A Debiased Machine Learning Framework for Optimizing Price Promotion within E-commerce

Bo Zhang
Beyond Inc
Midvale, Utah, USA
zhangbo.eco@gmail.com

Sergazy Nurbavliyev
Beyond Inc
Midvale, Utah, USA
sergazy.nurbavliyev@gmail.com

Vaidyanath Areyur Shanthakumar
Beyond Inc
Midvale, Utah, USA
vareyurshanthakumar@beyond.com

Stephen Merrill
Beyond Inc
Midvale, Utah, USA
stephen.merrill0@gmail.com

Raymond J. Pan
Beyond Inc
Midvale, Utah, USA
rpan@beyond.com

Komson Chanprapan
Beyond Inc
Midvale, Utah, USA
econkc@gmail.com

Abstract

Machine learning models have been widely adopted for price promotion and dynamic pricing in retail. However, one key challenge these models often overlook is the inherent bias in the training data, which is typically observational rather than experimental. This bias can lead to confounded estimates of how discounts influence sales volumes, obscuring true causal relationships. Inspired by the fixed-effects regression model, widely used in panel data analysis for more reliable causal inference from observational data, we introduce the *Delta Method*. This approach leverages de-meaned variables in machine learning models to isolate within-product variation and estimate product-level treatment effects. Using data from an online furniture retailer, we found that the *Delta Method* not only enhances out-of-sample prediction accuracy but, more importantly, provides a clearer and more interpretable understanding of the causal relationship between discount levels and sales at both the individual product and aggregate levels. Furthermore, a real-world experiment on the retailer's website demonstrated that the *Delta Method* led to a 3% increase in revenue and a 2% increase in profit compared to traditional methods, confirming its practical value for price promotion optimization.

CCS Concepts

• **Theory of computation** → *Theory and algorithms for application domains*; • **Applied computing** → *Economics, Machine Learning*.

Keywords

Price Promotion, Causal Inference, Delta Method, Demand Model, Revenue Optimization

KDD Applied Data Science (ADS) Track, Toronto, ON, Canada

ACM ISBN 978-1-4503-XXXX-X/2025/06
https://doi.org/XXXXXXX.XXXXXXX

ACM Reference Format:
Bo Zhang, Sergazy Nurbavliyev, Vaidyanath Areyur Shanthakumar, Stephen Merrill, Raymond J. Pan, and Komson Chanprapan. 2025. A Debiased Machine Learning Framework for Optimizing Price Promotion within E-commerce. In *KDD: Applied Data Science (ADS) Track, August 03-05, 2025, Toronto, ON, Canada.* ACM, Toronto, Canada, 8 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Many retailers rely on price promotions to attract new customers, increase user engagement, and increase sales volumes. However, identifying the optimal set of price promotions in a large assortment of products can be challenging. Traditional retail practices often rely on rule-based approaches, typically aiming for certain profit margins or revenue targets. Although straightforward, these methods provide limited insight into how promotions truly affect sales. Consequently, there is growing interest in data-driven strategies that adapt pricing decisions more dynamically [3, 6, 9, 11, 12, 17].

A common solution is the "predict and optimize" framework, which involves two steps. First, a predictive model estimates the lift in sales volumes under different promotion levels, revealing how sales might respond to varying discounts. Second, given these estimates, the retailer solves an optimization problem (e.g., a knapsack-type problem with budget constraints) to maximize total revenue or profit under budget constraints. Intuitively, when budget limits are present, deeper discounts should be assigned to products more sensitive to price changes, while shallower discounts suffice for less sensitive products.

However, this approach faces two key challenges. First, the "predict" component goes beyond simply forecasting demand. More importantly, it should assess how sales respond to different promotion levels, which is known as causal impact, as this insight is crucial for the "optimize" [9]. Many recent studies fail to discuss this issue, but it can be the key to the success of the algorithm [12, 17, 21]. Second, the model needs to be capable of measuring the product-level elasticity estimates. With large retailers managing millions of products, it is essential for algorithms to determine which products are highly responsive to promotions and which are less sensitive, thereby enabling more effective decision-making.

Establishing causal effects typically requires large-scale experiments or A/B tests. Unfortunately, conducting large-scale pricing experiments can be expensive and may face regulatory or ethical hurdles related to price discrimination. Moreover, experimental results can quickly become outdated due to seasonality and shifting market conditions. As a result, many studies are based on observational data instead of experimental data [6, 9, 12, 17]. However, relying on observational data introduces confounding factors, which can bias causal estimates unless appropriately addressed. Our study aims to mitigate this issue by proposing a new method to estimate the causal relationship between discounts and product sales using observational data.

To address this issue, we propose a new approach called the *Delta Method*. Inspired by fixed-effects regression models, which are commonly used in economics and social sciences [18], our approach focuses on within-product variation over time. Instead of comparing different products, we compare each product to itself at different points in time, effectively "differencing out" time-invariant confounders such as baseline product popularity or brand effects. Using de-meaning the variables, we expect to obtain a better causal impact of promotions on product sales. In addition, the use of non-linear machine learning models allows us to capture heterogeneous effects at the product level. By de-meaning the variables before inputting them into tree-based machine learning models, we retain the flexibility and predictive power of these non-linear methods while achieving a clearer causal interpretation.

