A Data-Driven Approach To Newsvendor Problems

Dimitris Bertsimas* Aurélie Thiele†

May 2005

Abstract

We propose an approach to the classical newsvendor problem and its extensions subject to uncertain demand that: (a) works directly with data, i.e., combines historical data and optimization in a single framework, (b) yields robust solutions and incorporates risk preferences using one scalar parameter, rather than utility functions, (c) allows for tractable formulations, specifically, linear programming problems, and (d) leads to closed-form solutions based on the ranking of the historical demands, which provide key insights into the role of the cost parameters. Numerical results are very encouraging.

1 Introduction

The question of maximizing profit in the presence of uncertain demand has long been central to revenue management. A classical example of such a setting is the newsvendor problem, which Porteus summarizes in [11]. As one of the building blocks of inventory theory, it has received much attention in the literature, often under the assumptions that the demand distribution is known exactly and that the decision-maker is risk-neutral. In practice however, the volatility of the demand for most perishable products makes it difficult to obtain accurate forecasts. In Scarf's words [14], "we may have reason to suspect that the future demand will come from a distribution which differs from that governing past history in an unpredictable way". This unpredictability provides a strong incentive for the decision-maker to implement robust solutions, which will perform well for a wide range of actual demand outcomes. The issue of imperfect

*Boeing Professor of Operations Research, Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, E53-363, Cambridge, MA 02139, dbertsim@mit.edu. Research partially supported by the Singapore-MIT Alliance.

†Department of Industrial and Systems Engineering, Lehigh University, Mohler Building Room 329, Bethlehem, PA 18015, aurelie.thiele@lehigh.edu. Research partially supported by the Singapore-MIT Alliance.

information has been addressed in the past by assuming that only the first two moments are known. In 1958, Scarf [14] derived the optimal ordering quantity for the classical newsboy problem with mean and variance given, and his work was later extended by Gallego et. al. [6, 9, 10]. However, such a worst-case methodology relies on the correct estimation of the first two moments and lacks a strong connection to risk preferences, which in practice play a key role in the choice of the solution implemented. Specifically, the decision-maker will often accept a smaller expected return if it also yields a decrease in the standard deviation of his revenue, i.e., he will prefer solutions which exhibit less variability, although riskier approaches can potentially bring higher profits. In this paper, we refer to the process of trading off a high value of the expected return for a decrease in the downside risk as incorporating robustness. The first attempt in the literature to model risk preferences in the newsboy problem is due to Lau [7], who considers two alternative criteria to the expected revenue: the expected utility and the probability of reaching a prespecified profit target. More recently, Eeckhoudt et. al. [5] have revisited the framework based on the expected utility of the newsvendor. In practice, however, it is difficult to articulate a particular individual's utility.

The approach that we propose here departs from these earlier frameworks in two major aspects: (a) it is entirely data-driven, in the sense that we build directly upon the sample of available data instead of estimating the probability distributions, and (b) it does not rely on utilities but rather on a scalar parameter to incorporate robustness in the model. This scalar parameter corresponds here to a prespecified quantile of the revenue. In this framework, the value of the random variable is determined by computing the expected revenue below that quantile, that is, by removing (trimming) the instances of the profit above the quantile and taking the average over the remaining ones. The fraction of data points removed will be referred to as the trimming factor. Hence, the decision-maker focuses on a more conservative valuation of his revenue than the one provided by an expected-value approach, but is able to adjust the degree of conservatism by selecting the trimming factor appropriately. Although two-sided trimming has been extensively studied in statistics, for instance by Rousseeuw and Leroy [12], Ryan [13], Wilcox [16], to develop robust estimators, one-sided trimming has received little attention outside the field of portfolio optimization, where it was studied by Bertsimas et. al. in [2]. Levy and Kroll have characterized in [8] investor preferences in terms of quantile functions, and most importantly have shown that a strategy chosen to maximize the one-sided trimmed revenue is non-dominated, i.e., no other strategy would be preferred by all decision-makers with increasing and concave utilities. An appealing feature of this approach is that it allows for nonparametric estimators and tractable formulations (see Bertsimas et. al. [2]).

Specifically, the contributions of this paper are as follows:

1. We present an approach that works directly with data, not probability distributions, and does not require the estimation of parameters.

2. We incorporate robustness through a single scalar parameter that allows the decision-maker to adjust the trade-off between risk and return.

3. The data-driven approach is tractable, as it leads to convex problems when the profit function is concave, and in particular to linear programming problems when the profit function is concave and piecewise linear.

4. We develop data-driven models in a single-period setting under a wide range of assumptions and extend the framework to the finite-horizon case, i.e., the approach is generally applicable to many inventory problems.

5. We derive theoretical insights into the optimal solution and show that, under some relatively mild conditions, the optimal order is equal to a historical realization of the demand. Its index depends on the cost parameters and the degree of conservatism chosen by the decision-maker.

Structure of the paper: In Section 2, we introduce the general framework of data-driven optimization. We apply this technique to the classical newsvendor problem in Section 3, and consider several extensions in Section 4. We present computational experiments in Section 5. Finally, Section 6 contains some concluding remarks.

2 The Data-Driven Framework

In this section, we describe the data-driven framework, first in a general setting (Section 2.1) and then in the context of revenue management (Section 2.2).

2.1 Preferences and One-Sided Trimming

Let $q_a(X)$ be the a-quantile of the random variable X, i.e., $q_a(X) = \inf\{x | P(X \leq x) \geq a\}$, for any $a \in (0,1)$. Levy and Kroll have characterized in [8] investor preferences in terms of the quantile functions of their returns.

Theorem 2.1 (Levy and Kroll [8]) *$E[U(X)] \geq E[U(Y)]$ for all U increasing and concave if and only if $E[X|X \leq q_a(X)] \geq E[Y|Y \leq q_a(Y)]$ for any a in $(0,1)$, and we have strict inequality for some a.*

Therefore, a strategy chosen to maximize the tail conditional expectation $E[X|X \leq q_a(X)]$ is non-dominated. Equivalently, maximizing $E[X|X \leq q_a(X)]$ for a specific a guarantees that no other strategy can improve the value (expected utility) of the random variable for *all* risk-averse decision-makers. The tail conditional expectation $E[X|X \leq q_a(X)]$ has recently received much attention in the literature, because it satisfies properties that Artzner et. al. [1] have identified as particularly desirable for measures of risk. Risk measures verifying this set of axioms are said to be coherent. In contrast, many better-known measures such as standard deviation violate at least one of the axioms presented by Artzner et. al. [1].

Another appealing feature of this approach is that $E[X|X \leq q_a(X)]$ can be estimated non-parametrically, that is, without making any assumptions on the underlying probability distributions. In mathematical terms, let N be the total number of observations of the random variable X and let $x_1, \ldots, x_N$ be those observations. Furthermore, let $x_{(1)}, \ldots, x_{(N)}$ be the observations ranked in increasing order ($x_{(1)} \leq \ldots \leq x_{(N)}$). With $a \in [0,1]$ given, we define the trimming factor, i.e., the fraction of scenarios that are removed, as $\alpha = 1 - a$, and the number of scenarios left after trimming as $N_\alpha = \lfloor N\,(1-\alpha) + \alpha \rfloor$ so that there is no trimming at $\alpha = 0$ ($N_\alpha = N$) and that only the worst scenario remains at $\alpha = 1$ ($N_\alpha = 1$). (In other words, this allows N_α to span $\{1, \ldots, N\}$ when α spans $[0,1]$.) Then the quantity:

$$\frac{1}{N_\alpha} \sum_{k=1}^{N_\alpha} x_{(k)} \tag{1}$$

is a non-parametric estimator of $E[X|X \leq q_a(X)]$.

2.2 Data-Driven Revenue Management

In the context of revenue management, we focus on the random profit $\pi(\mathbf{Q}, \mathbf{D})$ associated with a vector of orders $\mathbf{Q}$ subject to uncertain demand $\mathbf{D}$, for which we have observations $\mathbf{d_1}, \ldots, \mathbf{d_N}$. It follows from the discussion in Section 2.1 that we will measure the value associated with the random profit $\pi(\mathbf{Q}, \mathbf{D})$ by computing:

$$\frac{1}{N_\alpha} \sum_{k=1}^{N_\alpha} \pi(\mathbf{Q}, \mathbf{d}_\cdot)_{(k)}, \tag{2}$$

where $\pi(\mathbf{Q}, \mathbf{d}.)_{(k)}$ is the k-th smallest $\pi(\mathbf{Q}, \mathbf{d_j})$. As a result, the problem:

$$\begin{aligned} \max \quad & \frac{1}{N_\alpha} \sum_{k=1}^{N_\alpha} \pi(\mathbf{Q}, \mathbf{d}.)_{(k)} \\ \text{s.t.} \quad & \mathbf{Q} \in \mathcal{Q}, \end{aligned} \tag{3}$$

will be at the center of our analysis, with $\mathcal{Q}$ the feasible set, assumed convex, for the orders $\mathbf{Q}$. The naive approach at this point would be to select a feasible $\mathbf{Q} \in \mathcal{Q}$, rank the $\pi(\mathbf{Q}, \mathbf{d_k})$, take the average over the N_α smallest profits, and re-iterate for another $\mathbf{Q}$ that would improve the trimmed profit. The following theorem establishes that Problem (3) has a tractable, convex formulation, provided that $\pi(\mathbf{Q}, \mathbf{D})$ is concave in $\mathbf{Q}$, and therefore can be solved much more efficiently than the naive approach suggests.

Theorem 2.2 (The Data-Driven Problem)
(a) Problem (3) is equivalent to:

$$\begin{aligned} \max \quad & \phi + \frac{1}{N_\alpha} \sum_{k=1}^{N} \psi_k \\ s.t. \quad & \phi + \psi_k \leq \pi(\mathbf{Q}, \mathbf{d_k}), \quad \forall k, \\ & \psi_k \leq 0, \qquad\qquad\qquad \forall k, \\ & \mathbf{Q} \in \mathcal{Q}, \end{aligned} \tag{4}$$

(b) Problem (4) is a convex problem if $\pi(\mathbf{Q}, \mathbf{D})$ is concave in $\mathbf{Q}$, and a linear programming problem if $\pi(\mathbf{Q}, \mathbf{D})$ is piecewise linear, concave in $\mathbf{Q}$ and $\mathcal{Q}$ is a polyhedron.

