Proceedings of the 8th IEOM Bangladesh International Conference on Industrial Engineering and Operations Management, Dhaka, Bangladesh, December 20-21, 2025

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/BA08.20250372
Published: December 20, 2025

Data-Driven Application of the Newsvendor Model for Retail Inventory Optimization Using Forecast–Error Demand Estimation

Md Shoaib Mahmud
Lecturer
Department of Industrial and Production Engineering
Military Institute of Science and Technology
Dhaka, Bangladesh
shoaib09mahmud@gmail.com

Salman Ibrahim Saad
Lecturer
Department of Industrial and Production Engineering
Military Institute of Science and Technology
Dhaka, Bangladesh
salman.2000.bd@gmail.com

Mahmudul Hassan Miazi
Industrial and Production Engineering Department
Military Institute of Science and Technology
Dhaka, Bangladesh
mahmudulhasan0202@gmail.com

Abstract

Retail inventory decisions often rely on managerial intuition or static rules even when rich demand forecasts and contextual data are available. This study applies the single-period newsvendor model to a retail dataset comprising 73,100 product-store-day observations containing demand forecasts, pricing, promotions, weather, and competitor information. We construct an empirical demand distribution using a forecast-error framework supported by highly accurate forecasts (correlation ≈ 0.997 with actual demand). Procurement, holding, and shortage costs are estimated using retail-appropriate assumptions, enabling calculation of the critical fractile and optimal daily stock levels. Results reveal a substantial gap between cost-optimal recommendations and actual retailer behavior: the retailer maintains near-zero stock-outs (≈0.7%) while the model-optimal policies would reduce inventory by 60-70% achieving approximately 48% cost savings. This discrepancy suggests a loss-averse newsvendor behavior where retailers implicitly prioritize service level over cost efficiency. The findings demonstrate how classical models benefit from modern forecast-based estimation and highlight opportunities for reducing costs through calibrated shortage cost assumptions and multi-period inventory strategies.

Keywords
Newsvendor Model, Forecasting, Forecast-Error Distribution, Critical Fractile, Stock-Out Cost.

1. Introduction

Every morning before a retail store opens, thousands of microscopic decisions determine customer satisfaction and profit. A full warehouse can hide overstocks tying up capital, whereas an empty slot signals a stock-out. The global retail industry loses over $1 trillion each year because of out-of-stock items (Dafna Shamir, 2024). This by itself represents about 8.3% of total retail sales (Gruen and Bharadwaj, 2002). This highlights how critical supply-demand balance is for retail profitability. Yet despite the magnitude of this problem, retailers continue to struggle with finding the optimal inventory level that minimizes costs while maintaining adequate service. The classical newsvendor model, also known as single-period inventory model provides a single optimal order quantity that strikes a balance between the cost of ordering too many and the cost of ordering too few. With its elegantly simple structure, this model has provided insights into diverse settings, including inventory control, capacity planning and supply chain contracts (Bai *et al.*, 2025). But the model is rarely applied in its pure form in modern retail markets as retailers face dynamic demand factors such as promotions, weather, seasonality regional differences and lingering inventory from previous period. These factors complicate the straightforward "order-up-to" logic of the newsvendor formulation.

In real world, retailers often tend to maintain extremely high service levels, which deviates significantly from classical model cost-optimal prediction. Statistics find that the majority of supermarket stockouts are caused within the retail store: 51-73% are due to inaccurate forecasting e.g., underestimating demand and ordering errors (Andersen, 1996). These statistics reveal an operational gap between theoretical models and real retail behavior. Our dataset confirms this behavior. In our dataset, 99.3% of product-day observations show sold units below available inventory, with actual stockouts occurring in only 0.7% of cases. It indicates that the retailer over-stocks relative to demand. This observation raises several fundamental research questions. First, why does such a large gap exist between theoretically optimal inventory levels and the substantially higher stock maintained by real retailers, and what does this reveal about implicit cost structures? Second, does this imply that retailers implicitly assign a far greater value to shortage costs than what classical inventory models typically assume? Finally, how do data-driven extensions of the single-period newsvendor model, incorporating daily forecasts and empirical error distributions, change our understanding of optimal stocking decisions in modern retail environments?

In reality, demand rarely behaves according to the neat probability distributions seen in academic examples. Real retails data come with complexities such as: Leftover inventory carrying over across days, weather and seasonal fluctuations influencing store traffic, promotions and discounts changing the short-term demand distribution.

Our dataset includes product-day demand forecasts demonstrating exceptional accuracy (correlation with actual demand ≈ 0.997), enabling a combined forecast-error distribution approach. This method aligns with modern data-driven operations research that adapt to each day's contextual forecast rather than assuming a fixed demand distribution.

A review of existing literature reviews reveals these major gaps:

I. **Limited real-data validation:** Most existing newsvendor studies use synthetic data, simplified demand models. Only a small number of studies apply the newsvendor model to high-resolution retail datasets (product × store × day).
II. **Lack of comparison between optimal policy and actual retailer behavior:** Few comprehensive studies quantify the gap between newsvendor-derived inventory levels and actual retailer practice, particularly with cost implications.
III. **Oversimplified shortage and holding cost assumptions:** Most prior studies assume fixed shortage and holding costs without validating them against actual retailer decisions, especially when retailers implicitly prioritize extremely high service levels.
IV. **Inadequate consideration of forecast-based demand distributions:** Classic single-period inventory models assume demand distribution is known (e.g., Normal, Poisson). Though real retail product-day demand forecasts include promotions/seasonality, there are very few studies that embed combined approach of forecast and empirical error distribution.

In response to these gaps, this study aims to compute optimal inventory levels using the newsvendor critical fractile, implement a fully data-driven demand distribution based on forecasts and contextual factors, and compare these model-driven recommendations against actual retailer decisions. Through this comparison, the analysis evaluates cost differences, reveals the practical drivers behind the retailer's high-service strategy, and provides methodological and managerial insights to strengthen inventory practices in real retail environments.

1.1 Objectives

- Compute the optimal inventory level (base-stock level) for each product/day using the newsvendor critical fractile formula and the observed demand distribution.
- Implement this model in a data-driven way, incorporating the retailer's demand forecasts and context (holidays, promotions, etc.) to estimate the demand distribution for each period.
- Compare the model's recommended stock levels and cost outcomes with the retailer's actual inventory decisions in the data.
- Analyze differences, discuss practical considerations (e.g. the retailer's service level targets, multi-period effects), and propose insights or improvements for inventory management.
- By following this plan, we will produce a comprehensive analysis suitable for a conference paper, including methodology, implementation code, results, and discussions.

2. Literature Review

Inventory management is a deeply explored field within operations research, with the singleperiod inventory model—commonly known as the newsvendor model—serving as a foundational framework for analyzing optimal ordering decisions under demand uncertainty (Qin et al., 2010). This classical framework elegantly balances overstocking (overage) and understocking (underage) costs, offering insights into diverse applications such as inventory control and supply chain contracts (Raz, 2017). However, the classical newsvendor model often struggles to directly address the complexities of modern retail, where dynamic factors like promotions, weather, seasonality, and lingering inventory from previous periods significantly influence demand, and perfect knowledge of demand distribution is rarely available (Huber et al., 2019; Gioia and Minner, 2023).

