



*Research article***Causal inference and quantile regression for measuring the impact of pricing strategies on business performance****Masad A. Alrasheedi***

Department of Management Information Systems, College of Business Administration, Taibah University, Madinah, Saudi Arabia

* **Correspondence:** Email: mrshedi@taibahu.edu.sa.

Abstract: Pricing strategies are critical to a business's success. Still, their effects are hard to study because many factors influence a product's price, including customer behavior, promotions, demand, inventory, and market diversity. This research presents a new approach to distribution-aware causal pricing optimization (DCPO) that moves beyond evaluating pricing strategies by averages. The new approach integrates the diversity of pricing strategies and their effects on different levels of business performance into one unified framework. The new approach provides four new methods for research. First, it models pricing strategies and their effects at the treatment level using response and propensity models, along with doubly robust and orthogonal estimation, to reduce confounding and modeling error. Second, the new approach models pricing strategy effects at lower and higher levels of business performance, including the mean and median. Third, a new approach to causal pricing DCPO, is presented. The new objective function evaluates pricing strategies by rewarding average causal business improvement and penalizing variability in treatment across quantiles of business performance. The new approach to pricing strategy evaluation is integrated into a unified framework and assessed along with other methods on the Dunnhumby and Rossmann retail datasets. The new framework yields average treatment-effect estimation errors of 0.084 and 0.089 for the Dunnhumby and Rossmann datasets, respectively. Additionally, the new framework achieves R^2 values of 0.958 and 0.953 for the Dunnhumby and Rossmann datasets, respectively. Cross-dataset evaluation increases treatment-effect estimation error by only 9.12% and decreases R^2 by 0.96%. On Dunnhumby, the framework yields an 11.48% revenue improvement, an 11.82% profit improvement, and an 8.54% conversion-rate improvement, with a DCPO of 0.931. Ablation, robustness, and statistical analyses corroborate the importance of the framework's main components.

Keywords: causal inference; pricing strategy; quantile regression; treatment-effect estimation; business analytics; retail pricing

Mathematics Subject Classification: 62P20, 68T05

1. Introduction

Pricing strategy is fundamental to business performance as it impacts demand, revenues, and profits, as well as factors such as competitive positioning. Digital transformation across retail and service industries has shifted pricing from a relatively fixed decision to a fluid optimization challenge that involves advanced transactional data, AI, and real-time market data. As a result, organizations adopt a fluid, data-based pricing approach that responds to changes in customer behavior, market uncertainty, competitor pricing, and inventory. Dynamic pricing has advanced the integration of machine learning with optimization to enhance the value of pricing decisions for the firm over the long run. Similarly, algorithmic pricing has become a promising area of research to facilitate intelligent pricing decisions in highly competitive and predominantly digital markets.

Predictive analytics has certainly advanced demand forecasting and pricing optimization, but determining the real effects of pricing strategies on business outcomes is still problematic. Pricing data relies on observational business data, which often accompanies price promotions, seasonality, diverse customers, and regional demand discrepancies. As a consequence, many predictive models discover statistical associations rather than real causal relationships. Overcoming this obstacle has recently been accomplished by integrating causal inference and pricing analytics to derive and measure the effects of pricing strategy interventions. Causal forecasting has shown an increased ability to predict price-demand relationships by decoupling causal dependencies from purely predictive relationships [1]. Causal experimental design has also been proposed to price policy optimization while providing reliable estimates of the effects of pricing interventions within the domain of real-world business limits [2]. Offline causal pricing and feature pricing have allowed businesses to set pricing policies based on the analysis of prior observational data without the need for costly causal experiments [3]. Overall, these innovations show that integrating causal inference into pricing strategy analysis provides a more trustworthy foundation than predictive models based on correlation.

Recent innovations in pricing analytics have utilized heterogeneous treatment-effect estimation and personalized decision-making. In many retail settings, Bayesian pricing frameworks can model a nonlinear demand function with varying price elasticities [4]. Pricing algorithms that utilize context and are designed to change dynamically have been developed to make optimal pricing decisions based on the current state of the customers and the market [5]. Similarly, personalized pricing systems have been developed to use real-world data to create pricing systems that improve customer satisfaction [6]. Additionally, algorithm-based pricing systems have been developed to address automated pricing in large-scale web-based markets. These pricing systems incorporate operational efficiency and address market regulations [7]. All of these systems and frameworks indicate a major change in the pricing systems and methodologies. In particular, these systems show that robust causal methodologies are increasingly needed to evaluate the impact of pricing systems in complicated business systems.

Recent advances in quantile treatment-effect estimation provide ample opportunity to construct pricing policies that account for the full distribution of customer and product outcomes. Pricing interventions can be evaluated at any point in the outcome distribution [8]. Furthermore, unlike the traditional policies learned through quantile analysis, the pricing optimization policies learned through the integration of policy learning and quantile analysis are designed to function within the bounds of real-world business practices [9]. Additionally, integrating deep learning tools with causal analysis offers great potential to address complex combinatorial treatment-effects in large-scale

business experiments while maintaining a focus on cutting-edge work with large, sparse datasets [10]. Finally, the integration of reinforcement learning facilitates real-time, adaptive retail pricing and promotional decision-making within a framework that accounts for the dynamic and uncertain nature of modern markets [11]. The confluence of these developments demonstrates the potential to address businesses' pricing evaluation needs better than current revenue optimization frameworks.

Inspired by recent advancements, this study seeks to synthesize analytical tools to incorporate causal inference and distribution-cognizant quantile modeling to assess the impact of pricing on business performance. The proposed framework brings together orthogonal treatment-effect estimation, confounder adjustment, quantile response modeling, and a distribution-aware causal pricing optimization (DCPO) to assess the average and heterogeneous pricing effects at the retail level. The framework seeks to address the pricing challenges to enhance business performance at the most desirable customer-product pricing combinations, all while controlling for confounding factors. The proposed methodology is assessed on the Dunhumby and Rossmann retail datasets and is tested against dominant pricing and causal inference methodologies to assess its effectiveness, robustness, and efficiency.

The main findings of this study are as follows:

- Integration of several previously separate approaches to orthogonal treatment-effect estimation, confounder adjustment, quantile response modeling, and a DCPO offers a unified and comprehensive causal pricing framework.
- A novel observational causal inference framework is introduced to overcome confounding bias in estimating pricing effects by modeling treatment assignment and the response to treatment-specific outcomes.
- A novel distribution-aware quantile modeling approach is designed to explain pricing effects of varying magnitude across different business performance quantiles as a more advanced form of decision support compared to average treatment-effect modeling.
- Experiments conducted using two large retail datasets validate claimed advantages of superior causal estimation accuracy, business performance, distributional robustness, computational efficiency, and generalization across datasets compared to existing methods in causal pricing and related domains.

The remainder of this paper is laid out as follows. Section 2 summarizes background research. Section 3 outlines the research methodology. Section 4 presents the research results and analyzes the findings. Section 5 provides a summary of the research and offers some advancements to consider for the continuation of this research.

2. Literature review

Effective pricing strategies are imperative to increase enterprise profitability, maintain clientele, and ensure competitiveness in the market over the long run. Pricing strategies employed to date are grounded in the manipulation of demand elasticity through the lens of econometrics, generalizing consumer behavior. However, the recent availability of large-scale transaction data from retail outlets has revealed extreme variations in consumer behavior. It has therefore encouraged the use of causal inference to identify the real effects of pricing from pricing correlations caused by confounding

variables. Double machine learning, paired with forecasting, has been used to conduct off-policy pricing analysis [1]. Additionally, to determine optimal pricing, estimate the effects of pricing, and maximize expected revenue, pricing experiments that integrate causal frameworks have been conducted [2]. The primary focus of these studies has been the development of decision support systems that provide pricing rationale, rather than pricing predictive models.

Heterogeneous treatment-effects can now be assessed with the recent advancements made in causal machine learning. Pricing policies can now be derived by utilizing offline causal learning made possible by the availability of historical observational data containing censored demand [3]. Pricing elasticity and variations in the sales response to pricing across different retail settings have been modeled using Bayesian frameworks [4]. Additionally, bandit frameworks have enabled the dynamic design of promotional pricing systems based on uncertainty about optimal pricing [5]. These frameworks have focused on pricing goods and services while incorporating customer behavior that is more sophisticated than traditional econometric approaches.

The ability to leverage customer behavior has generated great interest in personalized and algorithmic pricing. Numerous studies have shown that online and offline grocery pricing strategies differ greatly [12]. Demonstrating pricing policy segmentation for diverse customer groups, personalized pricing has been shown to employ context-sensitive promotional pricing [13]. Additionally, the structural framework of algorithmic pricing has been shown to affect consumers' price sensitivity, purchasing behavior, and even rivalry amongst retailers [12]. While these studies have made great strides in the field, the focus of the existing studies has been on personalized recommendations and revenue optimization, and they have not examined the pricing effects on different levels of business performance in a causal manner.

Analyzing the impact of prices on different business levels using the framework of quantile regression offers a unique opportunity to study the effects of heterogeneous pricing using conditional distributions instead of average outcomes [14]. Recently, the estimation of prices has involved assessing pricing policies to study the effects of pricing on distribution in a more refined manner using the concept of quantile treatment-effects [15]. Additionally, in the field of pricing policy, generalized causal forests offer the ability to study heterogeneous treatment-effects. At the same time, the structure of doubly robust estimation allows the researcher to study the effect of pricing policies using an integrated framework of pricing treatment and outcomes, even in a context of uncertainty, and with a greater degree of flexibility than has previously been offered in the pricing field [16]. While great headway has been made in the pricing field using these techniques, the ability to study pricing policies and their effects on disparate business outcomes has yet to be offered.

Advances in machine learning emphasize the importance of flexible nonlinear representations for capturing complex patterns in high-dimensional data. For instance, one paper used flexible representations to categorize heterogeneous online data [17]. Ensemble learning approaches, on the other hand, show that complex input-output relationships can be captured by combining a plurality of simultaneously learned representations [18]. Transformer-based feature learning has also been used to capture heterogeneous data [19]. More recent architectures that integrate multiple representations and attention show the same emphasis on modeling heterogeneity [20,21]. Clustering and transfer learning have also been used to model heterogeneity, given the variability in the conditions under which the data was observed [22,23]. Although these papers apply to domains other than retail pricing, they highlight a broader trend in the literature to move away from rigid representation systems and toward flexible

learning frameworks that capture non-linear, context-dependent relationships.

Flexible predictive modeling can capture patterns in data; however, it does not validate that an observed intervention leads to a change in the outcome. This is particularly relevant in business pricing, where price changes are observed alongside customer behavior, product attributes, promotional campaigns, inventory levels, seasonality, and past demand. It is possible to predict revenue or sales quite accurately, but the price change might still be misleading. Recent causal machine-learning research incorporates more flexible methods for nuisance function estimation alongside treatment-effect estimation to capture heterogeneous intervention effects while relaxing reliance on strong parametric assumptions about the response. This motivates the treatment-effect component of the current framework, where doubly robust and orthogonal estimation integrates propensity and treatment-outcome models, rather than interpreting predictive relations as causal pricing effects.

Another related development is the growing reliance on distribution-sensitive and quantile-based analysis when an average response cannot capture heterogeneous outcomes. In business data, customers and products in the lower, central, and upper regions of the performance distribution may respond differently to the same business intervention. Quantile regression helps overcome this limitation by estimating conditional outcomes beyond the conditional mean, revealing response patterns that average-effect models may conceal. This matters for pricing analytics because a business intervention that yields a favorable average revenue response may have a much smaller, or even a completely different, effect on low- and high-performing observations. Traditional quantile prediction, however, remains associational; that is, it does not help overcome the challenge of confounding unless it is specifically addressed. The authors propose a new framework that addresses this challenge by first estimating confounder-adjusted heterogeneous causal pricing effects and then describing their behavior across multiple conditional quantiles. In Table 1, we review the pricing literature focusing on causal pricing, pricing policies, and the effects of pricing on business outcomes, to highlight gaps in the existing literature.

While recent advancements have been made in causal inference, dynamic pricing, and the estimation of heterogeneous treatment-effects, some important gaps remain in pricing research. Most existing methods of pricing analytics consider average treatment-effects or the optimization of expected revenues. As a result, pricing analytics have a poor understanding of how a pricing policy may motivate different levels of performance across the business. Similarly, many methods of causal inference focus on estimating treatment-effects and ensuring bias, but do not focus on modeling shifts in business outcomes. On the other hand, many methods of pricing analytics and inference that focus on shifts across the price distribution do not address pricing policy evaluation and observation bias. Most methods are incapable of jointly estimating the pricing effects and their shifts across different customer/product segments. Finally, a pricing policy evaluation model that integrates causal effects, shifts across the price distribution, and is computationally and practically efficient has yet to be developed.

Table 1. Comparison of representative studies on causal inference and pricing analytics.

Study	Method	Contribution	Limitation
Causal forecasting [1]	Double ML + Transformer	Causal demand forecasting	No quantile analysis
Causal pricing experiments [2]	Causal experimental design	Balances revenue and treatment-effects	Requires controlled experiments
Offline pricing [3]	Offline causal learning	Pricing from observational data	Limited distributional analysis
Bayesian pricing [4]	Hierarchical Bayesian model	Models heterogeneous elasticity	Model assumptions required
Contextual bandits [5]	Contextual bandit learning	Adaptive promotional pricing	Exploration dependent
Personalized pricing [13]	Contextual bandits	Customer-specific pricing	Limited causal interpretation
Algorithmic pricing [12]	Algorithmic pricing	Studies automated pricing	No distributional treatment-effects
Double ML [24]	Orthogonal ML estimation	Robust causal estimation	Not pricing-specific
Generalized causal forests [16]	Generalized random forests	Captures nonlinear heterogeneity	No quantile modeling
Causal forests [25]	Random forests	Individual treatment estimation	High computational cost
Continuous treatments [26]	Doubly robust estimation	Continuous treatment modeling	Limited retail applications
Quantile regression [14]	Regression quantiles	Models heterogeneous outcomes	No causal framework
Quantile treatment-effects [15]	QTE estimation	Distributional treatment-effects	Requires strong assumptions
Doubly robust quantiles [27]	DR quantile estimation	Robust quantile estimation	Limited retail validation
Meta-learners [28]	T/X/S learners	Flexible causal estimation	Requires multiple models
Orthogonal random forest [29]	Orthogonal random forest	Debiased heterogeneous estimation	Complex implementation

These research directions describe an important methodological gap. Currently, machine learning models can represent heterogeneous data in numerous ways, distinguish interventional effects from predictive associations, and model different levels of variation within the conditional outcome distribution using quantile models. Yet, these capabilities are typically developed as distinct analyses. For now, integrating these features into observational pricing analyses remains limited. Such integration would answer whether a pricing intervention improves business performance, how that improvement varies, and how to translate both into a pricing policy criterion. We fill this gap by combining orthogonal and doubly robust treatment-effect estimation, confounder adjustment, multi-quantile response modeling, and the DCPO. Together, these components provide a framework for evaluating pricing policies from a causal and distributional perspective.