Proof: (a) For any vector $\mathbf{x}$ with ranked components $x_{(1)} \leq \ldots \leq x_{(N)}$, $\sum_{k=1}^{N_\alpha} x_{(k)}$ is the optimal solution of:

$$\begin{aligned} \min \quad & \sum_{k=1}^{N} x_k \, y_k \\ \text{s.t.} \quad & \sum_{k=1}^{N} y_k = N_\alpha, \\ & 0 \leq y_k \leq 1, \qquad \forall k. \end{aligned} \tag{5}$$

The feasible set of Eq. (5) is nonempty and bounded, therefore by strong duality, Eq. (5) is equivalent to:

$$\begin{aligned} \max \quad & N_\alpha \cdot \phi + \sum_{k=1}^{N} \psi_k \\ \text{s.t.} \quad & \phi + \psi_k \leq x_k, \qquad \forall k, \\ & \psi_k \leq 0, \qquad \forall k. \end{aligned} \tag{6}$$

Reinjecting Eq. (6) into Eq. (3) with $x_k = \pi(\mathbf{Q}, \mathbf{d_k})$ for all k yields Eq. (4).

(b) It follows immediately that Eq. (4) is a convex problem if π is concave in $\mathbf{Q}$. Moreover, if π is (concave) piecewise linear in $\mathbf{Q}$ and $\mathcal{Q}$ is a polyhedron, then Eq. (4) is a linear programming problem. □

As the revenue functions in the newsvendor problems considered in Sections 3 and 4 are piecewise linear with linear ordering constraints, Theorem 2.2 will allow us to derive tractable, linear programming formulations of the data-driven models.

3 The Classical Newsvendor Problem

We develop here the data-driven approach for the classical newsvendor problem, and provide some extensions to this framework in Section 4.

The purpose of the newsboy problem, in its most elementary version, is to find the optimal order for a perishable item in a single-period horizon, in order to maximize revenue. It has been at the center of a wide body of literature under the assumption of a known probability distribution for the demand, and the reader is referred to Porteus [11] for an introduction to this problem. We use the following notation:

c	the unit cost,
p	the unit selling price,
s	the unit salvage price,
$m = p/c - 1$	the markup factor,
$t = 1 - s/c$	the discount factor,
Q	the order quantity,
D	the random demand.

The random profit is given by:

$$\pi(Q, D) = p \min(Q, D) + s \max(0, Q - D) - c\,Q, \tag{7}$$

or, equivalently:

$$\pi(Q, D) = (p - c)\, Q + (p - s)\, \min(0, D - Q). \tag{8}$$

Let $d_1, \ldots, d_N$ be previous observations of the demand, and $d_{(1)}, \ldots, d_{(N)}$ the same observations ranked in increasing order, so that $d_{(1)} \leq \ldots \leq d_{(N)}$. We assume $d_k \geq 0$ for all k and integers. The data-driven approach described in Section 2 maximizes the sample profit over the $N_\alpha = \lfloor N\,(1-\alpha) + \alpha \rfloor$ worst cases, that is:

$$\max_{Q \geq 0}\ (p - c)Q + (p - s) \cdot \frac{1}{N_\alpha} \sum_{k=1}^{N_\alpha} (\min(0, d_{\cdot} - Q))_{(k)}. \tag{9}$$

(There is no need to impose an integrality constraint on Q since, as we show in Theorem 3.1, the optimal order will be one of the breakpoints of the objective function, i.e., one of the d_j, which are all integers.) Since $\min(0, D-Q)$ is nondecreasing in D, the k-th smallest $\min(0, D-Q)$ at Q given is equal to $\min(0, d_{(k)} - Q)$. Therefore, in this case the data-driven formulation follows immediately, without requiring the use of strong duality:

$$\max_{Q \geq 0}\ (p - c)\, Q + (p - s) \cdot \frac{1}{N_\alpha} \sum_{k=1}^{N_\alpha} \min(0, d_{(k)} - Q). \tag{10}$$

Theorem 3.1 (Optimal policy) *The optimal order Q^* solution of Problem (10) satisfies:*

$$Q^* = d_{(j)} \ \textit{with}\ j = \left\lceil \frac{m}{m+t}\, N_\alpha \right\rceil. \tag{11}$$

Proof: For $d_{(k)} \leq Q \leq d_{(k+1)}$, the optimum is reached at $Q = d_{(k)}$ if $p - c \leq (p-s)k/N_\alpha$ and at $Q = d_{(k+1)}$ otherwise. Similarly, for $Q \leq d_{(1)}$ (resp. $Q \geq d_{(N)}$), the optimum is reached at $d_{(1)}$ (resp. $d_{(N)}$). Hence, the optimum over $Q \geq 0$ is reached at $Q = d_{(j)}$ where $(j-1)/N_\alpha < (p-c)/(p-s) \leq j/N_\alpha$. The result follows from $(p-c)/(p-s) = m/(m+t)$. □

Remarks:

- The optimal order is available in closed form, and only depends on ranking the historical data appropriately.
- The risk-averse newsvendor orders less than his risk-neutral counterpart, as $N_\alpha \leq N$ for all α. The rank j of the demand chosen as optimal order is roughly divided by two if the number of scenarios N_α that the decision-maker wants to keep is divided by two.

Comparison with the optimal order for the worst-case distribution:

Scarf gives in [14] the optimal order for the newsboy problem when only the mean μ and the standard deviation σ of the demand are known and the objective is to maximize the expected profit for the worst-case distribution:

$$Q^0 = \mu + \frac{\sigma}{2}\left(\sqrt{\frac{m}{t}} - \sqrt{\frac{t}{m}}\right). \tag{12}$$

The optimal policies in Scarf's model and in the data-driven approach have in common that they do not depend on the unit cost c, but rather on the markup and discount factors m and t. We note that if $m = t$, the optimal order is equal to the mean in Scarf's formulation, and to $d_{(\lceil N_\alpha/2 \rceil)}$ in the data-driven framework. In particular, if $\alpha = 0$, the data-driven optimal order is equal to the *median*, rather than the mean. Therefore, using the historical data or assuming only that the first two moments are known can result in very different ordering strategies, in particular for skewed distributions.

4 Extensions

We consider here several extensions to the classical problem presented in Section 3: other cost structures (Section 4.1), specifically with (a) a holding/shortage cost, (b) recourse, and (c) a fixed ordering cost, multiple items (Section 4.2), and multiple time periods (Section 4.3). We establish the following key results:

- The problems can be formulated as linear programming problems (or mixed integer programming problems if we impose integrality constraints on the orders), and therefore remain tractable.

- As in the classical newsvendor problem, the optimal solution can be interpreted by ranking the historical demands and provides key insights into the impact of the cost parameters on the optimal strategy.

- The optimal order will be available either in closed form (Section 4.1) or as a function of auxiliary parameters, obtained when solving the mathematical formulations: dual variables in Section 4.2, next reorder time and index of the N_α scenarios included in the data-driven approach in Section 4.3.

4.1 Other Cost Structures

4.1.1 The model with holding and shortage cost

In this section, the newsvendor incurs a cost when he has items left in inventory at the end of the time period, and when he has a shortage of items. Let h^+, h^- be the unit holding and shortage cost, respectively. As before, the unit cost is c and the unit price is p. The random profit is now given by:

$$\pi(Q, D) = (p - c + h^-)\,Q + (p + h^+ + h^-)\min(0, D - Q) - h^- D. \tag{13}$$

Our goal is to maximize the trimmed mean of the profit:

$$\max_{Q\geq 0}\ (p - c + h^-)\,Q + \frac{1}{N_\alpha}\sum_{k=1}^{N_\alpha}\left[(p + h^+ + h^-)\min(0, d. - Q) - h^- d.\right]_{(k)}, \tag{14}$$

where for any $\mathbf{y} \in \mathcal{R}^n$, $y_{(k)}$ is the k-th smallest component of $\mathbf{y}$.

Theorem 4.1 (Optimal policy)
(a) *The optimal order Q^* in (14) is the solution of the linear programming problem:*

$$\begin{array}{lll}
\max & (p - c + h^-)Q + \phi + \dfrac{1}{N_\alpha}\displaystyle\sum_{k=1}^{N}\psi_k & \\
s.t. & \phi + \psi_k - (p + h^+ + h^-)\,Z_k \leq -h^- d_k, & \forall k,\\
 & Z_k + Q \leq d_k, & \forall k,\\
 & Z_k \leq 0,\ \ \psi_k \leq 0, & \forall k,\\
 & Q \geq 0. &
\end{array} \tag{15}$$

Moreover, $Q^* = d_{(j)}$ for some j.
(b) *Let $M_\alpha = \left\lceil \dfrac{p - c + h^-}{p + h^+ + h^-} N_\alpha \right\rceil$. Q^* satisfies:*

$$Q^* = \min\left\{ d_{(j)} | d_{(j)} \geq \frac{p + h^+}{p + h^+ + h^-}\, d_{(M_\alpha)} + \frac{h^-}{p + h^+ + h^-}\, d_{(N - N_\alpha + M_\alpha)} \right\}. \tag{16}$$

(c) *Let S_α be the set of the N_α worst-case scenarios at optimality, that is,* $\sum_{i=1}^{N_\alpha}\pi(Q, d.)_{(i)} =$

$\sum_{i \in S_\alpha} \pi(Q, d_i)$*, and let* $d_{(j)}^{S_\alpha}$ *the* j*-th lowest demand within that set. We have:*

$$Q^* = d_{(M_\alpha)}^{S_\alpha}, \tag{17}$$

where M_α *is defined in (b).*

Proof: (a) This follows from applying Theorem 2.2 to (14). At optimality, $Q = d_{(j)}$ for some j because the function to maximize in (14) is concave piecewise linear with breakpoints in the set (d_i).