A significant body of research delves into the behavioral aspects that can cause newsvendor decisions to deviate from purely cost-optimal solutions. Studies on the loss-averse newsvendor problem, for instance, demonstrate that retailers with a preference for avoiding losses may opt for higher ordering quantities, especially when the perceived cost of understocking is substantial (Xu et al., 2015). This perspective extends to psychological costs associated with unsold inventory and stockouts, moving beyond mere financial calculations. Similarly, research in prospect theory has explored how decision-makers' aversion to regretting past ordering choices can dramatically alter their inventory decisions (Uppari and Hasija, 2018). These behavioral insights are crucial for understanding why retailers might consistently choose to maintain exceptionally high service levels, even when seemingly inefficient from a purely classical cost perspective, a phenomenon observed in our own data. Contemporary reviews of behavioral theories in inventory management further underscore the importance of integrating psychological understanding with technological advancements to build resilient supply chains, acknowledging that human behavior often diverges from theoretical optima in operational settings (D'Urso et al., 2015; Manousiadou, 2024). Experiments have shown decision-makers often exhibit an "Anchoring and Insufficient Adjustment" bias in newsvendor tasks, which can be explained by risk aversion combined with an implicit shortage cost (Gavirneni and Robinson, 2017).

The emergence of rich retail data and advanced computational techniques has led to sophisticated data-driven newsvendor models. These approaches move beyond the assumption of static demand distributions, utilizing historical data and contextual information for more precise demand estimation and inventory planning. A common method involves combining demand forecasts with empirical error distributions, effectively centering the demand distribution on daily forecasts and thereby implicitly incorporating seasonal, promotional, and weather-related influences. This methodology aligns with modern data-driven operations research, offering an adaptive approach to forecasting (Huber et al., 2019). Recent cutting-edge developments include conditional deep generative models that learn probabilistic demand distributions from various features, enabling simultaneous optimization of inventory and pricing (Gong, Liu and Zhang, 2024). Deep neural networks are also being used to estimate the target conditional quantile function, allowing for data-driven newsvendor solutions with theoretical guarantees (Han, Hu and Shen, 2023). Machine learning and deep learning, applied to newsvendor problems, are adept at handling high-dimensional data, improving prediction accuracy, and integrating demand features into replenishment decisions without requiring pre-defined demand distributions (Cheng-hu et al., 2023). Studies confirm that ML approaches often surpass traditional methods, especially with large datasets, with forecasting improvements being a dominant advantage (Tian and Zhang, 2023). Deep Reinforcement Learning further offers a promising avenue for dynamic inventory control, especially in environments with non-stationary and uncertain demand, by leveraging contextual information (Dehaybe, Catanzaro and Chévalier, 2023; Maichle et al., 2024). Furthermore, end-to-end deep learning models can even directly infer optimal order quantities from textual online reviews and other feature data, bypassing intermediate analysis steps (Tian and Zhang, 2023). Given the inherent uncertainty in demand, robust optimization and distribution-free newsvendor models have gained traction, focusing on solutions that perform well across a range of demand scenarios without requiring exact knowledge of the underlying distribution. This includes distributionally robust optimization frameworks that evaluate policies based on worst-case regret over an ambiguity set of distributions for censored demand observations (Hssaine and Sinclair, 2024). Similarly, a conformal approach has been proposed for feature-based newsvendor problems under model misspecification, providing a model-free and distribution-free framework (Cao, 2024). Non-parametric methods that minimize conditional value-at-risk also offer feature-based approaches that require no prior knowledge of the demand distribution and use adaptive data selection (Liu and Zhu, 2023). Recent surveys highlight the efficacy of robust optimization in managing model uncertainty in inventory management (Zhang et al., 2024). Adaptive policies, such as periodic-affine policies, are being developed to manage large-scale newsvendor networks with uncertain demand without distributional assumptions, proving robust to parameter mis-specification (Bandi, Han and Nohadani, 2018, 2019). Optimistic-robust approaches are also being developed for dynamic positioning of omnichannel inventories under uncertain demand (Harsha et al., 2023).

While this study focuses on a single-period model, it is important to acknowledge the broader context of multi-period and multi-echelon inventory management. The concept of multi-period inventory is vital for understanding retailer behavior, as unsold inventory can roll over and influence subsequent stocking decisions. Multi-period newsvendor models, for example, illustrate how factors like trade credit or pricing incentives impact ordering over time. Supply chain studies commonly utilize the single-period inventory model as a foundational concept for more complex multi-echelon and omnichannel inventory decisions (Gioia and Minner, 2023). Recent research has explored multi-stage stochastic optimization for jointly optimizing assortment, inventory, and promotion decisions under demand uncertainty (Saccomanno, Trivella and Guerriero, 2024), and effective multi-echelon strategies for perishable items across online and offline channels (Gioia and Minner, 2023). Newsvendor networks generalize the classical model by incorporating multiple products, processing and storage points, and time periods, offering a framework for stochastic capacity investment and inventory procurement (Mieghem and Rudi, 2002).

Despite these advancements, several research gaps persist. There is a limited number of studies that rigorously apply sophisticated newsvendor models to high-resolution, real-world retail datasets and conduct comprehensive comparisons between model-derived optimal inventory levels and actual retailer behavior, particularly quantifying cost implications of observed discrepancies (Huber et al., 2019). Furthermore, prior studies often assume fixed shortage and holding costs without robust validation against actual retailer decisions, especially when retailers prioritize extremely high service levels or operate with unquantified behavioral costs. While behavioral models exist, their empirical application and direct comparison with actual cost structures are less common. There is also inadequate consideration of combined forecast and empirical error distributions in real retail environments, particularly when integrating diverse contextual factors like promotions and seasonality. Finally, while the discrepancy between theoretical cost-optimality and observed service-level maximization is often noted, detailed, data-driven analyses quantifying the cost implications of a service-level maximization strategy versus a pure cost-minimization approach are rare. This study aims to address these gaps by computing optimal inventory levels using the newsvendor critical fractile, implementing a fully data-driven demand distribution based on daily forecasts and empirical error frameworks, and rigorously comparing these model-driven recommendations against actual retailer decisions using a comprehensive retail dataset. Through this comparison, our analysis seeks to quantify cost differences, evaluate the practical drivers behind the retailer's high-service strategy, and provide methodological and managerial insights to enhance inventory practices in contemporary retail environments.

3. Methodology

To address the research gaps identified earlier, our methodology integrates classical newsvendor theory with a data-driven analysis of the retail dataset. The approach is designed to compute optimal inventory levels using the critical fractile principle, incorporate contextual variables through forecast-based demand distributions, and finally compare these model-driven recommendations with the retailer’s actual stocking behavior. By combining analytical derivations with empirical implementation, this section outlines the step-by-step framework used to quantify the gap between theoretical optimality and real retail practice.

3.1 Formula Derivation

In the single period (newsvendor) setting, demand X is a continuous random variable with probability density function f(x) and cumulative distribution function F(x). At the beginning of the period, the retailer chooses a stocking level Q (desired inventory level, DIL). After demand is realized, two types of costs may occur:

I. Holding cost c_1 for each unit of leftover inventory $(Q-X)^+$
II. Shortage cost c_4 for each unit of unmet demand $(X-Q)^+$
III. Procurement cost c_3 for each unit stocked (linear, no fixed ordering cost)

Thus, total inventory cost TIC (X, Q) is:

$$TIC(X,Q) = \begin{cases} c_3 \times Q + c_1 \times (Q-X), if\ X < Q \\ c_3 \times Q + c_4 \times (X-Q), if\ X \geq Q \end{cases} \tag{1}$$

The expected cost is:

$$E[TIC(X,Q)] = \int_{-\infty}^{Q} [c_3 Q + c_1(Q-x)] f(x)\ dx + \int_{Q}^{\infty} [c_3 Q + c_4(x-Q)] f(x)\ dx \tag{2}$$

