efficiency level, which generally varies by power plant and also by operational conditions (ramping up and down). The lack of information on these efficiency values and operational insights makes it challenging to capture realistic market price formation by accounting for each power plant in the above-mentioned manner. Given these necessary insights that are unavailable to the academic modeller and to the independent regulator, there is a need to derive cost functions that can sufficiently approximate bids and offers, to simulate market price outcomes. This thesis will subsequently attempt to derive such theoretical cost functions to simulate realistic market outcomes, while keeping the original, forward model both conceptually and computationally simple. The computational simplicity comes from the fact that the model aims to use single cost functions for aggregated sets of power generation, instead of simulating bids separately, as it takes place in real markets.

5.2. Choice of cost characteristics

Based on the available data, and to keep the model within a maintainable size, the following resolution is chosen. Aggregated cost functions of power plants are constructed on a technology level per each zone. Generation levels on a per technology and per zone basis are adequately reported on the ENTSO-e Transparency Platform, making it feasible to fit our proxy model with a similar resolution to the available observations through inverse optimisation. Generation on a unit level is also reported, but it is not equally available for all the CORE countries, making the data fitting problem considerably under-determined.

It is common to capture the cost functions of thermal generators with a quadratic cost curve, taking the following shape:

$$c(g) = \alpha + \beta g + \gamma g^2 \tag{5.1}$$

where c is the cost of power production dependent on the generation level and g is the generation output of a plant. The α constant component can be attributed to all fixed operational and maintenance costs. The β and γ coefficients on the one hand account for the fuel dependence of plants, capturing the fact that producing more power involves more fuel consumption, increasing the overall cost of production. The quadratic nature of the curve comes from the theoretical thermal input-output curve of generators, also involving the thermal cycle efficiency. The thesis defines this quadratic cost function for every technology in every zone separately, and reconstructing the coefficients of these function will be the subject of the inverse optimisation problem.

At this point, it is important to make two remarks. Firstly, hydro-power plants are being dynamically optimised across market time units, and practically cannot be captured with thermal cost curves. Including such a dynamic (and supposedly strategic) element in the market optimisation model is outside the scope of this thesis however, therefore it has been decided to keep the quadratic form for every technologies.

The second remark is about power plants whose fuel costs are heavily market-based, which means that these fuel prices are subject to considerable volatility that is external to the electricity markets. This is the case for natural gas, coal and oil fuelled generators. It is therefore important to correct the cost functions of these technologies with their respective known fuel prices at a given time, so that the inverse model does not try to capture this exogenous volatility within the cost function coefficients. The fuel price-corrected cost functions take the following form:

$$c(g) = \alpha + c_f(\beta_f g + \gamma_f g^2) \tag{5.2}$$

5.3. Available observation data

The goal of the inverse optimisation is to impute the power plant cost functions based on market outcomes as available observations. Based on the publications available on the ENTSO-e Transparency Platform, the JAO Publication Tool and on Gridwatch, the quantities that can be used as observations are presented in Table 5.1.

Quantity	Platform	Resolution	Time horizon
total load	ENTSO-e	zonal	day-ahead
			actual
generation	ENTSO-e	per technology	actual
		per unit	actual
renewable generation forecast	ENTSO-e	zonal	day-ahead
day-ahead prices	ENTSO-e	zonal	day-ahead
scheduled commercial exchanges	ENTSO-e	border	actual
forecasted transfer capacities	ENTSO-e	border	day-ahead
net positions	JAO	zonal	actual
CNEC publications	JAO	line	actual
UK total load	Gridwatch	zonal	actual
UK generation	Gridwatch	per technology	actual

Table 5.1: Published quantities available to use as observations in the inverse optimisation problem

The total load per zone is used as the demand in the supply-demand equality constraint. Since it is assumed to be known to the market operator during running the market clearing algorithm, the day-ahead version is used to account for this fact. As this thesis does not model renewable curtailment, the renewable generation forecasts (also day-ahead) are deducted from these total load values to obtain a residual demand, that will eventually be present in the supply-demand constraint. When modelling exchanges with the non-FBMC zones via ATC trades, the UK data needs to obtained from Gridwatch, as it is not present on ENTSO-e. When modelling ATC trades, the published data on forecasted transfer capacities is also necessary, which is also assumed to be known during the clearing. At this point of the modelling process, the flow-based domain is not in the focus yet, the PTDF and RAM values will be retrieved from the JAO CNEC publications, which are also assumed to be know to the market operator during the clearing.

Up to this point, only data has been mentioned that is already available to the market operator during clearing. While this data will aid to run the forward model, an inverse formulation requires also outcomes of the forward model as observations. The reported actual generation levels (per technology, or per unit), zonal net positions and zonal day-ahead prices can be used as such market outcomes. How these market observations are used or not used in the inverse formulation is discussed in the next section.

5.4. Inverse optimisation formulation

In order to reconstruct the cost functions, the forward flow-based model introduced in Chapter 3 needs an inverse reformulation in accordance with the primal-dual formulation discussed in Chapter 2. At a first glance, the goal here seems in line with the classic formulation: the cost function needs to be reconstructed knowing the resulting optimal values of the decision variables. These known optimal values are the actual generation levels and the resulting zonal net positions. With these observations at hand, one could already retrieve a set of non-zero cost functions that make these observed values optimal. The only problem here is that only the relative cost levels would be reconstructed here, that is able to replicate the merit order captured in the resulting generation levels. To obtain meaningful price forecasts, these cost functions need to be adjusted for realistic price levels that are reflected in the market outcomes. With that said, additional cost information needs to be encoded in the IO problem. The price information that is available here is the zonal day-ahead price, which corresponds to the $\boldsymbol{\lambda}$ dual variable in the forward problem. Due to the primal-dual formulation, this variable is also present in the IO model directly, therefore it makes sense to include the day-ahead prices as observations for $\boldsymbol{\lambda}$.

As a result of the considerations above, one has 3 types of decision variables observed: $\mathbf{g_0}$ generation levels, **np** net positions and $\boldsymbol{\lambda}_0$ shadow prices. Based on initial experimental runs, it can be stated early on that observing net positions either over-determines the problem in some scenarios, or does not result in improvements in others. Therefore the thesis proceeds with focusing only on $\mathbf{g_0}$ and $\boldsymbol{\lambda}_0$.

The next question to explore is how to observe these decision variables. Simply fixing the values of these variables as it is done in (2.5) yields an infeasible problem. This was expected, as the forward

model used here is clearly a stylized proxy model of the EUPHEMIA algorithm (from where the observations originate), having a considerably lower dimensionality. To start along the line of an L-norm formulation, one could consider keeping the generation level as a decision variable, an introducing an error term as their distance from the observed levels (as it was shown in (2.6)). There a bilinear term arose, due to the multiplication of two decision variables. Here the cost vector is already quadratic however, so further multiplying with the generation level will result in a non-convex NLP problem. Therefore, it is only feasible to turn the generation to a fixed observed parameter. Fortunately, the L-norm approach can be kept in the case of $\boldsymbol{\lambda}$, relaxing the initially encountered infeasible problem.

One could also introduce duality gap minimisation. Such an option would further relax the requirement of the IO problem to yield an optimal solution for the observations, but might also higher chance for IO to better fit the observed price points to the resulting curves (favouring regression over optimality). In Chapter 2 the $\mathbf{IO}_L$ and $\mathbf{IO}_d$ concepts were introduced separately, but these two approach can be combined in fact. With that in mind, the inverse formulation of the forward problem is presented in (5.3). One can also notice the addition of the quadratic cost expression in (5.3f), where $\odot$ represents the element-wise multiplication between the vectors. The source code for the inverse problem written in Julia is briefly presented in Appendix A

$$\min ||\boldsymbol{\lambda} - \boldsymbol{\lambda_0}||_L + \varepsilon_d \tag{5.3a}$$

primal constraints:

$$\mathbf{A_g} \cdot \mathbf{g} = \mathbf{b_{1,g}} \tag{5.3b}$$
$$\mathbf{A_e} \cdot \mathbf{np} = \mathbf{b_{1,e}} \tag{5.3c}$$
$$\mathbf{B_g} \cdot \mathbf{g} \leq \mathbf{b_{2,g}} \tag{5.3d}$$
$$\mathbf{B_e} \cdot \mathbf{np} \leq \mathbf{b_{2,e}} \tag{5.3e}$$
$$\mathbf{c} = \boldsymbol{\alpha} + \boldsymbol{\beta} \odot \mathbf{x} + \boldsymbol{\gamma} \odot \mathbf{x} \odot \mathbf{x} \tag{5.3f}$$

dual constraints:

$$\begin{bmatrix} \mathbf{A_g} & \mathbf{0} \\ \mathbf{0} & \mathbf{A_e} \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\lambda} \\ \boldsymbol{\kappa} \end{bmatrix} + \begin{bmatrix} \mathbf{B_g} & \mathbf{0} \\ \mathbf{0} & \mathbf{B_e} \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\nu} \end{bmatrix} = \begin{bmatrix} \mathbf{c} \\ \mathbf{0} \end{bmatrix} \tag{5.3g}$$

relaxed duality:

$$\mathbf{1}^T \mathbf{c} + \varepsilon_d = \mathbf{b_{1,g}}^T \boldsymbol{\lambda} + \mathbf{b_{1,e}}^T \boldsymbol{\kappa} + \mathbf{b_{2,g}}^T \boldsymbol{\mu} + \mathbf{b_{2,e}}^T \boldsymbol{\nu} \tag{5.3h}$$

5.5. Experiment design

At the time of conducting the experiments, data (the publications shown in Table 5.1) is available for 6 months, starting from 2022 September, up until 2023 February. This short time span is due to the fact that the CORE region has only been established in 2022 August, before which 8 of the bidding zones were not part of the flow-based domain. Considering reconstructing the price dynamics before this time might be informative in general, but not for the purpose of replicating the dynamics of the CORE region.

These 6 months have been split into a training and a validation set. On the training set the inverse optimisation is conducted, while on the validation set the forward model is run to assess the performance of the acquired cost functions. Given that these 6 months include two seasons, the months have been split in the following way: 2022 September, October, December and 2023 January are part of the training set, while 2022 November and 2023 February are the validation months. This way one has a validation month from both seasons.

5.5.1. Problem size and optimisation intervals

Now that the training data set is given, the discussion continues with the problem size of the inverse optimisation. The number of variables are presented in Table 5.2. The calculation includes the non-FBMC bordering countries as well. The error term ϵ_d is excluded from the list, as there is only one of that for the whole inverse problem. The number of decision variables adds up to 824 per market time unit (MTU hereafter), which is an hourly resolution in this case. The number of MTUs for the 4 training months is 2952 however, which results in a rather large problem space with 2 432 448 variables.

Decision variable type	Number per MTU
net position (np)	12
ATC exchange (ex)	$2 \cdot 11 = 22$
alpha	$16 \cdot 10 = 160$
beta	$16 \cdot 10 = 160$
gamma	$16 \cdot 10 = 160$
lambda	16
mu	$16 \cdot 10 = 160$
kappa	1
nu	$133 + 2 \cdot 11 = 155$
Total per MTU	824
Total in training set	$824 \cdot 2952 = 2432448$

Table 5.2: Derivation of problem size with regards to the number of decision variables in the IO problem

During the thesis, a personal computer with an AMD 12-Core 3081 Mhz processor and 32 GB of RAM was available to run the optimisation. This computational power was not enough to run the model for the whole training set at once within a considerable amount of time and was further constrained by the available memory. After experimenting with splitting the training data into smaller chunks, the following division emerged as the best solution. A window of 10 days has been chosen. This means that the cost curves are reconstructed for every 10 days of observed market data. This results in 18 sets of cost curve reconstructions, whose coefficients are simply averaged over the 18 sets to retrieve an overall reconstruction based on the whole training set.

5.5.2. Coefficient averaging validation

Splitting the optimisation into 10-daily runs and averaging all the resulting coefficients might bring about the following questions: Is it truly necessary to use all of the reconstruction sets? Would the forecasts not deliver better results if one only uses the most recently reconstructed cost curves (only for the last 10 days for example)?

The following preliminary experiment attempts to answer these questions. For this analysis, the concepts of training and validation sets are disregarded. Cost curves then are reconstructed for every subsequent 10 days within the overall data set. Then the forward model is run for every day in the data set, using the average of cost curves that are available up to that day. The available set of 10-daily reconstructions are not simply taken with their mean average, however. An exponential moving average (EMA) is applied instead, a common filter used in statistical time series analysis (Lucas and Saccucci, 1990). It has an exponential factor k with which the most recent reconstruction sets are given larger weights than older ones. With this k factor the extent can be controlled to which extent the forecast model uses older cost reconstructions. The smaller the k factor is, the more older cost curves are discarded. If $k = 1$, the mean average is retrieved again.

The results of running the forward model for the whole data set with different k values are displayed in Figure 5.2. The trend is clearly visible that the more data one uses from the past, the more representative cost functions can be obtained. This concludes that it is more important to have more data for fitting the curves than their recency.

5.5.3. Variations of IO approaches

For the L-norm approach, the two commonly used norms will be utilised: the L1 and L2 norms, as introduced in Section 2.4.1. Given these two norms, and their combination with duality gap minimisation, for variations of IO approaches are realised for the experiments: **IO**$_{\text{L1}}$ uses a standalone L1-norm minimisation, **IO**$_{\text{L1+d}}$ combines this with duality gap minimisation, **IO**$_{\text{L2}}$ and **IO**$_{\text{L2+d}}$ represent the same scenarios, but with the L2-norm. The above-described cost curve reconstruction is carried out for all these variations, and their relative performance is evaluated in Chapter 7.

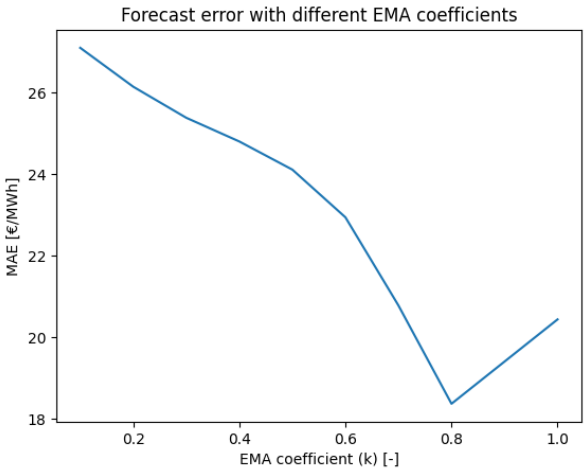

Figure 5.2: Mean average errors of day-ahead price forecast over the whole date set, with different k values

6

Flow-based Domain Inference

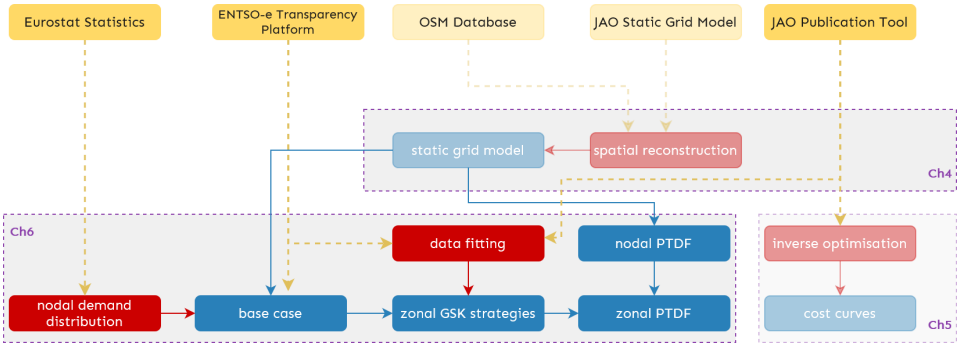

Figure 6.1: Modelling stage of flow-based domain inference

Following up on the reconstructed static grid model in Chapter 4, this chapter enters the flow-based domain. To visually follow the modelling process, one is referred to Figure 6.1, which highlights the current stage. The chapter starts with a perspective of short and long-term income effects of congestion management, and discusses how the need of a structural flow-based domain representation comes into picture. Subsequently, a D-2 nodal base case is developed, that will serve two purposes. The resulting reference flows are used to compute RAMs, while the nodal generation level outcomes will serve as inputs to some of the GSK strategies. The chapter proceeds with enumerating the GSK strategies that are expected to be potentially used by TSOs and discusses their calculation based on available market data and the prior base case runs. Having the potential GSK strategies pre-calculated and zonal PTDF observations for CNEs available, the discussing follows with formulating a data fitting optimisation to reveal which GSK strategies are most likely being used in each bidding zone. Having an identified GSK method for each zone, and the previously calculated reference flows, the complete flow-based calculation process is conducted: zonal PTDFs are retrieved from the nodal PTDFs and GSKs, critical network elements are selected based on the 5% rule, and combined with the computed RAMs, the flow-based constraints are formulated. The results are compared to the published flow-based data in three aspects. First examining the overlap between the selected CNEs by our model and the ones observed, and calculating the error between the calculated and observed PTDF and RAM values for these elements. The chapter concludes with comparing the Max Net Position and Max Border Exchange plots for the calculated and observed domain.

6.1. A need for structural flow-based domain representation

The immediate benefit of knowing the flow-based domain before the market clears is for market participants to predict day-ahead markets. In order to obtain accurate day-ahead forecasts, knowing the cross-border trading domain that constraints net positions is sufficient. As it was shown in Chapter 3 (and further illustrated in Figure 1.1), this is a set of linear inequalities. Forecasting these linear

inequalities can turn to be an abstract geometrical prediction algorithm, employing arbitrary machine learning models. While such models can take exogenous factors (renewable generation, seasonality, fuel prices) as inputs, they do not require a structural derivation of the grid model and TSO strategies over GSKs.

The above-mentioned approach benefits social welfare on the short term. This can be achieved via an improved resource allocation by better informed market participants who optimise generation assets. Day-ahead market forecasts are not only relevant to participants optimising their generation portfolio, but also to other traders optimising their positions in the forward and futures markets. Traders in the forward and futures market are interested in the effects of longer-term developments on the market. Such developments include planned generation or transmission line outages due to maintenance, or grid expansion. These structural changes can only be captured by an equally structural forecast model, where there is a direct correspondence between (the lack of) power injections in the grid, inactive transmission lines and the resulting flow-based domain. One can perceive that expressing such aspects within an abstract geometrical model is limited.

Up to now, only the resource optimisation of market participants has been touched upon. Capturing the flow-based domain in a tractable way has further benefits to regulators. As introduced in Chapter 2, every model constraint has a corresponding dual variable or shadow price. This shadow price for the supply-demand equality is the zonal market price. The shadow prices of flow-based net position constrains also have a corresponding meaning: the cost of congestion. More specifically, they represent the marginal increase in total production costs when relieving congestion (increasing the constraint value) by one MW (Green, 1998). This cost of congestion can give insights for TSOs and regulators in the value of expanding on the given grid elements. However, this congestion cost only represents a momentary value of a marginal expansion, whereas interconnector constructions takes long time and a substantial initial investment (De Vries et al., 2020). Therefore, in order to provide more adequate investment signals, the congestion costs should be obtained for a longer time horizon, which is potentially also affected by changing demand and generation fleet.

