The Data-driven Multi-item Newsvendor Problem: A Learning Theoretic Perspective

Jakob de Raaij **Ulrik Unneberg**

Abstract

We consider the multi-dimensional version of the newsvendor problem, frequently called the multi-item case, with fixed ordering costs, in a setting where the distribution of the demand is unknown and can only be learned from random samples (the data-driven case). It is known from [13] that in the 1-dimensional case, an almost optimal supply can be efficiently obtained from past samples with high probability using sample average approximation (SAA). We connect SAA to the PAC-learing framework and prove that the same holds in the multi-item case: An almost optimal supply can be efficiently learned from past samples with high probability.

1 Introduction

In the data-driven multi-item newsvendor problem, an agent has to decide on the supply of k different resources, $q = (q_1, ..., q_k)$. There are unknown stochastic demands $d = (d_1, ..., d_k)$ for the k different resources, where d is drawn from a distribution $\mathcal{D}$. In contrast to the much-studied normal setting where $\mathcal{D}$ is known, we assume that $\mathcal{D}$ is unknown, but a polynomial number of past samples drawn i.i.d. from $\mathcal{D}$ are available to the agent. The agent faces a holding cost h_i and a lost-sales penalty b_i for each unit of resource i that they supplied exceeding demand and missing to fulfill the demand, respectively. All costs are assumed to be non-negative. The agent's expected cost of a supply vector q is

$$c(q) = \mathop{\mathbb{E}}_{d\sim\mathcal{D}}\left[\sum_{i=1}^{k} b_i(d_i - q_i)^+ + h_i(q_i - d_i)^+\right]. \tag{1}$$

We furthermore assume that there is an upper bound, Q, on the total quantity of resources available. We define $B = \left\{q \in \mathbb{R}^k_{\geq 0} : \sum_{i=1}^{k} q_i \leq Q\right\}$ as the budget set of all possible supply vectors. We assume the agent to be risk-neutral cost-minimizing. As such, she wants to find the supply vector

$$q^* := \arg\min_{q\in B} c(q) \tag{2}$$

minimizing expected cost.

Some authors specify the (multi-item) newsvendor problem by the costs c_i of purchasing a unit of item i, (fixed) prices p_i at which a unit of item i is sold, and potentially salvages value of s_i for each unsold unit of i (i.e., in excess of demand), with $p_i \geq c_i \geq s_i \geq 0$. We note that our problem statement is equivalent to this setting up to a constant additive term if one sets $h_i = c_i - s_i$ and $b_i = p_i - c_i$.

The importance of studying the newsvendor problem exceeds the illustrative example of a newsvendor deciding which quantity of different newspapers to bring by a lot. It is a foundational problem for a single company trying to make optimal stocking decisions in a competitive market, where the demand is unknown but can be estimated. While the single-item newsvendor problem is a useful and well-known model for optimal stocking decisions for a single item, it fails to capture correlations in the demand. Correlations in the demand, however, are very likely in the real world: The state of the economy affecting consumers wealth or, say, the weather, affecting the number of shoppers on a street, adds positive correlation to demands, while consumers having limited financial means adds negative correlation to demands. Since the overall resources (or budget) of a company are limited, they cannot making optimal decisions for each item independently, so correlations need to be considered.

1.1 Related Literature

Given both its theoretical and practical importance, it is no wonder that a vast literature exists on the newsvendor problem. The newsvendor problem was first investigated by Edgeworth [8]. Its modern formulation in the single-item case is due to Arrow et al. [2]. The multi-item case was first considered by Hadley [11], and has been since extended in many ways, for example to risk-adverse consumers or products with substitutions. We refer to Turken et al. [20] for a contemporary summary of results when the distribution $\mathcal{D}$ is known.

Scarf [17] was the first to consider the newsvendor problem with an unknown demand distribution $\mathcal{D}$ in a distributional robust optimization setting: He considers the single-item case where only the first two moments of $\mathcal{D}$ are known and solves for the optimal supply under a worst-case distribution. Gallego and Moon [9] extend their analysis and consider the multi-item case, to which Moon and Silver [15] introduced fixed ordering costs. Recently, Wang et al. [24] gave a closed form solution for the multi-item newsvendor problem for a worst-case distribution with known first two moments.

However, one criticism of this distribution-robust optimization approach is that in practice a scenario where on knows the first two moments – mean and variance – of the demand distribution, but nothing else, is highly unlikely. Much more likely, the agent has access to past samples from the demand distribution. While she could use the samples to estimate mean and variance with high accuracy, much more information about $\mathcal{D}$ is contained in the samples which she can use in her advantage.

The idea of optimizing for an unknown distribution with access to samples is generally referred to as Sample Average Approximation (SAA) in Operations Research literature, and was pioneered by Kleywegt and Shapiro [12]. Levi et al. [13, 14] were the first to apply SAA to the single-item newsvendor problem in a setting very similar to the one we consider ('data-driven'). They obtain tight bounds on the relative error. In particular, they show that for a single iterm, a sample of size

$$\Theta\left(\frac{1}{\varepsilon^2}\frac{b+h}{\min\{b,h\}}\ln\left(\frac{1}{\delta}\right)\right) \tag{3}$$

is sufficient to guarantee that $c(\hat{q}) \leq (1+\varepsilon)c(q^*)$, where $\hat{q}$ is a supply amount minimizing empirical cost on the sample, with probability at least $1-\delta$ over the i.i.d. random selection of the sample. To the best of our knowledge, the multi-item newsvendor problem has not yet been considered in the data-driven/SAA setting.

The idea of approximately optimizing with respect to an unknown distribution by obtaining a sufficiently large sample and optimizing on the sample has been known in computer science for a long time as probably-almost-correct (PAC) learning, as defined by Valiant [21]. We refer to Anthony and Bartlett [1] for an extensive introduction to Learning Theory; we will cite this book for some of the foundational results in our analysis. Only very recently has Learning Theory, in particular VC-Theory (named after Vapnik and Chervonenkis [22]), been started to be used in inventory theory by Xie et al. [25]. To the best of our knowledge, this is currently the only paper explicitly using learning theoretic tools for an inventory problem. Given the close and natural connection between SAA and PAC learning, we believe that it will be fruitful to apply results from PAC learning to SAA-type problem.

Lastly, this section would be incomplete without mentioning that there also is a rich literature on behavioral aspects of the newsvendor problem. Schweitzer and Cachon [18] demonstrate that in experiments a human agent, given full information about the demand distribution, frequently miscalibrates and chooses a supply closer to the mean than optimal. Benzion et al. [4] seek to explain these results with risk-aversion and loss-aversion. Bolton and Katok [5] extend these experiments to a case where the human agents receive feedback on their decisions. They verify that humans tend to choose supply rates too close to the center given high cost-asymmetries and don't adapt much to feedback.

1.2 Results

In Section 2, we show how to model our problem in a classical learning setting and show that its Pseudo-dimension is bounded by the dual VC dimension of a class we call axis-aligned distorted L_1-cubes. We believe that studying the VC-dimension of this class and related class, as we elaborate in the section, is of independent interest. We prove a bound of $O(k^2)$ and conjecture that this bound can be improved to $k+1$. We use these results to give a polynomial bound on the sample complexity of learning a probably almost optimal supply policy.

In Section 3, we reprove that learning a probably almost optimal supply policy is PAC-learnable from polynomially many samples by unwrapping the Pseudo-dimension machinery and tailering it to this specific problem. We get a sample complexity bound that is better than the bound from Section 2 roughly by a factor of k, and asymptotically identical to the bound we would obtained from Section 2 if our conjecture on the dual VC dimension of a class we call axis-aligned distorted L_1-cubes was true.

In Section 4 we provide the corresponding polynomial time algorithm for finding the optimal supply given a polynomial-size sample of past supply vectors. This completes the proof that the multi-item newsvendor problem is efficiently PAC-learnable.