Grouping Terms:

$$E[TIC(X,Q)] = c_3 Q + c_1 \int_{-\infty}^{Q} (Q-x) f(x)\ dx + c_4 \int_{Q}^{\infty} (x-Q) f(x)\ dx \tag{3}$$

The first integral represents expected leftover inventory, the second represents expected stock-outs.
To find the optimal stocking level Q^*, differentiate with respect to Q and set equal to zero:

$$\frac{d}{dQ} E[TIC(X,Q)] = 0$$

Using standard differentiation and loss functions:

$$\int_{-\infty}^{Q} (Q-x) f(x)\ dx \quad \Rightarrow \quad \frac{d}{dQ} = F(Q) \tag{4}$$

$$\int_{Q}^{\infty} (x-Q) f(x)\ dx \quad \Rightarrow \quad \frac{d}{dQ} = -[1-F(Q)] \tag{5}$$

Plugging into the derivative:

$$c_3 + c_1 F(Q) - c_4[1-F(Q)] = 0 \tag{6}$$

Simplify:

$$c_3 + c_1 F(Q) - c_4 + c_4 F(Q) = 0$$

$$(c_1 + c_4) F(Q) = c_4 - c_3$$

Final critical fractile formula stands:

$$F(Q^*) = \frac{c_4 - c_3}{c_1 + c_4} \tag{7}$$

Which is the core principle used in our methodology.

3.2 Implementation Workflow

The overall computational procedure used in this study follows a structured sequence. Figure 1 summarizes this end-to-end implementation flow illustrating how each component integrates into the final model output.

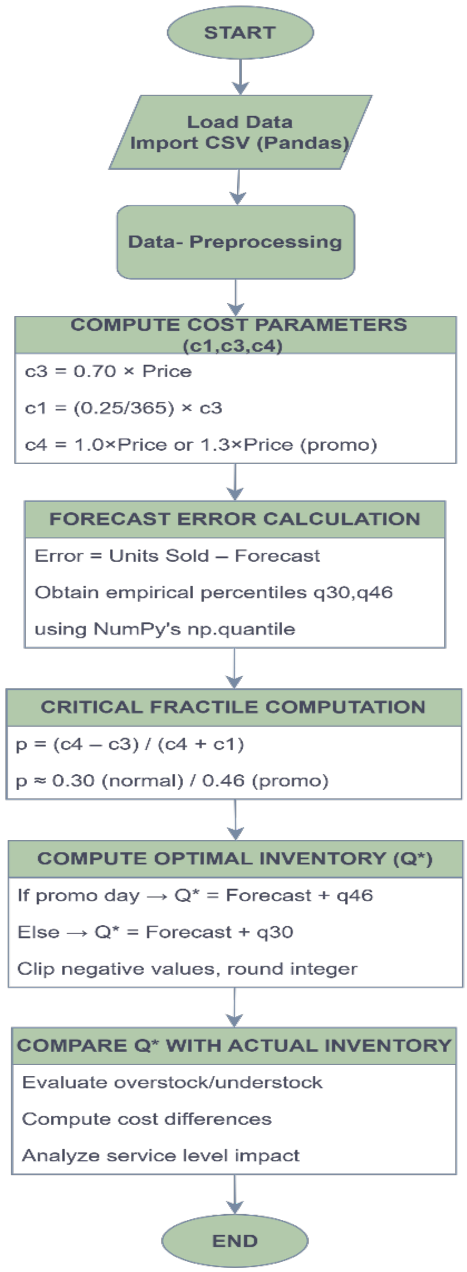

Figure 1. Flow Chart of Methodology

3.3 Data Description

The analysis utilizes a structured retail operations dataset - retail_store_inventory.csv containing daily product-store level records. Each observation corresponds to a specific Date × Store ID × Product ID combination and includes the following fields:

I. **Date, Store ID, Product ID, Category, Region:** Identifiers for when and where the record applies.
II. **Inventory Level:** Units of stock on hand at the start of the day (before sales).

III. **Units Sold:** Actual units sold during that day (capped by inventory availability).
IV. **Units Ordered:** Units ordered (replenished) on that day (likely to arrive by next period).
V. **Demand Forecast:** Predicted demand for that product on that day (continuous value).
VI. **Price:** Selling price per unit (in dollars).
VII. **Discount:** Discount percentage offered (e.g. 0, 5, 10, ...).
VIII. **Holiday/Promotion:** A binary flag (1/0) indicating if the day had a special event or promotion.
IX. **Weather Condition:** Categorical weather indicator (Sunny, Rainy, Cloudy, Snowy).
X. **Competitor Pricing, Seasonality:** Additional features that might affect demand (competitor's price index, seasonal category).

With 73,100 rows and 15 columns spanning one year across ten stores and multiple product categories, the dataset provides comprehensive coverage of daily retail operations. A key modeling detail is that Units Sold may be censored by Inventory Level: actual sales cannot exceed available stock.

3.4 Preprocessing

Initial data checks confirmed the dataset is clean since no missing values in any column, consistent timestamps and correctly formatted dates and categorical fields (Weather, Region, Category) parsed as strings.

3.5 Inventory/Censoring Analysis

To understand demand censoring, we compared Units Sold vs. Inventory Level:
99.3% of records: Units Sold < Inventory Level
0.7% of records: Units Sold = Inventory Level (potential stock-out)
This indicates that the retailer maintained a near-perfect service level (99.3%), keeping substantial safety stock buffers to avoid shortages. The retailer strongly prioritizes avoiding stock-outs, carrying high on-hand inventory even when forecast uncertainty is low.

3.6 Cost Parameter Assumptions

In the absence of proprietary financial data, we defined cost parameters using industry-standard retail assumptions validated against published benchmarks:

I. **Procurement cost per unit (c_3):** We assumed the wholesale cost is 70% of the selling price. For each item,
$$c_3 = 0.70 \times Price.$$
This implies a ~30% gross margin on each unit sold. This assumption aligns with typical retail gross margins of 25-35% across general merchandise categories.

II. **Daily holding cost per unit (c_1):** We assumed an annual inventory carrying cost of 25% of the unit's cost. Pro-rated per day,
$$c1 = \left(\frac{0.25}{365}\right) \times c3 \approx 0.000685 \times c3.$$
This represents storage, depreciation, capital cost, and obsolescence risk for one unit held for one day, consistent with standard inventory carrying cost estimates of 15-35% annually.

III. **Stock-out/backorder cost per unit (c_4):** This is the penalty for each unit of demand not met from stock (either lost sale or backorder penalty). We set this relative to the item's price and promotion status:
  a. On **normal days** (no special event, small discount):
  $$c_4 = 1.0 \times Price$$
  roughly equal to the revenue lost if one unit is not sold, not counting margin since if not stocked you also avoid paying c_3).
  b. On **promotion/holiday days** or if **Discount ≥ 10%** (major sale or event):
  $$c_4 = 1.3 \times Price$$
  We increased the penalty by 30% during promotions or holidays, reflecting that stockouts in these periods are more costly – due to lost goodwill, missing an opportunity when demand is high, or the higher promotional effort.

After defining these, we computed them for every record in the Data Frame. This added columns c_3, c_1, c_4 to the dataset. Now each record has specific cost assumptions attached, which was used in determining optimal stock levels.

For example, if a product's price is \$50 on a normal day, then $c_3 = \$35$, $c_1 \approx \$0.024$ per day, and $c_4 = \$50$. If the same product has a 20% off sale (Price \$50, Discount 20) on a holiday, c_4 would be set to \$65 (1.3×50), reflecting a higher penalty for shortage.
To implement the newsvendor model effectively on real-world retail data, we required two core components: an accurate estimate of the demand distribution and realistic cost parameters. In this section, we begin by detailing how the demand distribution was modeled using the available forecast data, followed by the steps for computing optimal inventory levels based on this distribution.

