Submitted to *Management Science*
manuscript (Please, provide the mansucript number!)

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Stochastic Optimization in Supply Chain Networks: Averaging Robust Solutions

(Authors' names blinded for peer review)

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 does not make distributional assumptions regarding the demand uncertainty and bypasses the challenge of generating scenarios. 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.

Key words: Stochastic Optimization, Simulation, Robust Optimization, Inventory Systems, Supply Chain

1. Introduction

Supply chain management has received significant attention both in industry and academia. Understanding and optimizing the performance of supply chain networks is particularly challenging given the uncertainty around the demand. Suppose that we are interested in modeling a measure of the system's performance $L(\pi, \boldsymbol{\omega})$, where π denotes a given ordering policy and $\boldsymbol{\omega}$ 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 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

$$\overline{L}(\pi) = \mathbb{E}_{\boldsymbol{\omega}}\left[L(\pi, \boldsymbol{\omega})\right]. \tag{1}$$

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

$$\overline{L} = \min_{\pi \in \boldsymbol{\Pi}} \overline{L}(\pi) = \min_{\pi \in \boldsymbol{\Pi}} \mathbb{E}_{\boldsymbol{\omega}}\left[L(\pi, \boldsymbol{\omega})\right], \tag{2}$$

where $\boldsymbol{\Pi}$ 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 Tayur (1995), Fu and Healy (1997) and Kapuscinsky and Tayur (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

$$\min_{\pi \in \Pi} \max_{\boldsymbol{\omega} \in \mathcal{U}} L(\pi, \boldsymbol{\omega}), \tag{3}$$

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 et al. (2017) 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

$$\min_{\pi \in \Pi} \max_{\boldsymbol{\omega} \in \mathcal{U}(\Gamma)} L(\pi, \boldsymbol{\omega}), \tag{4}$$

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 polices that are affine in prior demands under a robust optimization lens.

In a series of work reviewed in Bandi and Bertsimas (2012) performance analysis problems in a variety of areas are modeled 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) avoids the challenges of fitting probability distributions,

(b) eliminates generating scenarios to describe the states of randomness,

(c) demonstrates the power of our modeling framework to optimize expected performance.

In this paper, we demonstrate the merits of our approach in computing optimal base-stock levels and assessing the performance of affinely adaptive ordering policies for complex supply chain networks. 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-tailed and heavy-tailed arrival and service times. While our modeling framework builds off of the authors' previous work in Bandi et al. (2015a,b), it is worth mentioning the distinct contributions of this paper. Specifically,

(a) we demonstrate the tractability of *optimizing the expected performance* of complex networks, going beyond performance analysis and system simulation as presented in our earlier work,

(b) we show that, while treating the variability parameter as a random variable, optimizing the expected system performance results in a tractable robust optimization formulation that can be solved via a Bender's decomposition approach, and

(c) we apply our framework to supply chain management, notably the problem of optimizing inventory networks under demand uncertainty.

Structure: The remainder of this paper is structured as follows. Section 2 introduces our uncertainty set modeling assumptions and provides a synopsis of our proposed framework. Section 3 describes our methodology for obtaining optimal (s, S) policies in supply chain networks. Section 4 applies our framework to obtain optimal adaptive policies that minimize the expected total cost across supply chain networks. In section 5, we include concluding remarks.

2. Proposed Framework

We consider a supply chain network in which inventories are reviewed periodically and unfulfilled orders are backlogged. For simplicity, we assume zero lead times throughout the network; however, our framework can be easily applied to systems with non-zero lead times. 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, where an echelon n is defined as the set of all installations within the network that receive stock from some installation n, including installation n, and the links or edges between them. 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. To describe the system dynamics, we define the following sets.

$\mathcal{N}$ 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.

In the special case of a network with installations in series, and assuming that the items transit from installation n to installation $n-1$, then the sets $\mathcal{E}_n = \{n, n-1, \ldots, 1\}$, $\mathcal{S}_n = \{1\}$ and $\mathcal{L}_n = \{\ell_{n+1,n}\}$, where $\ell_{n+1,n}$ is the link between installation $n+1$ and n. To illustrate the echelon definitions for a more complex supply chain network, consider the nine-installation network presented in Figure 1. This network has nine echelons with

(1) $\mathcal{E}_1 = \{1,5,6,8,9\}$, $\mathcal{S}_1 = \{5,8,9\}$, and $\mathcal{L}_1 = \{(0,1),(2,5),(2,6),(3,5),(3,6),(4,6),(7,9)\}$
(2) $\mathcal{E}_2 = \{2,5,6,8,9\}$, $\mathcal{S}_2 = \{5,8,9\}$, and $\mathcal{L}_2 = \{(0,2),(1,5),(1,6),(3,5),(3,6),(4,6),(7,9)\}$
(3) $\mathcal{E}_3 = \{3,5,6,7,8,9\}$, $\mathcal{S}_3 = \{5,7,8,9\}$, and $\mathcal{L}_3 = \{(0,3),(1,5),(1,6),(2,5),(2,6),(4,6),(4,7)\}$
(4) $\mathcal{E}_4 = \{4,6,7,8,9\}$, $\mathcal{S}_4 = \{7,8,9\}$, and $\mathcal{L}_4 = \{(0,4),(1,6),(2,6),(3,6),(3,7),(5,8)\}$
(5) $\mathcal{E}_5 = \{5,8\}$, $\mathcal{S}_5 = \{5,8\}$, and $\mathcal{L}_5 = \{(1,5),(2,5),(3,5),(6,8)\}$
(6) $\mathcal{E}_6 = \{6,8,9\}$, $\mathcal{S}_6 = \{8,9\}$, and $\mathcal{L}_6 = \{(1,6),(2,6),(3,6),(4,6),(5,8),(7,9)\}$
(7) $\mathcal{E}_7 = \{7,9\}$, $\mathcal{S}_7 = \{7,9\}$, and $\mathcal{L}_7 = \{(3,7),(4,7),(6,9)\}$
(8) $\mathcal{E}_8 = \{8\}$, $\mathcal{S}_8 = \{8\}$, and $\mathcal{L}_8 = \{(5,8),(6,8)\}$

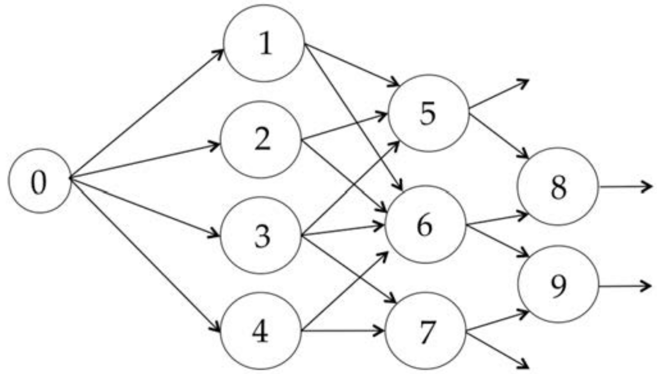

Figure 1 Nine-installation network with four sink nodes and nine echelons.

(9) $\mathcal{E}_9 = \{9\}$, $\mathcal{S}_9 = \{9\}$, and $\mathcal{L}_9 = \{(6,9),(7,9)\}$

To track the system's operation, we capture information about the stock available and the stock ordered at each echelon at the beginning of each time period as well as the demand at each installation sink throughout each time period. Specifically, we define the following notation.

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.

We can express the dynamics of the echelon inventories for all $n \in \mathcal{N}$ and $t = 0, \ldots, T-1$ as

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

where 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.,

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

Note that the ordering quantities $u_n^t = u_n^t(\pi, \boldsymbol{\omega})$, and therefore the amount of available stock $x_n^t = x_n^t(\pi, \boldsymbol{\omega})$, 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

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

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 four types of costs.

K_n Fixed cost of order at echelon n,

h_n Holding cost per unit of inventory hold at echelon n,

p_n Backorder penalty cost per unit of negative inventory at echelon n,

c_ℓ Variable cost per unit of order moved along edge $\ell \in \mathcal{L}$.

The total cost incurred in period t across the inventory network accounts for (1) the holding cost at each echelon, (2) the penalty cost associated with a shortage at each echelon, and (3) the fixed cost of ordering stock at each echelon, i.e.,

$$C_t(\pi, \boldsymbol{\omega}) = \sum_{\ell \in \mathcal{L}} c_\ell \cdot o_\ell^t + \sum_{n \in \mathcal{N}} \left[h_n \cdot \left(x_n^t\right)^+ + p_n \cdot \left(x_n^t\right)^- + K_n \cdot \mathbb{1}_{u_n^t > 0} \right], \tag{8}$$

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

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 the uncertainty via uncertainty sets whose size is controlled by a variability parameter,

(b) treat the variability parameters as random variables following some probability distribution,

(c) approximate the expected system behavior by averaging the worst case values, and

(d) employ the power of robust optimization to optimize the average system performance.

We next present a synopsis of our approach.

2.1. Demand Uncertainty

For the sake of simplicity, we assume that there is no demand seasonality and that the demand realizations are light-tailed in nature (i.e., the demand variance is finite). 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.

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

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

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

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 $\boldsymbol{\omega} = (\omega_k^0, \ldots, \omega_k^T)_{k \in \mathcal{S}}$ belong to the parametrized uncertainty set

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

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.