Finally, in Section 5, we present and analyze our empirical findings from implementing our algorithm and testing its convergence for different b_i, h_i, and $\mathcal{D}$.

2 The pseudo-dimension of demand policies

We first show how finding an optimal supply policy can be rephrased as a learning problem. We then show that the function class of optimal demand policies has finite Pseudo-dimension. The Pseudo-dimension is one of several ways to generalize VC-dimension from classification to regression and gives comparable guarantees on convergence and the performance of empirical risk minimization. A detailed introduction can be found in [1].

We define $f_q : \mathbb{R}^k_{\geq 0} \to \mathbb{R}_{\geq 0}$ as

$$f_q(d) = \sum_{i=1}^{k} b_i(d_i - q_i)^+ + h_i(q_i - d_i)^+$$

and let $\mathcal{F} = \left\{ f_q : q \in \mathbb{R}^k_{\geq 0} \right\}$ be the set of all such f_q and let $\mathcal{F}_B = \{f_q : q \in B\}$ be its restriction where q is within our budget set. Furthermore, for any function $f \in \mathcal{F}$ we define its error $\text{err}(f)$ and empirical error $\widehat{\text{err}}(f)$ to be it's expected value on the distribution $\mathcal{D}$ and an i.i.d. sample $S \sim \mathcal{D}^m$, respectively, as

$$\text{err}(f) = \mathop{\mathbb{E}}_{d \sim \mathcal{D}}[f(d)] \qquad\qquad \widehat{\text{err}}(f) = \frac{1}{m} \sum_{d \in S} f(d).$$

Thus, we can rephrase our problem, Equation (2), as finding

$$f_{q^*} := \arg \min_{f_q \in \mathcal{F}_B} \text{err}(f).$$

Essential for our approach is the following theorem relating the empirical error to the error.

Theorem 1 (Theorem 19.2 in [1], modified)**.** *Suppose that $\mathcal{F}$ is a class of functions mapping from a domain X into the interval $[0, u]$ of real numbers for $u \in \mathbb{R}_{>0}$, and that $\mathcal{F}$ has finite pseudo-dimension. For $0 < \varepsilon < u$, $0 < \delta < 1$, let S be an i.i.d. sample from $\mathcal{D}$ of size*

$$m = \frac{64u^2}{\varepsilon^2} \left(2\text{P}\textsc{dim}(\mathcal{F}) \ln\left(\frac{16u}{\varepsilon}\right) + \ln\left(\frac{2}{\delta}\right) \right).$$

Then, with probability at least $1 - \delta$ over drawing S, it holds that $|err(f) - \widehat{err}(f)| \leq \varepsilon$ for all $f \in \mathcal{F}$.

Thus, if we can show that $\mathcal{F}$ has small pseudo-dimension, it follows that we can learn a good supply policy q by optimizing over a small number of past samples.

Theorem 2. *For any $(b_i)_{i\in[k]}$, $(h_i)_{i\in[k]}$, it holds that* $\textsc{Pdim}(\mathcal{F}) = O(k^2)$.

To prove Theorem 2, we need some definitions and 2 lemmas. We first define the dual VC dimension and a class we call "axis-aligned distorted L_1-balls in $\mathbb{R}^k$". We then show that the Pseudo-dimension of $\mathcal{F}$ is upper bounded by the dual VC-dimension of such axis-aligned distorted L_1-balls in $\mathbb{R}^k$. We then use a well-known corollary from combinatorial algebraic geometry to bound the dual VC-dimension of this class.

Definition 3 (Dual VC-dimension, modified from [10], Chapter 47.1)**.** Let X be the instance space and let $\mathcal{C}$ be a set of classifiers $c : X \to \{0,1\}$. In an abuse of notation, we write $c = \{x \in X : c(x) = 1\} \subseteq X$. The dual VC dimension of a class $\mathcal{C}$, consisting of subspaces, denoted $\textsc{VCdim}^*(\mathcal{C})$, is the VC dimension of the class $\{C_x = \{c \in \mathcal{C} : x \in c\} : x \in X\}$ over instance space $\mathcal{C}$.

To make sense of the definition of the dual VC-dimension, we give the following example. Consider $X = \mathbb{R}^3$ and let $\mathcal{C}$ be the class of cubes. Then, the (normal) VC-dimension is the maximum number n of points $P = \left\{x^1, ..., x^n\right\} \subset X$ that can be selected such that for any subset $Q \subseteq P$, there exists a cube $c \in \mathcal{C}$ such that all points in P are in the cube c, but no point in Q is in cube c. That is, $P \cap c = P$ but $Q \cap c = \emptyset$. Now, the dual VC-dimension is the maximum number n of cubes $C = \left\{c^1, ..., c^n\right\}$ that one can arrange in $\mathbb{R}^3$ so that for any subset $Q \subseteq C$, there exists a point $x \in \mathbb{R}^3$ such that x is in exactly those cubes that are in Q but no other. That is, $x \in c$ for all $c \in Q$ but $x \notin c$ for all $c \in C \backslash Q$.

To add another illustrative example, for $X = \mathbb{R}^2$, the dual VC-dimension of $\mathcal{C}$ is the largest size of a complete Venn-diagram one can draw using shapes in $\mathcal{C}$. For example, as one can easily evaluate with pen and paper, the dual VC-dimension of circles is 3 – there is no way to arrange 4 circles as a complete Venn diagram. We recommend [11, Chapter 47] for a thorough treatment.

Definition 4 (Axis-aligned distorted L_p-balls)**.** For $p > 1$, the axis-aligned distorted L_p-ball in $\mathbb{R}^k$ with center $c \in \mathbb{R}^k$, radius $r \in \mathbb{R}_{\geq 0}$ and 'stretching factors' $s^+, s^- \in \mathbb{R}^k_{\geq 0}$ for $i = 1, ..., k$ is

$$\tilde{B}^k_p(c, r, s^+, s^-) = \left\{ x \in \mathbb{R}^k : \left(\sum_{i=1}^{k} \frac{((x_i - c_i)^+)^p}{s_i^+} + \frac{((c_i - x_i)^+)^p}{s_i^-} \right)^{1/p} \leq r \right\} \subset \mathbb{R}^k.$$

For any $s^+, s^- \in \mathbb{R}^k_{\geq 0}$, we let

$$\tilde{\mathcal{B}}^k_p(s^+, s^-) = \left\{ \tilde{B}^k_p(c, r, s^+, s_-) : c \in \mathbb{R}^k, r \in \mathbb{R}_{\geq 0} \right\}$$

be the set of all axis-aligned distorted L_p-balls in $\mathbb{R}^k$ with stretching factors s^+, s^-.

Note that axis-aligned distorted L_∞-balls are boxes, in the case where all s_i^+ and s_i^- are equal they are axis-aligned (hyper)cubes. Axis-aligned distorted L_2-balls where $s_i^+ = s_i^-$ for all i are axis-aligned ellipsoids, and if all s_i^+ and s_i^- are equal they are (hyper)spheres. In R^3, axis-aligned distorted L_1-balls where all s_i^+ and s_i^- are equal are octahedrons.

Lemma 5. *For any k, $(b_i)_{i\in[k]}$, $(h_i)_{i\in[k]}$, it holds that* $\textsc{Pdim}(\mathcal{F}) \leq \textsc{VCdim}^*(\tilde{\mathcal{B}}^k_1(h, b))$.

