

MIT Open Access Articles

Stochastic optimization in supply chain networks: averaging robust solutions

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published: <https://doi.org/10.1007/s11590-019-01405-0>

Publisher: Springer Berlin Heidelberg

Persistent URL: <https://hdl.handle.net/1721.1/131385>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Stochastic optimization in supply chain networks: averaging robust solutions

Cite this article as: Dimitris Bertsimas and Nataly Youssef, Stochastic optimization in supply chain networks: averaging robust solutions, Optimization Letters <https://doi.org/10.1007/s11590-019-01405-0>

This Author Accepted Manuscript is a PDF file of an unedited peer-reviewed manuscript that has been accepted for publication but has not been copyedited or corrected. The official version of record that is published in the journal is kept up to date and so may therefore differ from this version.

Terms of use and reuse: academic research for non-commercial purposes, see here for full terms. <https://www.springer.com/aam-terms-v1>

Author accepted manuscript

Optimization Letters manuscript No.
(will be inserted by the editor)

Stochastic Optimization in Supply Chain Networks: Averaging Robust Solutions

Dimitris Bertsimas · Nataly Youssef

Received: December 5, 2018

Abstract. We propose a novel robust optimization approach to analyze and optimize the expected performance of supply chain networks. We model uncertainty in the demand at the sink nodes via polyhedral sets which are inspired from the limit laws of probability. We characterize the uncertainty sets by variability parameters which control the degree of conservatism of the model, and thus the level of probabilistic protection. At each level, and following the steps of the traditional robust optimization approach, we obtain worst case values which directly depend on the values of the variability parameters. We go beyond the traditional robust approach and treat the variability parameters as random variables. This allows us to devise a methodology to approximate and optimize the expected behavior via averaging the worst case values over the possible realizations of the variability parameters. Unlike stochastic analysis and optimization, our approach replaces the high-dimensional problem of evaluating expectations with a low-dimensional approximation that is inspired by probabilistic limit laws. We illustrate our approach by finding optimal base-stock and affine policies for fairly complex supply chain networks. Our computations suggest that our methodology (a) generates optimal base-stock levels that match the optimal solutions obtained via stochastic optimization within no more than 4 iterations, (b) yields optimal affine policies which often times exhibit better results compared to optimal base-stock policies, and (c) provides optimal policies that consistently outperform the solutions obtained via the traditional robust optimization approach.

1 Introduction

Supply chain management has received significant attention both in industry and academia. Understanding and optimizing the performance of supply chain net-

F. Author

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139
dbertsim@mit.edu

S. Author

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139
E-mail: youssefn@mit.edu

works is particularly challenging given the uncertainty around the demand. Suppose that we are interested in modeling a measure of the system's performance $C(\pi, \omega)$, where π denotes a given ordering policy and ω represents the vector of uncertain demand. To evaluate and optimize the performance under demand uncertainty, two main avenues have been suggested in the literature: stochastic analysis and optimization describing the uncertainty probabilistically and robust optimization describing the uncertainty deterministically.

Stochastic Approach: The traditional stochastic approach relies on the modeling power of probability theory. Specifically, the demand at each time period is treated as a random variable governed by some posited probability distribution. The most common problem is to assess the expected performance and evaluate

$$\bar{C}(\pi) = \mathbb{E}_{\omega} [C(\pi, \omega)]. \quad (1)$$

Finding the optimal policy under the probabilistic assumptions gives rise to the following stochastic optimization problem

$$\bar{C} = \min_{\pi \in \Pi} \bar{C}(\pi) = \min_{\pi \in \Pi} \mathbb{E}_{\omega} [C(\pi, \omega)], \quad (2)$$

where Π denotes a set of ordering policies. The performance evaluation problem in Eq. (1) and the stochastic optimization problem in Eq. (2) may yield closed-form expressions and analytical solutions for rather simple supply chain systems and under simplifying distributional assumptions over the uncertain demand. For instance, the optimal order quantity for a single period installation that minimizes the expected total cost can be easily expressed as a quantile of the distribution associated with the uncertain demand. However, the more complex the system dynamics, the more challenging it is to derive closed-form expressions. The advances of computing power and memory over the past decades have sprung a wealth of computational techniques to solve such complex problems.

Taking a stochastic programming approach is challenging, given the need to generate scenarios that account for the complex interactions among random variables and the computational difficulties to solve stochastic programs with binary and integer decisions or generally non-linear functions. Simulation optimization has attempted to take advantage of the availability of computational resources and the power of simulation for evaluating functions. For a comprehensive overview of commonly used simulation optimization techniques, we refer the reader to the survey by Fu et al. (2005). Fu (1994), Glasserman and Tayyur (1995), Fu and Healy (1997) and Kapiscinsky and Tayyur (1999) have developed various gradient-based algorithms to study inventory systems. These methods work practically whenever the input variables are continuous and their success depends on the quality of the gradient estimator.

Stochastic optimization is a powerful tool when an accurate probabilistic description of the demand uncertainty is given. However, in many cases, this information is not available. Given this challenge, the field of robust optimization was born in the mid 1990s (see El-Ghaoui and Lebret (1997), El-Ghaoui et al. (1998), Ben-Tal and Nemirovski (1998) and Ben-Tal and Nemirovski (1999)) as an alternative approach for analyzing and optimizing systems under uncertainty.

Robust Optimization Approach: While stochastic optimization views uncertainty probabilistically, the field of robust optimization considers a deterministic model for demand uncertainty by assuming that the uncertain variables lie within some set, referred to as the “uncertainty set”. It then seeks to deterministically immunize the solution against all possible realizations of the uncertain variables satisfying the uncertainty set via a min-max approach (i.e., worst case) as follows