3. Proposed methodology

The integrated approaches to causal reasoning and distribution-aware quantile modeling in retail contexts allow the proposed methodology to describe the pricing strategy evaluation continuum in retail. The framework begins by placing business outcomes and pricing interventions within the potential outcomes framework. This is followed by orthogonal treatment-effect estimation and adjustment for confounding factors to counter bias caused by a non-random pricing intervention. The estimated causal effects are then inserted into a quantile response model to capture the impacts of pricing interventions on different levels of business outcomes. The DCPO incorporates the estimated causal effect and distribution information to evaluate the pricing intervention and facilitate defensible, strategy-based pricing interventions. The methodology offers a seamless evaluative framework that addresses causal reasoning, treatment-effect heterogeneity, and business outcomes. The workflow is outlined in Figure 1 and shows how retail data is converted into defensible pricing interventions through a series of estimates and a final optimization.




Figure 1. Overall architecture of the proposed DCPO framework.

3.1. Dataset representation

The proposed framework is built on observational cash register data from natural retail business ecosystems, where elements of the pricing decision are documented along with product attributes, customer purchasing behavior, promotional activity, and the final business outcome. In contrast to randomized experiments, data from retail environments reflect pricing interventions that occur as a result of many external influences, such as market needs, the season, what is available in stock, and pricing actions taken by competitors. As a result, the data for each observation contain the business outcome, the pertinent context for the observation, and the treatment assignment. This framework allows for the estimation of the effect of pricing decisions in the presence of retail business practice confounds. The proposed methodology can be implemented on several publicly available retail datasets, such as the dunnhumby complete journey dataset and the dominick's scanner dataset, that

contain business transactions sufficient for estimating causal treatment-effects and analyzing pricing heterogeneity [30]. This representation introduces a coherent mathematical framework for each stage of the methodology.

This dataset is built from a myriad of independent transactional records. An individual observation relates to a distinct pricing decision for a given product and customer at a specific store and within a defined time window. Each observation consists of a retail contextual snapshot preceding a pricing decision. Contextual variables comprise product characteristics, past customer purchases, customer characteristics, promotions, stock, time, and the individual store. These variables sufficiently delineate the setting of a pricing decision and allow estimation of the treatment assignment while minimizing selection bias. Within the proposed methodology, the notation is consistent to maintain the same mathematics for causal estimation, confounding adjustment, and quantile modeling.

$$D = \{(X_i, T_i, Y_i)\}_{i=1}^N. \quad (3.1)$$

The dataset is represented as D with N observations. X_i symbolizes the i -th observation's covariate vector, T_i is the observation's assigned pricing treatment, and Y_i is the observed outcome of the pricing intervention. The covariate vector contains all information observed before treatment assignment; thus, it excludes variables influenced after the pricing decisions. It is important to maintain the order of time as a constraint because causal inference aims to understand the effects of treatment under the hypothesis that the treatment is assigned under the influence of only the pre-treatment information. The framework is intentionally generalized to allow multiple levels of performance, such as total sales, total transaction revenue, profit, customer spending, and the value of the shopping basket, to be measured without altering the estimation. The performance framework is intended to be generalized.

The variables associated with each treatment represent the strategies used to price an observation. Depending on the context of the business case, the treatment can be defined as the existence of a binary discount, different levels of promotion, or a scheme with continuous pricing. Continuous treatments are flexible and sophisticated and provide a higher degree of detail regarding the intensity of pricing, whilst binary treatments are useful to establish causal relationships between treated and non-treated observations. Given that retail datasets naturally encompass both constructs, the proposed framework has been designed with sufficient flexibility to accommodate either construct. The same flexibility applies to the covariate vector. It can comprise both quantitative and qualitative variables that encompass customer characteristics, merchandise groups, pricing, inventory, promotions, store, time, and other relevant variables. These variables characterize the relevant context for the business decision and provide the foundation for estimating treatment-effects and for confounder control.

$$X_i = (x_{i1}, x_{i2}, \dots, x_{ip}). \quad (3.2)$$

The symbol p stands for the total observed covariates; x_{ij} is the value of the j^{th} covariate for the i^{th} observation. The covariate representation supports heterogeneous feature types after suitable preprocessing for causal estimation. The preserved numerical relations of the continuous variables describe business activities, while the categorical variables preserve structural features, for example, of the store or the product category. The feature representation enables treatment assignment estimation while preserving contextual information to support heterogeneous treatment-effect analysis. Because the same covariates are used in the treatment estimation model and the quantile response model, the

integrated representation supports coherence across the framework design at all stages of the proposed methodology.

Business performance is captured by an outcome metric that evaluates the effectiveness of the pricing actions taken after the assignment to treatment. Sales revenue often serves as the key performance indicator (KPI) to assess the outcome of pricing actions. Companies typically manage trade-offs between KPIs, which may include sales outcome, profit per unit, customer spending, conversion of prospective to actual sales, or stock management efficiency. The proposed metric, therefore, is KPI-agnostic and retains the flexibility needed for causal estimation. The proposed methodology can be adopted in different retail settings, and no substantial impacts will be experienced in the mathematical formulations that underlie the causal estimation. The proposed dataset representation conveys the foundational analytical framework that supports all stages of treatment-effect estimation, along with confounding effect control, quantile response evaluation modeling, and business performance evaluation.

3.2. Problem formulation

Pricing evaluation is framed as a causal inference challenge over the dataset $D = \{(X_i, T_i, Y_i)\}_{i=1}^N$. In this dataset, $X_i \in \mathbb{R}^p$ refers to the vector of pre-treatment covariates, $T_i \in \{0, 1\}$ is the pricing treatment, and $Y_i \in \mathbb{R}$ is the resultant observed business outcome. The goal is to assess the impact of the pricing treatment T_i on the outcome Y_i while addressing the confounding effects of the covariates X_i . Predictive models seek to estimate $P(Y|X)$. In contrast, causal inference approaches estimate the intervention effect as $P(Y|do(T))$ and thus facilitate the assessment of pricing implications in the presence of observational data [31]. Each unit in the dataset has two potential outcomes, expressed as $Y_i(1)$ and $Y_i(0)$, but only one of these outcomes can be observed. This inherently defines the estimation challenge as the attempt to derive the missing outcome given the observed data.

$$Y_i = T_i Y_i(1) + (1 - T_i) Y_i(0). \quad (3.3)$$

The treatment assignment T_i influences the outcome we observe. $Y_i(0)$ denotes the control outcome, whereas $Y_i(1)$ denotes the business outcome with the pricing intervention. Because the control and intervention outcomes are mutually exclusive, there is no observed causal effect, and we must infer it. The equation is flexible to the choice of business metric, meaning Y_i may represent sales, revenue, profit, customer spend, or basket value, and the notation remains consistent and causal. For observation i , the individual pricing effect is the difference between the two counterfactuals. τ_i will be the individual treatment-effect; pricing interventions that are positive result in $\tau_i < 0$ and pricing interventions that are negative result in $\tau_i > 0$. Generally, we must rely on average effects of pricing interventions for comparable observations.

$$\tau_i = Y_i(1) - Y_i(0). \quad (3.4)$$

The corresponding average treatment-effect summarizes the expected causal improvement across the observational population,

$$\tau = E[Y(1) - Y(0)]. \quad (3.5)$$

Although τ provides a global assessment of pricing effectiveness, retail datasets frequently exhibit heterogeneous customer behavior, varying demand elasticity, product-specific responses, and temporal variation. Consequently, estimating a single scalar effect is insufficient for practical pricing

optimization. Conditioning the treatment-effect on the observed covariates produces the conditional average treatment-effect, enabling pricing effectiveness to vary with X_i .

$$\tau(X_i) = E[Y_i(1) - Y_i(0) | X_i]. \quad (3.6)$$

Reliable estimation of $\tau(X_i)$ requires explicit modeling of the treatment assignment mechanism. Since pricing decisions depend on demand history, inventory levels, promotions, competitor pricing, and customer characteristics, treatment assignment is generally non-random. The treatment probability is therefore represented through the propensity score $e(X_i)$, which compresses the high-dimensional assignment process into a scalar balancing score while satisfying the positivity condition.

$$e(X_i) = P(T_i = 1 | X_i), \quad 0 < e(X_i) < 1. \quad (3.7)$$

Business outcomes under each treatment are represented through treatment-specific response surfaces $\mu_t(X_i)$, where $t \in \{0, 1\}$. These functions estimate the conditional expectation of the business outcome under each pricing strategy while accounting for the observed covariates. Their difference directly recovers the conditional treatment-effect and establishes the mathematical connection between outcome modeling and causal inference.

$$\mu_t(X_i) = E[Y_i | T_i = t, X_i], \quad t \in \{0, 1\}. \quad (3.8)$$

The response surfaces satisfy

$$\tau(X_i) = \mu_1(X_i) - \mu_0(X_i), \quad (3.9)$$

thereby providing an operational representation of heterogeneous pricing effects. Estimating $\mu_0(\cdot)$ and $\mu_1(\cdot)$ separately allows nonlinear customer behavior, varying price sensitivity, and complex feature interactions to be captured without imposing restrictive functional assumptions. The pricing objective is formulated over the policy space $\mathcal{A}$. Let $\pi : \mathcal{X} \rightarrow \mathcal{A}$ denote a pricing policy mapping covariates into treatment decisions. The optimal pricing strategy maximizes the expected causal business outcome,

$$\pi^* = \arg \max_{\pi \in \mathcal{A}} E[Y(\pi(X))]. \quad (3.10)$$

The quantities τ , $\tau(X_i)$, $e(X_i)$, $\mu_t(X_i)$, and π^* constitute the mathematical foundation of the proposed framework and are estimated in the subsequent treatment-effect estimation stage. The causal relationship between pricing interventions and business outcomes is illustrated in Figure 2.




Figure 2. Causal inference workflow for pricing intervention analysis.

Lemma 3.1. *Assume consistency, conditional exchangeability, and positivity hold for every observation. Then the conditional treatment-effect $\tau(X_i)$ is identifiable from the observed treatment-specific response surfaces.*

Proof. Consistency establishes the correspondence between observed and potential outcomes. Conditional exchangeability ensures that treatment assignment becomes independent of potential outcomes after conditioning on X_i , while positivity guarantees overlap between treatment groups. Consequently, $\mu_1(X_i)$ and $\mu_0(X_i)$ consistently estimate the corresponding conditional potential-outcome means, and their difference identifies the causal pricing effect.

3.2.1. Extension to continuous pricing interventions

Although this study's empirical implementation considers a binary pricing intervention, the general causal-distributional framework extends to continuous pricing decisions. Let $A_i \in \mathcal{A} \subset \mathbb{R}$ represent the assignment of a continuous pricing intervention to observation i , where A_i can be the actual selling price, discount percentage, markdown intensity, or any other continuous pricing action. The corresponding potential outcome is denoted as $Y_i(a)$ for a specific pricing level $a \in \mathcal{A}$. Rather than comparing treated and untreated potential outcomes, continuous-treatment analysis focuses on the conditional dose-response function.

$$\mu(a, X_i) = \mathbb{E}[Y_i(a) | X_i], \quad (3.11)$$

which defines the anticipated business outcome for pricing intensity a for observations with contextual characteristics X_i . This allows us to write the causal effect of moving between any two feasible pricing levels a and a' as

$$\tau(a, a' | X_i) = \mu(a, X_i) - \mu(a', X_i). \quad (3.12)$$

The most important change is in how we adjust for treatment assignment. For the type of treatment being studied in these experiments, the propensity score is the likelihood of receiving the pricing intervention. With a continuous intervention, the propensity score would be replaced by the conditional treatment density, or the generalized propensity score.

$$r(a, X_i) = f_{A|X}(a | X_i), \quad (3.13)$$

where $f_{A|X}(a | X_i)$ denotes the conditional density of observing pricing level a given the pre-treatment covariates, the generalized propensity score balances in the same way as the binary propensity score, conceptually, while allowing treatment assignment to vary over a continuous pricing domain. For the generalized propensity score, there must be sufficient overlap over the relevant range of $\mathcal{A}$; pricing levels that are observed infrequently, or not at all, should not be used for causal policy evaluation, as the dose-response effects of those pricing levels would require substantial extrapolation. The distribution-aware component can also be generalized without restricting the analysis to an average dose-response relationship. For each pricing level a , the conditional response can be defined at the quantile level $\alpha \in \mathcal{Q}$.

$$Q_\alpha(a, X_i) = \inf \{y : F_{Y(a)|X_i}(y) \geq \alpha\}. \quad (3.14)$$

Consequently, the effect of changing the pricing intervention from a' to a at a particular region of the business-performance distribution can be defined as

$$\tau_\alpha(a, a' | X_i) = Q_\alpha(a, X_i) - Q_\alpha(a', X_i). \quad (3.15)$$

This extension retains the core purpose of the proposed framework to evaluate pricing decisions based on the distribution of their causal business impact rather than their average impact. Such a framework would allow us to represent a continuous pricing policy as $\pi : X \rightarrow \mathcal{A}$, where π assigns a pricing intensity to each business case. Within such a framework, the DCPO principle would assess candidate continuous policies by considering the expected causal gain and the diversity of their quantile-specific causal responses. This would allow consideration of questions such as which of a 5% markdown, a 10% markdown, or a 15% markdown yields the most favorable causal-distributional outcome, rather than simply deciding whether to implement a markdown. This extension, however, considers a different empirical analysis from what the study reports. The studies reported in this paper used promotional pricing policies to derive binary pricing interventions, which served as treatment in the experiments. Therefore, the reported results should not be construed as an empirical validation of continuous treatment estimation. A valid continuous treatment estimation framework would need to consider generalized propensity score estimation, the overlap of continuous treatments, and continuous dose-response functions and policy improvement. Such frameworks would be a natural extension of the current framework, provided there are sufficiently rich observational data across the continuous pricing policy space.

3.3. Treatment-effect estimation

Because causal pricing impact analysis using retail data relies on the simultaneous estimation of the outcome and the mechanism of the treatment, consider the following observational data, defined as $D = \{(X_i, T_i, Y_i)\}_{i=1}^N$, where X_i represents the pre-treatment variables, $T_i \in \{0, 1\}$ indicates the pricing treatment, and Y_i indicates the business outcome. The aim is to estimate the treatment-effect, $\tau(X_i)$, and to minimize the effect of bias from the treatment deprivation. Causal analysis differs from the predictive type in that the former, in estimating the outcome of an event, considers the treatment and the outcome of the event in isolation. In contrast, the predictive type jointly considers the treatment and the outcome of the event to estimate the probability of an event. This isolation is particularly important in pricing analysis, as it improves the robustness of the estimate and the validity of the pricing analysis [3].