Note: By the central limit theorem, the expression

$$\frac{1}{\sqrt{|\mathcal{S}_n|}} \cdot \sum_{k \in \mathcal{S}_n} \frac{\sum_{\tau=0}^{t-1} w_k^\tau - t \cdot \mu_k}{\sigma_k \cdot \sqrt{t}}$$

follows a standard normal distribution for a big enough value of t, under the assumption that demand realizations are independent and identically distributed at each sink node $k \in \mathcal{S}$.

Under Assumption 1 and given an ordering policy π, the traditional robust optimization approach analyzes the worst case performance, e.g., the worst case cost, by solving the following optimization problem

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

The optimization problem in Eq. (9) effectively selects the scenario where the realizations of the random variables produce the worst performance. The selection of Γ dictates how much variability we allow the normalized sums to exhibit around zero. With higher variability, the uncertainty set includes more extreme scenarios which directly drive the worst case performance measure.

Instead of pre-selecting a specific value for Γ and carrying out a worst case performance analysis, we propose to treat the variability parameter Γ as a random variable and devise a methodology to model the average system behavior.

2.2. Performance Analysis

For a given ordering policy π, analyzing the expected cost $\overline{C}(\pi)$ entails understanding the dependence of the system on the demand uncertainty. Suppose that $C(\pi, \boldsymbol{\omega})$ is governed by a distribution F which can be derived from the joint distribution over the random variables $\boldsymbol{\omega}$. Then, we can express the expected cost as

$$\overline{C}(\pi) = \int \xi \mathrm{d}F(\xi).$$

For the purpose of our exposition, suppose that the distribution function is continuous. The inverse of $F(\cdot)$ then corresponds to the quantile function, which we denote by

$$Q(p) = F^{-1}(p) = \Big\{ q : F(q) = p \Big\} = \Big\{ q : \mathbb{P}\left(C(\pi, \boldsymbol{\omega}) \leq q\right) = p \Big\},$$

for some probability level $p \in (0,1)$. By a simple variable substitution, we can view the expected value as an "average" of quantiles,

$$\overline{C}(\pi) = \int_0^1 Q(p)\mathrm{d}p.$$

We can map each quantile value $Q(p)$ to a corresponding worst case value $\widehat{L}(\pi, \mathit{\Gamma})$. Let G denote the function that maps p to $\mathit{\Gamma}$ such that $Q(p) = \widehat{L}(\pi, \mathit{\Gamma})$, i.e.,

$$p = \mathbb{P}\left(C(\pi, \boldsymbol{\omega}) \leq \widehat{C}(\pi, \mathit{\Gamma}) \right) = F\left(\widehat{C}(\pi, \mathit{\Gamma}) \right) = G(\mathit{\Gamma}). \tag{10}$$

In this context, the expected value can be written as an average over the worst case values, with

$$\overline{C}(\pi) = \mathbb{E}_{\mathit{\Gamma}}\left[C(\pi, \mathit{\Gamma})\right] = \int \widehat{C}(\pi, \mathit{\Gamma})\, \mathrm{d}G(\mathit{\Gamma}). \tag{11}$$

Philosophically, our averaging approach distills the probabilistic information contained in the random variables $\boldsymbol{\omega}$ into $\mathit{\Gamma}$, hence allowing a significant dimensionality reduction of the uncertainty. This in turn yields a tractable approximation of the expected system cost by reducing the problem to solving a low-dimensional integral.

Note that knowledge of G allows us to compute the expected cost $\overline{C}(\pi)$ exactly; this however depends on the knowledge of the distribution function F. While feasible for simple systems, characterizing F, and therefore G, is otherwise challenging and is immediately dependent on the distributional assumptions over the random variables $\boldsymbol{\omega}$. Instead of deriving the exact distribution $G(\cdot)$, we propose an approximation $\widehat{G}(\cdot)$ inspired by the conclusions of probability theory and approximate the expected cost as

$$\overline{C}(\pi) \approx \int \widehat{C}(\pi, \mathit{\Gamma})\, \mathrm{d}\widehat{G}(\mathit{\Gamma}). \tag{12}$$

We next approximate the distribution of the variability parameter $\mathit{\Gamma}$ by considering a single installation system with a simple re-ordering policy.

Variability Distribution

Let us first consider the simple case of a multi-period single installation system that operates under a deterministic policy π, in which stock is replenished at the beginning of each time period with equal orders of $u=\mu$ units per time period, i.e., $u^t = u = \mu, \forall t = 1,\ldots,T$, where μ represents the mean of the demand at the sink node. Under this deterministic policy, the recursion in Eq. (7) becomes

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

For simplicity, let us assume that the system has no stock on hand at $t=0$, which results in

$$x^t = -\sum_{\tau=0}^{t-1} (\omega^\tau - \mu).$$

We are interested in minimizing the average cost of moving inventory across this system. To simplify the presentation, we now focus on the simple case where the holding and backorder penalty costs are equal, i.e., $h = p = \kappa$, and we assume that there is no fixed or variable costs. As a result, the cost function in Eq. (8) becomes

$$C_t(\pi, \boldsymbol{\omega}) = \kappa \cdot \left|x^t\right| = \kappa \cdot \left|\sum_{\tau=0}^{t-1} (\omega^\tau - \mu)\right|.$$

We make the following observations:

Cost Distribution: Under the assumption of independent and identically distributed demands over the time horizon T, and finite demand variance, the term x^t for large t follows a normal distribution centered around zero, with a standard deviation $\sigma \cdot \sqrt{t}$, where σ denotes the standard deviation of the demand at the sink node. As a result, the cost function $C_t(\pi, \boldsymbol{\omega})$ follows a half-normal distribution with

$$C_t(\pi, \boldsymbol{\omega}) \sim \kappa \cdot \sigma \cdot \sqrt{t} \cdot |Z|, \text{where } Z \sim \mathcal{N}(0,1), \text{for sufficiently large } t. \tag{13}$$

Note that looking at the expected cost under our scenario, we observe that

$$\overline{C}_t(\pi) = \mathbb{E}\left[C_t(\pi, \boldsymbol{\omega})\right] = \kappa \cdot \sigma \cdot \mu_{|Z|} \cdot \sqrt{t},$$

where $\sigma_{|Z|}$ denotes the mean of a standard half-normal random variable.

Robust Approach: We consider the worst case cost, formulated as follows:

$$\widehat{C}_t(\pi, \mathit{\Gamma}) = \max_{\boldsymbol{\omega} \in \mathcal{U}(\mathit{\Gamma})} \kappa \cdot \left|x^t\right| = \max_{\boldsymbol{\omega} \in \mathcal{U}(\mathit{\Gamma})} \kappa \cdot \left|\sum_{\tau=0}^{t-1} (\omega^\tau - \mu)\right| = \kappa \cdot \sigma \cdot \mathit{\Gamma} \cdot \sqrt{t}.$$

Applying Eq. (10), we obtain the following representation of the distribution of $\mathit{\Gamma}$

$$G(\mathit{\Gamma}) = \mathbb{P}\left(C_t(\pi, \boldsymbol{\omega}) \le \widehat{C}_t(\pi, \mathit{\Gamma})\right) = \mathbb{P}\left(\sigma \cdot \sqrt{t} \cdot |Z| \le \kappa \cdot \sigma \cdot \mathit{\Gamma} \cdot \sqrt{t}\right) = \mathbb{P}\left(|Z| \le \mathit{\Gamma}\right) = 2 \cdot \mathit{\Phi}(\mathit{\Gamma}) - 1,$$

where $\Phi(\cdot)$ denotes the cumulative distribution of a standard normal. For this simple scenario, the variability parameter Γ is viewed as a random variable following a half-normal distribution with

$$g(\Gamma) = \frac{\mathrm{d}G(\Gamma)}{\mathrm{d}\Gamma} = 2\phi(\Gamma) = \frac{\sqrt{2}}{\sqrt{\pi}} \cdot \exp\left(-\frac{\Gamma^2}{2}\right).$$

We employ the above approximation for the distribution of Γ throughout the remaining of this paper. Next, we continue with the simple scenario discussed above, but touch upon the analysis of the expected total cost over the time horizon T, which can be defined as

$$\overline{C}(\pi) = \mathbb{E}\left[\sum_{t=0}^{T} C_t(\pi, \boldsymbol{\omega})\right] = \sum_{t=0}^{T} \mathbb{E}\left[C_t(\pi, \boldsymbol{\omega})\right] \approx \kappa \cdot \sigma \cdot \mu_{|Z|} \cdot \sum_{t=0}^{T} \sqrt{t}, \tag{14}$$

where the approximation is due to the approximating the cost distribution at smaller values of t using the probabilistic central limit law. We note that, in practice, for demand distributions without heavy tails, the central limit law is a good approximation numerically starting $t = 30$.