$$\min_{\pi \in \Pi} \max_{\omega \in \mathcal{U}} C(\pi, \omega), \quad (3)$$

where \mathcal{U} denotes the demand uncertainty set. The tractability of the robust optimization problem depends on the choice of the uncertainty set. For example, Ben-Tal and Nemirovski (1998, 1999), El-Ghaoui and Lebret (1997) and El-Ghaoui et al. (1998) proposed linear optimization models with ellipsoidal uncertainty sets, whose robust counterparts correspond to conic quadratic optimization problems. Bertsimas and Sim (2003, 2004) proposed constructing polyhedral uncertainty sets that can model linear variables, and whose robust counterparts correspond to linear optimization problems. Furthermore, Bertsimas and Brown (2009) and Bertsimas et al. (2015) provide guidelines for constructing uncertainty sets from the historical realizations of the random variables using a data-driven approach. For a review of robust optimization, we refer the reader to Ben-Tal et al. (2009) and Bertsimas et al. (2011). The robust framework allows the system designer to adapt the analysis to their risk preferences. By parameterizing different classes of uncertainty sets, one can control the size of the uncertainty set, which provides a notion of a “budget of uncertainty” (see Bertsimas and Sim (2004)). This, in fact, allows the design to control the corresponding level of probabilistic protection, thus choosing the tradeoff between robustness and performance. In this setting, the problem is formulated as

$$\min_{\pi \in \Pi} \max_{\omega \in \mathcal{U}(\Gamma)} C(\pi, \omega), \quad (4)$$

where the variability parameter Γ reflects the degree of conservatism in the model. A growing body of research has applied the robust optimization paradigm to study supply chain networks. Bertsimas and Thiele (2006) and Bienstock and Özbay (2008) studied the performance of base-stock policies, and Ben-Tal et al. (2005), Kuhn et al. (2011), and Bertsimas et al. (2010) investigated policies that are affine in prior demands under a robust optimization lens.

In a recent series of work, Bandi and Bertsimas (2012, 2015a,b, 2014a,b) modeled performance analysis problems in a variety of areas as robust optimization problems. In the same spirit, Bandi et al. (2015a,b) presented a novel approach for modeling the primitives of queueing systems by polyhedral uncertainty sets inspired from the probabilistic limit laws and provided exact characterizations for the steady-state and transient performance analysis of queueing networks. The robust approach generates parametrized solutions (functions of the variability parameter) that matched the conclusions obtained via probabilistic analyses for simple systems and furnished tractable extensions to more complex systems. Capturing the choice of values for the variability parameters to reflect the average performance is however challenging.

Contributions: We propose a novel framework which takes advantage of the power of robust optimization in providing tractable solutions to approximate and

optimize the expected performance of supply chain networks. To model the demand uncertainty, we construct polyhedral sets that are inspired by the limit laws of probability and introduce a variability parameter that controls the size of these sets, and thus the level of probabilistic protection. At each level, we obtain worst case values which directly depend on the values of the variability parameter. We then treat the variability parameter as a random variable and portray the expected behavior by averaging the worst case values over the possible realizations of the variability parameter. This approach

- (a) reduces the dimensionality of the system's uncertainty,
- (b) eliminates generating scenarios to describe the states of randomness,
- (c) does not require simulation replications to evaluate the performance, and
- (d) demonstrates the use of robust optimization to evaluate and optimize expected performance.

In this paper, we demonstrate the merits of our approach in computing optimal base-stock levels. Note that Bandi et al. (2015b) have applied this framework to analyze the transient performance of multi-server queues and feedforward queueing networks, and obtained approximations that are comparable to simulated results for both light and heavy-tailed arrival and service times.

2 Proposed Framework

We consider a supply chain network in which inventories are reviewed periodically and unfulfilled orders are backlogged. For the remainder of this paper, we assume zero lead times and fixed costs throughout the network, no demand seasonality and that the demand realizations are light-tailed in nature (i.e., the demand variance is finite). We consider a T -period time horizon and, within each period, events occur in the following order: (1) the ordering decision is made at the beginning of the period, (2) demands for the period then occur and are filled or backlogged depending on the available inventory, (3) the stock availability is updated for the next period. We view the dynamics of the system from an echelon perspective. Specifically, each echelon is comprised a set of installations that receive stock from a source installation n , including all the links or edges between them (see Fig. 1).




Fig. 1 For this nine-installation network with 4 sink nodes, we consider nine echelons defined as follows. For instance $\mathcal{E}_1 = \{1, 5, 6, 8, 9\}$ and $\mathcal{S}_1 = \{5, 8, 9\}$.

This definition was first introduced by Clark and Scarf (1960) for a supply network in series, however it can be generalized for more complex networks. In the spe-

cial case of a network with installations in series, and assuming that the items transit from installation n to installation $n - 1$, then the sets $\mathcal{E}_n = \{1, 2, \dots, n - 1, n\}$, $\mathcal{S}_n = \{1\}$ and $\mathcal{L}_n = \{\ell_{n+1,n}\}$, where $\ell_{n+1,n}$ is the link between installation $n + 1$ and n . To describe the system dynamics, we define the following notation.

- \mathcal{N} Set of all installations within the supply chain,
- \mathcal{S} Set of all installations with external demand (sink nodes),
- \mathcal{L} Set of all links (edges) within the inventory network,
- \mathcal{E}_n Set of installations belonging to echelon n ,
- \mathcal{S}_n Set of sink installations at the n^{th} echelon. Note that $\mathcal{S}_n \subseteq \mathcal{S}$,
- \mathcal{L}_n Set of all links (or edges) supplying stock to the n^{th} echelon.

- x_n^t Stock available at the beginning of period t and echelon n ,
- u_n^t Total stock ordered at the beginning of period t at echelon n ,
- o_ℓ^t Stock ordered and moved along link $\ell \in \mathcal{L}$ at the beginning of period t ,
- ω_k^t Demand observed at sink $k \in \mathcal{S}$ throughout time period t .