We validate our approach using a data set of more than half million products spanning two months from a large e-commerce company. In the offline evaluation, we found that the machine learning model with the *Delta Method* method largely improves the causal effects of promotions on product sales. Specifically, we observe that the estimated elasticity is aligned with fundamental economic principles—higher discounts lead to higher sales—at both the product and aggregate levels. The out-of-sample prediction accuracy also slightly improves using the *Delta Method*. Based on the offline evaluation, we chose the LightGBM model with *Delta Method* as the demand model in the first stage and applied an integer programming to solve the product promotion levels at the second stage. To further evaluate the performance of the whole algorithm on the business, we also conducted a large-scale A/B test on millions of products for two weeks. The results show that the revenue significantly increases by 3% and the profit increases by 2% using the *Delta Method* compared to the current algorithm that we do not use *Delta Method*.

The main contributions are summarized as follows:

- We propose an end-to-end framework based on a debiased machine learning model and integer optimization model to suggest the price promotion levels for products, as shown in Figure 1.
- We employ the machine learning model with the *Delta Method* to improve the causal relationship between predicted product sales and discounts using observational data.
- We successfully deploy the proposed price promotion algorithm in real-world e-commerce retail. A/B test results reveal that the algorithm significantly boosts key metrics, including revenue and profit.

Figure 1: The production framework for the price promotion system

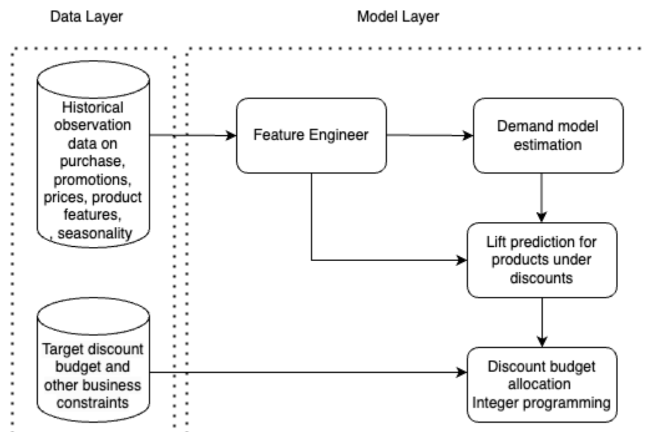

2 Related Work

2.1 Two Approaches to Pricing Optimization: Predict-Then-Optimize vs. Reinforcement Learning

Pricing optimization in e-commerce and retail has historically been divided into two major paradigms. The first, often referred to as predict-then-optimize, involves building a demand model to predict sales under various prices (or discounts) and then solving an optimization problem to maximize revenue or profit subject to practical constraints [5, 10, 20]. This paradigm has the advantage of model interpretability in the prediction step and can incorporate sophisticated optimization routines (e.g., knapsack-like budget constraints).

The second paradigm uses reinforcement learning (RL), where an agent iteratively explores different pricing policies and updates its strategy based on observed rewards [4, 14, 16, 22]. RL methods excel in highly dynamic, uncertain environments, but can be data-hungry and less transparent about how price sensitivities are derived. Although RL has found success in applications such as real-time bidding and ride-sharing surge pricing, it often requires extensive online experimentation to converge to optimal policies [14], which may be infeasible for many retailers who are wary of customer reactions or regulatory constraints.

In this paper, we focus on the predict-then-optimize framework, which is widely adopted in e-commerce pricing due to its flexibility and conceptual clarity. It allows for the integration of explicit causal or econometric tools into the predictive step before running an optimization routine.

2.2 Classical Regression and Machine Learning Methods

Before the advent of modern machine learning, many retailers and academics used classical regression methods–often linear models with price or discount as a main regressor–to estimate a single "average" elasticity across all products [6, 11, 19].

While such approaches are relatively easy to interpret, they tend to impose uniform price sensitivity across heterogeneous products, potentially over- or under-estimating the true responsiveness of individual items. Variations of regression-based methods have attempted to introduce group-level or segment-level elasticities [11], but this approach depends heavily on how the products are grouped, and the resulting elasticity estimates remain coarse.

With the emergence of machine learning, retailers began fitting more flexible models–such as random forests, gradient boosting (e.g., XGBoost, LightGBM), and neural networks–that can capture non-linearities and product-level heterogeneity [2, 3, 9, 12, 17, 21]. These ML models often achieve stronger predictive performance compared to linear regression and can, in principle, deliver individual-level elasticity or uplift estimates [7, 15]. However, *interpretability* becomes a concern: it is challenging to extract clear elasticity coefficients from a deep neural network or a complex ensemble of trees. Moreover, standard ML algorithms are typically agnostic to economic constraints, such as the law of demand, unless special structures–like monotonic constraints–are enforced [17].