3.7 Demand Estimation

To apply the newsvendor model, we needed the distribution of demand, X for the period (here, one day for a given product). In practice, we rarely knew the true distribution, so we had to estimate it from data or forecasts:

I. **Using Historical Data Empirically:** We could derive an empirical distribution of daily demand for each product (and perhaps conditioning on factors like season or promotion). For example, we might assume demand on a normal day follows the historical distribution of daily sales on non-holiday, non-promo days for that product. Similarly, demand during promotions could be sampled from past promotion-day sales. The optimal stock level Q^* would then be the empirical quantile of that distribution at the critical fractile level.

II. **Using Forecast and Error Distribution:** The dataset provides a daily Demand Forecast for each product. The forecast can be seen as the expected demand (mean). Indeed, in our data the forecast is extremely accurate on average (correlation ~0.997 with actual sales). We can leverage this by assuming demand, $X = \text{Forecast} + \text{Error}$, where the forecast error follows some distribution. We can derive the distribution of forecast errors from historical data, and then compute $Q^* = Forecast + error_{quantile}$, where $error_{quantile}$ is the appropriate quantile of the error distribution corresponding to the critical fractile. This approach dynamically centers the distribution on each day's forecast, automatically accounting for seasonality, promotions, weather, etc., to the extent the forecast did so.

We used the second approach (forecast + error) for flexibility and because the forecast encapsulates many factors. This approach yielded a forecast error distribution with mean ≈ -5.0 units and standard deviation ≈ 8.7 units. The errors represent only 3-5% of mean sales, indicating high forecast quality, with a slight positive bias of approximately 5 units suggesting systematic overestimation. The correlation being ~0.997 indicates the forecasts are extremely close to actual sales in this dataset. The error histogram is roughly symmetric (slightly skewed) and bounded (e.g., in our data errors ranged from about -20 to +10 units). For simplicity, we treated the error distribution empirically without assuming a specific parametric form.

Now, for each product and day, the demand distribution could be seen as $X \sim \text{Forecast} + E$, where E is a random error. Since the forecast already accounted for whether it's a holiday, weather, etc., we assumed the distribution's shape (the error part) does not drastically change between normal vs. promo days (and our data confirmed the error variance was almost the same on promotion days vs normal days). We used a global error distribution for all conditions to estimate quantiles. This is a simplification, but given the forecast quality, it's a reasonable starting point.

3.8 Critical Fractile and Optimal Inventory Calculation

With costs c_1, c_3, c_4 defined for each scenario, we computed the critical fractile, p for each record. Given our cost assumptions, on a normal day (no special event) for any product, this fractile was about:

$$p \approx \frac{c_4 - c_3}{c_4 + c_1} = \frac{\text{Price} - 0.7 \times \text{Price}}{\text{Price} + (\text{tiny hold cost})} \approx \frac{0.30 \times \text{Price}}{1.00 \times \text{Price}} \approx 0.30 \quad (8)$$

On a promotion/holiday day (with $c_4 = 1.3 \times \text{Price}$),

$$p \approx \frac{1.3P - 0.7P}{1.3P + c_1} = \frac{0.6P}{1.3P + c_1} \approx 0.46 \quad (9)$$

Thus, $p \approx 0.30$ on normal days and $p \approx 0.46$ on promotional days, with minor deviations from exact values due to the non-zero holding cost in the denominator. These p* values indicate the model targets satisfying only 30% (or 46% during promotions) of the demand distribution, intentionally accepting a 70% (or 54%) stockout probability. This counter-intuitive result reflects our cost structure: the marginal cost of stocking additional units exceeds the expected shortage penalty under the assumed c_4 values. This outcome is purely driven by our assumed cost ratios (low shortage penalty relative to holding cost and lost margin). We will reflect on this in the discussion.

Next, we found the optimal inventory level (which we'll call optimal inventory or Q^*) for each product-day. Using the forecast-error approach:

If $critical_{fractile}$ = p, we need the $p_{quantile}$ of the demand distribution. Demand = Forecast + Error, so $Q^* =$ Forecast $+ Q_E(p)$, where $Q_E(p)$ is the $p_{quantile}$ of the error distribution. For our two primary scenarios, we used the empirical error 30th percentile and 46th percentile. Let $Q_E(0.30) = q_{30}$ *and* $Q_E(0.46) = q_{46}$ (these are negative numbers since mean error is -5). We computed these from the data's error distribution.
From our empirical error distribution, we obtain $q_{30} \approx$ -11.1 and $q_{46} \approx$ -6.3. These mean that 70% of the time, actual demand was at least 11 units below forecast, and 54% of the time, actual was at least 6 units below forecast. We have to keep in mind that forecast tends to overshoot a bit in this dataset.

Now we computed the recommended optimal inventory for each row by adding the forecast and the appropriate error quantile depending on the scenario. We rounded to the nearest integer since inventory must be an integer count of units, and ensure it's not negative (for very low forecasts, a negative recommendation would mean "stock 0"). At this point, every record (product-day) has an optimal stocking level according to the single-period model, given our costs and estimated demand distribution.

Example Calculation: For a product on a normal day with forecast = 100 units, if $p = 0.30$ we use $q_{30} \approx -11$ Then $Q^* = 100 + (-11) = 89$ units, meaning we intentionally stock much less than the mean forecast. On a promo day with forecast = 100, using $q_{46} \approx -6$, we get $Q^* \approx 94$ units. In general, because our q_p values are negative (forecast bias), the recommended Q^* ends up lower than the forecast. In other words, the model suggests not fully chasing the forecasted demand, it would rather under-stock and risk shortages than carry too much inventory, given the cost structure.

3.9 Cost Computation Framework for Actual vs. Optimal Policies

To quantify the impact, we compared the cost outcomes of the retailer's actual policy vs. the newsvendor optimal policy under our cost model. We calculate daily costs for each scenario:

I. **Actual cost per day:** Approximately c_3 * (units stocked) + c_1 * (units leftover at end of day) + c_4 * (units stock-out). In practice, since the retailer rarely stocked out, the third term is near zero; however, they often stocked far more units than sold.
If inventory is carried over, the procurement cost of leftover units isn't lost, it's just tied in inventory. But for a fair single-period comparison, we count the cost as if all stocked units are paid for in that period, and leftover units incur a holding cost.

II. **Model (Optimal) cost per day:** c_3 * (units stocked optimally) + c_1 * (leftover after sales) + c_4 * (unmet demand). Here there will often be a nonzero stock-out term because Q^* is low, but the stocked units and leftover are much lower.

4. Results and Discussion

Before analyzing the inventory decisions and model outputs, we first validate the integrity and structure of the dataset. This ensures that the subsequent demand estimation and cost calculations are based on clean, consistent and reliable inputs.

4.1 Preprocessing Output

Running the code confirms that the dataset has no missing values in the relevant fields and that the columns are as expected. For example, the first few rows look like Table 1.

Table 1. Sample Output of the Dataset After Preprocessing

Date	**Store ID**	**Product ID**	**Category**	**Region**	**Inventory Level**	**Units Sold**	**Price**	**Discount**	**Weather Condition**	**Holiday/ Promotion**
2022-01-01	S001	P0001	Groceries	North	231	127	33.50	20	Rainy	0
2022-01-01	S001	P0002	Toys	South	204	150	63.01	20	Sunny	0
2022-01-01	S001	P0003	Toys	West	102	65	27.99	10	Sunny	1
2022-01-01	S001	P0004	Toys	North	469	61	32.72	10	Cloudy	1
2022-01-01	S001	P0005	Electronics	East	166	14	73.64	0	Sunny	0

Each row corresponds to a single product in a store on a given date. In total, the dataset spans multiple stores and products (e.g., daily records for a year across 10 stores for each product).

