

**PREDICT-THEN-OPTIMIZE VS. CONSTRAINT LEARNING:
AN EMPIRICAL COMPARISON FOR E-COMMERCE
DECISION-MAKING**

AHMET MELİK AKSOY

ÖZYEĞİN UNIVERSITY

AUGUST, 2025

**PREDICT-THEN-OPTIMIZE VS. CONSTRAINT LEARNING:
AN EMPIRICAL COMPARISON FOR E-COMMERCE
DECISION-MAKING**

by
Ahmet Melik Aksoy

Thesis
Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Science

in
Data Science

Advisor: Asst. Prof. Erinc Albey

Graduate School of Science and Engineering
Özyeğin University
İstanbul

August, 2025

PREDICT-THEN-OPTIMIZE VS. CONSTRAINT LEARNING: AN EMPIRICAL COMPARISON FOR E-COMMERCE DECISION-MAKING

Approved by:

Asst. Prof. Erinç Albey, Advisor
Department of Artificial Intelligence and Data
Science
Özyeğin University

Asst. Prof. Mehmet Önal
Department of Industrial Engineering
Özyeğin University

Prof. Mehmet Güray Güler
Department of Data Science and Analytics
Istanbul Technical University

Approval Date: August 5, 2025

To my wife and daughter...



DECLARATION OF ORIGINALITY

I hereby declare that I am the sole author of this thesis and that this is the true copy of my thesis, including the final revisions, approved by my thesis committee. All data and information have been obtained, produced, and presented in accordance with the rules of research ethics and principles of academic honesty. As required by these rules, to the best of my knowledge I have acknowledged ideas, thoughts, and any copyrighted material in accordance with the standard referencing rules. I certify that any part of this thesis has not been submitted for a degree or diploma in another educational institution. *Signature*

Ahmet Melik Aksoy

ABSTRACT

In today's rapidly evolving e-commerce landscape, effectively balancing price and commission strategies is vital for platforms aiming to maximize both merchant profitability and their own revenue. This MSc. thesis introduces an innovative approach by integrating advanced machine learning techniques with optimization frameworks, specifically comparing the traditional Predict-then-Optimize (PTO) method against the integrated Constraint Learning (CL) paradigm. Through rigorous simulation and empirical evaluation using real-world marketplace data, the proposed frameworks demonstrate clear benefits in strategic decision-making, efficiency, and robustness. Ultimately, this research paves the way for smarter, more dynamic pricing decisions in complex online retail ecosystems.

Keywords: Machine Learning, Constraint Learning, Predict-then-Optimize, Nash Bargaining, E-commerce

ÖZET

Günümüzün hızla değişen e-ticaret ortamında, fiyatlandırma ve komisyon stratejilerinin etkin biçimde dengelenmesi, hem platformların gelirini hem de satıcıların kârlılığını maksimize etmek açısından kritik önem taşımaktadır. Bu yüksek lisans tezi, ileri makine öğrenimi tekniklerini matematiksel optimizasyon yöntemleriyle bütünleştiren yenilikçi bir yaklaşım sunmaktadır. Bu kapsamda, geleneksel Tahmin Sonrası Optimizasyon (Predict-then-Optimize - PTO) yaklaşımı ile daha bütünleşik bir yöntem olan Kısıt Öğrenme (Constraint Learning - CL) paradigması karşılaştırılmaktadır. Gerçek pazar verileri üzerinden gerçekleştirilen kapsamlı simülasyon ve ampirik değerlendirme sonuçları, önerilen yöntemlerin stratejik karar verme, verimlilik ve dayanıklılık açısından açık avantajlar sağladığını ortaya koymaktadır. Sonuç olarak bu çalışma, karmaşık e-ticaret pazarlarında daha akıllı ve dinamik fiyatlandırma kararlarının önünü açmaktadır.

Anahtar Kelimeler: Makine Öğrenmesi, Kısıt Öğrenme, Tahmin Sonrası Optimizasyon, Nash Pazarlık Modeli, E-ticaret

ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to the Scientific and Technological Research Council of Turkey (TÜBİTAK) for providing financial support through the BİDEB (Bilim İnsanı Destek Programları Başkanlığı) scholarship program (BİDEB 2210-A Genel Yurt İçi Lisansüstü Burs Programı). This scholarship has enabled me to pursue my Master's degree and conduct this research. I would also like to extend my thanks to Hepsiburada, one of the leading e-commerce companies in Turkey, for providing the necessary support for this study. Without their support, this research would not have been possible.

Most importantly, I would like to express my sincere gratitude to my wife, Nihan Aksoy, for her unwavering support and encouragement throughout the research and education phases of this thesis. Her love, patience, and understanding were invaluable to me during this journey.

Lastly, I would like to thank my supervisor Prof. Erinç Albey and all the faculty members who have supported me throughout my studies. Their guidance and encouragement have been invaluable in my academic journey.

TABLE OF CONTENTS

DECLARATION OF ORIGINALITY	iv
ABSTRACT	v
ÖZET	vi
ACKNOWLEDGEMENTS	vii
LIST OF TABLES	xii
LIST OF FIGURES	xiii
LIST OF ACRONYMS AND ABBREVIATIONS	1
1 INTRODUCTION	3
1.1 Background and Motivation	3
1.2 Central Problem and Research Questions	5
1.3 Primary Objectives	6
1.4 Significance of the Study	7
1.5 Scope and Limitations	8
1.6 Thesis Structure	10
2 LITERATURE REVIEW	13
2.1 Price Elasticity of Demand Modeling in E-commerce	13
2.1.1 From Classical Approach to Machine Learning for Predictive Accuracy	13
2.1.2 Interpretable and High-Performance Models for Demand Prediction	14
2.2 The Strategic Landscape of E-commerce: Two-Sided Markets and Platform Competition	17
2.2.1 Modeling Conflicting Objectives in Platform Optimization	17
2.2.2 A Cooperative Game-Theoretic Approach: The Nash Bargaining Framework	20
2.3 Integrating Machine Learning with Mathematical Optimization	21
2.3.1 The Predict-then-Optimize Framework and its Limitations	21

2.3.2	The Emergence of Decision-Focused Learning	23
2.3.3	The Constraint Learning Framework: A Deep Integration	23
2.3.4	Summary and Research Gap	24
3	METHODOLOGY	26
3.1	Problem Formulation	26
3.1.1	Problem Statement	26
3.1.2	Constraints	26
3.1.3	Decision Variables	27
3.1.4	Objective Function	27
3.1.5	Full Mathematical Formulation	29
3.1.6	Environment Dynamics	29
3.2	Data Collection	30
3.2.1	Data Sources	30
3.2.2	Data Preprocessing	31
3.2.3	Data Organization	32
3.3	Demand Modeling	33
3.3.1	Rationale for a Machine Learning Approach	33
3.3.2	Selected Modeling Approaches	33
3.3.3	Justification of Chosen Modeling Strategies	34
3.4	Optimization Frameworks	35
3.4.1	The Integrated Framework (Constraint Learning)	35
3.4.2	The Decoupled Framework (Predict-then-Optimize)	37
3.5	Objective Function and Constraints	38
3.5.1	The Nash Bargaining Model	38
3.5.2	Operational Constraints	39
3.6	Demand Model Training and Evaluation	40
3.6.1	Training Protocol	40
3.6.2	Performance Metrics	40
3.6.3	Baseline Comparison	41
3.7	Optimization Evaluation Framework	41

3.7.1	Research Questions.....	41
3.7.2	The Evaluation Oracle.....	42
3.7.3	Performance Metrics	42
3.8	Experimental Setup	42
3.8.1	Hardware and Software Environment	42
3.8.2	Simulation Environment and Procedure.....	44
3.9	Limitations	45
4	RESULTS AND ANALYSIS	48
4.1	Introduction.....	48
4.2	Demand Model Performance Analysis.....	48
4.2.1	Predictive Accuracy Comparison	48
4.2.2	Analysis of Model Behavior and Prediction Errors.....	49
4.2.3	Analysis of Learned Price Elasticity.....	51
4.2.4	Feature Importance Analysis.....	53
4.3	Optimization Framework Performance Analysis	54
4.3.1	Computational Performance and Solver Robustness	54
4.3.2	Business Outcome Analysis	56
4.3.3	Analysis of Strategic Decisions	57
4.4	Discussion and Interpretation.....	59
4.4.1	Answering the Research Questions	59
4.4.2	Synthesis of Findings and Strategic Implications	60
4.5	Limitations and Recommendations	61
4.6	Practical Applications.....	62
4.7	Conclusion.....	62
5	CONCLUSIONS	63
5.1	Summary of Research.....	63
5.2	Contributions	63
5.3	Practical Implications	64
5.4	Future Research	65

5.5 Conclusion Statement	66
REFERENCES	68



LIST OF TABLES

Table 4.1: Aggregate Predictive Accuracy of Demand Models.....	49
Table 4.2: Solver Status by Framework.....	55



LIST OF FIGURES

Figure 4.1: Actual vs. Predicted Demand for a Sample SKU.	50
Figure 4.2: Demand Predictions for a Sample of SKUs.	51
Figure 4.3: Prediction Errors for a Sample of SKUs.	52
Figure 4.4: Distribution of Elasticity Across Models.	53
Figure 4.5: Feature Importance by Feature Group.	54
Figure 4.6: Solver Time by Framework.	55
Figure 4.7: Relative Merchant Profit Change by Framework.	56
Figure 4.8: Relative Platform Revenue Change by Framework.	57
Figure 4.9: Price Discount Ratio by Framework.	58
Figure 4.10: Commission Discount Ratio by Framework.	59

LIST OF ACRONYMS AND ABBREVIATIONS

ASP Average Selling Price

CL Constraint Learning

CPU Central Processing Unit

DBSCAN Density-Based Spatial Clustering of Applications with Noise

DFL Decision Focused Learning

DPCOP Dynamic Price and Commission Optimization Problem

EBM Explainable Boosting Machine

GAM Generalized Additive Model

GBDT Gradient Boosting Decision Tree

GDP Generalized Disjunctive Programming

GMV Gross Merchandise Value

GPU Graphics Processing Unit

KKT Karush-Kuhn-Tucker

KPI Key Performance Indicator

MAPE Mean Absolute Percentage Error

MINLP Mixed-Integer Non-Linear Program

MIO Mixed Integer Optimization

MIP Mixed Integer Programming

ML Machine Learning

MSE Mean Squared Error

NBS Nash Bargaining Solution

NLP Non-Linear Program

OLS Ordinary Least Squares

OMLT Optimization & Machine Learning Toolkit

PTO Predict-then-Optimize

RAM Random Access Memory

RMSE Root Mean Squared Error

SCIP Solving Constraint Integer Programs

SKU Stock Keeping Unit

SOS2 Special Ordered Set of Type 2

SPO Smart Predict-then-Optimize

1. INTRODUCTION

The main goal of this thesis is to develop and evaluate a framework for the Dynamic Price and Commission Optimization Problem (DPCOP) within a large scale e-commerce marketplace. The research explores the integration of machine learning for demand prediction and mathematical optimization for decision-making. At the core of this work, there is an investigation into two competing paradigms for integrating these disciplines: a decoupled Predict-then-Optimize (PTO) approach and a more integrated CL methodology.

This study aims to provide a data driven answer to a critical question for modern e-commerce platforms: how to design promotional campaigns that are mutually beneficial for both the platform and its third-party merchants. In order to achieve this goal, it is presented that a cooperative game-theoretic objective based on the Nash Bargaining Solution (NBS) might be utilized [1]. The performance of the developed frameworks is empirically evaluated in a comprehensive simulation environment. This evaluation framework might provide insights into the practical trade-offs between predictive accuracy, optimization complexity, solver robustness, and the generation of tangible business value.

1.1 Background and Motivation

The global commercial landscape has been fundamentally reshaped by the rise of e-commerce. The rising online marketplaces have evolved into complex ecosystems serving as the primary interface between consumers and sellers. Therefore, this environment is shaped by intense competition, where platforms try to capture market share while their member merchants pursue profitability and growth [2]. What appears at the core of the economic structure of these dynamic platforms is the concept of a two-sided market

[3]. It is argued that in this model, the marketplace platform functions as an intermediary, creating value by facilitating interactions between two distinct groups: consumers demanding goods and third-party merchants seeking customers. The success of the platform is therefore contingent on its ability to effectively satisfy the needs of both groups. The growth of one side might create a strong pull for the other [4]. For example, a wider range of products offered by sellers attracts more customers to the marketplace, which in turn makes the platform a more attractive place for merchants.