2.3 Experiments and Causal Inference

The ideal way to obtain product-level elasticity is often to run large-scale experiments–such as A/B testing–so that causal effects can be precisely estimated without bias [1, 7]. Indeed, reinforcement learning methods often rely on live experimentation to explore different price actions [14]. Yet, field experiments in e-commerce can be prohibitively expensive or legally delicate, especially if frequent price variations raise concerns about price discrimination. This tension leaves many retailers with observational data in which discounts are not randomly assigned but systematically determined by marketing teams, inventory needs, or product popularity.

Consequently, researchers have turned to econometric techniques—like fixed-effects, instrumental variables, and difference-in-differences—to approximate experimental conditions [11, 18]. Fixed-effects approaches, in particular, exploit within-product variation over time to filter out time-invariant confounders such as baseline popularity or brand perception [11, 18]. However, these techniques are often limited to linear models, making them less adaptable to complex, non-linear machine learning routines that can better capture heterogeneous relationships.

2.4 Our Contribution: A Product-Level and Causally Informed Demand Model

In this work, we bridge the gap between flexible machine learning and econometric techniques for causal inference. We propose the *Delta Method*, which combines de-meaning (from fixed-effects ideas) with machine learning models to estimate product-level treatment effects more accurately and transparently. Implementing these transformations in the machine learning models retains the flexibility of modern predictive algorithms while yielding elasticity estimates aligned with fundamental economic principles.

2.4.1 *Alignment with Economic Theory.*

De-meaning the discount and sales variables at the product level ensures that predicted relationships between discount depth and sales volume are consistent with the law of demand (monotonic trend), avoiding globally linear assumptions that can oversimplify consumer behavior.

2.4.2 *Granular, Product-Level Elasticities.*

Unlike classical regression approaches that rely on average or group-level elasticity measures, the *Delta Method* delivers individualized treatment effects. Retailers can thus identify which products are most price-sensitive, enabling more precise and profitable discount allocations.

2.4.3 *Reduced Dependence on Costly Experiments.*

Large-scale pricing experiments pose logistical, financial, and ethical challenges, particularly around price discrimination. Leveraging an observational dataset, our method approximates many benefits of random assignment while avoiding these hurdles, thereby offering robust causal insights drawn directly from real-world data.

Overall, the *Delta Method* elevates the predict-then-optimize paradigm by enhancing causal fidelity in demand estimation and preserving the accuracy of contemporary ML techniques. This hybrid framework is readily scalable to e-commerce systems with frequent price updates and large product assortments, providing interpretable, product-level elasticity estimates to guide legal and data-driven pricing strategies.

3 Method

This section delves into the key components of the price promotion system. The primary goal is to develop an algorithm that determines optimal price promotion levels for millions of products in upcoming promotional events. The assigned discounts should maximize revenue or profit while adhering to a budget constraint on promotional costs.

Our approach follows the conventional two-step "predict and optimize" framework. In the first step, we build a demand model to quantify the relationship between quantity sold and discount levels. Using this model, we estimate lift by predicting counterfactual demand across different discount levels. In the second step, we leverage these lift estimates in a mathematical optimization process to determine the optimal discount levels while satisfying predefined business constraints.

The next two sections provide a detailed breakdown of these two steps.

3.1 Demand Model

The demand model estimation is a critical component of the price promotion algorithm, designed to measure the sensitivity of product sales to price promotions. It comprises several key elements:

First, the model must provide accurate demand predictions, ensuring reliable forecasts of product quantity sold.

Second, it should capture the causal relationship between the product quantity sold and the discount level, effectively quantifying sales sensitivity to promotions. This is essential for generating counterfactual predictions—estimating demand under different promotion levels. According to the law of demand, the product quantity sold should increase as discount levels rise. Thus, the model should ideally exhibit a monotonic trend, where qty_{it} increases with dsc_{it}, aligning with standard economic theory.

Finally, the model must estimate the elasticity of product quantity sold to promotions at the product level. The core intuition behind the algorithm is to allocate deeper discounts to products that are more responsive to price reductions, maximizing the impact of promotions.

3.1.1 *Standard Method.*

Usually, the demand model is estimated using the traditional method has the basic structure of (1):

$$qty_{it} = f(dsc_{it}, X_{it}) \tag{1}$$

where qty_{it} represents the item quantity sold for product i on day t, dsc_{it} represents the discount level for product i on day t, and X_{it} represents a vector of covariates including product price, product categories, popularity of the product, day of week, and, and week of year. Here f is the demand function. As we mentioned in the related work section, previous studies [9, 12, 17, 21] have used regression models, tree-based models, and neural network models to predict demand.

Regression models are widely used in revenue management due to their simplicity and interpretability. They also benefit from well-established methods to address estimation bias. However, they estimate the average treatment effect across all products, and capturing group-level heterogeneity requires careful sample segmentation. The accuracy of this segmentation impacts results, and using group-level elasticity may limit the optimization algorithm's ability to differentiate product sensitivity to discounts. Additionally, linear models may have lower predictive accuracy than more complex models.