On the other hand, the worst case total cost of the time horizon is formulated as follows

$$\widehat{C}(\pi, \Gamma) = \max_{\boldsymbol{\omega} \in \mathcal{U}(\Gamma)} \sum_{t=0}^{T} C_t(\pi, \boldsymbol{\omega}) = \max_{\boldsymbol{\omega} \in \mathcal{U}(\Gamma)} \sum_{t=0}^{T} \kappa \cdot \left|\sum_{\tau=0}^{t-1} (\omega^\tau - \mu)\right| \tag{15}$$

$$\leq \sum_{t=0}^{T} \max_{\boldsymbol{\omega} \in \mathcal{U}(\Gamma)} \kappa \cdot \left|\sum_{\tau=0}^{t-1} (\omega^\tau - \mu)\right| \tag{16}$$

$$= \sum_{t=0}^{T} \left(\kappa \cdot \sigma \cdot \Gamma \cdot \sqrt{t}\right) = \kappa \cdot \sigma \cdot \Gamma \cdot \sum_{t=0}^{T} \sqrt{t}. \tag{17}$$

We next show that the bound in Eq. (16) is in fact tight. To see this, consider the demand values

$$\widehat{\omega}^t = \mu + \sigma \cdot \Gamma \cdot \left(\sqrt{t+1} - \sqrt{t}\right), \forall t = 1, \ldots, T. \tag{18}$$

We make the following observations:

(1) Feasibility: The above demand values belong to the uncertainty set, i.e., $\widehat{\boldsymbol{\omega}} = \{\widehat{\omega}^0, \ldots, \widehat{\omega}^T\} \in \mathcal{U}(\Gamma)$. This can be easily observed since

$$\frac{1}{\sigma \cdot \sqrt{t}} \cdot \left|\sum_{\tau=0}^{t-1} \omega^\tau - t \cdot \mu\right| = \frac{1}{\sigma \cdot \sqrt{t}} \cdot \left|\sum_{\tau=0}^{t-1} \left[\mu + \sigma \cdot \Gamma \cdot \left(\sqrt{\tau+1} - \sqrt{\tau}\right)\right] - t \cdot \mu\right|$$

$$= \frac{1}{\sigma \cdot \sqrt{t}} \cdot \left|\sigma \cdot \Gamma \sqrt{t}\right| = \Gamma, \ \forall t = 1, \ldots, T. \tag{19}$$

As a result, the demand values presented in Eq. (18) are feasible to the maximization problem in Eq. (15). Next, we show that it is in fact optimal.

(2) Optimality: For each $t = 1, \ldots, T$, the partial sum is met at equality in Eq. (19), therefore the demand values $\{\widehat{\omega}^0, \ldots, \widehat{\omega}^t\}$ are optimal solutions to the maximization problem

$$\max_{\boldsymbol{\omega} \in \mathcal{U}(\Gamma)} \kappa \cdot \left|\sum_{\tau=0}^{t-1} (\omega^\tau - \mu)\right|, \ \forall t = 1, \ldots, T.$$

As a result, $\widehat{\boldsymbol{\omega}}$ is an optimal solution to Eq. (16). Since Eq. (16) provides an upper bound on $\widehat{C}(\pi)$, and given that $\widehat{\boldsymbol{\omega}}$ is a feasible solution to the maximization problem in Eq. (15), we deduce that the bound in Eq. (16) is in fact tight, yielding

$$\widehat{C}(\pi,\Gamma)=\sum_{t=0}^{T}\kappa\cdot\left|\sum_{\tau=0}^{t-1}(\widehat{\omega}^{\tau}-\mu)\right|=\kappa\cdot\sigma\cdot\Gamma\cdot\sum_{t=0}^{T}\sqrt{t}. \tag{20}$$

Averaging over the worst case values with respect to Γ, under the assumption that Γ is a half-normal, yields

$$\mathbb{E}_{\Gamma}\left[\widehat{C}(\pi,\Gamma)\right]=\kappa\cdot\sigma\cdot\mu_{|Z|}\cdot\sum_{t=0}^{T}\sqrt{t}. \tag{21}$$

Notice that, under our distributional assumptions over the variability parameter, averaging the worst case values renders the same value as our approximation of the probabilistic expected cost presented in Eq. (14). We next discuss in further details how we use this framework to approximate the expected system performance.

Robust Approximation

We propose to approximate the expected cost as

$$\widetilde{C}\left(\pi\right)=\mathbb{E}_{\Gamma}\left[\widehat{C}\left(\pi,\Gamma\right)\right], \tag{22}$$

where Γ follows a half-normal distribution. Note that, for complex supply chain networks, the worst case cost may not be determined analytically. Therefore, we propose to approximate the expected value in Eq. (22) by discretizing the space of values that Γ can take on, giving rise to the following approximation

$$\mathbb{E}_{\Gamma}\left[\widehat{C}\left(\pi,\Gamma\right)\right]\;\approx\;\sum_{i\in\mathcal{I}}f_i\cdot\widehat{C}\left(\pi,\Gamma_i\right), \tag{23}$$

where $(\Gamma_i)_{i\in\mathcal{I}}$ denotes the values of Γ in the discretization $\mathcal{I}$, f_i denotes the corresponding density. To find the weights f_i, $i\in\mathcal{I}$, one could use methods for numerical integration, such as the Gaussian-Hermite quadrature, with

$$f_i=\frac{2^n n!}{n^2\left(H_{n-1}\left(\Gamma_i/\sqrt{2}\right)\right)^2},$$

where $n=2\,|\mathcal{I}|$ denotes the level of discretization, $H_{n-1}(\cdot)$ is the Hermite polynomial with degree n, and Γ_i denote the non-negative roots associate with H_n.

Note: The discretization need not include a large number of values to obtain a very accurate approximation of the integral. In fact, a Gaussian-Hermite approximation with $\mathcal{I}=5$ yields values within 1% relative to the exact expected value. This implies that we can achieve good approximations of average cost in our framework by evaluating the worst case performance for a small number of values of Γ.

2.3. Performance Optimization

To obtain an optimal ordering policy from a set of available ordering policies $\boldsymbol{\Pi}$, the traditional approach solves the following stochastic optimization problem

$$\overline{C} = \min_{\pi \in \boldsymbol{\Pi}} \mathbb{E}_{\boldsymbol{\omega}} \left[C(\pi, \boldsymbol{\omega}) \right].$$

Instead, we leverage the worst case values and cast the problem of finding an optimal policy as

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

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

$$\left\{ \begin{array}{ll} \min\limits_{\pi \in \boldsymbol{\Pi}} & \sum\limits_{i \in \mathcal{I}} f_i \cdot y_i \\ \text{s.t.} & y_i \geq C(\pi, \boldsymbol{\omega}) \quad \forall \boldsymbol{\omega} \in \mathcal{U}(\Gamma_i), \text{ and } \Gamma_i : i \in \mathcal{I} \end{array} \right\}. \tag{24}$$

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

$$\min_{\pi \in \Pi} \max_{\boldsymbol{\omega} \in \mathcal{U}(\Gamma)} C(\pi, \boldsymbol{\omega}) = \left\{ \begin{array}{ll} \min\limits_{\pi \in \Pi} & y \\ \text{s.t.} & y \geq C(\pi, \boldsymbol{\omega}) \quad \forall \boldsymbol{\omega} \in \mathcal{U}(\Gamma) \end{array} \right\}. \tag{25}$$

Both formulations in Eqs. (24) and (25) belong to the same class of problems. Our proposed approach in Eq. (24) 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. (24) 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. (24) for the optimal inventory policy. The algorithm maintains a working list $\widehat{\mathcal{U}}_i$ of demand patterns $\widehat{\boldsymbol{\omega}}^i = \{(\widehat{\omega}_k^0)^i, \ldots (\widehat{\omega}_k^T)^i\}_{k \in \mathcal{S}}$ 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. (24). 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)).

Note that, at a given iteration of the algorithm, the set $\widehat{\mathcal{U}}_i$ is finite as it is incrementally populated by the vectors of demand realizations $\widehat{\boldsymbol{\omega}}^i$. As a result, the size of the set $\widehat{\mathcal{U}}_i$ is equal to the number of iterations run thus far. The size of problem (DM) in Eq. (26) 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.

Algorithm **(Optimizing the Ordering Policy)**

Input: Accuracy level ϵ. Available ordering policies $\boldsymbol{\Pi}$.

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

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

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

$$LB \;=\; \min_{\pi \in \boldsymbol{\Pi}} \sum_{i \in \mathcal{I}} \left[f_i \cdot \max_{\boldsymbol{\omega} \in \widehat{\mathcal{U}}_i} \left\{ C\left(\pi, \boldsymbol{\omega}\right) \right\} \right], \tag{26}$$

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

$$UB_i = \max_{\boldsymbol{\omega} \in \mathcal{U}(\Gamma_i)} C\left(\pi^*, \boldsymbol{\omega}\right), \tag{27}$$

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 $\widehat{\boldsymbol{\omega}}^i$ to $\widehat{\mathcal{U}}_i$ for all $i \in \mathcal{I}$ and go to Step 1.

On the other hand, the size of problem (AP) in Eq. (27) 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. (24). 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".