However, at the core of this symbiotic relationship, there lies a structural tension. It can be seen that the major e-commerce platforms often play a dual role, not only as a marketplace operator, but also as a direct competitor to sellers who rely on their infrastructure [5]. As argued, this situation creates a significant challenge, which can be seen in practices such as self-preferencing, where a platform might algorithmically favor its own private label products in search results or promotional placements over others, even though the items offered by the independent merchant are superior or better priced [6]. This tension highlights a fundamental problem that is not merely technical but strategic. While this inherent dual role and competitive tension derived from it are fundamental in some e-commerce settings, this study specifically focuses on optimizing pricing and commission strategies for the third-party merchants. Thus, it does not attempt to model scenarios where the platform itself sells the same SKUs (Stock Keeping Unit) or the competitive dynamics arising from this kind of direct platform merchant competition.

Platform pricing and commission strategies act as powerful levers that can either worsen this conflict or help align incentives. At the core of designing such strategies is the complex task of estimating the price elasticity of demand for thousands of different products, a task that requires the use of advanced machine learning models. After building these predictive models, a critical question arises: how can we integrate these predictions into a decision-making framework? Two main philosophies for this quest have emerged

in the academic literature and in industry practice. The first approach, being more traditional, is PTO. In this paradigm, a data science team builds a machine learning model with the sole objective of maximizing predictive accuracy. Then, in a separate step, an operations research team takes these predictions as fixed inputs to solve an optimization problem [7]. The alternative is a more integrated approach, often called CL. This method attempts to unify the prediction and optimization steps, for instance, by embedding the machine learning model directly into the optimization problem's constraints [8].

This choice presents both a scientific and a pragmatic tension. The integrated CL (alternatively called Decision Focused Learning) approach is argued to be more elegant, as it might allow the optimizer to see the full structure of the prediction model, at least in theory [8]. The PTO approach, on the other hand, might often be characterized as more practical. It aligns with common organizational structures where data science and optimization are separate functions, thereby simplifying the complexity of the final optimization problem [7]. In the following chapters, it will be analyzed that the choice between these two paradigms is not just a technical detail but a fundamental decision with consequences for the quality and viability of the final solution, by creating certain trade-offs for the decision makers. This research is motivated by the need to understand these trade-offs between these two approaches in a simulated environment mimicking a real-world setting, using the problem of balancing platform and merchant objectives as a complex and relevant test case [6].

1.2 Central Problem and Research Questions

This research formally addresses the DPCOP. This problem involves the simultaneous determination of optimal prices and commission rates for a multitude of products within an e-commerce platform. The goal is to satisfy a dual objective of maximizing both platform revenue and merchant profit, subject to a series of operational and business

constraints that reflect real-world market conditions.

The investigation is guided by two primary research questions:

- How can machine learning and mathematical optimization be combined to dynamically set prices and commissions that are mutually beneficial for an e-commerce platform and its merchants?
- What are the performance trade-offs between a deeply integrated (CL) and a decoupled (PTO) optimization framework in this context?

1.3 Primary Objectives

To provide clear answers to the research questions, the following specific objectives were established:

Demand Model Development and Evaluation: The Initial objective is to develop and evaluate robust machine learning models capable of accurately predicting the price elasticity of demand. A rich set of features, such as historical sales and orders, pricing, categorical features, and campaign data is utilized for this endeavor. The aim is to capture the complex price elasticities of a wide range of products in a dynamic e-commerce environment. Therefore, the performance of these models is a critical prerequisite for the success of any subsequent optimization.

Formulation of a Cooperative Optimization Model: Formulating a dual objective optimization model that is capable of addressing the above mentioned platform merchant tension is another aim of this project. The NBS is employed for this purpose instead of applying a more traditional weighted sum. The cooperative game-theoretic characteristic of that approach is peculiarly chosen to find a fair and balanced outcome. The aim was to move beyond simplistic, single-sided profit maximization to promote a sustainable marketplace ecosystem.

Framework Implementation and Comparison: A central goal of this research is to implement and empirically compare two distinct optimization frameworks: an integrated CL framework and a decoupled PTO framework. This implementation provides a direct test of these two competing paradigms from the operations research literature in a practical setting.

Comprehensive Performance Analysis: The final objective is to conduct a comprehensive analysis of the results from the framework comparison based on three key dimensions: computational performance, technical robustness, and the quality of the generated business outcomes. The evaluation is performed within a controlled simulation environment to ensure a fair comparison between the implemented frameworks.

1.4 Significance of the Study

As already mentioned, the main aim of the study is to contribute to both academic literature and industry practice through several distinct avenues.

From an academic perspective, an empirical comparison of the integrated CL approach versus the decoupled PTO paradigm for a complex real-world e-commerce problem. Central to this is the key finding that the predictive quality and economic sensibility of the underlying machine learning model are more critical for generating business value than the mathematical exactness of the optimization method. It can be argued that the results clearly show that an approximate framework (PTO) guided by a superior predictive model generates better outcomes than an exact framework (CL) constrained by a less accurate model. The findings might provide a data driven insight into the practical trade-offs within the PTO literature [9]. Additionally, it not only compares these two distinct frameworks but also offers a different perspective than a traditional weighted sum approach. In terms of problem formulation, the choice of the NBS framework might offer a novel approach to modeling cooperative agents in this domain. It might be used for these types

of dual objective platform merchant optimization problems.

From a practical standpoint, the findings of the thesis might provide several valuable contributions. First, it presents the PTO-Advanced pipeline as a scalable and robust template for implementing a dynamic pricing and commission optimization system. In order to unify profitability with fairness, the framework is designed to serve both parties' needs. Therefore, it promotes a "win-win" outcome that would foster a healthier and long-term marketplace ecosystem. Second, the research offers critical insights into the real-world engineering trade-offs of such systems. By analyzing not only the key metrics determined to evaluate business utilities, the thesis attempts to understand each framework's effects on the solver performance and robustness. This might help practitioners to better comprehend the computational expense and fragility of the two frameworks in a comparative manner. As results might have shown, the integrated CL approach shows lower quality in these performance based metrics. Therefore, it provides a clear, data backed justification for pursuing the more robust PTO methodology in a production environment. Finally, the demand modeling analysis itself yields actionable business insights. It identifies the key drivers of sales and validates the importance of factors such as product visibility and campaign structure.

1.5 Scope and Limitations

While this study attempts to offer a comprehensive analysis, it is important to clearly define its scope and acknowledge its inherent limitations upfront. These boundaries should be carefully considered when it comes to interpreting the conclusions drawn in this research.

First, the evaluation is performed within a simulated environment rather than a live A/B test. This is a serious setback and might be seen as questionable towards the results of the thesis. However, the business outcomes are calculated using the most accurate

predictive model as an "oracle" since it is believed that this best performance model might replicate the real-world behavior in the best manner, given that it is not possible to conduct an A/B test within the scope of this thesis. Therefore, it should be accepted that the validity of the quantitative findings is contingent on the fidelity of this oracle, and a live, randomized experiment would be the definitive next step for validation. In relation to this, the choice of using the best performance model might be justified for this characteristic. The results show that the oracle model establishes the most reliable ground truth simulator available for this study.

Second, the analysis is conducted at the individual SKU level. What this means is that the optimization is done for each product independently and in an isolated way. Thus, the complex interdependencies among products are not considered in the scope of this thesis. Even though it is known that the substitution effect has an enormous impact on market dynamics, such as a discount on one product might cause a cannibalization effect on sales of a substitute product, it is not modeled in this study. The simplification was necessary in order to manage the already complex problem. However, this means that the framework does not perform a true portfolio level optimization considering the interdependencies between products. In addition to this limitation, the study does not attempt to model scenarios where the platform itself offers the same SKUs that it tries to optimize pricing and commission for the merchant within the campaign structure. The study of this competitive dynamics, therefore, is out of the scope of this work.

Third, the effectiveness of the data-driven models is dependent on sufficient historical data. This inherently limits the framework's applicability for newly launched products or for "long-tail" SKUs with sparse sales history, a well-known challenge often referred to as the "cold-start" problem [10].

Finally, this thesis illuminates a crucial trade-off between theoretical exactness and practical viability in optimization. The CL framework is theoretically elegant; however,

its formulation as a Mixed-Integer Non-Linear Program (MINLP) proved to be computationally intensive, with a solver failure rate of over 11% in the experiments. This contrasts with the robust and efficient PTO approach. The computational performance results are also specific to the open-source SCIP solver used. It is acknowledged that state-of-the-art commercial solvers could potentially improve the robustness of the CL framework, a possibility that is outside the scope of this particular study.

1.6 Thesis Structure

This thesis is organized into five chapters to provide a clear progression of the research. It starts from foundational concepts and ends with the empirical results and final conclusions.

Chapter II: Literature Review provides the academic context for this study by delving into a review of existing research. The chapter starts tracing the evolution of demand forecasting from traditional econometric models to modern machine learning techniques. It establishes the necessity of machine learning models for complex e-commerce applications. It then reviews the economic literature on two-sided markets and cooperative game theory. The aim here is to provide a theoretical foundation for the NBS objective function used in this work. Finally, it attempts to conduct a rigorous survey on the operations research literature, particularly about the two dominant paradigms for the intersection between machine learning and optimization, namely PTO and CL approaches. Finally, it identifies the specific research gap that this thesis aims to fill.

Chapter III: Methodology offers a detailed description of the specific methods and procedures employed in this thesis. The chapter provides details of the entire research pipeline. It starts with the process of data collection from the e-commerce platform and the following steps of data preprocessing, cleaning, and feature engineering. It then goes on to explain the development and training of the two machine learning demand models,

namely the Linear Tree and the Ensemble model. It also provides the complete mathematical formulation of both the CL and PTO frameworks. By doing that, the chapter also provides the formulation for the Hull Reformulation and SOS2 approximation techniques utilized in the above-mentioned frameworks. It concludes by defining the NBS objective function, its constraints and decision variables, and the experimental design of the simulation used for evaluation.

Chapter IV: Results and Analysis demonstrates the results obtained from the experiments and provides a detailed interpretation of their significance. It is structured in two main parts: first, it presents a detailed evaluation of the predictive performance of the price elasticity of demand models by comparing them against a baseline. The baseline log-log econometric model is also utilized to validate the effectiveness of alternative machine learning models. It then analyzes their learned economic behaviors by calculating the price elasticities of different products to understand their different behaviors in terms of demand prediction. Secondly, it attempts to provide a comparative analysis of the three optimization pipelines. For this purpose, it assesses these pipelines on key metrics including computational performance, solver robustness, and the financial business outcomes. These business outputs are calculated utilizing the simulations for both the platform and its merchants to have a better idea about how their effects differ for the involved parties.

Chapter V: Conclusions serves as the final chapter, summarizing the key findings and contributions of the research with a general perspective. It revisits the research objectives and reiterates the evidence-based solutions to the core research questions posed in the introduction. It continues by discussing the practical implications of the findings for e-commerce platforms. In the end, it offers concrete recommendations and proposes several specific directions for future research that can be built upon the blueprints presented in this thesis.

In summary, this thesis is carefully organized to offer a comprehensive exploration

of the DPCOP. It proceeds logically from introducing the problem and its context to discussing the research methodology. Then it attempts to present the empirical results and draw meaningful conclusions. The reader is guided through a structured journey that aims to provide a clear and data-driven answer to the central research questions.



2. LITERATURE REVIEW

2.1 Price Elasticity of Demand Modeling in E-commerce

2.1.1 *From Classical Approach to Machine Learning for Predictive Accuracy*

Early academic and industry research on price elasticity of demand was dominated by econometric methods. Precisely, the log-log linear regression model served as the most preferred template for a long time. The log-log model assumes a constant elasticity of price in demand estimations, which can be estimated using traditional Ordinary Least Squares (OLS) after applying a logarithmic transformation for both demand and price. The primary reason for using this model is due to its simplicity and direct interpretability of its price coefficient, which actually represents a constant measure of price elasticity across all price levels. This became a foundational tool for this kind of modeling, especially when the data is sparse and computational power is limited [11].

However, the simplicity of the log-log model is also its weakness when applied to the modern e-commerce landscape [12]. It was shown that the constant elasticity assumption might be overly restrictive, as in most real-world cases, the price elasticity of demand does not demonstrate a static characteristic, especially if one considers promotional factors, seasonal effects, and intense competition among rival e-commerce platforms [13]. Additionally, the traditional regression models struggle to handle the non-linear relationships and high dimensionality of the data due to their inherent linear characteristic. Several studies on the topic have shown that the static elasticity of the linear log-log model is not reliable, and there is a need for more flexible modeling techniques that are capable of capturing the nuanced dynamics of the digital marketplace. It is especially relevant in the e-commerce context as the dynamic nature of online marketplaces reveals significant challenges. As shown in one recent work, these traditional models perform poorly in common retail scenarios with low product price variation, a known issue that can lead to

unreliable or economically nonsensical elasticity estimates [14].