6.2. Base case

As discussed in Chapter 3, the base case calculations of TSOs are most likely experience-based, using empirical considerations and proprietary information on their own network. There is no insight on how realistic it is to approximate these empirical approaches by a nodal base case. Nevertheless, it still comes intuitively to follow such a nodal approach in this thesis, as other researchers in the field (Schonheit et al., 2021). Running another market optimisation algorithm fits the interpretable modelling framework, the already reconstructed transmission network and cost curves can directly be utilised here, the already retrieved market parameters (demand and renewable generation forecasts) can also be directly fed into the base case. Only the distribution of the zonal demand forecasts over the nodes need extra consideration.

6.2.1. Formulation

Based on Schonheit et al. (2021), a nodal base case can be mathematically formulated in the following way.

$$\min \sum_{t,z} c_z^{\phi}(G_{t,p}^{\phi,D-2}) \tag{6.1a}$$

subject to

$$G_{t,p}^{D-2} \leq g_p^{\max} \qquad \forall t \in T \quad \forall p \in P \tag{6.1b}$$

$$c_z^{\phi}(G_{t,p}^{\phi}) = a_0^{z,\phi} + f_{\phi}(b_1^{z,\phi} \cdot G_{t,p}^{\phi} + b_2^{z,\phi} \cdot G_{t,p}^{\phi^2}) \qquad \forall t \in T \quad \forall p \in P \tag{6.1c}$$

$$d_{t,n}^{r} + INJ_{t,n}^{D-2} = \sum_{p \in mp(n)} G_{t,p}^{D-2} \qquad \forall t \in T \quad \forall n \in N \tag{6.1d}$$

$$NP_{t,z}^{D-2} = \sum_{n \in mn(z)} INJ_{t,n}^{D-2} \qquad \forall t \in T \quad \forall z \in Z^{FB} \tag{6.1e}$$

$$F_{t,l}^{D-2} = \sum_{n \in N} (h_{l,n} \cdot DELTA_{t,n}^{D-2}) \qquad \forall t \in T \quad \forall l \in L \tag{6.1f}$$

$$INJ_{t,n}^{D-2} = \sum_{q \in N} (l_{n,q} \cdot DELTA_{t,q}^{D-2}) \qquad \forall t \in T \quad \forall n \in N \tag{6.1g}$$

$$DELTA_{t,v}^{D-2} = 0 \qquad \forall t \in T \quad \forall v \in V \subset N \tag{6.1h}$$

$$F_{t,l}^{D-2} \leq cap_l \cdot (1 - frm_l) \qquad \forall t \in T \quad \forall l \in L \tag{6.1i}$$

$$F_{t,l}^{D-2} \geq -cap_l \cdot (1 - frm_l) \qquad \forall t \in T \quad \forall l \in L \tag{6.1j}$$

$$\tag{6.1k}$$

The total cost of generation is minimised based on the quadratic cost curves that were derived previously. The nodal supply-demand balance is enforced by the conventional power generation, renewable generation and nodal injections. Renewable generation is not considered on a nodal level, they are subtracted from the zonal demand at the beginning, leaving $d_{t,n}^r$ as a residual demand. The nodal injections and line flows are determined based on the physical properties of the network (nodal and line susceptances). $h_{l,n}$ is the element of the line susceptance matrix $\mathbf{H} \in \mathbb{R}^{|L| \times |\mathcal{N}|}$, while $l_{n,q}$ is the element of the nodal susceptance matrix $\mathbf{L} \in \mathbb{R}^{|\mathcal{N}| \times |\mathcal{N}|}$. Their calculation is shown in 6.2 and 6.3 respectively, where **K** is the incidence matrix of the network, and **B** is the diagonal susceptance matrix for the lines:

$$\mathbf{H} = \mathbf{B}\mathbf{K}^{\mathbf{T}} \tag{6.2}$$

$$\mathbf{L} = \mathbf{K}\mathbf{H}^{\mathbf{T}} \tag{6.3}$$

Net positions are calculated from the injections. A limit is put on the line flows based on the physical capacity of the line and the Flow Reliability Margins (FRM). The primary output of the base case are the reference flows (used to calculate the Remaining Available Margins) and reference net positions (used in some GSK calculation strategies and to enforce the flow-based constraints in the market clearing algorithm).

6.2.2. Nodal demand distribution

One of the cornerstone limitations of using a nodal base case to model the real-world electricity system is that there is no publicly available demand data for nodes, only for the zones as a whole. To fill this gap, this thesis suggests to distribute the available zonal demand over nodes weighted by historical GDP and population information retrieved for regions where the nodes lie. This is based on the observation in Robinius et al. (2017) that these two indicators correlates with electricity consumption within the region. It is noted that such an approach is a considerable simplification, which further ignores concentrated industrial areas. The implementation of this demand distribution is described below.

First, the regional resolution is determined, based on the *Nomenclature of territorial units for statistics* (NUTS), published by European Commission (2020). The NUTS 2 resolution level is chosen to be appropriate for this application, which comprise the basic regions for the application of regional policies. This mainly corresponds to a county/province-level, having sufficient publications on both GDP and population size.

GDP and population size data is subsequently gathered for all the NUTS 2 regions with the CORE, using the latest available data set, which is from 2021 (European Commission, 2020). Normalised weights to each NUTS 2 region are assigned, which are proportional to the product of the GDP and population size of that region (to account for the total gross domestic product produced within that region). The electricity network nodes are then mapped to the NUTS 2 regions they reside in, by using their geographical locations and geocoding via OpenStreetMap. If only one node lies within a NUTS 2 region, the weight of the region is directly assigned to the node. If multiple nodes preside in the region, its weight is evenly distributed over these nodes. After this procedure, every node has a weight that sums up to 1 within a region. If the total zonal demand is multiplied by these weights, a nodal demand value is obtained.

6.2.3. Feasibility problem

After test runs of the base case, the problem formulated in (6.1) shows to be infeasible for the input data. This is checked for both for the NUTS-based demand distribution as well as a naive even demand distribution among nodes. The infeasibility lies in the line transmission limits, as if they are lifted, the optimisation completes successfully. In order to be able to still generate approximations for reference flows, the lines constraints are relaxed. However, the least amount of relaxation is pursued, which can be implemented via introducing a separate feasibility problem, as described below.

A feasibility problem is formulated in (6.4), which precedes running the main base case (6.1).

$$\min \sum_{l} \varepsilon_l \tag{6.4a}$$

subject to

$$(6.1b) - (6.1h) \tag{6.4b}$$

$$F_{t,l}^{D-2} \leq cap_l \cdot (1 - frm_l) + \varepsilon_l \qquad \forall t \in T \quad \forall l \in L \tag{6.4c}$$

$$F_{t,l}^{D-2} \geq -cap_l \cdot (1 - frm_l) - \varepsilon_l \qquad \forall t \in T \quad \forall l \in L \tag{6.4d}$$

$$\tag{6.4e}$$

The feasibility problem defines an error term for each line capacity. The sum of these error terms are subsequently minimised while the problem stays feasible. This way, one can make sure that the extent of relaxation keeps to a minimum, while making the originally infeasible problem feasible. To achieve the latter, the derived error terms are simply added to the capacities of lines in the original base case formulation.

6.3. Generation Shift Key strategies

In practice, several Generation Shift Key strategies can be employed by TSOs. Finck et al. (2018) enumerates 6 GSK strategies based on previous TSO publications. They range from flat strategies, only weighing nodes by the number of generators present there, or the maximum generation capacity available, to more marginal strategies, weighing nodes based on the available generation capacity that is calculated from the base case. The aspect with which these strategies are extended in this thesis, is the inclusion of production outages. Instead of considering the total installed capacities, the GSK values are corrected for the actual availability of plants based on planned outages. Subsequently, these strategies are introduced briefly, then the section is concluded with formulating a strategy selection for this model.

GSK1: flat strategy

$$GSK_{n,z}^{1} = \frac{n_{Gen,n}}{n_{Gen,z}} \qquad \forall n \in N, \quad z \in Z \tag{6.5}$$

The keys are simply defined here by the number of available generation units located at a given node with respect to all the generation units in the zone.

GSK2: flexible flat strategy

$$GSK_{n,z}^{2} = \frac{n_{flex,n}}{n_{flex,z}} \qquad \forall n \in N, \quad z \in Z \tag{6.6}$$

Similar to the previous case, the keys are defined here by the number of flexible and available generation units (gas, oil and pump storage) located at a given node.

GSK3: maximum installed capacity

$$GSK_{n,z}^{3} = \frac{P_{max,n}}{P_{max,z}} \qquad \forall n \in N, \quad z \in Z \tag{6.7}$$

The keys here are defined by the maximum available generation capacity located at a given node with respect to the maximum available generation capacity available in the whole zone.

GSK4: generation level from base case

$$GSK^{5}_{n,z,t} = \frac{P_{gen,n,t}}{P_{gen,z,t}} \qquad \forall n \in N, \quad z \in Z \tag{6.8}$$

The keys in this case are defined by the generation levels at a given node following a base case model run.

GSK5: available free capacity from base case

$$GSK^{6}_{n,z,t} = \frac{n_{free,n,t}}{n_{free,z,t}} \qquad \forall n \in N, \quad z \in Z \tag{6.9}$$

These keys are also dependent on the outcome of a base case model run, and defined by the number of generation units at a given node that still have free generation capacities after running the base case.

6.4. Critical network elements and contingencies

This thesis defines the PTDFs and hence the flow-based constraints for single line elements. In practice, TSOs are required to include one more aspect in their CNE reports towards the market operator, which is the so-called N-1 contingency. Considering critical network elements within the clearing algorithm has an obvious security of supply objective: it constraints trades within the physical boundaries of the transmission system. The N-1 standard goes one step further in this respect. It requires TSOs to always keep the network in an operating state such that it can withstand the case when one single element is unexpectedly removed from the system (due to an unplanned outage). The N-1 protocol applies when specifying CNEs as well. Critical network elements are always reported with regards to a scenario when a nearby network connection is removed from the system, giving the specified CNE an even higher sensitivity to certain injections and thus larger potential for congestion. This outage scenario-specific CNE record also bears a name, which will be used for further reference: Critical Network Elements and Contingencies (CNEC).

TSOs reporting CNECs instead of CNEs has considerable consequences for the to-be-conducted data fitting as well in this thesis, complicating the process further. Since the published PTDFs are calculated for an N-1 version of the overall network (removing one of the elements), one needs to account for this too while calculating the zonal PTDFs through the GSKs to be matched with the observations. This requires identifying all the outage scenarios considered by TSOs within the studied 6-month period, and pre-calculating the nodal PTDFs for each of these scenarios. One can do so by removing the respective row from the the incidence matrix **K** and row-column pair in the diagonal susceptance matrix **B** for the network element to exclude. The resulting contingency scenario-specific nodal PTDF is hereby referred to as $\mathbf{P^N_s}$.

6.5. Strategy inference

For the purpose of the strategy inference, the zonal PTDFs can be expressed in the following way:

$$\mathbf{P^Z_{s,t}} = \mathbf{P^N_s}\boldsymbol{\sigma}_z^T\mathbf{GSK}^{\sigma}_t \tag{6.10}$$

where $\boldsymbol{\sigma}_z \in \mathbb{R}^{5\times 1}$ is strategy preference vector defined for every bidding zone, representing the relative preference weights over the strategies. The latter means that it needs to sum up to 1 for every zone. $\mathbf{GSK}^{\sigma}_t \in \mathbb{R}^{5\times|\mathcal{L}|\times|\mathcal{Z}^F|}$ is the set of pre-calculated GSK matrices for a given timestamp. Notice that the elements of $\boldsymbol{\sigma}_z$ are not enforced to be discrete. This means that rather than appointing a single strategy for a zone, $\boldsymbol{\sigma}_z$ tells the relative fitness of the strategies to represent the actual TSO choices. It can still be the case that one strategy emerges as clearly dominant (clearly resembling the TSO choices the best), but when the optimisation algorithm cannot decide over which strategy fits the best, the weights can be distributed rather equally.

In order to infer $\boldsymbol{\sigma}_z$, a constrained quadratic program is presented in (6.11). The quadratic nature comes from the choosing a least-square formulation of the objective function in order to fit the model to the observations, while the constraint comes from the summing-up-to-1 criterion of $\boldsymbol{\sigma}_z$.

$$\min_{\boldsymbol{\sigma_z}} \sum_{(l,s,t)\in\mathcal{LST}} ||\mathbf{PTDF}^{\mathbf{T}}_{l,s,t} - (\mathbf{P^N_s}\boldsymbol{\sigma}_{\boldsymbol{z}}^{\mathbf{T}}\mathbf{GSK}^{\sigma}_t)^{\mathbf{T}}_{l,:}||^2_2 \tag{6.11a}$$

$$\mathbf{1}^{\mathsf{T}}\boldsymbol{\sigma}_z = 1 \tag{6.11b}$$

$$\tag{6.11c}$$

where $\mathbf{PTDF}_{l,s,t} \in \mathbb{R}^{1\times|\mathcal{Z}^F|}$ denotes the zonal PTDF observations for the given CNEC (PTDF value for $\forall z \in \mathcal{Z}^F$ over line l. $()_{l,:}$ denote selecting the l-th row of the matrix. It is important to note that PTDFs are not observed for all combinations of $l \in \mathcal{L}$, $s \in \mathcal{S}$ and $t \in \mathcal{T}$, the available observation subset is denoted by $\mathcal{LST}$. The source code for the strategy inference written in Julia is briefly presented in Appendix B

6.6. Quantifying distance between domains

In order to assess the quality of fit after the strategy inference, the difference between inferred PTDF values and the observed ones can be quantified with absolute and relative distance metrics. However, these metrics will only inform about PTDF-level mismatches, while in this thesis the quality of the overall reconstructed flow-based domain is relevant. To quantify the distance between two flow-based domains, the approach of Zad et al. (2021) is utilised. Their work quantifies the geometric distance between two domains as presented in (6.12). The two domain is annotated by FB_A and FB_B, each consisting of a set of vertices A and B respectively. To calculate the distance between FB_A and FB_B, first all the vertices are considered in A, and for each the distance from the closest vertex $b \in B$ is calculated. The sum of these distances results in the overall distance measure between the domains.

$$d(FB_A, FB_B) = \sum_{a\in A} \min \sqrt{(a-b)^2} \tag{6.12}$$

The size and positions of domains can vary border by border across the zones. In order to be able to compare these distance metrics across different borders, the modification of the original form is suggested, to include relative terms, as presented in (6.13). In the adjusted formulation, the resulting distance values are taken relative to the sum of distances of vertices $b \in B^A$ to the origin. B^A comprised the vertices that were selected to be the closest to $a \in A$ in the nominator.

$$d(FB_A, FB_B) = \frac{\sum_{a\in A} \min \sqrt{(a-b)^2}}{\sum_{b\in B^A} \min \sqrt{(b-0)^2}} \tag{6.13}$$

The distance measure described above will be used as an error metric to quantify the difference between the inferred and observed flow-based domains, considered for each border.

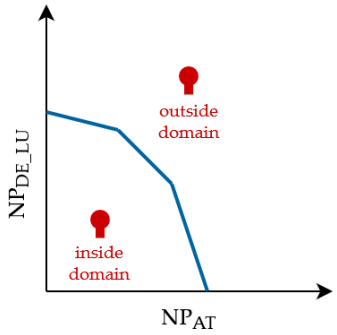

Figure 6.2: Observed flow inclusion rate illustration

6.7. Quantifying observed flow inclusion of inferred domain

In the previous section, geometrical distance (and hence similarity and dissimilarity) between the observed and inferred domains are quantified. However, to quantify the ability to forecast flows that are actually happening in the market, it is more interesting to consider the area within the feasibility domain, where flows are actually observed. Put in other words, in order to forecast cross-border flows, one is more interested to check whether the observed flows are actually allowed by the inferred domain. For

this purpose, the metric of f_{ir} *observed flow inclusion rate* is defined here, as the fraction of observed flows per MTU that fall within the inferred flow-based domain in the examined time period. Keeping ourselves to the projected two-dimensional space of net positions of two-bordering countries, the observed flow is a point in this space, having two respective net position coordinates. The computational task that needs to be carried out is to check in each MTU, whether this observed net position pair falls within the flow-based domain represented by a two-dimensional polytope, or not. The concept is further illustrated in Figure 6.2.

7

Results

This chapter presents the model results that led to forecasting day-ahead prices and flows. Before diving deeper, these model forecasts, which form the final objective of the thesis are introduced in figures 7.1 and 7.2. Figure 7.1 shows the achieved mean relative errors of day-ahead prices forecasts for each bidding zone in the CORE region, relative to observed prices. For each zone, only the best realized results are displayed, ranging from the best of 7.15% in Germany-Luxembourg, until 14.95% in Romania. The best realised results come from different critical network element selections for each zone, which are elaborated in Section 7.4. As one can see, $\alpha = 10\%$ dominates, but the best overall zonal forecast is actually obtained by $\alpha = 5\%$.

In a similar manner, Figure 7.2 shows the mean relative error of cross-border flow forecasts, following the day-ahead market clearing. Given that flows are centred around zero, and can also be zero for longer time periods, the errors are presented relative to the maximum observed flows across the respective borders. The label *n/a* generally indicates that no border flows were observed in the chosen validation period. In the case of Romania, no border flows were forecasted. These flow forecasts are apparent to have a larger spread than price forecasts: ranging from 8.0% between Croatia and Slovenia, until 38% between Belgium and the Netherlands.

Following a broad overview, the chapter proceeds with the analysis of results produced in different stages of the modelling process. First, it examines the effects of the different IO approaches on the cost curve reconstruction. That is followed by presenting the results of the GSK strategy inference and showing the performance of this GSK-based inference in replicating the flow-based domain. The second part of the chapter assesses the performance of the assembled model: day-ahead forecasts delivered by jointly utilising the recovered cost curves and inferred flow-based domain. This last part will give further insights into the dependence of forecasting prices and cross-border flows on the inferred flow-based domain. But first, experimental scenarios are introduced for CNE selection, that will be used to uncover the sensitivity of forecast results to selecting the flow-based domain.

7.1. CNE selection scenarios

An important bridge between the zonal PTDF calculations and constraining the market clearing in the flow-based domain is the aspect of selecting critical network elements. As introduced in Chapter 3, the selection is controlled by the α factor, capturing a level of transmission line sensitivity to cross-border exchanges that is still included in the constraints. In this section, the effects of CNE selection is considered on the outcomes of the market clearing runs.

For α, the scenarios of 10%, 5% and 4% are considered. TSOs are known to operate around these values, with 50Hertz operating at 5% for example (50Hertz et al., 2017). The set of selected CNEs is known not to differ significantly between the range of 10% and 5% as investigated by Schonheit et al. (2021). To isolate the potential errors coming from the CNE selection itself, an additional scenario is added: the same set of CNEs are used that were published by the TSOs in the given MTU, but the calculated PTDF and RAM values are used for these elements. Finally, a naive scenario is added, where the published PTDF constraints are used that were reported the day before. Considering the latter scenario will show whether an intuitive but naive approach can outperform the flow-based domain

inference, with regards to the market clearing results.

For the sake of brevity and relevance, mainly the $\alpha = 10\%$ and 5% as well as the naive scenarios are used throughout the results presentation. In Section 7.4, $\alpha = 4\%$ and the TSO selection are also reported to illustrate that both of them stay on the same performance level as $\alpha = 5\%$ in both price and flow forecasting.