For all $n \in \mathcal{N}$ and $t = 0, \dots, T - 1$,

$$x_n^{t+1} = x_n^t + u_n^t - \sum_{k \in \mathcal{S}_n} \omega_k^t = x_n^0 + \sum_{\tau=0}^t u_n^\tau - \sum_{k \in \mathcal{S}_n} \sum_{\tau=0}^t \omega_k^\tau, \quad (5)$$

x_n^0 denotes the initially available stock at echelon n , and the ordering quantity at each echelon is simply the sum of all stock ordered from the edges feeding into the n^{th} echelon, i.e.,

$$u_n^\tau = \sum_{\ell \in \mathcal{L}_n} o_\ell^\tau. \quad (6)$$

Note that the ordering quantities $x_n^t = x_n^t(\pi, \omega)$, and therefore the amount of available stock $u_n^t = u_n^t(\pi, \omega)$, are functions of the ordering policy π and the demand vector. For a single-installation system, the available stock level at the beginning of time $t + 1$ is a function of the sum of the demand realizations at that installation over the time horizon

$$x^{t+1} = x^t + u^t - \omega^t = x^0 + \sum_{\tau=0}^t u^\tau - \sum_{\tau=0}^t \omega^\tau. \quad (7)$$

The high-dimensional nature of modeling the demand uncertainty probabilistically and the complex dependence of the system on the random variables highlight the difficulty of analyzing and optimizing the expected total cost across the supply chain network. Instead of taking a probabilistic approach, we propose a framework that builds upon the robust optimization framework to approximate the expected system behavior. Specifically, we

- (a) model uncertainty via sets whose size is controlled by a variability parameter,
- (b) treat variability parameters as random variables following some distribution,
- (c) approximate the expected behavior by averaging the worst case values, and
- (d) leverage robust optimization to optimize the system's average performance.

We next present a synopsis of our approach.

2.1 Demand Uncertainty

At installation k , we denote the demand mean by μ_k and the demand standard deviation by σ_k , which could be inferred from historical data. Instead of describing the uncertainty in the demand using stochastic processes, we leverage the partial sums in Eq. (5) and propose polyhedral uncertainty sets inspired by the limit laws of probability. Given that we are interested in modeling the amount of holding stock $(x_n^t)^+ = \max(0, x_n^t)$ and the backorder quantity $(x_n^t)^- = -\min(0, x_n^t)$, we wish to upper and lower bound the partial sums in Eq. (5). We therefore propose to constrain the absolute value of the partial sums and introduce a single variability parameter Γ . Note that one can define variability parameters for each echelon; however, in this effort, we propose to constrain the absolute value of the partial sums and introduce a single variability parameter Γ and show in Section 3.3 that this choice results in a good accuracy and tractability relative to stochastic solutions.

ASSUMPTION 1. We assume that the demand at the sink nodes satisfies.

- (a) For inventory systems with a single sink node, the demand realizations $\omega = (\omega^0, \dots, \omega^T)$ belong to the parametrized uncertainty set

$$\mathcal{U}(\Gamma) = \left\{ (\omega^0, \dots, \omega^T) \left| \frac{1}{\sigma \cdot \sqrt{t}} \cdot \left| \sum_{\tau=0}^{t-1} \omega^\tau - t \cdot \mu \right| \leq \Gamma, \quad \forall t = 1, \dots, T+1 \right. \right\},$$

where $\Gamma \geq 0$ is a parameter that controls the degree of conservatism, μ and σ respectively denote the mean and the standard deviation of the demand realizations at the sink node.

- (b) For inventory systems with multiple sink nodes, the demand realizations $\omega = (\omega_k^0, \dots, \omega_k^T)_{k \in \mathcal{S}}$ belong to the parametrized uncertainty set

$$\mathcal{U}(\Gamma) = \left\{ (\omega_k^0, \dots, \omega_k^T)_{k \in \mathcal{S}} \left| \frac{1}{\sqrt{|\mathcal{S}_n|}} \cdot \left| \sum_{k \in \mathcal{S}_n} \frac{\sum_{\tau=0}^{t-1} \omega_k^\tau - t \cdot \mu_k}{\sigma_k \cdot \sqrt{t}} \right| \leq \Gamma, \quad \forall n \in \mathcal{N}, t = 1, \dots, T+1 \right. \right\},$$

where $\Gamma \geq 0$ is a parameter that controls the degree of conservatism, μ_k and σ_k respectively denote the mean and the standard deviation of the demand at the sink node k .

2.2 Performance Analysis

For a given ordering policy π , analyzing the expected performance $\bar{L}(\pi)$ entails understanding the dependence of the system on the demand uncertainty. Suppose that $L(\pi, \omega)$ is governed by a distribution F which can be derived from the joint distribution over the random variables ω . Instead, inline with Bandi et al. (2015b), we approximate the expected value via averaging the worst case values, i.e.,

$$\bar{C}(\pi) \approx \mathbb{E}_\Gamma \left[\widehat{C}(\pi, \Gamma) \right] \tag{8}$$

where Γ is a random variable following some probability density and $\widehat{C}(\pi, \Gamma)$ is the worst case inventory cost defined as

$$\widehat{C}(\pi, \Gamma) = \max_{\omega \in \mathcal{U}(\Gamma)} C(\pi, \omega). \quad (9)$$

To motivate our selection for the density of Γ in the remainder of this manuscript, we analyze the Lindley-type recursion of the shortfall in a multi-period single-installation system operating under a base-stock policy π . Assuming orders are placed at the beginning of each time period to restore the inventory to a target level S , with a maximum order capacity of κ per period, the shortfall at time $t + 1$ is written recursively as

$$\begin{aligned} L_{t+1} &= S - x^{t+1} = \max\left(S - x^t + \kappa - \omega^t, \omega^t\right) = \max\left(L_t + \omega^t - \kappa, \omega^t\right) \\ &= \omega^t + \max_{0 \leq \tau \leq t-1} \left\{ \sum_{i=\tau}^{t-1} (\omega^i - \kappa), 0 \right\} = \omega^t + M_t, \end{aligned} \quad (10)$$

where x^t is the amount of stock available at the beginning of period t and the ordering quantity at the beginning of time period t can be expressed as $\min(\kappa, S - x^t)$. A standard property of the Lindley recursion implies M_t is the maximum of the random walk Δ_τ . Note that M_t is well approximated by a reflected Brownian motion with drift $(\mu - \kappa)$ and variance σ^2 , implying

$$\mathbb{P}\left(L_{t+1} \leq \ell | \omega^t\right) = \mathbb{P}\left(M_t \leq \ell - \omega^t\right) \approx 2 \cdot \Phi\left(\frac{\ell - \omega^t - (\mu - \kappa)t}{\sigma\sqrt{t}}\right) - 1. \quad (11)$$

By following the approach in Bandi et al. (2015b) for transient queueing applications, we consider the worst case shortfall conditioned on ω^t as follows

$$\widehat{L}_{t+1}(\Gamma) = \omega^t + \max_{\omega \in \mathcal{U}(\Gamma)} \max_{\tau} \left\{ \sum_{i=\tau}^{t-1} (\omega^i - \kappa), 0 \right\} = \omega^t + \Gamma \cdot \sigma\sqrt{t} + (\mu - \kappa)t. \quad (12)$$

Given Eq. (12), we can rewrite the conditional probability in Eq. (11) as

$$\mathbb{P}\left(L_{t+1} \leq \widehat{L}_{t+1}(\Gamma) | \omega^t\right) = 2 \cdot \Phi\left(\frac{\Gamma \cdot \sigma\sqrt{t}}{\sigma\sqrt{t}}\right) - 1 = 2 \cdot \Phi(\Gamma) - 1.$$

As a result, we approximate the density of Γ by a half-normal. In the remainder of the paper, we show that, while our approximation of the density of Γ stems from a simple example, the half-normal density yields solutions that are comparable to the ones obtained via stochastic assumptions.

