RT2 - Retail Analytics:

Data-driven Newsvendor

Ben Lowery

May 7, 2022

1 Introduction

In 2015, the United Kingdom's retail sector produced 240,000 tonnes of surplus food (WRAP 2017). Of this, only 13% was redistributed to people or provided as animal feed, leaving a vast majority to contribute to the growing problem of food waste. A key issue to arise this century is food wastes contribution to climate change, with a study by Poore & Nemecek (2018) indicating food production is responsible for 26% of global greenhouse emissions and 24% of this figure is attributed to food waste. While household and manufacturing are the most responsible sectors for these figures, reduction in food waste for the retail sector is still an achievable and important goal. Conversely, retailers also need to keep in mind customer demand and avoid ordering too few products in fear of producing excess food waste, leading to unsatisfied customers and lost profits.

For this challenge in determining optimal inventory levels, satisfying demand and avoiding such situations, companies turn to retail analytics and leverage the vast amounts of retail data; harnessing this to predict trends, outcomes and make informed decisions on demand forecasting and inventory optimisation. Further, a key consideration within inventory optimisation is the planning horizon. These range from short, single period models investigating lone purchases, to infinite order-up-to models (Fisher 2009). Given the vast literature in the field, we restrict ourselves and focus primarily on this first problem, with models of perishable products over a single selling period and purchases for the product solely being made before a given selling season.

An area where this situation is prevalent is the Bakery sector. For large Bakery chains, or supermarkets containing baked sections, we envision a situation in which orders must be placed after the close of business for the next day. Given the placement of distribution centres, reordering during the day is not viable. While the shelf life for the majority of products do not last more than one day this leads to potential food waste, however, not enough food being ordered then leads to out-of-stock scenarios and lost sales for the day. Such scenario was investigated by Huber et al. (2019) with a large Germany bakery chain.

Comparisons can also be drawn between the issues faced with baked goods retailers, and newspapers. In which daily newspapers, also having the shelf-life of a single day period, need to be ordered by newsvendors who in turn need to estimate the demand they expect to face, which is almost sure to be clouded in uncertainty. This newspaper analogy lends itself as inspiration to the most widely used model in inventory management.

The Newsvendor (alternatively newsboy or single-period) problem has origins attributed to Edgeworth (1888) and his seminal work "A Mathematical Theory of Banking". The paper led to the subsequent establishment of inventory theory, although contemporary literature on the Newsvendor topic commonly uses the problem statement of Arrow et al. (1951) as a starting point for exploration and implementation.

The set up for the Newsvendor model leads to a stochastic optimisation problem that is simplistic and well investigated, making it suitable for a wide array of problem settings within inventory management. The key underlying assumptions are a company wanting to sell a lone product over a singe period of time which has uncertain demand, and for orders to be placed before it begins to be sold. A more formal definition is provided in the following subsection, alongside overviews of its possible solutions.

1.1 Classical Newsvendor Problem

We outline the classical Newsvendor problem under cost minimisation and service level frameworks as follows: consider the sale of a perishable item with uncertain demand (D) over a single period time frame, in which we need to determine an order quantity $Q = \mu_D + SI$. Consisting of mean demand (μ_D) and a safety stock level (SI).

Solutions aim to create a balancing act between ordering too much safety stock and preventing out-of-stock (or stockout) situations, with penalties incurred for causing either scenario. Under a cost minimisation framework, the penalties are expressed by an overage cost c_o per unit if we order too much product, and an underage cost c_u per unit if we fail to meet demand. The goal is to select the optimal order quantity such that it minimises the total expected costs, which is given by Equation (1).

$$\min_{Q\geq 0} \mathbb{E}\left[c_u (D-Q)^+ + c_o (Q-D)^+\right], \tag{1}$$

where $(x)^+ = \max(0, x)$.

While many models focus on the objective of cost minimisation (e.g. Levi et al. 2007, 2015, Ban & Rudin 2019, Huber et al. 2019, Ban et al. 2019), the problem can also be formulated under service level constraints (Beutel & Minner 2012, van der Laan et al. 2019), which are used in inventory management to measure performances of systems. Under a service level framework, one such performance measure is the non-stockout probability. Here, let $\alpha \in (0, 1)$ denote the probability of a stockout event, it then follows that $(1 - \alpha)$ yields the non-stockout probability. We simply wish to minimise expected excess inventory subject to the constraint that - by setting our inventory level Q - stochastic demand D is met with prescribed probability $(1 - \alpha)$ (van der Laan et al. 2019). An objective function for this is given in the following Equation (2),

$$\min_{Q\geq 0} \left\{\mathbb{E}\left[(Q-D)^+\right] : \mathbb{P}[Q \geq D] \geq 1-\alpha\right\}. \tag{2}$$

In an idealised scenario, the Newsvendor model assumes full knowledge of the demand distribution, with cumulative distribution function (cdf) F. For cost minimisation, the optimal inventory level Q^* is given by,

$$Q^*_{CM} = F^{-1}\left(\frac{c_u}{c_u + c_o}\right) = \min\left\{t : F(t) \geq \frac{c_u}{c_u + c_o}\right\}, \tag{3}$$

where F^{-1} is the inverse cdf. And under a service level framework,

$$Q^*_{SL} = F^{-1}(1-\alpha). \tag{4}$$

For further information on the basic problem, and derivation of the critical fractile under cost minimisation, one can refer to Arikan (2011). Equations (3) and (4) are equivalent, and thus we can set the values of c_o and c_u based on knowledge of $(1-\alpha)$, with the converse also being true. While the classical Newsvendor provides satisfying and simply derived results, it has been long accepted that the demand distribution being known is not realistic in practical settings. Due to these models being too simple and assumptious to capture the intricacies of real world data (Fisher 2009). There exists a wide range of potential solutions for when the demand distribution is unknown to help make more appropriate and reflective real-world decisions, most falling under a class of "data-driven" models.

1.2 Data-driven inventory management

Early literature that addressed the issues of not knowing the full demand distribution came from Scarf (1957) and extensions by Gallego & Moon (1993). They assumed only partial knowledge of the demand distribution, specifically the mean and standard deviation. They employed a min-max method which sought to maximise expected profit against the worst possible distribution. Gallego & Moon (1993) noted its conservative approach and ease of implementation, and offered extensions to scenarios in which secondary purchasing opportunities could be made; as well as exploring a multi-item case. A min-max approach is simple, easy to implement, and provides closed form results; however, optimising over the worst case distributions requires accurate estimation of the mean and variance (Bertsimas & Thiele 2005). While these parameters can be estimated from the observed data, in real world scenarios the potential limited set of past observations leads to unlikely and infeasible solutions.

The data-driven literature has since expanded in recent years to include a variety of approaches. These are usually split into two categories: (1) parametric models, where the distribution family is known but parameters need to be estimated. We cover this in Section 2.1. (2) Non-parametric models, which are free from distributional assumptions and encompass a wide array of methods, such as Sample Average Approximation (Levi et al. 2007, 2015), Robust Optimisation (Scarf 1957, Gallego & Moon 1993, Bertsimas & Thiele 2005, Wang et al. 2015) and Machine Learning (Huber et al. 2019, Ban & Rudin 2019). These are all discussed in Sections 2.2, 2.3 and 5.1 respectively.