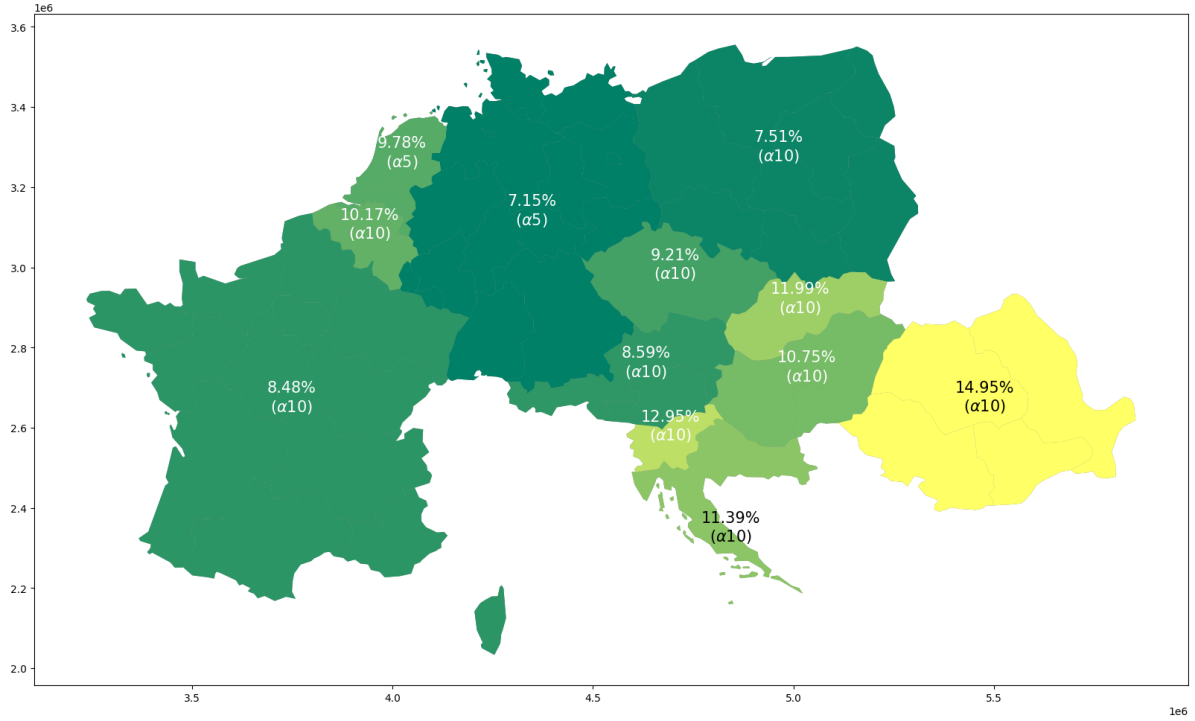

Figure 7.1: Achieved mean relative errors of day-ahead price forecasts relative to observed prices, presented per zone

7.2. Cost curve reconstruction

Two examples of resulting cost curves are presented in figures 7.3a and 7.3b. Both in the case of the Dutch natural gas curve, and the Polish coal curve, a more steep elevation can be observed for the L1-based reconstructions. This steeper elevation would assume higher peak prices, which indeed will make a significant difference in the eventual forecast results.

To separately assess the performance of the cost curves on the price forecast results, the forward model was run on the validation data set with the obtained cost functions, while still using the originally published PTDF and RAMs. On figures 7.4 and 7.5, the results are presented for the 4 IO approaches: $\mathbf{IO}_{\text{L1}}$, $\mathbf{IO}_{\text{L1+d}}$, $\mathbf{IO}_{\text{L2}}$, $\mathbf{IO}_{\text{L2+d}}$. One can make three observations from these charts. First of all, it is apparent that the obtained cost functions are not able to capture the extreme price volatility in November, still in the middle of an unprecedented energy crisis. Especially in the second half of November, due to unnaturally high gas prices, combined with a very high clean spark spread (the profit margin of operating a gas-fired power plant, having bought the fuel, as defined in Abadie and Chamorro (2009)). As price dynamics were starting to return to a more normal setting in February, the results of the forward model became substantially better comparable to the observed prices. As a second remark, duality gap minimisation does not seem to add to the performance of the cost functions, as the resulting price time series are highly overlapping between the cases of having and not having duality gap minimisation.

As a final remark, there is a substantial difference between the performance of the L1 and the L2 norm. This is due to the fact that a smaller amount of 10-daily cost reconstructions were used in the L2 case, which seems to affect the representational power of the averaged cost functions. (One can refer back to Section 5.5.2, where it was concluded that the more 10-daily optimisation results one uses, the better the reconstruction becomes.) A smaller amount of 10-daily optimisation results for the L2 case was caused by the inverse optimisation runs not being able to return a feasible or globally

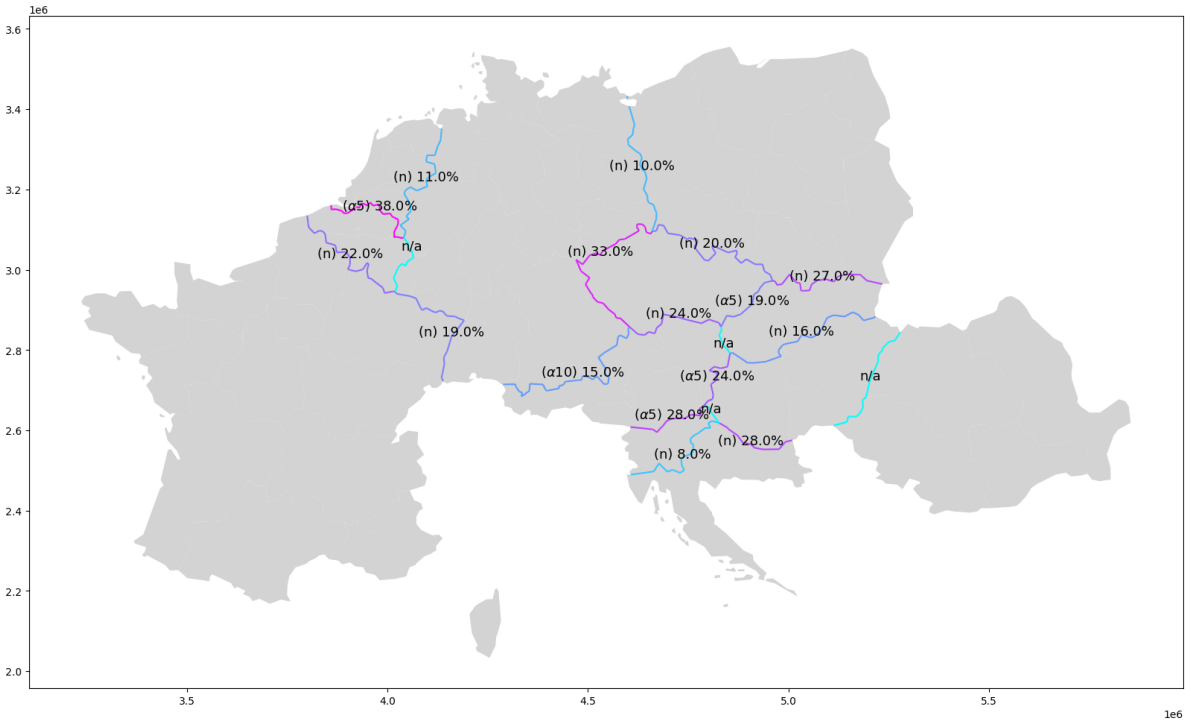

Figure 7.2: Achieved mean relative errors of cross-border flow forecasts relative to maximum observed flows, presented per zone

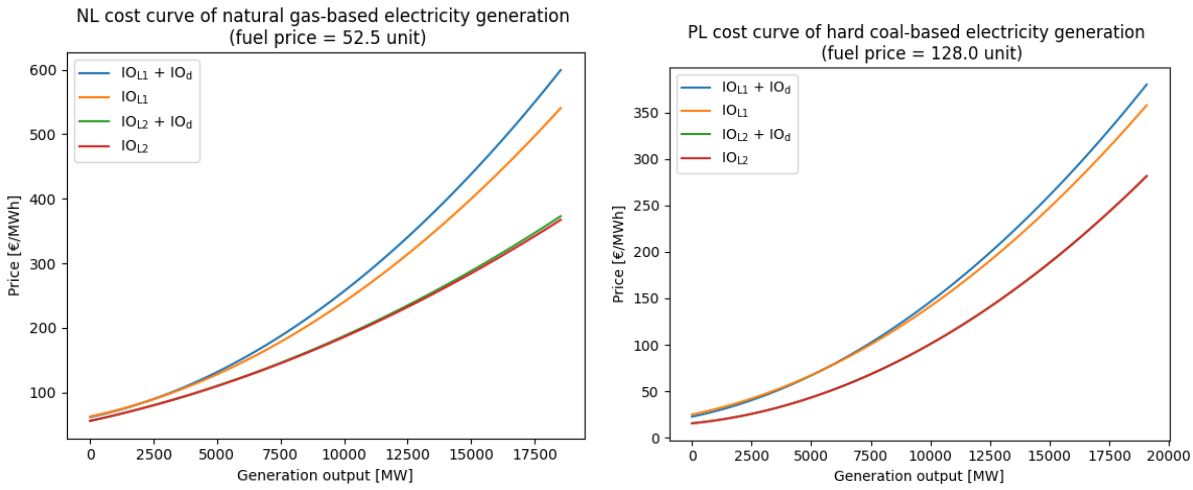

(a) For natural gas-fired plants in the Netherlands

(b) For coal plants in Poland

Figure 7.3: Aggregated cost curve reconstructions with each IO approach

optimal solution within the specified time limit (180 seconds). The time limit was set due to time and computational constraints during the thesis.

Figure 7.6 summarizes the mean relative error obtained for the the different IO approaches across all the CORE zones, while running the forward model to forecast day-ahead prices. On the basis of L1 having a slightly better performance, and striving for a simpler model (without a duality gap minimisation term), $\mathbf{IO}_{L1}$ is selected to be used in the final model assembly.

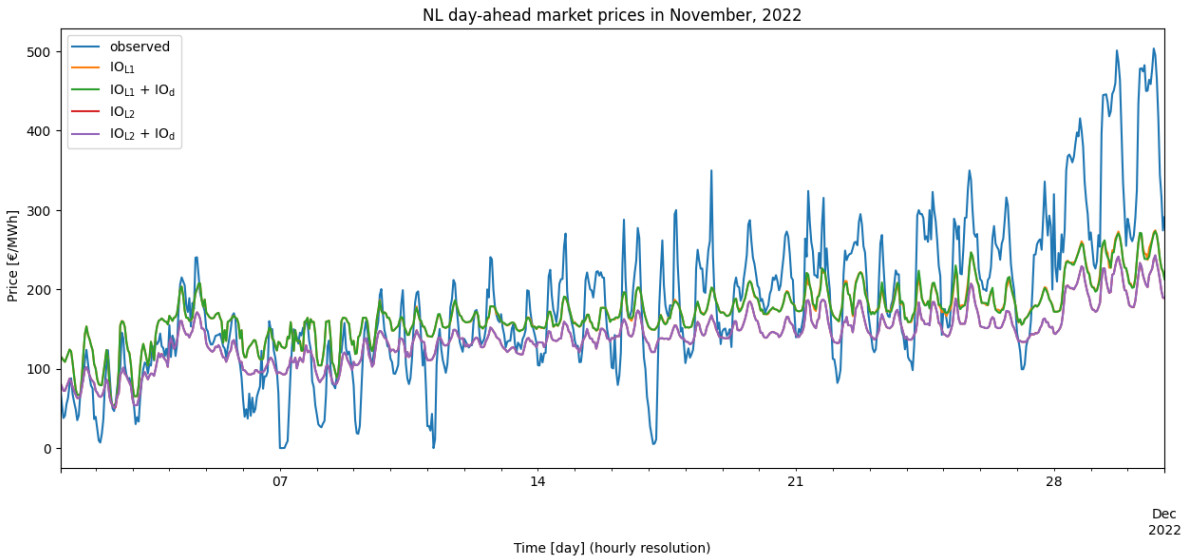

Figure 7.4: Day-ahead price forecasts results in November, 2022 for the Netherlands, using the originally published PTDF constraints for the respective hours

7.3. Performance of the flow-based domain inference

Assessing the performance of the flow-based domain inference is subsequently divided into two parts. First, the ability of the identified GSK strategies to reconstruct observed zone-to-line PTDFs is presented. This assessment further makes it possible to compare the quality of fit with existing literature. More important in this thesis however, is the extent to which the flow-based domain itself can be replicated. To address this objective, the second part of the section compares the resulting flow-based domains following inference to available observations.

7.3.1. Performance of GSK strategies

One can see the identified GSK strategies for each zone in Figure 7.7, following the flow-based domain inference run. It is observed that the flexible flat strategy (GSK2) fits to most of the zones with respect to their published PTDF values over time. In some cases, GSK1 and GSK3 have some shares as well in the matches. The base-case based strategies (GSK4 and 5) are mapped to Austria and Poland, whereas Slovenia also has some share in GSK5.

Figures 7.8 and 7.9 show examples of observed zonal PTDF values over time, compared to their calculated counterparts based on the respective GSK strategies. These plots are presented for given CNECs (with a null-contingency scenario), presenting their assigned PTDF values with respect to a given zone. One can conclude that the levels of the PTDF values are generally well-captured, suggesting a successful calculation of nodal PTDFs from the underlying grid model and the relevance of the GSK strategies.

On the other hand, the volatility of the values does not seem to be sufficiently captured. Figure 7.9 was chosen as a characteristic representation of this issue, but this observation can be made for several CNECs. When one sees a sudden jump in the second half of the plot in Figure 7.9, the curve calculated from GSK3 also seem to follow the trend, although to a much lower extent. The change in GSK3 over time is known to stem only from the extension of the original formulation (Finck et al., 2018) to include reported production unit outages (of flexible units in this case). The fact that this aspect can capture the trend in changing PTDF values shows the relevance of including this factor and also proves that this is most likely accounted for by TSOs. While the trend seems to be followed, the extent is rather insufficient, however. This either means that production outages are considered with a larger factor in TSO calculations, or other (mutually) reinforcing factors are also considered that are not included in the flat GSK strategies here. Such factors might involve the production outages of renewable generation, or planned outages of transmission network elements, neither of which are included in the flat strategies or the base case runs.

Finally, the performance of each GSK strategies are presented in Figure 7.10. The chart shows

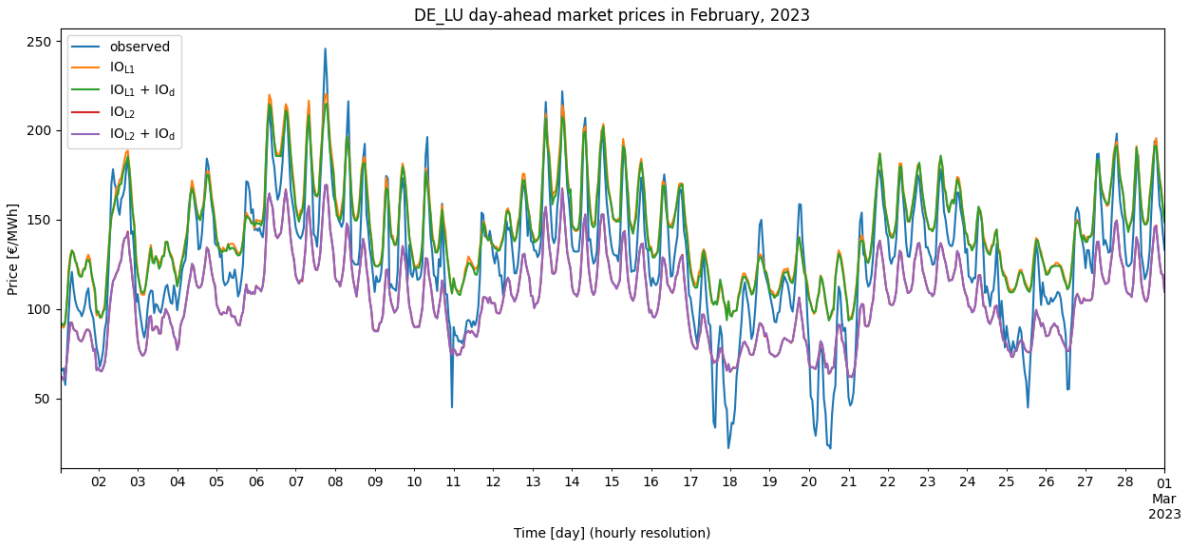

Figure 7.5: Day-ahead price forecasts results in February, 2023 for Germany, using the originally published PTDF constraints for the respective hours

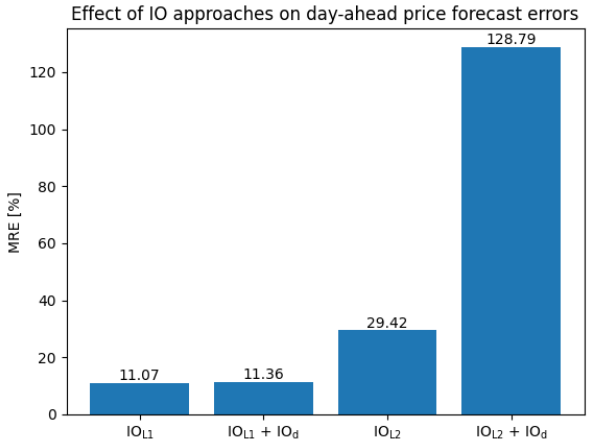

Figure 7.6: Mean relative error of day-ahead price forecasts for different IO approaches, using the originally published CNECs as constraints

how the calculated PTDF values of reported CNECs compare to their observed values. The errors are grouped by the identified GSK strategies that are paired with the respective PTDF calculations, and are presented with their mean relative error (with respect to the observations). One can see that while GSK2 was mapped most frequently to observations, GSK4 delivered a better-quality fit. This shows that involving a base case approach was beneficial to better model the flow-based domain of Austria. A more general claim on its performance cannot be deduced however, as this performance is only represented by one zone (which helps the case of fitting).