(b) The slope of the profit function is: $(p - c + h^-) - \frac{1}{N_\alpha}(p + h^+ + h^-) \cdot |\{i \in S(Q), d_i \leq Q\}|$, where $S(Q)$ is the set of indices of the N_α smallest $(p + h^+ + h^-)\min(0, d_i - Q) - h^- d_i$ at Q given. It is easy to show that for any $i \in S(Q)$ and any k such that $d_k \leq d_i \leq Q$, $k \in S(Q)$ as well. Similarly, for any $i \in S(Q)$ and any k such that $d_k \geq d_i \geq Q$, $k \in S(Q)$. Hence, $S(Q)$ consists of the indices of $d_{(1)}, \ldots, d_{(M_\alpha)}$ and $d_{(N-N_\alpha+M_\alpha+1)}, \ldots, d_{(N)}$, for some $0 \leq M_\alpha \leq N$, with $d_{(m_\alpha)} \leq Q \leq d_{(N-N_\alpha+M_\alpha+1)}$. The slope of the trimmed profit function is then proportional to $\frac{p - c + h^-}{p + h^+ + h^-} N_\alpha - M_\alpha$, and at optimality M_α is equal to $\left\lceil \frac{p - c + h^-}{p + h^+ + h^-} N_\alpha \right\rceil$. We now have to determine the optimal value of Q.

Let $f_i^j = (p + h^+ + h^-)\min(0, d_{(i)} - d_{(j)}) - h^- d_{(i)}$ be the profit realized when $Q = d_{(j)}$ and $D = d_{(i)}$, for all i and j. The optimal M_α is the greatest integer less than or equal to N_α such that $f_{M_\alpha}^j \leq f_{N-N_\alpha+M_\alpha}^j$. (Otherwise, we would remove M_α from $S(Q)$ and add $N - N_\alpha + M_\alpha$ instead.) Plugging the expressions of $f_{M_\alpha}^j$ and $f_{N-N_\alpha+M_\alpha}^j$ yields:

$$(p + h^+)d_{(M_\alpha)} - (p + h^+ + h^-)d_{(j)} \leq -h^- d_{(N-N_\alpha+M_\alpha)}. \tag{18}$$

Combining the previous results, Equation (16) follows immediately.

(c) Considering only the scenarios in S_α, we inject $N = N_\alpha$ into Equation (16). □

Remarks:

- If the shortage cost h^- is very large, then $M_\alpha = N_\alpha$ and $Q^* = d_{(N)}$, that is, the optimal order is equal to the largest demand, independently of the choice of N_α, as the main goal of the newsvendor is to have enough items in all possible scenarios.

- If the holding cost h^+ is very large, then $M_\alpha = 1$ and $Q^* = d_{(1)}$, that is, the optimal order is equal to the smallest demand, independently of the choice of N_α, as the main goal of the newsvendor is now to avoid holding items at the end of the time period.

- For $\alpha = 0$, we obtain the data-driven version of the optimal policy obtained when the exact distribution of the demand is known (see Gallego et. al. [6]):

$$Q = \arg\min \left\{ y | P(D \leq y) \geq \frac{p - c + h^-}{p + h^+ + h^-} \right\}. \tag{19}$$

4.1.2 The model with recourse

We now consider the classical newsvendor problem described in Section 3, under the additional assumption that, if the amount ordered at the beginning of the time period is insufficient to satisfy all the customers, the decision-maker places a second order at a unit cost of $c' = c(1+e)$ with $0 < e < m$ once the demand has been realized. Therefore, the demand is always met. The profit becomes:

$$\pi(Q, D) = p\,D + s\,\max(0, Q - D) - c\,Q - c'\,\max(0, D - Q), \tag{20}$$

or equivalently:

$$\pi(Q, D) = (c' - c)\,Q + (c' - s)\min(0, D - Q) + (p - c')D. \tag{21}$$

The profit function is identical to Equation (13) the model with holding and shortage cost, with $c' = p + h^-$ and $s = -h^+$. Therefore, all the results presented in Section 4.1.1 apply here. Moreover, if $\alpha = 0$, i.e., we maximize the sample profit, it is straightforward to show that the optimal policy with recourse is the same as the optimal policy without recourse where the markup factor m has been replaced by the cost premium e. This property is also verified by the optimal policy when the decision-maker seeks to maximize the expected profit for the worst-case distribution with mean and variance given (see Moon and Gallego [9].)

4.1.3 The model with fixed ordering cost

In this section, we consider the case where a fixed cost A is incurred whenever an order is made. Let $I \geq 0$ be the initial inventory. All the other assumptions remain the same as in Section 3. The random profit is here:

$$\pi(Q, D) = -A\,1_{\{Q>0\}} + p\,\min(Q + I, D) + s\,\max(Q + I - D, 0) - c\,Q, \tag{22}$$

or, defining $S = Q + I$:

$$\tilde{\pi}(S, D) = -A\,1_{\{S>I\}} + c\,I + (p - c)S + (p - s)\,\min(0, D - S). \tag{23}$$

Since $c\,I$ is a constant and $\min(0, D - S)$ is increasing in D, the data-driven approach solves:

$$\max_{S \geq I} \; -A\,1_{\{S>I\}} + (p-c)S + \frac{(p-s)}{N_\alpha} \sum_{k=1}^{N_\alpha} \min(0, d_{(k)} - S), \tag{24}$$

with $d_{(1)} \leq \ldots \leq d_{(N)}$.

Theorem 4.2 (Optimal policy) *It is optimal to order $S^* - I$ if $I \leq s^*$ and 0 otherwise, with:*

$$S^* = d_{(j)} \text{ where } j = \left\lceil \frac{m}{m+t} N_\alpha \right\rceil, \tag{25}$$

and:

$$s^* = S^* - \frac{\dfrac{A}{c} + \dfrac{m+t}{N_\alpha} \displaystyle\sum_{i=k+1}^{j} (S^* - d_{(i)})}{m - (m+t)\,\dfrac{k}{N_\alpha}}, \tag{26}$$

where k is such that:

$$m\,d_{(k)} + \frac{m+t}{N_\alpha} \sum_{i=1}^{k-1} (d_{(i)} - d_{(k)}) \leq -\frac{A}{c} + m\,S^* + \frac{m+t}{N_\alpha} \sum_{i=1}^{j} (d_{(i)} - S^*) < m\,d_{(k+1)} + \frac{m+t}{N_\alpha} \sum_{i=1}^{k} (d_{(i)} - d_{(k+1)}). \tag{27}$$

Proof: We have seen in Section 3 that the optimal solution for $A = 0$ is $S^* = d_{(j)}$ with $j = \left\lceil \frac{m}{m+t} N_\alpha \right\rceil$, yielding an objective value of:

$$K_\alpha(S^*) = c\left[m\,S^* + \frac{(m+t)}{N_\alpha} \sum_{i=1}^{j-1} (d_{(i)} - S^*)\right]. \tag{28}$$

If we choose not to order, the objective value becomes:

$$K_\alpha(I) = c\left[m\,I + \frac{(m+t)}{N_\alpha} \sum_{i=1}^{N} \min(0, d_{(i)} - I)\right]. \tag{29}$$

We want to find $s = \arg\min\{I | K_\alpha(I) \geq -A + K_\alpha(S^*)\}$, the threshold value for the inventory that makes ordering optimal (note that s here is not the salvage value. The only parameter related to the salvage value in this section is the discount factor t.) It is easy to see that $s \leq S$. Since $K_\alpha(x)$ is piecewise linear, nondecreasing for $x \leq S$, we first identify which interval $[d_{(k)}, d_{(k+1)})$ s belongs to by solving in k: $K_\alpha(d_{(k)}) \leq -A + K_\alpha(S) < K_\alpha(d_{(k+1)})$. Then $K_\alpha(s) = c\left[m\,s + \frac{m+t}{N_\alpha} \sum_{i=1}^{k} (d_{(i)} - s)\right]$ is linear in s and the value of s for which $K_\alpha(s) = -A + K_\alpha(S^*)$ follows from simple algebraic manipulations. □

Comparison with the optimal policy for the worst-case distribution: For $\alpha = 0$, we can compare the optimal policy with the one obtained by Moon and Gallego [10] when only the mean and the standard deviation of the distribution are known, namely:

$$S^0 = \mu + \frac{\sigma}{2}\left(\sqrt{\frac{m}{t}} - \sqrt{\frac{t}{m}}\right), \quad s^0 = \mu + \frac{(m-t)\widehat{A} - (m+t)\sqrt{\widehat{A}^2 - m\,t\,\sigma^2}}{2mt}, \tag{30}$$

where $\widehat{A} = \sigma\sqrt{m\,t} + \frac{A}{c}$. The differences between S^* and S^0 have been described in Section 3. Here we focus on $S^* - s^*$ and $S^0 - s^0$.

- For A/c large, $k = 0$ and:
$$S^* - s^* \approx S^0 - s^0 \approx \frac{A}{m\,c}. \tag{31}$$
In this case, both approaches leads to a similar value of $S - s$.

- For A/c small, $k = j - 1$ and:
$$S^* - s^* \approx \frac{A}{c\left(m - (m+t)\frac{j-1}{N}\right)}. \tag{32}$$
In the case where only mean and variance are known:
$$S^0 - s^0 \approx \frac{m+t}{(mt)^{3/4}} \cdot \sqrt{\frac{A\sigma}{2c}}. \tag{33}$$
Here, the data-driven approach yields very different results from the case where only the first two moments are known.

An appealing feature of the results presented throughout this section is that the optimal order is available in closed form. Furthermore, it is linked to the historical realizations of the demand, which highlights the impact of robustness (expressed through the trimming factor) on the optimal strategy. The closed-form expression was possible because there was only one source of uncertainty. In what follows, we consider cases with several random variables: multiple items (Section 4.2) or multiple periods (Section 4.3). Although a closed-form expression is no longer attainable, the parametric expressions we derive allow for a deeper understanding of the structure of the optimal policy, and in particular still characterize the optimal order as a past realization of the demand.

4.2 Multiple Items

We first consider an extension of the newsboy problem to multiple items in presence of a budget constraint. This case is also referred to in the literature as “stochastic product mix”. We use the following notations for item i, $i = 1, \ldots, n$:

c_i	the unit cost,
p_i	the unit selling price,
s_i	the unit salvage price,
$m_i = p_i/c_i - 1$	the markup factor,
$t_i = 1 - s_i/c_i$	the discount factor,
Q_i	the order quantity,
D_i	the random demand.

