Research Article GU J Sci 37(4): 1853-1869 (2024) DOI: 10.35378/gujs.1334184

Gazi University

Journal of Science

http://dergipark.gov.tr/gujs

Data-Driven Mechanisms for a Newsvendor Problem: A Case Study

Afsin SANCAKTAROGLU, Burak GOKGUR*, Ayse KOCABIYIKOGLU

Sabancı Business School, Sabancı University, Orta Mahalle, Tuzla 34956, İstanbul, Türkiye

Highlights

- This paper develops a data-driven approach for the Newsvendor problem.
- The proposed approach is tested on a real data set.
- The value of incorporating data into an operational decision framework is investigated.
- The impact of data-driven mechanisms on reducing waste is revealed on a real data set.

Article Info

Received: 28 July 2023
Accepted: 17 May 2024

Keywords

Newsvendor
Machine learning
Food waste

Abstract

Reducing food waste is paramount for a sustainable future as its implications are important to achieving sustainable development goals set by the United Nations. In many industry groups, the public awareness of reducing food waste that may potentially emerge along firms' operations has grown. In the era of Big Data, one of the most pursued exercises of this escalating attention on reducing food waste is to utilize artificial intelligence techniques to incorporate sustainability concerns into the decision framework. Many firms embrace machine learning methods to build effective decision mechanisms that help make efficient and sustainable decisions. In this study, we analyze the impact of blending machine learning approaches with demand forecasting and order quantity decisions for a firm operating in a setting where the market demand is random, and the demand structure is not observable to the firm. The performance of the methodology is evaluated on sunflower seed demand data taken from Tadım company. Our results suggest that the joint consideration of forecasting and ordering decisions using the quantile regression approach can lead the firm to decrease its operational cost by 8,11% on average.

1. INTRODUCTION

Food waste has reached unprecedented levels. According to the study conducted by [1], about 133 billion pounds of the total food produced (equivalently, around 31%) could not be brought for consumption in 2010, delineating the fact that how resources along various tiers of the supply chain are inefficiently used and allocated. To reduce food waste and manage their operations in line with the agenda adopted by the United Nations, many firms have attempted to construct data-driven decision mechanisms that help make both efficient and sustainable decisions.

In particular, for products with limited life cycles, the mismatch between supply and demand is one of the primary reasons for accelerating the accumulation of food waste along a supply chain. Matching supply with demand lies at the core of operations management; when the demand is uncertain, firms make their ordering decisions considering the trade-off between having excess units in inventory and falling short of demand: if the order quantity exceeds the realized demand, firms have extra units that cannot be sold, and thus, these units become waste. Otherwise, firms do not match supply with demand and lose the potential profit that could have been earned if the order quantity had been higher.

Operational inefficiencies occur especially when the demand is uncertain, and the demand distribution parameters are unknown, corresponding to the case many firms have encountered in practice as opposed to the case where distribution parameters are assumed to be known as many theoretical models consider (i.e.,

*Corresponding author, e-mail: burak.gokgur@sabanciuniv.edu

traditional newsvendor problem). The need for following a holistic approach and involving machine learning approaches in decision mechanisms emerges at this stage where the demand distribution cannot be fully characterized (i.e., data-driven newsvendor problem). In the era of big data, the accessibility of customer data helps the demand characterization problem and elevates the fit of inventory models to real-world settings [2].

We formulate the problem as a data-driven newsvendor problem. In the presence of uncertain demand with unknown distribution parameters, we pursue two methodologies to estimate demand structure and determine order quantity, which we dub as sequential approach and joint approach, respectively. The reason behind taking the two mechanisms into account is to quantify and highlight the impact of integrating demand forecasting with ordering decisions on the profit that the firms could achieve. The sequential approach comprises two steps. In the first step, the future demand is estimated using historical demand patterns and other features. While generating the demand estimates, we consider both traditional and machine learning methods and select the procedure with the minimum estimation errors. The second step is determining the order quantity. In this step, we employ data-driven optimization methods that do not need any inputs about the structure of demand distribution, unlike traditional optimization approaches. In particular, we utilize approaches such as Sample Average Approximation (SAA) at the optimization stage of this approach. The integrated approach takes a holistic view and combines demand forecasting and data-driven optimization: it determines the optimal order quantity using historical demand patterns and features.

We partner with a leading Turkish packaged nuts company, i.e., Tadım, to establish a data-driven mechanism that aims to reduce food waste by matching supply with demand as best as possible. Complying with the short-lived nature of the firm's products, the proposed mechanism generates demand estimates for weekly demand. In this context, we would like to answer the following questions: What do machine learning methods bring about to the firm regarding food waste? What is the value of coordinating demand forecasting and optimization stages?

Our numerical experiment suggests that machine learning methods perform better than conventional methods in forecasting, thereby lowering the potential waste that occurred due to inefficiency at this stage. In particular, Random Forest and Linear Regression methods yield results with 10% higher accuracy than those obtained by the reference method. The firm decreases the total cost in the optimization stage using machine learning methods. On average, Quantile Regression (QR) achieves the best results that yield a 6% improvement in the total cost incurred by the firm.

The organization of the paper is as follows: section 2 provides a summary of the relevant literature and the contributions of this study. In section 3, we delineate the problem setting and present the structure of the dataset we work on. Sections 4 and 5 examine the approaches we employ to solve the problem and compare the impact of incorporating machine learning methods into the decision-making framework by comparing the results with those obtained by traditional methods. In section 6, we provide a summary of our findings and conclude with avenues for future research.

2. LITERATURE REVIEW

The primary objective is to quantify the impact of employing machine learning methods in the decision-making framework on a dataset of a firm facing uncertain demand with an unknown distribution parameter. Hence, we focus on the studies that examine the data-driven newsvendor problem in different contexts in the relevant literature. We cluster the existing studies with respect to the path they take in the development of solution approaches into two classes: i) the sequential approach (i.e., estimating demand and optimizing order quantity separately) and ii) the joint approach (i.e., estimating demand and optimizing order quantity jointly).

Studies that fall into the first cluster, though few, focus on either developing a solution approach only for optimizing the order quantity or building a two-stage method composed of demand estimation and optimization. [3] suggests a sampling-based algorithm to address the single-period newsvendor problem where the demand structure is unknown and develops bounds on the number of samples needed to ensure that the expected cost obtained by the sampling-based policies is close to that of the optimal policies; the developed bounds are distribution-free. For a similar problem setting considered in [3], [4] derives a new analytical bound for the relative regret of the sample average approximation (SAA); the bound is tighter than the already existing bounds. Through computational experiments, the authors show the empirical accuracy of the bound, and conclude that the performance of SAA is dependent on the demand distribution's weighted mean spread. [5] studies the data-driven newsvendor problem when both the firm and consumers learn the product's value over time, revealing that two-sided learning lowers the firm's optimal inventory level compared to the case of one-sided learning.