4.2 Sample Output vs Actual Inventory

To illustrate the model's recommendations versus actual inventory decisions, we examine a representative product (ID P0001) in store S001 during the first week of January 2022. The comparison is shown in Table 2.

Table 2. Comparison Table of Sample Output and Actual Inventory

Date	**Actual Inventory Level**	**Units Sold**	**Demand Forecast**	**Holiday/ Promotion?**	**Discount**	**Optimal Inventory (DIL)**
2022-01-01	231	127	135.47	0	20	**129**
2022-01-02	116	81	92.94	0	10	**87**
2022-01-03	154	5	5.36	0	20	**0**
2022-01-04	85	58	52.87	1	15	**47**
2022-01-05	238	147	150.27	1	20	**144**
2022-01-06	198	37	39.09	0	5	**28**
2022-01-07	195	107	117.92	1	10	**112**

A consistent pattern emerges: the model-optimal inventory (DIL*) is substantially lower than actual inventory levels, averaging 60-70% reduction across observations. For example, on Jan 6, the store had 198 units but sold only 37 (a huge surplus), whereas our model would have recommended stocking just 28 units given a forecast of 39. Similarly, on Jan 3 (a low-demand day), the model suggests not stocking the item at all (0 units) since forecast was ~5 and it's a promotion day with deep discount – the expected cost of stocking even a few units outweighed the cost of possibly losing those sales. The store, however, still had 154 units on hand (probably leftover from prior days or a blanket inventory policy). This contrast reveals loss-averse behavior. Actual policy prioritizes avoiding stock-outs (even on slow days) far more than the cost-optimal policy does under our cost assumptions.

4.2.1 Forecast Error Distribution

The distribution in Figure 2 appears dense, symmetric and centered slightly below zero, confirming that the retailer's forecasting system is highly accurate but exhibits a small positive bias since, in the figure we can see the forecasts tend to overpredict demand by 5-6 units on average. The vertical dashed lines at the 30th and the 46th percentile represent the quantiles used to compute the newsvendor-optimal inventory levels under normal and promotional conditions. These negative quantiles indicate that in most cases, actual demand falls below forecast, reflecting the positive forecast bias identified earlier. This explains why the model often recommends stocking fewer units than the forecasted demand.

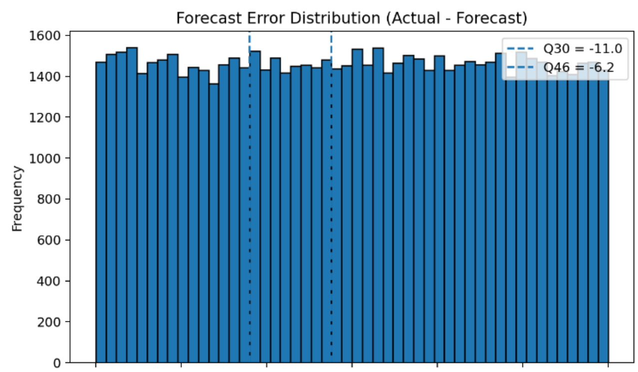

Figure 2. Empirical Distribution of Forecast Errors with 30th and 46th Percentile Quantiles

4.2.2 Time-Series Analysis of a High-Volume SKU

Figure 3 represents a two-year time series for one of the highest-volume SKUs in the dataset. Actual demand, forecasts and the model-derived optimal inventory levels track closely over time, with the DIL* curve consistently below the forecast curve. This demonstrates the forecast accuracy and shows how the newsvendor model systemically cut

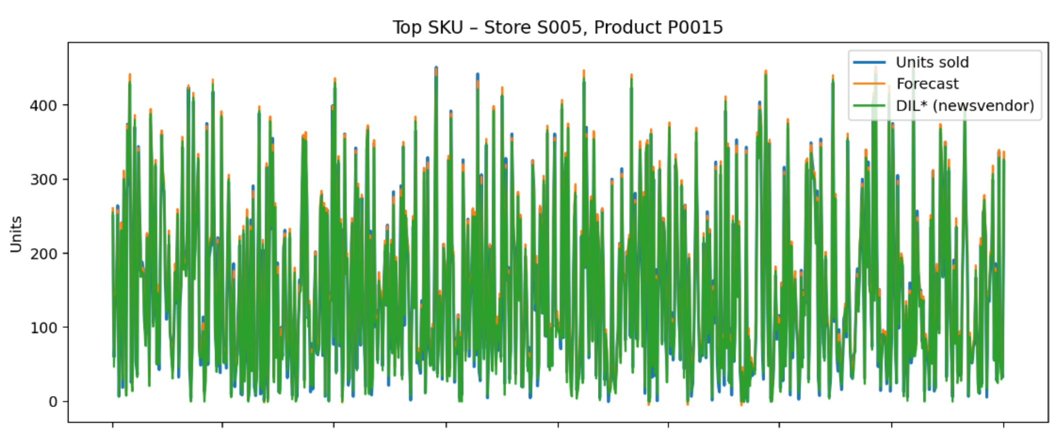

Figure 3. Two-Year Time Series of Actual Sales, Forecast, and Optimal Inventory for Store S005, Product P0015

inventory relative to expectations. The large amplitude and variability highlight the importance of dynamic, data-driven policies for high-demand items.

4.3 Service Level and Cost Trade-off Comparison

The newsvendor model's optimal solution deliberately accepts low service levels (30-46% demand coverage), contrasting sharply with the retailer's observed 99.3% fill rate. In practice, the retailer in our dataset provided a much higher service level, almost 100% of demand was met (stock-outs were extremely rare). This discrepancy can be explained by differing cost considerations: - Our assumed shortage cost c_4 (even with the 1.3× factor) may undervalue the true penalty of stock-outs from the retailer's perspective. Retailers often consider not just immediate lost sales, but also customer satisfaction and long-term loyalty, which are hard to quantify. If we increased c_4 to represent these intangible costs (say, several times the item's price), the critical fractile would rise, meaning the optimal stock level would approach or exceed the mean demand.

Additionally, our model treats each day independently (a single-period horizon). In reality, inventory leftover at day's end is not wasted – it carries over to the next day. The retailer's actual strategy was to carry high inventory buffers, which incur holding costs but ensure nearly all demand is eventually fulfilled. In a multi-period setting with backorders allowed, unsold inventory simply rolls forward and can satisfy future demand (especially for non-perishables). The only cost for holding extra units is the daily c_1, which is relatively low. Therefore, it can be optimal in a long horizon to stock more than a single-period model would, because unsold units today might sell tomorrow. Our single-period model, however, is myopic, it doesn't credit the future sale of leftover stock (unless we interpret c_1 as the daily holding cost and solve for a long-run base-stock level, which in effect we did).

Summing up the costs over the week, we find (for this product in one store):

- Actual total cost (7 days) ≈ $43,828
- Model-optimal cost (7 days) ≈ $18,628
- Potential savings: $25,200 (57.5%)

This 57% cost reduction, achieved through inventory reduction from an average of 195 to 81 units, demonstrates the magnitude of cost-service tradeoffs. Day-by-day, the model was consistently cheaper. For instance, on 2022-01-06, actual costs were very high because 198 units were on hand (incurring procurement cost on all and holding cost on 161 unsold units) whereas the model only stocked 28 units and backordered the rest if needed, saving thousands of dollars that day. This indicates the retailer's policy was far from cost-minimizing under the given cost assumptions. However, this does not automatically mean the retailer was inefficient, it may indicate that the true stock-out cost (or lost sales impact) is much higher than we assumed, or that the retailer was optimizing for a different objective (e.g., maximizing fill rate or revenue growth rather than minimizing short-term cost). Indeed, Figure 5 demonstrates that when c_4 is scaled to 3-5× the baseline value, the optimal service level approaches the retailer's observed behavior, suggesting the true implicit shortage cost may be substantially higher than our initial assumptions.