The budget constraint is:

$$\sum_{i=1}^{n} c_i \, Q_i \leq B, \tag{34}$$

where B is the total budget. The random profit can then be written as:

$$\pi(\mathbf{Q}, \mathbf{D}) = \sum_{i=1}^{n} c_i \left[m_i \, Q_i + (m_i + t_i) \, \min(0, D_i - Q_i) \right]. \tag{35}$$

The data observed in realization k, $k = 1, \ldots, N$, is $(d_1^k, \ldots, d_n^k)$, where d_i^k is the demand for item i. We consider the following problem:

$$\begin{aligned} \max \quad & \sum_{i=1}^{n} c_i \, m_i \, Q_i + \frac{1}{N_\alpha} \sum_{k=1}^{N_\alpha} \left(\sum_{i=1}^{n} c_i \, (m_i + t_i) \min(0, d_i^{\cdot} - Q_i) \right)_{(k)} \\ \text{s.t.} \quad & \sum_{i=1}^{n} c_i \, Q_i \leq B, \\ & Q_i \in \mathcal{Z}^+, \ \forall i. \end{aligned} \tag{36}$$

Theorem 4.3 (Optimal policy)

(a) The optimal orders $\mathbf{Q}^*$ *are solutions of the mixed integer programming problem:*

$$\begin{aligned} \max \quad & \sum_{i=1}^{n} c_i \, m_i \, Q_i + \phi + \frac{1}{N_\alpha} \sum_{k=1}^{N} \psi_k \\ s.t. \quad & \sum_{i=1}^{n} c_i \, Q_i \leq B, \\ & \phi + \psi_k - \sum_{i=1}^{n} c_i \, (m_i + t_i) Z_i^k \leq 0, \quad && \forall k, \\ & Z_i^k + Q_i \leq d_i^k, && \forall i, k, \\ & Z_i^k \leq 0, \ \psi_k \leq 0, \ Q_i \in \mathcal{Z}^+, && \forall i, k. \end{aligned} \tag{37}$$

(b) If we relax the integrality constraint on $\mathbf{Q}$*, let* $\lambda \geq 0$ *be the optimal Lagrangian multiplier for the budget constraint. For any* i*, exactly one of the following three cases holds:*

(i) $m_i < \lambda$**:** *Then we do not order item* i *(*$Q_i^* = 0.$*)*

(ii) $m_i > \lambda$**:** *Then* Q_i^* *satisfies:*

$$Q_i^* = (d_i)_{(j_i)} \text{ with } j_i = \left\lceil \frac{m_i - \lambda}{m_i + t_i} N_\alpha \right\rceil, \tag{38}$$

i.e., the optimal order for item i *is equal to the optimal order in the single-item problem with the new markup factor* $m_i' = m_i - \lambda$ *and the new discount factor* $t_i' = t_i + \lambda$*.*

(iii) $m_i = \lambda$**:** *Then if there is only one such* i*,* Q_i^* *satisfies:*

$$Q_i^* = \frac{B}{c_i} - \sum_{j \neq i} \frac{c_j}{c_i} Q_j^*, \tag{39}$$

i.e., after ordering the items j *for which* $m_j > \lambda$*, we use the remaining budget to order item* i*. If several items* i *have the same markup factor and this factor is equal to the Lagrangian multiplier* λ*, then the remaining budget is split among these items* i *(so that* $Q_i^* < (d_i)_{(1)}$*).*

Proof: (a) follows from Theorem 2.2. Moreover, if the integrality constraint on $\mathbf{Q}$ is relaxed, we can solve Problem (36) using a Lagrangian multiplier approach. Let $\lambda \geq 0$ be the optimal Lagrangian multiplier assigned to the budget constraint. The stochastic mix problem then becomes separable in the items and the subproblem for item i is equivalent to a single-item problem with a modified markup factor $m_i' = m_i - \lambda$, and a modified discount factor $t_i' = t_i - \lambda$. If $m_i' > 0$, we use the results of Section 3 to conclude. If $m_i' < 0$, the newsvendor obviously should not order, as he cannot sell the item at a premium. If $m_i' = 0$, the newsvendor can use his remaining budget (once he has ordered items with positive modified markup factor) to order such items. This proves (b). □

This approach can be extended to other cases with multiple sources of uncertainty, for instance random yield, where only some of the items ordered are in good enough condition to be sold, and random demand at the salvage value, where some of the remaining items may not find a buyer, even at the lower salvage price. In particular, the formulation remains a linear programming problem (or a mixed integer programming problem) and is thus tractable (see Thiele [15]).

4.3 Multiple Time Periods

We now extend the data-driven methodology to a multi-period inventory problem, where the decision-maker can order and sell the item over a finite time horizon and incurs a holding or shortage cost at the end of each period. Let X_t^k be the state of the inventory at the beginning of period t in scenario k. Let also d_t^k be the demand during period t in scenario k and $Q_t \geq 0$ the order at time t, common to all scenarios. The inventory follows linear dynamics:

$$X_{t+1}^k = X_t^k + Q_t - d_t^k, \quad \forall k, \ \forall t. \tag{40}$$

The profit realized during period t in scenario k is equal to, with the same notations as in Section 4.1 in the holding/shortage case:

$$\pi_t(X_t^k, Q_t, d_t^k) = p \, \min(d_t^k, X_t^k + Q_t) - c\, Q_t - \max(h^+ X_{t+1}^k, -h^- X_{t+1}^k), \quad \forall k, \ \forall t, \tag{41}$$

or equivalently:

$$\pi_t(X_t^k, Q_t, d_t^k) = p\, d_t^k - c\, Q_t - \max(h^+ X_{t+1}^k, -(p + h^-)\, X_{t+1}^k), \quad \forall k, \ \forall t. \tag{42}$$

Theorem 4.4 (The data-driven problem) *The data-driven problem can be formulated as a linear programming problem:*

$$\begin{array}{lll}
\max & -c \displaystyle\sum_{t=0}^{T-1} Q_t + \phi + \frac{1}{N_\alpha} \sum_{k=1}^{N} \psi_k & \\
s.t. & \phi + \psi_k + \displaystyle\sum_{t=0}^{T-1} Z_{t+1}^k \leq p \sum_{t=0}^{T-1} d_t^k, & \forall k, \\
& Z_{t+1}^k - h^+ \displaystyle\sum_{s=0}^{t} Q_s \geq h^+ \left(X_0 - \sum_{s=0}^{t} d_s^k \right), & \forall k, \ \forall t, \\
& Z_{t+1}^k + (p + h^-) \displaystyle\sum_{s=0}^{t} Q_s \geq -(p + h^-) \left(X_0 - \sum_{s=0}^{t} d_s^k \right), & \forall k, \ \forall t, \\
& \psi_k \leq 0, \ Q_t \geq 0, & \forall k, \ \forall t.
\end{array} \tag{43}$$

Proof: Theorem 2.2 yields the following data-driven counterpart for the multi-period newsboy problem:

$$\begin{aligned} \max \quad & -c\sum_{t=0}^{T-1} Q_t + \phi + \frac{1}{N_\alpha}\sum_{k=1}^{N}\psi_k \\ \text{s.t.} \quad & \phi + \psi_k \le \sum_{t=0}^{T-1}\left(p\, d_t^k - \max(h^+ X_{t+1}^k, -(p+h^-)X_{t+1}^k)\right), \quad && \forall k, \\ & X_{t+1}^k = X_0 + \sum_{s=0}^{t}(Q_s - d_s^k), && \forall k,\ \forall t, \\ & \psi_k \le 0,\ Q_t \ge 0, && \forall k,\ \forall t. \end{aligned} \tag{44}$$

We then use the closed-form expression of X_{t+1}^k and introduce auxiliary variables Z_{t+1}^k with $Z_{t+1}^k \ge h^+ X_{t+1}^k$ and $Z_{t+1}^k \ge -(p+h^-)X_{t+1}^k$ to obtain (43). □

We next provide some insight into the structure of the optimal solution at time 0, under the assumption that the set of worst cases is "nondegenerate", i.e., at optimality the profit of the $(N_\alpha+1)$-st worst scenario is strictly higher than the profit of the (N_α)-th worst one. We also assume $Q_0 > 0$ (it is optimal to order at the current time period). Let $\tau = \arg\min\{t \ge 1 | Q_t > 0\}$ be the next time period an order is made ($\tau = T$ by convention if no order is made after time 0), and let $\left[\sum_{s=0}^{\tau-1} d_s^{\cdot}\right]_{(k)}^{S_\alpha}$ be the k-th smallest $\sum_{s=0}^{\tau-1} d_s^k$ among the set S_α of the N_α worst-case scenarios. (The numerical value for τ and the worst-case scenarios are obtained by solving Problem (43).)

Theorem 4.5 (The data-driven solution) *We have:*

(a) If $\tau < T$ and $\dfrac{p+h^-}{p+h^++h^-}N_\alpha$ is not an integer,

$$Q_0 = -X_0 + \left[\sum_{s=0}^{\tau-1} d_s^{\cdot}\right]_{(k)}^{S_\alpha} \quad \textit{with } k = \left\lceil \frac{p+h^-}{p+h^++h^-}N_\alpha \right\rceil. \tag{45}$$

(b) If $\tau = T$ and $\dfrac{p+h^- - c}{p+h^++h^-}N_\alpha$ is not an integer,

$$Q_0 = -X_0 + \left[\sum_{s=0}^{T-1} d_s^{\cdot}\right]_{(k)}^{S_\alpha} \quad \textit{with } k = \left\lceil \frac{p+h^- - c}{p+h^++h^-}N_\alpha \right\rceil. \tag{46}$$

Proof: The dual of Problem (43) is:

$$\begin{aligned}
\min \quad & \sum_{t=0}^{T-1}\sum_{k=1}^{N}\left\{H_{t+1}^k\left[p\,d_t^k + h^+\left(X_0 - \sum_{s=0}^{t} d_s^k\right)\right] + P_{t+1}^k\left[p\,d_t^k - (p+h^-)\left(X_0 - \sum_{s=0}^{t} d_s^k\right)\right]\right\} \\
\text{s.t.} \quad & H_{t+1}^k + P_{t+1}^k - S_k = 0,\ \forall k, t, & : Z_{t+1}^k \\
& \sum_{k=1}^{N} S_k = 1, & : \phi \\
& S_k \le \frac{1}{N_\alpha},\ \forall k, & : \psi_k \\
& \sum_{t\ge s}\sum_{k=1}^{N}\left((p+h^-)P_{t+1}^k - h^+ H_{t+1}^k\right) \le c,\ \forall s, & : Q_s \\
& \mathbf{S}, \mathbf{H}, \mathbf{P} \ge \mathbf{0}.
\end{aligned} \tag{47}$$

where we have indicated the corresponding primal variables on the right, or equivalently:

$$\begin{aligned}
\min \quad & \sum_{t=0}^{T-1}\sum_{k=1}^{N}\left\{H_{t+1}^k\left[p\,d_t^k + h^+\left(X_0 - \sum_{s=0}^{t} d_s^k\right)\right] + P_{t+1}^k\left[p\,d_t^k - (p+h^-)\left(X_0 - \sum_{s=0}^{t} d_s^k\right)\right]\right\} \\
\text{s.t.} \quad & H_{t+1}^k + P_{t+1}^k \le \frac{1}{N_\alpha},\ \forall k, \\
& \sum_{k=1}^{N}\left(H_{t+1}^k + P_{t+1}^k\right) = 1,\ \forall t, \\
& \sum_{t\ge s}\sum_{k=1}^{N}\left((p+h^-)\,P_{t+1}^k - h^+ H_{t+1}^k\right) \le c,\ \forall s, \\
& \mathbf{H}, \mathbf{P} \ge \mathbf{0}.
\end{aligned} \tag{48}$$

Because of the assumption of "nondegeneracy" of the set of worst-case scenarios, at optimality: (a) exactly $N_\alpha - 1$ of the ψ_k are strictly negative, and these indices k are those of the worst $N_\alpha - 1$ scenarios, and (b) ϕ is equal to the (data-driven part of the) N_α-th worst profit, which corresponds to a unique scenario. Therefore, we have at optimality by complementarity slackness:

1. if k is not among the N_α worst-case scenarios, $H_{t+1}^k = P_{t+1}^k = 0$ for all t.