Most studies in the relevant literature investigate joint consideration of demand estimation and order quantity optimization. [6] develops a maximum entropy-based approach that allows decision-makers to form an expectation about the demand distribution and update it over time as new data becomes available and shows that the closed-form solution they derive for the updating mechanism generalizes the traditional Bayesian approach. Besides, numerical experiments conducted in the study of [6] reveal the importance of incorporating partial distributional information into the mechanism. [7] deals with the pricing and order quantity problem of a retailer in a setting where the demand function is unknown. The authors develop a maximin framework and, to solve the maximin model, they propose a two-sided cutting surface algorithm; they use a real dataset to test the proposed algorithm. They conclude that a risk-averse retailer opts to achieve less expected profit when the true demand information is lacking. Besides, they state that the number of iterations needed by the proposed algorithm is dependent on the initial region of uncertainty for the demand function. [8] builds a single-step learning mechanism to solve a data-driven newsvendor problem, where the decision-maker has limited information about demand features and presents performance bounds of the proposed approaches. In particular, algorithms are developed based on the empirical risk minimization principle and kernel-weight optimization, respectively. The authors test the algorithms using a large data set and conclude that the proposed algorithms produce better results than the best practice benchmark by around 25% in terms of the out-of-sample cost, and the kernel-weight optimization-based algorithm is faster than the empirical risk minimization principle-based algorithm. [2] works on gleaning value from data in the context of a data-driven newsvendor problem; the authors, particularly, develop solution methods relying on Machine Learning and Quantile regression; based on the numerical experiments conducted in the study, they conclude that machine learning techniques perform better than traditional methods when the data set is sufficiently large. [9] develops a method that relies on artificial neural networks for the newsvendor problem and studies the impact of a machine-learning-based approach in demand forecasting on optimizing order quantity level. The authors conclude that the proposed method brings up to about 30% cost savings. [10] proposes a solution framework that utilizes deep learning techniques for the data-driven newsvendor problem. Based on numerical experiments, the authors conclude that their approach performs better than existing methods, mainly when demand volatility is high. [11] develops a robust optimization approach to make decisions for the newsvendor problem with an unknown demand distribution and tests the performance of the proposed method using real data sets. [11] shows the proposed approach yields greater profits than the level existing approaches can achieve.

To summarize, this study contributes to the relevant literature by employing various learning and optimization methods in the decision-making framework and quantifying its impact on the firm's objective. Besides, we unravel the value of using data-driven methods by comparing them with model-based counterparts. In particular, the contributions of this study can be stated as follows: i) the comparison of the performance of sequential and integrated policies in terms of the total costs that they bring about on a new case study, and ii) the revelation of the value of employing data-based operational policies over a new case study.

3. PROBLEM SETTING

Each week, the firm must choose an inventory level for the sunflower seed to be roasted (i.e., the sunflower seed is the raw material in the roasted sunflower production) before the demand for the roasted sunflowers is realized. The demand depends on various external factors and is subject to change; so, it is uncertain, i.e., it is not known in advance.

Because the final product is short-lived, the firm determines its production schedule for each week separately and does not prefer carrying inventory to future periods. It follows the lot-for-lot (L4L) policy as the production strategy. The goals of the firm are twofold: i) to match supply (i.e., the sunflower seed) with demand (i.e., the roasted sunflower) at the targeted service level and ii) to minimize the amount of excess waste that occurs due to mismatch between supply and demand. In this context, the firm experiences two different cases: excess supply (i.e., overage) and excess demand (i.e., underage). In the case of excess supply, the firm would incur higher inventory-related costs as the demand falls short of the supply and higher production-related costs due to a greater volume of production. Otherwise, when excess demand is present, the firm cannot fully satisfy the demand and would miss the potential profit it would receive if it had a higher inventory level for the raw material. Besides, the firm would incur substitution-related costs that emerge due to customers who purchase another product when the firm's product is stocked out. Matching supply with demand requires minimizing the total expected cost the firm would incur in both cases of excess supply and excess demand. Hence, we formulate the firm's single-period inventory problem as the Newsvendor problem as follows:

$$\min_{q\geq 0} E[C(q)] \coloneqq E[C(q;D)]. \tag{1}$$

Equation (1) shows the objective function of the firm, where q denotes the order quantity the firm chooses for the product, D represents uncertain demand, and $C(q;D)$ corresponds to the function, comprised of shortage and overage costs, that depends on q and the random demand D. Because demand D is not known in advance, the firm minimizes the expected total cost of shortage and overage. The function $C(q;D)$ can be given by:

$$C(q;D) \coloneqq b(D-q)^{+} + h(q-D)^{+}, \tag{2}$$

where b is the unit cost associated with being unable to satisfy demand (i.e., underage cost), and h is the unit cost of having excess inventory (i.e., overage cost). In Equation (3), if the cumulative demand distribution function (CDF) of the random variable D, denoted by F, is known, the optimal solution, as shown in [12], is to determine the order quantity by the critical fractile as follows:

$$q^{*} = \inf\left\{y: F(y) \geq \frac{b}{b+h}\right\}. \tag{3}$$

In practice, the demand distribution cannot be easily characterized. In most cases, decision-makers need to learn the structure of the demand distribution to deal with the inventory problem. Utilizing historical demand data helps revamp this obstacle by replacing the true theoretical expectation with a sample average expectation through machine learning methods [2]. In the following sections, we present a data-driven solution approach to predict demand and determine the optimal order quantity.

4. DATA

Our access to the uncensored demand dataset for sunflower seeds covers 167 weeks between January 2018 and March 2021, provided in Figure 1.