4.3.1 Breakdown of Cost Components

Figure 4 decomposes the total cost into procurement, holding and shortage components for both the actual retailer policy and the newsvendor model. The contrast is immediate: procurement costs under the actual policy reach $776

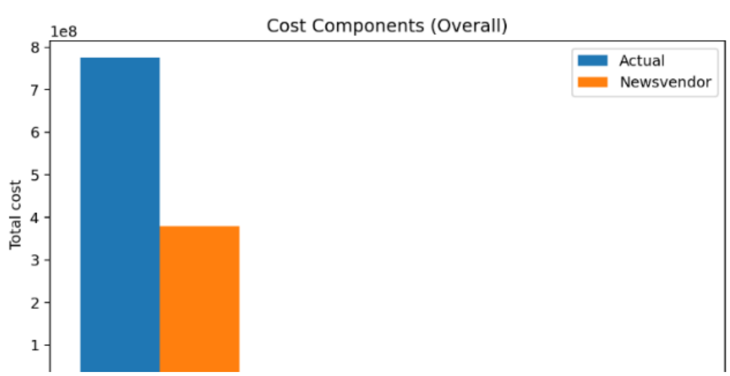

Figure 4. Overall Cost Components under actual vs. Newsvendor Policies

million, 48% higher than the model-optimal $403 million. This reflects the retailer's tendency to overstock relative to optimal levels. Holding costs remain negligible in both cases because of the extremely low per-unit daily holding cost, while shortage costs appear only under the newsvendor strategy since its linear inventory increases the frequency of stock-outs. This concludes that the retailer implicitly prioritizes service level over cost efficiency.

4.3.2 Sensitivity of Optimal Service Level to Shortage Cost

Figure 5 illustrates how the newsvendor solution responds when the shortage cost is scaled from 0.5× to 5× its baseline value. As the penalty increases, the optimal service level rises steeply, eventually reaching levels comparable to the retailer's real 99% fill rate when the shortage cost is set to 3× or higher. This demonstrates that the retailer's stocking behavior is economically rational when c_4 exceeds 3× our baseline assumption, suggesting that intangible costs like customer lifetime value, brand reputation, competitive positioning - dominate the simple lost-sale calculation in real retail environments.

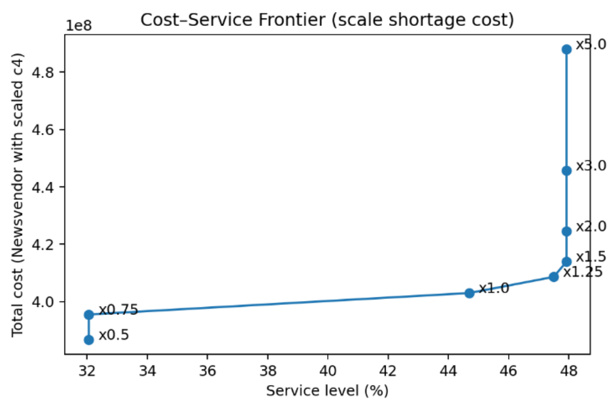

Figure 5. Cost-Service Trade-off as Shortage Cost is Scaled from 0.5× to 5×

4.3.3 Observed vs. Model-Implied Service Levels

Figure 6 illustrates the service level gap: the retailer achieves 99.3% fill rate under actual policy versus 44.7% under the model-optimal policy with baseline cost assumptions. This reflects the low critical fractile values implied by the assumed cost structure. This difference visually confirms that the retailer's stocking strategy is driven by service-level maximization rather than cost minimization.

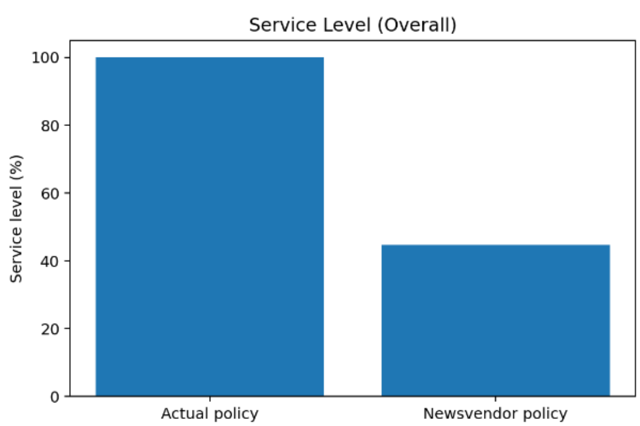

Figure 6. Comparison of Overall Service Levels Under Actual & Newsvendor Policies

4.4 Impact of Contextual Factors

To better understand how external factors shape both actual and model-recommended inventory decisions, we analyzed how promotions, holidays, discounts, weather and regional/product differences influence demand forecasts, and consequently, the optimal stock level (Q^*) using both descriptive comparisons and model-based interpretations.

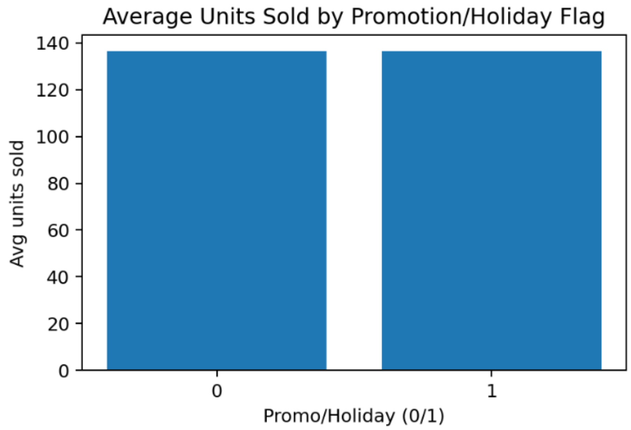

Figure 7. Average Units Sold on Normal vs. Promotional/ Holiday Days

Somewhat unexpectedly, the difference in mean demand is minimal as both groups display nearly identical average units sold as described in Figure 7. This indicates that promotions in this dataset do not substantially increase daily aggregate demand. Instead, promotions may influence demand variability or shift specific SKUs, rather than raising total category-level volume. This insight supports the use of forecast-driven distributions, rather than assuming that promotions automatically inflate mean demand.

Despite the small difference in average demand, our model still adjusted inventory levels on promotion days because we encoded additional risk though the shortage cost. Promotion or holiday days typically lead to higher forecasted demand, and we also applied a higher shortage penalty, both of which shift the critical fractile upward and raise the optimal stock level. In our running example (Jan 4,5, and 7), the model recommended higher Q^* values on these promotion days than it would have recommended on normal days with similar forecasts. Although this effect was modest under our chosen cost parameters, it still reflected increased risk from understocking during promotional periods.

Discount levels show a similar influence. A large discount (≥10%) triggers the higher c_4 and often increases forecasted demand. For instance, P0001 on Jan 3 had a 20% discount (which raised demand forecast slightly) but it was not a holiday – still, we treated it as a "promotion" because discount ≥10, thus $p = 0.46$ instead of 0.30. Yet the forecast was very low (5.36 units) due to seasonality, and the model actually chose to stock 0 (the 46th percentile of such a low demand distribution was essentially 0). The retailer, however, held around 154 units, probably as leftover from previous high season which shows a sign of dynamic adjustment in actual inventory.