As we move forward, it should be noted data-driven inventory management is interpreted differently amongst literature in the area. For example, the definition given by Bertsimas & Thiele (2005) says data-driven models possess the quality that 'we build directly upon the sample of available data instead of estimating the probability distributions', implying models falling under (1) would not be part of the data-driven paradigm. Alternatively, in a paper focusing on the comparison of parametric and non-parametric models, Ban et al. (2019) use a more literal sense of data-driven models, encompassing any model in which a decision maker chooses to base a decision on available historic data, which covers parameter estimation based on historic data. We use the latter's definition moving forward.

With regards to the steps taking to solve inventory level models, data-driven approaches are split into two tasks, *demand estimation* and *inventory optimisation*. These can be tackled sequentially or in an integrated fashion. Huber et al. (2019) and Ban & Rudin (2019) discussed the merits of a two-step and integrated approach, investigating in what situations one paradigm is better than another.

Another issue regarding the data-driven Newsvendor - heavily influenced by the plethora of new data available - are the effects external features have on demand. Particularly in retail settings, given many influences on the demand, and with access to this data, researchers have aimed to find ways to incorporate this into the model. The so called *feature-based* models usually extend from featureless counterparts. We cover this in detail in Section 3, before comparing these methods to purely demand driven approaches in a Numerical Study (Section 4). We first investigate the parametric and non-parametric featureless data-driven frameworks.

2 Data-driven Inventory models

In recent years, focus has primarily been on leveraging historic data to infer on making optimal inventory decisions. These *data-driven* models can split into into parametric and non-parametric approaches. Parametric models use data to estimate parameter values for what decision makers assume is the underlying distribution, in practicality these models are often made in misspecification (Ban et al. 2019). In contrast, non-parametric empirical models don't make any parametric assumptions on the demand, and are completely reliant on the data at hand and empirical estimates. Parallel to these are robust optimisation methods, which look to incorporate levels of robustness in the model. We explore parametric models from both a Frequentist and Bayesian setting, before investigating non-parametric empirical and robust approaches.

2.1 Parametric Models

Decision makers may find themselves in a situation in which the class of the demand distribution is known, or can be heavily inferred from the observations, but the parameters themselves need to be estimated. For such scenario, a parametric approach is preferable (also known as model-based), and these are broadly separated into Frequentist and Bayesian methods. For example, a Frequentist approach would be using maximum likelihood estimation to attain estimates. While, a Bayesian method assumes unknown parameters follows a prior distribution, and the posterior derived from this prior using data observations. The posterior for the demand distribution is then constructed based on the aforementioned posterior. We reserve the next two subsections to introduce both a Frequentist and Bayesian method, assuming access to N independent and identically distributed (i.i.d) historic demand observations $\mathbf{D}(N) = \{d_1, ..., d_N\}$.

2.1.1 Method of Moments

Assume our sample of historic observations for demand follows some distribution F with density function $f_D(x; \boldsymbol{\theta})$, where the parameters $\boldsymbol{\theta} = (\theta_1, ..., \theta_k)$ are unknown and need to be estimated. Define the true j^{th} moment of the distribution as a function of the parameters,

$$m_j = \mathbb{E}\left[D^j\right] \equiv g_j(\theta_1, ..., \theta_k), \tag{5}$$

and the sample j^{th} sample moment as,

$$\hat{m}_j = \frac{1}{N}\sum_{i=1}^{N} d_i^j. \tag{6}$$

In the *method of moments* (MM) approach, we assume the true and sample moments match and set $m_j = \hat{m}_j$ for $j = 1, ..., k$. This results in the set of simultaneous equations:

$$\frac{1}{N}\sum_{i=1}^{N} d_i = g_1(\hat{\theta}_1, ..., \hat{\theta}_k),$$
$$...$$
$$\frac{1}{N}\sum_{i=1}^{N} d_i^k = g_k(\hat{\theta}_1, ..., \hat{\theta}_k)$$

in which the solution, $\hat{\boldsymbol{\theta}}_{MM} = (\hat{\theta}_1, ..., \hat{\theta}_k)$ gives our parameter estimates. We then solve the inventory problem (Equation 1) and decide on an optimal decision rule $\hat{Q}_{MM}$. This is given by Equation (7),

$$\hat{Q}_{MM} = F^{-1}\left(\frac{c_u}{c_u + c_o} \,\middle|\, \hat{\boldsymbol{\theta}}_{MM}\right). \tag{7}$$

Alternatively, instead of MM, maximum likelihood estimation is another widely used frequentist parametric approach, or quasi-maximum likelihood estimation if the model is misspecified (Ban et al. 2019).

2.1.2 Bayesian Parametric

In a Bayesian approach, we begin by reflecting our beliefs in the unknown parameters $\boldsymbol{\theta}$ by specifying a prior distribution, $f_{\text{prior}}(\cdot)$. This is selected before any data is seen so it is up to the decision makers judgement or other factors that can influence a decision. Then, based on the data observations $\mathbf{D}(N)$, we also define a likelihood function $\mathcal{L}(\theta|\mathbf{D}(N))$ and find a posterior distribution through Bayes rule:

$$f_{\text{post}}(\boldsymbol{\theta}|\mathbf{D}(N)) = \frac{f_{\text{prior}}(\boldsymbol{\theta})\mathcal{L}(\boldsymbol{\theta}|\mathbf{D}(N))}{\int_{\Theta} f_{\text{prior}}(t)\mathcal{L}(t|\mathbf{D}(N))\text{d}t}, \tag{8}$$

where $\Theta \subseteq \mathbb{R}^k$ is the parameter space for the k parameters. The posterior for the demand distribution is derived from the posterior predictive distribution, i.e.,

$$f\left(x|\mathbf{D}(N)\right) = \int_{\Theta} \mathcal{L}(\boldsymbol{\theta}|x) f_{\text{post}}(\boldsymbol{\theta}|\mathbf{D}(N))\text{d}\boldsymbol{\theta}. \tag{9}$$

This Equation (9) then forms the density of the distribution F. And we solve the Newsvendor model and obtain the optimal order quantity,

$$\hat{Q}_B = F^{-1}\left(\frac{c_u}{c_u + c_o}\right). \tag{10}$$

Parametric models provide simple, closed form estimates for optimal order quantities in the face of demand uncertainty. They tackle this in a two stage problem of first forecasting demand through parameter estimation, and then solving an inventory problem akin to the classical Newsvendor problem. There are several issues with this approach however. First, obvious caveats are with the need of some input from the decision maker on distributional assumptions, if a decision maker was faced with total uncertainty on how the historical observations came about, the resulting models under these frameworks will likely be poor fits; they are not truly 'data-driven'. Frequentist methods like the method of moments can suffer from significantly higher inventory levels than needed (Beutel & Minner 2012). While Bayesian methods have proven hard to parsimoniously update prior distributions (Levi et al. 2007).