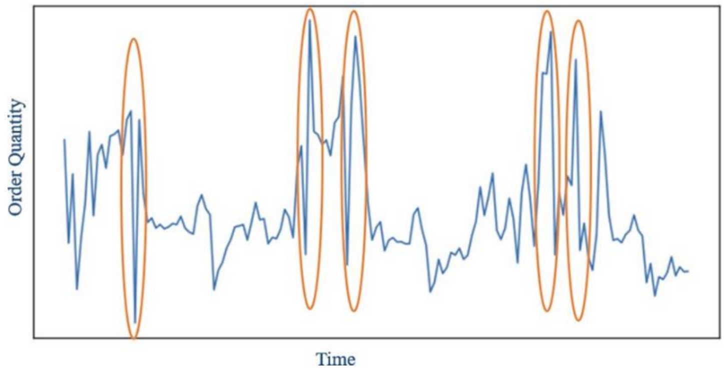

Figure 1. *Weekly Demand in Tonnes*

Figure 1 demonstrates that the weekly demand over approximately three years has fluctuated. The firm has experienced demand shocks at some intervals. In the figure, we mark these dramatic increases and decreases with circles. A closer look at Figure 1 reveals that the periods during which demand shocks have been observed actually coincide with the weeks right before, during, and right after Eid. Since Eid is an official holiday in Turkey, the demand piles up immediately before and after this period. The company executes a particular working program for these periods; the production schedule is built together with the sales and supply chain departments.

After observing the general structure of the aggregate demand data, it is helpful to analyze the shape of its distribution that will serve as input toward the proposed data-driven mechanism. Figure 2 demonstrates the distribution of the aggregate demand data with respect to particular intervals of tonnes. We can state that the distribution shape shown in Figure 2 is right-skewed.

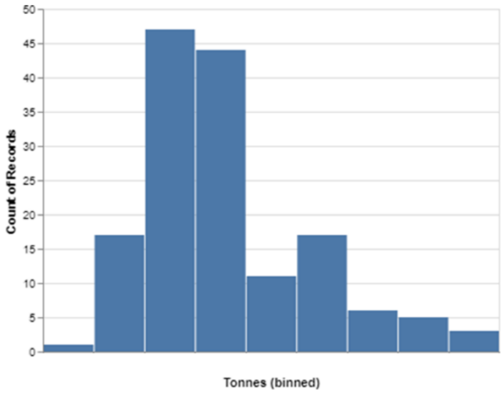

Figure 2. *Histogram of Weekly Demand*

In order to verify the skewness of the data, we apply the Chi-square test and obtain the p-value of 7.47E − 08, implying that the distribution does not satisfy the requirements for normality. Besides, after further investigating and receiving information from the firm, we notice that the inventory levels are kept at different levels and vary across seasons.

Figure 3 presents the seasonal demand distributions over three years. After clustering the aggregate demand data into seasons, we apply the Chi-square test on seasonal demand distributions. We obtain p-values of 0.176, 0.637, 0.001, and 0.549 for the spring, summer, autumn, and winter seasons, respectively. The resulting p-values imply that, except for the distribution belonging to the autumn season, demand distributions observed in the remaining three seasons comply with the requirements for the normal distribution, which can also be underpinned by the shape of the seasonal distributions provided in Figure 3.

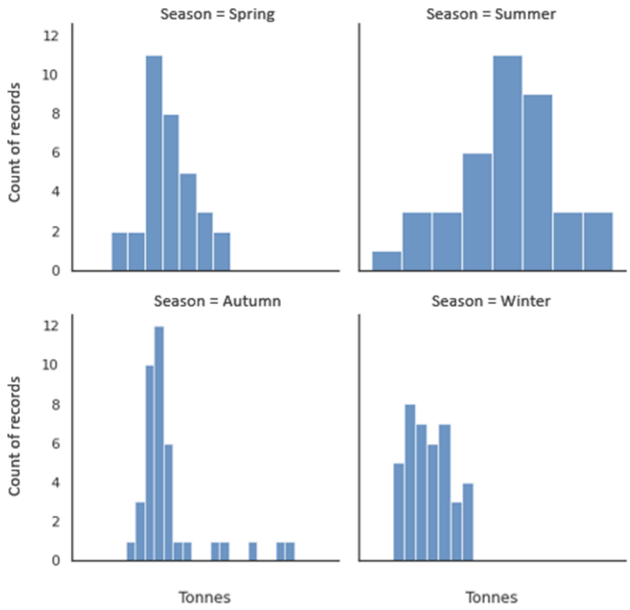

Figure 3. *Seasonal Demand Distributions*

After thoroughly examining the dataset, presented in Figures 1, 2, and 3, we exclude those above three standard deviations and label them as outliers. Analyzing the weeks in which outliers are eliminated reveals an important observation associated with them: fill rates from distribution centers fell below 70% in the previous week. Being unable to supply sufficient units of products that would satisfy the demand of the previous week caused an excessive rise in demand realized in the relevant week, thereby creating anomalies due to the orders that fell short of demand during harvest. Hence, we extract these outliers from the data.

Finally, we consider internal and external features that could affect the demand to be incorporated into machine learning models. Table 1 presents the list of these explanatory variables that are related to the firm's operational performance, calendar, and transactional data.

Table 1. *Features of the machine learning framework*

Features	**Data Type**	**Description**
Lag1-Lag16	Continuous	Demand that realized in the last between one and sixteen weeks, given in terms of tonnes.
Holiday	Nominal	It equals 1 if that particular week includes a religious holiday (Kurban or Ramazan); otherwise, it is set to 0.
Ramazan	Nominal	It equals 1 if that particular week includes Ramazan; otherwise, it is set to 0.
NewYear	Nominal	It equals 1 if that particular week includes New Year; otherwise, it is set to 0.
Before_holiday	Nominal	It equals 1 if that particular week is followed by a religious holiday; otherwise, it is set to 0.
After_holiday	Nominal	It equals 1 if that particular week follows a religious holiday; otherwise, it is set to 0.
Bnewyear	Nominal	It equals 1 if that particular week is one week before New Year week; otherwise, it is set to 0.
Workdays	Ordinal	The number of days the company runs to meet demand.
LagStonnes	Continuous	The supply amount in tonnes the company has in previous weeks.
LagFullfilment	Continuous	The supply amount in percentage the company has in previous weeks.

LagChange	Continuous	The realized change between the consecutive last two weeks.
LagChangeBinary	Nominal	It equals 1 if the percentage change between the consecutive last two weeks is positive; otherwise, it is set to 0.
Seasonality1-2	Continuous	Seasonal description

Table 1 presents all the features to make the model more interpretable. In Section 4.1, we determine the number of features that sufficiently explain the variations in demand.

4.1. Best Subset Selection

In this procedure, a separate least square regression is fit for each possible combination of the explanatory p features presented in Table 1. As provided in [13], the steps of the procedure can be given as follows:

Step 1. Let M_0 denote the *null model*, containing no predictors. M_0 predicts the sample mean for each observation.
Step 2. For $k = 1,2, \dots, p$:

- *a)* Fit all $\binom{p}{k}$ models that contain exactly *k* predictors.
- *b)* Select the *best* among these $\binom{p}{k}$ models, and label it M_k. "*best*" is defined as having the smallest residual sum of squares (RSS), or, equivalently, the largest R^2.