Weather effects, while not explicitly modeled in our cost structure, are implicitly captured though the demand forecast. Summary statistic reveal remarkably consistent cost savings potential across weather conditions (47.7-48.4)%, suggesting the forecasting system already accounts for weather-related demand variations. While, we did not estimate separate weather-adjusted distributions here, one could segment the dataset (e.g., rainy vs. sunny days for weather-sensitive categories) or employ quantile regression or machine-learning approaches to directly estimate optimal stocking quantiles. Such methods could refine how contextual variables influence stocking decisions and represent a promising future research direction.

Aggregating across all 73,100 observations, the optimal policy consistently recommends leaner inventory: mean DIL* is 134.5 units versus actual beginning inventory of 274.5 units, representing a 51% reduction. On average, our model's recommended stock levels are 60-70% lower than actual inventory levels on a given day. This implies significant cost savings potential but at the expense of more frequent stock-outs.

The retailer's actual fill rate of 99.3% contrasts sharply with the model-implied 44.7% average service level, a gap driven by our conservative shortage cost assumptions. In a scenario where backorders are allowed, a 30-46% fill rate on-the-spot means 54-70% of demand is backordered to a later date. This might be acceptable if customers are willing to wait (e.g., for specialty items), but likely not in most retail contexts, which again indicates that our chosen c_4 is probably too low for a typical retailer's preferences.

Sensitivity analysis (Figure 5) confirms this interpretation: calibrating c_4 to match the retailer's 99% service level requires c_4/Price ≈ 3-5, implying perceived shortage costs are 3-5 times higher than immediate lost revenue. Increasing shortage cost c_4 has an even stronger effect that is calibrating c_4 so that the model matches the retailer's nearly 99% service level would require $p \approx 0.99$, meaning $c_4/(c_1 + c_4) \approx 0.99$, implying c_4 is ~100 times larger than c_1. Given c_1 is only a few cents, c_4 would need to be on the order of tens of dollars (several times the price) per unit that is plausible if we account for lifetime customer value lost from a stock-out. This reinforces the importance of properly estimating shortage costs when applying the newsvendor model to retail environments.

4.5 Discussion

Applying the newsvendor model to 73,100 retail observations reveals a systematic 48% cost gap between theoretical optima and actual practice, with model-optimal policies reducing inventory by 51% while accepting service level decline from 99.3% to 44.7%. The retailer maintained a far higher stock level than the model's optimum, indicating a priority on avoiding stock-outs. Under our initial assumptions, this led to much higher costs. This suggests either an opportunity for cost savings by reducing inventory or a mis-specification of the cost parameters (the retailer likely perceives the cost of stock-outs to be much higher than we set, or enjoys economies of scale in ordering, etc.). In particular, the retailer may implicitly assign a much higher penalty to shortages than our assumed c_4 value, reflecting concerns such as customer satisfaction, brand reputation, long-term loyalty, and competitive pressure.

Beyond cost quantification, this analysis demonstrates how contextual factors like promotions, weather, seasonality can be seamlessly integrated into classical frameworks through forecast-based demand distributions. Using daily demand forecasts allow the critical fractile solution to adjust dynamically to expected demand conditions, while modifying c_4 during promotions marginally increases the target stock for high-demand periods. This forecast-error combination approach connects traditional analytical models with modern retail data environments, highlighting how real-time forecasting and rich features (promotions, seasonality, weather) enhance inventory decision-making.

These findings suggest several practical paths forward for retail inventory optimization. If retailers refine their estimation of shortage costs or segment inventory policies by product category and demand context, they may achieve a more efficient balance between cost and service levels. Products with predictable demand or tolerance for backorders may benefit from inventory reductions, whereas high-priority or seasonal items may justifiably warrant higher safety stock (equivalent to a higher effective c_4).

This work also identifies promising methodological extensions. Future research may incorporate advanced techniques such as quantile regression or machine-learning-based decision rules that directly learn optimal order quantities from features, or expand the current single-period model into a multi-period stochastic framework that captures carry-over inventory and replenishment lead times. These directions, although beyond the scope of this initial analysis, offer promising paths for enhancing retail inventory optimization in data-rich environments.

5. Conclusion and Future Work

In summary, our research plan demonstrates how to implement and analyze a single-period inventory (newsvendor) model using real retail data. We outlined the steps - from data preprocessing and cost modeling to computing optimal stock levels and comparing with actual decisions - and provided code snippets for each step. Analyzing 73,100 product-store-day observations, we identify a systematic 48% cost gap between theoretical cost-optimal inventory levels and actual retailer practice, driven by the retailer's preference for near-perfect service levels (99.5%) over cost minimization. The results reveal that retailers implicitly assign shortage costs 3-5 times higher than immediate lost revenue, reflecting intangible factors like customer lifetime value and brand reputation that dominate inventory decisions in competitive retail markets.

This study contributes both methodologically - demonstrating forecast-error distribution effectiveness, and managerially - quantifying the cost of service-level maximization strategies and identifying conditions under which inventory reduction is feasible. These insights benefit academic academics seeking to understand theory-practice gaps and practitioners evaluating whether high service-level strategies justify their costs or whether selective inventory optimization could capture savings without sacrificing competitive position.

Future scope of this study can be as follows:

I. **Refine Cost Estimates:** If possible, gather or estimate actual costs (e.g., how much lost sales truly cost in terms of customer attrition, or how holding cost scales with inventory and time). Adjust the model with these refined costs to see if the gap closes.
II. **Service Level Constraints:** Introduce minimum service level constraints (e.g., 95% fill rate) and re-optimize to identify compromise policies balancing cost efficiency with service requirements.
III. **Segmentation:** Apply the model separately for different product categories or regions. Some products (like perishable groceries) might warrant a different treatment than others (like electronics) due to shelf-life or substitution effects. We could also examine if competitor pricing or weather conditions systematically shift the demand distribution and adjust the stocking decisions accordingly (for example, heavy snow might dramatically cut store traffic, so the distribution of demand on snowy days is different).
IV. **Data-driven Policies:** Train quantile regression or machine learning models to directly predict target percentiles conditional on daily features, enabling fully data-driven newsvendor policies that adapt to context without distributional assumptions.
V. **Validate on Profit/Revenues:** Translate the cost results into profit or service metrics. For example, simulate a period (e.g., a year) under the newsvendor policy vs. actual: how much would profit increase, and how many sales would be delayed/lost? This would make the case analysis more concrete for a business audience.

Ultimately, this work demonstrates that classical optimization models remain highly relevant when properly calibrated to real operational contexts and augmented with modern forecasting capabilities.

Funding Statement
This project did not receive any external financial support; it was fully funded by the authors.

References
Andersen Consulting, Where to look for incremental sales: The retail problem of stockout merchandise, *The Coca Cola Research Council*, Atlanta, 1996.