The (s, S) ordering policy is proven optimal for simple stochastic inventory systems. In 1960, Scarf (1960) proved that base-stock policies are optimal for a single installation model. Clark and Scarf (1960) extended the result for serial supply chains without capacity constraints and showed that the optimal ordering policy for the multiechelon system can be decomposed into decisions based on the echelon inventories. Karlin (1960) and Morton (1978) showed that base-stock policies are optimal for single-state systems with non-stationary demands. Federgruen and Zipkin (1986) generalized the analysis to a single-stage capacitated system, and Rosling (1989) extended the analysis of serial systems to assembly systems. Further work has been done to extend, refine and generalize the optimality results of base-stock policies; see Langenhoff and Zijm (1990), Sethi and Cheng (1997), Muharremoglu and Tsitsiklis (2008), Huh and Janakiraman (2008). Determining the optimal policy for general supply chain networks is a challenging problem. It involves a complex stochastic optimization problem with a high-dimensional state space. This sparked interest in simulation-based approaches, notably the work of Glasserman and Tayur (1995) and Fu (1994).

In reality, we only have access to historical demand realizations, and it is not immediately clear which distribution drives the source of uncertainty. In that regard, Scarf (1958), Kasugai and Kasegai (1961), Gallego and Moon (1993), Graves and Willems (2000) developed distribution-free approaches to inventory theory. Bertsimas and Thiele (2006) first took a robust optimization approach to inventory theory and have shown that base-stock policies are optimal in the case of serial supply chain networks. Bienstock and Özbay (2008) presented a family of decomposition algorithms aimed at solving for the optimal base-stock policies using a robust optimization approach. Rikun (2011) extended the robust framework framework introduced by Bienstock and Özbay (2008) to compute optimal (s, S) policies in supply chain networks and compared their performance to optimal policies obtained via stochastic optimization.

In this section, we employ the methodology we proposed in Section 2 to compute optimal base-stock policies that minimize the average cost within the inventory network, without making distributional assumptions regarding the demand uncertainty.

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

$$u_n^t = u_n^t(\mathbf{s}, \mathbf{S}, \boldsymbol{\omega}) = \begin{cases} S_n - x_n^t, & \text{if } x_n^t \leq s_n, \\ 0, & \text{otherwise,} \end{cases} \tag{28}$$

where $x_n^t = x_n^t(\mathbf{s}, \mathbf{S}, \boldsymbol{\omega})$ 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. (24). Specifically, we consider the following formulation

$$\left\{ \begin{array}{lll} \min\limits_{(\mathbf{s},\mathbf{S})} & \sum\limits_{i \in \mathcal{I}} f_i \cdot y_i & \\ \text{s.t.} & y_i \geq C(\mathbf{s}, \mathbf{S}, \boldsymbol{\omega}) & \forall \boldsymbol{\omega} \in \mathcal{U}(\varGamma_i) \text{ and } \varGamma_i : i \in \mathcal{I} \end{array} \right\}, \tag{29}$$

where the total cost across the inventory network is given by

$$C(\mathbf{s}, \mathbf{S}, \boldsymbol{\omega}) = \sum_{t=0}^{T} \sum_{\ell \in \mathcal{L}} c_\ell \cdot o_\ell^t + \sum_{t=0}^{T} \sum_{n \in \mathcal{N}} \left[h_n \cdot (x_n^t)^+ + p_n \cdot (x_n^t)^- + K_n \cdot \mathbb{1}_{u_n^t > 0} \right], \tag{30}$$

with o_ℓ^t, x_n^t, and u_n^t are functions of $(\mathbf{s}, \mathbf{S}, \boldsymbol{\omega})$, for all values of n and t. We solve the problem in Eq. (29) 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 $\mathcal{U}(\varGamma_i)$ and retrieve the optimal solution $\widehat{\boldsymbol{\omega}}^i$ that drives the worst case value. For a given $\varGamma_i$, problem (AP) in Eq. (27) can be re-written as

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

Note that problem (AP) is a non-concave maximization problem and the optimal solution $\widehat{\boldsymbol{\omega}}^i$ may not occur at a corner point of the uncertainty set $\mathcal{U}(\varGamma_i)$. Furthermore, the structure of the ordering policy involves non-convex ordering constraints.

By introducing the following two sets of auxiliary binary variables

$$y_n^t = \left\{ \begin{array}{ll} 1, & \text{if } x_n^t \leq s_n \\ 0, & \text{otherwise} \end{array} \right\}, \quad \text{and} \quad z_n^t = \left\{ \begin{array}{ll} 1, & \text{if } x_n^t > 0 \\ 0, & \text{otherwise} \end{array} \right\},$$

we can formulate problem (AP) as a mixed integer optimization problem which can be solved relatively efficiently using available optimization solvers. Constraints (29)-(30) linearize the term associated with the amount of holding stock $(x_n^t)^+$, constraints (31)-(32) linearize the term associated with the amount of backordered stock $(x_n^t)^-$, and constraints (33)-(35) provide a linear description of the dynamics associated with a base-stock policy.

$$
\begin{aligned}
\max_{\boldsymbol{\omega}\in\mathcal{U}(\Gamma_i)} \quad & \sum_{t=0}^{T}\sum_{\ell\in\mathcal{L}} c_\ell \cdot o_\ell^t + \sum_{t=0}^{T}\sum_{n\in\mathcal{N}} \left[h_n\cdot (x_n^t)^+ + p_n\cdot (x_n^t)^- + K_n\cdot y_n^t\right] \\
\text{s.t.} \quad & \forall t=0,\dots,T \text{ and } n\in\mathcal{N}: \\
& x_n^{t+1} = x_n^t + u_n^t - \sum_{k\in\mathcal{S}_n}\omega_k^t, \\
& u_n^t = \sum_{\ell\in\mathcal{L}_n} o_\ell^t, \\
& x_n^t \le (x_n^t)^+ \le x_n^t + M\cdot(1-z_n^t), && (31)\\
& 0 \le (x_n^t)^+ \le M\cdot z_n^t, && (32)\\
& -x_n^t \le (x_n^t)^- \le -x_n^t + M\cdot z_n^t, && (33)\\
& 0 \le (x_n^t)^- \le M\cdot(1-z_n^t), && (34)\\
& -M\cdot y_n^t \le x_n^t - s_n \le M\cdot(1-y_n^t), && (35)\\
& -M\cdot(1-y_n^t) \le u_n^t - (S_n - x_n^t) \le M\cdot(1-y_n^t), && (36)\\
& 0 \le u_n^t \le M\cdot y_n^t, && (37)\\
& y_n^t, z_n^t \in\{0,1\}. && (38)
\end{aligned}
$$

Note that we may devise an algorithm to approximately solve problem (AP); see for instance the work by Bienstock and Özbay (2008).

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 $\widehat{\mathcal{U}}_i$ with the optimal solutions $\widehat{\boldsymbol{\omega}}^i$ that we obtain from solving the i^{th} adversarial problem (AP) at each iteration of the algorithm. Mathematically, we formulate problem (DM) in Eq. (26) as

$$
\left\{\begin{array}{ll}
\min\limits_{(\mathbf{s},\mathbf{S})} & \sum\limits_{i\in\mathcal{I}} f_i\cdot q_i \\
\text{s.t.} & q_i \ge C\left(\mathbf{s},\mathbf{S},\widehat{\boldsymbol{\omega}}^i\right), \quad \forall \widehat{\boldsymbol{\omega}}^i\in\widehat{\mathcal{U}}_i, i\in\mathcal{I}
\end{array}\right\}, \qquad (39)
$$

where the total cost across the inventory network is given by Eq. (30).

Note that the size of problem (DM) grows with the number of iterations needed for the algorithm to converge. For a small number of iterations, solving the integer optimization problem may not constitute a challenge. In fact, as our computations suggest, the algorithm converges within an accuracy of 2% in no more than 4 iterations.

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 Figure 2).

Instance (1): single installation ($|\mathcal{N}| = |\mathcal{S}| = 1$) with normal/lognormal distributed demand, mean $\mu = 100$, and standard deviation $\sigma = 30$ (unless otherwise specified)

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.

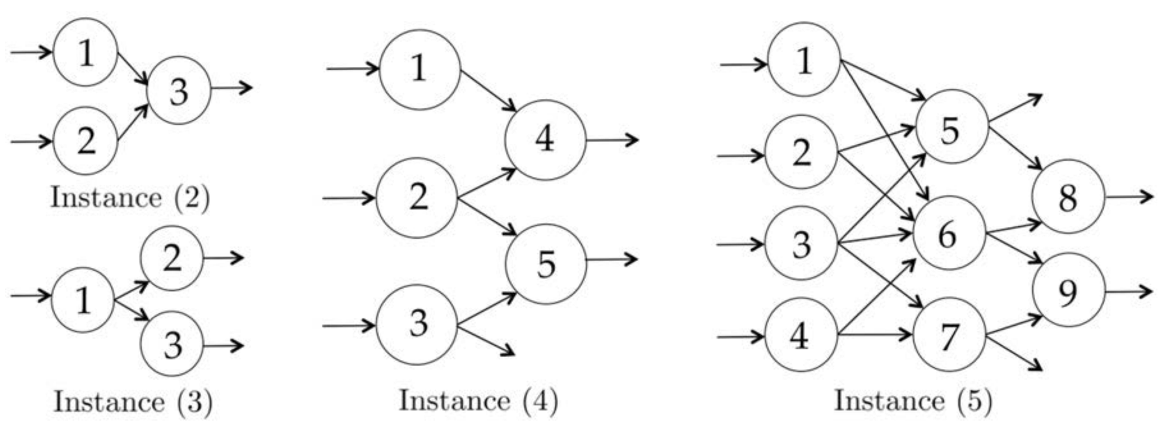

Figure 2 Network instances with three installations in Instances (2) and (3), five installations in Instance (4), and nine installations in Instance (5).