2.2 Sample Average Approximation

For the reasons discussed at the end of Section 2.1, alongside other scenarios such as the underlying distribution being too difficult to infer, or we have access to large swathes of data in which we can construct effective empirical distributions, a non-parametric approach may be preferred. Non-parametric models can be purely data-driven and require no inference or knowledge of demand distributions, leading to versatile implementations and utilising powerful modern methods; like those from Machine Learning. In the last 20 years, there has been a significant rise in non-parametric approaches, with the recent review article by de Castro Moraes & Yuan (2021) compiling these endeavours.

A simple and widely implemented non-parametric model is *Sample Average Approximation* (SAA). The method is used for stochastic optimisation problems and only requires samples from a distribution either by means of Monte Carlo methods or historical data. For the purpose of the Newsvendor the latter option is taken, and for a tutorial on the topic, see Kim et al. (2015).

Suppose we do not have knowledge of the underlying demand distribution but we do have access to a set of N historic demand observations, $\mathbf{D} = \{d_1, ..., d_N\}$. We say our historic demand are samples from the (unknown) true distribution, and each of these samples in $\mathbf{D}$ occurs with probability $\frac{1}{N}$. Under a cost minimisation framework, we replace the expectation seen earlier in Equation (1), with its weighted average across samples and new objective function given as,

$$\min_{Q \geq 0} \frac{1}{N} \sum_{i=1}^{N} \left(c_u \left(d_i - Q \right)^+ + c_o \left(Q - d_i \right)^+ \right). \tag{11}$$

Given we are dealing with assumed random samples, the optimal solution to the SAA problem for cost minimisation is in the form of a random variable $\hat{Q}_{SAA}$. We can see that each individual sample is analogous to the deterministic counterpart under a known distribution function, i.e., its solution will be the $\frac{c_u}{c_u+c_o}$ quantile on the individual sample. Hence, $\hat{Q}_{SAA}$ can be calculated by finding the $\frac{c_u}{c_u+c_o}$ quantile of the samples (Levi et al. 2007). Specifically,

$$\hat{Q}_{SAA} = \hat{F}^{-1}\left(\frac{c_u}{c_u + c_o}\right) = \min\left\{t : \hat{F}(t) \geq \frac{c_u}{c_u + c_o}\right\}, \tag{12}$$

in which $\hat{F}$ denotes the empirical cdf of the samples and is given by Equation (13),

$$\hat{F}(t) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}(d_i \leq t). \tag{13}$$

The SAA method was studied heavily in the context of the Newsvendor by Levi et al. (2007, 2015). Several theoretical guarantees were proven including solutions converging to the optimal value as the number of observations reach infinity. Another important finding were bounds on the observations needed to replicate results when the demand is fully realised (using a cumulative regret metric). These were initially investigated for general demand distributions in Levi et al. (2007), resulting in uninformative and conservative results. Subsequently, Levi et al. (2015) tightened the bounds by making the problem distribution specific.

2.3 Robust Optimisation

Robust optimisation is an alternative to the methods covered thus far, and a growing presence in the literature. Given unpredictability in the demand, decision makers have incentive to turn to robust solutions to handle problems with uncertain data (Ben-Tal et al. 2009). Robust optimisation, in which we have no knowledge of the distribution, is a methodology that takes a worst-case approach, leading to conservative estimates (Huber et al. 2019). While this is often attributed as a criticism, modern methods offer ways in which we can reduce the conservative estimates - through tolerance parameters - and robust optimisation remains a popular approach for Newsvendor models. Robust optimisation frameworks are unified in the construction of an uncertainty set (denoted $\mathcal{U}$) of possible values that can be taken; the goal is then to optimise. The general form of a robust optimisation problem is given by,

$$\min_{q\in\mathcal{Q}} \max_{u\in\mathcal{U}} \left\{h(\mathbf{q},\mathbf{u}) : \mathbf{g}(\mathbf{q},\mathbf{u}) \leq 0,\ \forall u \in \mathcal{U}\right\}. \tag{14}$$

Breaking down the various components of Equation (14): $\mathcal{Q}$ denotes the decision space, $\mathcal{U} \subset \mathbb{R}^d$ is the uncertainty set for parameters $\mathbf{u}$, $h(\cdot)$ and $\mathbf{g}(\cdot)$ being the random cost function and the vector of random functions respectively. The only distributional knowledge we require in this setting is the support of the random vector $\mathbf{u}$. This contrasts stochastic optimisation problems, such as those seen with SAA in Equation (11), where the true distribution is based entirely on empirical measurements of the historical data. Robust models serve as an advantage over SAA type methods which are purely data-driven, as they can mitigate any adverse effects that can occur when taking the whole data set into account, such as the potential for misguided decisions (Wang et al. 2015).

A modification of the Robust Optimisation, where we have a distribution P that we know lies in a family of distributions $\mathcal{P}$, provides an ambiguous variation of the model and given as,

$$\min_{q\in\mathcal{Q}} \max_{P\in\mathcal{P}} \left\{\mathbb{E}\left[h(\mathbf{q},\tilde{\mathbf{u}})\right]\right\}. \tag{15}$$

In the above, $\mathcal{P}$ is regarded as the ambiguity set and is induced by the random vector $\tilde{\mathbf{u}}$ (Rahimian & Mehrotra 2019). Special cases of this robust optimisation problem have been implemented in the literature. The work of Scarf (1957) and the distribution free Newsvendor approach - in which we have access to or estimate the first two moments - is one such example. The method falls under the category of *distributionally robust optimisation* (DRO), with Equation (15) being a generalisation of such optimisation set up. Non parametric Data-driven models with robustness came into prominence with a paper by Bertsimas & Thiele (2005), which removed any need to estimate parameters. The method incorporated techniques such as trimming to reduce the significance of outliers, the trimming factor allowed robustness in determining how conservative the decision maker wants the estimate to be. The result was a closed form solution, and an optimal ordering quantity corresponding to an empirical quantile of the data observations.

More recently, a modified version of DRO, *likelihood robust optimisation* (LRO), was introduced and applied to Newsvendor contexts by Wang et al. (2015). The authors modified the set of possible distributions, $\mathcal{P}$, to only contain those which reach a certain level of likelihood with the data available. We look at both the method of Scarf (also known as the min-max method) and LRO in the reminder of this sub-section.

2.3.1 Scarf's min-max method

The DRO approach of Scarf (1957) looked to maximise profit over worst case distributions. Analogously, under cost minimisation we seek to minimise worst case expected costs by considering worst case distributions. The method also requires the first two moments, but no other distribution assumptions. In a data-driven scenario, we calculate estimates for the parameters $\hat{\mu}$ and $\hat{\sigma}$ from the set of data observations $\mathbf{D}$. Collectively, we have in Equation (16) an expression for the min-max problem:

$$\min_{Q \geq 0} \quad \max_{D \in \mathcal{D}(\hat{\mu}, \hat{\sigma})} \mathbb{E}\left[c_u (D - Q)^+ + c_o (Q - D)^+\right], \tag{16}$$

where the function $\mathcal{D}(\mu, \sigma)$ is the set of distributions with mean μ and standard deviation σ. One should be able to see Equation (16) is just a simplified case of Equation (15). For this problem, the optimal order quantity, $\hat{Q}_{SF}$ is given by,

$$\hat{Q}_{SF} = \hat{\mu} + \frac{\hat{\sigma}}{2}\left(\sqrt{\frac{c_u}{c_o}} - \sqrt{\frac{c_o}{c_u}}\right). \tag{17}$$

This is known as *Scarf's rule*, and for an elegant proof the reader is referred to Gallego & Moon (1993).