Therefore, the shortcomings of traditional econometric models precipitated a paradigm shift towards the use of machine learning (ML) models in similar tasks [12]. This transition represents a movement of focus from parameter estimation towards a focus on prediction. As argued in the literature, ML provides a new set of tools for practitioners and researchers, specifically for the task of generating accurate predictions ($\hat{y}$) rather than estimating structural parameters (β) [15]. On that, numerous studies have demonstrated the consistent superiority of ML models in terms of predictive accuracy, due to the fact that non-parametric methods are particularly designed to learn complex and non-linear relationships and high order interactions from data [13]. This empirical predominance is also confirmed in this study, where both the Linear Tree and Ensemble models significantly outperform the log-log baseline.

This shift also introduced a new challenge that has become a central theme in the ML literature, known as the black box problem. Many of the powerful predictive models might have an internal logic that is ambiguous to the human eye [16]. The inherent nature of price elasticity studies, on the other hand, might require a better comprehension of which factors are more effective in sales, even though it might mean less accurate models. This tension between predictive accuracy and model interpretability has paved the way for the development of more sophisticated models [12].

2.1.2 Interpretable and High-Performance Models for Demand Prediction

One significant area of research has focused on developing models that are both powerful and transparent. Linear (Optimal) Trees have emerged as a promising approach in this regard. Their partitioning logic is good for helping the model to understand the

underlying non-linear relations, and it is coupled with the modeling power of linear regression by fitting a separate linear model at each leaf node [17]. Apart from their usability in terms of understanding segment specific and heterogeneous elasticities, their direct applicability to prescriptive analytics was a key contribution in this line of research, especially when the close relation of demand estimation and price optimization is considered. As shown in recent works, the mathematical structure of an optimal regression tree can be formulated as a set of mixed integer optimization (MIO) constraints, which makes their direct integration into optimization problems quite possible [18]. This optimization-friendly structure is therefore the primary reason behind the choice of linear trees for the integrated CL framework evaluated in this thesis.

In parallel to the pursuit of maximizing predictive accuracy, another class of models has become quite dominant, namely, Gradient Boosting Decision Trees (GBDTs). Highly efficient implementations like LightGBM are widely recognized as state of the art for prediction tasks on tabular data, hence, extensively used in e-commerce demand forecasting applications [19]. In relation to its achievement in providing the lowest prediction error in many cases, LightGBM has become one of the core components of high performance Ensemble model used in this thesis.

In order to directly address the black box dilemma without sacrificing performance, recent works have produced a new class of glass box models. One of the most prominent of these models is the Explainable Boosting Machine (EBM) [20], which modernizes the mostly used but traditional Generalized Additive Model (GAM) [21] with techniques from gradient boosting. A GAM models an outcome as a sum of individual and potentially non-linear functions of each feature:

$$g(\mathbb{E}[y]) = \beta_0 + f_1(x_1) + f_2(x_2) + \dots \quad (2.1)$$

An additive characteristic of EBM is inherently interpretable as the contribution of each feature can be isolated independently. EBMs use a specialized cyclic boosting algorithm

to learn these functions and are capable of learning pairwise interaction terms. Also, several studies have shown that EBM can achieve predictive accuracy that is quite competitive with best performing black box models [22]. In this study, the EBM is also a key component of the high-performance Ensemble model.

A prevailing theme in modern ML applications is the use of ensemble models due to its capability of combining the predictions from multiple and diverse models to generate a final prediction that is more robust and accurate in terms of predictive power and stability compared to their constituent models. The individual errors and biases of the base models tend to cancel each other out due to the ensemble model's behaviour of averaging the outputs of different models [23]. This approach has been proven to be particularly effective in e-commerce, where researchers have applied ensemble models to handle complex prediction problems. Its advantage of handling the problems characterized by high volatility and complex relationships makes it a powerful tool for managing bias variance trade-off and reducing the risk of overfitting [24]. It is shown that the effectiveness of an ensemble is greatest when the constituent models are diverse in their structure and type of errors they make [23], so it clearly is a strong motivation for using such ensembles in this thesis.

While the literature has extensively explored the predictive capabilities of these models and the challenges of the e-commerce data, it might be argued that there is still a crucial need for works that intend to execute a direct and empirical comparison of these different approaches in terms of their performance and stability when faced with real world e-commerce data. The practical trade-offs exist between each model choice, and one of the main goals of this thesis is to present empirical results and new theoretical approaches in this regard.

2.2 The Strategic Landscape of E-commerce: Two-Sided Markets and Platform Competition

The preceding review established the evolution of modeling techniques for achieving high performance in demand predictions. However, there is a need for translating these forecasts into optimal actions to surface the ultimate value of such models, which is referred to as the transition from predictive analytics to prescriptive analytics. The former mainly focuses on the question of what might happen, whereas the latter is centered around the question of what we should do. As articulated in the current literature, prescriptive analytics leverages predictive models to recommend specific and optimized decisions that should align with business objectives so as to bridge the gap between insight and action [25]. It is particularly critical in e-commerce domain where pricing and promotional campaigns are the primary levers of revenue management.

2.2.1 Modeling Conflicting Objectives in Platform Optimization

Before jumping into constructing a prescriptive framework, one should understand the unique economic environment of modern e-commerce marketplaces. Online marketplaces are complex ecosystems that are arguably best comprehended through the lens of two-sided market theory [26]. A two-sided market is a platform that acts as an intermediary between two distinct parties: the consumers and merchants. As pointed out in previous parts of this study, the value of the platform increases when more merchants join, and conversely, the value of the platform for a merchant also increases as more customers join. Due to these network effects, the platform's primary issue is to make both sides satisfied [26].

This situation makes the platform's pricing structure even more critical and strategic lever. This structure has become further complicated when a significant source of tension has risen as a result of the platform's dual role as both a marketplace operator

and a direct competitor to the merchants. This dual role has significant implications that must be accounted for in any realistic optimization model. As recent studies have shown that platforms can engage in self-preferencing where they algorithmically favor their own products over potentially superior third-party offerings [27]. Furthermore, the information asymmetry between the platforms and their sellers can create a situation where platforms leverage this advantage to identify successful products to imitate and launch under their own brand [28].

Within this competitive context, the commission rate charged to sellers is transformed from a simple service fee into a powerful strategic weapon that necessitates careful engineering. The choice of selling the optimized product under its own brand or offering a competitive commission to lower the price is out of this study, but it remains an important topic to delve into. Consequently, due to the scope of this work, it does not attempt to model this competition that would arise if the e-commerce platform were also selling the identical SKUs to the customers.

Apart from this tension, the commission rate offered to the merchant still remains as an important concept that must be included in the objective function as a revenue channel for the platform and a cost for the merchant. Considering that the optimal price and commission levels are part of a promotion structure, an ideal optimization problem must include the merchant side, since it is the merchants' choice of whether to accept or reject these offers generated by the platform itself.

Consequently, we have a structural tension inherent in an e-commerce platform, which requires a careful design of a multi-objective optimization problem. The platform, on one hand, seeks to maximize its own revenue, whereas the merchants seek to maximize their individual profits. The literature on joint price and commission optimization has often employed game theoretic models, typically framing the interaction as a Stackelberg game where the platform acts as a leader by setting the commission rate and sellers

respond with their pricing decisions [29]. When formulating the objective function for these problems, the common approach in operations research is the weighted sum method [30]. This technique makes the multiple objectives scalar into a single objective function by assigning a subjective weight to each party's utility, formulated as:

$$U = w_{\text{platform}} \cdot \pi_{\text{platform}} + w_{\text{merchant}} \cdot \pi_{\text{merchant}} \quad (2.2)$$

However, despite its common use, the weighted sum method suffers from several drawbacks that might render it unsuitable. As widely discussed in multi-objective optimization literature, a primary limitation is the inherent subjectivity of weights. There is no objective and principled method for determining the ideal weights that represent an optimal balance between different groups' interests. It is quite possible that each decision maker will choose different weights based on their own biases and priorities, which might lead to arbitrary and potentially inequitable outcomes. Furthermore, the weighted sum method has significant limitations, which are well documented in the literature. It might be unable to find solutions on non-convex portions of the Pareto optimal front, meaning that if the set of all efficient trade-offs between platform and merchant utility is non-convex, there may exist optimal compromise solutions that the weighted sum method is mathematically incapable of identifying, no matter how the weights are determined. Another studied weakness is that the method also suffers from instability. A small change in the weights might lead to a large jump to a completely different solution on the Pareto front. A sustainable platform relies on building stable, predictable, and trustworthy relationships with its merchants to encourage long-term participation. Therefore, the weighted sum method undermines this goal as it might create a compelling need for a solution concept founded on robust principles [30].

2.2.2 A Cooperative Game-Theoretic Approach: The Nash Bargaining Framework

As an alternative to the weighted sum method due to its weaknesses, this study posits the NBS. It is a foundational concept in cooperative game theory, offering theoretically a more sound and robust framework as an answer to the dual objective DP-COP. Cooperative game theory focuses on how a surplus generated in a collaborative way should be divided equally among parties instead of modeling the strategic actions of self-interested players. The proposed solution provides a unique outcome for a bargaining problem where two parties are involved by identifying the allocation that maximizes both players' utility gains relative to a disagreement point [1]. The objective is generally formulated as:

$$\max (\pi_{\text{platform}} - d_{\text{platform}}) \cdot (\pi_{\text{merchant}} - d_{\text{merchant}}), \quad (2.3)$$

However, we adopt the log-sum NBS objective formulated as:

$$\max \log(\pi_{\text{platform}}) + \log(\pi_{\text{merchant}}), \quad (2.4)$$

There are several reasons behind this choice. First, the log sum formulation is numerically more stable and avoids the instability that can occur when one party's profit approaches its threshold. Secondly, the log transformation normalizes scales between the two parties, which can be essential given the fact that their possibly disparate magnitudes in e-commerce contexts. Lastly, maximization of the sum of logs is mathematically equivalent to maximizing the geometric mean of the profits. Therefore, it preserves the fairness properties of NBS while it also yields a formulation that is better suited to real-world computation.

The NBS framework is not a purely technical tool but a proven tool with a rich history of application in operations management for analyzing how collaborative firms divide profits and align incentives. For instance, foundational study on supply chain contracts makes research on the problem of surplus division between manufacturers and retailers.

It sets the stage for formal bargaining models [31]. Building on this work, many studies use the NBS to find an optimal equilibrium in pricing and wholesale contract terms. Their approach is often compared with non-cooperative models [32]. Cooperative game theory is used to design joint promotional strategies and determine optimal cost sharing policies between supply chain partners in the online retail domain [33]. More directly, the NBS is employed to analyze negotiations between e-commerce platforms and their manufacturers. The studies underscore its relevance to the platform merchant dynamic, which is also a central point in this study [34]. The framework's versatility is further highlighted in studies that use both Nash and other bargaining models, such as the Rubinstein bargaining contract, in order to coordinate complex objectives such as green initiatives in supply chains [35]. The successful application of this framework in diverse but relevant contexts provides strong motivation for its use in this study to resolve the tension between the e-commerce platform and its sellers.

2.3 Integrating Machine Learning with Mathematical Optimization

It is not sufficient to only generate accurate price elasticity of demand estimations for a successful application of prescriptive analytics. In the end, the ultimate goal is to translate these predictions into optimal insights to take effective actions. This proliferated a rapidly evolving field of research at the intersection of ML and optimization studies that mainly focused on developing methods for integrating predictive models into optimization frameworks [25, 36]. The following section reviews two dominant paradigms that have emerged from this literature, namely the PTO approach and the integrated CL approach, to establish the methodological context for the central question of this work.

2.3.1 The Predict-then-Optimize Framework and its Limitations

The most widely adopted method for combining prediction and optimization is the two stage and decoupled framework known in the literature as PTO. In this paradigm,

two tasks are performed in a sequential manner, and these tasks are handled independently from each other. At the first stage, an ML model is trained with the sole objective of maximizing predictive accuracy and its main focus is to minimize a standard statistical loss function such as MSE. Secondly, the point predictions of this model are treated as fixed and deterministic parameters in an optimization problem. This problem is then solved to come up with a decision. Due to its practical appeal, it aligns with conventional organizational structures where data science and optimization teams function as distinct units. It also allows for the use of state-of-the-art predictive models without further complicating the final optimization task, even though these models themselves are highly complex in nature.