In contrast, tree-based models and neural networks provide higher predictive accuracy and better capture product-level sensitivity. However, they are less suitable for estimating causal effects.

3.1.2 *The Delta Method.*

To solve the above problem, we propose to use *Delta Method*. *Delta Method* was inspired by the "de-meaning" concept from fixed-effects regression models. Fixed-effects regression models are widely used for causal inference with longitudinal or panel data in many social science domains [18]. One advantage of these models is the ability to control for time-invariant variables, allowing them to better quantify the causal effect compared to basic regression models.

A basic regression model can be written as equation (2), while a fixed-effects regression model can be written as equation (3). From the fixed-effects regression model representation, it can be seen that the respective mean values across time are subtracted from both the dependent variable (y) and the independent variable (x). Intuitively, this results in all of the variance for the dependent (y) and the independent (x) variables being associated with the deviation from the mean. If the product ID is used as a fixed effect in the fixed-effects regression model, all the variance will be associated with the deviation from the mean quantity sold for each product under different discount levels.

$$y_{it} = X_{it}\beta + \alpha_i + u_{it} \tag{2}$$

$$y_{it} - \overline{y}_i = (X_{it} - \overline{X}_i)\beta + u_{it} - \overline{u}_i \tag{3}$$

For the *Delta Method*, we also apply this idea to the other models by subtracting the mean from the key dependent and independent variables (quantity sold and discount level) as shown in equation (4).

$$qty_{it} - \overline{qty}_i = f(dsc_{it} - \overline{dsc}_i, X_{it} - \overline{X}_i, X_{it}) \tag{4}$$

3.2 Optimization

After training the model using historical data, we generate the predicted product quantity sold (qty_i) given the current price (p_i) under different discounts (dsc_i). These lift estimates are used in a constrained optimization, which ultimately suggests the optimal discounts for all the products. In our operation, dsc can only be one of the following discounts: {0%, 5%, 10%, 15%, 20%, 25%}.

The discount optimization problem is formulated to maximize the overall revenue. Following [9], we assume there is no substitution effect among different products. In future work, one might extend the demand function to incorporate cross-terms that capture product interactions, turning the optimization into a multi-product problem where each product's discount affects others.

In our setting, we will ultimately assign one of the discounts from {0%, 5%, 10%, 15%, 20%, 25%} to each product. This suggests the usage of a multiple choice knapsack problem (MCKP) [13] with the goal of maximizing the revenue under the budget constraint. The full optimization can therefore be denoted as follows in equation (5):

$$\begin{aligned} \max_{dsc_i} \quad & \sum_i p_i \cdot qty_i(p_i, dsc_i) \\ \text{s.t.} \quad & \sum p_i \cdot qty_i(p_i, dsc_i) \cdot dsc_i < \text{Budget} \end{aligned} \tag{5}$$

where p_i is the current price of product i. To clarify, $\sum p_i \cdot qty_i(p_i, dsc_i) \cdot dsc_i$ represents the absolute discount cost multiplied by predicted demand. This cost must remain under the allocated budget.

Due to budget constraints, we are not able to select all the products at their optimal discount. Therefore, we face the problem of distributing the budget and selecting discounts under budget constraints. Despite this constraint, we can use this optimization method to explore these suboptimal limitations. Here, we consider the three different ways to optimize the target with limited budget.

The first method is to optimize revenue with a fixed budget, the second method is optimizing profit (revenue less discount cost) with a fixed budget, and the final method is optimizing revenue with a fixed discount rate.

We selected the first option for two reasons. First, the business goal is to generate more revenue without hurting the profit margin. Optimizing revenue instead of profit meets the goal. Second, using a fixed budget is consistent with how the budget is planned by the business. However, we may experiment with the fixed rate option in the future, which may help us to better control the budget.

Currently, we only consider a fixed number of discount options, as described above. Therefore, we are solving an integer programming problem. The integer programming is almost identical to linear programming. We solve this integer programming problem using the pulp package in Python. Finally, since the discount set is discrete, this problem remains tractable for millions of products by leveraging efficient CBC solvers and parallel processing.

4 Evaluation

In this section, we discuss the data that we use, the evaluation methods, and the performance of the models.

4.1 Data

The data we are using come from a leading e-commerce retailer selling furniture and home products. The data includes four main components: the purchase records for the products, the daily price and discounts for the products, the category and subcategories of the products, and the seasonality variables such as the day of the week and the week of the year. From these features, we constructed a daily-product level dataset.

As we mentioned in the method session, the outcome variable is the daily product quantity sold. The key dependent variable is the product discount of the day. For each product, one of the following discounts {0%, 5%, 10%, 15%, 20%, 25%} is assigned for each discount campaign. Each campaign may last anywhere from 2 days to 2 weeks.