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 $\overline{C}(S)$, and (b) approximate the average cost using our framework by applying Eq. (23) and the discretization corresponding to $|\mathcal{I}| = 5$ (see Table 1), and report the approximated cost $\widetilde{C}(S)$.

Table 1 Associated costs of interest.

Framework[†]	Average Cost
Our Approach	$\widetilde{C}(S) = \mathbb{E}_{\boldsymbol{\Gamma}}\left[\widehat{C}(S, \boldsymbol{\Gamma})\right]$
Stochastic Approach	$\overline{C}(S) = \mathbb{E}_{\boldsymbol{\omega}}\left[C(S, \boldsymbol{\omega})\right]$

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

Figure 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 assess the performance of our approach and the effect of the demand behavior on our solutions. To do so, we compute the optimal base-stock policy $(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}})$ under our approach. We also evaluate the optimal policy $(\widehat{\mathbf{s}}, \widehat{\mathbf{S}})$ obtained via the traditional robust optimization approach (using Eq. (25)) for different values of Γ. We compare the solutions from our framework and the traditional robust optimization approach with the optimal policy $(\overline{\mathbf{s}}, \overline{\mathbf{S}})$ obtained for the stochastic system given some distributional assumptions on the demand at the sink node. To evaluate the performance of policies $(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}})$ and $(\widehat{\mathbf{s}}, \widehat{\mathbf{S}})$ against policy $(\overline{\mathbf{s}}, \overline{\mathbf{S}})$, we compute the following quantities. Note that the expected values are taken with respect to some particular demand distribution. We report the relative percent errors with respect to the stochastic optimal cost, i.e.,

$$\frac{\widetilde{C} - \overline{C}}{\overline{C}} \times 100 \quad \text{and} \quad \frac{\widehat{C} - \overline{C}}{\overline{C}} \times 100.$$

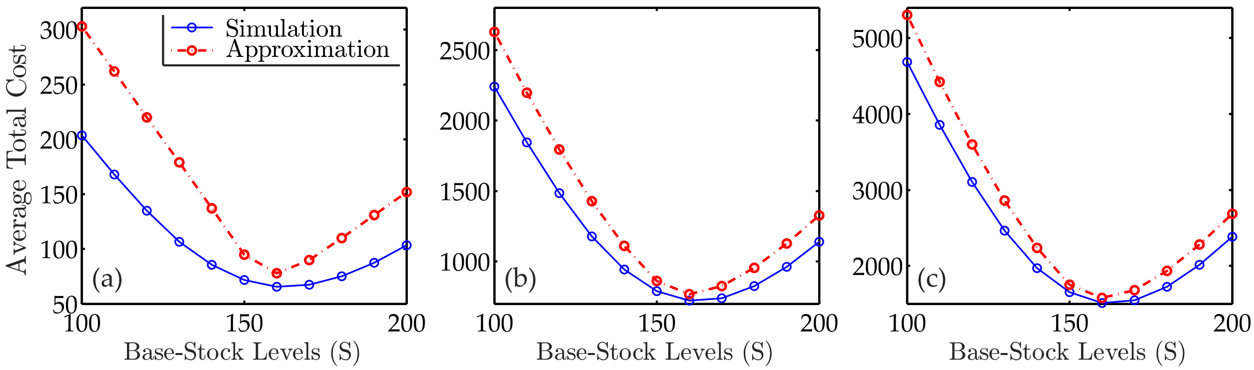

Figure 3 Simulated (solid line) versus approximated values (dotted line) for a single installation with an order-up-to policy, demand mean $\mu = 150$, standard deviation $\sigma = 30$, holding cost $h = \$2$ and penalty cost $p = \$4$, and zero fixed cost. Simulated values computed for normally distributed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

Table 2 Solutions and associated costs of interest.

Framework	Optimal Policy	Average Cost
Our Approach	$(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}})$	$\widetilde{C} = \mathbb{E}_{\boldsymbol{\omega}}\left[C(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}}, \boldsymbol{\omega})\right]$
Robust Approach[†]	$(\widehat{\mathbf{s}}, \widehat{\mathbf{S}})$	$\widehat{C} = \mathbb{E}_{\boldsymbol{\omega}}\left[C(\widehat{\mathbf{s}}, \widehat{\mathbf{S}}, \boldsymbol{\omega})\right]$
Stochastic Approach	$(\overline{\mathbf{s}}, \overline{\mathbf{S}})$	$\overline{C} = \mathbb{E}_{\boldsymbol{\omega}}\left[C(\overline{\mathbf{s}}, \overline{\mathbf{S}}, \boldsymbol{\omega})\right]$

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

To illustrate our results, we consider the example of Instance (2) with three echelons and a single sink node with time horizon $T = 8$, demand mean $\mu = 100$. Figure 4 compares the percent relative errors obtained using our framework and the robust approach ($\Gamma = 2$ and $\Gamma = 3$) versus stochastic optimization. We report the errors for various values of $\sigma \in [10, 100]$ with four different demand distributions at the sink node (normal, lognormal, gamma and uniform distributions). 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 $\Gamma = 1, 2, 3$ yield larger errors for all considered instances. Note that the effect of variability is more visible for lognormal and gamma distributed demand.

Impact of Network Size

We consider the network instances depicted in Figure 1 and use our framework to obtain the optimal inventory policy $(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}})$. We then assess the performance of our solution to the optimal inventory policy $(\overline{\mathbf{s}}, \overline{\mathbf{S}})$ obtained in the stochastic setting under some given distributional assumptions around the demand behavior. We compute the solution percent error

$$\frac{\widetilde{C} - \overline{C}}{\overline{C}} \times 100,$$

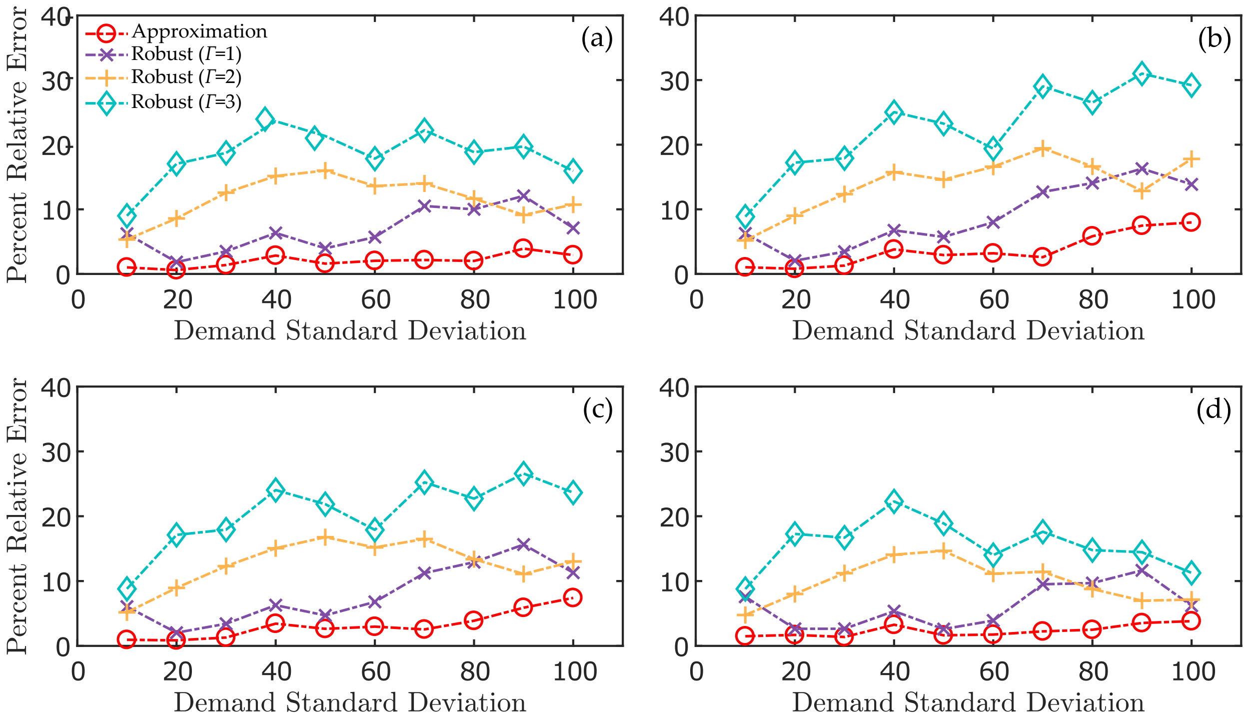

Figure 4 **Percent errors associated with implementing the solutions given by our approximation and the robust optimization approach ($\Gamma = 1, 2, 3$) relative to implementing the optimal stochastic solution. Errors are depicted for Instance (2) with demand mean $\mu = 100$, $T = 8$, while varying the demand standard deviation in the range of $[10, 100]$. Panel (a)-(d) compare the performance to the stochastic instance with the demand at the sink node following (a) normal distribution, (b) a lognormal distribution, (c) a gamma distribution, and (d) a uniform distribution, respectively.**

where $\overline{C}$ and $\widetilde{C}$ are defined in Table 3. Table 4 compares the computational performance of our approach for Instances (1)-(5) with the performance of the sample average approximation (SAA).