2.3.2 Likelihood Robust Optimisation

Unsatisfied with the two moment DRO approach such as Scarf's method, Wang et al. (2015) raised two key issues. First, the discarding of important information within the data by simply taking into account just the mean and variance. Second, taking a worst-case approach caused suffering in the performance of more likely scenarios due to the methods inherit conservative nature. In response, they proposed to redefine how a distribution set is selected, using the likelihood function over moment information.

Consider a set of historical demand observations $\mathbf{D}$, where we round each observation to the nearest integer. The support of the underlying, but unknown, demand distribution is given by $\mathcal{S} = \{1, .., , n\}$. We have the opportunity to truncate n if we would like to save computational time or remove possible anomalous demand data. Denote N_i as the number of observations, from the set $\mathbf{D}$, which is equal to i, it then follows that the total number of observations, $N = \sum_{i=1}^{n} N_i$. We then construct a distribution set, $\mathbb{D}$, which is much like the family of distributions $\mathcal{P}$ seen earlier. This is given by Equation (18),

$$\mathbb{D}(\gamma) = \left\{\mathbf{p} = (p_1, ..., p_n) : \sum_{i=1}^{n} N_i \log(p_i) \geq \gamma,\ \sum_{i=1}^{n} p_i = 1,\ p_i \geq 0,\ \forall i\right\}. \tag{18}$$

The parameter γ governs the threshold for some desired level of likelihood to be reached by a potential distribution, and the first constraint in the set $\mathcal{D}(\gamma)$ ensures the worst case distribution is selected.

Regarding selection for γ, Wang et al. (2015) proposed the following heuristic:

$$\gamma^* = \sum_{i=1}^{n} N_i \log\left(\frac{N_i}{N}\right) - \frac{1}{2}\chi^2_{n-1,1-\beta}, \tag{19}$$

where χ^2 is the chi-squared distribution with $(n-1)$ degrees of freedom, and β selected such that we cover the true distribution with probability $(1-\beta)\%$.

For the Newsvendor problem under cost minimisation, by combining Equation (18), with the general structure of a DRO problem in Equation (15), we have the two-stage deterministic problem as follows,

$$\begin{aligned} \min_{Q\geq 0} \ \max_{\mathbf{p}} & \sum_{i=1}^{n} p_i \left(c_u(d_i - Q)^+ + c_o(Q - d_i)^+\right) \\ \text{s.t.} \ & \sum_{i=1}^{n} N_i \log(p_i) \geq \gamma, \\ & \sum_{i=1}^{n} p_i = 1, \\ & p_i \geq 0, \ \forall i. \end{aligned} \tag{20}$$

For the objective function, we replace the expectation with a weighted sum governed by the probabilities from a distribution set. The inner problem maximises the worst case distribution and cost, while the outer problem calculates the stocking quantity. For implementation, a tractable single stage optimisation problem can be formed from Equation (20) as,

$$\begin{aligned} \min_{Q,\lambda_1,\lambda_2,\mathbf{y}} \ & \lambda_2 + \lambda_1 \left(\sum_{i=1}^{n} N_i \log N_i - N - \gamma\right) + N\lambda_1 \log(\lambda_1) - \sum_{i=1}^{n} N_i \lambda_1 \log(y_i) \\ \text{s.t.} \ & c_u(d_i - Q) + y_i \leq \lambda_2 \ \ \forall i, \\ & c_o(Q - d_i) + y_i \leq \lambda_2 \ \ \forall i, \\ & \lambda_1 \geq 0, \ \mathbf{y} \geq 0. \end{aligned} \tag{21}$$

The model in Equation (21) is then solved to obtain an optimal stocking quantity, $\hat{Q}_{LRO}$.

We can also modify the problem to make sure all probability distributions selected carry the same mean and variance and the historical observations. To achieve this, we add two constraints to the model, namely,

$$\begin{aligned} & \sum_{i=1}^{n} p_i d_i = \hat{\mu}, \\ & \sum_{i=1}^{n} p_i d_i^2 = \hat{\mu}^2 + \hat{\sigma}^2, \end{aligned} \tag{22}$$

for the sample mean $\hat{\mu}$ and sample variance $\hat{\sigma}^2$. Fixing the mean and variance was shown to perform better than both the original LRO variant and Scarf's min-max method in the case of underlying Normal and Exponential distributions (Wang et al. 2015). They also found in the case of Scarf's method when looking at an asymmetrical distribution such as the Exponential, performance issues suffered greatly in comparison to LRO methods; the reasoning was with only taking two moments, the skew in distribution could not be incorporated into the model. Another recent adaptation of DRO is Functionally Robust Optimisation, introduced by Hu et al. (2019). Here, the uncertainty set consists of non-parametric demand functions and does not assume that the form of a model function must be predetermined. They allow a decision maker to modify their risk appetite based on a risk-reward tradeoff by tuning model parameters.

3 Feature-based Newsvendor

Thus far, the methods discussed are all featureless, and based purely off taking demand observations at face value. This may not be a realistic assumption to make in certain contexts. For example, within retail, focusing just on historic demand observations has lead to unsatisfactory analysis when forecasting demand (Beutel & Minner 2012), this being due to strong dependency on external factors. Recalling our Bakery motivation from the introduction, we may find ourselves with a higher demand for seasonal products such as hot-cross buns during the Easter period, surges in general demand during holidays and weekends, or store location and weather affecting customers willingness to travel in certain conditions. Thus, some data-driven approaches consider how these exogenous variables can be incorporated into the decision making process.

The general framework for a feature based Newsvendor readily extends its featureless counterpart. We define a feature vector $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^m$, where $\mathcal{X}$ is the feature space. We assume the decision maker has obtained a set of historical data, $S_n = \{(d_1, \mathbf{x}_1), ..., (d_N, \mathbf{x}_N)\}$, and the demand decision is now a function $Q(\cdot) : \mathcal{X} \to \mathbb{R}$. The cost minimisation formulation of the Newsvendor is modified to optimise the conditional expected cost function, ie.,

$$\min_{Q(\cdot) \geq 0} \ \mathbb{E}\left[c_u \left(D(\mathbf{x}) - Q(\mathbf{x})\right)^+ + c_o \left(Q(\mathbf{x}) - D(\mathbf{x})\right)^+ |\mathbf{x}\right]. \tag{23}$$

The decision space $\mathcal{Q}$, in which $Q(\cdot)$ is based in, can be selected in different ways. A linear relationship (Beutel & Minner 2012, van der Laan et al. 2019, Ban & Rudin 2019) is a popular approach and constructed as:

$$\mathcal{Q} = \left\{ Q : \mathcal{X} \to \mathbb{R} : Q(\mathbf{x}) = \mathbf{q}^T \mathbf{x} = \sum_{j=1}^{m} q_j x_j \right\}. \tag{24}$$

We then optimise to find these $\mathbf{q}$ coefficients. Huber et al. (2019) notes a linearity assumption causes restrictions on the underlying functional relationships, however, if the relationship is too complex then this could lead to overfitting. With this being said, non-linear relationships have been constructed in linear models through incorporating them as additional features (e.g. van der Laan et al. 2019, Ban & Rudin 2019). While a purely non-linear decision space was suggested by Huber et al. (2019).