4.2 Model Performance

In the offline evaluation, we use two main approaches to assess the model's performance with historical data. The key advantage of the *Delta Method* is its ability to provide a more accurate estimation of the causal relationship between discounts and product quantity sold. Additionally, we evaluate the model's prediction accuracy, a standard practice in machine learning.

To compare the *Delta Method* with the standard approach, we test six models: a standard regression model, LightGBM model, and neural network model, alongside their counterparts using the *Delta Method*. In the regression model, the standard approach applies a basic regression, while the *Delta Method* uses a fixed-effects regression. For LightGBM and neural networks (which combine linear and ReLU layers), the model structures remain consistent across both methods.

4.2.1 Causality. To evaluate the model's causal validity, we check whether the results align with the law of demand—where higher discounts lead to increased product quantity sold. We conduct a detailed comparison at both the aggregate and product levels.

First, we analyze the relationship between product quantity sold and discount levels. Table 1 presents the regression results. Column (1) shows the average product quantity sold at various discount levels, while Columns (2) and (3) examine the impact of different discount levels on sales. Column (1) indicates that products with a 5% discount have the highest average sales, whereas those with a 25% discount have lower sales. Since discounts are not assigned randomly, this basic regression captures existing correlations rather than causal effects. Column (2) further suggests that a 5% discount is the most effective, contradicting the law of demand. However, Column (3), which incorporates fixed effects, corrects this bias and reveals a proper monotonic relationship—higher discounts lead to higher product quantity sold.

One possible explanation is that more popular products are more likely to receive a 5% discount under the current pricing strategy. The basic regression results reflect differences across products, while the fixed-effects regression better isolates the causal impact by controlling for within-product variation over time. This suggests that the fixed-effects model more accurately captures the causal relationship between discounts and sales.

We also compare aggregate predicted results between the standard approach and the *Delta Method* using LightGBM and the neural network model. Table 2 shows the predicted average product quantity sold at different discount levels from the LightGBM model. Similar to the linear regression results, the standard method does not follow a monotonic trend, while the *Delta Method* correctly indicates that higher discounts lead to higher product quantity sold. Table 3 presents the neural network results, showing the same pattern.

Table 1: Regression results

	sample average (1)	standard regression (2)	FE regression (3)
dsc=0%	0.084		
dsc=5%	0.543	0.449^{***}	0.206^{***}
dsc=10%	0.418	0.345^{***}	0.252^{***}
dsc=15%	0.448	0.367^{***}	0.307^{***}
dsc=20%	0.308	0.221^{***}	0.369^{***}
dsc=25%	0.199	0.121^{***}	0.382^{***}

Note: Column (1) shows the average product quantity sold at different discounts. Column (2) and (3) show the coefficients of discounts in the standard and fixed-effect (FE) regression models, which represent the impact of the discount on sales. Significance levels: $^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

Table 2: LightGBM: Standard vs Delta

	sample average (1)	standard LightGBM (2)	delta LightGBM (3)
dsc=0%	0.084	0.266	0.224
dsc=5%	0.543	0.372	0.230
dsc=10%	0.418	0.299	0.246
dsc=15%	0.448	0.313	0.286
dsc=20%	0.308	0.202	0.320
dsc=25%	0.199	0.202	0.346

Note: Column (1) shows the average product quantity sold at different discounts. Column (2) and (3) show prediction product quantity sold under different discounts from the standard LightGBM model and the LightGBM with **Delta Method**

Table 3: Neural network: Standard vs Delta

	sample average (1)	standard network (2)	delta network (3)
dsc=0%	0.084	0.164	0.099
dsc=5%	0.543	0.227	0.180
dsc=10%	0.418	0.260	0.260
dsc=15%	0.448	0.274	0.330
dsc=20%	0.308	0.316	0.393
dsc=25%	0.199	0.294	0.451

Note: Column (1) shows the average product quantity sold at different discounts. Column (2) and (3) show prediction product quantity sold under different discounts from the standard neural network model and the neural network model with **Delta Method**.

The machine learning model incorporating the **Delta Method** has demonstrated superior performance in capturing causal relationships at an aggregated level. We anticipate a similar pattern

Figure 2: The quantitative relationship between discount levels and product quantity sold for 5 randomly selected products from the standard LightGBM

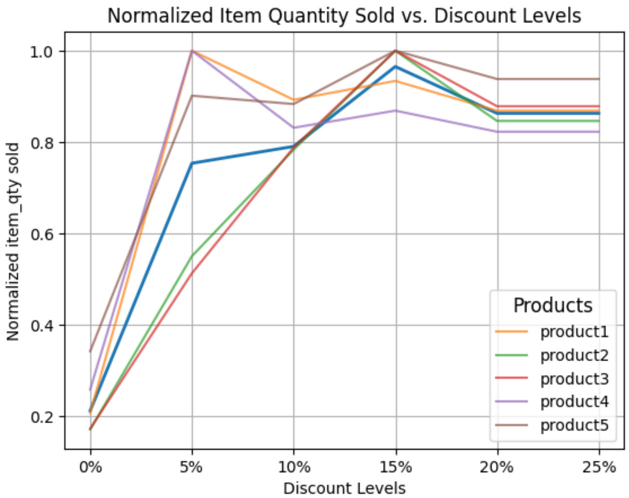

Figure 4: The quantitative relationship between discount levels and product quantity sold for 5 randomly selected products from the standard neural network model