However, despite its ease of use in optimization problems, a body of literature claims a fundamental flaw in the PTO framework. The objective of the predictive model is misaligned with the objective of the decision making task. As articulated in the literature, ML models trained to minimize prediction error do not account for how those predictions will be used in decision making [7]. This disconnection problem is crucial since not all prediction errors are equal from a decision making perspective. Even a small error on a parameter to which the optimization problem is highly sensitive might lead to significant errors in terms of optimal decision; conversely, a large error on an insensitive parameter may have a negligible impact on the final decision quality.

It is demonstrated that this misalignment can lead to worse decisions as the model might learn to be accurate on average, but not for specific instances that have crucial importance for making good decisions [37]. It is studied that improving predictive accuracy does not guarantee better decision quality in various domains. In one particular example, it is argued that the choice of decision making strategy often had a greater impact on outcomes than the predictive model's training objective [9].

2.3.2 The Emergence of Decision-Focused Learning

This fundamental challenge has given rise to the field of Decision Focused Learning (DFL), which seeks to meld the pipeline by making the ML model aware of the optimization task during training [7, 37, 38]. The central innovation of DFL (also referred to as end-to-end learning) is to create an end-to-end pipeline that aligns the training of the predictive model with the final decision making task [39]. Instead of only minimizing a purely statistical loss, DFL methods train the predictive model to minimize a decision aware loss function that tries to measure the quality of the decision induced by the model’s predictions. In the literature, it is often called a decision error or regret, meaning that the difference in objective value between the decision made with the predicted parameters and the true optimal decision that would have been made with perfect information.

There are studies labeling this new framework as the Smart Predict-then-Optimize (SPO) framework, and in this structure, they formally define the SPO loss function measuring the decision error [7]. However, due to the non-convex and discontinuous characteristics of this loss function, which make it computationally challenging to optimize directly, they derive a convex surrogate loss function, namely SPO+loss. It is shown that this surrogate function is statistically consistent and computationally tractable for a wide range of optimization problems. There are other methodologies developed by several researchers that can be categorized as DFL methods, such as KKT conditions [36]. These DFL approaches have been shown to yield superior decision quality, particularly when the predictive model is misspecified or the data is complex [36, 38, 39, 40, 41].

2.3.3 The Constraint Learning Framework: A Deep Integration

A particularly powerful implementation of the DFL philosophy is CL. It is a methodology seeking to embed the ML model itself as a set of constraints within the mathematical optimization problem. As detailed in the literature, this approach creates a

unified Mixed Integer Optimization (MIO) problem that can be solved holistically. This deep integration allows the optimizer to see the complete structure of the predictive model. This characteristic enables the optimization problem to find a globally optimal solution with respect to the learned relationships [8]. The practical value of this approach has been demonstrated in real world applications. In a recent study, the CL approach has been used to optimize problems in healthcare, nutrition [42], and energy [43, 44] domains.

The feasibility of the CL framework depends on the mathematical representability of the ML model. However, this theoretical superiority comes with important practical challenges. Not all models can be translated into a format that an MIO solver can understand. This paved the way to significant research into embedding different ML model classes into optimization problems. However, even a simple ML model like a decision tree might increase the size and complexity of the resulting optimization problem, often leading to a Mixed Integer Non-Linear Program (MINLP). As documented, these problems can be computationally expensive to solve and the solver might suffer from robustness issues such as solver failures and long solution times [42]. This characteristic might make it harder to implement it in real world applications that might require timely decisions.

2.3.4 Summary and Research Gap

In summary, the literature on integrating ML and optimization has branched into two main philosophies. On the one hand, there is the classical approach of PTO, treating each phase as a separate entity, and on the other hand, there is also the decision-focused approach seeking to unify the prediction and decision processes. While the theoretical properties and performance of these approaches have been studied in various contexts, there is still a significant research gap in their empirical comparison, particularly for real-world e-commerce problems such as the DPCOP. It remains an open question how the practical benefits of the PTO approach weigh against the theoretical exactness of the CL

approach. This study is an attempt to fill this critical gap. The results of this thesis might contribute to the literature by providing empirical evidence on the relative advantages of these competing philosophies and offering insights for both academics and practitioners on the trade-offs when it comes to designing prescriptive analytics systems, especially in the e-commerce domain.



3. METHODOLOGY

3.1 Problem Formulation

This research addresses the problem of the DPCOP within the context of an e-commerce platform acting in a large scale environment. It involves determining the optimal price and commission rates for a multitude of products (SKUs) over a given period. The DPCOP is subject to various business and market constraints in response to volatile demand movements, competitor actions, and internal campaign activities.

3.1.1 Problem Statement

The main objective of the DPCOP is to simultaneously maximize the platform's revenue and the individual merchant's profit. The goal is to develop an intelligent optimization framework that determines ideal price and commission levels for the sellers in the marketplace. The platform, therefore, offers this ideal price and commission for a given SKU to the merchant with the aim of ensuring a sustainable ecosystem that benefits both the platform and its merchants in a balanced manner.

3.1.2 Constraints

The DPCOP is subject to several constraints that must be considered during the optimization process:

- **Price Constraints:** The first of these couple of constraints is the price constraint. The selling price for each product is bounded by a defined range established for the campaign. The maximum price is an estimated historical price; therefore, the merchants are offered the opportunity to reduce their prices, as in the campaign structure. This estimated price is derived from past order data and competitor pricing. This past data serves as the reference point for the campaign price reduction. On the other hand, the minimum price is determined by the cost estimation of the

platform. This estimation accounts for commissions charged to the merchants, retail costs of similar products that are offered by the platform, and market dynamics. Its primary goal is to ensure merchant profitability.

- **Commission Constraints:** The commission rate charged to the sellers by the platform is also bounded. The maximum commission is set to the previously applied rate. Similar to price bounds, the aim here is to offer a reduction in commissions to motivate the merchant to participate in the campaign mechanism. The minimum is a default value. This default is strategically determined by the platform to maintain its revenue floor.
- **Non-negativity Constraints:** The predicted demand for each product, y_{it} , must be non-negative ($y_{it} \geq 0$) to ensure feasible and meaningful sales levels. Additionally, both the platform's and merchant's profits must remain strictly positive ($\pi_{\text{platform},it} > 0$, $\pi_{\text{merchant},it} > 0$), as required by the logarithmic NBS objective.

3.1.3 *Decision Variables*

The decision variables in the DPCOP are the price to be set for each product and the commission rate to be applied to each sale. Let p_{it} represent the price for the item i in time period t , and c_{it} represent the commission rate for the item i in time period t . The objective is to optimize these two decision variables for each SKU.

3.1.4 *Objective Function*

In this study, the objective function is formulated as an NBS to achieve a balanced solution that is mutually beneficial for both the e-commerce platform and its merchants. The NBS approach aims to maximize the product of the two parties' profits (or utilities), which is achieved in practice by maximizing the sum of the logarithms of their respective

profit functions:

$$\max \log(\pi_{\text{platform}}) + \log(\pi_{\text{merchant}}) \quad (3.1)$$

where π_{platform} and π_{merchant} represent the profits of the platform and the merchant, respectively.

The platform's revenue is modeled as a function of the price and the commission rate, while the merchant's profit is determined by the price, commission rate, and estimated product cost. This cooperative game-theoretic objective would ensure that the optimization seeks a fair equilibrium. As in this optimization framework, this equilibrium does not favor a single party at the expense of the other, thus it promotes a more stable and sustainable marketplace.

The use of the logarithmic transformation, rather than directly maximizing the product of profits, brings several practical and theoretical advantages. First, the log transformation enhances numerical stability. This is especially relevant when someone considers that profit values span several orders of magnitude, which is often the case in e-commerce settings. It compresses extreme values and reduces the risk of computational overflow or underflow in nonlinear solvers. Second, maximizing the sum of logs is mathematically equivalent to maximizing the geometric mean of the profits. Thereby it preserves the fairness at the core of the NBS. Third, the log-sum objective is concave and differentiable for positive profits, which improves the tractability and convergence of the optimization process and ensures that the resulting solution is robust even when profits approach their lower bounds. Finally, the log transformation normalizes the scale between the platform and merchant profits, preventing the solution from being dominated by the party with larger absolute profits.

In summary, the cooperative nature and log-transformed application of the NBS function would enable the optimization framework to be fair to both parties. It therefore efficiently balances the interests of both the platform and its merchants. This, in return,

would lead to more stable outcomes for the entire marketplace ecosystem.

3.1.5 Full Mathematical Formulation

Let y_{it} denote the predicted demand for item i at time t , p_{it} the price, c_{it} the commission rate, and k_{it} the unit cost. The profit functions for the platform and merchant are:

$$\pi_{\text{platform},it} = y_{it} \cdot p_{it} \cdot c_{it}$$

$$\pi_{\text{merchant},it} = y_{it} \cdot [p_{it} \cdot (1 - c_{it}) - k_{it}]$$

The optimization problem for each product and period is:

$$\begin{aligned} & \underset{p_{it}, c_{it}}{\text{maximize}} && \log(\pi_{\text{platform},it}) + \log(\pi_{\text{merchant},it}) \\ & \text{subject to} && p_{it}^{\min} \leq p_{it} \leq p_{it}^{\max} \\ & && c_{it}^{\min} \leq c_{it} \leq c_{it}^{\max} \\ & && y_{it} \geq 0 \\ & && \pi_{\text{platform},it} > 0 \\ & && \pi_{\text{merchant},it} > 0 \end{aligned}$$

where $y_{it} = f(p_{it}, \dots)$ is the predicted demand from the machine learning model, and all other terms are as previously defined. The non-negativity constraint on y_{it} ensures that predicted demand is always feasible, and the positivity constraints on profits ensure that the log objective is well-defined.

3.1.6 Environment Dynamics

The e-commerce environment is highly dynamic. This dynamic nature is characterized by fluctuating demand patterns, shifting competitor prices, and the influence of various promotional campaigns. Customer demand can be affected by seasonality, market trends, and product visibility (e.g., view counts). Competitor pricing directly impacts

price elasticity and sales volume. Furthermore, internal decisions, such as running specific campaigns, alter the demand landscape. The optimization model must interact with this dynamic environment to make optimal decisions that account for these real-time changes.

By clearly formulating the DPCOP and its key components, this research establishes a solid framework for developing the optimization approach, including its implementation. The model's structure is tailored to address the dynamic and competitive nature of e-commerce, offering a novel solution to enhance pricing and commission strategies in modern online marketplace operations.

3.2 Data Collection

Data collection plays a pivotal role in the training phase of the price elasticity of demand prediction models that underpin the DPCOP. This subsection outlines the process of data collection, the sources for the data used in the research, and the steps involved in data preprocessing.

3.2.1 Data Sources

The data required for this research were obtained from various sources within the e-commerce platform's data warehouse. The key data sources include:

- **Historical Sales Data:** Transactional data for each SKU, including order dates, quantities sold, and Gross Merchandise Value (GMV), were collected. Then, this data was analyzed to better understand sales performance and demand patterns over time.
- **Listing and Pricing Information:** Historical listing prices and calculated Average Selling Prices (ASP) for each SKU were gathered to analyze price movements and their impact on sales.

- **Competitor Prices:** Data on competitor pricing for identical products were collected to provide a benchmark for the platform’s own pricing strategies. It is also utilized to understand the competitive landscape.
- **Campaign Details:** Information regarding promotional campaigns, including campaign type, discount amounts, and conditions, was acquired to model the impact of marketing activities on product demand.
- **Product View Data:** Data on product view counts were collected as a crucial indicator of customer interest and website traffic. Sales data is hugely affected by the traffic that the marketplace acquires, especially during heavy campaign periods. Therefore, analyzing sales from a view perspective is essential for accurately modeling demand.
- **Categorical Data:** Categorical information, such as product category, sub-category, and brand, was also included to enrich the feature set for demand modeling.

3.2.2 Data Preprocessing

Before the data could be used for modeling, it underwent several preprocessing steps to ensure its quality, consistency, and suitability for analysis. These steps included:

- **SKU Selection:** An initial filtering process was applied to select SKUs with sufficient historical sales data. This step is crucial for constructing a reliable price elasticity model, as sparse data can lead to inaccurate demand predictions.
- **Data Cleaning and Outlier Removal:** The collected data was cleaned to handle missing values and erroneous entries. An outlier removal process, using the DBSCAN clustering algorithm [45], was applied to the price and quantity data for each SKU to identify and exclude anomalous data points that could skew the demand models.