Note: For complex supply chain networks, the worst case cost may not be determined analytically. Therefore, we propose to approximate the expected value in Eq. (8) by discretizing the space of values that Γ can take on, giving rise to the following approximation

$$\mathbb{E}_\Gamma \left[\widehat{C}(\pi, \Gamma) \right] \approx \sum_{i \in \mathcal{I}} f_i \cdot \widehat{C}(\pi, \Gamma_i), \quad (13)$$

where $(\Gamma_i)_{i \in \mathcal{I}}$ denotes the values of Γ in the discretization \mathcal{I} , f_i denotes the corresponding density, and $\widehat{C}(\pi, \Gamma_i)$ denotes the worst case performance given the demand $\omega \in \mathcal{U}(\Gamma_i)$. To find the weights f_i , $i \in \mathcal{I}$, one could use methods for numerical integration, such as the Gaussian-Hermite quadrature (see Abramowitz and Stegun (1972)).

2.3 Performance Optimization

A major consideration in the study of inventory systems consists of determining optimal policies that minimize the average cost of moving inventory across the supply chain network. We consider three types of costs: the holding cost per unit of inventory hold at echelon n denoted by h_n , the backorder penalty cost per unit of negative inventory at echelon n denoted by p_n , and the variable cost per unit of order moved along edge $\ell \in \mathcal{L}$ denoted by c_ℓ . The total cost incurred in period t across the inventory network accounts for the holding cost at each echelon and the penalty cost associated with a shortage at each echelon, i.e.,

$$C_t(\pi, \omega) = \sum_{\ell \in \mathcal{L}} c_\ell \cdot o_\ell^t + \sum_{n \in \mathcal{N}} \left[h_n \cdot (x_n^t)^+ + p_n \cdot (x_n^t)^- \right], \quad (14)$$

where the terms $(x_n^t)^+ = \max(0, x_n^t)$ and $(x_n^t)^- = -\min(0, x_n^t)$ denote the holding and the backordered stock, respectively. Note that the amount of stock ordered $u_n^t = u_n^t(\pi, \omega)$ and the amount of stock available $x_n^t = x_n^t(\pi, \omega)$ depend on the policy π and the demand realizations.

To obtain an optimal ordering policy from a set of available ordering policies Π , the traditional approach solves the following stochastic optimization problem

$$\bar{C} = \min_{\pi \in \Pi} \mathbb{E}_\omega [C(\pi, \omega)],$$

where $C(\pi, \omega)$ denotes the overall cost summed over all time periods. Instead of taking stochastic approach, we leverage the worst case values and cast the problem of finding an optimal policy as

$$\min_{\pi \in \Pi} \mathbb{E}_\Gamma [\hat{C}(\pi, \Gamma)] \approx \min_{\pi \in \Pi} \sum_{i \in \mathcal{I}} f_i \cdot \hat{C}(\pi, \Gamma_i)$$

where $\hat{C}(\pi, \Gamma_i)$ denotes the worst case total cost of moving inventory through the entire time horizon, given the demand $\omega \in \mathcal{U}(\Gamma_i)$. The above optimization problem can be cast as a robust optimization problem with the following re-formulation

$$\left\{ \begin{array}{l} \min_{\pi \in \Pi} \sum_{i \in \mathcal{I}} f_i \cdot q_i \\ \text{s.t. } q_i \geq C(\pi, \omega) \quad \forall \omega \in \mathcal{U}(\Gamma_i), \text{ and } \Gamma_i : i \in \mathcal{I} \end{array} \right\}. \quad (15)$$

We note that, in the traditional robust optimization setting, the designer selects a particular value of Γ reflecting their risk preference and solves the resulting problem

$$\min_{\pi \in \Pi} \max_{\omega \in \mathcal{U}(\Gamma)} C(\pi, \omega) = \left\{ \begin{array}{l} \min_{\pi \in \Pi} q \\ \text{s.t. } q \geq C(\pi, \omega) \quad \forall \omega \in \mathcal{U}(\Gamma) \end{array} \right\}. \quad (16)$$

Both formulations in Eqs. (15) and (16) belong to the same class of problems. Our proposed approach in Eq. (15) therefore conserves the desirable tractability of the robust optimization approach, while exploring different levels of protection against uncertainty.

Note: The size of the robust optimization problem in Eq. (15) depends on the level of discretization over the space of possible values that Γ can take on.

Quadrature methods help numerically approximate the value of a definite integral with few possible evaluations. Using such methods ensures a good level of precision while keeping control over the size of the discretization set \mathcal{I} .

We propose a variant of the generic algorithm developed by Bienstock and Özbay (2008) to iteratively solve Eq. (15) for the optimal inventory policy. The algorithm maintains a working list $\hat{\mathcal{U}}_i$ of demand patterns $\hat{\omega}^i = \{(\hat{\omega}_k^0)^i, \dots, (\hat{\omega}_k^T)^i\}_{k \in \mathcal{S}}$ that satisfy the uncertainty set $\mathcal{U}(\Gamma_i)$, for all $i \in \mathcal{I}$. At every iteration, we increment the list while computing an upper bound U and a lower bound L on the value of the problem in Eq. (15). The algorithm is stopped whenever the difference between the upper and lower bounds becomes small enough. This algorithm is inspired by the Bender's decomposition method, commonly used in the stochastic optimization literature (see Higle and Sen (1996)).

ALGORITHM (Optimizing the Ordering Policy)

Input: Accuracy level ϵ . Available ordering policies Π .