2. if k is among the N_α worst-case scenarios,

   (i) if $X_{t+1}^k < 0$, $H_{t+1}^k = 0$ and $P_{t+1}^k = \frac{1}{N_\alpha}$,

   (ii) if $X_{t+1}^k > 0$, $H_{t+1}^k = \frac{1}{N_\alpha}$ and $P_{t+1}^k = 0$,

   (iii) if $X_{t+1}^k = 0$, $H_{t+1}^k + P_{t+1}^k = \frac{1}{N_\alpha}$ for $H_{t+1}^k,\ P_{t+1}^k \ge 0$.

3. if $Q_s > 0$ at time s,

$$\sum_{t \geq s} \sum_{k=1}^{N} \Big((p + h^-) \, P_{t+1}^k - h^+ \, H_{t+1}^k \Big) = c. \tag{49}$$

4. let assume that $Q_0 > 0$ and $Q_1 > 0$. Equation (49) becomes:

$$(p + h^-) \sum_{k=1}^{N} P_1^k - h^+ \sum_{k=1}^{N} H_1^k = 0. \tag{50}$$

   Let S_α^-, resp. S_α^0, S_α^+, be the set of scenarios in S_α for which $X_1^k < 0$, resp. $X_1^k = 0$, $X_1^k > 0$. (We have: $|S_\alpha^-| + |S_\alpha^0| + |S_\alpha^+| = N_\alpha$.) Equation (50) becomes, using the expressions of H_1^k and P_1^k derived above:

$$|S_\alpha^+| + |S_\alpha^0| = \frac{p + h^-}{p + h^+ + h^-} \, N_\alpha + \frac{h^+}{p + h^+ + h^-} \sum_{k \in S_\alpha^0} p_k, \tag{51}$$

   where the p_k are numbers in $[0, 1]$. Assuming that $\frac{(p + h^-) \, N_\alpha}{p + h^+ + h^-}$ is not an integer, so that there will be at least one scenario k in S_α such that $X_1^k = 0$, this yields:

$$|S_\alpha^+| \leq \left\lfloor \frac{p + h^-}{p + h^+ + h^-} \, N_\alpha \right\rfloor \text{ and } |S_\alpha^+| + |S_\alpha^0| \geq \left\lceil \frac{p + h^-}{p + h^+ + h^-} \, N_\alpha \right\rceil, \tag{52}$$

   In particular, the k-th smallest demand in S_α, $(d_0)_{(k)}^{S_\alpha}$ with $k = \left\lceil \frac{p + h^-}{p + h^+ + h^-} \, N_\alpha \right\rceil$, corresponds to a scenario in S_α^0, i.e., $X_1^k = 0$. Equivalently,

$$Q_0 = -X_0 + (d_0)_{(k)}^{S_\alpha} \text{ with } k = \left\lceil \frac{p + h^-}{p + h^+ + h^-} \, N_\alpha \right\rceil. \tag{53}$$

5. In the general case where no control is implemented before time τ with $\tau < T$ the time of the first order after time 0, and if $Q_0 > 0$ and $\frac{(p + h^-) \, N_\alpha}{p + h^+ + h^-}$ is not an integer, the situation is equivalent to the case where $Q_1 > 0$ if we aggregate the demand from time 0 to time $\tau - 1$. Therefore, we have:

$$Q_0 = -X_0 + \left[\sum_{s=0}^{\tau - 1} d_s^{\cdot} \right]_{(k)}^{S_\alpha} \text{ with } k = \left\lceil \frac{p + h^-}{p + h^+ + h^-} \, N_\alpha \right\rceil, \tag{54}$$

6. If $Q_0 > 0$, time 0 is the last time period where a control is implemented and $\frac{p + h^- - c}{p + h^+ + h^-} \, N_\alpha$

is not an integer, a similar analysis yields:

$$Q_0 = -X_0 + \left[\sum_{s=0}^{T-1} d_s^{\cdot}\right]_{(k)}^{S_\alpha} \text{ with } k = \left\lceil \frac{h^+ + c}{p + h^+ + h^-} N_\alpha \right\rceil, \tag{55}$$

□

Remark: As in the single-period case, the decision-maker will order an amount equal to the largest, resp. smallest, demand among the worst-case scenarios if the shortage, resp. holding, cost is very large.

An important consequence of the results in this section is that the optimal policy in the data-driven, robust framework is basestock, as it brings the total amount of stock on hand and on order up to a certain level. In particular, we have expressed the threshold in terms of the historical cumulative demand for a specific scenario. These results are similar in structure to those obtained by Bertsimas and Thiele [3] using a robust optimization approach with polyhedral uncertainty sets.

5 Computational Experiments

5.1 Generalities

In this section, we provide numerical evidence that the data-driven approach leads to robust solutions, as well as empirical guidelines to select the trimming factor. We analyze how trimming the data affects statistics on the profit, with a focus on the change in the following measures:

(i) the sample conditional value-at-risk:

$$\rho = \mathrm{E}[\mathrm{X}|X \leq q_\alpha(X)], \tag{56}$$

with $X = \text{Profit}(Q_\alpha)$, Q_α is the optimal solution to the data-driven problem for a given α, and $q_\alpha(X)$ is the α-quantile of the random variable X,

(ii) the standard deviation to mean ratio, i.e., the coefficient of variation of the profit:

$$\lambda = \frac{\text{Std}[\text{Profit}(Q_\alpha)]}{\mathrm{E}[\text{Profit}(Q_\alpha)]}, \tag{57}$$

(iii) the ratio of the relative decreases in the standard deviation and mean:

$$\nu = \frac{\Delta\text{Std}[\text{Profit}(Q_\alpha)]}{\Delta\text{Mean}[\text{Profit}(Q_\alpha)]}, \tag{58}$$

where we have defined:

$$\Delta\text{Std}[\text{Profit}(Q_\alpha)] = \frac{\text{Std}[\text{Profit}(Q_0)] - \text{Std}[\text{Profit}(Q_\alpha)]}{\text{Std}[\text{Profit}(Q_0)]}, \tag{59}$$

and:

$$\Delta\text{Mean}[\text{Profit}(Q_\alpha)] = \frac{\text{Mean}[\text{Profit}(Q_0)] - \text{Mean}[\text{Profit}(Q_\alpha)]}{\text{Mean}[\text{Profit}(Q_0)]}. \tag{60}$$

The choice of these measures is justified as follows. The conditional value at risk is at the core of the data-driven approach, as the method's objective is to maximize its sample value over the historical realizations of the demand. It is therefore natural to compute its value in our numerical experiments and study its dependence on the trimming factor. The coefficient of variation is another intuitive way to measure risk, as it represents the variability of the profit per unit of expected return. Finally, the ratio of the relative decreases measures how much trimming the data changes the standard deviation per unit of change in the expected profit. In particular, if the ratio is greater than 1, the relative decrease in the standard deviation is greater than the relative decrease in the mean profit, which indicates a good trade-off between risk and return.

5.2 Numerical results

We implement the data-driven approach for the classical newsvendor problem described in Section 3. Our goal is to determine how the trimming factor should be selected, and to investigate the impact on this choice of the actual demand distribution, its coefficient of variation, the number of data points available and the cost parameters.

We generate the data according to three possible distributions: gamma, lognormal and gaussian, of given mean μ (here, $\mu = 100$) and of standard deviation $cv \cdot \mu$, where cv is the coefficient of variation of the demand. cv varies between 0.1 and 1, resp. 0.3, if the distribution is gamma or lognormal, resp. gaussian. (We do not consider greater coefficients of variation when the distribution is gaussian, as they would yield negative demands.) The number of historical data available is $N = 20$, 50, 100, 500 and the trimming factor α takes values between 0 and 1, specifically, $\alpha = 0.05 \cdot \beta$, with $\beta \in \{0, 1, \ldots, 20\}$. The ordering cost is set to $c = 10$. We take $p \in \{12, 14, 16, 18, 20\}$ and $s \in \{1, 3, 5, 7, 9\}$. We generate $N + 1$ data points according to the chosen distribution, use the first N values to derive the optimal order for α given and compute the corresponding profit realized when the demand is equal to the $(N+1)$-st value. We reiterate this experiment $5,000$ times, and compute the sample average and standard deviation of the profit as well as the ratios given in Equations (56), (57) and (58).

We highlight the most significant conclusions of the experiments below. The numerical results are available as an online appendix [4].

Influence of the distribution and its coefficient of variation

Mean and standard deviation of the profit

Figures 1-3 represent the mean (on the left panel) and the standard deviation (on the right panel) of the profit as a function of the trimming factor, for gaussian, resp. gamma, lognormal distributions. The coefficient of variation of the demand varies between 0.1 and 0.9. Here, we have chosen $N = 50$, $p = 14$ and $s = 7$.