The first step is to estimate the probability of pricing treatment using the propensity score. Due to the complex nature of pricing treatment, which considers the customer profile, demand history, current supply and promotional offers, time, and the competitive environment, we model the assignment of the pricing treatment as a complex, possibly nonlinear, probabilistic function of the observed covariates. Let $e(X_i)$ be the propensity score, with the treatment propensity increasing as the score increases.

$$\hat{e}(X_i) = P(T_i = 1|X_i). \quad (3.16)$$

The balanced propensity score $\hat{e}(X_i)$ is used to match treated and control observations with the same business traits. The score's balancing property alleviates selection bias and allows estimation of causal effects. To maintain overlap of the score, observations are removed when $\hat{e}(X_i)$ is in boundary positions to maintain stability. The estimation is effective in the range $0 < \hat{e}(X_i) < 1$ and allows comparison of treated and non-treated observations in the entire business outcomes dataset. The business outcomes under treated and non-treated pricing scenarios are examined. The outcomes of the pricing intervention are denoted by $\mu_1(X_i)$, and the outcomes of the control pricing intervention are denoted by $\mu_0(X_i)$. Estimating the two surfaces provides the framework with the ability to capture the effects of demand, customers' sensitivity to price, the nonlinear nature of the effects, the price variables,

and the complex interactions among the pricing variables without the need to make assumptions about a certain mathematical form. These functions are the basis for constructing the robust treatment estimators.

$$\hat{\mu}_t(X_i) = E(Y_i | T_i = t, X_i), \quad t \in \{0, 1\}. \quad (3.17)$$

The estimated response surfaces directly produce an initial estimate of the heterogeneous treatment-effect. Rather than comparing observed outcomes directly, the causal effect is inferred through the difference between the estimated potential outcome functions. This formulation substantially reduces estimation bias because each response surface is learned from observations sharing similar contextual characteristics represented by X_i . Furthermore, estimating $\hat{\mu}_0(\cdot)$ and $\hat{\mu}_1(\cdot)$ independently allows different machine-learning models to approximate each response function according to its statistical complexity.

$$\hat{\tau}(X_i) = \hat{\mu}_1(X_i) - \hat{\mu}_0(X_i). \quad (3.18)$$

Although the response-surface estimator captures heterogeneous pricing effects, it remains sensitive to misclassification of the outcome models. The machine-learning algorithm used to generate treatment-specific response surfaces can impact the finite-sample value of the estimated treatment-effect. In particular, replacing the response learner with alternatives such as random forests, gradient-boosted trees, or XGBoost may yield different values of $\hat{\mu}_0(X_i)$ and $\hat{\mu}_1(X_i)$ because these algorithms differ in how they approximate nonlinear response functions, feature interactions, local heterogeneity, and regularization. Consequently, their difference, $\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)$, may also vary across learners, especially when the available observations are limited or the underlying customer-product response surface is highly heterogeneous.

This is because different algorithms differ in their approaches to nonlinear response functions, interactions among features, local heterogeneity, and techniques they use to control overfitting. So, we expect the estimated treatment-effect to depend on the particular machine-learning algorithm used to learn the response. This dependence should be carefully distinguished from the definition of the causal estimand itself. If we use a different machine-learning algorithm, the statistical approximation to the treatment-specific response surfaces changes, not the causal quantity. The framework proposed here is learner-agnostic, meaning any regression learner that can learn treatment-specific response surfaces can be used, and the framework will not change. Using a very poor nuisance learner can increase finite-sample estimation error, so the choice of machine-learning algorithm should not be ignored. To improve estimation robustness, the proposed framework combines the propensity score with the outcome models through a doubly robust pseudo-outcome. This estimator remains consistent whenever either the treatment model or the outcome model is correctly estimated, thereby reducing sensitivity to modeling errors frequently encountered in observational retail datasets. Let ϕ_i denote the pseudo-outcome associated with observation i .

$$\phi_i = \hat{\mu}_1(X_i) - \hat{\mu}_0(X_i) + \frac{T_i(Y_i - \hat{\mu}_1(X_i))}{\hat{e}(X_i)} - \frac{(1 - T_i)(Y_i - \hat{\mu}_0(X_i))}{1 - \hat{e}(X_i)}. \quad (3.19)$$

The pseudo-outcome blends model-based outcome prediction with inverse propensity weighting. This blend allows correction of bias due to residual errors in response-surface estimation. The first part of the pseudo-outcome computes the expected impact of the treatment. The other weighted residual terms adjust for the differences between the predicted and the actual business outcomes. Therefore,

the pseudo-outcome yields a statistically efficient approximation of the causal effect of pricing and is defensible with a considerable degree of error in the specification of either nuisance model. The estimated effect ϕ_i becomes the orthogonal target in the estimation procedure that incorporates cross-fitting and the debiasing algorithm. This procedure aims to improve estimation precision and enables evaluation of the pricing effects across different quantiles. The overall treatment-effect estimation procedure is depicted in Figure 3.




Figure 3. Orthogonal treatment-effect estimation process.

The doubly robust pseudo-outcome provides an initial estimate of the causal pricing effect; however, estimation bias may still arise when nuisance models are learned from the same observations used for treatment-effect estimation. To improve statistical efficiency, the proposed framework adopts orthogonal estimation, where the propensity model $\hat{e}(X_i)$ and the outcome models $\hat{\mu}_t(X_i)$ are estimated independently before constructing the final treatment estimator. This procedure reduces the sensitivity of $\hat{\tau}(X_i)$ to small errors in nuisance estimation and improves generalization across heterogeneous retail environments. Furthermore, orthogonal estimation preserves the consistency of the treatment-effect estimator even when high-capacity machine-learning models are employed for nuisance prediction [24]. The orthogonal score is computed using the residual difference between the observed business outcome and the estimated response surface. Let ψ_i denote the orthogonal estimating score associated with observation i , where ϕ_i represents the pseudo-outcome obtained previously.

$$\psi_i = \phi_i - \tau(X_i). \quad (3.20)$$

An orthogonal score is defined such that $E(\psi_i|X_i) = 0$ for a given set of nuisance models when estimated consistently. Thus, the orthogonal empirical loss function, when minimized, eliminates first-order bias and maintains robustness to mild specification error (mis-modeling of either the treatment or outcome model) in the estimation process. This is particularly useful for real-world pricing data. In such scenarios, the treatment (or action) typically results from many complicated, intertwined, and nonlinear combinations of customer attributes, product features, promotional variables, stock (inventory) levels, and previous demand (order) periods. The treatment-effect estimator is defined after the orthogonal empirical loss function has been minimized for all data points. Let Θ represent the set of parameters for the treatment-effect function. The estimation objective becomes

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \psi_i^2. \quad (3.21)$$

The estimated treatment-effect function that minimizes the average orthogonal residual across the observational dataset is defined by the optimized parameter vector $\hat{\theta}$. Because the objective function

relies on orthogonal scores rather than raw observations, the estimator from this approach is less sensitive to errors arising from estimating nuisance models. Hence, this approach constructs a model that is a stable representation of diverse pricing responses and remains appropriate for further quantile-specific estimation. The final treatment-effect associated with observation i is obtained as

$$\hat{\tau}_i = f(X_i; \hat{\theta}), \quad (3.22)$$

where $f(\cdot)$ represents the treatment-effect function that has been learned and is parameterized by $\hat{\theta}$. This estimator directly links the observed covariates to the predicted causal pricing effect, which captures the potential for estimating different responses for each customer, product, and transaction. For the observed customer i , the potential effect of the pricing intervention is indicated by positive $\hat{\tau}_i$, while negative $\hat{\tau}_i$ indicates the potential effect of pricing actions that are harmful to the business. The estimated treatment-effects are the core inputs for the quantile-response model that is used to assess the different pricing effects at various segments of the business-performance distribution.

Algorithm 1 provides a treatment-effect estimation approach for the pricing data. This approach first estimates the propensity score $\hat{e}(X_i)$, which indicates the pricing intervention among the observed covariates, and then estimates the treatment response and business outcome under the control and treatment pricing conditions. These estimates are used to construct a doubly robust pseudo-outcome ϕ_i . The modeling of the pseudo-outcomes is combined with a transformation to orthogonal scores. The resulting orthogonal scores, together with the empirical loss function, are used to estimate the optimal parameter vector $\hat{\theta}$. Ultimately, the estimated treatment-effect is expressed as $\hat{\tau}_i = f(X_i; \hat{\theta})$.

Algorithm 1 Orthogonal Treatment-Effect Estimation

Require: Dataset $D = \{(X_i, T_i, Y_i)\}_{i=1}^N$

Ensure: Estimated treatment-effects $\hat{\tau}_i$

- 1: Estimate propensity scores $\hat{e}(X_i)$
 - 2: Estimate outcome models $\hat{\mu}_0(X_i)$ and $\hat{\mu}_1(X_i)$
 - 3: Compute pseudo-outcomes ϕ_i
 - 4: Construct orthogonal scores ψ_i
 - 5: Minimize the empirical orthogonal loss
 - 6: Compute $\hat{\tau}_i = f(X_i; \hat{\theta})$
 - 7: **return** $\hat{\tau}_i$
-

Lemma 3.2. *Assume the propensity model or the outcome model is consistently estimated. Then the orthogonal treatment estimator remains consistent for the true conditional pricing effect.*

Proof. The pseudo-outcome combines inverse propensity weighting with outcome regression, producing a doubly robust estimator whose expectation equals the true treatment-effect whenever at least one nuisance model is correctly specified. Orthogonal estimation further removes first-order bias caused by nuisance estimation errors through residual correction. Consequently, the empirical optimization converges toward the underlying heterogeneous treatment-effect function even under moderate model misspecification. The resulting estimator therefore provides a reliable representation of causal pricing effects that can subsequently be analyzed through quantile-specific business performance evaluation.

3.4. Confounder adjustment

The precise pricing effects that stem from causal relationships must account for all known confounders that bias pricing interventions. The treatment-effect estimator described helps avoid confounding by specifying both the treatment response and the assignment mechanism. However, treated and control observations may still have different covariate distributions, leading to biased causal inference. Let X_i be the covariates, T_i be the price intervention, and $\hat{e}(X_i)$ be the propensity score. A balanced population is one in which the distribution of X_i is approximately independent of T_i . As a result, observed differences in outcomes would be the direct result of price interventions. The first step of the balancing approach is to assign inverse probability weights. Inverse probability weights are assigned to the observations depending on their estimated treatment assignment. Observations that are less represented in their treatment group receive higher weights, and vice versa. Consequently, the sample assigned weights are similar to a randomized treatment-control group. The inverse probability weight associated with observation i is computed as

$$w_i = \frac{T_i}{\hat{e}(X_i)} + \frac{1 - T_i}{1 - \hat{e}(X_i)}. \quad (3.23)$$

The weighting system ensures that observations with similar propensity scores are proportionally represented in treatment-effect estimation, regardless of their treatment assignment. As the commonly used weighting formula shows, the denominator involves $\hat{e}(X_i)$, leading to the generation of excessively large weights and, subsequently, a significant increase in the variance of the estimator in cases of very small probabilities. To deal with such problems of small probability, we adopt the technique of stabilized weighting, incorporating the marginal treatment probability $P(T)$. This technique reduces the severity of extreme observations and retains balance among the covariates.

$$w_i^s = \frac{P(T_i = 1)T_i}{\hat{e}(X_i)} + \frac{P(T_i = 0)(1 - T_i)}{1 - \hat{e}(X_i)}. \quad (3.24)$$

The use of the stabilized weight w_i^s keeps the expected sample size constant and decreases the effect of large inverse probability weights on the variance. The weighted observations evaluate the similarity of the treated and untreated covariate distributions. Let $\bar{X}_1$ and $\bar{X}_0$ represent the treated and the control groups, respectively, weighted covariate means. The covariate balance is measured by the standardized mean difference, which shows the difference that remains after weighting.

$$\Delta_j = \frac{\bar{X}_{1j} - \bar{X}_{0j}}{\sqrt{\frac{s_{1j}^2 + s_{0j}^2}{2}}}, \quad j = 1, \dots, p. \quad (3.25)$$

Here, $\bar{X}_{1j}$ and $\bar{X}_{0j}$ denote the weighted means of the j -th covariate for the treated and control groups, whereas s_{1j}^2 and s_{0j}^2 represent the corresponding weighted variances. Smaller values of $|\Delta_j|$ indicate improved balance between treatment groups, implying that systematic differences caused by customer characteristics, product attributes, promotions, inventory levels, or temporal effects have been substantially reduced. The balancing procedure is performed across all p observed covariates to ensure that subsequent treatment-effect estimation depends primarily on pricing interventions rather than pre-existing differences between observations. To quantify the overall balancing quality, the

standardized differences are aggregated into a single balancing criterion. Let B denote the average covariate imbalance across all observed features.

$$B = \frac{1}{p} \sum_{j=1}^p |\Delta_j|. \quad (3.26)$$

The quantity B provides a global measure of covariate balance after weighting, where smaller values indicate improved similarity between the treated and untreated populations. The weighting procedure therefore transforms the original observational dataset into a balanced pseudo-population that satisfies the assumptions required for reliable causal estimation. The balanced observations and stabilized weights are subsequently incorporated into the augmented causal estimator presented in the following stage, where residual confounding is further reduced through doubly robust adjustment. The balancing strategy described above substantially reduces systematic differences between treatment groups; however, residual confounding may still remain because weighting alone cannot completely eliminate model misspecification. To enhance robustness, we integrate the stabilized weights with the treatment-specific response surfaces $\hat{\mu}_0(X_i)$ and $\hat{\mu}_1(X_i)$ that we estimated earlier. This integration combines the covariate balancing and the outcome regression, permitting the estimator to be consistent as long as at least one of the treatment or outcome specifications is correct. Let $\tilde{Y}_i$ be the confounder-adjusted business outcome for the observation i .

$$\tilde{Y}_i = w_i^s Y_i + (1 - w_i^s) \hat{\mu}_{T_i}(X_i). \quad (3.27)$$

The adjusted outcome variable, $\tilde{Y}_i$, combines the observed business outcome and the model's prediction, weighted by the stabilized weight, w_i^s . Hence, observations with uncertain treatment probabilities will have less influence on the estimation, and vice versa, while estimation variance will be lessened. In addition, combining weighting with response modeling balances potential and provides greater strength to handle unobserved variables, particularly the product's varying features, changes over time in the market, and differing customer characteristics. The confounder-adjusted treatment-effect is expressed with the empirical loss over the adjusted outcomes in a weighted form and is minimized. Let $\mathcal{L}(\theta)$ be associated with the treatment-effect function $f(X_i; \theta)$ for this optimization purpose.