We formulate one such feature based approach in the remainder of this section. Motivated by the inefficiencies of inventory planning within the retail sector, Beutel & Minner (2012) suggested the integration of causal demand forecasting and inventory optimisation. They focus primarily on the issue of safety stocks, and use a Econometrics approach based on ordinary least squares (OLS) to model the relationship between features and demand. We formulate demand in this way as follows: Consider historical demand observations within $\mathbf{D} = \{d_1, ..., d_n\}$ as a linear function of m explanatory variables and an error term ($\mathbf{u}$). In matrix notation this is given by,

$$\mathbf{D} = \mathbf{Xq} + \mathbf{u}, \tag{25}$$

where $\mathbf{q}$ are the coefficients for m explanatory variables, and $\mathbf{X}$ a $m \times n$ feature matrix of observations. We set the first row of the matrix, $X_{0,i}$ equal to 1 to allow for an intercept term in the model. For a specific observation, i, we have,

$$d_i = q_0 + \sum_{j=1}^{m} q_j X_{j,i} + u_i. \tag{26}$$

The optimal order quantity, Q, is derived from the feature vector with known value, $X_{j,0}$, e.g., the next days sales price or weather. The order quantity is given in Equation (27).

$$Q = q_0 + \sum_{j=1}^{m} q_j X_{j,0}. \tag{27}$$

The feature dependent linear program under cost minimisation seeks to optimise decision variables, q_i for each observation i. Inventory levels (v_i) and satisfied demands (s_i) are also indirectly decision variables. Historic demand d_i and feature matrix $X_{j,i}$ for $j = 1, .., m$ are known quantities. The resulting formulation for the linear program is given by Equation set (28) to (32) as follows,

$$\min_{\mathbf{q},\mathbf{v},\mathbf{s}} \quad \sum_{i=1}^{n} (c_o v_i + c_u(d_i - s_i)) \tag{28}$$

$$\text{s.t.} \quad v_i \geq \sum_{j=0}^{m} q_j X_{ji} - d_i, \quad i = 1, ..., n \tag{29}$$

$$s_i \leq d_i, \quad i = 1, ..., n \tag{30}$$

$$s_i \leq \sum_{j=0}^{m} q_j X_{ji} \quad i = 1, ..., n \tag{31}$$

$$\mathbf{s}, \mathbf{v} \geq 0, \mathbf{q} \in \mathbb{R}. \tag{32}$$

The objective function (28) sums the holding and penalty costs for each demand observation, whether we have a surplus with leftover inventory v_i, or shortage $(d_i - s_i)$. The constraint (29), alongside the objective function, determines excess inventory v_i corresponding to demand d_i. While constraints (30) and (31) ensure the sales quantity, s_i, is equal to the minimum of the demand observation, d_i, and the supply.

This problem is tractable and can be easily implemented, with the resulting coefficient values substituted into Equation (27) to obtain the optimal order quantity $\hat{Q}_{FB}$. In other literature, the formulation is referred to as Empirical Risk Minimisation (Ban & Rudin 2019), or the hindsight approach when formulating under service level constraints (van der Laan et al. 2019). The method was found to under perform in comparison to the parametric method of moments, and another method based on ordinary least squares, when the sample size was small. However under model mis-specification, they found their approach greatly outperformed comparable methods.

4 Numerical Study

To compare a selection of featureless methods to a feature method, and computationally analyse the findings we've discussed thus far in the literature, we consider an example in which the demand is expressed a function of the price. Specifically we base the study off Beutel & Minner (2012) and van der Laan et al. (2019), quantifying the methods based on achieved service levels and average inventory levels. Elsewhere, researchers have used a regret metric (Levi et al. 2007, 2015), average cost increases (Huber et al. 2019), and expected profit/cost (Bertsimas & Thiele 2005, Wang et al. 2015, Ban et al. 2019) to evaluate methods. We select our metrics in particular as they have relevant applicability to what real world decision makers are likely looking for in the retail sector. We have generated the data ourselves, with code found in the **Data Availability** section.

4.1 Setup

We assume the true demand depends on a single factor, price, and each historic demand observation is modelled by,

$$d_i = a - bp_i + \eta, \tag{33}$$

for price $p_i \sim U[0,1]$, market size and slope a and b respectively, with $\eta \sim \mathcal{N}(0, \sigma)$. a and b are selected uniformly on intervals $[1000, 2000]$ and $[300, 500]$ respectively. For each instance of a problem, a and b stay fixed and are unknown to the decision maker. To select σ, we recover this from the coefficient of variation formula at mean price p,

$$CV = \frac{\sigma}{\mu_D},$$

for mean demand μ_D. We fix $CV = 0.3$ for the majority of the study, before varying this later in Section 4.2.2. For the underage cost c_u and overage cost c_o, recall from Section 1.1 that we can recover these from setting a desired service level, given by Equation (34),

$$\frac{c_u}{c_u + c_o} = 1 - \alpha. \tag{34}$$

If say, we fix the service level for the non-stockout probability to $(1 - \alpha) \times 100\% = 90\%$ and the overage cost as $c_o = 1$, then we recover $c_u = 9$.

A table of selected methods is given by Table 1. We aim to compare a method from each of the family of approaches we've explored in this report. For the method of moments, we assume an underlying normal distribution. While for the known true function, we use:

$$Q_{true} = \mu_D + SI. \tag{35}$$

Here, $a - b \cdot 0.5$ is the mean demand (given the mean price is 0.5 if picked uniformly on $[0, 1]$) and $F^{-1}\left(\frac{c_u}{c_u + c_o}\right) \cdot \sigma$ is the safety stock level. In this scenario, the CDF is standard normal.

Featureless	Method of Moments (parametric)	Section 2.1.1
	SAA (non-parametric empirical)	Section 2.2
	Scarf's Rule (non-parametric robust)	Section 2.3.1
Feature Based	Feature based LP	Section 3
	Known	Equation (35)

Table 1: Methodologies.

Each method is trained on n samples, then the metrics are created by testing on 100,000 out of sample observations. We calculate the metrics as so,

1. Average service levels: For the non-stockout probability service level, we find the proportion of out of sample observations that are below the optimal decision rule.

2. Average inventory levels: For each out of sample observation, we observe the demand and find the difference between the optimal decision rule and the demand, averaging over all the samples.

Given the initial set up, the methodologies are tested on varying training sample sizes and analysing low and high variance data.

4.2 Results

To add clarification to the numerical study, we run through a simple example of a single instance. For this run through on n=200 training samples:

$$a = 1348.217,\ b = 313.41,\ c_u = 1,\ c_o = 9,\ CV = 0.3, \text{ and } \sigma = 357.4536.$$

Figure 1a graphs each rule from the methods outlined in Table 1, with Figure 1b giving the average inventory levels after 100 iterations of testing and training on randomly generated demand samples. From this single instance we see immediately the conservative nature of Scarf's rule that was discussed in the literature. In fact, on this instance all (sans the feature based LP method) show conservative inventory levels, with SAA and Feature based LP performing the best. Method of Moments and Scarf's also appear to have tighter bounds on the spread of the inventory levels.