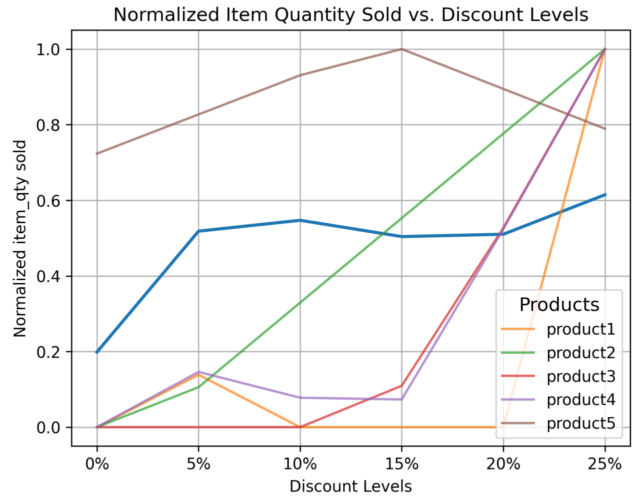

Figure 3: The quantitative relationship between discount levels and product quantity sold for 5 randomly selected products from the LightGBM with Delta Method

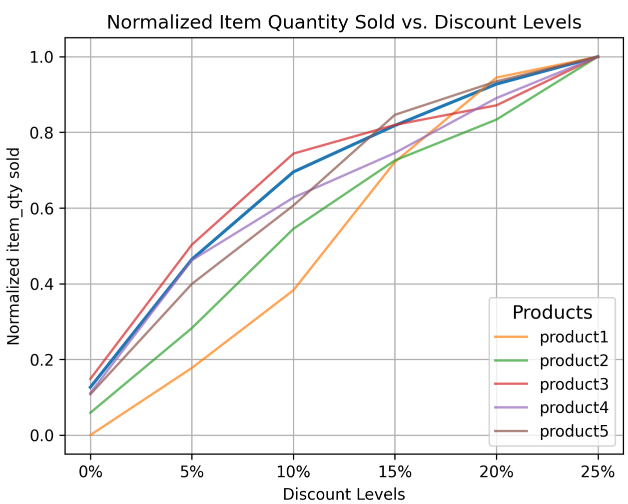

Figure 5: The quantitative relationship between discount levels and product quantity sold for 5 randomly selected products from the neural network model with Delta Method

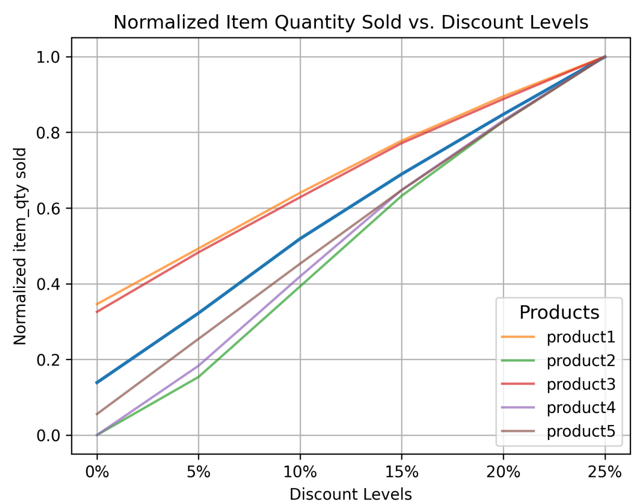

at the product level for most items, which is crucial for the algorithm's success. Both the nonlinear LightGBM and neural network models can estimate product-level elasticity. Using these models, we predict the product quantity sold under varying discount levels while holding other features constant.

Figures 2 and 3 depict the results for the standard LightGBM model and the LightGBM model with the **Delta Method**, respectively. To better illustrate trends, we normalize item quantities to the range [0,1] and randomly select five products from the sample. Notably, the LightGBM model with the **Delta Method** captures a more reasonable relationship between discounts and product quantity sold compared to the standard LightGBM model. Additionally, we observe heterogeneous discount effects across different products.

A similar pattern emerges with a simple two-layer neural network model. Figures 4 and 5 show the results for the standard neural network and the neural network incorporating the **Delta Method**. Notably, predictions from the **Delta Method**-enhanced neural network align with the law of demand, indicating that higher promotions lead to increased sales.

Overall, the **Delta Method** significantly improves causal inference, allowing for more reliable predictions of product quantity sold under different discount levels.

4.2.2 Prediction Accuracy. In addition to evaluating the quantitative relationship between discount levels and product quantity sold , we also assess the prediction accuracy of each model.