Output: Optimal policy π^* for the inventory network.

Step 0. Initialize lower bound $LB = 0$, upper bound $UB = +\infty$, and $\hat{\mathcal{U}}_i = \emptyset$, for all $\Gamma_i : i \in \mathcal{I}$.

Step 1. Solve the decision maker's problem (DM) and let π^* to be its optimal solution.

$$LB = \min_{\pi \in \Pi} \sum_{i \in \mathcal{I}} \left[f_i \cdot \max_{\omega \in \hat{\mathcal{U}}_i} \{C(\pi, \omega)\} \right], \quad (17)$$

Step 2. For each $i \in \mathcal{I}$, solve the adversarial problem (AP) and let $\hat{\omega}^i$ be its optimal solution.

$$UB_i = \max_{\omega \in \mathcal{U}(\Gamma_i)} C(\pi^*, \omega), \quad (18)$$

Step 3. Set the upper bound $UB = \sum_{i \in \mathcal{I}} f_i \cdot UB_i$.

Step 4. If $U - L < \epsilon$, exit. Else, add the vector $\hat{\omega}^i$ to $\hat{\mathcal{U}}_i$ for all $i \in \mathcal{I}$ and go to Step 1.

Note that, at a given iteration of the algorithm, the set $\hat{\mathcal{U}}_i$ is finite as it is incrementally populated by the vectors of demand realizations $\hat{\omega}^i$. As a result, the size of the set $\hat{\mathcal{U}}_i$ is equal to the number of iterations run thus far, compared to the exponential size of the uncertainty set $\mathcal{U}(\Gamma_i)$. The size of problem (DM) in Eq. (17) grows with the number of iterations. However, if convergence occurs within a few number of iterations (as shown in Section 3), the size of problem (DM) is kept small. On the other hand, the size of problem (AP) in Eq. (18) is a function of the size of the inventory network. Bienstock and Özbay (2008) present an approximation that uses simple combinatorial arguments which proves more efficient than solving the integer optimization problem. Since the size of \mathcal{I} need not be large to obtain good approximations, the number of problems (AP) that we would need to solve is relatively small. In the stochastic programming framework, Bender's decomposition is used to reduce the large deterministic equivalent to a number of smaller problems that can be solved independently. In our case, the usefulness of the decomposition algorithm lies in reducing the combinatorial complexity of the problem in Eq. (15). We next apply our framework to study generalized inventory networks with base-stock and affinely adaptive ordering policies.

3 Optimizing Base-Stock Policies

The analysis and optimization of (s, S) inventory policies has received considerable attention since the 1950s. The seminal work of Arrow et al. (1951) introduced the multistage periodic review inventory model, where the inventory is reviewed once every period and a decision is made to place an order, if a replenishment is necessary. The (s, S) inventory policy establishes a lower (minimum) stock point s and an upper (maximum) stock point S . When the inventory level on hand drops below s , an order is placed “up to S ”.

3.1 Problem Formulation

We define s_n and S_n to be the lower (minimum) and the upper (maximum) stock points, respectively, at echelon n . In vector form, we refer to the base-stock levels as (\mathbf{s}, \mathbf{S}) across the network’s echelons. Given a set of echelon base-stock levels (s_n, S_n) , the ordering quantity at each time period t at echelon n is given by

$$u_n^t = u_n^t(\mathbf{s}, \mathbf{S}, \boldsymbol{\omega}) = S_n - x_n^t \text{ if } x_n^t \leq s_n, \text{ and } 0 \text{ otherwise,} \quad (19)$$

where x_n^t denotes the stock available at the beginning of time t at echelon n .

Finding the optimal base-stock levels in our framework calls for solving a robust optimization problem of the form of Eq. (15) via decomposition as presented in Section 2 by solving iteratively (a) the adversarial problems (AP), and (b) the decision maker’s problem (DM).

Adversarial Problems: In our setting, problem (AP) consists of solving for the worst case cost given the parameterized uncertainty set $\mathcal{U}(I_i)$ and retrieve the optimal solution $\hat{\boldsymbol{\omega}}^i$ that drives the worst case value. For a given I_i , problem (AP) in Eq. (18) can be re-written as

$$\begin{aligned} \max_{\boldsymbol{\omega} \in \mathcal{U}(I_i)} \quad & \sum_{t=0}^T \sum_{\ell \in \mathcal{L}} c_\ell \cdot o_\ell^t + \sum_{t=0}^T \sum_{n \in \mathcal{N}} \left[h_n \cdot (x_n^t)^+ + p_n \cdot (x_n^t)^- + K_n \cdot \mathbb{1}_{u_n^t > 0} \right] \\ \text{s.t.} \quad & x_n^{t+1} = x_n^t + u_n^t - \sum_{k \in \mathcal{S}_n} \omega_k^t, \quad \forall t, n, \\ & u_n^t = \sum_{\ell \in \mathcal{L}_n} o_\ell^t, \quad \forall t, n, \\ & u_n^t = S_n - x_n^t \text{ if } x_n^t \leq s_n, \text{ and } 0 \text{ otherwise} \quad \forall t, n. \end{aligned}$$

Note that problem (AP) is a non-concave maximization problem and the optimal solution $\hat{\boldsymbol{\omega}}^i$ may not occur at a corner point of the uncertainty set $\mathcal{U}(I_i)$. Furthermore, the structure of the ordering policy involves non-convex ordering constraints. By introducing auxiliary variables, we can formulate problem (AP) as a mixed integer optimization problem (MIO) which can be solved relatively efficiently using available optimization solvers.

Decision Maker’s Problem: At each iteration of the algorithm, problem (DM) consists of finding the best base-stock policy, given a finite collection of demand realizations stored thus far. Specifically, for each index $i \in \mathcal{I}$, we populate the set $\hat{\mathcal{U}}_i$ with the optimal solutions $\hat{\boldsymbol{\omega}}^i$ that we obtain from solving the i^{th} adversarial

problem (AP) at each iteration of the algorithm. Mathematically, we formulate problem (DM) in Eq. (17) as

$$\left\{ \begin{array}{l} \min_{(\mathbf{s}, \mathbf{S})} \quad \sum_{i \in \mathcal{I}} f_i \cdot q_i \\ \text{s.t.} \quad q_i \geq C(\mathbf{s}, \mathbf{S}, \hat{\omega}^i), \quad \forall \hat{\omega}^i \in \hat{\mathcal{U}}_i, i \in \mathcal{I} \end{array} \right\}. \quad (20)$$

3.2 Computational Results

We investigate the performance of our framework relative to simulation and examine the effect of the system's parameters, i.e., time horizon, demand distribution and variability, and network size on the accuracy of our solutions. We consider five network topologies (see Fig. 1).