Proof. Let $D = \left\{d^1, ..., d^n\right\}$ be a set of $n = \textsc{Pdim}(\mathcal{F})$ points, with respective thresholds $t_1, ..., t_n$, such that D is being pseudo-shattered by $\mathcal{F}$ with witnesses $t_1, ..., t_n$. For all $q \in \mathbb{R}_{\geq 0}k$, we denote by $g_q : \mathbb{R}^k_{\geq 0} : \{-1, 1\}$ the function $g_q(d) = \text{sign}(t_i - f_q(d))$, where we define $\text{sign}(0) = 1$. Thus, D being pseudo-shattered with these witnesses implies that for any labeling $(\ell_1, ..., \ell_n) \in \{-1, 1\}^n$, there exists a $q \in \mathbb{R}^k_{\geq 0}$ such that $-g_q(d) = \ell_i$ for all $i \in [n]$.

For any $j \in [n]$ and $q \in \mathbb{R}^k_{\geq 0}$, we know that $g_q(d^j) = 1$ if and only if

$$f_q(d^j) = \sum_{i=1}^{k} b_i(d_i^j - q_i)^+ + h_i(q_i - d_i^j)^+ \leq t_j.$$

Fixing $q_m = d_m^j$ for all $m \in [k] \backslash \{i\}$ for some $i \in [k]$, we see that $q_i \in [d_i^j - t_j/b_i, d_i^j + t_j/h_i]$ is equivalent to $g_q(d^j) = 1$. By the piecewise linearity of the constraints, this implies that $g_q(d^j) = 1$

is equivalent to q being in the convex polytope in $\mathbb{R}^k$ defined by the $2k$ points

$$\left\{(d_1^j, ..., d_{i-1}^j, d_i^j + t_j/h_i, d_{i+1}^j, ..., d_k^j), (d_1^j, ..., d_{i-1}^j, d_i^j - t_j/b_i, d_{i+1}^j, ..., d_k^j)\right\}_{i\in[k]},$$

which we denote P^j. Note that $P^j = \tilde{B}_1(d, t_j, 1/h, 1/b)$, where we denote by $1/h = (1/h_i)_{i\in[k]}$ and $1/b = (1/b_i)_{i\in[k]}$.

Since D is shattered, we know that for any labeling ℓ there exists a q_ℓ such that $g_{q_\ell}(d^j) = -\ell_j$ for all $j \in [n]$, i.e., such that $q_\ell \in P^j$ if and only if $\ell_j = -1$. This implies that $P = \left\{P^1, ..., P^n\right\}$ is an arrangement of n axis-aligned distorted L_1-cubes in $\mathbb{R}^d$ such that for any $P' \subseteq P$, there exists a point $q \in \mathbb{R}^d$ such that $q \in P^i$ for all $P^i \in P'$ but $q \notin P^i$ for all $P^i \in P \backslash P'$. Thus, $n \leq \text{VC}\textsc{dim}^*(\tilde{\mathcal{B}}_1^k(h, b))$. $\square$

Finding the exact dual VC-dimension of axis-aligned distorted L_1-balls is an open problem. Indeed, there are many interesting open problems regarding the (dual and normal) VC-dimension of L_p balls. For example, while the VC-dimension of axis-aligned L_2 balls ($k + 1$, [7]) and axis-aligned L_∞ balls ($\lfloor(3k + 1)/2\rfloor$, [6]) in $\mathbb{R}^k$ are known, the VC dimension for other axis-aligned L_p balls, including L_1, is an open problem. There exist convex shapes in $\mathbb{R}^d$, $d > 2$, for which just their homothets (i.e. scaled and translated versions) already have infinite VC-dimension [16], so it is not a given that all L_p balls will have polynomial or even finite VC-dimension. To the best of our knowledge, far less research has been carried out on the dual VC-dimension of L_p balls, so even more interesting open problems remain there.

We give a proof that the dual VC-dimension of axis-aligned distorted L_1 balls in $\mathbb{R}^d$ is at most quadratic in d.

Lemma 6. *For any k, $(b_i)_{i\in[k]} > 0$, $(h_i)_{i\in[k]} > 0$, it holds that* $\text{VC}\textsc{dim}^*(\tilde{\mathcal{B}}_1^k(h, b)) = O(k^2)$.

Proof of Lemma 6. Let $P = \left\{P^1, ..., P^n\right\}$ be n axis-aligned distorted L_1 balls in $\mathbb{R}^k$ such that for any subsets $P' \subseteq P$ there exists a point $x \in \mathbb{R}^k$ such that $x \in P^i$ for all $P^i \in P'$ but $x \notin P^i$ for all $P^i \in P\backslash P'$.

Each axis-aligned distorted L_1 ball is the union of exactly 2^k hyperplanes in $\mathbb{R}^k$. Thus, P defines $n2^k$ hyperplanes in $\mathbb{R}^k$. It is known (see, for example [10, Chapter 28, Corollary 28.1.2]) that the number of k-cells (i.e. a maximal connected region in $\mathbb{R}^k$ not intersected by any hyperplane) due to m hyperplanes is $O(m^k)$. Thus, in our case there are $O(n^k 2^{(k^2)})$ k-cells. However, since each cell is either fully included or fully excluded from each L_1 ball, we know that the number of subsets $P' \subseteq P$ such that there exists a point $x \in \mathbb{R}^k$ such that $x \in P^i$ for all $P^i \in P'$ but $x \notin P^i$ for all $P^i \in P\backslash P'$ is upper bounded by the number of cells. We get that

$$2^n \geq O(n^k 2^{(k^2)}).$$

Thus, we see that there exists a constant c such that for $n = cd^2$ the above inequality is false. It follows that for P to be a set as described above, $n = O(k^2)$. $\square$

Our proof is rather 'wasteful' in a way that gives hope for a tighter upper bound to be found in the future. We conjecture that the actual dual VC-dimension is $d + 1$:

Conjecture 7. *For any k, $(b_i)_{i\in[k]} > 0$, $(h_i)_{i\in[k]} > 0$, it holds that* $\text{VC}\textsc{dim}^*(\tilde{\mathcal{B}}_1^k(h, b)) = k + 1$.

We now have the tools to prove Theorem 2.

Proof of Theorem 2. Follows directly from Lemmas 5 and 6. $\square$

We can now combine Theorems 1 and 2 to give a bound on the sample complexity of learning an optimal supply vector q:

Theorem 8. *Assume that there exists a $D > 0$ such that $\Pr_{d\sim\mathcal{D}}\left(\|d\|_\infty \leq D\right) = 1$. Let $L = \max\{b_1, ..., b_k, h_1, ..., b_k\}$. For any $0 < \varepsilon < u$, $0 < \delta < 1$, given an i.i.d. sample S from $\mathcal{D}$ of size*

$$m = O\left(\frac{L^2(kD + Q)^2}{\varepsilon^2}\left(k^2 \ln\left(\frac{1}{\varepsilon}\right) + \ln\left(\frac{1}{\delta}\right)\right)\right),$$

let $\hat{q}$ be the budget that minimizes $c(q)$ on S, i.e.

$$\hat{q} = \arg\min_{q \in B} \frac{1}{m} \sum_{i=1}^{k} b_i (d_i - q_i)^+ + h_i (q_i - d_i)^+ .$$

Recall that q^* is the actually best supply for $\mathcal{D}$, i.e.

$$q^* = \arg\min_{q \in B} c(q).$$

Then, with probability at least $1 - \delta$, it holds that

$$c(\hat{q}) \leq c(q^*) + \varepsilon.$$

Proof. Note that

$$\begin{aligned}
\max_{q \in B, d \in \mathcal{D}} f_q(d) &= \max_{q \in B, d \in \mathcal{D}} \sum_{i=1}^{k} b_i (d_i - q_i)^+ + h_i (q_i - d_i)^+ \\
&\leq \max_{q \in B, d \in \mathcal{D}} \sum_{i=1}^{k} \max\{b_i, h_i\} \cdot (d_i + q_i) \\
&\leq L \sum_{i=1}^{k} (D + q_i) \leq L(kD + Q).
\end{aligned}$$

Next, note that obviously $\textsc{Pdim}(\mathcal{F}_B) \leq \textsc{Pdim}(\mathcal{F})$. Thus, the theorem follows from Theorem 1, setting $u = L(kD + Q)$. $\square$

3 Unpacking the pseudo-dimension machinery: an improved bound

In this section, we take a step back and reprove a more tight version of Theorem 8, unwrapping the Pseudo-dimension machinery and applying it to this specific problem. We use the tools from the proof of Theorem 1, see for example [1], and show that when adapting them to our specific problem, we can improve the bound on m to the bound we would obtain from Theorem 8 if our Conjecture 7 was true (up to some changes in the constants).