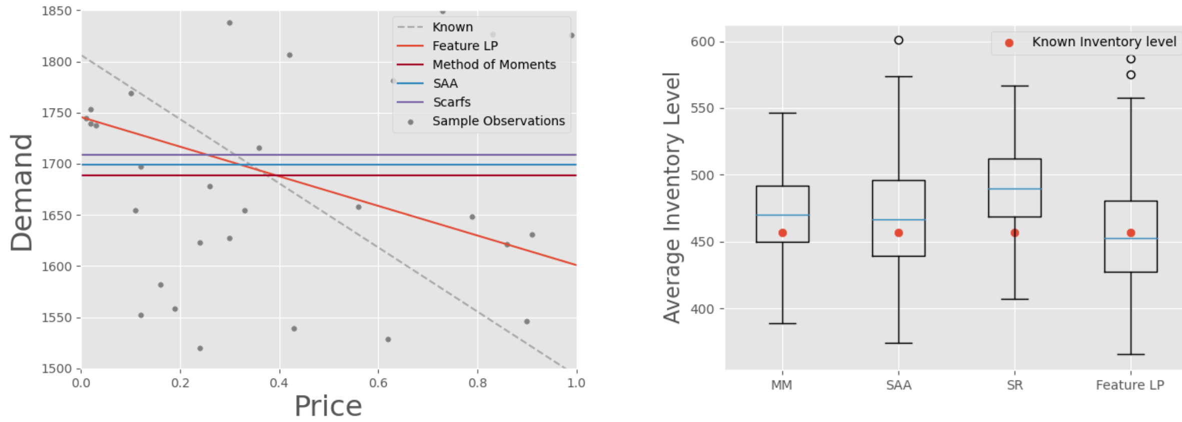

(a) Q decision rules (b) Average inventory Levels

Figure 1: Single instance

We can now test different situations and aspects a decision maker may be interested in evaluating. For the following tests, 100 instances are generated and results averaged out for each metric.

4.2.1 Sample size changes

Here, we fix $CV = 0.3,\ c_u = 1,\ c_o = 9$ and range the training sample size $n \in [10, 200]$ in intervals of 10. Results for service levels and average inventory levels are given in Figure 2.

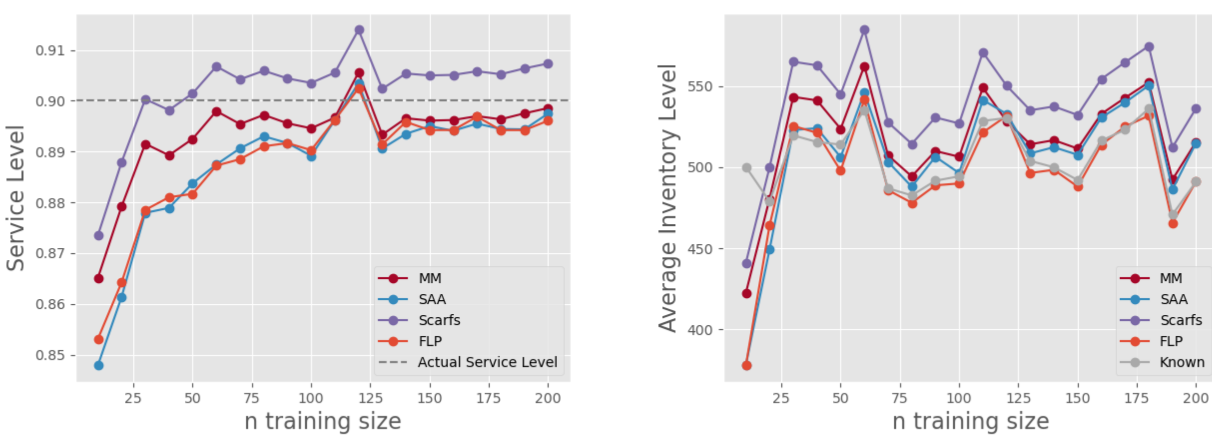

(a) Service Level (b) Average inventory Levels

Figure 2: n ranges

For service levels we again see Scarf's rule giving conservative results, however this leads to the best service levels for for low training sample sizes. Conversely, SAA and the Feature based LP are almost indistinguishable for the service levels, and both perform rather poorly for low training sample sizes. Note, there is a seemingly anomalous result for $n = 110$ which is likely as a result of not doing enough simulations. This could be alleviated with more computational power and time. Turning to average inventory levels in Figure 2b, all methods require a relatively low sample size to attain reasonably close to actual inventory levels (~ 20) before suffering quite poor results for Scarf and Method of Moments. The Feature based LP clearly has the best results and is found to be consistently within $1-3\%$ of the actual inventory level for $n > 20$.

4.2.2 Low and high variance data

We experiment with modifying the coefficient of variation between the values of $CV = 0.3$ and $CV = 0.5$ at 0.05 increments, with 200 training samples. Here, we just focus on the Average Inventory Level metric, as the service level did not change a meaningful amount when varying CV levels. Figure 3 gives plots of the Average Inventory Level and its standard deviations. In particular, Figure (3a) gives the difference between the specific method and value for the known function.

We see that the Feature Based LP is definitively better than other methods for all variations in the spread of the data. The difference is especially stark for low variance data. With this being said, the standard deviation of the Feature Based LP is the highest. Scarf's conservative estimates still shine through in this variation.

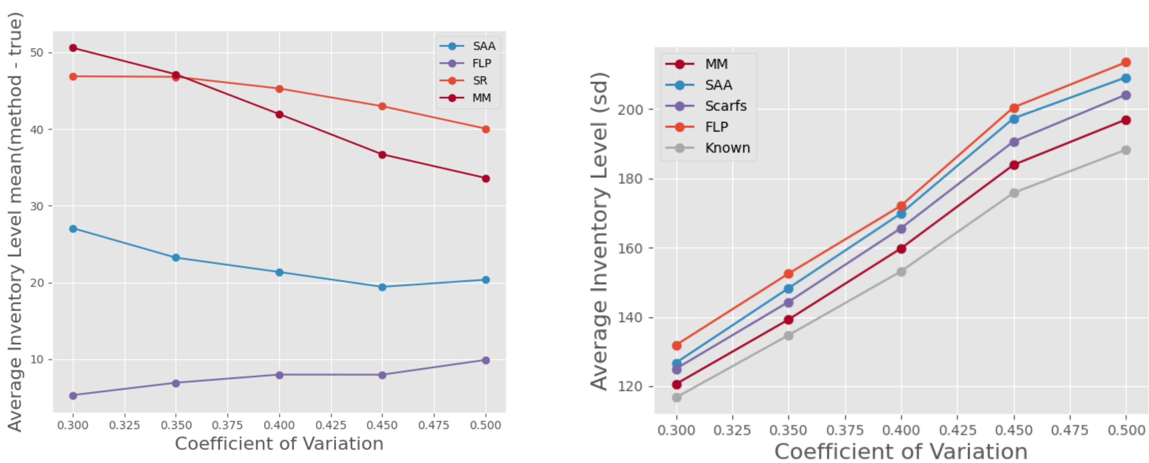

(a) difference in Average Inventory Level between true and each respective method.