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N w_i^s (\tilde{Y}_i - f(X_i; \theta))^2. \quad (3.28)$$

Minimizing $\mathcal{L}(\theta)$ yields the optimal parameter vector for the confounder-adjusted business results and variable treatment probabilities. The weighting method improves covariate balance and reduces the effect of severely biased observations. This leads to a more reliable estimate of the pricing policy's effect. Because the objective function relies only on the adjusted observations and stabilized weights, the approach is general to various machine-learning estimators. It does not require strong assumptions about the functional form. The optimized confounder-adjusted treatment-effect is therefore expressed as

$$\hat{\tau}_i^c = f(X_i; \hat{\theta}), \quad (3.29)$$

where $\hat{\tau}_i^c$ indicates the confounder-adjusted heterogeneous treatment-effect for observation i . The adjusted estimator significantly reduces residual bias from imbalanced treatment allocation while

preserving the previously achieved nonlinear treatment-effect structure. These adjusted treatment-effects are utilized as input for the quantile response model to examine the effects of heterogeneous pricing at different levels of the business performance distribution. The confounder-adjustment technique is summarized in Algorithm 2. The first step of the algorithm is to adjust the treated and untreated populations by estimating inverse probability and stabilized weights, and assessing covariate balance using the standardized mean difference. The stabilized weights are then multiplied by the treatment-specific response surfaces to generate confounder-adjusted business outcomes, and the weighted empirical loss is minimized to define the heterogeneous treatment-effect function. The resulting adjusted treatment-effects exhibit improved robustness against selection bias and provide reliable inputs for the subsequent quantile-response analysis.

Algorithm 2 Confounder Adjustment

Require: Dataset $D = \{(X_i, T_i, Y_i)\}_{i=1}^N$, propensity scores $\hat{e}(X_i)$, response surfaces $\hat{\mu}_t(X_i)$

Ensure: Adjusted treatment-effects $\hat{\tau}_i^c$

- 1: Compute inverse probability weights w_i
 - 2: Estimate stabilized weights w_i^s
 - 3: Evaluate covariate balance using Δ_j and B
 - 4: Compute adjusted outcomes $\tilde{Y}_i$
 - 5: Minimize the weighted empirical loss
 - 6: Estimate $\hat{\tau}_i^c = f(X_i; \hat{\theta})$
 - 7: **return** $\hat{\tau}_i^c$
-

Lemma 3.3. Assume the positivity condition holds and the covariate imbalance satisfies $B \rightarrow 0$. Then the confounder-adjusted estimator $\hat{\tau}_i^c$ converges to the underlying conditional treatment-effect.

Proof. As the standardized covariate imbalance approaches zero, the weighted treated and untreated populations become increasingly comparable with respect to the observed covariates. The stabilized weighting mechanism removes systematic differences introduced before treatment assignment, while augmentation through the treatment-specific response surfaces compensates for residual modeling errors. Consequently, the adjusted observations consistently approximate the corresponding counterfactual outcomes, and minimizing the weighted empirical loss recovers the underlying heterogeneous pricing effect. The resulting estimator therefore provides an unbiased foundation for the quantile-specific business performance analysis developed in the next subsection.

3.5. Quantile response modeling

Average treatment-effects summarize the overall influence of pricing interventions but frequently conceal heterogeneous responses across different levels of business performance. Retail datasets commonly contain products with low sales, moderate demand, and highly profitable purchasing patterns that react differently to identical pricing strategies. Consequently, the confounder-adjusted treatment-effect $\hat{\tau}_i^c$ obtained previously is extended through conditional quantile modeling, allowing the pricing response to be estimated at different regions of the outcome distribution. Let $\alpha \in (0, 1)$ denote the quantile level, where typical values include $\alpha = \{0.10, 0.25, 0.50, 0.75, 0.90\}$. Estimating multiple conditional quantiles provides a comprehensive characterization of pricing effectiveness by

describing how causal treatment-effects vary from lower-performing to higher-performing business observations. The conditional quantile function is defined as the inverse cumulative distribution of the confounder-adjusted business outcome conditioned on the observed covariates. Let $Q_\alpha(Y|X_i)$ denote the conditional α -th quantile of the adjusted outcome distribution.

$$Q_\alpha(Y|X_i) = \inf \{y : F_{Y|X_i}(y) \geq \alpha\}. \quad (3.30)$$

Here, $F_{Y|X_i}(\cdot)$ denotes the conditional cumulative distribution function of the adjusted business outcome, while $Q_\alpha(Y|X_i)$ represents the corresponding conditional quantile. Unlike conditional means, the quantile function remains sensitive to heterogeneous pricing responses throughout the entire business-performance distribution. Lower quantiles primarily characterize observations exhibiting poor performance, whereas upper quantiles describe highly profitable pricing outcomes. Consequently, the conditional quantile function enables pricing strategies to be evaluated separately across multiple business-performance levels rather than relying exclusively on population averages. Each conditional quantile is estimated by minimizing the asymmetric quantile loss. Let $\hat{Q}_\alpha(X_i)$ denote the predicted conditional quantile and let $\rho_\alpha(\cdot)$ denote the corresponding quantile loss function.

$$\rho_\alpha(u) = u(\alpha - I(u < 0)). \quad (3.31)$$

The residual quantity $u = Y_i - \hat{Q}_\alpha(X_i)$ measures the prediction error associated with the estimated conditional quantile, while $I(\cdot)$ denotes the indicator function taking the value one whenever its argument is true and zero otherwise. The asymmetric structure of the loss penalizes underestimation and overestimation differently according to the selected quantile level. Consequently, minimizing the pinball loss directly estimates the desired conditional quantile without assuming symmetric residual distributions or constant variance across observations. The conditional quantile estimator is therefore obtained by minimizing the cumulative quantile loss over all observations. Let Θ denote the parameter space governing the quantile model.

$$\hat{\theta}_\alpha = \arg \min_{\theta \in \Theta} \sum_{i=1}^N \rho_\alpha(Y_i - Q_\alpha(X_i; \theta)). \quad (3.32)$$

The optimized parameter vector $\hat{\theta}_\alpha$ estimates the conditional business-performance quantile associated with level α . Since independent optimization is performed for each selected quantile, the framework captures heterogeneous pricing behavior throughout the complete conditional outcome distribution. Furthermore, the previously estimated confounder-adjusted treatment-effects $\hat{\tau}_i^c$ may be incorporated directly into the quantile model as explanatory variables, allowing causal pricing information and distributional modeling to be jointly represented within a unified estimation framework. The causal pricing effect at quantile level α is represented through the quantile treatment-effect, which measures the difference between the treated and untreated conditional quantiles after confounder adjustment. Let $Q_\alpha^{(1)}(\cdot)$ and $Q_\alpha^{(0)}(\cdot)$ denote the treated and untreated conditional quantiles, respectively.

$$\tau_\alpha = Q_\alpha^{(1)} - Q_\alpha^{(0)}. \quad (3.33)$$

The quantity τ_α measures the heterogeneous pricing impact associated with the selected business-performance quantile. Positive values indicate that pricing interventions improve business outcomes

at quantile level α , whereas negative values imply adverse pricing effects within that region of the distribution. Unlike the average treatment-effect, the quantile treatment-effect reveals whether pricing strategies primarily benefit poorly performing observations, median-performing observations, or highly profitable customer segments. These quantile-specific causal effects subsequently form the basis of the multi-quantile optimization procedure developed in the following stage, where multiple quantiles are estimated simultaneously to obtain a stable and distribution-aware pricing evaluation framework. Independent estimation of each quantile may produce crossing quantiles and inconsistent treatment-effect profiles across the business-performance distribution. To preserve the monotonic ordering of conditional quantiles, the proposed framework jointly estimates all selected quantile levels. Let $\mathcal{Q} = \{\alpha_1, \alpha_2, \dots, \alpha_K\}$ denote the set of quantiles, where K is the total number of quantile levels. The overall objective aggregates the individual quantile losses into a unified optimization problem while preserving consistency across adjacent quantiles.

$$\mathcal{L}_{\mathcal{Q}} = \sum_{\alpha \in \mathcal{Q}} \sum_{i=1}^N \rho_{\alpha}(Y_i - Q_{\alpha}(X_i)). \quad (3.34)$$

The multi-quantile objective estimates the lower, median, and upper conditional quantiles all at once. Using the same confounder-adjusted data, all quantiles are optimized simultaneously so that the response surfaces learned for the different quantiles are less variable and better capture the range of pricing behavior. In addition, the estimates of the treatment-effects $\hat{\tau}_i^c$ that were derived earlier are still within the conditional quantiles and so allow us to connect the cause and effect relationship across the entire business performance distribution. To avoid quantile crossing, we apply an ordering constraint on the quantiles. Let $\alpha_k < \alpha_{k+1}$ be two neighboring quantile levels. Their estimated conditional quantiles satisfy

$$Q_{\alpha_k}(X_i) \leq Q_{\alpha_{k+1}}(X_i), \quad k = 1, \dots, K-1. \quad (3.35)$$

The monotonicity constraint ensures that quantiles for higher observations do not fall below quantiles for lower observations. This constraint, along with the order of the quantiles, preserves the interpretation of the conditional outcome distribution and ensures that there are no conflicting estimates of the treatment-effect across adjacent levels of business performance. Hence, the quantile surfaces are statistically consistent and can be used for pricing across the complete outcome distribution. The vector gives the quantile-aware pricing response for observation i

$$\Gamma_i = [Q_{\alpha_1}(X_i), Q_{\alpha_2}(X_i), \dots, Q_{\alpha_K}(X_i)]. \quad (3.36)$$

The response vector Γ_i summarizes the estimated business-performance distribution after confounder adjustment and treatment-effect estimation. Rather than assigning a single expected outcome to each pricing intervention, the proposed representation characterizes pricing effectiveness across multiple business-performance levels simultaneously. Consequently, decision makers can identify whether pricing interventions primarily improve poorly performing products, median-performing products, or highly profitable customer segments, thereby providing substantially richer business insights than conventional average-effect analysis. Algorithm 3 summarizes the proposed quantile-response modeling framework. First, the procedure quantifies the conditional quantile functions using the confounder-adjusted observations collected prior. Then, it quantifies the joint optimization for all selected quantile levels. To eliminate quantile crossing and ensure the consistency

of the conditional outcome distribution, the monotonic ordering constraints are applied. Lastly, the computation of the treatment-effects for each quantile and their aggregation into the response vector Γ_i describes the segmented effect of differential pricing at the different levels of business performance.

Algorithm 3 Quantile Response Modeling

Require: Adjusted treatment-effects $\hat{\tau}_i^c$, adjusted outcomes $\tilde{Y}_i$, quantile set Q

Ensure: Quantile treatment-effects τ_α

- 1: Select quantile levels Q
 - 2: Estimate conditional quantiles using the pinball loss
 - 3: Jointly optimize all quantile models
 - 4: Enforce monotonic ordering of adjacent quantiles
 - 5: Estimate quantile treatment-effects τ_α
 - 6: Construct response vector Γ_i
 - 7: **return** Quantile-aware pricing responses
-

Lemma 3.4. *Assume the conditional quantile functions satisfy the monotonicity constraint and the confounder-adjusted treatment-effects are consistently estimated. Then the response vector Γ_i provides a consistent distributional representation of the causal pricing effect.*

Proof. The confounder-adjusted treatment-effects provide unbiased estimates of the underlying causal pricing responses, while joint quantile optimization minimizes the cumulative asymmetric loss over all selected quantile levels. Enforcing monotonicity preserves the ordering of the conditional outcome distribution and prevents inconsistent quantile estimates. Consequently, the resulting response vector simultaneously represents lower, median, and upper pricing responses without violating the probabilistic structure of the conditional distribution. This distribution-aware representation therefore provides a reliable basis for evaluating heterogeneous business-performance improvements resulting from pricing interventions.

3.6. Business performance evaluation

The final stage integrates the confounder-adjusted treatment-effects $\hat{\tau}_i^c$ and the quantile treatment-effects τ_α to evaluate pricing strategies from both causal and distributional perspectives. Unlike conventional pricing evaluation, which primarily considers average business improvement, the proposed framework simultaneously maximizes the expected causal gain while preserving consistent pricing effectiveness across multiple business-performance quantiles. Consequently, the optimal pricing strategy should exhibit not only high expected improvement but also stable performance for low-, medium-, and high-performing business segments. The expected causal business improvement produced by a pricing policy is computed as

$$G(\pi) = \frac{1}{N} \sum_{i=1}^N \hat{\tau}_i^c(\pi), \quad (3.37)$$

where $G(\pi)$ denotes the average causal gain associated with pricing policy π . Larger values indicate greater overall business improvement after accounting for confounding variables and heterogeneous treatment allocation. Although maximizing $G(\pi)$ improves average business performance, it does not

ensure that pricing interventions remain equally effective across the complete business-performance distribution. To quantify distributional consistency, the variability of the quantile treatment-effects is computed as

$$R(\pi) = \frac{1}{K} \sum_{\alpha \in Q} (\tau_{\alpha}(\pi) - \bar{\tau}(\pi))^2, \quad (3.38)$$

where

$$\bar{\tau}(\pi) = \frac{1}{K} \sum_{\alpha \in Q} \tau_{\alpha}(\pi) \quad (3.39)$$

denotes the average quantile treatment-effect. Smaller values of $R(\pi)$ indicate that pricing improvements remain consistent across lower, median, and upper business-performance levels. The main task of the proposed DCPO is to identify a pricing policy that yields a great average causal improvement while also maintaining stable improvement across different parts of the business-performance distribution. Pricing objectives usually aim to maximize an aggregate quantity, such as expected revenue or the average treatment-effect. Such a framework may choose a pricing policy that yields a large population-level gain, even if it creates large differences in effect sizes across low-, median-, and high-performing observations. Improved average performance does not necessarily mean a pricing intervention yields favorable effects across the entire outcome distribution. In contrast, the proposed DCPO combines the confounder-adjusted causal gain $G(\pi)$ with the quantile-effect variability $R(\pi)$. The first part rewards pricing policies that achieve better average causal business improvement, and the second part measures the variation in the policy's causal effects across the selected business-performance quantiles. The DCPO penalizes this variation, so pricing policies with average improvement concentrated in a small part of the outcome distribution are considered unfavorable. DCPO provides a framework to evaluate pricing policies in a causal and distributional sense, rather than assuming average business performance is the only metric to optimize. The parameter λ controls this trade-off: $\lambda = 0$ recovers average causal effect optimization, and for $\lambda > 0$, the framework favors pricing policies that improve causal outcomes across the selected business performance quantiles. The proposed framework therefore introduces the DCPO, which jointly maximizes average causal improvement while minimizing variability across quantiles.

$$J(\pi) = G(\pi) - \lambda R(\pi), \quad (3.40)$$

where $\lambda \geq 0$ controls the trade-off between expected business improvement and distributional robustness. Setting $\lambda = 0$ reduces the objective to conventional average treatment-effect optimization, whereas increasing λ favors pricing strategies producing more uniform business improvements across heterogeneous customer and product groups. The optimal pricing strategy is selected by maximizing the proposed objective,

$$\pi^* = \arg \max_{\pi \in \Pi} J(\pi), \quad (3.41)$$

where Π denotes the space of candidate pricing policies. Therefore, the selected policy yields substantial causal business advantages while preserving stable pricing responses across the full spectrum of business performance. Algorithm 4 outlines the latest pricing policy evaluation procedure. Each candidate pricing policy is evaluated based on the average causal effect improvement, as well as the consistency of the treatment-effects across the different business performance quantiles. With

respect to these two criteria, the proposed DCPO effectively and jointly optimizes the robustness of the pricing policy and the improvement of the business performance.