Theorem 9. *Assume that there exists a $D > 0$ such that $\Pr_{d \sim \mathcal{D}}\big(\|d\|_\infty \leq D\big) = 1$. Let $L, D, S, q^\star, \hat{q}$ be defined as Theorem 8. For any $0 < \varepsilon < u$, $0 < \delta < 1$, if m satisfies*

$$m \;\geq\; \frac{18L^2\,(kD + Q)^2}{\varepsilon^2} \Big[k\,\log\Big(1 + \tfrac{6LkQ}{\varepsilon}\Big) + \log\tfrac{2}{\delta}\Big], \tag{4}$$

then with probability at least $1 - \delta$,

$$|c(\hat{q}) - c(q^\star)| \leq \varepsilon.$$

Proof. With lost–sales and holding costs $b_i, h_i > 0$, we define

$$\ell(q, d) := \sum_{i=1}^{k} \Big[b_i (d_i - q_i)^+ + h_i (q_i - d_i)^+\Big],$$

For m i.i.d. samples $d^{(1)}, \ldots, d^{(m)}$ let

$$c(q) := \mathbb{E}\,[\ell(q, d)], \qquad \hat{c}_m(q) := \frac{1}{m} \sum_{j=1}^{m} \ell\big(q, d^{(j)}\big).$$

$$q^\star := \arg\min_{q \in B} c(q), \qquad \hat{q}^{(m)} := \arg\min_{q \in B} \hat{c}_m(q).$$

Since $(x)^+ \le |x|$ and $q_i \ge 0$,

$$\ell(q,d) \le \sum_{i=1}^{k} \max(b_i,h_i)\,(|d_i| + q_i) \le L \sum_{i=1}^{k} (D + q_i) \le L\,(kD + Q).$$

Define

$$M := L\,(kD + Q). \tag{5}$$

For any $q, q' \in B$ and any d,

$$|\ell(q,d) - \ell(q',d)| \le \sum_{i=1}^{k} \max(b_i,h_i)\,|q_i - q_i'| \le L\,\|q - q'\|_1.$$

Taking expectations preserves the constant, so

$$|c(q) - c(q')| \;\le\; L\,\|q - q'\|_1 \qquad \forall q, q' \in B. \tag{6}$$

Fix accuracy $\varepsilon \in (0,1)$ and set

$$\eta := \frac{\varepsilon}{6Lk}. \tag{7}$$

Construct a Cartesian grid on $[0,Q]^k$ with spacing η and keep only points inside B

$$\mathcal{N} := \big\{ q \in B : q_i \in \{0, \eta, 2\eta, \dots\} \big\}.$$

Each coordinate has at most $1 + Q/\eta$ grid values, hence

$$|\mathcal{N}| \le \Big(1 + \frac{Q}{\eta}\Big)^k = \Big(1 + \frac{6LkQ}{\varepsilon}\Big)^k. \tag{8}$$

For any $q \in B$ choose $q_{\mathcal{N}} \in \mathcal{N}$ by rounding each coordinate down to the nearest grid point. Then $0 \le q_i - q_{\mathcal{N},i} < \eta$, so

$$\|q - q_{\mathcal{N}}\|_1 < k\eta. \tag{9}$$

For a fixed q the variables $X_j := \ell(q, d^{(j)})$ are i.i.d. in $[0,M]$. Hoeffding's inequality yields

$$\Pr\Big\{ |\hat{c}_m(q) - c(q)| > \tfrac{\varepsilon}{6} \Big\} \le 2\exp\Big(-\tfrac{2m(\varepsilon/6)^2}{M^2} \Big). \tag{10}$$

Apply (10) to every $q \in \mathcal{N}$ and union–bound using (8). If

$$m \;\ge\; \frac{18M^2}{\varepsilon^2} \Big[k\,\log\Big(1 + \tfrac{6LkQ}{\varepsilon}\Big) + \log\tfrac{2}{\delta} \Big], \tag{11}$$

then with probability at least $1 - \delta$ the event

$$\mathcal{E} := \Big\{ \forall q \in \mathcal{N} : |\hat{c}_m(q) - c(q)| \le \tfrac{\varepsilon}{6} \Big\}$$

occurs. Let $q \in B$ and choose $q_{\mathcal{N}}$ as in (9). Lipschitzness (6) gives

$$|c(q) - c(q_{\mathcal{N}})| \le Lk\eta = \tfrac{\varepsilon}{6}, \qquad |\hat{c}_m(q) - \hat{c}_m(q_{\mathcal{N}})| \le \tfrac{\varepsilon}{6}.$$

On $\mathcal{E}$, by definition we have

$$|\hat{c}_m(q_{\mathcal{N}}) - c(q_{\mathcal{N}})| \le \frac{\varepsilon}{6},$$

and by the triangle inequality,

$$|\hat{c}_m(q) - c(q)| \le |\hat{c}_m(q) - \hat{c}_m(q_{\mathcal{N}})| + |\hat{c}_m(q_{\mathcal{N}}) - c(q_{\mathcal{N}})| + |c(q_{\mathcal{N}}) - c(q)| \le \frac{\varepsilon}{6} + \frac{\varepsilon}{6} + \frac{\varepsilon}{6} = \frac{\varepsilon}{2}.$$

This holds for any $q \in B$, so

$$\sup_{q \in B} |\hat{c}_m(q) - c(q)| \le \frac{\varepsilon}{2}. \tag{12}$$

Within $\mathcal{E}$, since $c(q^\star) = \min_{q \in B} c(q)$,

$$\begin{aligned} 0 \le c\big(\hat{q}^{(m)}\big) - c(q^\star) &= c\big(\hat{q}^{(m)}\big) - \hat{c}_m\big(\hat{q}^{(m)}\big) + \hat{c}_m\big(\hat{q}^{(m)}\big) - \hat{c}_m(q^\star) + \hat{c}_m(q^\star) - c(q^\star) \\ &\le \frac{\varepsilon}{2} + 0 + \frac{\varepsilon}{2} = \varepsilon. \end{aligned}$$

Since $\mathcal{E}$ happens with probability at least $1 - \delta$, the theorem follows. $\square$

4 Learning algorithm

In this section we present a deterministic, polynomial–time routine that solves the sample–average approximation of the capacity–constrained multi–site newsvendor. Given the costs (b_i, h_i), a capacity budget Q and a set of m demand observations for each of the k locations, the procedure returns the order vector $\hat{q}^{(m)}$ that minimises the empirical objective (14). The algorithm uses a one–dimensional Lagrangian search after the m observations at every site are sorted once, a logarithmic number of evaluations of the aggregate order function $S(\lambda)$ isolates the critical multiplier λ^*. A final greedy fill distributes any left–over capacity across at most k sites. All operations therefore require at most $O(km \log m)$ time and linear memory, so the method scales gracefully with both the sample size and the number of locations. Algorithm 1 collects the steps.