- Instance (1):** single installation ($|\mathcal{N}| = |\mathcal{S}| = 1$) with normal/lognormal distributed demand, mean $\mu = 100$, and standard deviation $\sigma = 30$.
- Instance (2):** three-installation network with a single sink node ($|\mathcal{N}| = 3, |\mathcal{S}| = 1$) with gamma/uniform distributed demand, mean $\mu_3 = 100$, and standard deviation $\sigma_3 = 30$ (unless otherwise specified),
- Instance (3):** three-installation network with two sink nodes ($|\mathcal{N}| = 3, |\mathcal{S}| = 2$) with demand mean $(\mu_2, \mu_3) = (100, 50)$, standard deviation $(\sigma_2, \sigma_3) = (30, 25)$, and two possible distributional inputs: (a) gamma distributed demand at both sinks, and (b) normal demand at sink 2 and lognormal demand at sink 3,
- Instance (4):** five-installation network with three sink nodes ($|\mathcal{N}| = 5, |\mathcal{S}| = 3$) with demand mean $(\mu_3, \mu_4, \mu_5) = (100, 50, 120)$, standard deviation $(\sigma_3, \sigma_4, \sigma_5) = (30, 25, 40)$, and two possible distributional inputs: (a) lognormal distributed demand at all sinks, and (b) normal, gamma and uniform distributed demand at sinks 3, 4, and 5, respectively,
- Instance (5):** nine-installation network ($|\mathcal{N}| = 9, |\mathcal{S}| = 4$) with the following demand mean $(\mu_5, \mu_7, \mu_8, \mu_9) = (100, 50, 120, 80)$ and standard deviation $(\sigma_5, \sigma_7, \sigma_8, \sigma_9) = (30, 25, 40, 80)$, and two possible distributional inputs: (a) uniform distributed demand at all sinks, and (b) normal, lognormal, gamma and uniform distributed demand at sinks 5, 7, 8, and 9, respectively.




Fig. 2 Network topologies.

Impact of Time Horizon

We consider an instance with a single installation and assume that the fixed cost is zero. In this case, it is well-known that an order-up-to policy is optimal. This is a special case of the (s, S) policy where $s = S$, i.e., an order up to S is placed when the inventory position drops below S . For some given value of S , we (a) simulate the average total cost over T time periods using 10,000 simulation replications of normally distributed demand and report the simulated cost, and (b) approximate the average cost using our framework by applying Eq. (13) and report the approximated cost.




Fig. 3 Simulated (solid line) versus approximated values (dotted line) for a single installation with an order-up-to policy, demand mean $\mu = 150$, standard deviation $\sigma = 30$, holding cost $h = \$2$ and penalty cost $p = \$4$, and zero fixed cost. Simulated values computed for normally distributed demand. Panels (a)–(c) correspond to time horizons (a) $T = 1$, (b) $T = 12$, and (c) $T = 24$.

Fig. 3 compares the simulated values to our approximations for various values of S for a single installation for (a) $T = 1$, (b) $T = 12$, and (c) $T = 24$. Our approximation is closer to simulated values for larger time periods. This is expected given that our uncertainty set in Assumption 1(a) and our approximation of the choice of distribution for the variability parameter Γ rely on the accuracy of the central limit theorem. Furthermore, Figure 3 shows that both simulation and approximation point to similar values of S that minimize the average cost. It is around the optimal order-up-to policy that our approximation yields results that are closest to simulation. The percent errors relative to the optimal simulated values are 19.2%, 6.5% and 4.4% for $T = 1$, $T = 12$ and $T = 24$, respectively.

Impact of Demand Variability

We next compute the optimal base-stock policy (\tilde{s}, \tilde{S}) under our approach. We also evaluate the optimal policies (\hat{s}, \hat{S}) obtained via the traditional robust optimization approach (using Eq. (16)) for different values of Γ . To evaluate the performance of policies (\tilde{s}, \tilde{S}) and (\hat{s}, \hat{S}) against policy (\bar{s}, \bar{S}) , we compute the following quantities. We report the relative errors with respect to the stochastic optimal cost as

$$\frac{\tilde{C} - \bar{C}}{\bar{C}} \times 100 \quad \text{and} \quad \frac{\hat{C} - \bar{C}}{\bar{C}} \times 100.$$

Note that the stochastic solution is computed via a sample-average approximation approach. To illustrate our results, we consider the example of Instance (2) with

Table 1 Solutions and associated costs of interest.

Framework	Optimal Policy	Average Cost
Our Base-Stock Approach	(\tilde{s}, \tilde{S})	$\tilde{C} = \mathbb{E}_{\omega} [C(\tilde{s}, \tilde{S}, \omega)]$
Robust Base-Stock Approach [†]	(\hat{s}, \hat{S})	$\hat{C} = \mathbb{E}_{\omega} [C(\hat{s}, \hat{S}, \omega)]$
Stochastic Approach	(\bar{s}, \bar{S})	$\bar{C} = \mathbb{E}_{\omega} [C(\bar{s}, \bar{S}, \omega)]$

[†] Computed as a function of a given value of Γ .

time horizon $T = 8$, demand mean $\mu = 100$. Fig. 4 compares the percent relative errors obtained using our framework and the robust approach ($\Gamma = 2$ and $\Gamma = 3$) versus stochastic optimization for optimizing base-stock policies. Our approximation compares well with the stochastic solutions. The errors are generally negligible for lower values of σ and tend to increase slightly for larger values of σ , though not exceeding 10%. The robust approach for $\Gamma = 2$ and $\Gamma = 3$ yield larger errors for all considered instances. Note that the effect of variability is more visible for lognormal and gamma distributed demand (see Fig. 4).