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.

Computational Performance

Similarly to the observations made by Bienstock and Özbay (2008), the iterative algorithm converges to good solutions within a few iterations.

Figure 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. Figure 6 shows that the fast convergence of the algorithm is carried through for networks of different sizes.

Table 3 Errors (%) relative to the stochastic solution.

Instance	Demand‡	Solution Percent Error† $T=6$	$T=9$	$T=12$
(1)	N	0.33	0.41	1.19
	L	4.67	4.85	4.85
(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.

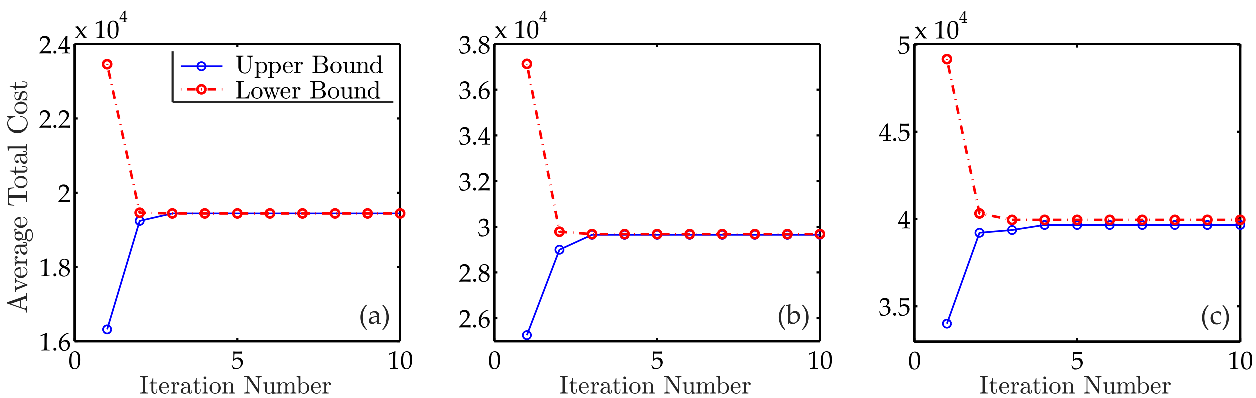

Figure 5 Evolution of the lower (solid line) and upper (dotted line) bounds through the iterative algorithm. Panels (a), (b) and (c) correspond to Instance (4) with an (s, S) policy and variable cost for $T=6$, $T=9$ and $T=12$, respectively.

Table 4 Runtimes (seconds) comparing our approach† to the sample average approximation (SAA)‡.

Instance	$T=6$ Iterations	$T=6$ Runtime	$T=6$ SAA	$T=9$ Iterations	$T=9$ Runtime	$T=9$ SAA	$T=12$ Iterations	$T=12$ Runtime	$T=12$ SAA
(1)	4	2.0	51.1	5	5.1	1,062.5	4	22.0	5,062.8
(2)	4	7.0	1,412.5	2	18.3	>7,500	4	489.7	>7,500
(3)	3	7.2	>7,500	3	75.5	>7,500	3	448.9	>7,500
(4)	4	27.2	>7,500	3	269.1	>7,500	3	1,112.7	>7,500
(5)	3	87.8	>7,500	3	1,185.7	>7,500	3	1,527.2	>7,500

† Convergence to within 2% gap between the lower and upper bound
† Convergence to within 2% gap between the lower and upper bound with 100 samples.

3.3. Comparison with the Distributionally Robust Framework

In this section, we compare our methodology to the distributionally robust optimization framework , where, similarly to our approach, only partial information about the demand's mean and

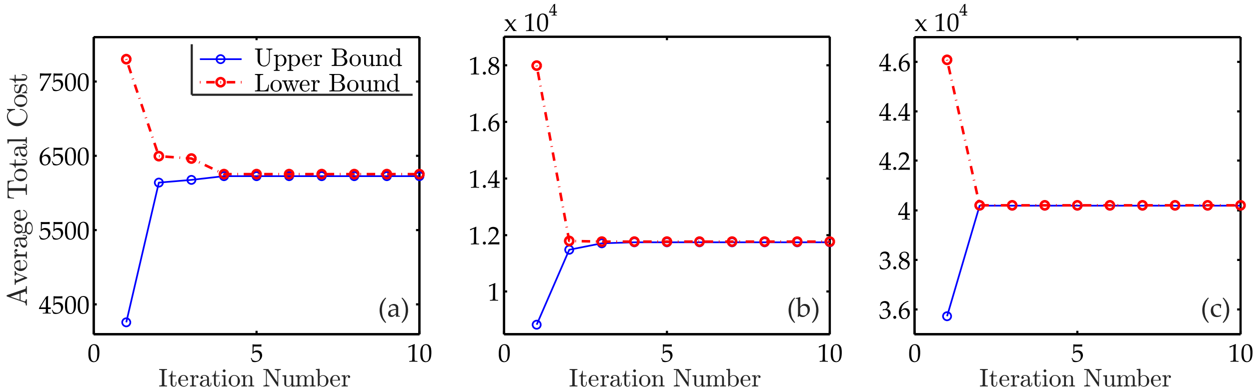

Figure 6 Evolution of the lower (solid line) and upper (dotted line) bounds through the iterative algorithm. Panels (a), (b) and (c) correspond to an inventory network with a horizon $T=8$, a (s,S) policy, and zero variable costs for Instance (2), Instance (4) and Instance (5), respectively.

variability is leveraged to solve the inventory problem at hand. Specifically, we suppose that we have knowledge of the following entities regarding the demand distribution

(a) Support, i.e., known range of demand values at the sink nodes,

(b) Mean μ_k, and therefore first moment, of the demand at sink node $k \in \mathcal{S}$, and

(c) Standard deviation σ_k, and therefore second moment μ_k, of the demand at sink node $k \in \mathcal{S}$.

Mathematically, the distributionally robust framework solves the following problem

$$\min_{\pi \in \Pi} \max_{H \in \mathcal{H}} C(\pi, \boldsymbol{\omega}), \tag{40}$$

where H denotes the distribution of the demand vector $\boldsymbol{\omega} = (\omega_k^0, \ldots, \omega_k^T)_{k \in \mathcal{S}}$ and $\mathcal{H}$ denotes the family of demand distributions such that, at each sink node $k \in \mathcal{S}$, and for all $0 \le t \le T$,

$$\begin{cases} \displaystyle\int_{\omega_k^t} \omega_k^t dH_k(\omega_k^t) = 1, \\ \displaystyle\int_{\omega_k^t} \omega_k^t dH(\omega_k^t) = \mu_k, \\ \displaystyle\int_{\omega_k^t} (\omega_k^t)^2 dH(\omega_k^t) = \nu_k, \\ \omega_k^t \in [\mu_k - \gamma \cdot \sigma_k, \mu_k + \gamma \cdot \sigma_k], \end{cases} \tag{41}$$

where γ is a user-input parameter that controls the support of the demand values. Note that, for simplicity, we have assumed that the demand is identically distributed at all time periods. Similarly to Section 2, the problem in Eq. (40) can be solved using a Bender's decomposition framework as detailed in Thorsen and Tao (2015). Note that, using the first and second moment information to model the demand uncertainty under the distributionally robust framework, the resulting adversarial problem under the Bender's decomposition approach is formulated as a mixed integer quadratic optimization (MIQO) problem, compared to the mixed integer linear optimization problems in Eq. (27) that we solve for the selected basestock and affine policies under both the traditional and our proposed robust optimization frameworks.

Computational Performance

We next assess the performance of the distributionally robust framework against our proposed approach. To do so, we compute the optimal base-stock policies $(\mathbf{s}^{\text{DRO}}, \mathbf{S}^{\text{DRO}})$ and $(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}})$ under the distributionally robust framework and our approach, respectively. We compare the solutions from our framework and the traditional robust optimization approach with the optimal policy $(\overline{\mathbf{s}}, \overline{\mathbf{S}})$ obtained for the stochastic system given some distributional assumptions on the demand at the sink node. To evaluate the performance of policies $(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}})$ and $(\widehat{\mathbf{s}}, \widehat{\mathbf{S}})$ against policy $(\overline{\mathbf{s}}, \overline{\mathbf{S}})$, we compute the following quantities.

Table 5 Solutions and associated costs of interest.

Framework	Optimal Policy	Average Cost
Our Approach	$(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}})$	$\widetilde{C} = \mathbb{E}_{\boldsymbol{\omega}}\left[C(\widetilde{\mathbf{s}}, \widetilde{\mathbf{S}}, \boldsymbol{\omega})\right]$
Distributionally Robust†	$(\mathbf{s}^{\text{DR}}, \mathbf{S}^{\text{DR}})$	$C^{\text{DR}} = \mathbb{E}_{\boldsymbol{\omega}}\left[C(\mathbf{s}^{\text{DR}}, \mathbf{S}^{\text{DR}}, \boldsymbol{\omega})\right]$
Stochastic Approach	$(\overline{\mathbf{s}}, \overline{\mathbf{S}})$	$\overline{C} = \mathbb{E}_{\boldsymbol{\omega}}\left[C(\overline{\mathbf{s}}, \overline{\mathbf{S}}, \boldsymbol{\omega})\right]$

† Computed as a function of a given value of γ.