4.1 Derivation

Let k locations be indexed by $i \in \{1, \ldots, k\}$. We ship a non–negative quantity q_i to each site, obeying the capacity

$$\sum_{i=1}^{k} q_i \;\le\; Q. \tag{13}$$

Random demand is $d = (d_1, \ldots, d_k) \sim \mathcal{D}$, unknown. Lost–sales and holding costs at location i are $b_i > 0$ and $h_i > 0$. Draw m independent demands $d^{(1)}, \ldots, d^{(m)} \in \mathbb{R}^k_{\ge 0}$ and denote $d_i^{(j)}$ the i–th component of scenario j. Replacing the expectation in (1) by an empirical mean gives the convex program

$$\min_{q \in B} \hat{c}(q), \qquad \hat{c}(q) := \frac{1}{m} \sum_{j=1}^{m} \sum_{i=1}^{k} \big[b_i (d_i^{(j)} - q_i)^+ + h_i (q_i - d_i^{(j)})^+ \big], \tag{14}$$

and define for a fixed m,

$$\hat{q}^{(m)} := \arg\min_{q \in B} \hat{c}(q). \tag{15}$$

Introduce a multiplier $\lambda \ge 0$ for the capacity

$$L(q, \lambda) := \hat{c}(q) + \lambda \Big(\sum_{i=1}^{k} q_i - Q \Big). \tag{16}$$

Because the sites couple only through the capacity, write the *indexed* single–site term

$$\phi_i(q_i, \lambda) = \frac{1}{m} \sum_{j=1}^{m} \Big[b_i \big(d_i^{(j)} - q_i \big)^+ + h_i \big(q_i - d_i^{(j)} \big)^+ \Big] + \lambda\, q_i, \tag{17}$$

$$L(q, \lambda) = \sum_{i=1}^{k} \phi_i(q_i, \lambda) - \lambda Q. \tag{18}$$

where the common multiplier $\lambda \ge 0$ enforces the global capacity.For any numbers x, y

$$(x - y)^+ = (x - y) \mathbf{1}\{x > y\}, \qquad (y - x)^+ = (y - x) \mathbf{1}\{x \le y\}.$$

Applying this identity to (17), after collecting the linear terms in q_i, gives

$$\phi_i(q_i, \lambda) = \Big[h_i\, \hat{F}_i(q_i) - b_i \big(1 - \hat{F}_i(q_i) \big) + \lambda \Big]\, q_i - C_i, \tag{19}$$

where

$$\hat{F}_i(q_i) = \frac{1}{m} \#\{ j : d_i^{(j)} \le q_i \}, \qquad C_i = \frac{h_i}{m} \sum_{d_i^{(j)} \le q_i} d_i^{(j)} - \frac{b_i}{m} \sum_{d_i^{(j)} > q_i} d_i^{(j)}.$$

Sort the m observations of $d_i^{(j)}$ at site i,

$$s_{i,1} \le s_{i,2} \le \cdots \le s_{i,m}, \qquad s_{i,0} = -\infty,\; s_{i,m+1} = +\infty. \tag{20}$$

On a slice $(s_{i,t}, s_{i,t+1})$ exactly t samples satisfy $d_i^{(j)} \le q_i$. Inside that interval $\phi_i(q_i, \lambda)$ is linear with slope

$$\frac{\partial \phi_i}{\partial q_i}(q_i, \lambda) = -b_i\big(1 - \hat{F}_i(q_i)\big) + h_i\, \hat{F}_i(q_i) + \lambda. \tag{21}$$

Setting (21) to zero yields the empirical *critical-fractile* condition

$$(b_i + h_i)\, \hat{F}_i\big(q_i(\lambda)\big) = b_i - \lambda, \qquad 0 \le \lambda \le b_i. \tag{22}$$

Equivalently,

$$\boxed{\hat{F}_i\big(q_i(\lambda)\big) = \frac{b_i - \lambda}{b_i + h_i}}$$

Since $\hat{F}_i$ jumps by $1/m$ at every order statistic, the minimal q_i satisfying (22) is

$$q_i(\lambda) = \begin{cases} s_{i,\lceil m\, p_i(\lambda)\rceil} & \text{if } \lambda < b_i, \\ 0 & \text{if } \lambda \ge b_i, \end{cases} \qquad p_i(\lambda) = \frac{b_i - \lambda}{b_i + h_i}. \tag{23}$$

Define the total order size

$$S(\lambda) = \sum_{i=1}^{k} q_i(\lambda). \tag{24}$$

Each $q_i(\lambda)$ is non-increasing in λ, hence $S(\lambda)$ is piece-wise constant, non-increasing, with $S(0) = \sum_i q_i(0)$ and $S(\lambda) = 0$ for $\lambda \ge \max_i b_i$. The complementary–slackness equation for the capacity is

$$\lambda^* = \max\{\lambda \in [0, \max_i b_i] : S(\lambda) \le Q, \}. \tag{25}$$

If $Q \ge S(0)$ the capacity does not bind; then $\lambda^* = 0$ and $\hat{q}_i^{(m)} = q_i(0)$ for every site. Otherwise, the monotonicity of S guarantees at least one $\lambda^* \in (0, \max_i b_i]$ satisfying (25). Such a λ^* is found efficiently by bisection because each evaluation of $S(\lambda)$ uses the pre-sorted arrays and costs $O(k \log m)$ time.

With λ^* obtained from (25) set initially

$$\hat{q}_i^m \leftarrow q_i(\lambda^*), \qquad i = 1, \dots, k. \tag{26}$$

Since $S(\lambda)$ is piece-wise constant and jumps each time $\lceil m p_i(\lambda)\rceil$ passes an integer, we might get a remainder $L \leftarrow Q - S(\lambda^*)$.

At the multiplier λ^* every component $q_i(\lambda^*)$ lies on a "plateau" of the piece–wise linear cost curve $\phi_i(\,\cdot\,, \lambda^*)$. Inside such a plateau the slope is constant and given by (21):

$$\phi_i'(q, \lambda^*) \;=\; -b_i\big(1 - \hat{F}_i(q)\big) + h_i \hat{F}_i(q) + \lambda^* \;=\; (b_i + h_i)\big[\hat{F}_i(q) - p_i(\lambda^*)\big].$$

Hence the marginal change in sample cost of shipping one additional unit to site i can differ across sites and can even be negative. Only the indices for which $q_i^{\uparrow} > q_i^{\downarrow}$ (the *jump set* J) can still accept more inventory without changing λ^*. To attain the minimum of (14) subject to the capacity, the left–over L must therefore be poured into the cheapest positions first, i.e. into those sites whose slope $\phi_i'(\,\cdot\,, \lambda^*)$ is most negative. The ordering used in Step 4, by the lexicographic key $(b_i, \bar{d}_i)$, is a low–cost surrogate for that slope: for fixed $\hat{F}_i$ the marginal benefit of an extra unit is monotone in $b_i + h_i$, and $\bar{d}_i$ provides a stable tie–break. After the greedy fill either L becomes zero (some sites in J still have room) or some subset of sites in J has reached their upper edges (and $L = 0$). In both cases complementary-slackness remains satisfied and the resulting allocation is the unique minimiser on the capacity face $\sum_i q_i = Q$. The vector $\hat{q}^{(m)} = (q_1^{(m)}, \dots, q_k^{(m)})$ minimises the sample-average objective (14) while meeting the capacity constraint. The algorithm is given in Algorithm 1.

4.2 Correctness of Algorithm 1

We now confirm that the routine indeed returns an optimal solution of (14). Recall that for every site i the empirical cost $q_i \mapsto \phi_i(q_i, \lambda)$ defined in (17) is the average of m piecewise linear functions, and is therefore convex. In particular,

$$\frac{\partial^2}{\partial q_i^2}\big[b_i(d_i^{(j)} - q_i)^+ + h_i(q_i - d_i^{(j)})^+\big] \;=\; (b_i + h_i)\, \delta(q_i - d_i^{(j)}) \;\ge 0,$$

where δ denotes the Dirac measure. Summing over the m scenarios shows that the (generalised) Hessian of the sample objective $\hat{c}$ is positive semidefinite, so $\hat{c}$ is convex in q. Consequently the Karush–Kuhn–Tucker (KKT) [19, 23] conditions are necessary and sufficient for optimality.