Step 3. Select a single best model from among $M_0, \dots, M_p$ using cross-validated prediction error, C_p, (AIC) (i.e., Akaike Information Criterion), BIC (Bayesian Information Criterion), or adjusted R^2.

The best subset selection method suffers from computational burden. The number of models that need to be investigated quickly increases in the number of explanatory features (i.e., d). For instance, when we increase the number of features to p from $\mathrm{p} - 1$, the additional number of models that we need to analyze equals $\sum_{\mathrm{k}=1}^{\mathrm{p}} \binom{\mathrm{p}}{\mathrm{k}} - \sum_{\mathrm{k}=1}^{\mathrm{p}-1} \binom{\mathrm{p}-1}{\mathrm{k}} = (2^{\mathrm{p}} - 1) - (2^{\mathrm{p}-1} - 1) = 2^{\mathrm{p}} - 2^{\mathrm{p}-1}$.

In total, there are 29 different variables that we can use, and putting all of those provided in Table 1 into the best subset selection method would produce more than 536 million (i.e., 2^{29}) different models. To circumvent this computational cumbersome, we first simultaneously ran all the variables and removed those with p-value levels above 0.05. This way enabled us to decrease the number of variables. For example, in the case of having 10 variables in the model, the number of models we excluded equals $2^{29} - 2^{10} > 536$ million, saving a significant amount of time.

Another problem that we should be cautious of is overfitting. We select two models with seven and ten variables, presented in Table 2, to prevent the problem of overfitting after implementing indirect and validation set approaches on train and test data [13].

Table 2. *Selected models with input variables*

Model	**Input Variables**
7 variables	Lag1, Lag2, Lag6, Lag11, Holiday, Aholiday, Bholiday
10 variables	Lag1, Lag2, Lag6, Lag11, Holiday, Aholiday, Bholiday, LagFullfill-ment, NewYear, Bnewyear

The selected variables in the two models reveal that the recent demand and the holiday periods help explain the variation. As shown in Step 3 of the procedure, we employ the Akaike Information Criterion (AIC), Bayesian Information Criterion (BIC), C_p, and adjusted R^2 statistics to measure the quality of these two models [14]. For a model with d estimators, the C_p estimate can be given by: $C_p = \frac{1}{n}(RSS + 2d\hat{\sigma}^2)$, where RSS is the residual sum of squares and $\hat{\sigma}^2$ is an estimate of the variance of the error, and n is the number of observations in the training set. Likewise, the AIC is another estimator of the prediction error that deals with the model's trade-off between complexity and simplicity. AIC can be expressed as follows: $AIC = \frac{1}{n\hat{\sigma}^2}(RSS + 2d\hat{\sigma}^2)$. As their expressions suggest, C_p and AIC are proportional to each other. Similar to C_p, the lower the AIC score, the better model we have.

The third estimator is BIC, which can be given by: $BIC = \frac{1}{n\hat{\sigma}^2}(RSS + \log(n)\, d\hat{\sigma}^2)$. This criterion adds a logarithmic penalty to the RSS. This estimator takes on a small value when the test error is low, so a lower BIC value is preferred. The final indirect estimation is the adjusted R^2 and can be calculated as follows: $adjusted\ R^2 = 1 - \frac{RSS/(n-d-1)}{TSS/(n-1)}$, where TSS is the total sum of squares, RSS is the residual sum of squares and n is the number of observations in the training set. Unlike C_p, AIC, and BIC, a larger value of $adjusted\ R^2$ indicates a lower test error.

Among these statistics, only BIC's result suggests the model with fewer variables. We implemented both because we could not see that one model was superior. Figures 4 and 5 demonstrate the suggested number of explanatory variables using C_p, AIC, BIC, and adjusted R^2, respectively.