- **Normalization:** Numerical data, such as price and view counts, were normalized to bring all parameters to a common scale. This prevents features with larger magnitudes from disproportionately influencing the model training process.
- **Encoding Categorical Data:** Categorical features within the dataset were converted into a numerical format using encoding techniques to ensure compatibility with the machine learning models.
- **Feature Selection:** For the more basic models developed in this study, a feature selection process was employed to choose the features with the most informational value, ensuring model parsimony and reducing the risk of overfitting.

3.2.3 Data Organization

Once preprocessed, the data was restructured into a suitable format for training the demand prediction models. The time-series data was organized on a weekly basis. Then, the lagged features and rolling means for price, quantity, and view count feature groups were engineered to capture temporal dependencies. The dataset was then divided into training, validation, and test sets. The training set was used to build the demand models, the validation set to tune hyperparameters, and the test set to assess the final performance of the models on unseen data.

These data preprocessing steps then laid the foundation for model development and also prepared the inputs of the optimization framework. These meticulous steps were taken to ensure the robustness of the model development and optimization model. The subsequent chapters will focus on the implementation and evaluation of this integrated system, providing valuable insights into the effectiveness of a data-driven approach to dynamic price and commission setting.

3.3 Demand Modeling

A critical component of the optimization framework is the accurate estimation of the price elasticity of demand. This requires a model capable of capturing the intricate relationship between price, product features, market conditions, and sales quantity. In this subsection, the rationale for using a machine learning approach is presented, and the two distinct modeling strategies chosen for this research are detailed and justified.

3.3.1 Rationale for a Machine Learning Approach

The dynamic e-commerce environment is characterized by complex, non-linear interactions between numerous variables. Factors such as competitor pricing, promotional activities, seasonality, and product visibility all influence consumer demand in ways that are difficult to capture with traditional statistical models. A simple linear regression, for instance, would fail to account for the nuanced, conditional relationships inherent in the data. Therefore, a data-driven machine learning approach is necessary to effectively model the price elasticity of demand, providing the predictive power required to make informed optimization decisions.

3.3.2 Selected Modeling Approaches

Two distinct modeling strategies were selected to serve different purposes within the overall optimization framework: a structured Linear Tree (alternatively Optimal Tree [18, 46]) model for a CL pipeline, and a high-performance ensemble model for a PTO pipeline.

3.3.2.1 Linear Tree

The Linear Tree model was chosen for its unique hybrid structure, which combines the non-linear partitioning capabilities of a decision tree with linear regression models at

its leaf nodes. This structure is exceptionally well-suited for representing the underlying demand relationship. The tree’s splits can capture complex, non-linear interactions between features, while the linear models in the leaves provide a clear, interpretable representation of price elasticity for specific segments of the data. Crucially, the well-defined mathematical form of a Linear Tree allows it to be embedded directly into a mathematical optimization model as a set of constraints. This compatibility is essential for the CL framework, enabling an integrated and exact optimization approach.

3.3.2.2 Ensemble Model

With the aim of establishing a benchmark for predictive performance, a more complex model was developed using an ensemble approach. This model combines the predictions of several powerful algorithms, including LightGBM [19], Explainable Boosting Machines (EBM) [20], and Linear (Optimal) Trees. Ensembling is a state-of-the-art technique for achieving highly accurate predictions in a robust manner since it leverages the diverse strengths of its constituent models and mitigates their individual weaknesses. The choice of an ensemble model is driven by the goal of maximizing predictive accuracy, which is paramount for the PTO pipeline. Its ability to produce more stable and realistic demand elasticity estimates makes it a powerful engine for this sequential optimization strategy.

3.3.3 Justification of Chosen Modeling Strategies

The selection of these two distinct modeling strategies is a deliberate choice to explore two different philosophies in integrated optimization, including the prediction component.

Due to its unique suitability for the CL paradigm, the Linear Tree model is justified. Its value lies not only in its predictive capability but also in its structural transparency, which allows the learned demand function to be explicitly represented within the

optimization problem. This enables the solver to find a globally optimal solution with respect to the learned model, therefore, it creates a tightly integrated system.

Conversely, the justification of using the Ensemble Model is due to its pursuit of maximum predictive accuracy. In a PTO workflow, the quality of the final decision is heavily dependent on the accuracy of the initial prediction, as will be demonstrated in the next sections of this study. By employing a state-of-the-art ensemble, this research aims to test the practical limits of this sequential approach, providing a powerful benchmark against which the CL method can be compared.

By developing and evaluating these two parallel approaches, this research provides a comprehensive analysis of different methods for tackling the DPCOP. The subsequent chapters will detail the implementation and empirical evaluation of these models within their respective optimization pipelines.

3.4 Optimization Frameworks

The core of this research lies in the development of two distinct optimization frameworks designed to solve the DPCOP and the comparison of these frameworks in an evaluation environment. These frameworks, namely CL and PTO, differ fundamentally in how they integrate the machine learning demand model with the mathematical optimization problem.

3.4.1 The Integrated Framework (Constraint Learning)

The CL framework represents a high-fidelity approach where the machine learning model is not treated as a black box but is instead embedded directly into the optimization problem as a series of mathematical constraints. This creates a unified Mixed Integer Non-Linear Program (MINLP) that the solver can address holistically.

3.4.1.1 Mathematical Formulation

The key to this framework is the formulation of the Linear Tree demand model in a way that is comprehensible to an optimization solver. The prediction of a Linear Tree for a given input feature vector x can be expressed as:

$$q(x) = \sum_{l \in \mathcal{L}} (w_l^T x + b_l) \cdot z_l(x) \quad (3.2)$$

where $\mathcal{L}$ is the set of all leaf nodes, w_l and b_l are the coefficients and intercept of the linear model in leaf l , and $z_l(x)$ is a binary indicator that is 1 if x falls into leaf l , and 0 otherwise.

To represent the "if-then" logic of the tree structure, we first use Generalized Disjunctive Programming (GDP), a high-level modeling paradigm. The tree is expressed as a large logical "OR" across all leaves:

$$\bigvee_{l \in \mathcal{L}} \begin{bmatrix} z_l \\ A_l x \leq d_l \\ q = w_l^T x + b_l \end{bmatrix} \quad (3.3)$$

subject to $\sum_{l \in \mathcal{L}} z_l = 1$. Here, z_l is a binary variable indicating the active leaf, and the matrix inequality $A_l x \leq d_l$ represents the path of splits from the root to leaf l .

While GDP is conceptually elegant, solvers require a standard MIP formulation. This is achieved via a Hull Reformulation, which converts the GDP into a tight and efficient set of linear constraints. This method is chosen over simpler Big-M formulations because it provides a stronger, tighter convex relaxation of the problem, which avoids numerical instability and generally leads to superior solver performance [47]. The reformulation works by disaggregating the original variables for each leaf. New continuous variables, $\hat{q}_l$ and $\hat{x}_l$, are introduced to represent the portion of the output and input variables associated with leaf l . The original variables are then reconstructed through

aggregation constraints:

$$q = \sum_{l \in \mathcal{L}} \hat{q}_l \quad (3.4)$$

$$x = \sum_{l \in \mathcal{L}} \hat{x}_l \quad (3.5)$$

Finally, the logic for each leaf is reformulated using these partial variables and the binary indicator z_l . For each leaf l , the linear model, path conditions, and variable bounds are rewritten as:

$$\hat{q}_l = w_l^T \hat{x}_l + b_l \cdot z_l \quad (3.6)$$

$$A_l \hat{x}_l \leq d_l \cdot z_l \quad (3.7)$$

$$x_{\min} \cdot z_l \leq \hat{x}_l \leq x_{\max} \cdot z_l \quad (3.8)$$

This system of linear equations effectively linearizes the tree's logic, allowing it to be solved as part of the larger MINLP.

3.4.2 The Decoupled Framework (Predict-then-Optimize)

The PTO framework follows a more traditional, two-stage approach. First, the trained machine learning model (either the Linear Tree or the more complex Ensemble Model) is used to predict a demand curve across a range of possible prices. Second, this predicted curve is used as a surrogate or lookup table within a separate optimization problem.

3.4.2.1 Mathematical Formulation

This decoupling simplifies the optimization problem's complexity. The complex, non-linear demand curve generated by the ML model is approximated by a series of linear segments. This is achieved by creating a Piecewise Linear Approximation of the demand function $q(p)$. Let the set of price breakpoints be $(P_1, P_2, \dots, P_K)$ and the corresponding predicted demands from the ML model be $(Q_1, Q_2, \dots, Q_K)$. Any price p and its

corresponding demand q on the approximated curve can be represented as a convex combination of these points:

$$p = \sum_{k=1}^K \lambda_k P_k \quad (3.9)$$

$$q = \sum_{k=1}^K \lambda_k Q_k \quad (3.10)$$

To ensure that the point (p, q) lies on a single line segment between two adjacent break-points, a set of non-negative weights λ_k is introduced, subject to two conditions:

1. $\sum_{k=1}^K \lambda_k = 1$
2. The set of weights $\{\lambda_1, \dots, \lambda_K\}$ is a Special Ordered Set of Type 2 (SOS2).

The SOS2 [48] constraint dictates that at most two λ_k values can be non-zero, and if two are non-zero, they must be adjacent in the set. This elegantly enforces the piecewise linear relationship, transforming the problem into a more tractable Non-Linear Program (NLP) that can be solved efficiently, albeit with the introduction of some approximation error.

3.5 Objective Function and Constraints

The optimization problem is defined by a central objective function that balances the interests of the marketplace and its merchants, subject to a set of operational constraints that ensure realistic and beneficial outcomes.

3.5.1 The Nash Bargaining Model

To ensure a fair and sustainable marketplace ecosystem, a simple weighted-sum objective was avoided, as it often leads to corner solutions where one party's profit is maximized at the complete expense of the other. Instead, this research adopts a cooperative game-theoretic approach based on the NBS [1]. The objective is to maximize the

geometric mean of the two parties' utilities, which is equivalent to maximizing the sum of the logarithms of their respective profit functions.

Let p_i and c_i be the price and commission for an SKU i . Let $q_i(p_i)$ be the predicted demand and cost_i be the product cost. The objective function is:

$$\text{maximize } \sum_i (\log(\pi_{\text{platform},i}) + \log(\pi_{\text{merchant},i})) \quad (3.11)$$

where:

$$\pi_{\text{platform},i} = c_i \cdot p_i \cdot q_i(p_i) \quad (3.12)$$

$$\pi_{\text{merchant},i} = (p_i(1 - c_i) - \text{cost}_i) \cdot q_i(p_i) \quad (3.13)$$

This formulation inherently seeks a balanced equilibrium. Depending on the strategic goals of the platform, weights can be introduced into the logarithmic terms to give precedence to one party over the other, while still avoiding extreme corner solutions.

3.5.2 Operational Constraints

The optimization is bound by a set of constraints derived from business logic. As the framework is intended for a campaign mechanism designed to boost sales, the new price and commission are constrained by their previous values. The time index t represents the current optimization period (i.e., the campaign week), while $t - 1$ refers to the most recent historical period for which data is available (i.e., the week immediately preceding the campaign).

For each SKU i at time t :

$$p_{i,t} \geq \text{cost}_i \quad (3.14)$$

$$p_{i,t} \leq p_{i,t-1} \quad (3.15)$$

$$c_{i,t} \geq c_{\min} \quad (3.16)$$

$$c_{i,t} \leq c_{i,t-1} \quad (3.17)$$

The lower price bound ensures merchant profitability, while the upper bound, $p_{i,t-1}$, enforces the discounting nature of a campaign by requiring the new price to be no higher than its pre-campaign level. Similarly, the commission is bounded by a strategic minimum, $c_{\min}$, and its value from the previous period, $c_{i,t-1}$, allowing the platform to offer a commission discount to sponsor the merchant’s price reduction.

3.6 Demand Model Training and Evaluation

Before evaluating the optimization frameworks, it is essential to establish the predictive quality of the underlying machine learning models. This section details the protocol for both training the demand models and assessing the models.

3.6.1 Training Protocol

With the aims of preventing data leakage and accurately simulating a real world scenario, the models were trained following a time-based splitting methodology. The dataset was divided into a training set for model training, a validation set for parameter tuning, and a final test set for performance evaluation. Hyperparameter tuning for both the Linear Tree and the various components of the Ensemble model was conducted using an automated grid search approach facilitated by the Optuna library.

3.6.2 Performance Metrics

The predictive accuracy of the demand models was evaluated using standard regression metrics. Specifically, the Mean Absolute Percentage Error (MAPE) and the Root Mean Squared Error (RMSE) were calculated on the held-out test set. These metrics provide a clear measure of the models’ ability to generalize to unseen data. The final performance results will be presented in the Chapter 4 (*Results*).