The algorithm chooses the Lagrange multiplier λ^* by the one–dimensional search (25). In Step 3 it constructs the vector $q^{\downarrow} = \big(q_i(\lambda^*)\big)_{i=1}^{k}$, which satisfies feasibility, $\sum_i q_i^{\downarrow} \leq Q$, stationarity, $0 \in \partial_{q_i} L(q^{\downarrow}, \lambda^*)$ for all i by (22), and complementary–slackness, $\lambda^* \big(\sum_i q_i^{\downarrow} - Q\big) = 0$. If $\sum_i q_i^{\downarrow} = Q$ we are done. Otherwise, Step 4 increases some of the $q_i^{\downarrow}$ inside their plateaus while keeping λ^* fixed, so properties (a)–(c) are preserved and the final vector $\hat{q}^{(m)}$ continues to satisfy all KKT conditions. Optimality of $\hat{q}^{(m)}$ therefore follows.

4.3 Uniqueness

Whenever at most one site lies on the plateau induced by λ^* the objective is strictly convex on the capacity face, so the algorithm's output is the unique minimiser. If several sites share the same plateau, the common slope of $\phi_i'(q, \lambda^*)$ equals zero, and shifting inventory among those sites preserves both feasibility and the KKT stationarity conditions. The objective therefore remains unchanged, and the set of optima forms a convex polytope. Step 4's cost–ordered greedy fill picks one representative point on this face, but any allocation that (i) keeps each q_i within its plateau $[q_i^{\downarrow}, q_i^{\uparrow}]$ and (ii) satisfies $\sum_i q_i = Q$ is equally optimal.

5 Experiments

Numerical experiments we performed includes the verification of the algorithm and implementation, and the experiments of convergence rates' dependence on the cost coefficients.

5.1 Verification of algorithm and implementation

5.1.1 Simple verification of unconstrained SAA newsvendor with symmetric penalty and demand

First, we simply verify that the unconstrained ($Q = \infty$) SAA newsvendor problem returns the empirical fractile solution (Littlewood's rule with empirical CDF) [3],

$$\hat{q}_i = s_{i, \lceil m \cdot \frac{1}{2} \rceil} = \hat{F}_i^{-1}(0.5)$$

which follows from our formula (23) with $\lambda = 0$ and $b_i = h_i$. This is the sample median of the demand distribution. By choosing a (truncated at zero to negative infinity) multivariate Gaussian distribution with means distributed uniformly on $[50, 150]$ and covariance matrix with entries < 1, the median should approximate the mean for large m, as it is symmetric with good approximation. We indeed see this in Figure 1.

5.1.2 Capacity-constrained SAA newsvendor

In the next experiment, we now change the capacity constraint, which now binds. We also modify the costs to be asymmetric (with high probability), with $h_i, b_i \sim \mathrm{Uniform}(0, 10)$. We again use a Gaussian distribution with means distributed uniformly on $[50, 150]$ and covariance matrix with entries < 1. We now observe more interesting patterns. is given in the title of the plot, and we see that it's binding. For instance, we observe that intuitively, where the lost sales cost are neglible, no allocation is given to this location. Observe Figure 2. The obtained capacity $S(\lambda) = \sum_{i=1}^{k} \hat{q}_i^{(m)}(\lambda)$ In one particular case (for the highest true mean), we see the allocation is slightly above the mean, since the lost sales cost is there dominating. Oppositely, for the location with the lowest true mean, the allocation is lower than the mean, since the holding cost is there dominating. However the skewness for the first case is much less than for the second case. This is a consequence of the binding capacity constraint - it's preferrable to allocate less. Both of these experiments give results that seems intuitive, and serve as a sanity check for our approach and implementation.

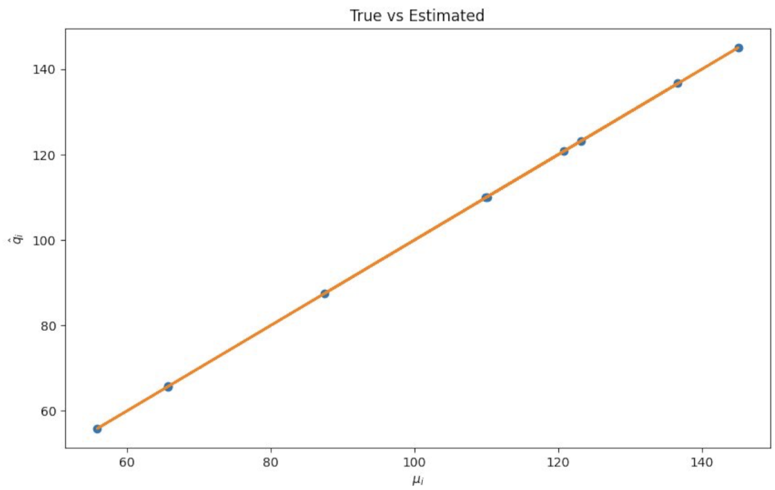

Figure 1: Simple verification of unconstrained SAA newsvendor

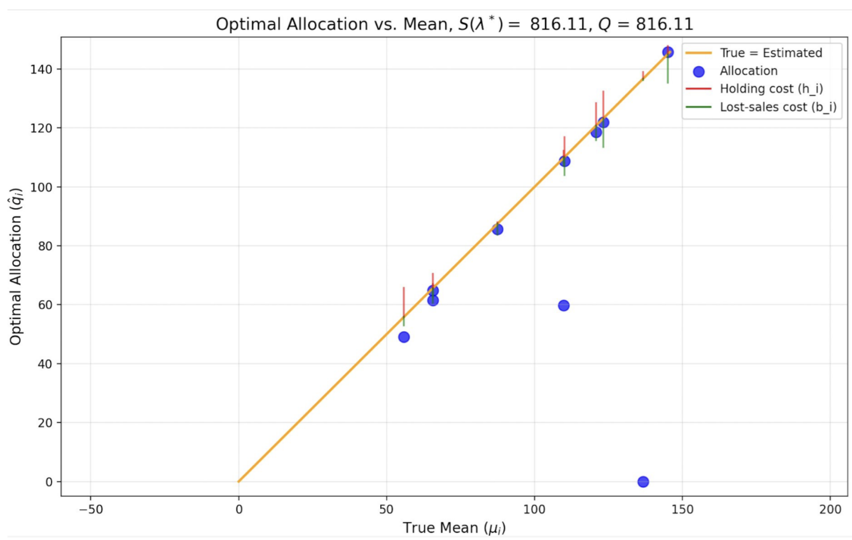

Figure 2: Capacity-constrained SAA newsvendor

5.2 h/b-ratio dependent convergence

We now want to perform some experiments that give some new insights related to the theory derived in Sections 2 and 3. Specifically, we want to relate our cost coefficients h_i and b_i to the value u in Theorem 1, to see if our relative crude bound for u used in Theorems 8 and 9 is asymptotically tight or not. Therefore, we are in this subsection interested in seeing how the convergence rate depends on different values of h_i, b_i. Specifically, we will vary their individual magnitudes, and their ratios. We both measure relative allocation error and relative cost error. These metrics are calculated by computing a reference solution for the allocation and cost, for a large m (in this case $m_{ref} = 2 \cdot 10^5$), and then computing the relative error for different m values. For each value m, we perform r independent runs, and take the average of the results to find the error metrics. For simplicity, we will here have a large capacity $Q = 10^6$, to avoid binding and isolate the effect of the cost coefficients.