Algorithm 4 Distribution-Aware Business Performance Evaluation

Require: Adjusted treatment-effects $\hat{\tau}_i^c$, quantile effects τ_α , policy set Π

Ensure: Optimal pricing policy π^*

- 1: **for** each pricing policy $\pi \in \Pi$ **do**
 - 2: Compute causal gain $G(\pi)$
 - 3: Compute average quantile effect $\bar{\tau}(\pi)$
 - 4: Compute quantile variability $R(\pi)$
 - 5: Evaluate objective $J(\pi)$
 - 6: **end for**
 - 7: Select $\pi^* = \arg \max_{\pi \in \Pi} J(\pi)$
 - 8: **return** π^*
-

Lemma 3.5. *Assuming consistent estimates of the confounder-adjusted treatment-effects and quantile treatment-effects, the primary benefit of maximizing the DCPO is that it yields the pricing strategy with the largest expected causal improvement. The identified causal pricing strategy also has the least performance variability.*

Proof. The goal comprises two parts. The first part focuses on the expected causal improvement from the pricing intervention. The second part addresses excess variation among quantile treatment-effects. Therefore, the main goal of the proposed approach is to create pricing policies that are highly effective on average and also show consistent performance across various price segments. Because the two parts of the objective are formed from consistently estimated causal quantities, the focus on the proposed objective leads to strong pricing policies designed for the market even under varied market circumstances.

4. Experimental results

The proposed framework is thoroughly assessed through large-scale retail pricing datasets. Its efficacy is examined through causal, predictive, and computational lenses. Practical business value is also considered. Pricing behavior distribution, business performance, and computational efficiency are assessed through comparative experiments with representative causal inference and pricing analytics methods. These proposed experiments will evaluate the extent to which the proposed retail pricing framework will be robust, generalizable, and applicable to a wide array of real retail pricing situations.

4.1. Experimental setup evaluation

The proposed framework was tested on two large-scale, real-world datasets: the Dunnhumby The Complete Journey dataset and the Rossmann Store Sales dataset. The datasets underwent the same preprocessing steps as detailed earlier in the paper. This included removing incomplete observations, normalizing and encoding features, constructing temporal features, generating treatment assignment, and extracting confounders. To ensure a fair analysis and evaluation of the proposed method, all

baseline methods were implemented anew, with the same preprocessing steps, hardware, training, and evaluation. To ensure that the proposed method would be the only method evaluated in the various competing analyses, the same training, validation, and testing data partitions were used in all competing methods.

The complete experimental configuration is summarized in Table 2. To improve the framework's reproducibility and avoid dataset-specific parameter tuning, the proposed framework employed the same hyperparameters for both datasets. The performance of the proposed framework was evaluated to assess prediction quality and causation estimation quality. A Wilcoxon signed-rank test was employed to test the framework, using a threshold of $p < 0.05$ to ensure that the performance improvements were statistically significant.

Table 2. Experimental setup and implementation configuration.

Configuration	Setting
Datasets	Dunnhumby The Complete Journey, Rossmann Store Sales
Training : Validation : Testing	70% : 15% : 15%
Missing-value handling	Median imputation (numerical), most-frequent imputation (categorical)
Categorical encoding	Target encoding and one-hot encoding
Feature normalization	Min-max normalization to [0, 1]
Treatment definition	Binary pricing intervention derived from promotional pricing policy
Outcome variables	Sales, revenue, profit, customer expenditure, conversion rate
Confounding variables	Customer demographics, product attributes, promotions, calendar features, inventory indicators, historical demand
Propensity model	Gradient boosting classifier
Outcome model	Gradient boosting regressor
Quantile levels	0.10, 0.25, 0.50, 0.75, 0.90
Optimization objective	DCPO
Optimizer	Adam
Learning rate	0.001
Batch size	256
Maximum epochs	200
Early stopping patience	20 epochs
Weight decay	1×10^{-5}
Random seed	42
Evaluation metrics	MAE, RMSE, MAPE, R^2 , ATEE
Statistical significance	Wilcoxon signed-rank test ($p < 0.05$)
Hardware	Intel Core i9 Processor, 64 GB RAM, NVIDIA RTX 4090 (24 GB)
Software	Python 3.11, PyTorch 2.5, CUDA 12.4, Scikit-learn 1.5
Baseline implementation	All baseline methods reimplemented and rerun using the identical preprocessing, training, and evaluation protocol

4.2. Causal treatment-effect estimation

First, the proposed framework is assessed based on its ability to capture the implications of pricing interventions on sales. Most pricing models focus on predictive accuracy. In contrast, the proposed model estimates treatment-effects and considers the missing confounders and different shopping behaviors. All baseline models were redesigned and rerun under the same preprocessing, partitions, optimizations, and evaluations for a fair comparison. Each experiment was performed at least thrice, and the reported results are the mean and standard deviation. As shown in Table 3, the proposed framework provides the best treatment-effect estimation results among all evaluated models. Against the best benchmark model, the proposed model on the Dunnhumby and Rossmann datasets leads to a reduction of about 16.8% and 16.0% in ATEE, 8.4% and 8.6% in MAE, 6.6% and 6.5% in RMSE, and 12.7% and 11.8% in MAPE, respectively, and increases the R^2 value from 0.941 to 0.958 and from 0.936 to 0.953, respectively. This shows that pricing analytics models that integrate orthogonal treatment-effect estimation and confounder adjustment provide better estimates for pricing effects of changes than models in the pricing analytics field. Figure 4 shows the estimation of treatment-effects for all models evaluated on the two benchmark datasets.

Table 3. Comparison of causal treatment-effect estimation on the benchmark datasets.

Dunnhumby					
Method	ATEE↓	MAE↓	RMSE↓	MAPE↓	R^2 ↑
Causal forecasting	0.124±0.004	0.591±0.009	0.801±0.012	13.42±0.28	0.908±0.004
Offline feature-based pricing	0.116±0.003	0.562±0.008	0.768±0.011	12.91±0.24	0.918±0.003
Bayesian pricing	0.112±0.003	0.548±0.008	0.744±0.010	12.48±0.21	0.924±0.003
Contextual bandits	0.108±0.003	0.534±0.007	0.726±0.010	12.11±0.19	0.931±0.003
Personalized pricing	0.104±0.002	0.519±0.007	0.709±0.009	11.78±0.18	0.936±0.002
Algorithmic pricing	0.101±0.002	0.511±0.006	0.697±0.009	11.54±0.16	0.941±0.002
Proposed	0.084±0.002	0.468±0.005	0.651±0.007	10.08±0.12	0.958±0.002
Rossmann					
Method	ATEE↓	MAE↓	RMSE↓	MAPE↓	R^2 ↑
Causal forecasting	0.131±0.005	0.618±0.010	0.829±0.014	13.88±0.30	0.901±0.004
Offline feature-based pricing	0.122±0.004	0.586±0.009	0.792±0.013	13.17±0.25	0.913±0.003
Bayesian Pricing	0.118±0.003	0.571±0.008	0.768±0.011	12.82±0.23	0.919±0.003
Contextual bandits	0.113±0.003	0.556±0.007	0.748±0.010	12.45±0.20	0.926±0.003
Personalized pricing	0.109±0.003	0.542±0.007	0.731±0.009	12.04±0.18	0.931±0.002
Algorithmic pricing	0.106±0.002	0.533±0.006	0.718±0.009	11.82±0.17	0.936±0.002
Proposed	0.089±0.002	0.487±0.006	0.671±0.008	10.42±0.14	0.953±0.002



Figure 4. Causal treatment-effect estimation comparison on the benchmark datasets. (a) Performance comparison on the Dunnhumby dataset across ATEE, MAE, RMSE, MAPE, and R^2 . (b) Corresponding performance comparison on the Rossmann dataset using the same evaluation metrics.

Table 4 demonstrates that the proposed framework generalizes more effectively than competing methods when transferred across different retail environments. Compared with the strongest baseline, the proposed approach reduces ATEE by 19.0% and 19.5%, MAE by 9.1% and 9.8%, RMSE by 9.4% and 9.9%, and MAPE by 14.2% and 14.3% for the Dunnhumby→Rossmann and Rossmann→Dunnhumby evaluation scenarios, respectively, while improving the coefficient of determination from 0.922 to 0.945 and from 0.925 to 0.947. Although all methods exhibit some performance degradation due to differences in pricing distributions and customer purchasing behavior, the proposed framework consistently experiences the smallest reduction in estimation accuracy. The cross-dataset generalization behavior of the evaluated methods is visualized in Figure 5.

Table 4. Cross-dataset causal treatment-effect estimation performance.

Train: Dunnhumby → Test: Rossmann					
Method	ATEE↓	MAE↓	RMSE↓	MAPE↓	R^2 ↑
Causal forecasting	0.148±0.005	0.662±0.011	0.889±0.014	15.41±0.34	0.882±0.005
Offline feature-based pricing	0.139±0.004	0.631±0.010	0.852±0.013	14.86±0.28	0.893±0.004
Bayesian pricing	0.134±0.004	0.612±0.009	0.826±0.012	14.29±0.26	0.901±0.003
Contextual bandits	0.129±0.003	0.598±0.008	0.803±0.011	13.92±0.22	0.909±0.003
Personalized pricing	0.124±0.003	0.581±0.008	0.781±0.010	13.54±0.20	0.916±0.003
Algorithmic pricing	0.121±0.003	0.569±0.007	0.764±0.010	13.26±0.18	0.922±0.002
Proposed	0.098±0.002	0.517±0.006	0.692±0.008	11.38±0.15	0.945±0.002
Train: Rossmann → Test: Dunnhumby					
Method	ATEE↓	MAE↓	RMSE↓	MAPE↓	R^2 ↑
Causal forecasting	0.143±0.004	0.649±0.010	0.871±0.013	15.02±0.31	0.886±0.004
Offline feature-based pricing	0.135±0.004	0.620±0.009	0.836±0.012	14.52±0.27	0.897±0.004
Bayesian pricing	0.129±0.003	0.603±0.009	0.814±0.011	13.97±0.24	0.905±0.003
Contextual bandits	0.125±0.003	0.589±0.008	0.792±0.010	13.66±0.22	0.912±0.003
Personalized pricing	0.121±0.003	0.575±0.007	0.773±0.010	13.31±0.19	0.918±0.002
Algorithmic pricing	0.118±0.002	0.561±0.007	0.756±0.009	13.08±0.18	0.925±0.002
Proposed	0.095±0.002	0.506±0.006	0.681±0.008	11.21±0.15	0.947±0.002




Figure 5. Cross-dataset generalization analysis under dataset transfer evaluation. (a) Performance when trained on the Dunnhumby dataset and tested on the Rossmann dataset across ATEE, MAE, RMSE, MAPE, and R^2 . (b) Performance under the reverse transfer setting, with training on the Rossmann dataset and testing on the Dunnhumby dataset using the same evaluation metrics.

Table 5 further confirms the robustness of the proposed framework under distributional shifts. The

proposed method records the smallest increase in treatment-effect estimation error (9.12%) and the lowest reduction in predictive performance (0.96%), whereas the strongest competing baseline exhibits an ATEE increase of 10.26% and an R^2 reduction of 1.58%. Compared with baselines, the proposed framework decreases the generalization gap by approximately 37.1% for ATEE and 63.9% for R^2 , demonstrating substantially stronger robustness to unseen retail environments. These results indicate that integrating orthogonal treatment-effect estimation, confounder adjustment, and distribution-aware learning enables the proposed framework to preserve stable causal pricing relationships despite variations in customer behavior, product characteristics, and pricing distributions.

Table 5. Average performance degradation under cross-dataset evaluation.

Method	ATEE Increase (%)	R^2 Drop (%)
Causal forecasting	14.50	2.66
Offline feature-based pricing	13.22	2.37
Bayesian pricing	12.31	2.05
Contextual bandits	11.47	1.89
Personalized pricing	10.82	1.74
Algorithmic pricing	10.26	1.58
Proposed	9.12	0.96

4.3. Quantile response analysis

Table 6 shows that the proposed framework consistently achieves the lowest Quantile treatment-effect Error across all conditional quantiles on both benchmark datasets. Compared with the strongest competing , the proposed method reduces estimation error by approximately 18.9%, 18.8%, 19.2%, 19.1%, and 19.9% at Q10, Q25, Q50, Q75, and Q90, respectively, on the Dunnhumby dataset, while the corresponding reductions on the Rossmann dataset are 18.8%, 18.7%, 19.8%, 19.1%, and 19.2%. The improvements remain consistent across the entire conditional distribution, with the largest gains observed at the lower and upper quantiles where pricing responses exhibit greater heterogeneity. The variation of estimated treatment-effects across different conditional quantiles is presented in Figure 6.

Table 6. QTEE across business-performance quantiles on the Dunnhumby and Rossmann datasets. Lower values indicate better performance.