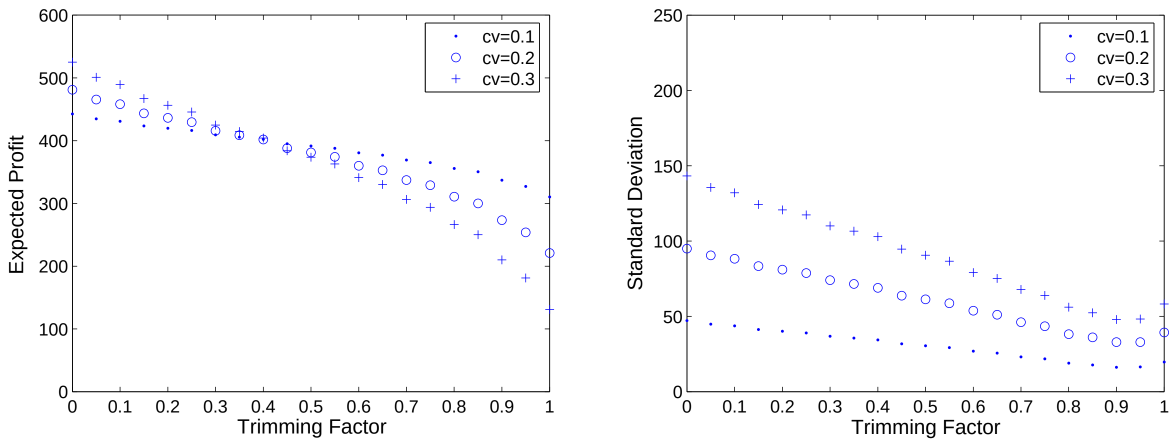

Figure 1: Expected value and standard deviation of the profit for gaussian distribution.

An important point is that the graphs exhibit the same qualitative trends for the different types of distributions. This is further illustrated on Figure 4, which represents the coefficient of variation of the profit as a function of the trimming factor for the three distributions. The coefficient of variation of the demand is equal to 0.1 on the left panel and 0.2 on the right panel. As an example, when the coefficient of variation of the demand is equal to 0.1, the coefficient of variation of the profit decreases *for each distribution* from approximately 0.10 to 0.04 as α increases from 0 to 0.9. Unsurprisingly, mean and standard deviation decrease as the trimming factor increases (Figures 1-3). We observe that the standard deviation initially decreases more sharply as a function of the trimming factor than the expected profit, which is further indicated

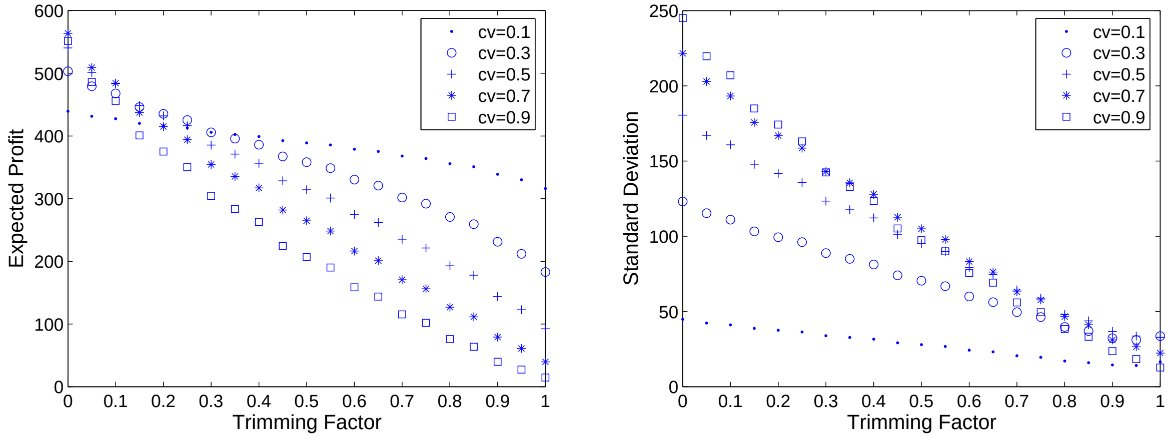

Figure 2: Expected value and standard deviation of the profit for gamma distribution.

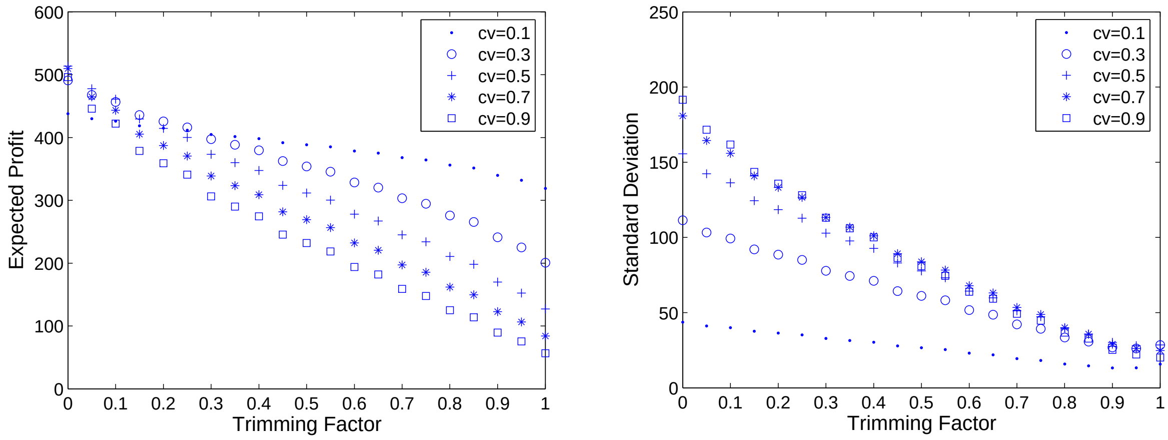

Figure 3: Expected value and standard deviation of the profit for lognormal distribution.

by the fact that the coefficient of variation of the profit decreases for $\alpha \leq 0.9$ on Figure 4. This in turn provides preliminary evidence that trimming the data does indeed increase the robustness of the solution.

Additional observations:

We also note from Figures 1-3 that the coefficient of variation of the demand has a greater impact on the expected profit when the trimming factor is large, and on the standard deviation when it is small. This matches intuition: when the trimming factor is large, the data-driven approach recommends ordering an amount equal to the smallest demand realized so far, and thus makes it extremely likely that all the items ordered will be bought in the future. This removes uncertainty and brings the standard deviation of the profit close to zero. On the other

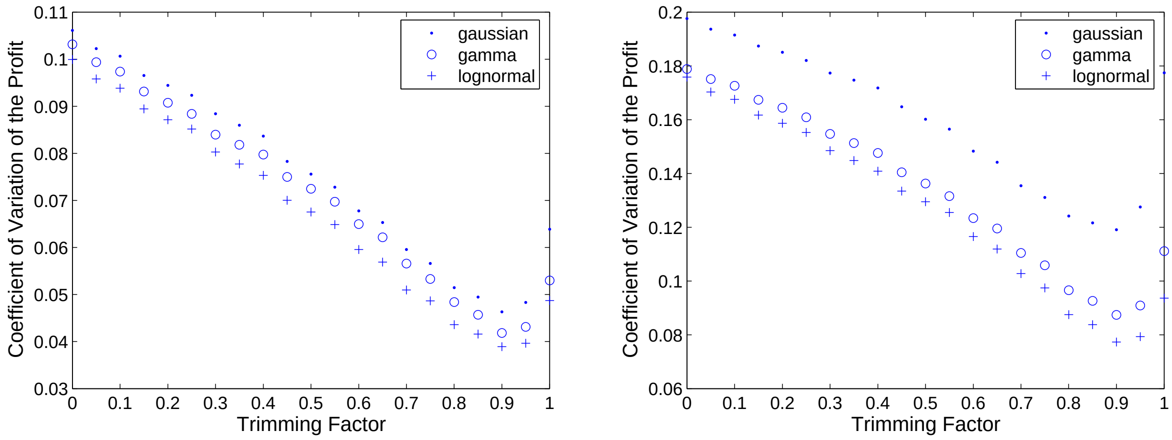

Figure 4: Coefficient of variation of the profit.

hand, the expected profit will then be proportional to the order, that is, the smallest demand, whose value varies widely depending on the coefficient of variation of the demand.

Finally, we observe that, when the trimming factor is below a threshold value of about 0.4 if the distribution is gaussian and 0.15 otherwise, the expected profit increases with the coefficient of variation of the demand, and decreases with higher trimming factors. An explanation is that more volatility yields some higher realizations of the demand, which are not (yet) trimmed when α is small. This incites the decision-maker to order more and to capture more profit opportunities. Moreover, the plots for the various coefficients of variation of the demand remain close from each other when the trimming factor is below the threshold. In other words, the coefficient of variation of the demand does not play an important role on the decrease in the expected profit for small values of the trimming factor. We will see in Figures 5-7 that this is also the case of the conditional value-at-risk.

Conditional value-at-risk and ratio of the relative decreases

Figures 5-7 represent the conditional value-at-risk (on the left panel) and the ratio of the relative decreases (on the right panel), as a function of the trimming factor, for gamma, resp. gaussian, lognormal distributions and values of the coefficient of variation of the demand between 0.1 and 0.9. As before, $N = 50$, $p = 14$ and $s = 7$.

The qualitative trends remain the same for the different distributions, suggesting that the actual distribution type does not overly affect the performance of the data-driven approach. Furthermore, we observe that the coefficient of variation of the demand does not significantly affect the conditional value-at-risk of the profit when the trimming factor is below some threshold.

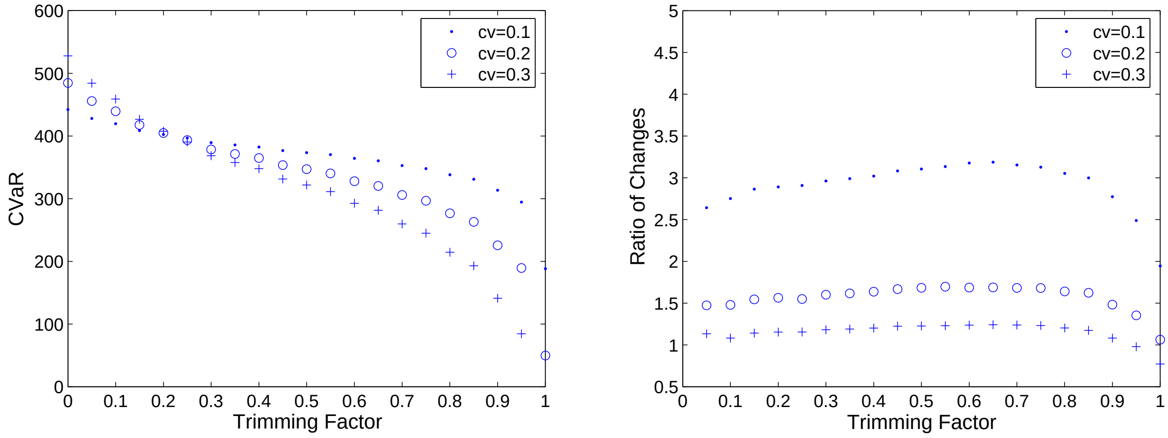

Figure 5: CVaR and ratio of relative decreases for gaussian distribution.