Table 7.1 summarizes the error metrics of the calculated PTDFs for all the 6 available months. Besides MAE and MRE, the mean error relative to the standard deviation of the data set is also presented. As argued by Puiu and Hauser (2021), this metric is useful to focus on the quality of the fit. If the deviation in the data set is large, the fitting ability generally reduces, therefore one should penalise the model less. The error values obtained in this thesis are compared to the ones presented in the work of Puiu and Hauser (2021). In their work, they only focused on the Central Western European region, and they rather aimed at completing the available data for the whole flow-based domain in a given MTU, as opposed to delivering a generalized model which can be used for forecasting. They did not focus on GSK strategies, but to improve the data fit given one generalized GSK calculation. Within this context, it is understandable that their work yielded better results regarding fitting PTDFs, but it is important to

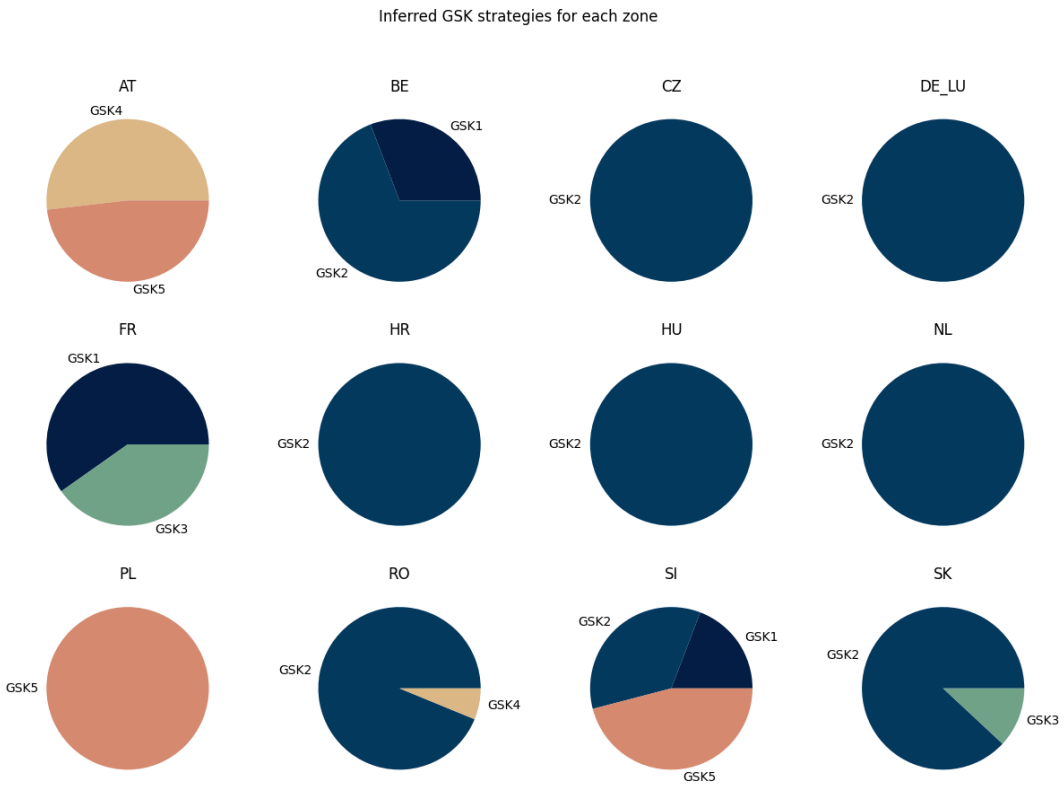

Figure 7.7: Inferred GSK strategies for each zone (with their relative share in case of multiple matches)

	MAE	**MRE**	**mean error relative to std. deviation**
Thesis	0.0161	29.68%	19.62%
Puiu and Hauser (2021)	0.0119	22.10%	14.00%

Table 7.1: Comparison of error metrics of the flow-based domain inference with literature

acknowledge that the results obtained here are comparable to state-of-the-art results.

7.3.2. Reconstructed flow-based domains

This sub-section compares the reconstructed flow-based domains to available observations. The difference between two domains is captured by a distance measure introduced in Chapter 6. These distance metrics are reported for the $\alpha = 10\%$ and $\alpha = 5\%$ CNE selection scenarios, as well as the naive approach. To account for the variable quality of reconstructions for the different borders, the distances are reported on each border for the domains projected onto each respective zone pairs. The overview of the results is presented in Figure 7.11. The scenarios that yield the best results on a give border are highlighted with orange. Among these, the overall best performing reconstruction is denoted with green, while the worst with red.

One can see that the naive approach outperforms the inference in replicating the flow-based domain, in almost all of the cases. Only on the Polish-Czech and Polish-German borders the $\alpha = 5\%$ selection delivers better results, where the day-to-day changes in the observed flow-based domains are considerable. One can get further insights on how these differences geometrically appear by looking at figures 7.13a and 7.13b. It is apparent that $\alpha = 10\%$ can result in a larger domain, surrounding the observed one, while $\alpha = 5\%$ delivers significantly smaller domains, mostly contained by the observed one. Further comments can be made on the shapes of these domains when they are considered together with the resulting cross-border flows, presented in Section 7.4.1.

For now, concluding from the area sizes of the domains, one can say that $\alpha = 5\%$ delivers closer fits

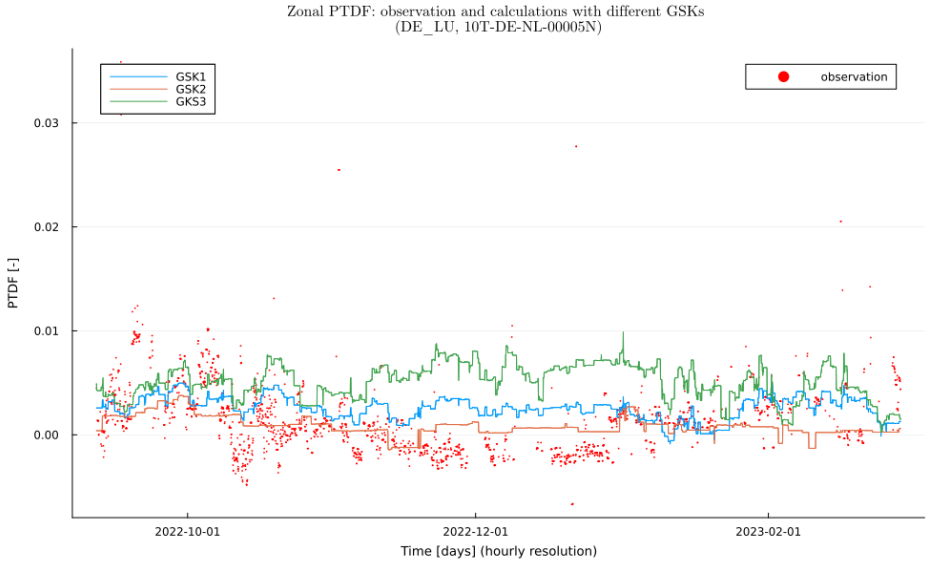

Figure 7.8: Example of observed PTDFs over time for a given CNE between Germany and the Netherlands, along with the calculated zonal PTDF values using different GSK strategies (showing its PTDF for Germany)

because it results in domains closer in size to the observations. The shrinking in domain size comes naturally if one considers the dynamics of CNE selection. With a lower α threshold more lines will be selected, which means more constraints. These linear inequality constraints represented by half-spaces on the 2D plane can slice off more from the polyhedron representing the domain. In the case of the reconstruction, more constraints mean potentially more error-prone PTDF-RAM pairs, where the errors in the calculated RAMs have the most effect on how much the domain is being sliced off. The increased flow-restrictions are not necessarily desirable that comes with the smaller sizes for the $\alpha = 5\%$ selection, as this undermines the original aim of FBMC to flexibly provide as much transfer capacity as possible.

The conflict between having closer resembling domains in size, and the ability to forecast flows actually occurring on the market is further captured when the observed flow inclusion rate is considered. Such an overview is presented in Figure 7.12, showing the inclusion rates for each border, calculated for the validation period. One can immediately perceive the general trend that with $\alpha = 10\%$, almost all occurring flows can be theoretically captured (due to its considerably larger size), while for $\alpha = 5\%$ the ability to capture observed flows is significantly reduced (despite the resulting domains are closer in size to the observed ones). The large variability of the naive approach (0.98 shows for PL-SK, but 0.34 for CZ-DE) shows that while the naive approach overall outperforms the inference, its adequacy cannot be proved in a general sense. There are clearly borders (such as the borders of Germany) where day-to-day changes in the flow-based domains are too large to rely on day-before settings as future estimations.

7.4. Effects of CNE selection on market outcome forecasts

At this stage, the forecast model developed in this thesis is assembled. The calculated and selected zonal PTDFs are used as flow-based constraints in the forward model, which utilises the cost functions reconstructed via inverse optimisation. To uncover the sensitivity of forecasts to selecting the flow-based domain, the results for forecasting cross-border flows and prices are presented, with respect to different CNE selections.

7.4.1. Cross-border flows

The cross-border flow forecast error metrics for the validation period of February 2023 are presented on figures 7.18a and 7.18b. The assigned cross-border flows after the clearing can be directly calculated from the resulting net positions and the zonal PTDFs, by considering the lines that cross the respective borders. Figure 7.18a shows that on average, the smaller the α value is (and hence the more CNEs is

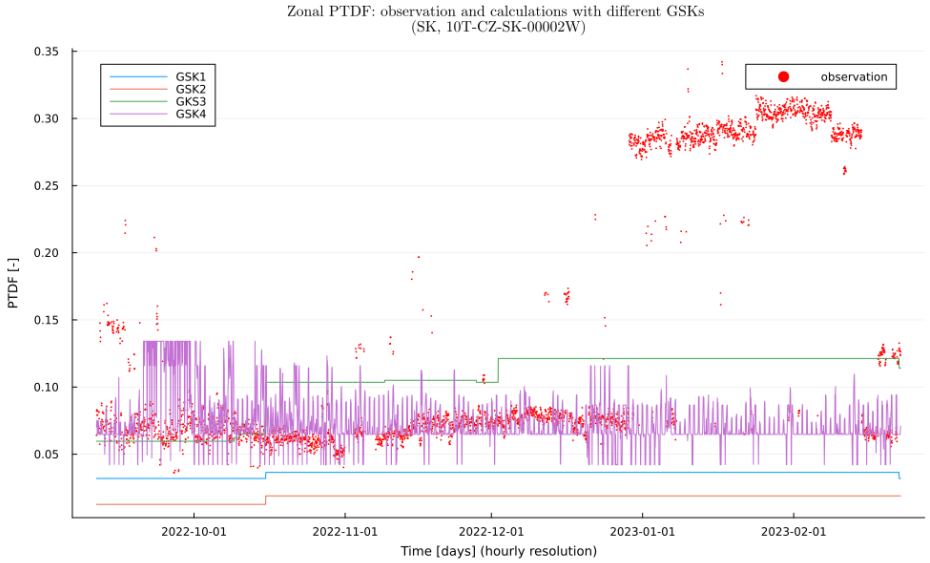

Figure 7.9: Example of observed PTDFs over time for a given CNE between Slovakia and Czech Republic, along with the calculated zonal PTDF values using different GSK strategies (showing its PTDF for Slovakia)

considered), the better the model performs on forecasting cross border flows. Looking at the box plots on Figure 7.18b adds to the picture: the median values of forecast errors are quite similar over different α values, it is rather that a lower α value performs good more consistently (the boxes and the tails are shorter).

This consistency mentioned above can further vary border by border. Figure 7.16 shows that at $\alpha = 10\%$, flows are captured considerable better on the German-Austrian border than for $\alpha = 5\%$. At the same time, $\alpha = 5\%$ performs better in the case of Figure 7.13b, where the flow forecasts are considerably off towards Croatia for $\alpha = 10\%$. One might attempt to capture the difference between these two cases with respect to the constraining flow-based domains. By looking back on figures 7.13a and 7.13b, one can see that the area size difference between the 5% and 10% cases remain consistent. In Figure 7.11, it can also be read that $\alpha = 5\%$ has a lower distance to the observed domain in both cases. The only difference one might notice by looking at the geometries is the relative positioning of the overlapping domains. While overall being closer to the observed domain on Figure 7.13a, the 5% domain seems to shift more to the bottom-quarter, highly constraining flows from Germany to Austria. This can also be observed on time series in Figure 7.16.

The above phenomenon is backed the observed flow inclusion rate, if one looks back on Figure 7.12. Even though the domains generated by $\alpha = 5\%$ are closer in size to the observed domains, they exclude a significant amount of observed flows (82%) that are attempted to be recovered on the validation data. This can be seen on Figure 7.16, where flows from Germany to Austria are barely captured. However, Figure 7.13b shows that the observed flow inclusion rate cannot be ultimately relied on either, when it comes to quantifying the quality of flow-based inference on forecasting cross-border flows. While there $\alpha = 10\%$ has almost full potential (99%) to include all observed flows, the forecasts considerably overshoot towards the direction of Croatia. This can be explained by the elongated part of the larger $\alpha = 10\%$ on Figure 7.13b in the upper half of the net position space.

The above results showcase the limitations of available similarity metrics of flow-based domains to capture the ability of inferred flow-based domains to forecast cross-border flows. It can be concluded that the size of the inferred domain comprises a trade-off between optimality and feasibility towards observed flows. Inferred domains should be sufficiently large enough to be able to recover observed flows, but also sufficiently small for forecasted flow levels to match with observed ones (avoiding overshooting). While the distance metric identified in the literature (Zad et al., 2021) emphasized the optimality part and the newly constructed observed flow inclusion rate emphasized the feasibility part, the two metrics used together could not capture the inherent trade-off between the two aspects.

There are no substantial improvements after further decreasing α from 5%. Simply using the reported TSO selections of network elements yields similar results to the 5% or 4% scenarios. Finally,

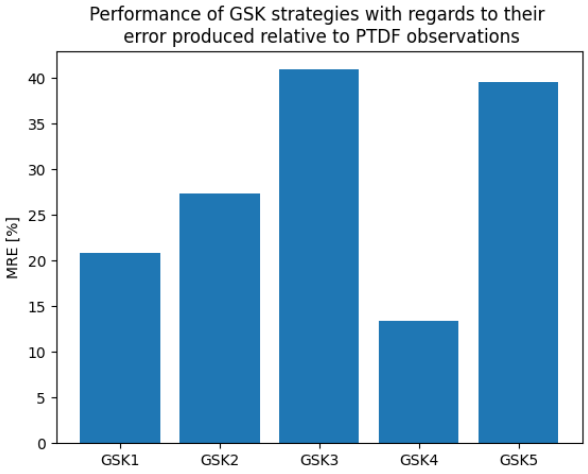

Figure 7.10: Mean relative error of the identified GSK strategies relative to PTDF observations after mapping them to bidding zones

simply taking the naive approach seems to outperform the flow-based domain inference conducted in this thesis. To take the overall picture, both its MAE, median error and error variance are lower than the previously presented scenarios. If one refers back to Figure 7.2, the naive approach delivers the best results for most of the borders. Exceptions lie around Austria (borders with Germany, Croatia and Hungary), on the Czech-Slovakian border, and on the Dutch-Belgian border (although with a highest observed error).

7.4.2. Market prices

The day-ahead price forecast error metrics for the validation period of February 2023 are presented on figures 7.18a and 7.18b. While lower α values lead to generally better flow forecast results, higher α values deliver better price forecast results. Taking a look at the price forecast time series for zones separately (and also by interpreting the box plot in Figure 7.18b), one can put this observation to a more concrete form: with $\alpha = 10\%$ the model performs more consistently good, while with lower α values price volatility is captured substantially better in some zones (Figure 7.19), whereas substantially worse in other zones (Figure 7.20). One could argue that by increasing the number of critical network elements, lines are more likely to be included that are actually present in the published selection and that are significantly affect congestion (and price volatility). This argument does not stay true however, if one compares figures 7.5 and 7.19. It shows that with the $\alpha = 5\%$ inferred domain, the price valleys are better captured than with the originally published flow-based domain. Due to its variability over borders, this performance difference between the different α scenarios is most likely structural in nature, which is yet to be unravelled. Nevertheless, there is apparent value in optimising for which α to use for respective zonal forecasts.

Finally, a comparison to the naive approach is worth mentioning here too. When it comes to price forecast errors, the naive approach does not deliver better results than the inferred domain. With regards to the mean relative errors, it is on the same level as $\alpha = 10\%$. The different performance over zones is true here as well, however. One can see on Figure 7.19 for example, that $\alpha = 5\%$ or 4% clearly outperforms the naive approach.

Given that $\alpha = 10\%$ has the most consistent good performance, the general price forecast analysis will proceed with this selection scenario.

7.5. Comparison to commercial model

In order to position the performance of the developed structural model in this thesis, its delivered day-ahead market price forecasts are compared to the results of a commercially available algorithm. It was accessed via Cross Options (the external party co-hosting this thesis), and the source of the algorithm is kept hidden, due to confidentiality reasons. The commercial model is known to be commonly used within the energy trading industry, and it is reported to model the CORE flow-based market coupling

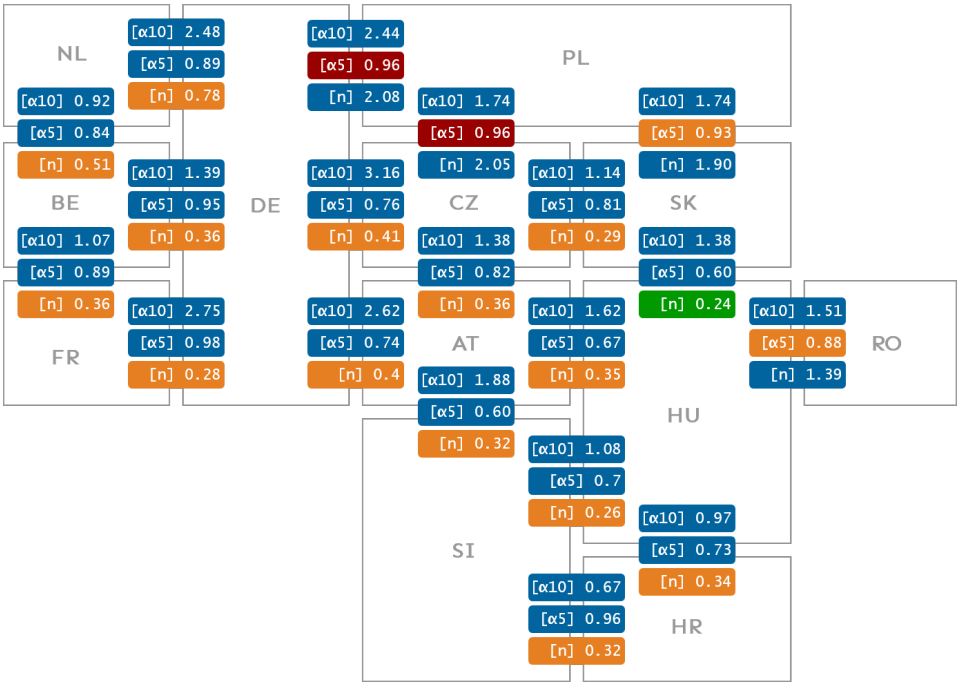

Figure 7.11: Overview of distances between the inferred and observed flow-based domains for each border in the CORE region

with a set of underlying machine learning models (employing random forests and neural networks).

In February 2023 as the considered validation period, the model developed here is showed to outperform the commercial one (albeit with a small margin) on average, as presented on Figure 7.21a. The median of the relative error for the thesis model is around 3%, while the commercial one approaches 0%. The width and the tails of the box plot (Figure 7.21b) for the commercial one are noticeably wider. The latter observation can be attributed to the fact that the machine learning model focuses on capturing price volatility, which sometimes results in substantial overshooting. This overreaction can be observed on Figure 7.22.

Finally, the zonal overview of forecast results, when it is compared to the commercial model is presented in Figure 7.23. The best forecast result is taken for each zone (among $\alpha = 5\%$ and 10%), and is compared to the commercial forecast result available for that zone. A negative value means that the thesis model performed better in those zones, wheres a positive value indicates the better performance of the commercial one. The improvements are the most considerable in the Germany-Luxembourg zone, while also notable in the Netherlands, Belgium, and the Czech Republic. The only countries where the commercial model consistently performs better are Slovenia and Romania.