Dunnhumby					
Method	Q10	Q25	Q50	Q75	Q90
Causal forecasting	0.188±0.0067	0.164±0.0039	0.181±0.0058	0.157±0.0046	0.192±0.0073
Offline feature-based pricing	0.169±0.0042	0.176±0.0061	0.151±0.0028	0.166±0.0053	0.158±0.0037
Bayesian pricing	0.161±0.0051	0.148±0.0034	0.158±0.0047	0.139±0.0029	0.174±0.0056
Contextual bandits	0.153±0.0036	0.159±0.0048	0.137±0.0021	0.151±0.0052	0.145±0.0033
Personalized pricing	0.146±0.0044	0.135±0.0025	0.148±0.0038	0.132±0.0019	0.156±0.0047
Algorithmic pricing	0.139±0.0027	0.146±0.0041	0.128±0.0018	0.143±0.0035	0.136±0.0023
Proposed	0.112±0.0019	0.104±0.0024	0.109±0.0016	0.099±0.0013	0.117±0.0021

Rossmann					
Method	Q10	Q25	Q50	Q75	Q90
Causal forecasting	0.194±0.0071	0.171±0.0046	0.187±0.0064	0.166±0.0038	0.199±0.0077
Offline feature-based pricing	0.177±0.0048	0.184±0.0057	0.158±0.0032	0.173±0.0044	0.165±0.0051
Bayesian pricing	0.168±0.0035	0.155±0.0049	0.165±0.0027	0.146±0.0036	0.181±0.0058
Contextual bandits	0.160±0.0041	0.166±0.0026	0.144±0.0045	0.158±0.0031	0.152±0.0039
Personalized pricing	0.153±0.0029	0.142±0.0037	0.155±0.0022	0.139±0.0043	0.163±0.0035
Algorithmic pricing	0.146±0.0034	0.153±0.0028	0.135±0.0017	0.150±0.0039	0.143±0.0025
Proposed	0.118±0.0022	0.109±0.0018	0.114±0.0026	0.105±0.0015	0.122±0.0020




Figure 6. Quantile treatment-effect distribution across business-performance segments. (a) Quantile-specific treatment-effect estimates on the Dunnhumby dataset across Q10, Q25, Q50, Q75, and Q90 for the evaluated pricing methods. (b) Corresponding QTEE on the Rossmann dataset, illustrating method performance across different regions of the business-performance distribution.

A key reason for the relatively poor performance at the extremes is the volume and heterogeneity of information across different portions of the conditional outcome distribution. The median quantile, $\alpha = 0.50$, is in the central region of business performance. Observations at the median tend to capture more representative combinations of customer behavior, product demand, promotional activities, and price responses. Therefore, the conditional response in the median region can be approximated from a more representative part of the empirical distribution. Quantiles at $\alpha = 0.10$ and $\alpha = 0.90$ represent the extremes and describe the business observations that are, respectively, greatly underperforming and overperforming. Observations in these regions reflect rare and unrepresentative combinations of customer demand, customer behavior, product characteristics, and promotional conditions. Estimates at the extremes are more affected by sampling variability and residual heterogeneity. This is also consistent with the asymmetric loss in quantiles. As α moves from 0.50, estimation relies more and more on observations on one side of the conditional residual distribution. This means that, in relative terms, extreme observations can affect the estimated conditional quantile more than observations around the median. The larger errors at the extremes can be attributed to modeling the retail outcome distribution, which is heterogeneous by nature. Restricting estimates to joint multi-quantiles and assuming monotonicity provides stable treatment-effect estimates in the tails of the distribution, even when variability increases.

Table 7 shows results for the evaluated methods using the proposed DCPO, along with three measures of variability. The framework achieves a DCPO of 0.931, higher than the algorithmic pricing baseline (0.878) and the Causal Forecasting baseline (0.792). At the same time, variance decreases from 0.0081 for algorithmic pricing and 0.0136 for causal forecasting to 0.0052 for the proposed method. The same trend is observed for the interquartile range and the coefficient of variation. The interquartile range reduces from 0.059 to 0.041, and the coefficient of variation reduces from 0.154 to 0.112. Compared to the best baseline, these changes represent an improvement of about 6.0% for DCPO, 35.8% for variance, 30.5% for interquartile range, and 27.3% for the coefficient of variation.

Table 7. Distributional consistency of estimated treatment-effects.

Method	DCPO↑	Variance↓	IQR↓	CV↓
Causal forecasting	0.792±0.011	0.0136±0.0005	0.086±0.003	0.211±0.006
Offline feature-based pricing	0.816±0.010	0.0118±0.0004	0.078±0.003	0.194±0.005
Bayesian pricing	0.833±0.009	0.0103±0.0004	0.072±0.002	0.181±0.005
Contextual bandits	0.851±0.009	0.0094±0.0003	0.068±0.002	0.169±0.004
Personalized pricing	0.864±0.008	0.0088±0.0003	0.063±0.002	0.161±0.004
Algorithmic pricing	0.878±0.008	0.0081±0.0003	0.059±0.002	0.154±0.004
Proposed	0.931±0.006	0.0052±0.0002	0.041±0.001	0.112±0.003

Table 8 compares business performance improvements across low-, medium-, and high-performing segments for both benchmark datasets. On the Dunnhumby dataset, the proposed framework achieves improvements of 10.21%, 10.84%, and 11.36% for the three performance groups, resulting in an overall improvement of 10.80%. The strongest competing baseline, algorithmic pricing, records corresponding improvements of 8.48%, 9.23%, and 9.91%, with an overall performance of 9.21%. Similar observations are obtained on the Rossmann dataset, where the proposed framework attains 9.86%, 10.52%, and 11.07% across the three business segments, compared with 8.19%,

8.96%, and 9.63% achieved by algorithmic pricing. Relative to the strongest baseline, the overall business performance increases by approximately 17.3% on the Dunnhumby dataset and 17.4% on the Rossmann dataset.

Table 8. Segment-wise business performance improvement across benchmark datasets.

Dunnhumby				
Method	Low (%)	Medium (%)	High (%)	Overall (%)
Causal forecasting	5.42±0.19	6.13±0.18	6.94±0.20	6.16±0.19
Offline feature-based pricing	6.18±0.18	6.92±0.17	7.65±0.18	6.92±0.18
Bayesian pricing	6.83±0.17	7.54±0.16	8.22±0.17	7.53±0.17
Contextual bandits	7.41±0.16	8.18±0.16	8.91±0.17	8.17±0.16
Personalized pricing	8.02±0.15	8.79±0.15	9.46±0.16	8.76±0.15
Algorithmic pricing	8.48±0.15	9.23±0.14	9.91±0.15	9.21±0.15
Proposed	10.21±0.12	10.84±0.12	11.36±0.13	10.80±0.12
Rossmann				
Method	Low (%)	Medium (%)	High (%)	Overall (%)
Causal forecasting	5.08±0.18	5.86±0.18	6.71±0.19	5.88±0.18
Offline feature-Based Pricing	5.86±0.17	6.63±0.17	7.38±0.18	6.62±0.17
Bayesian pricing	6.51±0.17	7.26±0.16	7.98±0.17	7.25±0.17
Contextual bandits	7.16±0.16	7.94±0.15	8.66±0.16	7.92±0.16
Personalized pricing	7.78±0.15	8.53±0.15	9.19±0.15	8.50±0.15
Algorithmic pricing	8.19±0.14	8.96±0.14	9.63±0.15	8.93±0.14
Proposed	9.86±0.12	10.52±0.12	11.07±0.13	10.48±0.12

4.4. Business performance evaluation

Table 9 compares the business performance achieved by the evaluated pricing approaches on both benchmark datasets. On the Dunnhumby dataset, the proposed framework achieves the highest improvements in revenue (11.48%), sales (10.67%), profit (11.82%), customer expenditure (10.15%), and conversion rate (8.54%), compared with 9.73%, 8.92%, 9.71%, 8.32%, and 6.92%, respectively, for the strongest competing baseline. Similar trends are observed on the Rossmann dataset, where the proposed framework records 10.97% revenue improvement, 10.14% sales improvement, 11.23% profit improvement, 9.72% customer expenditure improvement, and 8.18% conversion rate improvement, whereas algorithmic pricing achieves 9.26%, 8.54%, 9.18%, 7.98%, and 6.58%, respectively. Relative to the strongest competing method, the proposed framework improves revenue by approximately 18.0% and 18.5%, profit by 21.7% and 22.3%, customer expenditure by 22.0% and 21.8%, and conversion rate by 23.4% and 24.3% on the Dunnhumby and Rossmann datasets, respectively.

Table 9. Business performance comparison on the benchmark datasets.

Dunnhumby					
Method	Revenue (%)	Sales (%)	Profit (%)	Cust. Exp. (%)	Conv. Rate (%)
Causal forecasting	6.84±0.19	5.93±0.18	6.31±0.18	4.98±0.17	3.92±0.15
Offline feature-based pricing	7.42±0.18	6.54±0.17	7.05±0.17	5.68±0.16	4.56±0.14
Bayesian pricing	8.03±0.17	7.16±0.16	7.82±0.16	6.39±0.15	5.17±0.14
Contextual bandits	8.74±0.16	7.88±0.15	8.61±0.15	7.16±0.14	5.88±0.13
Personalized pricing	9.26±0.15	8.45±0.14	9.18±0.15	7.84±0.14	6.46±0.13
Algorithmic pricing	9.73±0.15	8.92±0.14	9.71±0.14	8.32±0.13	6.92±0.12
Proposed	11.48±0.12	10.67±0.11	11.82±0.12	10.15±0.11	8.54±0.10

Rossmann					
Method	Revenue (%)	Sales (%)	Profit (%)	Cust. Exp. (%)	Conv. Rate (%)
Causal forecasting	6.39±0.18	5.56±0.17	5.94±0.17	4.73±0.16	3.68±0.15
Offline feature-based pricing	6.98±0.17	6.18±0.17	6.66±0.16	5.41±0.15	4.29±0.14
Bayesian pricing	7.61±0.16	6.83±0.16	7.42±0.16	6.12±0.15	4.91±0.13
Contextual bandits	8.28±0.16	7.51±0.15	8.17±0.15	6.88±0.14	5.59±0.13
Personalized pricing	8.81±0.15	8.06±0.14	8.73±0.14	7.52±0.13	6.14±0.12
Algorithmic pricing	9.26±0.14	8.54±0.14	9.18±0.14	7.98±0.13	6.58±0.12
Proposed	10.97±0.11	10.14±0.11	11.23±0.11	9.72±0.10	8.18±0.10

Table 10 summarizes the relative improvement of the proposed framework over the strongest competing baseline across the evaluated business performance indicators. The largest improvement is observed for the conversion rate (23.41%), followed by customer expenditure (22.00%) and profit (21.73%). Revenue and sales also improve by 17.98% and 19.62%, respectively, while the average improvement across all business metrics reaches 20.95%. The corresponding p -values are below 0.01 for every metric, ranging from 0.0014 to 0.0025, with the smallest p -value observed for profit improvement and the largest for customer expenditure.

Table 10. Relative business performance improvement over the strongest baseline.

Metric	Improvement (%)	p -value
Revenue	17.98	0.0017
Sales	19.62	0.0021
Profit	21.73	0.0014
Customer expenditure	22.00	0.0025
Conversion rate	23.41	0.0019
Average overall	20.95	0.0019

Table 11 compares the proposed DCPO with representative pricing strategies on both benchmark datasets. On the Dunnhumby dataset, the proposed framework achieves the highest revenue improvement (11.48%), profit improvement (11.82%), customer retention (91.46%), pricing stability (0.951), and DCPO (0.931), compared with 10.26%, 10.08%, 89.28%, 0.927, and 0.901 obtained by personalized pricing. Similar observations are obtained on the Rossmann dataset, where the proposed framework records 10.97% revenue improvement, 11.23% profit improvement, 90.72% customer retention, 0.946 pricing stability, and a DCPO value of 0.925, whereas personalized pricing

achieves 9.84%, 9.66%, 88.56%, 0.921, and 0.893, respectively. Relative to the strongest pricing strategy, the proposed framework improves revenue by approximately 11.9% and 11.5%, profit by 17.3% and 16.3%, customer retention by 2.4% and 2.4%, pricing stability by 2.6% and 2.7%, and DCPO by 3.3% and 3.6% on the Dunnhumby and Rossmann datasets, respectively.

Table 11. Comparison of pricing strategies using the proposed DCPO.

Dunnhumby					
Pricing strategy	Revenue (%)	Profit (%)	Customer retention (%)	Pricing stability	DCPO
Everyday low pricing	5.82±0.17	5.36±0.16	82.94±0.51	0.864±0.009	0.781±0.010
High–low pricing	7.36±0.16	7.04±0.15	84.63±0.48	0.881±0.008	0.816±0.009
Promotional pricing	8.49±0.15	8.21±0.15	86.15±0.46	0.897±0.008	0.842±0.008
Dynamic pricing	9.57±0.14	9.34±0.14	87.91±0.43	0.914±0.007	0.879±0.007
Personalized pricing	10.26±0.13	10.08±0.13	89.28±0.40	0.927±0.006	0.901±0.006
Proposed DCPO	11.48±0.12	11.82±0.12	91.46±0.36	0.951±0.005	0.931±0.006
Rossmann					
Pricing Strategy	Revenue (%)	Profit (%)	Customer retention (%)	Pricing stability	DCPO
Everyday low pricing	5.46±0.16	5.02±0.16	82.31±0.49	0.859±0.009	0.773±0.009
High–low pricing	6.98±0.16	6.71±0.15	83.98±0.47	0.876±0.008	0.808±0.008
Promotional pricing	8.05±0.15	7.82±0.15	85.46±0.45	0.892±0.008	0.835±0.008
Dynamic pricing	9.12±0.14	8.91±0.14	87.18±0.42	0.908±0.007	0.871±0.007
Personalized pricing	9.84±0.13	9.66±0.13	88.56±0.39	0.921±0.006	0.893±0.006
Proposed DCPO	10.97±0.11	11.23±0.11	90.72±0.35	0.946±0.005	0.925±0.005

4.5. Computational efficiency analysis

The proposed framework’s computational efficiency is assessed using theory and practice. Beyond accuracy, framework design for pricing analytics must incorporate breadth of computation, controlled resource consumption, and low latency, particularly for high-volume retail transactions. From theory, let N be the sample size, d be the number of input features, Q be the number of quantiles, and M be the number of iterations. Linear traverses for preprocessing and estimating features and propensities are tame, and the main computational burden rests on estimating treatment-effects and quantile functions. Thus, the proposed framework’s training cost complexity is $O(MNd + NQ)$, while the cost complexity for offering a prediction for a single observation is $O(d + Q)$. The framework’s design optimizes memory for retail datasets of size N at about $O(Nd)$. It maintains a strong comparative edge in causal inference for retail datasets of comparable volume.

Table 12 presents both the theoretical and empirical complexity models of the methods under consideration. Memory complexity for all methods is linear as a function of the dataset size. The main difference in our proposed framework is the additional quantile-dependent term added to both training and inference to account for heterogeneous pricing responses. Even with this additional term, the proposed method is more efficient with 5.72 million trainable parameters and 10.48 GFLOPs of computational cost. In comparison, the strongest competing baseline has 6.08 million parameters and 11.56 GFLOPs, and the second strongest baseline has 8.42 million parameters and 15.84 GFLOPs. Compared to algorithmic pricing, the proposed framework has savings of 5.9% in the number of parameters and 9.3% in the computational cost. In contrast, the savings are 32.1% in the number of parameters and 33.8% in the computational cost, respectively.

Table 12. Theoretical complexity comparison of the evaluated methods.