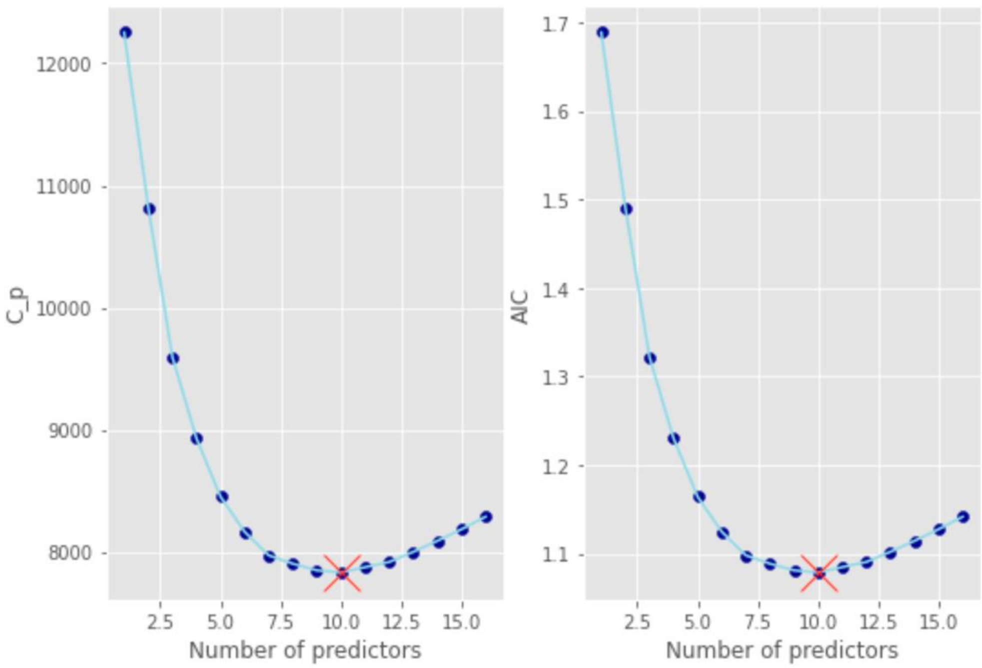

Figure 4. *Subset Selection Using C_p and AIC*

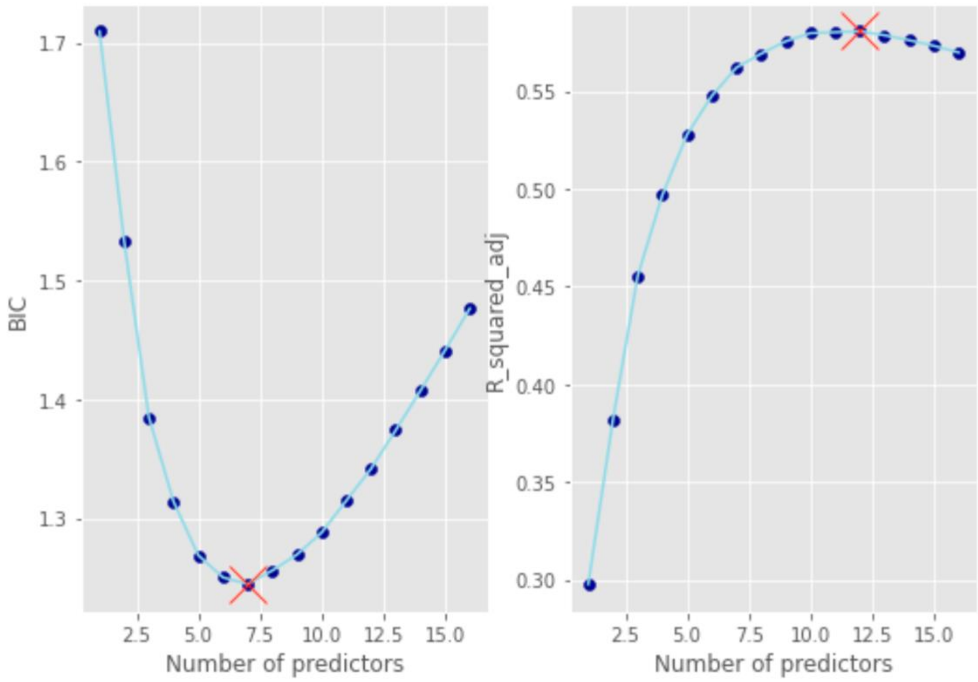

Figure 5. *Subset Selection Using BIC and adjusted* R^2

In Figure 4, C_p and AIC suggest a 10-variable model. In Figure 5, BIC produces a 7-variable model, and the model proposed by adjusted R^2, on the other hand, includes the number of variables between 10 and 12; because we want to refrain the model from being overfitted, we take the number of variables as 10 suggested by the adjusted R^2 criterion.

In the next section, we focus on the stage where we present the proposed optimization models for the firm's inventory problem. First, we present the model that rests on some assumptions on demand. Then, we move on to the model that is data-driven.

5. METHODOLOGY

This section explains the solution methodology followed in this study. First, we detail the sequential approach composed of two steps. In the first step, we employ forecasting models to estimate future demand. We introduce a model-based optimization method using assumptions regarding demand and cost structure [13]. Besides, we present a data-driven approach using Sample Average Approximation (SAA). Then, we focus on developing an integrated data-driven demand estimation and optimization model using quantile regression (QR).

5.1. Sequential Approach

This section details the stages through which the approach is constructed. Firstly, we zero in on demand estimation for the focal product and then the cost-minimizing (optimal) order quantity. The sequential approach requires first estimating the mean demand and the error distribution and then finding the optimal order quantity by solving the newsvendor problem.

The demand estimation stage considers traditional forecasting and machine learning-based forecasting methods. We first present five traditional approaches followed by the firm, which are Naive, Seasonal Naive (S-Naive), Median, Seasonal Median (S-Median), and Moving Average (MA):

1. *Naive* uses the realized sales observed in the last period as the forecast for the current period.
2. *Seasonal Naive* uses the realized sales observed in the same period a season ago as the forecast for the same period of the current season. For instance, in the prediction of the second fall week of the year, the realized demand for the same week of the previous year is used.

3. *Median* chooses the order quantity level separating the higher half from the lower half as future demand in the determined sample.
4. *Seasonal Median* clusters the available data sets into seasons and then chooses the seasonal order quantity levels separating the higher half from the lower half as the expected demand for each season.
5. *Moving Average* takes the average of the order quantities placed in the most recent predetermined number of periods before the current period. In the train dataset, the chosen level produces the best result for the weeks between two and sixteen.

When we turn to machine learning approaches that we employ in the study and compare them with traditional methods, we choose six procedures that have been heavily performed in the relevant literature. We introduce these methods as follows:

1. *Linear Regression* analyzes the linear relationship between independent and dependent variables. It has been applied to predict quantitative responses [14].
2. *Random Forest* builds independent tree-structured vectors that are independent and identically distributed in the same forest [15]. This approach can be used in regression and classification problems.
3. *XGBoost (eXtreme Gradient Boosting)* uses a gradient-boosted decision tree algorithm. Scalability and computational power can be considered as the distinguishing features of XGBoost [16]. What differentiates XGBoost from Random Forest is the way in which trees are constructed: trees are sequentially created in XGBoost, whereas multiple trees are produced in parallel in Random Forest. Besides, in XGBoost, the result is a total of outcomes from all the trees; however, in Random Forest, the outcome is determined via the majority vote.
4. *LightGBM* uses learning algorithms in a Gradient Boosting Decision Tree (GBDT) framework by utilizing GOSS (i.e., Gradient-based One-Side Sampling helps eliminate some of the data points while maintaining the prediction accuracy at a certain level) and EFB (i.e., Exclusive Feature Bundling helps reduce the computational burden) [17]. In the case of extensive data sets, LightGBM produces promising results.
5. *LSTM (Long Short-Term Memory)* generates the best outcomes by employing time series data and eliminating gradient problems [18]. In this approach, advanced RNN (i.e., recurrent neural network) stores memory states, which helps obtain information that enables us to understand the possible scenarios that will be realized in the next period.

The model-based optimization approach requires a particular forecast error distribution whose mean and standard deviation are estimated using past prediction errors [2]. We obtain these parameters in the estimation stage. To find the best quantity provided in Equation (3), the cost structure (i.e., h and b) and the forecast error distribution must be known. As we have presented in Section 4, the seasonal demands, except for autumn, comply with the requirements for the normal distribution, so we adopt normal distribution. However, in most business scenarios, the actual demand distribution has not been regularly predicted [19]. Using data-driven approaches attenuates the multiplier effect of a prediction error arising at the demand forecasting stage that may negatively impact the scale of a mistake to be made in the optimization phase; so leading to better results.