(b) Plot of standard deviations of Average Inventory Level.

Figure 3: CV changes

4.3 Discussion

These results aim to highlight the merits and drawbacks with different approaches to the data-driven Newsvendor problem. They also highlighted the advantages of using more than one metric to assess results. As an example, looking at inventory levels, the Feature based LP performed exceptionally, but lagged behind in service levels, and exhibited issues under small sample sizes. What we did not see however was a significant advantage of using the Feature based LP over featureless methodology. This could be due to the simple underlying assumptions of the data being normal. In future, we could consider modifying these underlying assumptions and investigating performance under model mis-specification. This could be especially insightful for the parametric method of moments that requires a family of distribution to be specified, and Scarf's rule if we have bimodal data. Some literature has already investigated this issue with Beutel & Minner (2012) considering the cases of heteroscedasticity (changing variance throughout the data) and a gamma demand distribution. Further research may also be needed in replicating results and independently repeating the study. This is due to the peculiarity in the performance between SAA and Feature Based LP, as they appear rather close to each other. Meanwhile, the method of moments appear to perform better here than in other computational studies (e.g. Beutel & Minner 2012).

5 Current and Future research areas

The data-driven Newsvendor is an active field, with regular developments in methodological, theoretical and computational aspects. We explore three active research areas of interest, looking at expansion of the toolset of data-driven methods, as well as the application of data-driven methods to these specific Newsvendor cases. Providing the current state of the art, and discussing future research directions in each respective area.

5.1 Machine Learning models

Contemporary developments in data-driven models heavily incorporate Machine Learning (ML) aspects, with a surge in ML approaches in recent years (de Castro Moraes & Yuan 2021). As a result of the rise in data available to retailers on all different facets of the inventory process, an obvious area of focus is the feature-based Newsvendor. Oroojlooyjadid et al. (2020) notes some of the state-of-the-art methodologies employed by decision makers, while themselves suggesting a Deep Neural Network based approach. One such topic is combining quantile regression (QR) with machine learning approaches such as neural networks; which have been used previously on predicting rainfall, drug activities and evaluating value-at-risk. Given the solution to Newsvendor problems often involve selecting a certain quantile of the cumulative demand distribution, it is natural to employ quantile regression methods, and ML-QR methods have been implemented by Huber et al. (2019), Ban & Rudin (2019), Oroojlooyjadid et al. (2020), Cao & Shen (2019) among others. The systematic literature review by de Castro Moraes & Yuan (2021) notes that their is work still to be done in QR-ML models, specifically in regards to modifying the loss function when considering various model constraints.

Elsewhere, model paradigms of separated and integrated demand forecasting and inventory optimisation (recall Section 1.2) have been subject of interest within ML models. Two recent papers to outline the merits and drawbacks of the approaches were Ban & Rudin (2019) and Huber et al. (2019). Ban & Rudin (2019) implemented one-step, distribution free machine learning algorithms, and commented on the reasoning behind not employing a separated approach; citing the problematic nature of demand model specification in high-dimensions. This has the effect of amplifying errors when moving into an optimisation stage. Alternatively, Huber et al. (2019) says the situations in which a separated or integrated approach is superior, remains an open problem. The paper considers both separate and integrated methodology and employs ML approaches in the demand estimation stage such as Neural Networks and Decision Trees. The also note the equivalence of an integrated approach and QR, and conclude that ML methods perform very well as long as sufficiently large datasets are available.

A modern problem with ML models is the lack of interpretability. Since these approaches are black-box, it is hard for the decision maker to derive how an algorithm achieves a given output from the original data (Oroojlooyjadid et al. 2020). Consequently this brings up issues of trust and transparency in the models (Lipton 2016). One argument, highlighted by Huber et al. (2019), posits the idea that if the metrics used to evaluate methods are good enough, then these should outweigh potential issues with interpretability. They use an example of forecast accuracy measurements being easy to attain in baked goods as justification. The issue, as well as being spread out broadly within ML literature, is also likely problem specific and more research is needed to evaluate the merits of black-box model, or alternatively develop more intepreatable models.

5.2 Service Level Constraints

In the introduction to the Newsvendor model at the beginning of the report, we briefly formulated the model under service level constraints. Often retailers are bounded by "Service level agreements" which are performance based contracts that manage supplier relationships; with some 91% of organisations using service level agreements of some sort (Liang & Atkins 2013). Therefore, given its wide use in practice, optimisation with respect to service level constraints should perhaps be more widely considered than it currently is. Very few models assess the problem under service level constraints, with most preferring the classical cost minimisation or profit maximisation. In addition to the service level of non-stockout probability that was used in both the introduction and Section 4, another popular service level is the fill rate. This pertains to satisfying a fraction of demand, which had been pre-determined.

In the feature based Newsvendor literature, there exists only three current papers on the issue: Beutel & Minner (2012), van der Laan et al. (2019), Ye & Yang (2021). The first of these is a modification on the Feature Based LP we introduced in Section 3. While, van der Laan et al. (2019) use similar assumptions but take a robust optimisation approach. They observe service level constraints are chance constraints, a special robust case approach of stochastic optimisation (see Rahimian & Mehrotra 2019, for more details). In the third model, Ye & Yang (2021) highlights the conservative results from the robust optimisation methods, and themselves produce a Machine Learning approach using K-nearest Neighbours.

Given the lack of papers, an exciting research path would be to further explore and apply methods to a service level framework. van der Laan et al. (2019) proposes applying their robust approaches to a fill rate service level, as they solely focused on the non-stockout probability. While SAA and parametric models could also be applied to the service level constraint problem.

5.3 Censored demand

If a retailer faces a stockout situation, they may not be able to estimate the lost sales faced. This in turn leads to right-censored data, with the only knowledge of the system being demand exceeded the inventory. Nahmias (1994) discusses the long history of investigating the censored demand problem, with parameter estimation of censored data going back to the early 20^{th} century. They suggest a method for parameter estimation with normally distributed demand, however focusing on order-up-to models rather than the Newsvendor. In order-up-to models, the decision maker periodically reviews data and updates stock back to a desired demand level. Ding et al. (2002) developed a Bayesian Markov decision process for the multi-period Newsvendor, finding performance improvements over Bayesian methods at the time, as well as upper bounds on performance. Lau & Hing-Ling Lau (1996) proposed an early non-parametric approach to the censored demand problem.

Sachs & Minner (2014) extends Beutel & Minner (2012) by considering demand censoring in the feature based Newsvendor context. They provided a data-driven approach that was found to perform better than the existing parametric and non-parametric models, including that of Nahmias (1994) and Lau & Hing-Ling Lau (1996). They suggest extensions to the multi-period case, noting product substitutions in the event of a stockout leads to artificially inflated sales of other products. Meanwhile, Huber et al. (2019) attempts to circumvent the issue of lost sales and demand censoring by using intra-day sales patterns of point-of-sales data. Given the abundance of data available to decision makers, this appears to be a preferable direction.