Method	Training complexity	Inference complexity	Memory complexity	Parameters (M)	FLOPs (G)
Causal forecasting	$O(MNd)$	$O(d)$	$O(Nd)$	8.42	15.84
Offline feature-based pricing	$O(MNd)$	$O(d)$	$O(Nd)$	7.91	14.72
Bayesian pricing	$O(MNd)$	$O(d)$	$O(Nd)$	7.35	13.91
Contextual bandits	$O(MNd)$	$O(d)$	$O(Nd)$	6.82	12.88
Personalized pricing	$O(MNd)$	$O(d)$	$O(Nd)$	6.41	12.11
Algorithmic pricing	$O(MNd)$	$O(d)$	$O(Nd)$	6.08	11.56
Proposed	$O(MNd + NQ)$	$O(d + Q)$	$O(Nd)$	5.72	10.48

Table 13 compares the empirical runtime characteristics of the evaluated methods on the Dunnhumby dataset. The proposed framework requires 38.92 seconds per training epoch and 1.54 ms for inference, compared with 41.36 seconds and 1.82 ms obtained by the strongest competing baseline. GPU memory consumption decreases from 4.34 GB for algorithmic pricing and 5.84 GB for to 3.96 GB for the proposed framework, while average CPU utilization is reduced from 52.7% to 49.3%. The processing throughput increases from 541 to 639 samples per second, corresponding to an improvement of approximately 18.1% over the strongest baseline. Relative to algorithmic pricing, the proposed framework reduces training time by 5.9%, inference latency by 15.4%, GPU memory consumption by 8.8%, and CPU utilization by 6.5%.

Table 13. Runtime performance on the benchmark datasets.

Dunnhumby					
Method	Training (s/Epoch)	Inference (ms)	GPU (GB)	CPU (%)	Throughput
Causal forecasting	53.41±0.72	2.31±0.05	5.84±0.08	61.8±1.2	432±8
Offline feature-based pricing	49.86±0.66	2.18±0.05	5.42±0.07	59.4±1.1	456±7
Bayesian pricing	47.33±0.61	2.07±0.04	5.08±0.06	57.3±1.0	479±7
Contextual bandits	44.81±0.58	1.95±0.04	4.73±0.06	55.6±1.0	506±7
Personalized pricing	42.94±0.54	1.88±0.04	4.51±0.05	54.1±0.9	528±6
Algorithmic pricing	41.36±0.52	1.82±0.03	4.34±0.05	52.7±0.9	541±6
Proposed	38.92±0.48	1.54±0.03	3.96±0.04	49.3±0.8	639±5

Table 14 provides an overview of the proposed framework's scalability when the dataset is expanded from 50K to 2M observations. Training takes a maximum of 4681.9 seconds with the largest dataset, compared to 118.6 seconds with the smallest. Inference latency is impacted far less, going from 1.42 ms to 1.84 ms. GPU memory consumption grows from 2.31 GB to 6.18 GB. Conversely, the system throughput is reduced from 702 to 571 observations per second across the various dataset sizes. The measured speedup for the smallest dataset was 11.3% and increased to 14.4% for the largest.

Table 14. Scalability analysis of the proposed framework.

Dataset size	Training time (s)	Inference (ms)	GPU Memory (GB)	Throughput (Samples/s)	Speedup (%)
50K	118.6±1.5	1.42±0.03	2.31±0.03	702±8	11.3±0.4
100K	234.9±2.3	1.48±0.03	2.86±0.04	684±7	11.9±0.4
250K	586.7±4.8	1.55±0.03	3.41±0.04	651±7	12.6±0.5
500K	1172.8±8.6	1.63±0.04	3.95±0.05	624±6	13.2±0.5
1M	2344.1±15.2	1.71±0.04	4.82±0.06	596±6	13.8±0.6
2M	4681.9±28.5	1.84±0.05	6.18±0.08	571±5	14.4±0.6

Table 15 summarizes the performance comparison of the proposed model against the strongest competing baseline. The throughput improvements are the most notable, with a 18.11% gain. A 15.38% reduction follows this in inference time. Additional improvements include reductions in computational cost by 9.34% in FLOPs, 8.76% in GPU memory, 6.45% in CPU time, 5.92% in the number of trainable parameters, and 5.90% in training time. Considering all proposed baseline comparisons, the strongest competing baseline shows a reduction in overall computational cost of 10.84%.

Table 15. Overall computational efficiency improvement over the strongest baseline.

Metric	Improvement (%)
Training time reduction	5.90±0.31
Inference time reduction	15.38±0.44
GPU Memory reduction	8.76±0.28
CPU utilization reduction	6.45±0.26
Throughput increase	18.11±0.47
Parameter reduction	5.92±0.23
FLOPs reduction	9.34±0.29
Overall computational gain	10.84±0.33

4.6. Ablation analysis

The proposed framework for causal treatment-effect estimation, confounder adjustment, quantile response modeling, and DCPO has been evaluated. The effects of removing each component on the performance of the complete framework have been shown in Table 16. We assess the proposed framework by ablating each component as follows: treatment-effect estimation, confounder adjustment, quantile response modeling, and the DCPO. It discusses the removal of each component from the full framework. Removing treatment-effect estimation causes the largest performance drop. In this case, revenue improvement decreased from 11.48% to 8.84% (a loss of 2.64 percentage points), and profit improvement decreased from 11.82% to 9.06% (a loss of 2.76 percentage points). Simultaneously, the absolute treatment-effect expectations (ATEE) and the quantile treatment-effect expectations (QTEE) increased from 0.084 and 0.101 to 0.132 and 0.141, respectively, and DCPO decreased from 0.931 to 0.861. Figure 7 shows the effects on performance of the various components of the framework.

Table 16. Overall component-wise ablation analysis.

Configuration	Revenue (%)	Profit (%)	ATEE	QTEE	DCPO
w/o Treatment-effect estimation	8.84±0.16	9.06±0.15	0.132±0.003	0.141±0.003	0.861±0.007
w/o Confounder adjustment	9.27±0.15	9.54±0.15	0.118±0.003	0.129±0.003	0.882±0.006
w/o Quantile Response Modeling	9.74±0.14	10.02±0.14	0.096±0.002	0.138±0.003	0.894±0.006
w/o DCPO	10.48±0.13	10.74±0.13	0.091±0.002	0.112±0.002	0.903±0.006
Complete Framework	11.48±0.12	11.82±0.12	0.084±0.002	0.101±0.002	0.931±0.006




Figure 7. Ablation contribution analysis of the proposed framework. (a) Revenue degradation from removing individual components of the framework, illustrating the contribution of treatment-effect estimation, confounder adjustment, quantile response modeling, and DCPO optimization to performance. (b) Component-wise ablation performance for revenue, profit, ATEE, QTEE, and DCPO metrics. (c) Relative ablation performance across the evaluation metrics with respect to the complete framework. (d) Relative contribution from each component based on the observed revenue degradation.

Table 17 compares the contribution of each framework component on the Dunnhumby and Rossmann datasets separately. On the Dunnhumby dataset, the complete framework achieves the highest revenue improvement of 11.48% and a DCPO value of 0.931, whereas removing the treatment-effect estimation module reduces these values to 9.06% and 0.866, respectively. Similar reductions are observed after removing the confounder adjustment, quantile response modeling, and DCPO components, resulting in revenue values of 9.52%, 10.04%, and 10.73%, respectively. Comparable trends are obtained on the Rossmann dataset, where the complete framework records a revenue improvement of 10.97% and a DCPO value of 0.925, compared with 8.61% and 0.856 after removing the treatment-effect estimation module.

Table 17. Dataset-wise component ablation analysis.

Configuration	Dunnhumby		Rossmann	
	Revenue (%)	DCPO	Revenue (%)	DCPO
w/o Treatment-Effect Estimation	9.06±0.15	0.866±0.007	8.61±0.15	0.856±0.007
w/o Confounder Adjustment	9.52±0.14	0.885±0.006	9.01±0.14	0.878±0.006
w/o Quantile Response Modeling	10.04±0.13	0.898±0.006	9.43±0.13	0.890±0.006
w/o DCPO	10.73±0.13	0.907±0.005	10.22±0.12	0.900±0.005
Complete Framework	11.48±0.12	0.931±0.006	10.97±0.11	0.925±0.005

Table 18 shows changes to individual business performance indicators when components of the proposed framework are removed. The largest business metric losses from component removal are due to the absence of the treatment-effect estimation module. Business metric losses due to this module's removal are a revenue loss in 2.64%, a loss of 2.47% in sales, a loss of 2.76% in profit, a loss of customer expense of 2.58%, and a loss of 2.31% in the conversion rate. Losses of 1.95% to 2.28% in business metrics, as the confounder adjustment module is removed, are the second largest. Removal of quantile response modeling results in customer expense loss of 1.92%, and losses in other business metrics are marginally larger than those resulting from the removal of the DCPO, which are in the range of 0.88% to 1.08%. Of the business metrics evaluated, profit loss due to the removal of any module is the largest, and conversion rate loss is the smallest.

Table 18. Component contribution to individual business performance indicators.

Removed Component	Revenue Loss (%)	Sales Loss (%)	Profit Loss (%)	Cust. Exp. Loss (%)	Conv. Rate Loss (%)
Treatment-Effect Estimation	2.64±0.08	2.47±0.08	2.76±0.09	2.58±0.08	2.31±0.07
Confounder Adjustment	2.21±0.07	2.08±0.07	2.28±0.08	2.16±0.07	1.95±0.06
Quantile Response Modeling	1.74±0.06	1.63±0.06	1.80±0.06	1.92±0.06	1.56±0.05
DCPO	0.98±0.04	0.91±0.04	1.08±0.04	1.02±0.04	0.88±0.03

Table 19 presents the computational overhead introduced by each component of the proposed framework. The complete framework requires 38.92 seconds for training, 1.54 ms for inference, 3.96 GB of GPU memory, and 5.72 million trainable parameters, while achieving a throughput of 639 samples per second. Removing the treatment-effect estimation module reduces the training time to 31.74 seconds and inference latency to 1.32 ms, followed by the configurations without confounder adjustment, quantile response modeling, and DCPO. A progressive increase in computational requirements is observed as additional modules are incorporated, with training time increasing from 31.74 seconds to 38.92 seconds and GPU memory consumption increasing from 3.18 GB to 3.96 GB. The parameter count increases from 4.96 million for the reduced configuration without treatment-effect estimation to 5.72 million for the complete framework, while throughput decreases from 708 to 639 samples per second.

Table 19. Computational contribution of individual framework components.

Configuration	Training (s)	Inference (ms)	GPU (GB)	Parameters (M)	Throughput
w/o Treatment-Effect Estimation	31.74±0.42	1.32±0.03	3.18±0.04	4.96	708±6
w/o Confounder Adjustment	34.18±0.44	1.38±0.03	3.36±0.04	5.18	691±6
w/o Quantile Response Modeling	36.42±0.46	1.46±0.03	3.74±0.04	5.46	662±6
w/o DCPO	37.21±0.47	1.49±0.03	3.88±0.04	5.61	648±5
Complete Framework	38.92±0.48	1.54±0.03	3.96±0.04	5.72	639±5

Table 20 presents the progressive development of the proposed framework by incrementally adding each component to the baseline pricing model. The baseline configuration achieves a revenue improvement of 7.82%, profit improvement of 8.04%, ATEE of 0.151, QTEE of 0.168, and DCPO of 0.791. Adding the treatment-effect estimation module increases revenue and profit to 9.06% and 9.31%, respectively, while reducing ATEE to 0.119 and QTEE to 0.151. The inclusion of confounder adjustment further improves revenue to 9.88% and profit to 10.15%, with additional reductions in ATEE (0.101) and QTEE (0.132). Incorporating quantile response modeling increases revenue and profit to 10.73% and 11.01%, while reducing QTEE to 0.109. The complete framework with the DCPO achieves the highest revenue improvement (11.48%), profit improvement (11.82%), and DCPO value (0.931), together with the lowest ATEE (0.084) and QTEE (0.101) values among all evaluated configurations.

Table 20. Progressive construction of the proposed framework.

Framework Configuration	Revenue (%)	Profit (%)	ATEE	QTEE	DCPO
Baseline Pricing Model	7.82±0.18	8.04±0.17	0.151±0.004	0.168±0.004	0.791±0.009
+ Treatment-Effect Estimation	9.06±0.16	9.31±0.16	0.119±0.003	0.151±0.003	0.844±0.008
+ Confounder Adjustment	9.88±0.15	10.15±0.15	0.101±0.002	0.132±0.003	0.883±0.007
+ Quantile Response Modeling	10.73±0.13	11.01±0.13	0.091±0.002	0.109±0.002	0.907±0.006
+ DCPO	11.48±0.12	11.82±0.12	0.084±0.002	0.101±0.002	0.931±0.006

Table 21 summarizes the incremental contribution of each newly introduced component during the progressive construction of the framework. The treatment-effect estimation module provides the largest initial improvement, increasing revenue by 1.24%, profit by 1.27%, reducing ATEE by 21.19%, and increasing DCPO by 0.053. The addition of confounder adjustment contributes further improvements of 0.82% in revenue and 0.84% in profit, while reducing ATEE by 15.13% and increasing DCPO by 0.039. Quantile response modeling provides additional gains of 0.85% in revenue and 0.86% in profit, together with a 9.90% reduction in ATEE and a 0.024 increase in DCPO. The final integration of the DCPO contributes an additional 0.75% revenue improvement, 0.81% profit improvement, 7.69% ATEE reduction, and 0.024 increase in DCPO.

Table 21. Incremental contribution of each framework component.

Component	Revenue	Profit	ATEE	DCPO
Treatment-Effect Estimation	+1.24	+1.27	−21.19	+0.053
Confounder Adjustment	+0.82	+0.84	−15.13	+0.039
Quantile Response Modeling	+0.85	+0.86	−9.90	+0.024
DCPO	+0.75	+0.81	−7.69	+0.024

4.7. Statistical significance analysis

Although the experimental results consistently demonstrate the superiority of the proposed framework across multiple datasets and evaluation metrics, statistical hypothesis testing is required to verify that the observed improvements are not caused by random experimental variation. Therefore, several complementary statistical analyses are performed, including the Wilcoxon signed-rank test, Friedman test, Nemenyi post-hoc test, Cohen's effect size analysis, and confidence interval estimation. All statistical analyses are performed using the three independent experimental runs reported throughout the previous subsections with a significance level of $\alpha = 0.05$. Table 22 reports the Wilcoxon signed-rank test results between the proposed framework and each competing method. All pairwise comparisons produce p -values below the significance threshold of 0.05, with values ranging from 0.0008 to 0.0063. The smallest p -value is obtained for the comparison with (0.0008), followed by offline feature-based pricing (0.0014), Bayesian pricing (0.0021), contextual bandits (0.0036), personalized pricing (0.0042), and algorithmic pricing (0.0063). Each comparison is categorized as statistically significant according to the adopted hypothesis testing criterion.

Table 22. Wilcoxon signed-rank test comparing the proposed framework with competing methods.