Now, we turn our attention to explaining how we adopt the sample average approximation (SAA) method, a robust data-driven technique with uncertainty [20], to the problem proposed in this study. The SAA is generally employed to solve stochastic optimization problems. In general, it is a promising approach for the two problem settings: 1) the actual demand distribution is known but computationally cumbersome, and 2) the actual demand distribution is unknown and computational evaluation is easy, like the problem we study, newsvendor, [4]. In SAA, empirical data replaces the demand distribution assumptions. In our case, using historical data as presented in the study of [4], we can express the problem as follows:

$$\min_{q\geq 0} \hat{R}\big(q;\boldsymbol{D}(n)\big) = \frac{1}{n}\sum_{i=1}^{n}[b(D_i - q)^+ + h(q - D_i)^+], \tag{4}$$

where the notation ^ indicates quantities are estimated from data. $\hat{R}\big(q;\boldsymbol{D}(n)\big)$ denotes the total expected cost of the firm when the samples of the demand are drawn from the empirical distribution. $\boldsymbol{D}(n)$ is the historical demand with n representing the total independent samples of the demand and D_i indicating the i^{th} demand point, q is the forecasted quantity, b is the unit-understocking cost, and h is the unit-overstocking cost (interested readers on the SAA approach are referred to [21]).

In Equation (4), obtaining close to the optimal order quantity depends entirely on demand estimation and cost ratio, making forecasting accuracy an important element in choosing the efficient order quantity [22]. In SAA, to determine the optimal order quantity, we need to incorporate the service level quantile of the empirical demand distribution. If the demand distribution was assumed to own a certain structure, such as normal distribution with mean μ and standard deviation σ, the service level quantile could have been given by the following expression: $\inf\left\{y: 1 - F(y, \theta) \geq \frac{b}{b+h}\right\}$, where θ represents parameters of the normal distribution (i.e., μ and σ), and then adding this level to the mean forecast value would have given the optimal order quantity. When the true demand distribution (i.e., F) is unknown to the firm, the service quantile can be calculated using the demand data as follows: $\inf\left\{y: 1 - \hat{F}(y, \boldsymbol{x}) \geq \frac{b}{b+h}\right\}$, where $\hat{F}(\cdot)$ represents the empirical demand distribution and $\mathbf{x}$ corresponds to the vector of demand features we have provided in Table 1; the summation of this level and the point forecast gives the optimal order quantity, as shown in the studies conducted by [2] and [8].

5.2. Integrated Approach

This approach directly incorporates the forecasting model into the order quantity optimization stage [2]. As discussed in [22], machine learning techniques can be used to build learning platforms for estimation and optimization in operations problems like newsvendor.

In the integrated approach, the optimal order quantity is estimated from the feature data by treating the optimal order quantity q^* of the standard newsvendor model provided in Equation (3) defined in terms of the feature data $\mathbf{x}$. [8] develops a linear programming model for this problem. [2] extends the approaches put forward in [8] by accommodating non-linear relationships. In line with the method proposed in [2], we predict the optimal order quantity from the feature data. The objective function we consider in the integrated approach can be given by:

$$\min_{\phi} \frac{1}{n} \sum_{i=1}^{n} \left[b\big(D_i - q_i(\phi, \mathbf{x}_i)\big)^+ + h(q_i(\phi, \mathbf{x}_i) - D_i)^+ \right], \tag{5}$$

where $q_i(\phi, \mathbf{x}_i)$ is the order quantity obtained by employing the machine learning method in period (or demand point) i, $i \in \{1, 2, \ldots, n\}$, $\mathbf{x}_i$ corresponds to the vector of the features, and ϕ represents the vector of parameters of the learning method, i.e., Linear, Artificial Neural Networks, and Decision Trees. Also, D_i indicates the i^{th} demand point, b is the unit-understocking cost, and h is the unit-overstocking cost.

Following the approach proposed in [2], the optimization problem provided in Equation (5) can be reformulated as a non-linear mathematical model that aims to minimize the empirical total cost of understocking and overstocking. Solving the non-linear model, based on the empirical data, gives us the total cost-minimizing parameters for the learning method (i.e., ϕ^*). After the training phase, the optimal order quantity for period i corresponds to the quantile forecast with $q_i^*(\phi^*, \mathbf{x}_i)$.

6. RESULTS

In this section, we first present results regarding the performances of the forecasting model and optimization approach. Then, we turn our attention to comparing the sequential and integrated approaches that form the base of this study.

To analyze the impact of machine learning methods, we consider various conventional forecasting methods that serve as a benchmark in comparison. Because the number of observations we have is relatively low, we consider 70% of the data as a train set and 30% as a test set. Then, we run all the methods and compare their performances on both sets. All the models are coded in Google Colab [23]. We employ two models with seven variables and ten variables, respectively, for linear regression (i.e., LinearReg7 and LinearReg10), XgBoost (i.e., XgBoost7 and XgBoost10), Random Forest (i.e., RandomForest7 and RandomForest10) and LightGBM (i.e., LgbmDT7 and LgbmDT10) using the inputs generated by the best subset selection method. In these methods, the train and test sets are chosen randomly, and the same observations are used. Besides, to prevent overfitting, we use cross-validation to determine the number of leaves and learning rates in tree-based ensemble methods.

In LSTM, 70% of the set is allocated to the train set, and the remaining portion is considered as the test set since the observations must be time series. Also, LSTM only takes tonnage data as the input parameter. On the other hand, in the ensemble method, which combines the existing models, we choose the method that yields the best result. Trying different combinations of machine learning methods, we seek to find the best possible ensemble model. Hence, from these attempts, we conclude that LightGBM and Random Forest models with seven variables achieve the best results.

In the model-based approach (relying on normally distributed seasonal demand) and the data-based approach (employing the SAA method), overage and underage costs taken from the company and the forecasting results are used. Besides, we perform the integrated model using the QR approach. Then, we examine the total expected cost the firm incurs with each strategy.

6.1. The Analysis of the Forecasting Methods

To assess the forecasting methods, we use the following criteria: root mean squared error (RMSE), mean absolute percentage error (MAPE), and mean absolute error (MAE). Table 3 compares the performance of forecasting methods using RMSE, MAPE, and MAE. Table 3 shows that ML methods yield better results than conventional methods in estimating demand.