While our findings below provide interesting insights, we realized that looking at the *relative* errors does not provide all the insight we hoped for for our *absolute* error bounds. We hope to rerun the experiments in the future, plotting the absolute error, to continue this investigation.

5.2.1 Increasing b, h fixed

In the following experiments, we fix $h = 1$ and vary b from 1 to 1000 (letting $h_i, b_i = h, b, \forall i$), to see how this affects the convergence rate. In Figure 3 and 4, we consider a symmetric distribution (a triangular distribution with support $[0, 100]$) and in Figure 5 and 6, we consider the non-symmetric log-normal distribution. We see in Figure 3 that for the triangular distribution, the relative allocation error seem independent of the ratio. This is not the case for the relative cost error in Figure 4, which is higher for the higher ratios. The rate of convergence in the log log plots seem constant for all ratios in both figures. Moving on to the log-normal distribution in Figure 5 and 6, we see that both errors are higher for the higher ratios, but here the convergence rate does not seem constant in the log log plots. We note that the relative cost values for larger m in general is smaller for the log normal distribution, especially for the lower ratios. This is also the case for the allocation error, which is higher for the lower ratios. For the higher ratios, this does in general not hold, especially for the higher ratios.

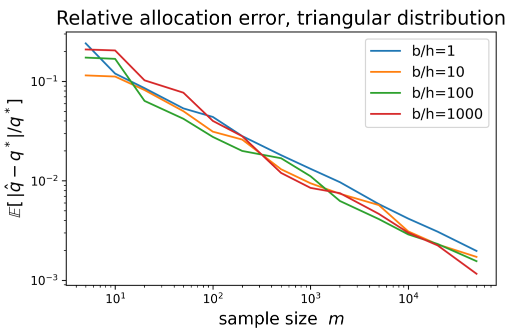

Figure 3: Relative allocation error, triangular distribution

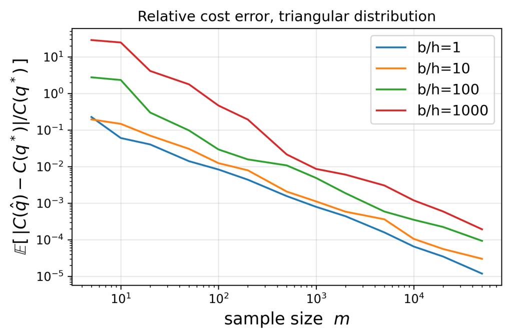

Figure 4: Relative cost error, triangular distribution

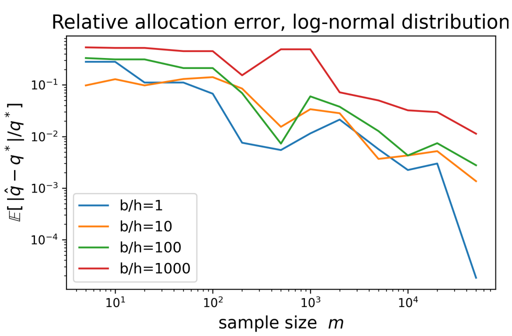

Figure 5: Relative allocation error, log-normal distribution

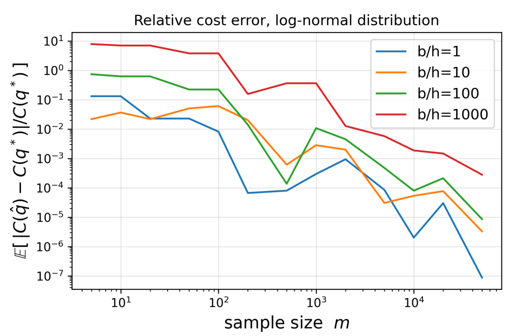

Figure 6: Relative cost error, log-normal distribution

5.2.2 Varying b/h ratios

Here, we vary the ratio by varying both b_i and h_i, and observe how this affects the convergence rate. We again consider the triangular and log-normal distributions (Figure 7 and 8 triangular, and Figure 9 and 10 log-normal). Starting with the triangular distribution, we here observe in Figure 7 that a large h and small b gives much higher allocation error than the other three cases, which have similar performance. In the relative cost error plot, we have larger errors for either large h/b or large b/h (reciprocal). Having $h = b$ gives the same relative cost error independent of magnitude of h and b. This is also prominent in the log-normal distribution 10, which shows the same tendency. The allocation error is here similar, by being constant for constant ratio. This differs from the triangular distribution, where the allocation error is higher for a high h and low b. Across all four plots, the same ratio gives identical relative allocation error and cost error.

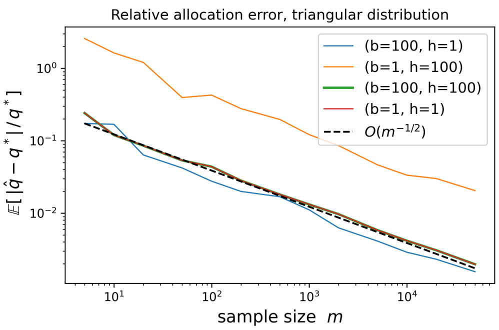

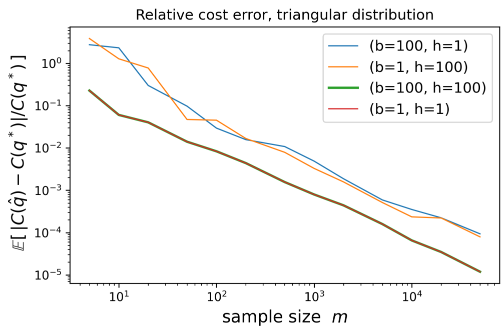

Figure 7: Relative allocation error, triangular distribution

Figure 8: Relative cost error, triangular distribution

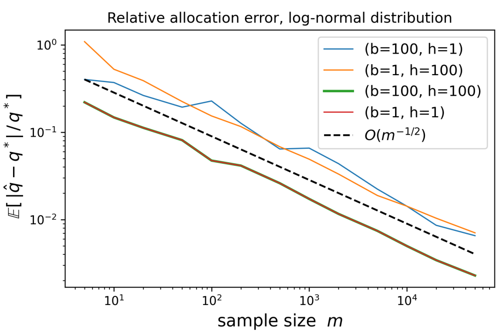

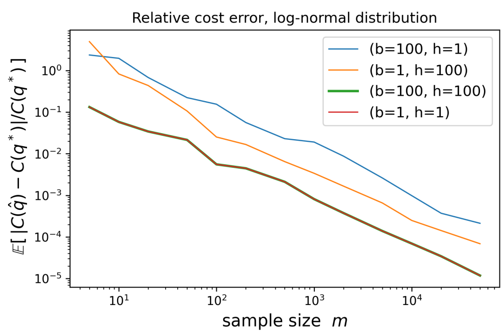

Figure 9: Relative allocation error, log-normal distribution

Figure 10: Relative cost error, log-normal distribution

5.3 Interpreting the results

Raising the lost–sales cost b while keeping h fixed increases this Lipschitz constant and makes the objective more sensitive to quantile estimation error. This explains why the relative cost error in the triangular experiments grows with b even though the allocation error remains essentially unchanged. The latter is possible because a symmetric demand distribution couples the left– and right–tail estimation errors. When d is symmetric around its mean, the empirical p-quantile and the empirical $(1-p)$-quantile tighten at exactly the same rate, so the allocation vector converges at $O(m^{-1/2})$ for every choice of $b/(b+h)$.

For the log–normal demand the heavier right tail makes high quantiles much harder to estimate accurately. As soon as $b/h > 1$ the critical fractile $p = b/(b+h)$ moves further into that tail and the relative allocation error grows, while the corresponding cost error rises even more sharply because the larger b amplifies any residual lost–sales errors.