3.6.3 Baseline Comparison

To validate the necessity of the chosen machine learning approaches, their performance was compared against a common econometric baseline: a log-log linear regression model [11]. This model regresses the logarithm of demand against the logarithm of price, a standard method where the resulting coefficient directly represents the price elasticity of demand. This comparison serves to demonstrate the superior predictive power of the more complex models in capturing the non-linearities of the marketplace environment.

3.7 Optimization Evaluation Framework

The primary goal of this research is to compare the effectiveness of the different optimization frameworks. This requires an experimental design that can fairly assess the quality of the decisions produced by each pipeline.

3.7.1 Research Questions

The experimental design aims to answer two fundamental questions:

1. **Methodology Comparison:** For a given machine learning model (the Linear Tree), how does a piecewise approximation approach (PTO) compare to a full mixed-integer integration (CL) in terms of solution quality and computational performance?
2. **Model Impact Comparison:** Does using a predictively superior machine learning model within an approximation framework (PTO-Advanced) lead to better business outcomes than using a simpler model, even when the simpler model is optimized exactly (CL)?

3.7.2 *The Evaluation Oracle*

To ensure a fair and unbiased comparison, the price and commission decisions generated by all frameworks are evaluated using a single "ground truth" simulator. The advanced Ensemble model, having been demonstrated to possess the highest predictive accuracy on the test set, was designated as this "oracle." When evaluating a decision from any framework, the resulting demand is calculated using this oracle model. This approach ensures that the final reported business metrics reflect the most accurate possible estimate of real-world performance, isolating the impact of the optimization methodology from the predictive accuracy of the model used within it.

3.7.3 *Performance Metrics*

The frameworks are compared across two dimensions: solver performance and business outcomes.

- **Solver Performance:** This is measured by the solver status (e.g., optimal, feasible, failed) and the computational time required to find a solution for each SKU.
- **Business Outcomes:** These are the KPIs evaluated by the oracle, including the relative change in platform profit, the relative change in merchant profit, and the strategic decisions made regarding price and commission discount ratios.

3.8 Experimental Setup

The empirical evaluation was conducted in a carefully defined environment to ensure the reliability and reproducibility of the results.

3.8.1 *Hardware and Software Environment*

The experiments are conducted on a computing system with the following specifications:

- CPU: The system is powered by an Apple M2 Pro processor, whose multi-core architecture provides sufficient capability for the model training and optimization tasks performed in this research.
- RAM: The system is equipped with 16 GB of integrated memory, which is adequate to accommodate the datasets and memory requirements of the various machine learning and optimization models.
- GPU (Not Used): While a powerful Graphics Processing Unit (GPU) can be employed to accelerate neural network training, it was not utilized in this thesis. The chosen models, such as gradient boosted trees and linear models, are primarily CPU-bound, and their training processes did not necessitate GPU acceleration.

The software environment includes:

- Operating System: The experiments are conducted on a Unix-based macOS operating system.
- Python: The Python programming language is used for implementing the entire end-to-end pipeline, as it offers a comprehensive ecosystem of libraries for data science and operations research.
- Machine Learning Libraries: Standard libraries including Scikit-learn, LightGBM, and Interpret are utilized for the development, training, and analysis of the demand prediction models.
- Optimization Libraries: The mathematical optimization models are constructed using Pyomo [49], an open-source Python-based modeling language. The OMLT (Optimization and Machine Learning Toolkit [50]) is a key component used to embed the trained Linear Tree model into the CL framework. All optimization problems are solved using the open-source SCIP (Solving Constraint Integer Programs

[51]) solver.

3.8.2 *Simulation Environment and Procedure*

The evaluation of the optimization frameworks is conducted within a carefully constructed simulation environment. This environment uses the historical data from the held-out test set as a static representation of the market conditions for a given week. This approach allows for a controlled, repeatable experiment to conduct a counterfactual analysis: "What would the business outcome have been if the decisions from each optimization framework had been applied during that period?"

The experimental procedure is executed iteratively for each individual SKU in the test set, ensuring that the evaluation is comprehensive across a diverse range of products. The procedure for a single SKU is as follows:

1. **Data Preparation and Contextualization:** The specific feature vector corresponding to the SKU and the time period is extracted from the test set. This vector contains all non-price and non-commission features that describe the market context at that moment (e.g., competitor prices, campaign details, view counts). Data transformations, such as scaling, that were fitted on the training data are applied to this vector to ensure consistency.
2. **Parallel Optimization:** The optimization problem for the SKU is solved independently by each of the three frameworks:
 - CL with the embedded Linear Tree model.
 - PTO using the Linear Tree model with a piecewise approximation.
 - PTO-Advanced using the high-performance Ensemble model with a piecewise approximation.

3. Oracle-Based Metric Calculation: The optimal price and commission pair determined by each of the three frameworks is then passed to the pre-defined evaluation oracle (the high-performance Ensemble model). The oracle calculates the "true" expected demand for that decision. This step is critical as it standardizes the evaluation, ensuring that the business outcomes for all frameworks are judged against the same, most accurate available representation of reality. Based on this oracle-predicted demand, the final business KPIs (platform profit, merchant profit) are computed.
4. Performance and Metadata Logging: For each optimization run, critical performance metadata is recorded. This includes the solver's termination status (e.g., optimal, infeasible) and the total computational time. This data is essential for evaluating the practical viability and robustness of each framework.

Finally, the results from all individual SKU runs are aggregated. This aggregated dataset allows for a comprehensive, statistical comparison of the frameworks' performance across the entire product pool, providing the basis for the conclusions drawn in the subsequent chapters.

3.9 Limitations

While the proposed optimization frameworks hold great promise, it is essential to recognize and address potential limitations that may affect the methodology and research findings. The following are the identified limitations:

- Simulation Fidelity: The primary limitation is that the evaluation is conducted in a simulated environment, not through a live A/B test. The validity of the findings is therefore contingent on the accuracy of the designated oracle model as a perfect proxy for real-world customer behavior. Any inherent bias or inaccuracy in the oracle will directly influence the calculated business outcomes.

- **Data Availability:** The success of any data-driven approach depends heavily on the quality and availability of data. While a SKU selection process was implemented to ensure model reliability, this inherently limits the applicability of the framework. The models cannot be applied to new products or "long-tail" SKUs with insufficient sales history to build a robust demand model.
- **Computational Complexity and Scalability:** The CL framework, which formulates a MINLP, is computationally demanding. While feasible for the scale of this research, deploying this for millions of SKUs in a real-time production environment would present significant scalability challenges. The PTO framework offers a more scalable alternative, but at the cost of approximation error.
- **Solver Capabilities:** This research utilizes SCIP, a high-quality open-source solver. However, leading commercial solvers such as Gurobi or CPLEX often provide superior performance, including faster solve times and more robust handling of complex numerical problems. The use of a commercial solver might alter the computational performance comparison, potentially reducing the number of solver failures observed in the CL framework and narrowing the speed gap with the PTO approach.
- **Generalization to New Market Conditions:** The demand models are trained on historical data. Their predictive performance may degrade if the market undergoes significant structural changes not represented in the training data, such as major shifts in consumer behavior, the entry of a disruptive competitor, or unforeseen macroeconomic events.
- **Real-World Implementability:** Moving from this research prototype to a production-grade system involves additional challenges. Practical considerations include integration with live data streams, ensuring low-latency decision-making required by

the business, monitoring for model drift, and gaining organizational trust and acceptance for an automated decision-making system.

- **Platform's Dual Role as Direct Seller:** An important limitation is that this study solely focuses on optimizing prices and commissions for the third-party merchants within the e-commerce marketplace. Therefore, it does not account for the complex dynamics, potential conflicts of interest, or strategic interactions that arise when the platform itself offers the same or similar SKUs for sale, as a result, directly competing with its merchants. Modeling such scenarios, including strategic decisions like self-preferencing or product imitation, falls outside the current scope of this thesis.

To mitigate these limitations, this research is presented with transparency regarding its scope. The conclusions drawn are framed within the context of the simulated environment, and the discussion emphasizes the potential impact of these limitations on the findings. By acknowledging these points, this research aims to provide valuable insights into the benefits and challenges of integrating advanced machine learning and optimization, enabling future studies to build upon this foundation and address the identified limitations.

4. RESULTS AND ANALYSIS

4.1 Introduction

In this chapter, the empirical results of the research are presented and the analysis of these results will be provided. The chapter begins by evaluating the predictive performance of the machine learning models. As mentioned before, the models were developed to estimate the quantity demanded. This initial analysis is important for validating the core components of the optimization frameworks and establishing a basis for a fair comparison between them. Following the model evaluation, the chapter demonstrates a comprehensive performance analysis of the three distinct optimization pipelines: CL, PTO, and PTO-Advanced. The frameworks are compared on both their computational efficiency and, most importantly, their ability to generate business outcomes, as measured by a robust simulation environment. The analysis aims to provide clear, data driven answers to the core research questions outlined in the methodology.

4.2 Demand Model Performance Analysis

It can be argued that the effectiveness of any optimization framework that relies on predictive inputs is fundamentally constrained by the quality of those predictions. Therefore, this section is dedicated to evaluating the performance of the developed demand models. The goal is to quantify their predictive accuracy, analyze their behavior, and justify the selection of the high-performance Ensemble model as the unbiased evaluation oracle for the subsequent optimization experiments.

4.2.1 Predictive Accuracy Comparison

The three demand models—the econometric Log-Log Baseline, the structured Linear Tree, and the high-performance Ensemble model—were evaluated using MAPE and RMSE as the primary performance metrics. The aggregated results are presented in

Table 4.1.

Table 4.1: *Aggregate Predictive Accuracy of Demand Models.*

Model	MAPE (%)	RMSE
Log-Log Baseline	38.71	95.12
Linear Tree	33.15	66.20
Ensemble Model	29.07	53.94

As the results indicate, the machine learning models significantly outperform the econometric baseline. The Linear Tree model achieves a MAPE of 33.15%, a relative improvement of 14.4% over the Log-Log Baseline. The improvement is even more pronounced in RMSE, where the Linear Tree's score of 66.20 represents a 30.4% reduction in error. This validates the necessity of a machine learning approach to capture complex market dynamics. The Ensemble model demonstrates superior performance, with a MAPE of 29.07% and an RMSE of 53.94. This constitutes a substantial improvement over the baseline (24.9% in MAPE, 43.3% in RMSE) and also improves upon the Linear Tree model by a relative 12.3% in MAPE and 18.5% in RMSE. The high performance of the Ensemble model confirms its ability to most accurately model customer demand, justifying its designation as the evaluation oracle for subsequent experiments.

4.2.2 Analysis of Model Behavior and Prediction Errors

To provide a visual analysis of the performance differences beyond aggregate metrics, the behavior of the models was analyzed visually. Figure 4.1 presents scatter plots of the predictions from both the Linear Tree and Ensemble models against the actual observed demand for a sample SKU. The plots show that while both models capture the general trend, the predictions from the Ensemble model exhibit a tighter clustering around the ideal $y=x$ line, indicating lower variance and higher accuracy on an individual prediction level.

Figure 4.1: *Actual vs. Predicted Demand for a Sample SKU.*





Figure 4.2 further illustrates this by comparing the models' predictions to the actual demand curve for a sample of SKUs, sorted by their actual demand. This visualization highlights that the Ensemble model's predictions (orange line) track the actual demand (black line) more closely and with greater stability than the Linear Tree (blue line) and especially the volatile Baseline model (green line).

Figure 4.2: *Demand Predictions for a Sample of SKUs.*




The nature of the prediction errors is examined in Figure 4.3. This plot shows the raw prediction errors for a sample of SKUs. It can be observed that the Baseline model's errors have the highest variance, while the Ensemble model's errors are more consistently centered around zero, which supports the findings from the aggregate metrics in Table 4.1.

4.2.3 Analysis of Learned Price Elasticity

In addition to predictive accuracy, the economic plausibility of the learned relationships is a crucial aspect of a demand model. Figure 4.4 compares the distribution of


Figure 4.3: *Prediction Errors for a Sample of SKUs.*



the learned price elasticities from the Linear Tree and Ensemble models.

The violin plot reveals a significant difference in the learned elasticity distributions. The Linear Tree model exhibits a wider distribution with a longer tail, suggesting it occasionally learns extremely high elasticity values for certain segments of the data. In contrast, the Ensemble model learns a much tighter and more constrained distribution of elasticities. This "smoothing effect," a common benefit of ensembling, suggests that the Ensemble model produces more robust and economically realistic elasticity estimates, making it less susceptible to overfitting on small data segments. This difference in learned characteristics is hypothesized to have a direct impact on the optimization results analyzed in the subsequent sections. The more extreme elasticity values learned by the Linear Tree are expected to guide the CL and PTO frameworks toward more aggressive price reductions. Conversely, the smoother and more conservative elasticity profile of the Ensemble model is expected to guide the PTO-Advanced framework toward more balanced and moderate optimal decisions.