In general, much like service-level constraints, the data-driven literature is rather sparse on the issue of demand censoring. Further research outside of that mentioned so far would likely be extending current Newsvendor approaches to a censored demand setting; much in the same vein of Sachs & Minner (2014).

6 Conclusion

This paper introduced a framework to select the inventory quantity of perishable products under uncertain demand, with specific focus on developing methods leveraging purely historical data observations. By drawing comparisons between perishable items and the issues faced by Newsstands, the Newsvendor model was introduced. The classical method utilised a known demand distribution, which is an impractical assumption. Thus an array of data-driven methodologies were discussed and compared, as well as the current state of the art in the field and contemporary research area. The primary application - and focus of this paper - was in the retail sector, however due to the generality of the framework, researchers have extended the model beyond this context. Arikan (2011) cites fast fashion, consumer electronics, and revenue management. The methods developed for the Newsvendor setting have also been adapted to other contexts. For example, Wang et al. (2015) applied their LRO method to the portfolio Selection Problem in addition to the Newsvendor.

Although we introduced methods with emphasis on the practical and real world situations, some liberties have still been taken that prevents fully realistic scenarios. One direction with pertinent applications to the real situations faced by retailers is multi-period and multi-item settings. The literature provide many instances of having a desire in their conclusions to extend the work to such situations. Future work could discuss and investigate these methods in both theoretical and computational studies.

Another extension of this report is further analysis into where each approach is most relevant. It is not yet known if non-parametric data-driven methods outperform their parametric counterparts (de Castro Moraes & Yuan 2021). Thus, more computational and theoretical studies will be required to offer more insights into this question.

Data Availability

Code used in this report can be accessed at the following GitHub link: https://github.com/BenSLowery/MResCode/tree/main/601/RT2.

References

Arikan, E. (2011), *A Review of the Newsvendor Model*, Peter Lang AG, pp. 21–32.

Arrow, K. J., Harris, T. & Marschak, J. (1951), 'Optimal inventory policy', *Econometrica* **19**(3), 250–272.

Ban, G., Gao, Z. & Taigel, F. (2019), 'Model mis-specification in newsvendor decisions: A comparison of frequentist parametric, bayesian parametric and nonparametric approaches', *SSRN Electronic Journal* .

Ban, G.-Y. & Rudin, C. (2019), 'The big data newsvendor: Practical insights from machine learning', *Operations Research* **67**(1), 90–108.

Ben-Tal, A., El Ghaoui, L. & Nemirovski, A. S. (2009), *Robust optimization*, Princeton series in applied mathematics, Princeton University Press, Princeton. OCLC: ocn318672208.

Bertsimas, D. & Thiele, A. (2005), A data-driven approach to newsvendor problems, Technical report, Operations Research Center, MIT.

Beutel, A.-L. & Minner, S. (2012), 'Safety stock planning under causal demand forecasting', *International Journal of Production Economics* **140**(2), 637–645. Sixteenth internationalworkingseminaronproductioneconomics, Innsbruck, 2010.

Cao, Y. & Shen, Z.-J. M. (2019), 'Quantile forecasting and data-driven inventory management under nonstationary demand', *Operations Research Letters* **47**(6), 465–472.

de Castro Moraes, T. & Yuan, X.-M. (2021), Data-driven solutions for the newsvendor problem: A systematic literature review, *in* 'Advances in Production Management Systems. Artificial Intelligence for Sustainable and Resilient Production Systems', Springer International Publishing, pp. 149–158.

Ding, X., Puterman, M. L. & Bisi, A. (2002), 'The censored newsvendor and the optimal acquisition of information', *Operations Research* **50**(3), 517–527.

Edgeworth, F. Y. (1888), 'The mathematical theory of banking', *Journal of the Royal Statistical Society* **51**(1), 113–127.

Fisher, M. (2009), 'Rocket science retailing: The 2006 philip mccord morse lecture', *Operations Research* **57**(3), 527–540.

Gallego, G. & Moon, I. (1993), 'The distribution free newsboy problem: Review and extensions', *The Journal of the Operational Research Society* **44**(8), 825–834.

Hu, J., Li, J. & Mehrotra, S. (2019), 'A data-driven functionally robust approach for simultaneous pricing and order quantity decisions with unknown demand function', *Operations Research* **67**(6), 1564–1585.

Huber, J., Müller, S., Fleischmann, M. & Stuckenschmidt, H. (2019), 'A data-driven newsvendor problem: From data to decision', *European Journal of Operational Research* **278**(3), 904–915.

Kim, S., Pasupathy, R. & Henderson, S. G. (2015), *A Guide to Sample Average Approximation*, Springer New York, New York, NY, pp. 207–243.

Lau, H.-S. & Hing-Ling Lau, A. (1996), 'Estimating the demand distributions of single-period items having frequent stockouts', *European Journal of Operational Research* **92**(2), 254–265.

Levi, R., Perakis, G. & Uichanco, J. (2015), 'The data-driven newsvendor problem: New bounds and insights', *Operations Research* **63**(6), 1294–1306.

Levi, R., Roundy, R. O. & Shmoys, D. B. (2007), 'Provably near-optimal sampling-based policies for stochastic inventory control models', *Mathematics of Operations Research* **32**(4), 821–839.

Liang, L. & Atkins, D. (2013), 'Designing service level agreements for inventory management', *Production and Operations Management* **22**(5), 1103–1117.

Lipton, Z. C. (2016), 'The mythos of model interpretability'.

Nahmias, S. (1994), 'Demand estimation in lost sales inventory systems', *Naval Research Logistics (NRL)* **41**(6), 739–757.

Oroojlooyjadid, A., Snyder, L. V. & Takáč, M. (2020), 'Applying deep learning to the newsvendor problem', *IISE Transactions* **52**(4), 444–463.

Poore, J. & Nemecek, T. (2018), 'Reducing food’s environmental impacts through producers and consumers', *Science* **360**(6392), 987–992.

Rahimian, H. & Mehrotra, S. (2019), 'Distributionally robust optimization: A review'.

Sachs, A.-L. & Minner, S. (2014), 'The data-driven newsvendor with censored demand observations', *International Journal of Production Economics* **149**, 28–36. The Economics of Industrial Production.

Scarf, H. E. (1957), *A Min-Max Solution of an Inventory Problem*, RAND Corporation, Santa Monica, CA, pp. 201–209.

van der Laan, N., Teunter, R., Romeijnders, W. & Kilic, O. (2019), The data-driven newsvendor problem: Achieving on-target service levels, Workingpaper, University of Groningen, SOM research school.

Wang, Z., Glynn, P. W. & Ye, Y. (2015), 'Likelihood robust optimization for data-driven problems', *Computational Management Science* **13**(2), 241–261.

WRAP (2017), Quantification of food surplus, waste and related materials in the grocery supply chain, Technical report, WRAP.

Ye, Y. & Yang, X. (2021), The New Data-Driven Newsvendor Problem with Service Level Constraint, *in* S. Liu, G. Bohács, X. Shi, X. Shang & A. Huang, eds, 'Liss 2020', Springer Books, Springer, pp. 1009–1023.