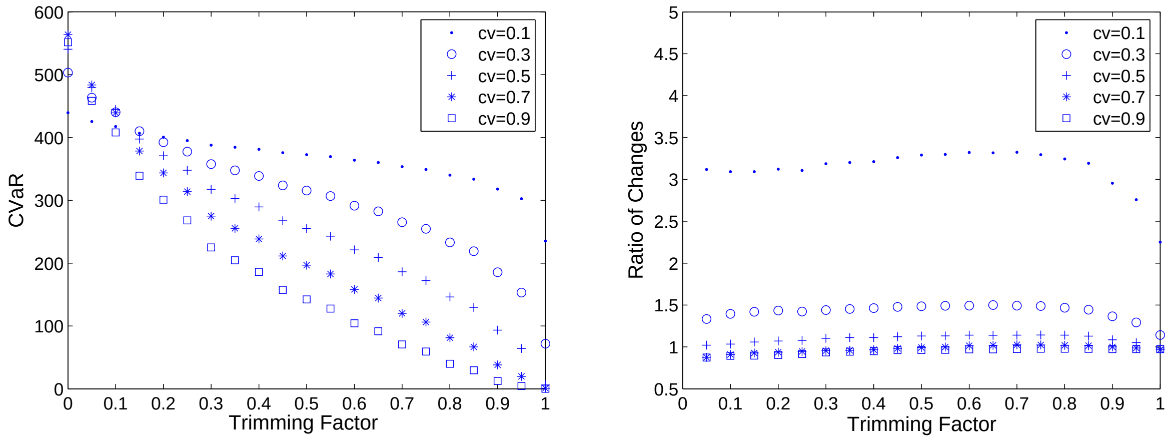

Figure 6: CVaR and ratio of relative decreases for gamma distribution.

The threshold in these numerical experiments is equal to 0.25 when the distribution is gaussian and 0.1 otherwise. Those values are slightly smaller than what we observed for the expected profit.

The ratio of the relative decreases, on the right panel of Figures 5-7, yields further insights. Most strikingly, the ratio remains approximately constant for all values of α below 0.8. In other words, the relative decrease in the standard deviation of the profit is roughly proportional to the relative decrease in the expected profit. The coefficient of proportionality appears to be mainly determined by the demand's coefficient of variation and decreases as the latter increases. For instance, a decrease of 10% in the expected profit will be associated with a decrease of about 30%, resp. 15%, 10% in the standard deviation of the profit if the coefficient of variation of the

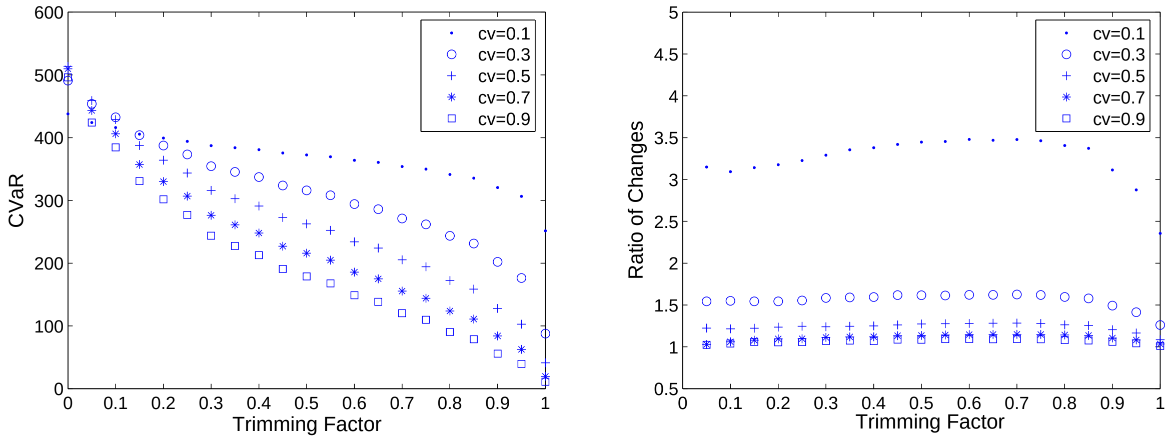

Figure 7: CVaR and ratio of relative decreases for lognormal distribution.

demand is 0.1, resp. 0.3, 0.9.

This further motivates the use of the data-driven approach to reduce uncertainty, as the standard deviation of the profit will be reduced by an amount greater than or similar to the decrease in the expected profit for a wide range of coefficients of variation, and the conditional value-at-risk decreases sharply for values of the trimming factor below 0.2. In order to preserve a good return, this motivates trimming the data by 10 to 20%.

Influence of the number of data points

Figure 8 shows the expected profit when the number of historical realizations of the demand is equal to $N = 20$ (left panel) and $N = 500$ (right panel). Figure 9 shows the conditional value-at-risk in those cases. We note that the performance of the method is not affected by the number of data points available, the only significant change being that the plots for $N = 500$ are smoother.

Influence of the cost parameters

Finally, we study the influence of the cost parameters on the data-driven approach. Figure 10 shows how the expected profit (on the left panel) and the conditional value-at-risk (on the right panel) evolve as a function of the trimming factor and the salvage value for $p = 14$. A key point is that, while the expected profit decreases steadily, the conditional value-at-risk decreases most sharply when the trimming factor is small, for all salvage values considered.

Figure 11, resp. Figure 12, Figure 13, shows the conditional value-at-risk of the profit as a function of the trimming factor for $p = 12$, resp. $p = 16$, $p = 20$. The left panel corresponds to $s = 1$ and the right panel corresponds to $s = 9$. We note again that trimming a small amount

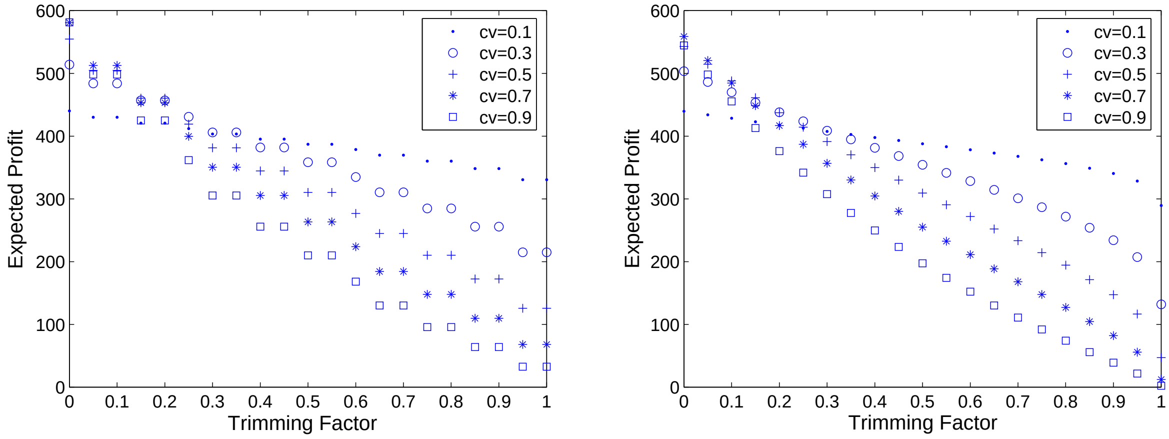

Figure 8: Expected profit for $N = 20$ (left) and $N = 500$ (right).

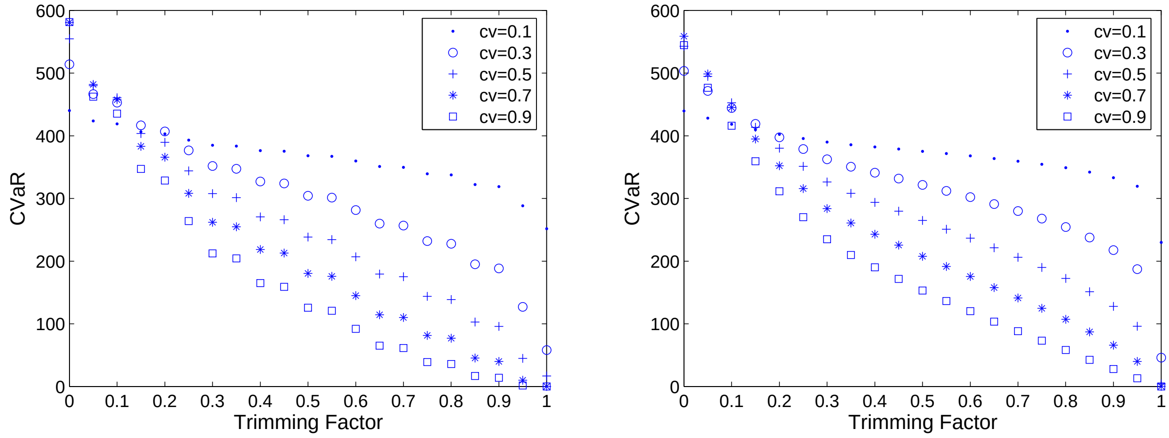

Figure 9: CVar for $N = 20$ (left) and $N = 500$ (right).

of the data ($\alpha \leq 0.20$) results in a sharp decrease in the conditional value-at-risk for all cost parameters considered. In order to reduce uncertainty without overly decreasing the expected profit, it thus seems appropriate to use a trimming factor between 10 and 15%.

5.3 Summary of Insights

We briefly summarize here the insights that we have gained from the numerical experiments.

- The performance of the data-driven approach is roughly independent of the actual demand distribution besides its first two moments.