Figure 4.4: *Distribution of Elasticity Across Models.*




4.2.4 Feature Importance Analysis

Finally, to understand the factors contributing to the Ensemble model's performance, its feature importance is analyzed. Figure 4.5 displays the top feature groups that the model identified as most influential in predicting demand.

The analysis confirms that the model learns logical relationships from the data. As expected, Price Features are by far the most dominant predictor of demand. This is followed by Sales Features and View Features, which act as crucial indicators of a product's current popularity and momentum. The inclusion of Campaign and Category Features further demonstrates the model's ability to incorporate contextual information into its predictions. This not only validates the model's predictive power but also provides confidence that its predictions are based on sound, interpretable market dynamics.

Figure 4.5: *Feature Importance by Feature Group.*



4.3 Optimization Framework Performance Analysis

Having established the performance characteristics of the underlying demand models, the analysis continues with the evaluation of the optimization frameworks themselves. This section assesses the frameworks on their computational performance and robustness, followed by a detailed analysis of the business value and strategic nature of the decisions they produce.

4.3.1 Computational Performance and Solver Robustness

The ability to produce solutions reliably and in a timely manner is a critical aspect of any practical optimization system. The performance of the three frameworks—CL, PTO, and PTO-Advanced—was evaluated based on solver termination status and computational time.

Table 4.2 reveals a significant difference in robustness. The CL framework, which solves a complex MINLP, is the least robust, failing to find an optimal solution in over

Table 4.2: *Solver Status by Framework.*

Framework	Optimal	Failed	Timed Out	Success Rate
CL	2359 (88.8%)	255 (9.6%)	43 (1.6%)	88.8%
PTO	2628 (98.9%)	2 (0.1%)	27 (1.0%)	98.9%
PTO Advanced	2657 (100%)	0 (0.0%)	0 (0.0%)	100%

11% of cases due to failures or timeouts. The PTO framework, which solves a simpler NLP, is significantly more robust with a 98.9% success rate. The PTO-Advanced framework demonstrates perfect robustness, successfully finding an optimal solution for every single SKU. This highlights the practical challenges of implementing the theoretically exact CL approach.


Figure 4.6: *Solver Time by Framework.*

Figure 4.6 illustrates the distribution of solver times. The CL framework is not only the least robust but also the most computationally expensive, with the widest distribution and highest median solve time. This is an expected result due to the complexity of solving the embedded MINLP. The standard PTO framework is the fastest, benefiting

from the simplest model structure. The PTO-Advanced framework finds a middle ground, being slightly slower than the standard PTO due to the more complex surrogate curve it approximates, but still significantly faster and more reliable than the CL approach.

4.3.2 Business Outcome Analysis

The primary evaluation of an optimization framework is the business value it generates. The following analysis uses the oracle model to evaluate the expected profit changes for both the platform and its merchants based on the decisions from each framework.

Figure 4.7: *Relative Merchant Profit Change by Framework.*





Figure 4.7 shows the distribution of relative profit changes for merchants. This is arguably the most critical metric for a sustainable marketplace. The PTO-Advanced framework is the best performing framework in this regard, delivering the highest median profit increase for merchants and exhibiting a distribution that is almost entirely positive.

The CL and PTO frameworks, while offering a positive median, show a much wider and riskier distribution, with a significant number of SKUs resulting in substantial profit losses for the merchant. This suggests their aggressive strategies can negatively impact merchant profitability.

Figure 4.8: *Relative Platform Revenue Change by Framework.*



From the platform’s perspective (Figure 4.8), the results are more nuanced. The CL and PTO frameworks generate a higher median revenue increase. However, this comes at the cost of higher variance and a greater number of negative outcomes. The PTO-Advanced framework provides a more conservative but significantly more stable revenue increase, with a much tighter distribution and almost no instances of revenue loss. This presents a clear trade-off between risk and reward for the platform.

4.3.3 Analysis of Strategic Decisions

To understand the reasons for the different business outcomes, we analyze the nature of the strategic decisions they make regarding price and commission discounts.


Figure 4.9: *Price Discount Ratio by Framework.*



Figure 4.9 confirms the hypothesis formulated in Section 4.2.3. The CL and PTO frameworks, guided by the extreme elasticities learned by the Linear Tree model, recommend significantly more aggressive price discounts, with a median discount ratio around 14%. In a real-world scenario, such deep, uniform discounts might be rejected by merchants or may not be financially viable. The PTO-Advanced framework, guided by the smoother elasticity of the Ensemble model, proposes far more conservative and realistic price discounts, with a median of approximately 5%.

The commission discount strategy, shown in Figure 4.10, reflects the same pattern. To incentivize the deep price cuts, the CL and PTO frameworks must offer substantial commission discounts to merchants, which represents a direct cost to the platform. The PTO-Advanced framework, by recommending more moderate price changes, can achieve its goals with smaller commission discounts, thus preserving more platform revenue. This demonstrates a more balanced approach to negotiation.

Figure 4.10: *Commission Discount Ratio by Framework.*



4.4 Discussion and Interpretation

This section synthesizes the preceding results to answer the core research questions and discuss the strategic implications for the e-commerce marketplace.

4.4.1 Answering the Research Questions

The empirical results provide clear answers to the research questions posed in the methodology.

4.4.1.1 Methodology Comparison (CL vs. PTO)

For a given ML model (the Linear Tree), how does a piecewise approximation (PTO) compare to a full integration (CL)? The results show that while both frameworks produce strategically similar (i.e., aggressive) decisions, the PTO framework is significantly better in terms of practical implementation. It is faster, significantly more robust (99% vs 89% success rate), and avoids the computational complexity of the CL's MINLP

formulation. The approximation error introduced by the piecewise approach appears to be an acceptable trade-off for the large gains in speed and reliability.

4.4.1.2 Model Impact Comparison (CL vs. PTO-Advanced)

Does using a predictively superior model in an approximate framework (PTO-Advanced) lead to better outcomes than using a simpler model in an exact framework (CL)? The results clearly indicate that this is the case. The PTO-Advanced framework, despite using an approximation, generates far superior business outcomes. It produces more realistic and strategically sound decisions (moderate discounts), leads to higher and more stable profits for merchants, and provides a more reliable revenue stream for the platform. This finding is critical: the predictive accuracy and learned economic behavior of the underlying demand model are more important for generating high-quality business decisions than the mathematical exactness of the optimization method. An exact optimization of an inadequate or overly aggressive model, as the Linear Tree was shown to be, yields poor results.

4.4.2 *Synthesis of Findings and Strategic Implications*

The central finding of this research is that the PTO-Advanced framework represents the most effective approach for implementing a dynamic price and commission system among the alternatives tested. It successfully balances the trade-offs between predictive accuracy, optimization complexity, and business value.

The aggressive strategies produced by the CL and standard PTO frameworks, while appearing to generate high platform revenue in the median case, are strategically flawed. They rely on deep, often unrealistic price cuts that could damage merchant relationships and profitability, posing a long-term risk to the marketplace ecosystem. Their high variance also means they are unreliable.

In contrast, the PTO-Advanced framework, powered by the robust and economically sensible Ensemble model, embodies the principles of the NBS. It finds a "win-win" solution that provides a significant and reliable profit boost to merchants while ensuring a stable and positive revenue increase for the platform. This approach is not only more profitable in a holistic sense but is also more likely to be adopted and trusted by merchants, ensuring the long-term health and success of the marketplace.

4.5 Limitations and Recommendations

While this research provides a rigorous comparison of different optimization methodologies, it is important to acknowledge its limitations.

- **Simulation Fidelity:** The primary limitation is that the evaluation is conducted in a simulated environment, not a live A/B test. The results are therefore contingent on the accuracy of the oracle model as a proxy for real-world customer behavior.
- **Data Availability:** The models cannot be applied to new products or "long-tail" SKUs with insufficient sales history to build a robust demand model.
- **Solver Capabilities:** This thesis utilized the open-source solver SCIP. It is a limitation that no commercial solver, such as Gurobi or CPLEX, was applied. Commercial solvers often provide superior performance for complex MINLP problems. Their use could potentially improve the success rate and reduce the solution time of the CL framework, although the fundamental strategic differences driven by the underlying model would likely remain.
- **Static Environment:** The study assumes that the historical data used for training is representative of future market conditions. External factors not captured in the feature set, such as macroeconomic shifts or major changes in competitor strategy, could affect the real-world performance of the proposed frameworks.

It is recommended that future work should focus on deploying the most promising framework (PTO-Advanced) in a controlled, live A/B testing environment to validate these findings with real transaction data.

4.6 Practical Applications

The PTO-Advanced pipeline, in particular, offers a scalable and robust template for developing an automated decision support system. Such a system could be used to generate data driven campaign offers for thousands of merchants. It, therefore, seeks to optimize for the dual objectives of merchant profitability and platform revenue. This can lead to improved merchant engagement, higher campaign uptake, and a more profitable marketplace for all participants in an efficient manner.

4.7 Conclusion

In conclusion, this chapter has presented a comprehensive analysis of the results from the developed optimization frameworks. The initial evaluation confirmed the superior predictive accuracy of the Ensemble demand model. The subsequent comparative analysis of the optimization frameworks demonstrated that the PTO-Advanced pipeline, which leverages this superior model within an efficient approximation framework, delivers the best combination of computational robustness, strategic intelligence, and balanced business outcomes. The findings indicate that the quality of the machine learning model is the most critical component in a successful PTO system, and that the PTO-Advanced approach provides a practical tool for enhancing strategic decision making in a dynamic e-commerce environment.

5. CONCLUSIONS

5.1 Summary of Research

In this thesis, the DPCOP was attempted to be addressed by developing and evaluating an optimization framework for a large scale e-commerce marketplace. The aim was to create a system that could dynamically set prices and commissions to improve platform revenue and merchant profit at the same time. At the heart of this thesis, there is a comparison between the pricing optimization methodologies: an integrated CL approach and a decoupled PTO approach. Both of the methods were built upon underlying machine learning models, namely a Linear Tree and a high-performance Ensemble model. These models were designed to predict the quantity demanded for the given product. The models' final decision is guided by a NBS objective function to make sure the model optimizes for a balanced outcome, favoring both the e-commerce platform and the merchants. The frameworks were then evaluated in a comprehensive simulation environment.

5.2 Contributions

This study attempted to make several key contributions in the field of machine learning and optimization in the e-commerce context:

- *Framework Comparison:* The research tries to provide an elaborate and empirical comparison between an integrated optimization approach (CL), where the machine learning model is embedded as mathematical constraints, and a decoupled approach (PTO), where the model's predictions are used to create a linear approximation of the real demand to be able to use within the optimization framework. The results demonstrate the superior practical performance of the PTO approach in terms of computational performance and solver robustness.
- *Model Impact Analysis:* A central contribution is the finding that the predictive

quality and economic sensibility of the underlying machine learning model are more critical for generating high value business outcomes than the mathematical exactness of the optimization method. The study presented that the PTO-Advanced framework, with a more accurate Ensemble model, produced significantly better and more balanced results than the CL framework, which performed an exact optimization on a more basic but less accurate model.

- *Practical Application Template:* The research might be utilized to develop a practical plan for implementing a dual objective optimization system in similar optimization problems. The PTO-Advanced pipeline, which combines a high-performance demand model with an efficient piecewise approximation, might yield a scalable and effective solution for generating data driven campaign offers in a real-world e-commerce setting.

5.3 Practical Implications

The findings of this research have several practical implications for e-commerce marketplaces in several ways:

- *Improved Merchant Relations and Profitability:* The more balanced approach of the PTO-Advanced framework might lead to higher and, more importantly, more stable profits for the sellers in the marketplace. The platform can build trust and improve long-term merchant relationships by avoiding overly aggressive discount strategies, which can be argued that it is critical for the health of the marketplace ecosystem.
- *Data Driven Campaign Strategy:* The proposed framework might be used to automate and optimize the generation of campaign offers in a similar setup. The platform can provide data driven price and commission recommendations that are optimized to benefit both parties; as a result, it is more likely to achieve higher

campaign adoption rates and overall effectiveness compared to using static rules for campaign creation.

- *Platform Risk Management:* The analysis might show that the PTO-Advanced framework provides a more stable and predictable revenue stream for the platform. The platform can manage its revenue more effectively by avoiding the high variance and significant downside risk associated with the more aggressive strategies presented by the alternative frameworks.