Comparison	p -value	Decision
vs. Causal Forecasting	0.0008	Significant
vs. Offline Feature-Based Pricing	0.0014	Significant
vs. Bayesian Pricing	0.0021	Significant
vs. Contextual Bandits	0.0036	Significant
vs. Personalized Pricing	0.0042	Significant
vs. Algorithmic Pricing	0.0063	Significant

Table 23 presents the average rankings obtained from the Friedman test across all evaluated datasets and performance metrics. The proposed framework achieves the lowest average rank of 1.12, followed by algorithmic pricing with 2.34 and personalized pricing with 2.98. The remaining methods obtain higher average ranks, including contextual bandits (4.17), Bayesian pricing (5.28), offline feature-based pricing (6.14), and (7.00). The Friedman statistic is calculated as 31.82 with a corresponding p -value of 0.0003, which is below the selected significance level of 0.05. The ranking distribution places the proposed framework at the first position across the aggregated evaluation settings.

Table 23. Average Friedman rankings of the evaluated methods.

Method	Average Rank
Proposed	1.12
Algorithmic Pricing	2.34
Personalized Pricing	2.98
Contextual Bandits	4.17
Bayesian Pricing	5.28
Offline Feature-Based Pricing	6.14
Causal Forecasting	7.00
Friedman Statistic	31.82
p -value	0.0003

Table 24 summarizes the Nemenyi post-hoc comparisons between the proposed framework and the competing methods. The rank difference between the proposed framework and algorithmic pricing is 1.22, while larger differences are observed for personalized pricing (1.86), contextual bandits (3.05), Bayesian pricing (4.16), offline feature-based pricing (5.02), and (5.88). The obtained critical difference value is 1.09, and all pairwise rank differences exceed this threshold. Consequently, all comparisons are categorized as significant according to the Nemenyi test criterion. The ranking relationship among the evaluated approaches and the critical difference obtained from the Nemenyi post-hoc analysis are visualized in Figure 8.

Table 24. Nemenyi post-hoc comparison with the proposed framework.

Method	Rank Difference	Significant
Algorithmic Pricing	1.22	Yes
Personalized Pricing	1.86	Yes
Contextual Bandits	3.05	Yes
Bayesian Pricing	4.16	Yes
Offline Feature-Based Pricing	5.02	Yes
Causal Forecasting	5.88	Yes
Critical Difference		1.09




Figure 8. Statistical significance analysis of the proposed framework against competing pricing methods. (a) Critical difference (CD) diagram showing the Friedman average rankings of all evaluated methods, with the proposed framework achieving the best average rank. (b) Wilcoxon signed-rank test results comparing the proposed framework with each competing method based on the corresponding p -values. (c) Nemenyi post-hoc comparison showing the average-rank differences between the proposed framework and competing methods relative to the critical difference threshold. (d) Effect-size analysis using Cohen's d with 95% confidence intervals across revenue, profit, ATEE, QTEE, and DCPO.

Table 25 presents the Cohen's d effect sizes together with the corresponding 95% confidence intervals for the major evaluation metrics. The largest effect size is obtained for ATEE with a value of 2.11 and a confidence interval of [1.78, 2.43], followed by QTEE with an effect size of 1.97 and an interval of [1.65, 2.29]. The business performance metrics also exhibit large effect sizes, including profit (1.94, [1.63, 2.26]), revenue (1.86, [1.52, 2.19]), and DCPO (1.83, [1.49, 2.15]). All reported confidence intervals remain above the commonly accepted large-effect threshold, with no interval crossing zero.

Table 25. Effect size analysis of the proposed framework.

Metric	Cohen's d	95% CI
Revenue	1.86	[1.52, 2.19]
Profit	1.94	[1.63, 2.26]
ATEE	2.11	[1.78, 2.43]
QTEE	1.97	[1.65, 2.29]
DCPO	1.83	[1.49, 2.15]

4.8. Discussion

The results of the experiments show that integrating causal inference and distribution-aware quantile modeling improves over the standard strategy for pricing evaluation based on average treatment-effects and predictive performance. Estimating treatment-effects and responses along conditional quantiles gives the proposed framework the means to capture pricing strategies' effects and responses across the distribution. Treatment-effect estimation, business performance, and improved application of the method across different datasets beyond the original all indicate that the proposed method improves pricing strategies in different retail pricing contexts. The ablation studies and computational analysis indicate that each of the framework's building blocks performs a complementary role. Establishing a treatment-effect and removing confounding provide a causal framework for pricing strategies, and modeling response along quantiles provides a means to capture customer responses across the distribution. The DCPO incorporates all of these building blocks of the framework and, thus, provides significant business performance at a relatively low computational cost. The framework's construction and design also provide the means to capture customer responses along the distribution in the context of retail datasets of large customer volumes.

There are several limitations to this research that should be recognized. First, this framework is built on learning from observed retail data; therefore, it is assumed that the available covariates capture the most significant pricing decision-related confounding factors. In the presence of unobserved confounding factors, treatment-effects are likely to be biased. In addition, this study is based on only two publicly available retail data sets; this may fail to describe pricing behavior in other business domains or in more dynamic markets. This framework is built on the use of defined treatments to capture pricing interventions. Consequently, the framework is flexible to exclude the competitive nature of the pricing interventions, the constraints in the supply chain, and the most dynamic nature of the pricing interventions.

5. Conclusions

This study's main innovation is a unified causal-distributional framework for pricing evaluation. Rather than treating pricing as a purely predictive problem or choosing a policy based solely on its expected average business outcome, this framework assesses whether a pricing intervention improves outcomes and whether that improvement is consistent across different parts of the business-performance distribution. This framework is appropriate for observational retail data, where an average response that appears positive may still hide significantly negative or positive effects across customer, product, and performance segments. This framework also includes a few other important innovations. First, this framework uses treatment-specific outcome modeling and propensity score estimation, along with doubly robust and orthogonal treatment-effect estimation, to reduce the effects of observed confounding and nuisance-model estimation. Second, the framework uses heterogeneous treatment-effects to estimate price shifts at different business-performance levels, which extends pricing evaluation from a single, population-level effect to lower, median, and higher levels of the business-performance distribution. Third, the DCPO converts these causal and distributional estimates to a policy-evaluation criterion by balancing the average causal gain with the variation in quantile-specific treatment-effects. Therefore, DCPO's main contribution is not a new revenue-maximization criterion; it provides new pricing policies where estimated causal benefits are large on average and

stable across the business-performance distribution. Future research will investigate continuous pricing interventions, online and reinforcement-learning-based pricing optimization, multi-objective pricing strategies, incorporation of competitive market dynamics and supply-chain constraints, and foundation-model-driven causal analytics for adaptive pricing in large-scale business applications.

Use of Generative-AI tools declaration

The author declares they have not used any Artificial Intelligence (AI) tools in the creation of this article.

Conflict of interest

The author declares no conflicts of interest.

References

1. D. Schultz, J. Stephan, J. Sieber, T. Yeh, M. Kunz, P. Doupe, et al., Causal forecasting for pricing, *arXiv*, 2023. <https://doi.org/10.48550/arXiv.2312.15282>
2. D. Simchi-Levi, C. Wang, Pricing experimental design: Causal effect, expected revenue and tail risk, *Manag. Sci.*, **72** (2026), 1157–1174. <https://doi.org/10.1287/mnsc.2023.03209>
3. J. Tang, Z. Qi, E. Fang, C. Shi, Offline feature-based pricing under censored demand: A causal inference approach, *Manuf. Serv. Oper. Manag.*, **27** (2025), 535–553. <https://doi.org/10.1287/msom.2024.1061>
4. P. Aschersleben, W. J. Steiner, Dynamic pricing using flexible heterogeneous sales response models, *OR Spectr.*, **46** (2024), 29–72. <https://doi.org/10.1007/s00291-024-00756-0>
5. G. Uma Maheswari, S. Sethuraman, S. Ramanan, Contextual bandits for online markdown pricing for e-commerce, In: *Proceedings of the 7th joint international conference on data science & management of data (11th ACM IKDD CODS and 29th COMAD)*, 2024, 393–402. <https://doi.org/10.1145/3632410.3632448>
6. R. Miao, Z. Qi, C. Shi, L. Lin, Personalized pricing with invalid instrumental variables: Identification, estimation, and policy learning, *arXiv*, 2023. <https://doi.org/10.48550/arXiv.2302.12670>
7. A. Vomberg, C. Homburg, P. Sarantopoulos, Algorithmic pricing: Effects on consumer trust and price search, *Int. J. Res. Mark.*, **42** (2025), 1166–1186. <https://doi.org/10.1016/j.ijresmar.2024.10.006>
8. Y. C. Hsu, M. Huber, Y. M. Yen, Doubly robust estimation of direct and indirect quantile treatment-effects with machine learning, *arXiv*, 2023. <https://doi.org/10.48550/arXiv.2307.01049>
9. Z. Chen, S. Chen, Z. Qi, X. Chen, Z. Yang, Quantile-optimal policy learning under unmeasured confounding, *arXiv*, 2025. <https://doi.org/10.48550/arXiv.2506.07140>
10. Z. Ye, Z. Zhang, D. J. Zhang, H. Zhang, R. Zhang, Deep learning-based causal inference for large-scale combinatorial experiments: Theory and empirical evidence, *Manag. Sci.*, **72** (2026), 5611–5634. <https://doi.org/10.1287/mnsc.2024.04625>

11. Y. Xia, S. Narayanamoorthy, Z. Zhou, J. Mabry, Simulation-based benchmarking of reinforcement learning agents for personalized retail promotions, *arXiv*, 2024. <https://doi.org/10.48550/arXiv.2405.10469>
12. D. Aparicio, Z. Metzman, R. Rigobon, The pricing strategies of online grocery retailers, *Quant. Mark. Econ.*, **22** (2024), 1–21. <https://doi.org/10.1007/s11129-023-09273-w>
13. C. Glaeser, O. Yapar, Strategic consumer behavior and personalized pricing: A contextual bandit approach for retailers, *Kelley Sch. Bus. Res. Pap.*, 2025. <https://doi.org/10.2139/ssrn.5313650>
14. R. Koenker, G. Bassett Jr., Regression quantiles, *Econometrica*, **46** (1978), 33–50. <https://doi.org/10.2307/1913643>
15. S. Firpo, C. Pinto, Identification and estimation of distributional impacts of interventions using changes in inequality measures, *J. Appl. Econom.*, **31** (2016), 457–486. <https://doi.org/10.1002/jae.2448>
16. Y. Zhang, H. Li, G. Ren, Estimating heterogeneous treatment-effects in road safety analysis using generalized random forests, *Accid. Anal. Prev.*, **165** (2022), 106507. <https://doi.org/10.1016/j.aap.2021.106507>
17. S. N. Yousafzai, H. Shahbaz, A. Ali, A. Qamar, I. M. Nasir, S. Tehsin, et al., X-News dataset for online news categorization, *Int. J. Intell. Comput. Cybern.*, **17** (2024), 737–758. <https://doi.org/10.1108/IJICC-04-2024-0184>
18. S. N. Yousafzai, I. M. Nasir, S. Tehsin, D. S. Malik, I. Keshta, N. L. Fitriyani, et al., Multi-stage neural network-based ensemble learning approach for wheat leaf disease classification, *IEEE Access*, **13** (2025), 30101–30116. <https://doi.org/10.1109/ACCESS.2025.3541347>
19. S. N. Yousafzai, I. M. Nasir, S. Tehsin, N. L. Fitriyani, M. Syafrudin, FLTrans-Net: Transformer-based feature learning network for wheat head detection, *Comput. Electron. Agric.*, **229** (2025), 109706. <https://doi.org/10.1016/j.compag.2024.109706>
20. S. N. Yousafzai, I. M. Nasir, O. Saidani, R. Ghodhbani, Y. Gu, M. Syafrudin, et al., A multi-scale simplicial transformer with graph attention for facial emotion recognition, *Ain Shams Eng. J.*, **16** (2025), 103584. <https://doi.org/10.1016/j.asej.2025.103584>
21. S. Tehsin, A. Hassan, F. Riaz, I. M. Nasir, DaGAM-Trans: Dual graph attention module-based transformer for offline signature forgery detection, *Results Eng.*, **27** (2025), 106425. <https://doi.org/10.1016/j.rineng.2025.106425>
22. S. N. Yousafzai, F. N. Al-Wesabi, H. Alsolai, S. A. Ebad, I. M. Nasir, E. Fadhal, et al., Advanced clustering and transfer learning based approach for rice leaf disease segmentation and classification, *PeerJ Comput. Sci.*, **11** (2025), e3018. <https://doi.org/10.7717/peerj-cs.3018>
23. S. N. Yousafzai, I. M. Nasir, S. Tehsin, M. A. Khan, J. Ahmad, W. Boulila, A fusion framework of transformer and CNN for non-small cell lung cancer classification, *Int. Conf. Smart Syst. Emerg. Technol.*, 2025, 162–173. https://doi.org/10.1007/978-3-031-91235-1_15
24. A. Ahrens, V. Chernozhukov, C. Hansen, D. Kozbur, M. Schaffer, T. Wiemann, An introduction to double/debiased machine learning, *arXiv*, 2025. <https://doi.org/10.48550/arXiv.2504.08324>

25. S. Wager, S. Athey, Estimation and inference of heterogeneous treatment-effects using random forests, *J. Am. Stat. Assoc.*, **113** (2018), 1228–1242. <https://doi.org/10.1080/01621459.2017.1319839>
26. E. H. Kennedy, Z. Ma, M. D. McHugh, D. S. Small, Non-parametric methods for doubly robust estimation of continuous treatment effects, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **79** (2017), 1229–1245. <https://doi.org/10.1111/rssb.12212>
27. S. Tang, Y. Fan, X. Liu, Doubly robust estimation for quantile treatment effect, *Statistics*, **60** (2026), 422–448. <https://doi.org/10.1080/02331888.2025.2560504>
28. S. R. Künnel, J. S. Sekhon, P. J. Bickel, B. Yu, Metalearners for estimating heterogeneous treatment effects using machine learning, *Proc. Natl. Acad. Sci. USA*, **116** (2019), 4156–4165. <https://doi.org/10.1073/pnas.1804597116>
29. M. Oprescu, V. Syrgkanis, Z. S. Wu, Orthogonal random forest for causal inference, *arXiv*, 2019. <https://doi.org/10.48550/arXiv.1806.03467>
30. N. B. Bhavikatta, N. V. R. Goluguri, Multi dataset LSTM analysis for accurate and sustainable supply chain forecasting, *Discov. Sustain.*, **7** (2026), 508. <https://doi.org/10.1007/s43621-026-02828-3>
31. J. Pearl, Causal inference in statistics: An overview, *Stat. Surv.*, **3** (2009), 96–146. <https://doi.org/10.1214/09-SS057>



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)