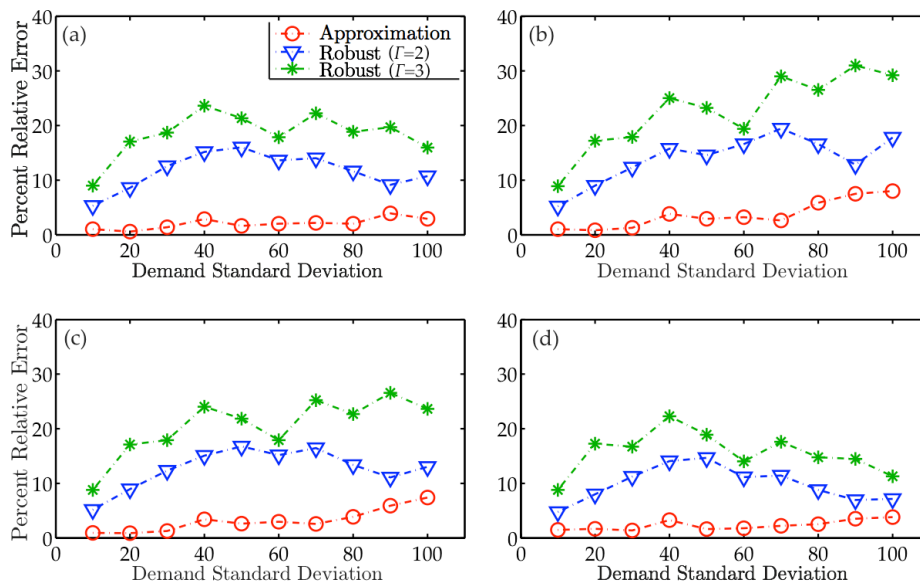


Fig. 4 Percent errors associated with implementing the solutions given by our approximation and the robust optimization approach ($\Gamma = 2$ and $\Gamma = 3$) relative to the optimal stochastic solution. Errors are depicted for Instance (2) with demand mean $\mu = 100$, $T = 8$, while varying the demand standard deviation in the range of $[10, 100]$. Clockwise from the top left panel, we compare the performance to the stochastic instance with the demand at the sink node following a normal, lognormal distribution, gamma, and uniform distributions, respectively.

Impact of Network Size

Table 2 compares the performance of our approach in optimizing base-stock policies for Instances (2)-(5). The solution percent errors generally lie within 5%, suggesting that our approach yields solutions that perform well compared to the stochastic optimal solution for a variety of networks and demand distributions.

Table 2 Errors (%) relative to the stochastic solution.

Instance	Demand [‡]	Base-Stock Solution Percent Error [†]		
		$T = 6$	$T = 9$	$T = 12$
(2)	G	2.28	2.83	2.05
	U	2.33	2.43	1.86
(3)	G	2.64	3.23	2.42
	N,L	3.44	9.38	2.16
(4)	L	2.79	3.37	4.72
	N,G,U	2.41	2.94	4.32
(5)	U	2.07	1.77	1.43
	N,L,G,U	2.05	1.81	1.33

[†] Convergence within 2% gap between the lower and upper bounds.
MIO gap of 2% and 120s time limit allowed for each MIO problem.
[‡] N, L, G, and U stand for normal, lognormal, gamma and uniform.

Computational Performance

Similarly to the observations made by Bienstock and Özbay (2008), the iterative algorithm converges to good solutions within a few iterations. Fig. 5 shows that, for instance (4) with time horizons ranging from $T = 6$ to $T = 12$, the algorithm converges to the solution within 4 iterations.




Fig. 5 Evolution of the lower (solid line) and upper (dotted line) bounds through the iterative algorithm. Panels (a), (b) and (c) correspond to Instance (4) with an (s, S) policy and variable cost for $T = 6$, $T = 9$ and $T = 12$, respectively.

Table 3 Number of iterations and runtime (in seconds).

Instance	$T = 6$		$T = 9$		$T = 12$	
	Iterations	Runtime	Iterations	Runtime	Iterations	Runtime
(2)	4	7.0	2	18.3	4	489.7
(3)	3	7.2	3	75.5	3	448.9
(4)	4	27.2	3	269.1	3	1,112.7
(5)	3	87.8	3	1,185.7	3	1,527.2

[†] Convergence to within 2% gap between the lower and upper bound

4 Concluding Remarks

In this paper, we presented a novel framework inspired by the robust optimization framework to analyze and optimize the expected performance of supply chain networks under uncertainty. Inspired by the probabilistic limit laws, we proposed

Table 4 Errors (%) relative to the stochastic solution for various network instances.

Instance	Demand [‡]	Objective Percent Error [†]			Solution Percent Error [†]		
		$T = 6$	$T = 9$	$T = 12$	$T = 6$	$T = 9$	$T = 12$
(1)	N	3.39	2.09	2.66	0.33	0.41	1.19
	L	1.94	0.97	1.78	4.67	4.85	4.85
(2)	G	3.15	3.12	3.75	2.28	2.83	2.05
	U	2.59	2.43	2.85	2.33	2.43	1.86
(3)	G	5.78	7.41	8.55	2.64	3.23	2.42
	N,L	6.61	8.13	9.56	3.44	9.38	2.16
(4)	L	6.41	9.08	7.43	2.79	3.37	4.72
	N,G,U	5.38	7.80	6.45	2.41	2.94	4.32
(5)	U	16.4	16.4	16.7	2.07	1.77	1.43
	N,L,G,U	18.3	17.9	18.5	2.05	1.81	1.33

[†] Convergence within 2% gap between the lower and upper bounds.

MIO gap of 2% allowed for the decision maker and adversarial problems.

Time limit is set to 120s for the decision maker and adversarial problems.

[‡] N, L, G, and U stand for normal, lognormal, gamma and uniform.

to model randomness via polyhedral uncertainty sets whose size is controlled by a variability parameter. While the traditional robust approach fixes the value of the variability parameter, we treated it as a random variable. Consequently, we devised a methodology that leverages the robust optimization approach and approximates the expected performance via averaging the worst case values. We demonstrated the applicability of our methodology to study and optimize base-stock for complex supply chain networks. Our computations suggest that our approach obtained base-stock levels whose expected performance matches that of optimal base-stock levels obtained via stochastic optimization.