We report the relative percent errors with respect to the stochastic optimal cost, i.e.,

$$\frac{C^{\text{DR}} - \overline{C}}{\overline{C}} \times 100 \quad \text{and} \quad \frac{\widehat{C} - \overline{C}}{\overline{C}} \times 100.$$

To illustrate our results, Table 8 compares the percent errors relative to the stochastic solutions for Instances (3), (4), and (5) for time periods $T = 6, 9$, and 12. Note that, for larger networks, our approach produces solutions whose expected costs are notably lower than the distributionally robust solutions, under similar runtime and gap restrictions.

4. Optimizing Affine Policies

In this section, we investigate *disturbance-affine policies* and compare their effect on the system's average cost. Disturbance-affine policies are expressed as affine parameterizations in the historically observed demand. Put differently, we express the amount ordered at the beginning of a given time period t as an affine function of the demand realizations observed up until time $t-1$. Such policies belong to the general class of *decision rules* which have originally been introduced in the context of stochastic programming by Charnes et al. (1958) and Garstka and Wets (1974).

Ben-Tal et al. (2004b) extended the robust optimization framework to dynamic settings and explored the use of disturbance-affine policies in allowing the decision maker to adjust their strategy given the information revealed over time. Within the robust optimization framework, affine policies

Table 6 Errors (%) relative to the stochastic solution for basestock policies.

Instance	Demand‡	Our Approach †			Distributionally Robust ($\gamma = 3$)†		
		$T = 6$	$T = 9$	$T = 12$	$T = 6$	$T = 9$	$T = 12$
(3)	G	2.64	3.23	2.42	2.84	11.21	2.44
	N,L	3.44	9.38	2.16	3.43	10.92	3.10
(4)	L	2.79	3.37	4.72	12.65	2.88	8.83
	N,G,U	2.41	2.94	4.32	9.71	1.65	5.97
(5)	U	2.07	1.77	1.43	8.41	6.84	7.43
	N,L,G,U	2.05	1.81	1.33	9.73	7.56	8.19

† Convergence within 2% gap between the lower and upper bounds for the Bender's Decomposition algorithms. MIO and MIQO gap of 2% with 120s time limit allowed for each MIO (600s total for the five adversarial problems under our approach), and 600s time limit allowed for the MIQO problem under the Distributionally Robust framework.

‡ N, L, G, and U stand for normal, lognormal, gamma and uniform demand distributions.

have gained much attention due to their tractability; depending on the class of the nominal problem, the optimal policy parameters can be solved via linear, quadratic, conic or semidefinite programs (see Löfberg (2003), Kerrigan and Maciejowski (2004), Ben-Tal et al. (2004a)). Empirically, Ben-Tal et al. (2005) and Kuhn et al. (2011) have reported that affine policies perform excellently and have shown many instances in which they were optimal. Bertsimas et al. (2010) proved the optimality of disturbance-affine control policies for one-dimensional, constrained, multistage robust optimization and showed that these results hold for the finite-horizon case with minimax objective. In particular, Bertsimas et al. (2010) have shown that, under the robust setting, affine policies are optimal for a single-product, single-echelon, multi-period supply chain with zero fixed costs.

Instead of taking a worst case approach, we propose to study the average performance of affine policies for supply chain networks. We employ the methodology we introduced in Sections 2 and 3 to compute optimal affine parameterizations and compare their performance with the solutions obtained via the traditional robust optimization approach. Furthermore, we assess the value of affine policies as opposed to base-stock policies for generalized inventory networks.

4.1. Problem Formulation

Under an affine policy, we represent the echelon order quantities at the beginning of time period t as a function of the historical demand observed by that echelon until time $t-1$. We define

$$u_n^t = \beta_{n,0}^t + \sum_{k \in \mathcal{S}_n} \sum_{j=1}^{t} \beta_{n,j}^t \cdot \omega_k^{t-j}, \tag{42}$$

where the vector $\boldsymbol{\beta}_n^t = \left\{\beta_{n,j}^t, j = 0, \ldots, t\right\}$ denote the affine parameters associated with echelon n at time t.

Note: We can simplify the model by expressing the ordering cost as an affine function of a subset of demand realizations. For instance, we can invoke the past τ time periods with $\boldsymbol{\beta}_n^t = \{\beta_{n,j}^t, j = 0, \ldots, \tau\}$ and obtain the following functional form

$$u_n^t = \beta_{n,0}^t + \sum_{k \in \mathcal{S}_n} \sum_{j=1}^{\tau} \beta_{n,j}^t \cdot \omega_k^{t-j}. \tag{43}$$

Finding the optimal affine parameters in our framework calls for solving a robust optimization problem of the form of Eq. (24). Specifically, we consider the following problem formulation

$$\left\{ \begin{array}{lll} \min\limits_{\boldsymbol{\beta}} & \sum\limits_{i \in \mathcal{I}} f_i \cdot y_i & \\ \text{s.t.} & y_i \geq C(\boldsymbol{\beta}, \boldsymbol{\omega}) & \forall \boldsymbol{\omega} \in \mathcal{U}(\mathit{\Gamma}_i) \text{ and } \mathit{\Gamma}_i : i \in \mathcal{I} \end{array} \right\}, \tag{44}$$

where the vector $\boldsymbol{\beta} = \{\boldsymbol{\beta}_n^t, \forall n, t\}$ and the total inventory cost is given by

$$C(\boldsymbol{\beta}, \boldsymbol{\omega}) = \sum_{t=1}^{T} \sum_{\ell \in \mathcal{L}} c_\ell \cdot o_\ell^t + \sum_{t=1}^{T} \sum_{n \in \mathcal{N}} \left[h_n \cdot (x_n^t)^+ + p_n \cdot (x_n^t)^- + K_n \cdot \mathbb{1}_{u_n^t > 0} \right], \tag{45}$$

with o_ℓ^t, x_n^t, and u_n^t are functions of $(\boldsymbol{\beta}, \boldsymbol{\omega})$, for all values of n and t. We solve the problem in Eq. (44) via decomposition as presented in Section 4.2.4 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 $\mathcal{U}(\mathit{\Gamma}_i)$ and retrieve the optimal solution $\widehat{\boldsymbol{\omega}}^i$ that drives the worst case value. For a given parameter $\mathit{\Gamma}_i$ and a vector $\boldsymbol{\beta}_n^t = \left\{\beta_{n,j}^t, j = 0, \ldots, \tau\right\}$, for all n and t, problem (AP) in Eq. (27) can be re-written as

$$\begin{array}{lll} \max\limits_{\boldsymbol{\omega} \in \mathcal{U}(\mathit{\Gamma}_i)} & \sum\limits_{t=0}^{T} \sum\limits_{\ell \in \mathcal{L}} c_\ell \cdot o_\ell^t + \sum\limits_{t=0}^{T} \sum\limits_{n \in \mathcal{N}} \left[h_n \cdot (x_n^t)^+ + p_n \cdot (x_n^t)^- + K_n \cdot \mathbb{1}_{u_n^t > 0} \right] & \\ \text{s.t.} & x_n^{t+1} = x_n^t + u_n^t - \sum\limits_{k \in \mathcal{S}_n} \omega_k^t, & \forall t = 0, \ldots, T, \\ & u_n^t = \sum\limits_{\ell \in \mathcal{L}_n} o_\ell^t, & \forall t = 0, \ldots, T, \\ & u_n^t = \beta_{n,0}^t + \sum\limits_{k \in \mathcal{S}_n} \sum\limits_{j=1}^{\tau} \beta_{n,j}^t \cdot \omega_k^{t-j}, & \forall t = 0, \ldots, T. \end{array}$$

Problem (AP) is a non-concave maximization problem and the optimal solution $\widehat{\boldsymbol{\omega}}^i$ may not occur at a corner point of the uncertainty set $\mathcal{U}(\mathit{\Gamma}_i)$. Problem (AP) can be cast as a mixed integer optimization (MIO) problem and solved relatively efficiently using available optimization solvers. Similarly to the case of base-stock policies, we introduce two sets of auxiliary binary variables to formulate problem (AP) as a mixed integer optimization problem

$$y_n^t = \left\{ \begin{array}{ll} 1, & \text{if } u_n^t > 0 \\ 0, & \text{otherwise} \end{array} \right\} \text{ and } z_n^t = \left\{ \begin{array}{ll} 1, & \text{if } x_n^t > 0 \\ 0, & \text{otherwise} \end{array} \right\}.$$

Note that, given the affine structure of the ordering policy, the problem above is easier to solve compared to the adversarial problem that we obtain for base-stock policies.