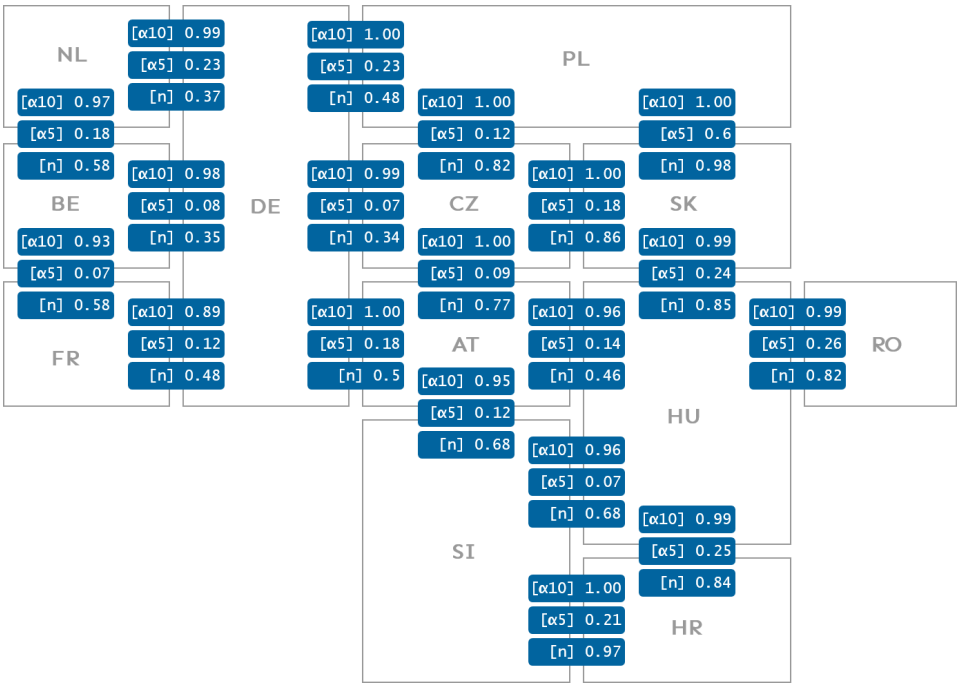

Figure 7.12: Overview of observed flow inclusion rate for each border in the CORE region

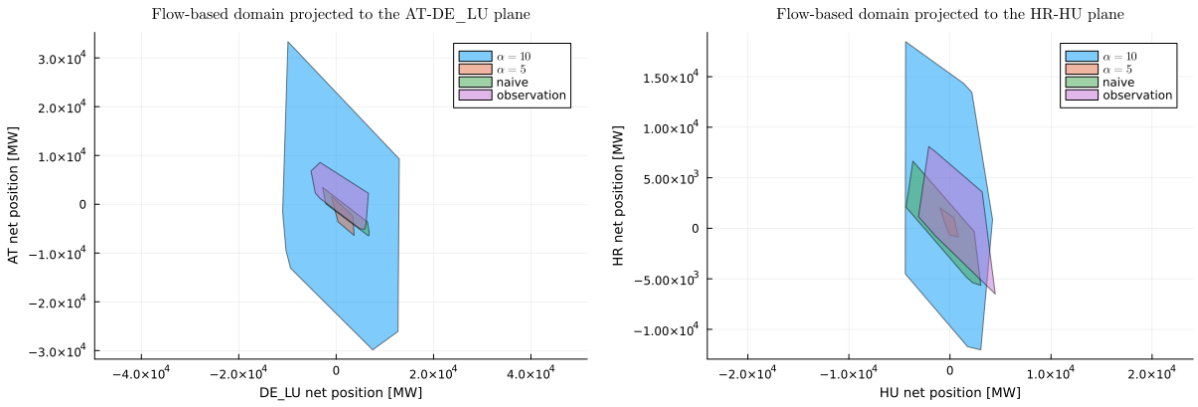

(a) Between Austria and Germany-Luxembourg

(b) Between Croatia and Hungary

Figure 7.13: 2D representation of the flow-based domain across two zones, with respect to their net positions (at a randomly selected MTU in February 2023)

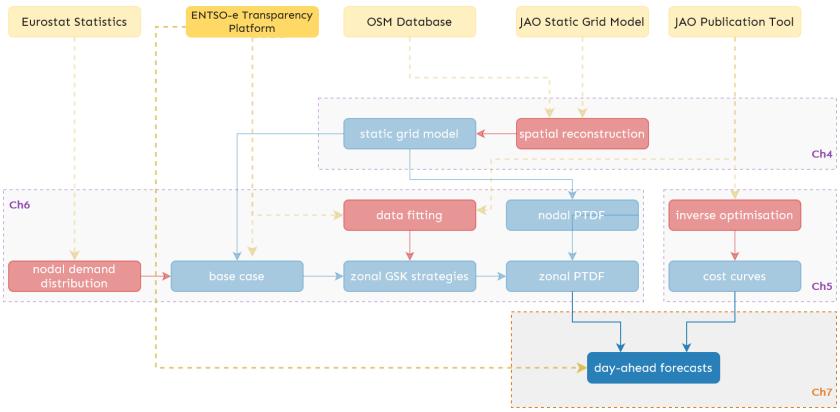

Figure 7.14: Modelling stage of assembly

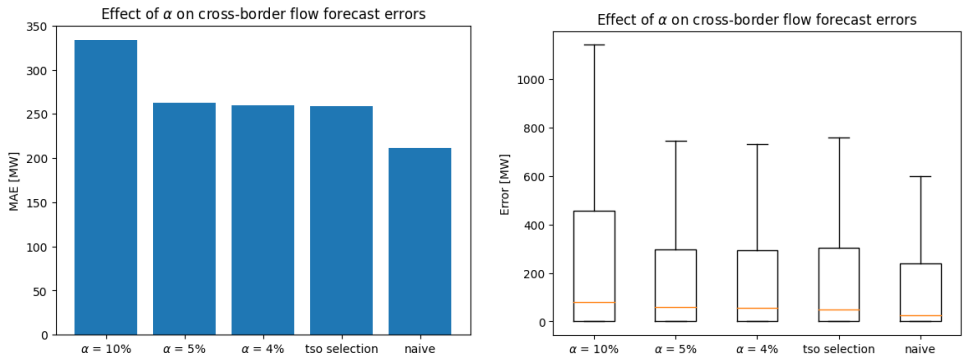

(a) Mean absolute errors

(b) Box plot of errors

Figure 7.15: Comparison of cross-border flow forecast errors for different flow-based domains and CNE selections

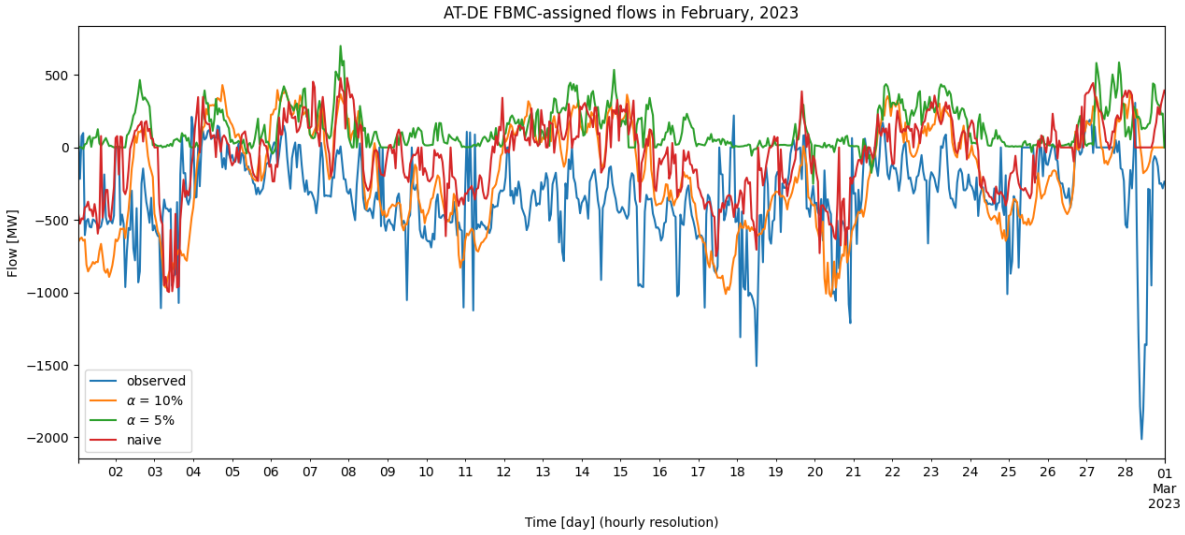

Figure 7.16: Forecasting cross-border flows between Germany and Austria with the inferred and naive flow-based domains

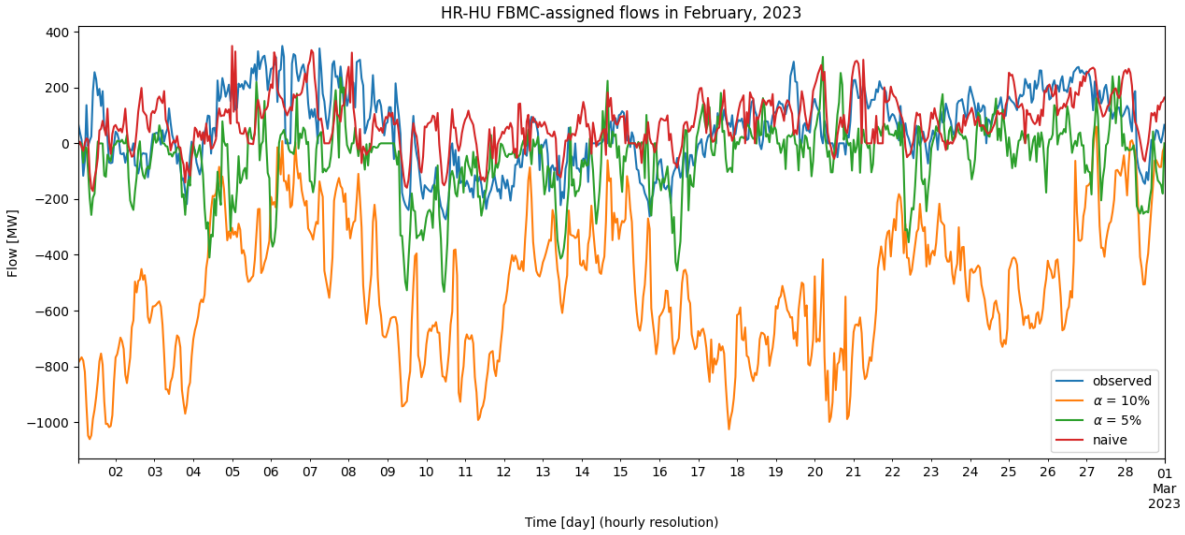

Figure 7.17: Forecasting cross-border flows between Croatia and Hungary with the inferred and naive flow-based domains

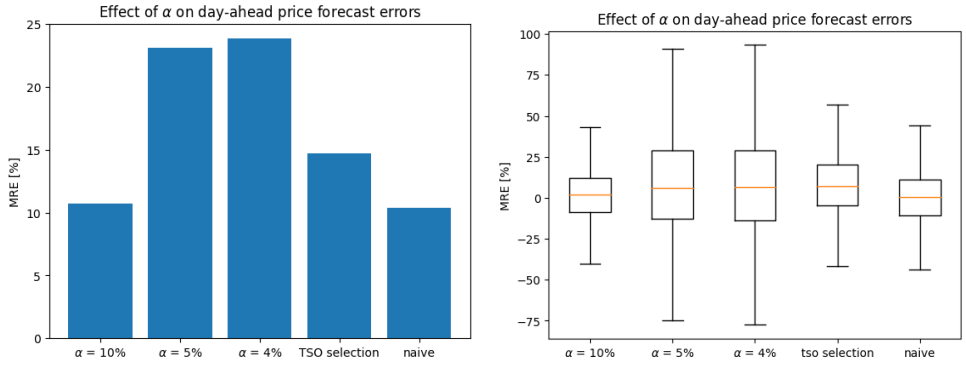

(a) Mean relative errors

(b) Box plot of errors

Figure 7.18: Comparison of day-ahead price forecast errors for different flow-based domains and CNE selections

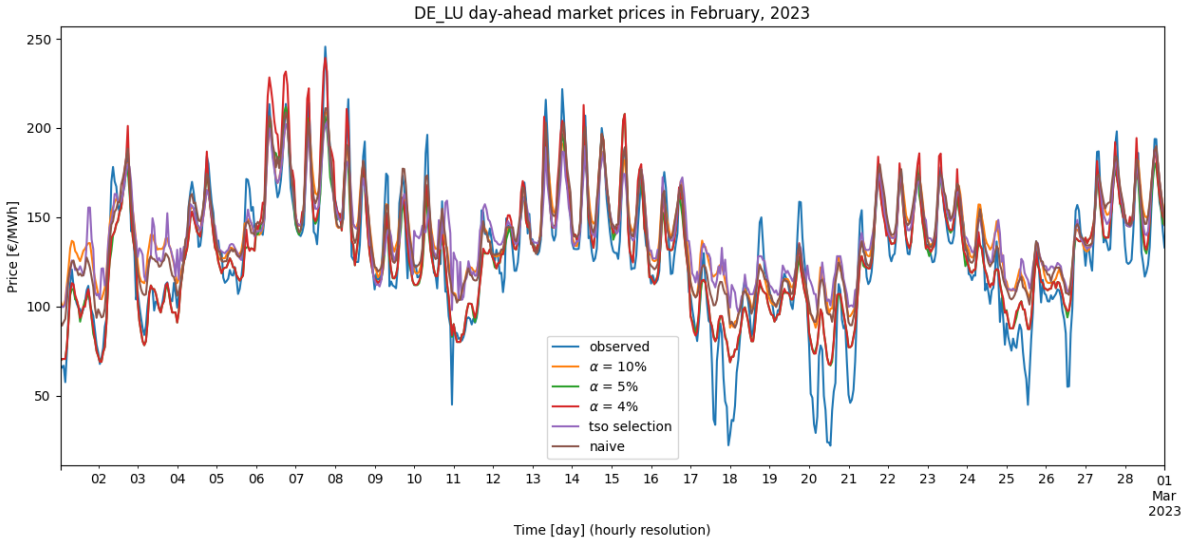

Figure 7.19: Day-ahead market price forecasts in Germany for different CNE selections

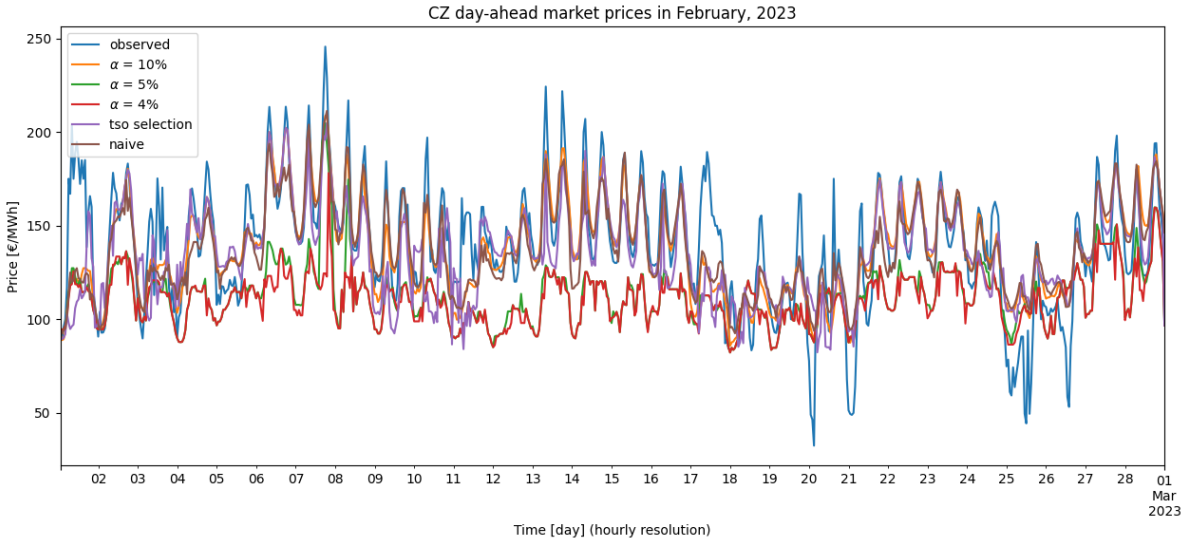

Figure 7.20: Day-ahead market price forecasts in the Czech Republic for different CNE selections

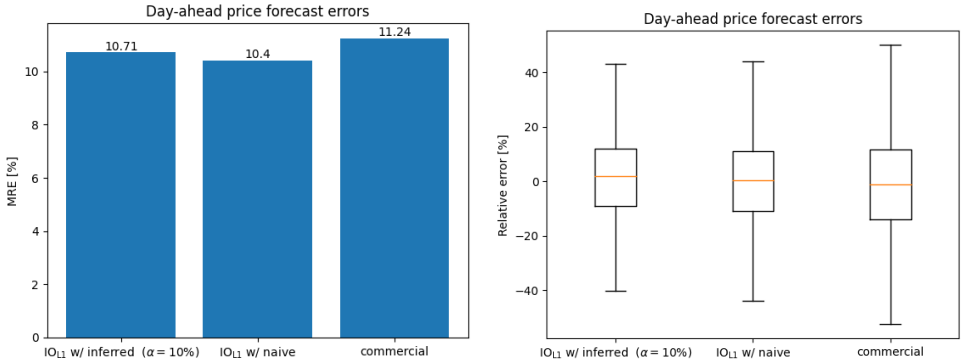

(a) Mean relative errors

(b) Box plot of relative errors

Figure 7.21: Comparison of day-ahead price forecast errors for the model developed in this thesis to a commercial one