Bandi, C., Han, E. and Nohadani, O., Sustainable inventory with robust periodic-affine policies and application to medical supply chains, *arXiv preprint*, 2018. https://doi.org/10.48550/arxiv.1806.06744

Bandi, C., Han, E. and Nohadani, O., Sustainable inventory with robust periodic-affine policies and application to medical supply chains, *Management Science*, vol. 65, no. 10, p. 4636, 2019. https://doi.org/10.1287/mnsc.2018.3152

Cao, J., A conformal approach to feature-based newsvendor under model misspecification, *arXiv preprint*, 2024. https://doi.org/10.48550/arxiv.2412.13159

Cheng-hu, Y., Li, X., Zhang, Y. and Wang, Z., A data-driven newsvendor problem: A high-dimensional and mixed-frequency method, *International Journal of Production Economics*, vol. 266, Article 109042, 2023. https://doi.org/10.1016/j.ijpe.2023.109042

D'Urso, D., Zizzo, D.J., Camerer, C.F. and Murphy, R.O., Modelling human behaviour in newsvendor game, *IFAC-PapersOnLine*, vol. 48, no. 3, p. 610, 2015. https://doi.org/10.1016/j.ifacol.2015.06.149

Dehaybe, H., Catanzaro, D. and Chévalier, P., Deep reinforcement learning for inventory optimization with non-stationary uncertain demand, *European Journal of Operational Research*, vol. 314, no. 2, p. 433, 2023. https://doi.org/10.1016/j.ejor.2023.10.007

Gavirneni, S. and Robinson, L. W., Risk aversion and implicit shortage cost explain the anchoring and insufficient adjustment bias in human newsvendors, *Operations Research Letters*, vol. 45, no. 3, p. 191, 2017. https://doi.org/10.1016/j.orl.2017.01.005

Gioia, D. G. and Minner, S., On the value of multi-echelon inventory management strategies for perishable items with on-/off-line channels, *Transportation Research Part E: Logistics and Transportation Review*, vol. 180, Article 103354, 2023. https://doi.org/10.1016/j.tre.2023.103354

Gong, S., Liu, H. and Zhang, X., Deep generative demand learning for newsvendor and pricing, *arXiv preprint*, 2024. https://doi.org/10.48550/arxiv.2411.08631

Gruen, T. W. and Bharadwaj, S., Retail out-of-stocks: A worldwide examination of extent, causes and consumer responses, 2002.

Han, J., Hu, M. and Shen, G., Deep neural newsvendor, *SSRN Electronic Journal*, 2023. https://doi.org/10.2139/ssrn.4582188

Harsha, P., Aydin, G., Çalışkan, E. and Keskinocak, P., An optimistic-robust approach for dynamic positioning of omnichannel inventories, *arXiv preprint*, 2023. https://doi.org/10.48550/arxiv.2310.12183

Hssaine, C. and Sinclair, S. R., The data-driven censored newsvendor problem, *arXiv preprint*, 2024. https://doi.org/10.48550/arxiv.2412.01763

Huber, J., Minner, S., Wagner, S. M. and Zhuang, J., A data-driven newsvendor problem: From data to decision, *European Journal of Operational Research*, vol. 278, no. 3, p. 904, 2019. https://doi.org/10.1016/j.ejor.2019.04.043

Liu, C. and Zhu, W., Newsvendor conditional value-at-risk minimisation: A feature-based approach under adaptive data selection, *European Journal of Operational Research*, vol. 313, no. 2, p. 548, 2023. https://doi.org/10.1016/j.ejor.2023.08.043

Maichle, M. J., Schiffer, M. and Schmid, V., A general-purpose deep reinforcement learning approach for dynamic inventory control, *SSRN Electronic Journal*, 2024. https://doi.org/10.2139/ssrn.4754937

Manousiadou, A., Behavioral theories in inventory management and forecasting: Driving supply chain resilience, *Journal of Supply Chain Management Science*, vol. 5, p. 143, 2024. https://doi.org/10.59490/jscms.2024.7758

Mieghem, J. A. V. and Rudi, N., Newsvendor networks: Inventory management and capacity investment with discretionary activities, *Manufacturing & Service Operations Management*, vol. 4, no. 4, p. 313, 2002. https://doi.org/10.1287/msom.4.4.313.5728

Qin, Y., Wang, R., Vakharia, A. J., Chen, Y. and Seref, M. M. H., The newsvendor problem: Review and directions for future research, *European Journal of Operational Research*, vol. 213, no. 2, p. 361, 2010. https://doi.org/10.1016/j.ejor.2010.11.024

Raz, G., Managing inventories: The newsvendor model, *SSRN Electronic Journal*, 2017. https://doi.org/10.2139/ssrn.2974971

Saccomanno, F. P., Trivella, A. and Guerriero, F., Integrated sales planning for in-store retail: A multi-stage stochastic optimization approach, *SSRN Electronic Journal*, 2024. https://doi.org/10.2139/ssrn.4977009

Shamir, D., The hidden costs of poor inventory management: How much are you really losing?, *ToolsGroup*, 2024.

Singh, , 2009.

Tian, Y. and Zhang, C., An end-to-end deep learning model for solving data-driven newsvendor problem with accessibility to textual review data, *International Journal of Production Economics*, vol. 265, Article 109016, 2023. https://doi.org/10.1016/j.ijpe.2023.109016

Uppari, B. S. and Hasija, S., Modeling newsvendor behavior: A prospect theory approach, *Manufacturing & Service Operations Management*, vol. 21, no. 3, p. 481, 2018. https://doi.org/10.1287/msom.2017.0701

Xu, X., Li, Y., Chen, Y. and Chen, X., Optimal decisions for the loss-averse newsvendor problem under CVaR, *International Journal of Production Economics*, vol. 164, p. 146, 2015. https://doi.org/10.1016/j.ijpe.2015.03.019

Zhang, D., Sun, X., Li, Y. and Wang, J., Robust optimization approaches in inventory management: Part A—the survey, *IISE Transactions*, vol. 57, no. 7, p. 818, 2024. https://doi.org/10.1080/24725854.2024.2381713

Biographies

Md Shoaib Mahmud is a Lecturer in the Department of Industrial and Production Engineering at the Military Institute of Science and Technology (MIST), Dhaka, Bangladesh. His research focuses on data-driven operations research and risk analytics, including supervised graph learning for resilient transportation networks, prescriptive analytics for urban flash-flood response, equity-aware ambulance dispatch planning, and reverse-logistics and inventory optimization in supply chains. He has co-authored multiple conference and journal publications in these areas, linking advanced analytics with disaster resilience, healthcare operations, and service systems.

Salman Ibrahim Saad is a Lecturer in Industrial & Production Engineering at the Military Institute of Science and Technology (MIST), Bangladesh. He earned his B.Sc. in Industrial and Production Engineering with summa cum laude honors in 2024. His research combines machine learning with optimization techniques to solve real-world supply chain challenges. His undergraduate thesis developed innovative electric vehicle routing models using Gurobi, demonstrating significant cost and vehicle reductions through split-delivery approaches. At MIST, Salman teaches operations research, statistics, and data analytics using industry tools like CPLEX, Python, and SPSS. He gained practical experience automating systems at Nestlé and Unilever Bangladesh. As former President of MIST's IEOM Student Chapter, he organized professional development programs for students.

Mahmudul Hassan Miazi is a fourth-year undergraduate student in the Department of Industrial and Production Engineering at the Military Institute of Science & Technology (MIST), Bangladesh, graduating in 2026. He has consistently ranked among the top students in his program, earning the MIST Dean's List Certificate for three consecutive terms. His academic interests include supply chain optimization, machine learning applications and ergonomics. Mahmudul completed his industrial attachment at British American Tobacco Bangladesh in the Primary Manufacturing Department, contributing to the OXYLOW nitrogen-flushing packaging line for export operations. He has also led an academic project on designing a modular multi-machine workstation for flexible small-scale mechanical work. He presented his work on ergonomics face mask design at ICMEAS 2025. Mahmudul's current thesis research focuses on fiber reinforced foamed concrete composites. Alongside the academics, he was the champion of the HULT Prize 2024 – On-Campus Round with ALGAEVA, a biodegradable algae-based rubber concept.