5.4 Future Research

Even though this research attempted to demonstrate the potential of the proposed frameworks, several directions for future research might be explored to build upon these findings:

- *Live A/B Testing:* The most critical next step might be to move beyond the simulation environment to test the frameworks in the presence of real-world data. Deployment of the frameworks presented in the thesis in a live and controlled A/B test is crucial to validate the findings of this thesis against real-world merchant behavior and to measure the true business impact on merchant and platform KPIs.
- *Modeling SKU Interdependencies:* This research optimized each SKU independently. A significant area for future work is to extend the framework to a multi-SKU or portfolio level optimization. This would require modeling cross-price elasticities to account for complex interactions such as product substitution (e.g. cannibalization effect) and complementarity, as their inclusion would lead to more globally optimal decisions for the platform.
- *Advanced Solver Integration:* The performance of the optimization frameworks was limited by the capabilities of the open-source SCIP solver. Future research could

create an environment to reevaluate the CL approach using a state-of-the-art commercial solver, such as Gurobi or CPLEX. This would contribute to determining whether superior solver technology can overcome the robustness and speed issues, potentially making the exact CL method more competitive with the PTO approach.

- *Dynamic and Online Learning:* The demand models in this thesis were trained in a static dataset in a batch manner. Future work could explore the integration of dynamic and online learning models. These models might be structured to be updated more frequently to adapt to rapid shifts experienced in market conditions, competitor actions, or customer behavior in terms of both demand modelling and optimization. Eventually, this more advanced approach might lead to more responsive and accurate decision-making over time.
- *Platform as Direct Seller and Competitive Dynamics:* The current framework only focuses on optimizing for the third-party merchants; therefore, it treats the platform primarily as an intermediary. Future research could extend this model to incorporate scenarios where the e-commerce platform itself sells the same SKUs. This additional dynamic would result in a situation where the platform was directly competing with its own merchants. This would involve modeling strategic decisions such as the platform's self-preferencing behaviors, product imitation, and pricing strategies for its own private label products. Such an extension would likely require more advanced game-theoretic approaches beyond the current cooperative NBS framework to capture these complex competitive dynamics and could offer insights into how platforms balance their marketplace and retail functions.

5.5 Conclusion Statement

In conclusion, the thesis successfully developed and evaluated frameworks for solving the DPCOP. The research demonstrated that a PTO approach, when coupled with a

high-performance machine learning model, provides a robust and more scalable solution. The research findings presented the superiority of the PTO-Advanced framework in delivering balanced and strategically sound results. These results would benefit both the merchants and the platform. The key finding is that the quality of the underlying machine learning model is the most critical factor for success in such a system. It also substantiates that having a better-performing machine learning model is more important than the mathematical exactness of the optimization algorithm itself. This research provides a strong foundation and a practical guide for e-commerce platforms seeking to implement data driven and automated decision-making systems. The framework presented in this thesis might be utilized as a blueprint for developing such systems.

REFERENCES

- [1] J. F. Nash, “The bargaining problem,” *Econometrica*, vol. 18, p. 155, Apr. 1950.
- [2] B. Niu and J. Zhao, “E-commerce industry: A comprehensive analysis of its competitive advantages,” *Highlights in Business, Economics and Management*, vol. 15, p. 93–98, June 2023.
- [3] G. Parker, M. Van Alstyne, and S. Choudary, *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*. W. Norton, 2016.
- [4] T. R. Eisenmann, G. Parker, and M. W. Van Alstyne, “Strategies for two-sided markets,” *Harvard Business Review*, vol. October, 2006. Available at SSRN: <https://ssrn.com/abstract=2409276>.
- [5] P. Andreoli-Versbach and J. Gans, “Interplay between amazon store and logistics,” *European Competition Journal*, vol. 20, p. 460–497, Mar. 2024.
- [6] O. Hartzell and A. Haupt, “Platform preferencing and price competition i: Evidence from amazon.” <https://ssrn.com/abstract=5126918>, Jan 2025. Available at SSRN: <https://ssrn.com/abstract=5126918> or <http://dx.doi.org/10.2139/ssrn.5126918>.
- [7] A. N. Elmachtoub and P. Grigas, “Smart “predict, then optimize”,” *Management Science*, vol. 68, p. 9–26, Jan. 2022.
- [8] D. Maragno, H. Wiberg, D. Bertsimas, I. Birbil, D. den Hertog, and A. Fajemisin, “Mixed-integer optimization with constraint learning,” *Operations Research*, vol. 73, 12 2023.
- [9] T. Vanderschueren, T. Verdonck, B. Baesens, and W. Verbeke, “Predict-then-optimize or predict-and-optimize? an empirical evaluation of cost-sensitive learning strategies,” *Information Sciences*, vol. 594, pp. 400–415, 2022.
- [10] Z. Fatemi, M.-T. T. Huynh, E. Zheleva, Z. Syed, and X. Di, “Mitigating cold-start forecasting using cold causal demand forecasting model,” *ArXiv*, vol. abs/2306.09261, 2023.
- [11] J. M. Wooldridge, *Introductory Econometrics: A Modern Approach*. Cengage Learning, 6th ed., 2016.
- [12] H. R. Varian, “Big data: New tricks for econometrics,” *Journal of Economic Perspectives*, vol. 28, p. 3–28, May 2014.

- [13] S. Athey and G. W. Imbens, “Machine learning methods that economists should know about,” *Annual Review of Economics*, vol. 11, p. 685–725, Aug. 2019.
- [14] K. Safonov, “Neural network approach to demand estimation and dynamic pricing in retail,” 2024.
- [15] S. Mullainathan and J. Spiess, “Machine learning: An applied econometric approach,” *Journal of Economic Perspectives*, vol. 31, p. 87–106, May 2017.
- [16] A. Barredo Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, and F. Herrera, “Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai,” *Information Fusion*, vol. 58, pp. 82–115, 2020.
- [17] D. Bertsimas and J. Dunn, “Optimal classification trees,” *Machine Learning*, vol. 106, no. 7, pp. 1039–1082, 2017.
- [18] S. Ikeda, N. Nishimura, N. Sukegawa, and Y. Takano, “Prescriptive price optimization using optimal regression trees,” *Operations Research Perspectives*, vol. 11, p. 100290, 2023.
- [19] G. Ke, Q. Meng, T. Finley, *et al.*, “Lightgbm: A highly efficient gradient boosting decision tree,” in *Neural Information Processing Systems*, 2017.
- [20] H. Nori, S. Jenkins, P. Koch, and R. Caruana, “Interpretml: A unified framework for machine learning interpretability,” *ArXiv*, vol. abs/1909.09223, 2019.
- [21] T. Hastie and R. Tibshirani, “Generalized additive models,” *Statistical science*, vol. 1, no. 3, pp. 297–310, 1986.
- [22] Y. Lou, R. Caruana, J. Gehrke, and G. Hooker, “Accurate intelligible models with pairwise interactions,” in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 623–631, 2013.
- [23] T. Hastie, R. Tibshirani, J. H. Friedman, and J. H. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, vol. 2. Springer, 2009.
- [24] K. H. Lee, M. Abdollahian, S. Schreider, and S. Taheri, “Supply chain demand forecasting and price optimisation models with substitution effect,” *Mathematics*, vol. 11, p. 2502, May 2023.
- [25] D. Bertsimas and N. Kallus, “The power and limits of predictive approaches to observational data-driven optimization: The case of pricing,” *INFORMS Journal on Optimization*, vol. 5, no. 1, pp. 110–129, 2023.
- [26] J.-C. Rochet and J. Tirole, “Two-sided markets: a progress report,” *The RAND journal of economics*, vol. 37, no. 3, pp. 645–667, 2006.

- [27] A. Hagi, T. Teh, and J. Wright, “Should platforms be allowed to sell on their own marketplaces?,” *The RAND Journal of Economics*, vol. 53, p. 297–327, May 2022.
- [28] F. Zhu and Q. Liu, “Competing with complementors: An empirical look at amazon.com,” *Strategic management journal*, vol. 39, no. 10, pp. 2618–2642, 2018.
- [29] M. Asghari, H. Afshari, S. Mirzapour Al-e-hashem, A. M. Fathollahi-Fard, and M. A. Dulebenets, “Pricing and advertising decisions in a direct-sales closed-loop supply chain,” *Computers & Industrial Engineering*, vol. 171, p. 108439, 2022.
- [30] K. Miettinen, *Nonlinear multiobjective optimization*, vol. 12. Springer Science & Business Media, 1999.
- [31] G. P. Cachon and M. A. Lariviere, “Supply chain coordination with revenue-sharing contracts: strengths and limitations,” *Management science*, vol. 51, no. 1, pp. 30–44, 2005.
- [32] A. Sharma, G. Dwivedi, and A. Singh, “Game-theoretic analysis of a two-echelon supply chain with option contract under fairness concerns,” *Computers & Industrial Engineering*, vol. 137, p. 106096, 2019.
- [33] G. Peng, X. Tian, Q. Chen, *et al.*, “Online cooperative promotion and cost sharing policy under supply chain competition,” *Mathematical Problems in Engineering*, 2016.
- [34] Z.-P. Li, J.-J. Wang, S. Perera, and J. J. Shi, “Coordination of a supply chain with nash bargaining fairness concerns,” *Transportation Research Part E: Logistics and Transportation Review*, vol. 159, p. 102627, 2022.
- [35] H. Song and X. Gao, “Green supply chain game model and analysis under revenue-sharing contract,” *Journal of Cleaner Production*, vol. 170, pp. 183–192, 2018.
- [36] J. Mandi, J. Kotary, S. Berden, M. Mulamba, V. Bucarey, T. Guns, and F. Fioretto, “Decision-focused learning: Foundations, state of the art, benchmark and future opportunities,” *Journal of Artificial Intelligence Research*, vol. 80, p. 1623–1701, Aug. 2024.
- [37] B. Wilder, B. Dilkina, and M. Tambe, “Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, p. 1658–1665, July 2019.
- [38] J. Liu, T. Xu, J. He, and C. Fang, “Smart predict-then-optimize method with dependent data: Risk bounds and calibration of autoregression,” *arXiv preprint arXiv:2411.12653*, 2024.
- [39] J. Kotary, F. Fioretto, P. Van Hentenryck, and B. Wilder, “End-to-end constrained optimization learning: A survey,” *arXiv preprint arXiv:2103.16378*, 2021.

- [40] G. Zarpellon, *Machine learning algorithms in Mixed-Integer Programming*. PhD thesis, Polytechnique Montréal, May 2020.
- [41] A. Capitaine, M. Haddouche, E. Moulines, M. I. Jordan, E. Boursier, and A. Durmus, “Online decision-focused learning,” 2025.
- [42] A. O. Fajemisin, D. Maragno, and D. den Hertog, “Optimization with constraint learning: A framework and survey,” *Eur. J. Oper. Res.*, vol. 314, pp. 1–14, 2021.
- [43] L. Sang, Y. Xu, and H. Sun, “Encoding carbon emission flow in energy management: A compact constraint learning approach,” *IEEE Transactions on Sustainable Energy*, vol. 15, pp. 123–135, 2023.
- [44] M. Dolányi, K. Bruninx, J. Toubreau, and E. Delarue, “Capturing electricity market dynamics in strategic market participation using neural network constrained optimization,” *IEEE Transactions on Power Systems*, vol. 39, pp. 533–545, 2024.
- [45] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *Knowledge Discovery and Data Mining*, 1996.
- [46] J. W. Dunn, *Optimal trees for prediction and prescription*. PhD thesis, Massachusetts Institute of Technology. Operations Research Center; Sloan School of Management, Cambridge, MA, 2018.
- [47] I. E. Grossmann and F. Trespalacios, “Systematic modeling of discrete-continuous optimization models through generalized disjunctive programming,” *AIChE Journal*, vol. 59, p. 3276–3295, May 2013.
- [48] E. M. L. Beale and J. J. H. Forrest, “Global optimization using special ordered sets,” *Mathematical Programming*, vol. 10, pp. 52–69, 1976.
- [49] W. E. Hart, J.-P. Watson, and D. L. Woodruff, “Pyomo: modeling and solving mathematical programs in python,” *Mathematical Programming Computation*, vol. 3, pp. 219–260, 2011.
- [50] F. Ceccon, J. Jalving, J. Haddad, A. Thebelt, C. Tsay, C. D. Laird, and R. Misener, “Omlt: Optimization & machine learning toolkit,” *J. Mach. Learn. Res.*, vol. 23, pp. 349:1–349:8, 2022.
- [51] D. Anderson, K. Bestuzheva, *et al.*, “The scip optimization suite 7.0,” tech. rep., Optimization Online, Mar 2020. Optimization Online preprint.