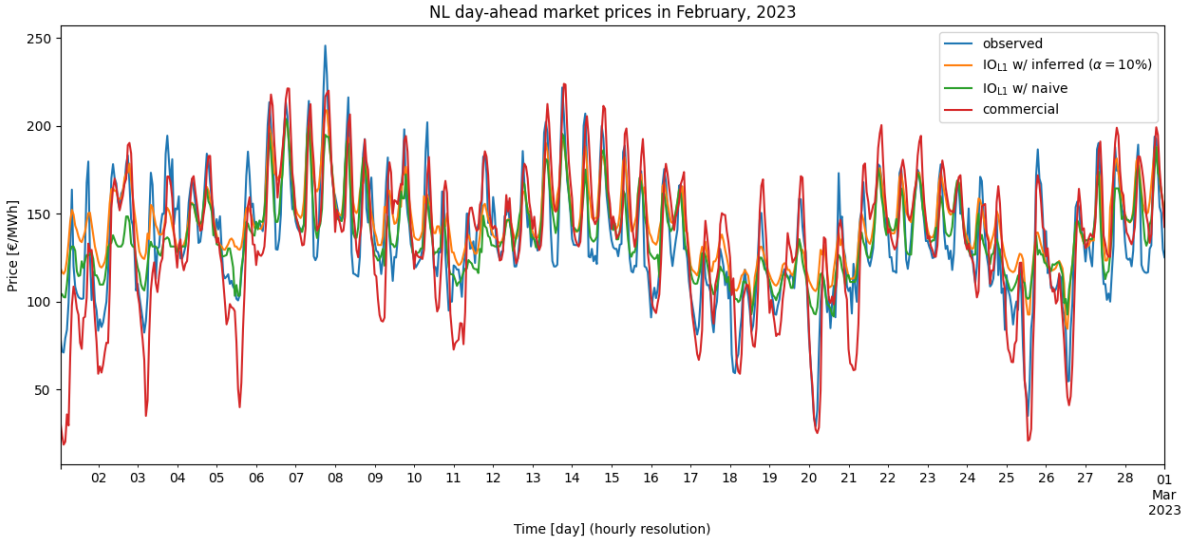

Figure 7.22: Day-ahead market price forecasts in the Netherlands

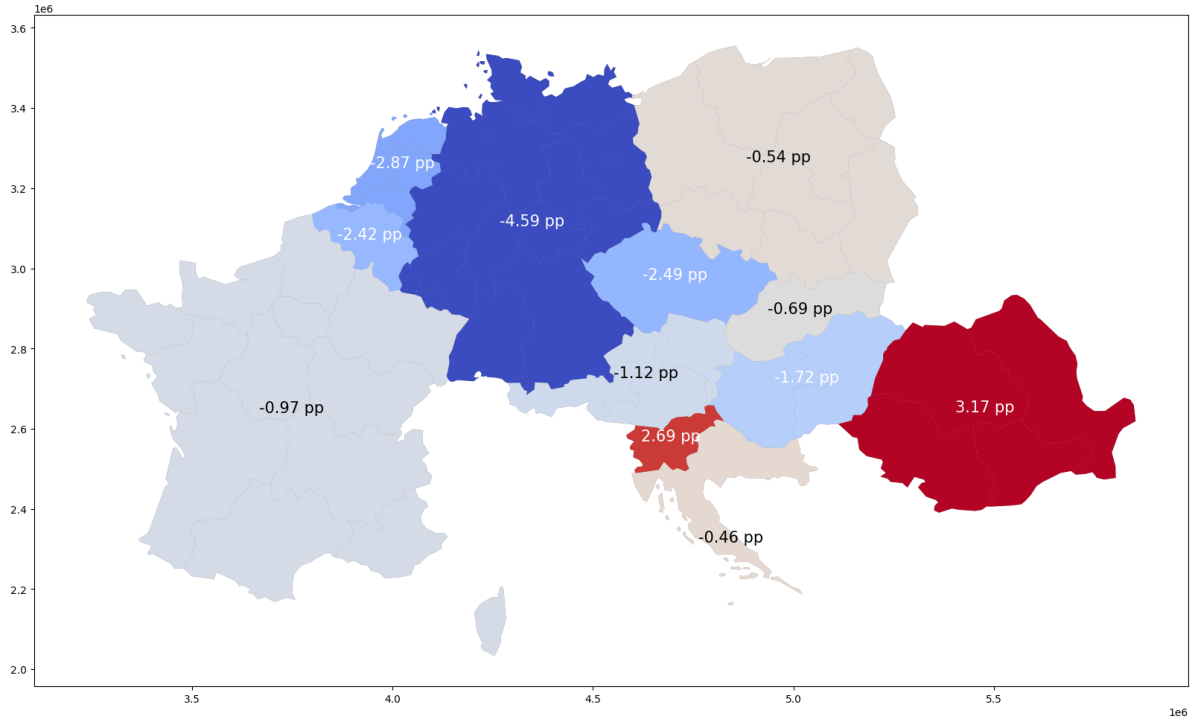

Figure 7.23: Best delivered day-ahead market price forecast errors relative to the commercial ones

8

Discussion

After quantitatively analysing the results, this chapter presents a qualitative discussion on the academic relevance of the products of this thesis, as well as its applicability to Cross Options, the company co-hosting the project. The chapter starts with revisiting the products of the spatial static grid model, serving as a foundation for multiple parts of the modelling process. It proceeds with interpreting the findings on the flow-based domain inference, which is followed by conclusions made on the results of the cost curve reconstruction. The chapters ends with motivating future academic work on further utilising inverse optimisation for capturing real-world electricity markets.

8.1. Static grid model

A spatial reconstruction of the static grid was carried out for the CORE region, in order to be able to match spatially distributed power plants to their closest network nodes (substations). This mapping was necessary for the GSK strategy calculations. Some remarks are made here on the data quality encountered during the reconstruction process.

First of all, after constructing a network from the Static Grid Model compiled by JAO (JAO, 2023), it turned out that it is not completely connected. For larger disconnected components, the connecting lines were looked up and completed from the respective TSO publications, while smaller residual chunks (mainly in France) were simply disposed from the network. The main reason an otherwise extensive, already compiled European grid model (such as in Hörsch et al., 2018) could not have been utilised, is that the transmission lines included there did no have an EIC identifier recorded for them. These identifiers are necessary to pair the network elements with their respective PTDF and RAM observations published by JAO (JAO, 2022b). The JAO Static Grid Model contains these identifiers, there only the spatial reconstruction needed to be conducted for the previously mentioned reasons. When it was attempted to match the observations with the network elements via their EIC identifiers, it turned out that the correspondence is not always unique, however. The same EIC values represent the cross-border tielines on both sides of the border, whereas constraints are published for both sides and directions separately, under the same EIC. Subsections of some network elements with different lengths are also published under the same EIC, so it is not always clear which of the network elements the observed PTDF and RAM values correspond to. It is either advised for ENTSO-e to revise their EIC naming system, or suggested for JAO to set up an unique identifier for their network element registry, which they use both in their Static Grid Model and in the Publication Tool (JAO, 2022b).

Guidelines and reports on the way maximum power capacities for lines are calculated by TSOs are also lacking. Since only voltage ratings and maximum current ratings are published, a calculation is necessary to determine maximum power flows that eventually constrain the exchanges. Based on Tufon et al. (2009), the thermal limit has been used for transmission lines shorter than 100 km, and the stability limit has been used for lines longer than that (as they were introduced in Chapter 6). Calculating the maximum transmission capacities this way has showed to consistently undershoot the observed maximum values of reported CNEs on the JAO Publication Tool by 5-15%. This was eventually one of the reasons that an additional feasibility check needed to be added to the base case runs.

In the end, the spatially reconstructed static grid model presented in this thesis provides a bridge

between unlabelled transmission network maps and publications on the state of the transmission system. Within the thesis, the achieved CORE grid model made it possible to calculate Generation Shift Key strategies for the nodes, to experiment with a nodal base case, and to calculate nodal PTDFs of N-1 outage scenarios. It can serve as a foundation for any further research involving these above-mentioned aspects. The grid model also serves important visualisation purposes for Cross Options: the reported outages, the participation of given lines in cross-border exchanges (given their reported RAMs and PTDFs) can be all presented on the map in a visually intuitive way.

8.2. Flow-based domain inference

The results of the flow-based domain inference exhibited the limited expressiveness of the identified Generation Shift Key strategies. It was shown that correcting the GSKs for the variable availability of production capacity at a given node (due to planned production outages) was indeed a step towards shifts in PTDF values over time. These shifts have been followed to a lesser extent with the calculated GSKs however, indicating either that TSOs incorporate these outages with a larger factor, or that other mutually reinforcing factors have led to such a larger shift, not included in this thesis.

Besides production outages, the base case-driven GSK strategies further include the temporal variability of renewable generation and potentially the changing merit order based on varying fuel prices. While they seem to induce the volatility that is observeable in the published PTDF values, they also fail to capture larger shifts. This potentially indicates the need that planned transmission line outages need to be accounted for as well, and that it is beneficial to explore more sophisticated, regression-based approaches, such as in Schonheit (2019), where several other factors are taken into the GSK calculation. Such factors include the day-ahead prices and net positions for a given zone in recent days, expected renewable yield, etc. The only limitation to fully utilise the latter approach is the need to have unit-level generation reported to conduct the regression, as the response variable for the above-mentioned regressors is the generation level. Such publications are still limited on the ENTSO-e Transparency Platform (ENTSO-E, 2023).

The results of the base case runs have not been analysed separately in the previous chapter. The reason for this is that the obtained RAM values from the generated reference flows were not directly comparable to the observed ones. They still contributed to an adequate flow-based domain reconstruction, but this does not tell much about the performance of the base case, rather about the relative robustness of the resulting flow-based domain to the quality of the RAMs. The only reason the base case calculation was kept in the thesis was the need to retrieve RAM values in a structurally similar way to the rest of the modelling process. Since the base cases of TSOs are reported to be mainly experience-based, there is no justifying evidence the nodal base case is the best, or if it is good approach at all. The structure of the static grid model, the calculated capacities of the lines, the assumed demand distribution (both even and GDP-based) have led to an infeasible problem. The infeasibility lies in the transmission constraints, either the grid connections or their capacities being too restrictive. This infeasibility was resolved by first formulating a feasibility problem, which relaxed the capacity constraints, giving them a slack that was eventually minimised. While this led to a feasible and optimal solution in the end, the slack values inevitable altered flows where line capacities were increased to accommodate the feasibility of the problem.

When looking at the resulting flow-based domains in a geometrical space projected to 2D, one can further reason about the features of these domains relevant for checking their similarity. Taking a distance measure suggested by literature might be adequate to assess clusterisation of domains (Zad et al., 2021), but as shown here, it is unable to capture the ability to replicate cross-border flows. For that purpose, it is more informative to take the distance between the centre of mass of the respective domains. This can be an intuitive revelation, if one considers the effects a domain shift can have. As can be perceived in Figure 7.16, if the domain moves too much to a given quarter, flows in a certain direction are completely restrained. Wheres if one considers a difference only in magnitude, only the magnitude of the flows changes, the trend can still be preserved.

As identified in the previous chapters, immediate improvements can be made to both the price and flow forecasts if on top of the current model, α is optimised for each zone separately. To improve the flow-based domain inference itself, two tracks are suggested. The more ambitious one is to bring the flow-based domain inference as well under the hood of inverse optimisation, which is further described in the next section. The other track is a clusterisation approach, stemming from the results that on

average, the domain inference does not give better results when it is compared to simply reusing the PTDFs from the day before. An amount of PTDF clusters can be identified via Principal Component Analysis (Jolliffe and Cadima, 2016) from the historically published observations, for a given set of exogenous factors. The exogenous factors can comprise weather conditions, outage scenarios, spark spreads (Abadie, 2021) (to account for fuel prices), among others. Based on the current expected values for these exogenous factors, the closest PTDF cluster can be found that will be used for the model run. This clustering method is not necessarily able to account for line outages, although if the outage is planned in advance and spans a long enough period, previous observations during the outage can be reused. A benefit of the clustering approach however, is keeping the N-1 security criterion involved, as historical observations are directly reused. Neither the current inference process, nor a future IO-based approach could contain a contingency selection within reasonable limits. As one has no information on which contingencies TSOs keep at sight, it would be necessary to enumerate all the scenarios for taking one network element out of the equation, which could easily make the problem space explode.

8.3. Cost curve reconstruction

From the results of the cost curve reconstruction, one can draw two conclusions. First, quadratic cost functions aggregated on a power plant technology level seemed to sufficiently capture price volatility. This is only true for the period when exogenous market conditions started to normalize after following the energy crisis. In the second half of November 2022, when an extreme spread between the marginal costs of gas and fuel was observed, even the fuel-adjusted cost curves were not able to capture the unusually large electricity price volatility.

It is important to revisit the fact that while the quadratic curves were inspired by the thermal cycle efficiency of thermal plants (Durvasulu and Hansen, 2018), the obtained curves here do not represent an average characteristic of a plant within the given technology. Instead, they capture the price-level composition of the generation stack in the given bidding zone, for the total amount of installed capacity.

Secondly, an inverse optimisation formulation proved to perform well with generating these cost curves based on observed generation levels and observed shadow prices of the supply-demand equality constraints. The inverse optimisation model achieves two things at the same time to retrieve the cost curves: first it successfully identifies the marginal type in each zone and for each hour, even if a plant (technology) that sets the price in a zone is from another zone. After identifying the marginal type, it pairs the observed shadow price with it, and carries out a regression with the identified volume-price data pairs for a given technology. And here lies the expressiveness of the inverse formulation. While conceptually the formally described process needs is carried out to obtain a cost curve, mathematically it all happens automatically, given an appropriate primal-dual formulation of the problem.

The last remark in this section is a practical observation. It soon became apparent that day-ahead market prices within the CORE FBMC region could not have been captured without also accounting for exports and imports to and from non-FBMC countries, via ATC exchanges. For these exchanges one also needs to now the import prices, therefore if completeness is pursued, one also needs to reconstruct the cost curves for these non-FBMC zones. Once Spain, the UK, North Italy and Switzerland were included in the model, the price levels sufficiently improved within the CORE region. This was especially true for day-ahead prices in France, where Spain were frequently a price setter in the considered time period.

8.4. Implications for stakeholders

The impacts of this thesis are further discussed from the perspective of actors, depicted in Figure 8.1. Blue boxes represent the elements of the technical subsystem, whereas green boxes depict elements of the institutional subsystem. The orange boxes show the relevant stakeholders within the context of the thesis. The arrows depict a direction of influence. Transmission System Operators involves all the TSOs within the CORE region. Market operators are the Nominated Electricity Market Operators (NEMOs) of European power exchanges (such as EPEX SPOT and Nord Pool). Data service providers are the ones reporting data on electricity generation, transmission, market outcomes and on the flow-based domain in a standardized manner (JAO and ENTSO-e in this case). Electricity producers involve companies with production assets (such as the Dutch Eneco or Swedish Vattenfall). Large consumers comprise energy-intensive industrial sites. Energy traders relevant here involve companies

trading either on forward markets or day-ahead markets, being active within or across any of the bidding zones within the CORE region. The category of regulators mainly consists of ACER, the umbrella agency of European energy regulators, setting the agenda of electricity market design (and thus zonal configurations as well).

Regulators, such as ACER can utilise the potential of closing the gap between academic models and real-world market outcomes, as presented in this thesis. Inverse optimisation has proven to be able to deliver artefacts (in this case realistic cost functions) that equips otherwise stylized models to produce results that are realistic for the current electricity system. Such simplistic models have the benefit to run considerably faster than the full-fledged implementation of the EUPHEMIA algorithm, making it possible to run full-factorial experiments on different market scenarios (such as different zonal configurations). Such simulation experiments can shed further light on the potential effect of structural market changes on the current system.

For Transmission System Operators, this work presents a level of understanding of the flow-based domain that market participants can obtain from the publicly available market data. It displays a sensitivity of flow-based domain reconstruction to the prescribed α for CNE selection. The thesis renders the base case as a cornerstone limitation of flow-based domain reconstructions, showing for example, that the more error-prone RAM values are included within the CNE selection, the more the reconstructed domain shrinks compared to the observed one. A nodal base case with limited network data fails to capture the empirical choices TSOs make for their reference flows and RAMs, which calls for more insights on how these calculations take place, or a better standardized calculation scheme.

For electricity producers, large consumers and energy traders (especially ones operating across bidding zones), this work has the potential to deliver better forecasts on day-ahead zonal prices and cross-border flows to better optimise their assets. Currently, its reliability rather lies in delivering price forecasts and their spreads between zones, than cross-border flows. However, its interpretable and fundamental nature makes it useful for speculative analysis: examining how specific transmission line outages, production outages, or other exogenous factors change the market outcomes.

The role of Data Service Providers is to facilitate transparency, an adequate open access to market data. This work is a feedback on how, and to what extent market outcomes can be reconstructed based on their publications. The beginning of this chapter provided some additional suggestions on how to improve the quality of reported data. Although for market operators there are no direct implications within the thesis, they will be potentially affected by better informed traders on the spot market, as well as potential design changes within FBMC, commissioned by regulators and TSOs.

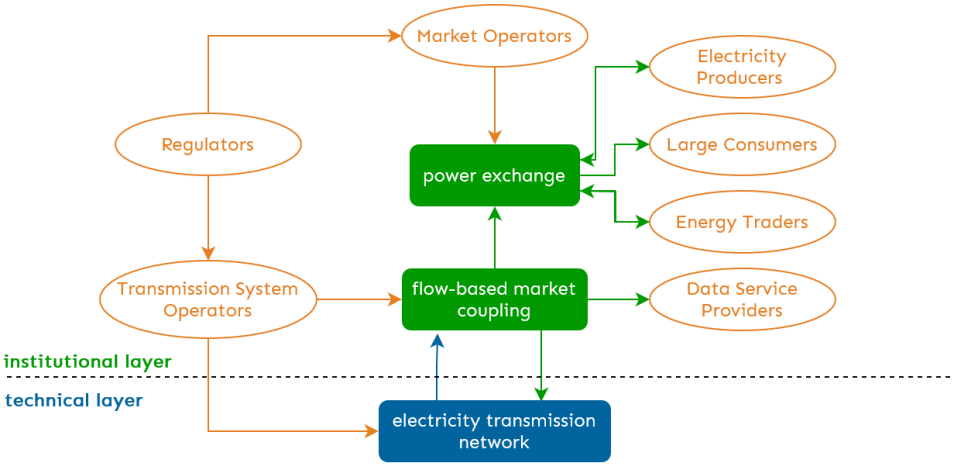

Figure 8.1: The identified socio-technical system evolving around flow-based market coupling

Finally, besides aiding stakeholders separately, the delivered model can act as an object of interaction among the actors within the socio-technical system. When it comes to communication among TSOs and regulators, this interaction can be an informative one. Regulators might use simulation outcomes of the model to advise TSOs on network congestion management. At the same time, TSOs might share their flow-based domain calculation strategies with regulators, that they can subsequently use in their model calculation. Both regulators and TSOs do not have access to cost formation of mar-

ket participants. Therefore, having an approximation about the market cost structure is essential for them to be able to focus on other rules and objectives of the market (network congestion and welfare). When it comes to interaction among market participants, it is rather a strategic one. While participants can fill in their own knowledge in the model (known cost functions of owned assets), they can use the model to reveal cost functions of surrounding competitors using the IO framework. This latter aspect was also explored in Ruiz et al. (2013).

9

Conclusion

Upon identifying a need to better understand real-world electricity market dynamics in the newly formed CORE region, this thesis set out on a journey to model flow-based market coupling in a data-driven, yet interpretable and tractable manner. After reviewing the relevant literature and theory in Chapter 2 and 3, the modelling process started with developing a spatial static grid model in Chapter 4. Following that, a cost curve reconstruction for power plants has been conducted on a zonal and technology level, by exploring different variations of inverse optimisation in Chapter 5. The different stages of flow-based domain inference were discussed in Chapter 6. In chapters 7 and 8 respectively, the results were first quantitatively analysed, then qualitatively interpreted. The report concludes with formally answering the main research question, elaborated by the answers for sub-questions set out in Chapter 1.

First, the main research question is addressed. Within the considered aspects, recovering realistic cost curves proved to be the most crucial part in replicating electricity market dynamics in the newly formed CORE region. It is important to mention here, that recovering the cost structures of bordering non-FBMC zones was equally cardinal in delivering good results. Inverse optimisation was successfully proven to be a suitable tool for recovering such cost curves. When it comes to market price forecasts, the model with well-reconstructed cost curves was shown to be robust against the quality of the inferred flow-based domain. This means that even though the flow-based domain inference under-performed the expectations of replicating the observed domain, the overall performance of the model still remained outstanding for prices. When it comes to forecasting cross-border flows, on average, a naively selected flow-based domain outperforms the GSK-based domain inference implemented in this thesis. The latter results motivate to examine further factors that were not considered in the zonal PTDF calculations (such as transmission outages). Or, as an alternative, to potentially depart from GSK calculations and rather focus on recovering the flow-based domain as a constraint set within the clearing algorithm, similarly via inverse optimisation. The latter approach would lose the tractability of including transmission line outages and examining the effects of specific lines in the physical grid, however.