When we change b and h simultaneously while keeping their ratio constant, the allocation and cost errors in both demand families collapse onto identical trajectories, Figures 7-10. This corroborates that the *ratio* rather than the absolute values governs the shape of the objective surface, at least when errors are measured in relative terms. Extreme ratios ($h \ll b$ or $h \gg b$) tilt the objective so strongly that small sample fluctuations in the favoured tail translate into large changes in q. Hence both cost and allocation errors are largest in those regimes.

Large asymmetries in the penalty parameters demand substantially more samples to reach a given accuracy, especially when demand is skewed. Conversely, if the cost parameters are comparable in magnitude or the demand distribution is light–tailed, the empirical newsvendor show better convergence rate and the performance appears almost insensitive to b/h.

References

[1] Martin Anthony and Peter L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999.

[2] Kenneth J Arrow, Theodore Harris, and Jacob Marschak. Optimal inventory policy. *Econometrica*, 19(3):250–272, 1951.

[3] Hayriye Ayhan, Jim Dai, R. D. Foley, and Joe Wu. Newsvendor notes. Course notes, 2004. ISyE 3232 Stochastic Manufacturing & Service Systems.

[4] Uri Benzion, Yoram Cohen, Ron Peled, and Tal Shavit. Decision-making and the newsvendor problem: An experimental study. *Journal of the Operational Research Society*, 59(9):1281–1287, 2008.

[5] Gary E Bolton and Elena Katok. Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing & Service Operations Management*, 10(3):519–538, 2008.

[6] Christian J. J. Despres. The vapnik-chervonenkis dimension of cubes in $\mathbb{R}^d$, 2017.

[7] R.M Dudley. Balls in rk do not cut all subsets of k + 2 points. *Advances in Mathematics*, 31 (3):306–308, 1979.

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

[9] Guillermo Gallego and Inmoo Moon. The distribution free newsboy problem: Review and extensions. *Journal of the Operational Research Society*, 44(8):825–834, 1993.

[10] Jacob E. Goodman and Joseph O'Rourke, editors. *Handbook of discrete and computational geometry*. CRC Press, Inc., USA, 1997.

[11] G. (George) Hadley. *Analysis of inventory systems*. Prentice-Hall international series in management. Prentice-Hall, Englewood Cliffs, N.J., 1963.

[12] Anton J. Kleywegt and Alexander Shapiro. *The Sample Average Approximation Method for Stochastic Discrete Optimization*. Humboldt-UniversitÃ¤t zu Berlin, Mathematisch-Naturwissenschaftliche FakultÃ¤t II, Institut fÃ¼r Mathematik, 1999.

[13] Retsef Levi, Robin O. Roundy, and David B. Shmoys. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research*, 32(4): 821–839, 2007.

[14] Retsef Levi, Georgia Perakis, and Joline Uichanco. The data-driven newsvendor problem: New bounds and insights. *Operations Research*, 63(6):1294–1306, 2015.

[15] I Moon and E A Silver. The multi-item newsvendor problem with a budget constraint and fixed ordering costs. *The Journal of the Operational Research Society*, 51(5):602–608, 2000.

[16] Márton Naszódi and Steven Taschuk. On the transversal number and vc-dimension of families of positive homothets of a convex body. *Discrete Mathematics*, 310(1):77–82, 2010.

[17] Herbert Scarf. A minimax solution of an inventory problem. In *Studies in the Mathematical Theory of Inventory and Production*, pages 201–209. Stanford University Press, 1958.

[18] Maurice E Schweitzer and Gérard P Cachon. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, 46(3):404–420, 2000.

[19] Daniel Tabak and Benjamin C. Kuo. *Optimal Control by Mathematical Programming*. Prentice-Hall, Englewood Cliffs, NJ, 1971.

[20] Nazli Turken, Yinliang Tan, Asoo Vakharia, Lan Wang, Ruoxuan Wang, and Arda Yenipazarli. *The Multi-Product Newsvendor Problem: Review, Extensions and Directions for Future Research*, volume 176, pages 3–39. 05 2012.

[21] Leslie Valiant. A theory of the learnable. *Comm. ACM*, 11(27):1143–1142, 1967.

[22] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.

[23] Pratkasem Vesarach. Optimization for communications and networks handout b4. Course handout, 2016. August Semester, Course CRCC2.

[24] Yi Wang, Li Xiao, and Xing-Gang Luo and. On the optimal solution of a distributionally robust multi-product newsvendor problem. *Journal of the Operational Research Society*, 74 (12):2578–2592, 2023.

[25] Yaqi Xie, Will Ma, and Linwei Xin. Vc theory for inventory policies, 2024.

Algorithm 1 Capacity–constrained SAA newsvendor

Require: demand samples $\{d_i^{(j)}\}_{j=1,\dots,m}^{i=1,\dots,k}$, costs (b_i, h_i), capacity Q
Ensure: optimal order quantities $\hat{q}_i^{(m)}$ with $\sum_i \hat{q}_i^{(m)} = Q$

Pre-processing
for $i = 1$ **to** k **do**
  sort the m samples to obtain order statistics $(s_{i,1} \le \cdots \le s_{i,m})$
  $\bar{d}_i \leftarrow \frac{1}{m} \sum_{j=1}^{m} d_i^{(j)}$ ▷ sample mean
end for

Auxiliary functions
function $q_i(\lambda)$ ▷ empirical critical-fractile rule
  if $\lambda \ge b_i$ **then return** 0
  else
    $p \leftarrow \dfrac{b_i - \lambda}{b_i + h_i}$
    $\ell \leftarrow \max\{1,\ \lceil m\, p \rceil\}$
    return $s_{i,\ell}$
  end if
end function
define $S(\lambda) = \sum_{i=1}^{k} q_i(\lambda)$

Step 1: check if capacity binds
if $S(0) \le Q$ **then** ▷ unconstrained optimum feasible
  return $\hat{q}_i^{(m)} = q_i(0)$ for all i
end if

Step 2: bisection to bracket the last step
$\lambda_{\min} \leftarrow 0, \quad \lambda_{\max} \leftarrow \max_i b_i$
while $\lambda_{\max} - \lambda_{\min} > \varepsilon_\lambda$ **do**
  $\lambda \leftarrow \frac{\lambda_{\min} + \lambda_{\max}}{2}$
  if $S(\lambda) > Q$ **then**
    $\lambda_{\min} \leftarrow \lambda$
  else
    $\lambda_{\max} \leftarrow \lambda$
  end if
end while
$\lambda_\downarrow \leftarrow \lambda_{\max}$ ▷ $S(\lambda_\downarrow) \le Q$
$\lambda_\uparrow \leftarrow \lambda_{\min}$ ▷ $S(\lambda_\uparrow) > Q$

Step 3: base allocation below the step
for $i = 1$ **to** k **do**
  $q_i^\downarrow \leftarrow q_i(\lambda_\downarrow), \quad q_i^\uparrow \leftarrow q_i(\lambda_\uparrow)$
end for
$L \leftarrow Q - \sum_i q_i^\downarrow$ ▷ left-over capacity, $0 < L \le \sum_i (q_i^\uparrow - q_i^\downarrow)$

Step 4: distribute left-over inside the plateau
$J \leftarrow \{\, i \mid q_i^\uparrow > q_i^\downarrow \,\}$ ▷ jump sites
sort J by descending pair $(b_i, \bar{d}_i)$
for $i \in J$ **do** ▷ cost-aware tie-break
  $\Delta \leftarrow \min\{\, q_i^\uparrow - q_i^\downarrow,\ L \}$
  $q_i^\downarrow \leftarrow q_i^\downarrow + \Delta$
  $L \leftarrow L - \Delta$
  if $L \le 0$ **then break**
  end if
end for

Step 5: return optimal allocation
$\hat{q}_i^{(m)} \leftarrow q_i^\downarrow$ for all i
return $\{\hat{q}_i^{(m)}\}_{i=1}^{k}$