Acknowledgements

The authors would like to thank Chaithanya Bandi and Vishal Gupta for insightful discussions and offering us constructive feedback about our framework.

References

- M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions*. Dover, 1972.
- K. Arrow, T. Harris, and J. Marschak. Optimal inventory policy. *Econometrica*, **XIX**:250–272, 1951.
- C. Bandi and D. Bertsimas. Tractable stochastic analysis via robust optimization. *Mathematical Programming*, 134(1):23–70, 2012.
- C. Bandi and D. Bertsimas. Optimal design for multi-item auctions: A robust optimization approach. *Mathematics of Operations Research*, 39(4):1012–1038, 2014a.
- C. Bandi and D. Bertsimas. Robust option pricing. *European Journal of Operational Research*, 239(3):842–853, 2014b.
- C. Bandi and D. Bertsimas. Channel coding via robust optimization, Part 1: The case of a single channel, submitted for publication. 2015a.
- C. Bandi and D. Bertsimas. Channel coding via robust optimization, Part 2: The multiple channel case, submitted for publication. 2015b.
- C. Bandi, D. Bertsimas, and N. Youssef. Robust queueing theory. *Operations Research*, 63(3): 676–700, 2015a.

- C. Bandi, D. Bertsimas, and N. Youssef. Robust transient analysis of multi-server queues and feedforward queueing networks, submitted for publication. 2015b.
- A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998.
- A. Ben-Tal and A. Nemirovski. Robust solutions to uncertain programs. *Operations Research Letters*, 25:1–13, 1999.
- A. Ben-Tal, B. Golany, A. Nemirovski, and J.-P. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7(3):248–271, 2005.
- A. Ben-Tal, L. El-Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, Princeton, NJ, 2009.
- D. Bertsimas and D. B. Brown. Constructing uncertainty sets for robust linear optimization. *Operations Research*, 57(6):1483–1495, 2009.
- D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98:49–71, 2003.
- D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.
- D. Bertsimas and A. Thiele. Robust and data-driven optimization: Modern decision making under uncertainty. *Tutorials in Operations Research*, pages 95–122, 2006.
- D. Bertsimas, D. Iancu, and P. Parillo. Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research*, 35(2):363–394, 2010.
- D. Bertsimas, D. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53:464–501, 2011.
- D. Bertsimas, V. Gupta, and N. Kallus. Data-driven robust optimization (submitted for publication). 2015.
- D. Bienstock and N. Özbay. Computing robust base-stock levels. *Discrete Optimization*, 5(2):389–414, 2008.
- A. Charnes, W. Cooper, and G. Symonds. Cost horizons and certain equivalents: An approach to stochastic programming of heating oil. *Management Science*, 4(3):235–263, 1958.
- A. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management Science*, 6:475–490, 1960.
- L. El-Ghaoui and H. Lebret. Robust solutions to least-square problems to uncertain data matrices. *SIAM Journal on Matrix Analysis and Applications*, 18:1035–1064, 1997.
- L. El-Ghaoui, F. Oustry, and H. Lebret. Robust solutions to uncertain semidefinite programs. *SIAM Journal on Optimization*, 9:33–52, 1998.
- A. Federgruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands, i: The average cost criterion. *Mathematics of Operations Research*, 11:193–207, 1986.
- M. Fu. Sample path derivatives for (s, s) inventory systems. *Operations Research*, 42:351–364, 1994.
- M. Fu and K. Healy. Techniques for simulation optimization: An experimental study on an (s, S) inventory system. *IIE Transactions*, 29(3):191–199, 1997.
- M. Fu, F. Glover, and J. April. Simulation optimization: a review, new developments, and applications. In *Simulation Conference, 2005 Proceedings of the Winter*, pages 83–95, Dec 2005.
- G. Gallego and I. Moon. The distribution free newsboy problem: Review and extensions. *Journal of Operational Research Society*, 44:825–834, 1993.
- S. J. Garstka and J.-B. Wets. On decision rules in stochastic programming. *Mathematical Programming*, 7(1):117–143, 1974.
- P. Glasserman and S. Tayyur. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science*, 41(2):263–281, 1995.
- S. Graves and S. Willems. Optimizing strategic safety-stock placement in supply chains. *Manufacturing & Service Operations Management*, 2(1):68–83, 2000.
- J. L. Higle and S. Sen. Stochastic decomposition. 1996.

- W. Huh and G. Janakiraman. A sample-path approach to the optimality of echelon order-up-to policies in serial inventory systems. *Operations Research Letters*, 36(5):547–550, 2008.
- R. Kapuscinsky and S. Tayyur. chapter Optimal policies and simulation based optimization for capacitated production inventory systems. Kluwer Academic Publishers, Boston, 1999.
- S. Karlin. Dynamic inventory policy with varying stochastic demands. *Management Science*, 6:231–258, 1960.
- H. Kasugai and T. Kasegai. Note on minimax regret ordering policy – static and dynamic solutions and a comparison to maximin policy. *Journal of Operations Research Society of Japan*, 3:155–169, 1961.
- D. Kuhn, W. Wiesmann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.
- L. Langenhoff and W. Zijm. An analytical theory of multi-echelon production/distribution systems. *Statistica Neerlandica*, 44(3):149–174, 1990.
- T. Morton. The non-stationary infinite horizon inventory problem. *Management Science*, 24:1474–1482, 1978.
- A. Muharremoglu and J. Tsitsiklis. A single-unit decomposition approach to multi-echelon inventory systems. *Operations Research*, 56:1089–1103, 2008.
- A. Rikun. *Applications of robust optimization to queueing and inventory systems*. PhD thesis, Massachusetts Institute of Technology, May 2011.
- K. Rosling. Optimal inventory policies for assembly systems under random demands. *Operations Research*, 37:565–579, 1989.
- H. Scarf. *Studies in the Mathematical Theory of Inventory and Production*, chapter A Min-Max Solution of An Inventory Problems, pages 201–209. Stanford University Press, Stanford, CA, 1958.
- H. Scarf. *Mathematical Methods in the Social Sciences*, chapter The Optimality of (s, S) policies in the dynamic inventory problem. Stanford University Press, Stanford, CA, 1960.
- S. Sethi and F. Cheng. Optimality of (s, S) policies in inventory models with markovian demand. *Operations Research*, 45(6):931–939, 1997.