Table 3. *The comparison of forecasting methods*

	Train Set			Test Set		
Method	**RMSE***	**MAPE****	**MAE***	**RMSE***	**MAPE****	**MAE***
Naive	156.84	0.39	91.76	169.51	0.35	108.72
SNaive	161.34	0.4	95.03	181.09	0.38	117.39
Median	160.64	0.35	112.51	173.07	0.41	124.76
SMedian	116.14	0.33	72.45	152.34	0.38	115.87
Moving Average	129.39	0.37	84.08	168.2	0.38	120.93
LinearReg7	84.61	0.21	63.87	88.9	0.22	62.66
XgBoost7	161.88	0.32	127.09	165.09	0.35	127.66
RandomForest7	**49.18**	**0.1**	**32.54**	98.74	0.23	65.74

LgbmDT7	129.19	0.35	105.04	123.11	0.32	93.48
LinearReg10	83.09	0.19	61.42	**87.23**	**0.21**	**60.48**
XgBoost10	158.36	0.32	127.21	166.19	0.37	130.22
RandomForest10	50.89	0.11	33.63	102.11	0.24	69.45
LgbmDT10	129.19	0.35	105.04	123.11	0.32	93.48
LSTM	148.35	0.32	97.33	149.85	0.38	111.2

*scale-dependent error measure (RMSE and MAE allow comparing data sets on the same scale)

**percentage-based error measure (MAPE, a scaled version of MAE, allows comparing data sets on different scales)

In Table 3, Random Forest generates the results with minimum deviations (shown in bold font) on the train set. Furthermore, Linear regression achieves the results with minimum variations (shown in bold font) on the test set. The change in the methods that achieve the best results between train and test sets underpin that Random Forest overfits the test data.

6.2. The Analysis of the Optimization Methods

In the optimization stage, we consider seasonal normalization (S-Norm) as an additional model-based method to adopt varying seasonal demand distributions. The model-based method assumes the underlying demand structure to follow a normal distribution. Following the interview with the firm's supply chain manager, the ratio of the overage cost to the underage cost (i.e., the critical fractile) is set to 1/3, and we set the service level (i.e., the proportion of demand that can be met from stock) to 0.75.

Tables 4 and 5 present the cost performance of the proposed approaches for the target service level on the train and the test sets, respectively. Tables 4 and 5 report the percentage cost increase for each strategy relative to the benchmark case where the firm is assumed to know the demand structure. For instance, in Table 4, estimating demand using the Naive method and optimizing inventory level with S-Norm increases the firm's total cost by 44.92% as opposed to the total cost if the firm had perfectly observed the demand and placed its orders accordingly. The larger the percentage value is, the greater the distance a method has from the optimal quantity.

Table 4. *Performance Analysis: Train Set*

	Optimization Methods		
Forecasting Methods	Δ Cost S-Norm (Model-based)	Δ Cost SAA (Sequential)	Δ Cost QR (Integrated)
Naive	44,92%	50,45%	42,48%
SNaive	42,29%	49,98%	45,48%
Median	74,56%	63,53%	62,84%
SMedian	40,94%	38,88%	35,77%
Moving Average	40,24%	45,99%	37,04%
LinearReg7	31,14%	36,85%	30,89%
XgBoost7	83,18%	47,09%	26,75%
RandomForest7	25,07%	19,59%	15,13%

LgbmDT7	59,61%	60,61%	32,76%
LSTM	59,65%	48,69%	41,42%
LinearReg10	30,99%	35,44%	30,30%
XgBoost10	82,38%	44,84%	25,38%
RandomForest10	25,19%	20,00%	15,37%
LgbmDT10	59,61%	60,61%	32,76%
Average	**49,98%**	**44,47%**	**33,88%**

In Table 4, we observe that, on average, the integrated approach brings about a cost 33,88% higher than the benchmark case, followed by the sequential approach with a cost increase of around 44,47% and the model-based approach with 49,98%. Based on the training set, the best-performing strategy is the integrated approach, leading to 31,2% (47,5%) cost reduction compared to the sequential (model-based) strategy. Besides, total cost decreases in all models when the RandomForest7 is employed at the forecasting stage.

In Table 4, the difference between the average ΔCost S-Norm and average ΔCost SAA gives the economic value of operationalizing data-driven mechanisms, amounting to a cost reduction of 11,02% (= $1-\frac{44{,}47\%}{49{,}98\%}$). Also, the difference between the average ΔCost SAA and average ΔCost QR identifies the economic impact of employing approaches that integrate estimation and optimization phases, corresponding to a cost reduction of 23,81% (= $1-\frac{33{,}88\%}{44{,}47\%}$).

Table 5. *Performance Analysis: Test Set*

	Optimization Methods		
Estimation Methods	Δ Cost S-Norm (Model-based)	Δ Cost SAA (Sequential)	Δ Cost QR (Integrated)
Naive	49,22%	62,21%	54,62%
SNaive	55,57%	62,05%	56,91%
Median	78,16%	71,73%	79,91%
SMedian	58,82%	58,40%	59,43%
Moving Average	61,65%	68,20%	62,77%
LinearReg7	30,36%	37,76%	30,83%
XgBoost7	91,66%	77,20%	46,76%
RandomForest7	36,26%	40,53%	37,36%
LgbmDT7	52,44%	52,53%	57,31%
LSTM	59,15%	63,97%	64,60%
LinearReg10	29,53%	37,83%	29,11%
XgBoost10	93,25%	56,69%	46,48%
RandomForest10	36,22%	44,14%	38,77%
LgbmDT10	52,44%	52,53%	57,31%
Average	**56,05%**	**56,13%**	**51,58%**

Table 5 reports that, on the test set, deviations of the results obtained by all three methods from the benchmark case increase. The integrated approach yields solutions that generate an average total cost 51,58% higher than the benchmark case, still being the best-performing policy among the other strategies. The sequential approach brings about a cost increase of around 56,13%, and the model-based approach increases the cost by 56,05%. Based on the test set, each policy reaches its best with different forecasting methods. The model-based approach yields the best result with the LinearReg10 estimation method. The integrated approach, too, produces the best results with the LinearReg10 method. Employing the LinearReg7 forecasting method leads to the sequential approach to achieve its best result.

In Table 5, the results the integrated approach puts forward bring about, on average, a cost of 8,11% (=1-$\frac{51,58\%}{56,13\%}$) less than that of the sequential approach and a cost of 7,98% (=1-$\frac{51,58\%}{56,05\%}$) less than that of the model-based approach. It is also noteworthy that the averages of the results the model-based and the sequential approaches produce are almost equal, representing a case where the sequential approach increases the impact of the error (as stated in [2] and [8]), leading to the marginal effect of data-driven operationalization being at its minimum.

On the test set, employing LinearReg10 and QR in the estimation and optimization stages, respectively, yields the firm a 29.11% cost deviation from the optimal level. In machine learning methods, explanatory variables increase prediction quality, implying that forecast performance significantly impacts total cost. For example, the results achieved via XgBoost are far from the optimum level, almost doubling the cost. On the other hand, similar to what we conclude in the estimation stage, Random Forest performs better on the train set; however, it is surpassed by linear regression on the same set, showing that tree-based methods perform worse than linear regression.