In this part, we aim to evaluate the performance of the proposed model by utilizing a comprehensive set of metrics to capture various aspects of its accuracy and reliability. These metrics are as follows:

- **Mean Absolute Percentage Error (MAPE)**: This metric measures the average percentage difference between the predicted and actual values, providing an intuitive sense of the prediction accuracy in relative terms. It is widely used to evaluate forecasting models, especially when the scale of the data varies.
- **Revenue-Weighted MAPE**: A variation of MAPE that weights errors based on revenue contribution. By prioritizing predictions for higher-revenue items, this metric ensures the model's performance is aligned with business priorities, focusing on areas with the greatest financial impact.
- **Pearson Correlation**: This statistic measures the linear relationship between predicted and actual values, with a value close to +1 indicating a strong positive correlation. It provides insights into how well the predicted trends align with the observed data.
- **Mean Absolute Error (MAE)**: MAE calculates the average magnitude of the errors between predictions and actual values, treating all deviations equally. It is useful for assessing the model's overall prediction accuracy in absolute terms, regardless of direction.
- **Median Absolute Deviation (MAD)**: MAD focuses on the median of the absolute errors, making it robust to outliers. This metric is particularly helpful for evaluating the model's performance in scenarios where extreme deviations are present.
- **Coefficient of Determination (R^2)**: Commonly known as R-squared, this metric quantifies the proportion of variance in the actual data that is explained by the model's predictions. A higher R^2 value indicates that the model effectively captures the variability in the data, demonstrating its explanatory power.

We assessed the model's performance on the test data, with the results presented in Table 4 and Table 5. As shown, models incorporating the **Delta Method** consistently outperform those using the standard approach across most evaluation metrics. The LightGBM model with **Delta Method** has the best performance in prediction accuracy using the test sample.

Given its strong performance in both causal inference and predictive accuracy, the LightGBM model with the **Delta Method** was selected as our demand model for online A/B testing and production deployment.

Table 4: Standard method model performance

	Regression	LightGBM	Neural Network
MAPE	1.914	0.852	1.116
weighted_MAPE	0.847	0.529	0.827
Person Correlation	0.230	0.757	0.195
MAE	6.872	3.875	3.927
MAD	4.610	1.605	2.032
R2	0.044	0.622	0.570

Table 5: Delta method model performance

	Regression	LightGBM	Neural Network
MAPE	1.188	0.836	0.798
weighted_MAPE	1.249	0.502	0.511
Person Correlation	0.636	0.840	0.831
MAE	5.508	3.712	3.687
MAD	2.371	1.554	1.549
R2	0.307	0.703	0.690

4.3 A/B testing

To evaluate the effectiveness of the new algorithm, we conducted a large-scale A/B test on an online retailer's platform. In this experiment, approximately 3 million products were randomly assigned to treatment and control groups. Products in the treatment group received discount assignments generated by the new algorithm, while those in the control group were assigned discounts using the existing algorithm.

The existing algorithm relies on an XGBoost model and does not account for causality, whereas the new algorithm incorporates the *Delta Method* to address this issue. Importantly, other key components—such as the optimization framework and the total discount budget—remained consistent across both groups.

To assess the impact of the new algorithm, we applied Lachenbruch's two-part test [8]. The results showed that the treatment group achieved a 3% increase in revenue and a 2% increase in profit compared to the control group. These improvements were statistically significant at the 5% level.

Based on these findings, we conclude that the new algorithm significantly improves key financial outcomes by addressing the causality issue inherent in the previous approach.

4.4 Deployment

Following the A/B test, the business decided to adopt the pricing promotions recommended by the machine learning algorithm. As a result, the algorithm was deployed into production. In addition to the components for data pulling, continuous model training, and inference illustrated in Figure 1, a monitoring system was implemented to track the algorithm's performance and project financial outcomes. Since its launch, the system has been functioning as expected.

5 Conclusion

In conclusion, price promotions remain a vital tool for retailers aiming to enhance revenue and profitability. While machine learning-based methods have become increasingly popular for optimizing price promotions, their application to historical data often results in models that are highly predictive but do not necessarily capture the causal relationship between price and quantity sold. On the other hand, traditional economic regression models, though interpretable, tend to focus on average treatment effects and exhibit lower predictive accuracy.

In this paper, we introduced a novel approach, the *Delta Method*, which synthesizes the strengths of both machine learning and traditional regression techniques. Drawing from the fixed-effects regression model commonly used in panel data analysis, we applied a de-meaning transformation to tree-based machine learning models. This allowed us to estimate product-level treatment effects, correcting for unobserved heterogeneity across products and focusing on variations over time within each product.

Our offline analysis, using data from an online furniture retailer, demonstrated that the *Delta Method* not only improved prediction accuracy on test data but also provided a more interpretable model of the relationship between discount levels and quantities sold. Furthermore, a real-world experiment on the retailer's website showed that the *Delta Method* outperformed traditional approaches, resulting in a 3% increase in revenue and a 2% increase in profit. These results underscore the importance of combining flexible machine learning models with careful causal inference methods to ensure both interpretability and higher profit outcomes in e-commerce. Future research might explore applying the *Delta Method* to more complex settings, such as scenarios with cross-product interactions or multi-period pricing decisions. Such extensions could further validate the adaptability and robustness of our approach across diverse retail environments.