After formulating an answer to the main research question, the findings are dissected to address each sub-question, starting with **SQ1**. In order to model the flow-based domain, a static grid representation is necessary that has sufficient spatial and physical characteristics recorded. Since there is no database that explicitly maps power plants to grid nodes, the geographical location of network nodes and power plants needs to be known, so that the mapping can be made based on geographical proximity. Based on this spatial correspondence, a relation can be made between nodal power injections and line flows. In order to calculate line flows, physical characteristics of the lines need to be known, including susceptance, voltage level, maximum rated current, resistance and reactance.

An inverse optimisation approach has been proven to be suitable for answering **SQ2**, as it was able to satisfy both stipulations of the research question. First, it is able to incorporate generation level observations to ensure fitting for the observed optimal outcomes. Secondly, day-ahead market prices can also be incorporated in the process, to make sure that with optimal generation outcomes, realistic price levels are paired in the fitting process. Using quadratic cost functions to represent output-based price levels in the objective function has shown the ability to capture price movements in periods where the same technology kept on being marginal.

For answering **SQ3**, first zonal Power Transfer Distribution Factors (PTDFs) were pre-calculated

with different GSK strategies identified in literature. Following that, the available PTDF observations for reported Critical Network Elements (CNEs) were matched with their calculated counterparts, to see which strategy calculation resembles the observations better. This process was carried out by formulating a constrained quadratic optimisation, and a closest matching GSK strategy was determined for each bidding zone. While the calculated PTDF values failed to capture the extent of large shifts in the observations, their trend were still discoverable in the calculated values. These shifts were attributed to production outages. Finally, the fact that most often GSK2 was identified as the closest matching strategy resonates with GSK monitoring studies (JAO et al., 2021). In these studies, multiple TSOs stated that they only include market-based flexible units in their calculations, which corresponds to GSK2. The results further show that it was not possible to gather sufficient information for calculating reference flows that come close to TSO publications.

Qualitatively analysing the inferred domains, it was shown that the CNE selection substantially affects the size and shape of the resulting domain. While the size of the domain intuitively affects the amount of trade that can happen across borders, the shape also turned out to be essential in delivering cross-border flow forecasts. The way the domain spreads across the quadrants of the net-position space affects the directions in which trades can actually take place. While quantifying the performance of the inference in delivering domains resembling observations, it has been found that an adequate metric to capture the similarity of domains with regards to resulting flows is lacking.

In order to answer **SQ4**, an access to a commercially available forecast model was obtained via Cross Options. The source of the algorithm is not disclosed due to confidentiality reasons. The commercial model is known to be commonly used among the energy traders, and it is reported to model the CORE flow-based market coupling with a set of underlying machine learning models (employing random forests and neural networks). In the ability to predict day-ahead market prices, the model developed in this thesis outperformed the commercial one on average, in the selected validation period. The spread of the relative errors is also smaller for the thesis model, which can be attributed to the fact that the machine learning model focuses on increasingly capturing price volatility, which sometimes results in substantial overreaction. It is worthwhile to emphasize again the period in which the proposed methodology was implemented and analysed. At the time of the thesis project, the CORE region had been formed for not more than 6 months, therefore data for training and validation was limited. Extreme gas prices caused by the energy crisis brought the volatility of electricity prices to higher levels than previously observed, making the observed period less representative on the long run.

9.1. Reflections

After evaluating the results and formulating conclusions, reflections are given on the research process itself. The greatest challenge of this project evolved around its data-intensity. The amount, non-consistency and incompleteness of data sources gave rise to a considerable data engineering challenge that was initially not accounted for in the thesis planning. Therefore, the analysis, and hence the results were significantly dependent on the quality of the input data.

In hindsight, a nodal base case should have been omitted from the modelling process, as it posed as the weakest link right from the beginning, with the known limitation that it is only an arbitrary approach to approximate reference flows derived by TSOs. The main consideration for a base case came from the necessity to have reference flows to use in the RAM calculations. However, no analysis has been conducted in the first place on the extent to which incorporating these reference flows actually affect the flow forecasts. This experiment could have been conducted, as both the maximum flow capacities and flow reliability margins are published alongside the RAMs for the reported CNECs, making it possible to calculate a RAM value without reference flows.

Finally, the modelling process has revealed the need for three different perspectives, that all contributed to the success of the project: A mathematically rigorous lens was needed to derive a feasible inverse formulation that can converge to an optimal solution. Secondly, a software engineering lens was necessary to derive a computationally efficient optimisation routine, as the model operated with large and sparse matrices. Finally, the view of a market expert was also essential to know where to look for input data, and to know what kind of assumptions one could make to navigate an overwhelming set of market variables.

9.2. Future work

Finally, directions for future work are provided. This thesis has proven the technique of inverse optimisation to be suitable for bridging the gap between stylized academic models and real-world markets. This was done by reconstructing realistic cost functions that equipped a simple welfare optimisation model with the ability to capture price dynamics in real markets. This was all achieved by maintaining expressiveness and simplicity, without the need to involve black-box approaches. A more in-depth research needs to follow to further explore the potential of inverse optimisation in reconstructing real-world electricity markets in a data-driven albeit still interpretable way. Going beyond imputing the structure of the objective function (cost functions of the decision variables), a number of theoretical works have already been dedicated to IO formulations that aim to retrieve constraint values (Chan and Kaw, 2019, Ghobadi and Mahmoudzadeh, 2021). Such formulations can be explored to potentially obtain the flow-based domain (RAMs and PTDFs) directly via inverse optimisation, and without involving base case and GSK approaches with limited applicability. In direct relation to this future direction, due to the identified limitation of currently available metrics, construction of new similarity metrics for flow-based domains are suggested, to better be able to capture the performance of new inference methods, with regards to recovering cross-border flows.

On a more short-term pathway however, the current methodology of flow-based domain inference is suggested to be revised. From the results, it was suspected that transmission line outages have a large effect on zonal PTDFs, which is still yet to be included in the calculations. The time frame of the research did not allow to test the performance of the advanced statistical GSK approach of Schonheit (2019) in the current model, which still might be valuable to explore.

Once one has well captured market prices and cross-border exchanges, the sensitivity of the system can be tested against structural changes. Such change can be the restructuring of bidding zones: splitting or merging existing markets. Upon running the model for different scenarios, both the resulting social welfare and the extent of cross border exchanges can be directly compared to the current operational settings. Specifically for the geographic restructuring of markets, the spatially reconstructed static grid model will also be practical while nominating new market zones.

References

50Hertz, Amprion, APG, Creos, ČEPS, ELES, Elia, HOPS, MAVIR, PSE, RTE, SEPS, TenneT, Transelectrica, & TransnetBW. (2017). *Explanatory note DA FB CC methodology for Core CCR* (tech. rep.). Retrieved May 30, 2023, from https://consultations.entsoe.eu/markets/core-da-ccm/user_uploads/explanatory-note-for-core-da-fb-cc-public-consultation_fv.pdf

Abadie, L. M. (2021). Current expectations and actual values for the clean spark spread: The case of Spain in the Covid-19 crisis. *Journal of Cleaner Production*, *285*, 124842. https://doi.org/10.1016/j.jclepro.2020.124842

Abadie, L. M., & Chamorro, J. M. (2009). Income risk of EU coal-fired power plants after Kyoto. *Energy Policy*, *37*(12), 5304–5316. https://doi.org/10.1016/j.enpol.2009.07.053

ACER. (2019). Decision No 02/2019 of the Agency for the Cooperation Of Energy Regulators. https://www.acer.europa.eu/Official_documents/Acts_of_the_Agency/Individual%20decisions/ACER%20Decision%2002-2019%20on%20CORE%20CCM.pdf

ACER. (2022). *DECISION No 11/2022 OF THE EUROPEAN UNION AGENCY FOR THE COOPERATION OF ENERGY REGULATORS of 8 August 2022 on the alternative bidding zone configurations to be considered in the bidding zone review process* (tech. rep.).

Bauknecht, D., & Brunekreeft, G. (2008). Chapter 13 - Distributed Generation and the Regulation of Electricity Networks. In F. P. Sioshansi (Ed.), *Competitive Electricity Markets* (pp. 469–497). Elsevier. https://doi.org/10.1016/B978-008047172-3.50017-9

Bichler, M., Buhl, H. U., Knörr, J., Maldonado, F., Schott, P., Waldherr, S., & Weibelzahl, M. (2022). Electricity Markets in a Time of Change: A Call to Arms for Business Research. *Schmalenbach Journal of Business Research*, *74*(1), 77–102. https://doi.org/10.1007/s41471-021-00126-4

Bradley, S., Hax, A. C., & Magnanti, T. (1977). Applied Mathematical Programming. Retrieved June 30, 2023, from https://www.semanticscholar.org/paper/Applied-Mathematical-Programming-Bradley-Hax/8a4ee083b23505df221410e6a2b41fc56fa250a6

Bromley, D. W. (1989). *Economic Interests and Institutions: The Conceptual Foundations of Public Policy*. New York; Oxford: Basil Blackwell. Retrieved June 30, 2023, from https://www.cambridge.org/core/product/identifier/S0266267100003126/type/journal_article

Bundesnetzagentur & Bundeskartellamt. (2021). *Monitoringbericht 2020* (tech. rep.).

Cannan, E., & Pigou, A. C. (1921). The Economics of Welfare. *The Economic Journal*, *31*(122), 206. https://doi.org/10.2307/2222816

Chan, T. C. Y., & Kaw, N. (2019). Inverse optimization for the recovery of constraint parameters [arXiv:1811.00726 [math]]. https://doi.org/10.48550/arXiv.1811.00726

Chatzigiannis, D. I., Dourbois, G. A., Biskas, P. N., & Bakirtzis, A. G. (2016). European day-ahead electricity market clearing model. *Electric Power Systems Research*, *140*, 225–239. https://doi.org/10.1016/j.epsr.2016.06.019

Chen, Y.-K., Koduvere, H., Gunkel, P. A., Kirkerud, J. G., Skytte, K., Ravn, H., & Bolkesjø, T. F. (2020). The role of cross-border power transmission in a renewable-rich power system – A model analysis for Northwestern Europe. *Journal of Environmental Management*, *261*, 110194. https://doi.org/10.1016/j.jenvman.2020.110194

Correljé, A., & Künneke, R. (2021). Institutional design in socio-technical systems (SEN1131) [Course handout].

De Vries, L. J., Correljé, A. F., Knops, P. A., & Van der Veen, R. (2020). Electricity: Market design and policy choices.

Dechezleprêtre, A., Nachtigall, D., & Venmans, F. (2023). The joint impact of the European Union emissions trading system on carbon emissions and economic performance. *Journal of Environmental Economics and Management*, *118*, 102758. https://doi.org/10.1016/j.jeem.2022.102758

de Marcos, R. A., Bello, A., & Reneses, J. (2019). Electricity price forecasting in the short term hybridising fundamental and econometric modelling. *Electric Power Systems Research*, *167*, 240–251. https://doi.org/10.1016/j.epsr.2018.10.034

Durvasulu, V., & Hansen, T. M. (2018). Market□based generator cost functions for power system test cases. *IET Cyber-Physical Systems: Theory & Applications*, *3*(4), 194–205. https://doi.org/10.1049/iet-cps.2018.5046

ENTSO-E. (2023). ENTSO-E Transparency Platform. Retrieved February 16, 2023, from https://transparency.entsoe.eu/

ENTSO-e. (2021). *European Electricity Transmission grids and the Energy Transition* (tech. rep.). https://eepublicdownloads.entsoe.eu/clean-documents/mc-documents/210414_Financeability.pdf

European Commission. (2018). COMMUNICATION FROM THE COMMISSION: A Clean Planet for all. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52018DC0773&from=EN

European Commission. (2020). *Statistical regions in the European Union and partner countries: NUTS and statistical regions 2021 : 2020 edition.* Publications Office. Retrieved June 6, 2023, from https://data.europa.eu/doi/10.2785/850262

Felten, B., Osinski, P., Felling, T., & Weber, C. (2021). The flow-based market coupling domain - Why we can't get it right. *Utilities Policy*, *70*, 101136. https://doi.org/10.1016/j.jup.2020.101136

Finck, R. (2021). Impact of Flow Based Market Coupling on the European Electricity Markets. *Sustainability Management Forum | NachhaltigkeitsManagementForum*, *29*(2), 173–186. https://doi.org/10.1007/s00550-021-00520-w

Finck, R., Ardone, A., & Fichtner, W. (2018). Impact of Flow-Based Market Coupling on Generator Dispatch in CEE Region [ISSN: 2165-4093]. *2018 15th International Conference on the European Energy Market (EEM)*, 1–5. https://doi.org/10.1109/EEM.2018.8469927

Galetto, C. (2022). The Ukrainian Conflict and the Energy Crisis: Sustaining the Energy Transition. Retrieved June 28, 2023, from https://www.iai.it/en/pubblicazioni/ukrainian-conflict-and-energy-crisis-sustaining-energy-transition

Ghobadi, K., & Mahmoudzadeh, H. (2021). Inferring linear feasible regions using inverse optimization. *European Journal of Operational Research*, *290*(3), 829–843. https://doi.org/10.1016/j.ejor.2020.08.048

Glachant, J.-M. (2012). Regulating Networks in the New Economy. *Review of Economics and Institutions*, *3*(1), 27. https://doi.org/10.5202/rei.v3i1.49

Green, R. (1998). Electricity Transmission Pricing: How Much Does it Cost to Get It Wrong? -. *CEEPR*. Retrieved June 16, 2023, from https://ceepr.mit.edu/workingpaper/electricity-transmission-pricing-how-much-does-it-cost-to-get-it-wrong/

Grimm, V., Martin, A., Schmidt, M., Weibelzahl, M., & Zöttl, G. (2016). Transmission and generation investment in electricity markets: The effects of market splitting and network fee regimes. *European Journal of Operational Research*, *254*(2), 493–509. https://doi.org/10.1016/j.ejor.2016.03.044

Hörsch, J., Hofmann, F., Schlachtberger, D., & Brown, T. (2018). PyPSA-Eur: An Open Optimisation Model of the European Transmission System [arXiv:1806.01613 [physics]]. *Energy Strategy Reviews*, *22*, 207–215. https://doi.org/10.1016/j.esr.2018.08.012

Huangfu, Q., & Hall, J. A. J. (2018). Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, *10*(1), 119–142. https://doi.org/10.1007/s12532-017-0130-5

Hurta, A., Žilka, M., & Freiberg, F. (2022). Impact of the splitting of the German–Austrian electricity bidding zone on investment in a grid-scale battery energy storage system deployed for price arbitrage with gray and green power in Austrian and German day-ahead power markets. *Energy Reports*, *8*, 12045–12062. https://doi.org/10.1016/j.egyr.2022.09.045

ICE. (2023). ICE Futures and Options | Product Guide. Retrieved June 28, 2023, from https://www.ice.com/products/81743160/Dutch-Spark-Spread-TTF

Ingham, H. (2023). COVID□19, the Great Recession and Economic Recovery: A Tale of Two Crises. *JCMS: Journal of Common Market Studies*, *61*(2), 469–485. https://doi.org/10.1111/jcms.13383

IPCC. (2022). Climate Change 2022: Impacts, Adaptation and Vulnerability | Climate Change 2022: Impacts, Adaptation and Vulnerability. https://www.ipcc.ch/report/ar6/wg2/

JAO. (2022a). Core FB MC | JAO S.A. Leading service provider for TSOs. Retrieved March 1, 2023, from https://www.jao.eu/core-fb-mc

JAO. (2022b). JAO Publication Tool. Retrieved February 16, 2023, from https://publicationtool.jao.eu/core/

JAO. (2022c). Successful go-live of the Core Flow-Based Market Coupling project. Retrieved March 1, 2023, from https://www.jao.eu/sites/default/files/2022-06/Core%20FB%20MC%20Successful%20Go-live.pdf

JAO. (2023). Static Grid Model. https://www.jao.eu/static-grid-model

JAO, Creos, Tennet, Amprion, RTE, TransnetBW, Elia, 50Hertz, & APG. (2021). *GSK Monitoring Study* (tech. rep.). Retrieved May 20, 2023, from https://www.jao.eu/sites/default/files/2022-04/CWE%20GSK%20Monitoring%20Study.pdf

Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *374*(2065), 20150202. https://doi.org/10.1098/rsta.2015.0202

Konstantelos, I., Pudjianto, D., Strbac, G., De Decker, J., Joseph, P., Flament, A., Kreutzkamp, P., Genoese, F., Rehfeldt, L., Wallasch, A.-K., Gerdes, G., Jafar, M., Yang, Y., Tidemand, N., Jansen, J., Nieuwenhout, F., van der Welle, A., & Veum, K. (2017). Integrated North Sea grids: The costs, the benefits and their distribution between countries. *Energy Policy*, *101*, 28–41. https://doi.org/10.1016/j.enpol.2016.11.024

Kristiansen, T. (2020). The flow based market coupling arrangement in Europe: Implications for traders. *Energy Strategy Reviews*, *27*, 100444. https://doi.org/10.1016/j.esr.2019.100444

Kröger, O., Coffrin, C., Hijazi, H., & Nagarajan, H. (2018). Juniper: An open-source nonlinear branch-and-bound solver in julia. *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 377–386.

KU Leuven, E. I. (2015). Cross-border electricity trading: Towards flow-based market coupling. Retrieved May 10, 2023, from https://set.kuleuven.be/ei/factsheet9/at_download/file

Künneke, R., Groenewegen, J., & Ménard, C. (2010). Aligning modes of organization with technology: Critical transactions in the reform of infrastructures. *Journal of Economic Behavior & Organization*, *75*(3), 494–505. https://doi.org/10.1016/j.jebo.2010.05.009

Lam, L. H., Ilea, V., & Bovo, C. (2018). European day-ahead electricity market coupling: Discussion, modeling, and case study. *Electric Power Systems Research*, *155*, 80–92. https://doi.org/10.1016/j.epsr.2017.10.003

Lovcha, Y., Perez-Laborda, A., & Sikora, I. (2022). The determinants of CO2 prices in the EU emission trading system. *Applied Energy*, *305*, 117903. https://doi.org/10.1016/j.apenergy.2021.117903

Lubin, M., Dowson, O., Garcia, J. D., Huchette, J., Legat, B., & Vielma, J. P. (2023). Jump 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*. https://doi.org/10.1007/s12532-023-00239-3

Lucas, J. M., & Saccucci, M. S. (1990). Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements. *Technometrics*, *32*(1), 1–12. https://doi.org/10.2307/1269835

Martins, J. R. R. A., & Ning, A. (2021). *Engineering Design Optimization* (1st ed.). Cambridge University Press. https://doi.org/10.1017/9781108980647

Matke, C., Medjroubi, W., & Kleinhans, D. (2016). SciGRID - An Open Source Reference Model for the European Transmission Network (v0.2). http://www.scigrid.de

Mohammed, O. O., Mustafa, M. W., Mohammed, D. S. S., & Otuoze, A. O. (2019). Available transfer capability calculation methods: A comprehensive review. *International Transactions on Electrical Energy Systems*, *29*(6). https://doi.org/10.1002/2050-7038.2846

Nemo. (2020). EUPHEMIA Public Description: Single Price Coupling Algorithm.