In summary, the results reported in Tables 4 and 5 emphasize the operational value the integrated approach brings to the firm. On average, the joint consideration of estimation and optimization phases engenders 23,81% and 8,11% cost improvement on the train and test sets, respectively. On the other hand, when we draw our attention to the marginal value of using data-driven mechanisms, we observe that, as discussed in [2] and [8], the sequential approach increases the impact of the error as opposed to the integrated approach: employing the sequential approach decreases the cost by 11,02% relative to the model-based scenario on the training set; whereas, it increases the cost too slightly by 0,14% on the test set.

7. CONCLUSION

This study considers the newsvendor problem with a single product in which the demand structure is estimated using learning algorithms. We propose two methods: i) predicting demand and determining the order quantity in sequence (i.e., sequential approach), and ii) the incorporation of estimation results with the optimization phase (i.e., integrated approach). We first use machine learning methods to estimate the demand and seasonal normalization and sample average approximation to optimize quantity sequentially. We then introduce integrated demand estimation and optimization based on machine learning and quantile regression.

The insights that can be taken from this study can be given by: when the underlying demand distribution is unknown, machine learning methods perform better than traditional methods. Besides, uniting forecasting and optimization phases into one framework (such as QR) leads to more effective decisions compared to the approach in which the two stages are sequentially executed. In the forecasting stage, the performance of Random Forest and Linear Regression is 10% better than the results achieved via the best-performing benchmark. In the optimization stage, we observe that the QR method yields solutions with 8,11% and 7,98% less costs than the sequential and model-based methods, respectively. Moreover, we find that the demand forecasting method is the key to performance, which aligns with the study of [2].

Our use of machine learning methods gives us a great advantage. Including variables such as holidays and fulfillment rate in the model increases the prediction power. Comparing it with traditional methods, we can observe the effect more clearly. Among the machine learning methods used in the study, Linear regression and Random Forest are the approaches with the highest performance. In contrast, the performance of techniques such as Xgboost and LightGBM is insufficient. This is because Xgboost and LightGBM need more training data to reach a reliable level. In our study, we have 167 data points.

Lastly, we point out future research directions. In this study, we partnered with a Turkish company, focusing on a single product. A direct line of future research is to cover all processes based on other products. Second, this study can be extended to a case where the firm manages its operations in a finite horizon that includes multiple periods. Finally, working with different companies could be another research area, contributing to the generalizability of the insights we put forward in this study.

ACKNOWLEDGEMENT

This manuscript is derived from the M.Sc. thesis of Afşin Sancaktaroğlu.

CONFLICTS OF INTEREST

No conflict of interest was declared by the authors.

REFERENCES

[1] Buzby, J.C., Wells, H.F., and Hyman, J., "The estimated amount, value, and calories of postharvest food losses at the retail and consumer levels in the United States", USDA-ERS Economic Information Bulletin, 121, (2014).

[2] Huber, J., Müller, S., Fleischmann, M., and Stuckenschmidt, H., "A data-driven newsvendor problem: from data to decision", European Journal of Operational Research, 278(3): 904-915, (2019).

[3] Levi, R., Roundy, R. O., and Shmoys, D.B., "Provably near-optimal sampling-based policies for stochastic inventory control models", Mathematics of Operations Research, 32(4): 821-839, (2007).

[4] Levi, R., Perakis, G., and Uichanco, J., "The data-driven newsvendor problem: new bounds and insights", Operations Research, 63(6): 1294-1306, (2015).

[5] Papanastasiou, Y., "Newsvendor decisions with two-sided learning", Management Science, 66(11): 5408-5426, (2020).

[6] Saghafian, S., and Tomlin, B., "The newsvendor under demand ambiguity: combining data with moment and tail information", Operations Research, 64(1): 167-185, (2016).

[7] Hu, J., Li, J., and Mehrotra, S., "A data-driven functionally robust approach for simultaneous pricing and order quantity decisions with unknown demand function", Operations Research, 67(6): 1564-1585, (2019).

[8] Ban, G.Y., and Rudin, C., "The big data newsvendor: practical insights from machine learning", Operations Research, 67(1): 90-108, (2019).

[9] Seubert, F., Stein, N., Taigel, F., and Winkelmann, A., "Making the newsvendor smart – order quantity optimization with ANNs for a bakery chain", AMCIS Proceedings, (2020).

[10] Oroojlooyjadid, A., Snyder, L.V., and Takac, M., "Applying deep learning to the newsvendor problem", IISE Transactions, 52(4): 444-463, (2020).

[11] Xu, L., Zheng, Y., and Jiang, L., "A robust data-driven approach for the newsvendor problem with nonparametric information", Manufacturing & Service Operations Management, 24(1): 504-523, (2022).

[12] Porteus, E., Foundations of Stochastic Inventory, Stanford University Press, (2002).

[13] Silver, E. A., and Pyke, D.F., Inventory and production management in supply chain, CRC Press, Taylor & Francis Group, (2017).

[14] James, G., Witten, D., Hastie, T., and Tibshirani, R., "An introduction to statistical learning: with applications in R", Springer, (2013).

[15] Breiman, L., "Random forests", Machine Learning, 45(1): 5-32, (2001).

[16] Chen, T., and Guestrin, C., "Xgboost: a scalable tree boosting system", Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, (2016).

[17] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T. Y., "Lightgbm: a highly efficient gradient boosting decision tree", Advances in Neural Information Processing Systems, 30, (2017).

[18] Kostadinov, S., "Recurrent neural networks with Python quick start guide: sequential learning and language modeling with tensorflow", PACKT Publishing Limited, (2018).

[19] Yue, J., Chen, B., and Wang, M.C., "Expected value of distribution information for the newsvendor problem", Operations Research, 54(6): 1128-1136, (2006).

[20] Bertsimas, D., Gupta, V., and Kallus, N., "Robust sample average approximation", Mathematical Programming, 171(1-2): 217-282, (2017).

[21] Kleywegt, A.J., Shapiro, A., and Homem-De-Mello, T., "The sample average approximation method for stochastic discrete optimization", SIAM Journal on Optimization, 12(2): 479-502, (2002).

[22] Qi, M., Mak, H., and Shen, Z.M., "Data-driven research in retail operations - a review", Naval Research Logistics, 67(8): 595-616, (2020).

[23] Sancaktaroğlu, A., "A Data-Driven Approach to Reduce Food Waste for a Consumer Goods Company", MSc. Thesis, Sabancı University, Sabancı Business School, İstanbul, 1-56, (2021).