- It is also roughly independent on the number of historical data points available. In par-

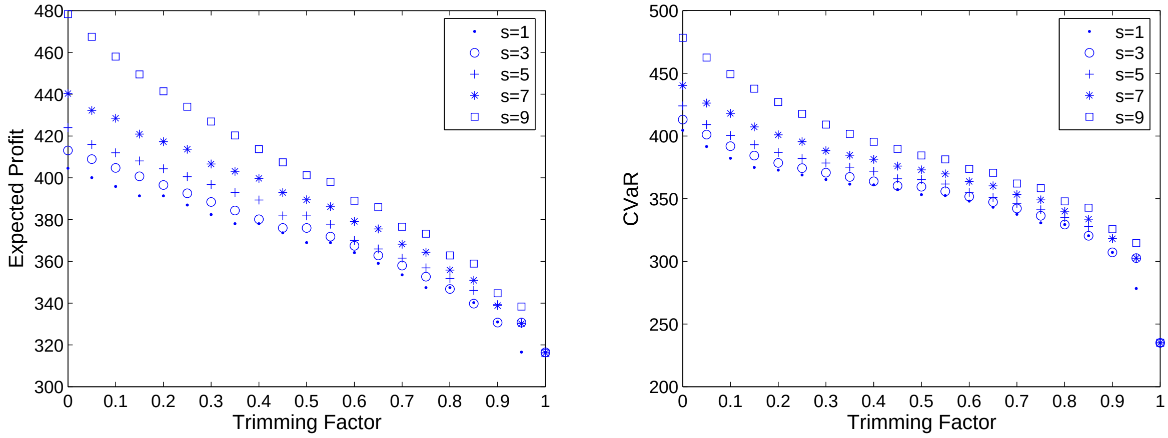

Figure 10: Expected profit and conditional value-at-risk for various salvage values.

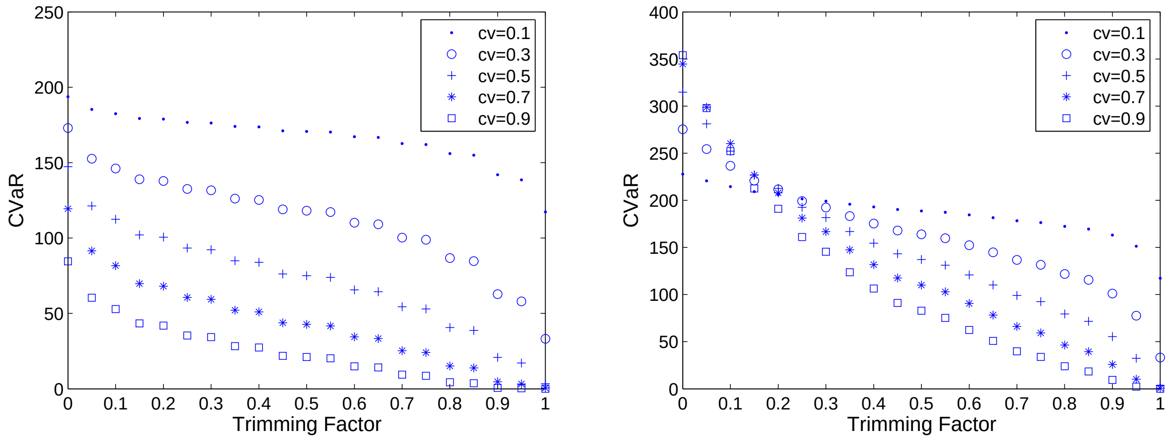

Figure 11: Conditional value-at-risk ($p = 12$).

ticular, the choice of the trimming factor does not depend on the number of past demand realizations.

- Standard deviation and conditional value-at-risk of the profit decrease as the trimming factor increases. The decrease is sharpest for small values of the trimming factor ($\alpha \leq 0.2$), while the decrease in the expected profit remains moderate on this range.

- A choice of the trimming factor between 0.1 and 0.15 yields a good trade-off between risk and return, i.e., between performance and conservatism.

- These insights hold for all cost parameters considered in the numerical study.

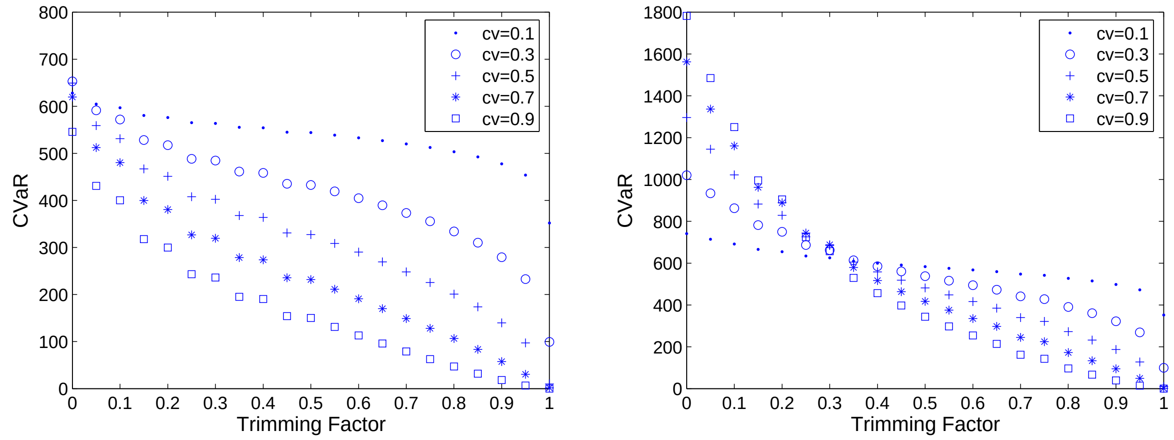

Figure 12: Conditional value-at-risk ($p = 16$).

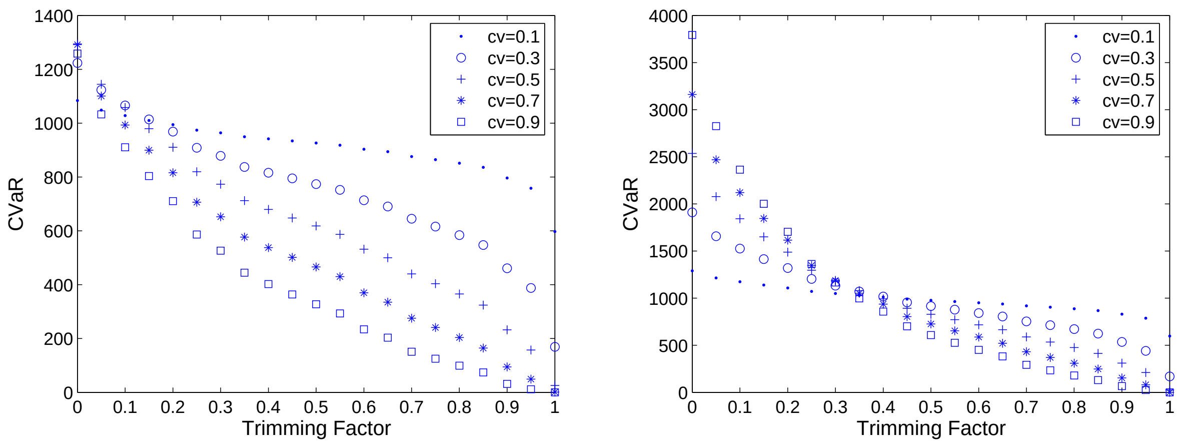

Figure 13: Conditional value-at-risk ($p = 20$).

6 Conclusions

We have proposed a robust approach to the newsvendor problem and its extensions that builds directly upon the historical data, without requiring any estimates of the probability distributions in the next time periods. This approach incorporates robustness through a single scalar parameter that can be adjusted to achieve an appropriate level of protection against uncertainty. Furthermore, the data-driven framework presented in this paper is connected to the theory of risk preferences, as it yields non-dominated strategies for risk-averse decision-makers. We have shown that the robust models could be formulated as linear programming problems, which has allowed us to derive structural properties of the optimal solutions, and these expressions have provided valuable insights into the impact of the problem parameters on the optimal policies.

The approach exhibits strong empirical performance in numerical experiments.

References

[1] Artzner, P., Delbaen, F., Eber, J.-M., Heath, D. (1999): Coherent measures of risk. *Mathematical Finance,* 9(3):203-228.

[2] Bertsimas, D., Lauprete, G., Samarov, A. (2004): Shortfall as a risk measure: Properties, optimization and applications. *Journal of Economic Dynamics and Control,* 28(7):1353-1381.

[3] Bertsimas, D., Thiele, A. (2005): A robust optimization approach to inventory theory, *Operations Research,* to appear.

[4] Bertsimas, D., Thiele, A. (2005): A data-driven approach to newsvendor problems, online appendix. Available at `http://www.lehigh.edu/~aut204/publications.html`.

[5] Eeckhoudt, L., Collier, C., Schlesinger, H. (1995): The risk-averse (and prudent) newsboy. *Management Science*, 41(5):786-794.

[6] Gallego, G., Ryan, J., Simchi-Levi, D. (2001): Minimax analysis for finite-horizon inventory models. *IIE Transactions,* 33:861-874.

[7] Lau, H.S. (1980): The newsboy problem under alternative optimization objectives. *Journal of the Operational Research Society,* 31(6):525-535.

[8] Levy, H., Kroll, Y. (1978): Ordering uncertain options with borrowing and lending. *The Journal of Finance,* 33:553-573.

[9] Moon, I., Gallego, G. (1993): The distribution-free newsboy problem: Review and extensions. *Journal of the Operational Research Society,* 44:825-834.

[10] Moon, I., Gallego, G. (1994): Distribution-free procedures for some inventory models. *Journal of the Operational Research Society,* 45:651-658.

[11] Porteus, E.L. (1990): Stochastic inventory theory. In D.P. Heyman and M.J. Sobel, editors, *Handbook in Operations Research and Management Science,* pages 605-652, Elsevier.

[12] Rousseeuw, P., Leroy, A. (1987): *Robust Regression and Outlier Detection,* Wiley.

[13] Ryan, T.P. (1996): *Modern Regression Methods*, Wiley.

[14] Scarf, H. (1958): A min-max solution of an inventory problem, in *Studies in the mathematical theory of inventory and production,* pages 201-209.

[15] Thiele, A. (2004): A robust optimization approach to supply chains and revenue management, PhD thesis, Massachusetts Institute of Technology.

[16] Wilcox, R. (1997): *Introduction to Robust Estimation and Hypothesis Testing,* Academic Press.