OpenStreetMap Contributors. (2017). Planet dump retrieved from https://planet.osm.org.

OpenStreetMap Wiki. (2023). Key:substation — openstreetmap wiki [[Online; accessed 1-June-2023]]. %5Curl%7Bhttps://wiki.openstreetmap.org/w/index.php?title=Key:substation&oldid=2460536%7D

Ovaere, M., Kenis, M., Van den Bergh, K., Bruninx, K., & Delarue, E. (2022). The Effect of Flow-Based Market Coupling on Cross-Border Exchange Volumes and Price Convergence in Central-Western European Electricity Markets. https://doi.org/10.2139/ssrn.4059778

Ovaere, M., Kenis, M., Van den Bergh, K., Bruninx, K., & Delarue, E. (2023). The effect of flow-based market coupling on cross-border exchange volumes and price convergence in Central Western European electricity markets. *Energy Economics*, *118*, 106519. https://doi.org/10.1016/j.eneco.2023.106519

Piccialli, V., & Sciandrone, M. (2018). Nonlinear optimization and support vector machines. *4OR*, *16*(2), 111–149. https://doi.org/10.1007/s10288-018-0378-2

Puiu, I. A., & Hauser, R. A. (2021). Principled Data Completion of Network Constraints for Day Ahead Auctions in Power Markets [arXiv:2106.04310 [cs, eess]]. https://doi.org/10.48550/arXiv.2106.04310

Robinius, M., Stein, F. t., Schwane, A., & Stolten, D. (2017). A Top-Down Spatially Resolved Electrical Load Model. *Energies*, *10*(3), 361. https://doi.org/10.3390/en10030361

Ruiz, C., Conejo, A. J., & Bertsimas, D. J. (2013). Revealing Rival Marginal Offer Prices Via Inverse Optimization. *IEEE Transactions on Power Systems*, *28*(3), 3056–3064. https://doi.org/10.1109/TPWRS.2012.2234144

Scholten, D., & Künneke, R. (2016). Towards the Comprehensive Design of Energy Infrastructures. *Sustainability*, *8*(12), 1291. https://doi.org/10.3390/su8121291

Schonheit, D. (2019). An Improved Statistical Approach to Generation Shift Keys: Lessons Learned from an Analysis of the Austrian Control Zone. *Zeitschrift für Energiewirtschaft*, *43*(3), 193–212. https://doi.org/10.1007/s12398-019-00261-w

Schonheit, D., Kenis, M., Lorenz, L., Most, D., Delarue, E., & Bruninx, K. (2021). Toward a fundamental understanding of flow-based market coupling for cross-border electricity trading. *Advances in Applied Energy*, *2*, 100027. https://doi.org/10.1016/j.adapen.2021.100027

Shiri, A., Afshar, M., Rahimi-Kian, A., & Maham, B. (2015). Electricity price forecasting using Support Vector Machines by considering oil and natural gas price impacts. *2015 IEEE International Conference on Smart Energy Grid Engineering (SEGE)*, 1–5. https://doi.org/10.1109/SEGE.2015.7324591

SPOT, E. (2022). European Market Coupling. Retrieved June 6, 2023, from https://www.epexspot.com/en/marketcoupling

Sun, S., Cao, Z., Zhu, H., & Zhao, J. (2019). A Survey of Optimization Methods from a Machine Learning Perspective [arXiv:1906.06821 [cs, math, stat]]. https://doi.org/10.48550/arXiv.1906.06821

Tennet. (2015). Flow-Based Methodology for CWE Market Coupling successfully launched. Retrieved March 1, 2023, from https://netztransparenz.tennet.eu/tinyurl-storage/detail/flow-based-methodology-for-cwe-market-coupling-successfully-launched/

Tufon, C., R, A., Kirby, B., Kueck, J., & Li, F. (2009). A tariff for reactive power, 1–7. https://doi.org/10.1109/PSCE.2009.4839932

Van den Bergh, K., & Delarue, E. (2014). DC power flow in unit commitment models. Retrieved February 16, 2023, from https://www.semanticscholar.org/paper/DC-power-ow-in-unit-commitment-models-Bergh-Delarue/a7e21c732536539f9a14f9c33ccb2ee5a864740d

Williamson, O. E. (2000). The New Institutional Economics: Taking Stock, Looking Ahead. *Journal of Economic Literature*, *38*(3), 595–613. Retrieved June 15, 2023, from https://www.jstor.org/stable/2565421

Zad, B. B., Toubeau, J.-F., Vatandoust, B., Bruninx, K., Grève, Z. D., & Vallée, F. (2021). Enhanced integration of flow-based market coupling in short-term adequacy assessment. *Electric Power Systems Research*, *201*, 107507. https://doi.org/10.1016/j.epsr.2021.107507

Zhao, Z., Wang, C., Nokleby, M., & Miller, C. (2017). Improving Short-Term Electricity Price Forecasting Using Day-Ahead LMP with ARIMA Models [arXiv:1801.02485 [eess]]. *2017 IEEE Power & Energy Society General Meeting*, 1–5. https://doi.org/10.1109/PESGM.2017.8274124

A

Inverse Optimisation

Below is an excerpt of the inverse optimisation source code for cost curve reconstruction, written in Julia, using the JuMP.jl optimisation library (Lubin et al., 2023), and by utilising the HiGHS linear solver (Huangfu and Hall, 2018).

```
using JuMP, HiGHS
using DataFrames, XLSX
using LinearAlgebra
using Statistics
using SparseArrays
using StatsBase

# setting up the solver engine
model = Model(HiGHS.Optimizer)
set_optimizer_attribute(model, "presolve", "on")
set_optimizer_attribute(model, "time_limit", 180.0)

# cost curve variables (coefficients)
@variable(model, c[1:(num_z+num_z_non_fbmc)*num_tech*num_t] >= 0)
@variable(model, alpha[1:(num_z+num_z_non_fbmc)*num_tech] >= 0)
@variable(model, beta[1:(num_z+num_z_non_fbmc)*num_tech] >= 0)
@variable(model, gamma[1:(num_z+num_z_non_fbmc)*num_tech] >= 0)

# quadratic curve constraints for different technologies (with different fuel)
for z in 1:(num_z+num_z_non_fbmc)
    for tech in 3:num_tech
        if tech == 7
            @constraint(model, gamma[num_tech*(z-1)+tech] + gamma[num_tech*(z-1)+tech] == 0)
        end
    end
    for t in 1:num_t
        for tech in 1:num_tech
            if tech == 1 || tech == 8 || tech == 9 || tech == 10 # without fuel-
                responsiveness
                @constraint(model, c[num_t*num_tech*(z-1)+num_t*(tech-1)+t] == alpha[num_tech
                    *(z-1)+tech] + (beta[num_tech*(z-1)+tech] * g_obs[num_t*num_tech*(z-1)+
                    num_t*(tech-1)+t] + gamma[num_tech*(z-1)+tech] * g_obs[num_t*num_tech*(z
                    -1)+num_t*(tech-1)+t]^2))
            elseif tech == 2 || tech == 3 || tech == 5 # coal
                @constraint(model, c[num_t*num_tech*(z-1)+num_t*(tech-1)+t] == coal_prices[t
                    ]/10 * (alpha[num_tech*(z-1)+tech] + beta[num_tech*(z-1)+tech] * g_obs[
                    num_t*num_tech*(z-1)+num_t*(tech-1)+t] + gamma[num_tech*(z-1)+tech] *
                    g_obs[num_t*num_tech*(z-1)+num_t*(tech-1)+t]^2))
            elseif tech == 4 # gas
                @constraint(model, c[num_t*num_tech*(z-1)+num_t*(tech-1)+t] == gas_prices[t
                    ]/10 * (alpha[num_tech*(z-1)+tech] + beta[num_tech*(z-1)+tech] * g_obs[
                    num_t*num_tech*(z-1)+num_t*(tech-1)+t] + gamma[num_tech*(z-1)+tech] *
                    g_obs[num_t*num_tech*(z-1)+num_t*(tech-1)+t]^2))
            elseif tech == 6 # oil
                @constraint(model, c[num_t*num_tech*(z-1)+num_t*(tech-1)+t] == oil_prices[t
                    ]/10 * (alpha[num_tech*(z-1)+tech] + beta[num_tech*(z-1)+tech] * g_obs[
```

```
                    num_t*num_tech*(z-1)+num_t*(tech-1)+t] + gamma[num_tech*(z-1)+tech] *
                    g_obs[num_t*num_tech*(z-1)+num_t*(tech-1)+t]^2))
            elseif tech == 7 # hydro
                @constraint(model, c[num_t*num_tech*(z-1)+num_t*(tech-1)+t] == alpha[num_tech
                    *(z-1)+tech] + beta[num_tech*(z-1)+tech] * g_obs[num_t*num_tech*(z-1)+
                    num_t*(tech-1)+t])
            end
        end
    end
end

# for L norms, replacing tha absolute expression in the objective
@variable(model, eps1[1:(num_z+num_z_non_fbmc)*num_t] >= 0)
@variable(model, eps2[1:(num_z+num_z_non_fbmc)*num_t] >= 0)

@variable(model, g[1:(num_z+num_z_non_fbmc)*num_tech*num_t] >= 0)
@variable(model, np[1:num_z*num_t])
@variable(model, atc_ex_1[1:num_atc_border*num_t] >= 0)
@variable(model, atc_ex_2[1:num_atc_border*num_t] >= 0)

# dual variables
@variable(model, lambda[1:(num_z+num_z_non_fbmc)*num_t])
@variable(model, lambda_exchange[1:num_t])

@variable(model, mu_gen[1:(num_z+num_z_non_fbmc)*num_tech*num_t] <= 0)
@variable(model, mu_exchange[1:(num_j*num_t+2*num_atc_border*num_t)] <= 0)

# for minimasing duality gap
@variable(model, epsilon_duality_abs)
@variable(model, epsilon_duality)

@constraint(model, g .== g_obs) # observe generation levels

# CONSTRUCTING MATRICES FOR THE INVERSE FORMULATION

A_balance = spzeros((num_z+num_z_non_fbmc)*num_t, (num_z+num_z_non_fbmc)*num_tech*num_t+num_z
    *num_t+2*num_atc_border*num_t) # contains g and np and atc (e/i)
prev_pos = (num_z+num_z_non_fbmc)*num_tech*num_t
for z in 1:(num_z+num_z_non_fbmc)
    for t in 1:num_t
        for tech in 1:num_tech
            A_balance[num_t*(z-1)+t, num_t*num_tech*(z-1)+num_t*(tech-1)+t] = 1
        end
    end
end

for z in 1:num_z
    for t in 1:num_t
        A_balance[num_t*(z-1)+t, prev_pos+num_t*(z-1)+t] = -1 # np
    end
end

prev_pos += num_z*num_t
for z in 1:(num_z+num_z_non_fbmc)
    for b in 1:num_atc_border
        for t in 1:num_t
            A_balance[num_t*(z-1)+t, prev_pos+num_t*(b-1)+t] = has_interconnector_ex_1(z, b)
                # atc exchange direction 1
        end
    end
end

prev_pos += num_atc_border*num_t
for z in 1:(num_z+num_z_non_fbmc)
    for b in 1:num_atc_border
        for t in 1:num_t
            A_balance[num_t*(z-1)+t, prev_pos+num_t*(b-1)+t] = has_interconnector_ex_2(z, b)
                # atc exchange direction 2
        end
    end
end
```

```
100
101 b1_balance = demand - ren_gen
102
103 B_gen = sparse(cat(Matrix(I, (num_z+num_z_non_fbmc)*num_tech*num_t, (num_z+num_z_non_fbmc)*
        num_tech*num_t), spzeros((num_z+num_z_non_fbmc)*num_tech*num_t, num_z*num_t + 2*
        num_atc_border*num_t); dims=(2)))
104 b2_gen = g_max_t
105
106 A_exchange = spzeros(num_t, (num_z+num_z_non_fbmc)*num_tech*num_t+num_z*num_t+2*
        num_atc_border*num_t)
107
108 prev_pos = (num_z+num_z_non_fbmc)*num_tech*num_t
109 for t in 1:num_t
110     for z in 1:num_z
111         A_exchange[t, prev_pos+num_t*(z-1)+t] = 1
112     end
113 end
114
115 b1_exchange = spzeros(num_t)
116
117 B_exchange_temp = spzeros(num_j*num_t, num_z*num_t)
118 ram = spzeros(num_j*num_t)
119 for t in 1:num_t
120     df_ptdf_h = df_ptdf[df_ptdf.DateTime .== timestamps[num_t_passed+t], :]
121     for j in 1:size(df_ptdf_h)[1]
122         for z in 1:num_z
123             B_exchange_temp[num_t*(j-1)+t, num_t*(z-1)+t] = df_ptdf_h[j, fbmc_zones[z]]
124             ram[num_t*(j-1)+t] = df_ptdf_h[j, :ram]
125         end
126     end
127 end
128
129 ram = convert(Vector{Float64}, ram)
130 B_exchange_ptdf = sparse(cat(spzeros(num_j*num_t, (num_z+num_z_non_fbmc)*num_tech*num_t),
        B_exchange_temp, spzeros(num_j*num_t, 2*num_atc_border*num_t); dims=(2)))
131
132 B_exchange_atc = spzeros(2*num_atc_border*num_t, (num_z+num_z_non_fbmc)*num_tech*num_t +
        num_z*num_t + 2*num_atc_border*num_t)
133 prev_pos = (num_z+num_z_non_fbmc)*num_tech*num_t + num_z*num_t
134
135 atc_1 = spzeros(num_atc_border*num_t)
136 for b in 1:num_atc_border
137     for t in 1:num_t
138         B_exchange_atc[num_t*(b-1)+t, prev_pos+num_t*(b-1)+t] = 1 # atc exchange direction 1
139         atc_1[num_t*(b-1)+t] = get_atc_ex_1(b, t)
140     end
141 end
142
143 prev_pos += num_atc_border*num_t
144
145 atc_2 = spzeros(num_atc_border*num_t)
146 for b in 1:num_atc_border
147     for t in 1:num_t
148         B_exchange_atc[num_t*(b-1)+t, prev_pos+num_t*(b-1)+t] = 1 # atc exchange direction 2
149         atc_2[num_t*(b-1)+t] = get_atc_ex_2(b, t)
150     end
151 end
152
153 b2_exchange = vcat(ram, remove_missing(atc_1), remove_missing(atc_2))
154 B_exchange = sparse(cat(B_exchange_ptdf, B_exchange_atc; dims=(1)))
155
156 # primal constraints
157 @constraint(model, balance, A_balance * vcat(g, np) .== b1_balance) # demand equality
158 @constraint(model, B_exchange * vcat(g, np, atc_ex_1, atc_ex_2) .<= b2_exchange) # net
        position, rams and ptdfs
159 @constraint(model, sum_z_np(np, num_t) .== 0)
160
161 # dual constraints
162 @constraint(model, cat(A_balance, A_exchange; dims=(1))' * vcat(lambda, lambda_exchange) .+
        cat(B_gen, B_exchange; dims=(1))' * vcat(mu_gen, mu_exchange) .== vcat(c, spzeros(num_z*
        num_t), spzeros(2*num_atc_border*num_t)))
```

```
163
164 # relaxed duality
165 @constraint(model, sum(c) .- b1_balance' * lambda .- b1_exchange' * lambda_exchange .- b2_gen
        ' * mu_gen .- b2_exchange' * mu_exchange == epsilon_duality_abs)
166 #@constraint(model, epsilon_duality_abs == 0) # strong duality theorem, for L norm only
167
168 @constraint(model, epsilon_duality_abs <= epsilon_duality)
169 @constraint(model, -1*epsilon_duality_abs <= epsilon_duality)
170 @constraint(model, lambda .- lambda_obs .== eps1 .- eps2) # for L1 norm
171
172 u = ones((num_z+num_z_non_fbmc)*num_t)
173
174 #@objective(model, Min, sum((lambda - lambda_obs) .^ 2)) # L2 norm only
175 #@objective(model, Min, eps1' * u + eps2' * u) # L1 norm only
176 @objective(model, Min, eps1' * u + eps2' * u + epsilon_duality) # L1 norm and duality gap
        minimisation
177 #@objective(model, Min, sum((lambda - lambda_obs) .^ 2) + epsilon_duality) # L2 norm and
        duality gap minimisation
178
179 optimize!(model)
```

B

GSK Strategy Inference

Below is an excerpt of the GSK strategy inference source code for fitting the observed zonal PTDFs, written in Julia, using the JuMP.jl optimisation library (Lubin et al., 2023), and by utilising the HiGHS linear solver (Huangfu and Hall, 2018) and MINLP solver Juniper (Kröger et al., 2018).

```
using JuMP, Ipopt, Juniper

# initializing MINLP solver
ipopt = optimizer_with_attributes(Ipopt.Optimizer)
minlp_solver = optimizer_with_attributes(Juniper.Optimizer, "nl_solver" => ipopt)
model = Model(minlp_solver)

# variable for GSK selection per zone
@variable(model, sigma[1:num_gsk_strategy, 1:12] >= 0)

# relative selection preference summing up to 1
@constraint(model, sum(sigma, dims=1) .== 1)

# simple plant to node map
M = zeros(size(O)[2], size(O)[1])
X = zeros(size(O)[3], size(GSK1_P)[2])
X[1, :] .= 1

for p in 1:size(GSK1_P)[2]
    M[:, p] = O[p, :, :] * X[:, p]
end

objective_function = 0

s = 1
# gathering CNEC observations for fitting
for cnec in eachrow(df_cnecs)
    cnec_zone_i = findall(zones .== cnec.zone)[1] - 2

    for obs in eachrow(df_ptdf[(df_ptdf.line_id .== cnec.line_id) .& (df_ptdf.contingency .==
         cnec.contingency) .& (df_ptdf.DateTime .>= start_date) .& (df_ptdf.DateTime .<
        end_date), :])
        edge = df_line_edge_map[df_line_edge_map.line_id .== cnec.line_id, :edge][1]
        t = findfirst(==(obs.DateTime), df_timestamps[(df_timestamps.DateTime .>= start_date)
             .& (df_timestamps.DateTime .< end_date), :].DateTime)

        # calculating zonal PTDFs for all GSK strategies
        PTDF_Z = [
            PTDF_N_C[cnec.contingency + 1] * M * GSK1_P[t, :, :];;;
            PTDF_N_C[cnec.contingency + 1] * M * GSK2_P[t, :, :];;;
            PTDF_N_C[cnec.contingency + 1] * M * GSK3_P[t, :, :];;;
            PTDF_N_C[cnec.contingency + 1] * M * GSK4_P[t, :, :];;;
            PTDF_N_C[cnec.contingency + 1] * M * GSK5_P[t, :, :];;;
        ]

        objective_function += sum((
```

```
            [values(obs[5:18])...] - [sigma[:, cnec_zone_i]' * PTDF_Z[edge + 1, i, :] for i =
                1:14]
        ) .^ 2) # L2-norm
    end

    s += 1
end

@NLobjective(model, Min, objective_function)
optimize!(model)

JuMP.value.(sigma)
```