Overall, our findings highlight the potential of the *Delta Method* as a powerful tool for improving the effectiveness of price promotions in e-commerce settings.

Acknowledgments

The authors would like to thank Beyond Inc for their invaluable support in helping us write this paper, providing the necessary resources and opportunities to pursue this research, and contributing to the broader community through their commitment to data-driven innovation. We also extend our gratitude to the colleagues and reviewers who offered constructive feedback and insights throughout the development of this work.

References

[1] Susan Athey and Guido W Imbens. 2015. Machine learning methods for estimating heterogeneous causal effects. *stat* 1050, 5 (2015), 1–26.

[2] Bobby Chen, Siyu Chen, Jason Dowlatabadi, Yu Xuan Hong, Vinayak Iyer, Uday Mantripragada, Rishabh Narang, Apoorv Pandey, Zijun Qin, Abrar Sheikh, et al. 2024. Practical Marketplace Optimization at Uber Using Causally-Informed Machine Learning. *arXiv preprint arXiv:2407.19078* (2024).

[3] Maxime C Cohen, Jeremy J Kalas, and Georgia Perakis. 2021. Promotion optimization for multiple items in supermarkets. *Management Science* 67, 4 (2021), 2340–2364.

[4] Arnoud V Den Boer. 2015. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science* 20, 1 (2015), 1–18.

[5] Adam N Elmachtoub and Paul Grigas. 2022. Smart "predict, then optimize". *Management Science* 68, 1 (2022), 9–26.

[6] Kris Johnson Ferreira, Bin Hong Alex Lee, and David Simchi-Levi. 2016. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & service operations management* 18, 1 (2016), 69–88.

[7] Pierre Gutierrez and Jean-Yves Gérardy. 2017. Causal inference and uplift modelling: A review of the literature. In *International conference on predictive applications and APIs*. PMLR, 1–13.

[8] Meisam Hejazinia, Majid Hosseini, and Bryant Sih. 2019. A/B Testing Measurement Framework for Recommendation Models Based on Expected Revenue. *arXiv preprint arXiv:1906.06390* (2019).

[9] Junhao Hua, Ling Yan, Huan Xu, and Cheng Yang. 2021. Markdowns in e-commerce fresh retail: a counterfactual prediction and multi-period optimization approach. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3022–3031.

[10] Shinji Ito and Ryohei Fujimaki. 2017. Optimization beyond prediction: Prescriptive price optimization. In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*. 1833–1841.

[11] Bowen Jiang, Jun Li, Gaobo Tan, and Lizhen Xu. 2022. Instrumenting for price elasticity: Machine learning and causal inference in panel data. *Management Science* (2022).

[12] Sajan Kedia, Samyak Jain, and Abhishek Sharma. 2020. Price optimization in fashion e-commerce. *arXiv preprint arXiv:2007.05216* (2020).

[13] Hans Kellerer, Ulrich Pferschy, David Pisinger, Hans Kellerer, Ulrich Pferschy, and David Pisinger. 2004. The multiple-choice knapsack problem. *Knapsack Problems* (2004), 317–347.

[14] Nazlı B Keskin and Assaf Zeevi. 2014. Dynamic Pricing with an Unknown Demand Model: Asymptotically Optimal Semi-Myopic Policies. *Operations Research* 62, 5 (2014), 1142–1167.

[15] Dugang Liu, Xing Tang, Han Gao, Fuyuan Lyu, and Xiuqiang He. 2023. Explicit Feature Interaction-aware Uplift Network for Online Marketing. *arXiv preprint arXiv:2306.00315* (2023).

[16] Xiao Liu. 2023. Dynamic coupon targeting using batch deep reinforcement learning: An application to livestream shopping. *Marketing Science* 42, 4 (2023), 637–658.

[17] Eleanor Loh, Jalaj Khandelwal, Brian Regan, and Duncan A Little. 2022. Promotheus: An End-to-End Machine Learning Framework for Optimizing Markdown in Online Fashion E-commerce. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3447–3457.

[18] Daniel L. Millimet. 2023. Fixed effects models: Theory and practice. *Journal of Economic Surveys* (2023).

[19] Kalyan T. Talluri and Garrett J. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Springer.

[20] Bart Van Parys and Bryan Wilder. 2023. From Predictions to Prescriptions in Non-Linear Learning. *Operations Research* (2023). Forthcoming or Early Access.

[21] Hong Xie, Yongkun Li, and John CS Lui. 2019. Optimizing discount and reputation trade-offs in e-commerce systems: Characterization and online learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 7992–7999.

[22] Jason Shuo Zhang, Benjamin Howson, Panayiota Savva, and Eleanor Loh. 2024. DISCO: An End-to-End Bandit Framework for Personalised Discount Allocation. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 33–49.