Decision Maker's Problem: At each iteration of the algorithm, problem (DM) consists of finding the best affine policy, given a finite collection of demand realizations stored thus far. Specifically, for each index $i \in \mathcal{I}$, we populate the set $\widehat{\mathcal{U}}_i$ with the optimal solutions $\widehat{\boldsymbol{\omega}}^i$ that we obtain from solving problem (AP) at each iteration of the algorithm. Mathematically, we formulate problem (DM) in Eq. (26) as

$$\left\{ \begin{array}{lll} \min\limits_{\boldsymbol{\beta}} & \sum\limits_{i \in \mathcal{I}} f_i \cdot q_i & \\ \text{s.t.} & q_i \geq C\left(\boldsymbol{\beta}, \widehat{\boldsymbol{\omega}}^i\right), & \forall \widehat{\boldsymbol{\omega}}^i \in \widehat{\mathcal{U}}_i, i \in \mathcal{I} \end{array} \right\}, \tag{46}$$

where the total cost is given by Eq. (45). For the generalized case where the fixed costs are non-zero, problem (DM) can be cast as an MIO whose size grows with the number of iterations. Our computations suggest that more iterations are needed to achieve a convergence within 5% for affine policies compared to base-stock policies. This suggests that affine policies are harder to solve for. However, they achieve lower costs, as shown in Section 4.2.

Note: For the case where the fixed costs are zero, we can implement the methodology provided by Ben-Tal et al. (2005) to formulate an approximation of (44) that can be cast as a linear optimization problem and achieve better tractability. For the case where the fixed costs are non-zero, we employ the generic decomposition algorithm presented in Section 4.2. However, one may investigate the performance of novel decomposition techniques such as the algorithms developed by Postek and Hertog (2016) and Bertsimas and Dunning (2016). We next evaluate the performance of affine policies and compare our solutions to those obtained for base-stock policies.

4.2. Computational Results

We investigate the performance of affine policies and examine the effect of the system's parameters on our solutions. For our computations, we consider the five network topologies presented in Figure 2. We assume throughout that the fixed costs are non-zero.

Impact of Demand Variability

We assess the performance of our approach and the effect of the demand behavior on our solutions. To do so, we apply our approach and compute the optimal affine policy $\widetilde{\boldsymbol{\beta}} = \{\widetilde{\boldsymbol{\beta}}_n^t, \forall n, t\}$, where $\widetilde{\boldsymbol{\beta}}_n^t = \{\widetilde{\beta}_{n,j}^t, j = 0, \ldots, \tau\}$, for all n and t. We also evaluate the optimal policy $\widehat{\boldsymbol{\beta}}$ obtained via the traditional robust optimization approach (using Eq. (25)). We compare the cost implied by the solutions from our framework and the traditional robust optimization approach to the optimal cost that we obtain using base-stock policies. In particular, we compute the following quantities.

Table 7 Solutions and associated costs of interest.

Framework	Optimal Policy	Average Cost
Our Affine Approach	$\widetilde{\boldsymbol{\beta}}$	$\widetilde{C}=\mathbb{E}_{\boldsymbol{\omega}}\left[C(\widetilde{\boldsymbol{\beta}},\boldsymbol{\omega})\right]$
Robust Affine Approach†	$\widehat{\boldsymbol{\beta}}$	$\widehat{C}=\mathbb{E}_{\boldsymbol{\omega}}\left[C(\widehat{\boldsymbol{\beta}},\boldsymbol{\omega})\right]$
Base-Stock Approach	$(\overline{\mathbf{s}},\overline{\mathbf{S}})$	$\overline{C}=\mathbb{E}_{\boldsymbol{\omega}}\left[C(\overline{\mathbf{s}},\overline{\mathbf{S}},\boldsymbol{\omega})\right]$

† Computed as a function of a given value of $\varGamma$.

Note that the expected values are taking with respect to some particular demand distribution. We report the relative percent errors with respect to the base-stock optimal cost, i.e.,

$$\frac{\widetilde{C}-\overline{C}}{\overline{C}}\times 100 \quad \text{and} \quad \frac{\widehat{C}-\overline{C}}{\overline{C}}\times 100.$$

Note that negative percent errors indicate that the optimal affine policy yields a lower cost compared to the optimal cost obtained under a base-stock policy.

To illustrate our results, we consider the example of Instance (2) with three echelons and a single sink node with time horizon $T=8$, demand mean $\mu=100$. Furthermore, we assume a fully affinely adaptive policy where $\tau=t$ (i.e., we invoke all past historical demand realizations for the affine parameterization).

Figure 7 compares the percent relative errors for the affine policies obtained using our framework and the robust approach ($\varGamma=2$ and $\varGamma=3$) versus the optimal base-stock policy obtained via stochastic optimization. We report the errors for various values of $\sigma\in[10,100]$ with four different demand distributions at the sink node (normal, lognormal, gamma and uniform distributions).

The optimal affine policy we obtain in our framework generates an average cost that is consistently below the optimal cost obtained under a base-stock policy (the associated percent errors are negative throughout). The benefits of implementing affine policies compared with base-stock policies are highlighted especially for the case of lower demand variability.

Furthermore, our approach yields solutions with lower average costs compared to the traditional robust optimization framework. While the robust approach with $\varGamma=2$ yields good solutions for lower demand variability, this does not carry through for higher demand variability.

Impact of Network Size

We consider the network instances depicted in Figure 2 and use our framework and the traditional robust approach (with $\varGamma=2$) to obtain the optimal affine policies $\widetilde{\boldsymbol{\beta}}$ and $\widehat{\boldsymbol{\beta}}$. We then assess the performance of our solution to the optimal inventory policy $(\overline{\mathbf{s}},\overline{\mathbf{S}})$ obtained in the stochastic setting under some given distributional assumptions around the demand behavior. We compute

$$\frac{\widetilde{C}-\overline{C}}{\overline{C}}\times 100 \quad \text{and} \quad \frac{\widetilde{C}-\overline{C}}{\overline{C}}\times 100,$$

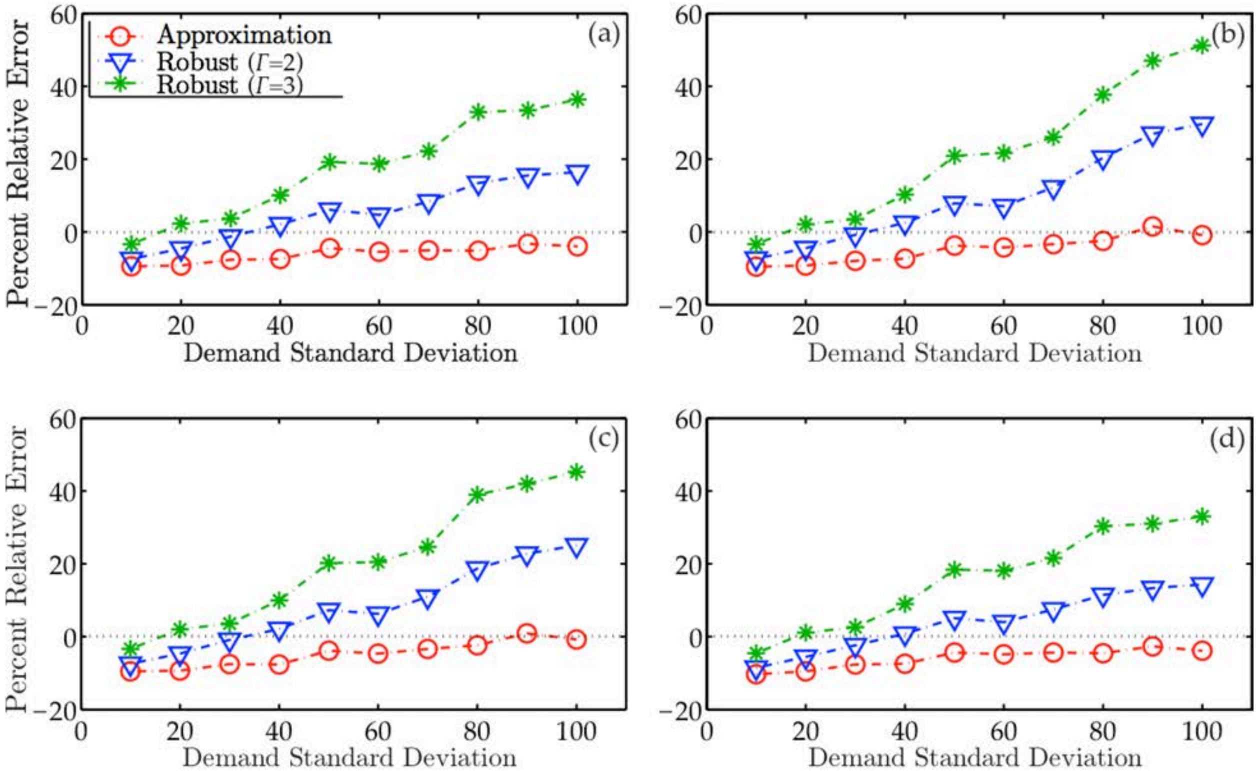

Figure 7 Percent errors of the average cost values implementing the solutions given by our approximation and the robust optimization approach ($\Gamma = 2$ and $\Gamma = 3$) relative to the optimal average cost implementing the optimal stochastic solution. Errors are depicted for Instance (2) with demand mean $\mu = 100$, $T = 8$, and zero variable costs, while varying the demand standard deviation in the range of $[10, 100]$. Panel (a)-(d) compare the performance to the stochastic instance with the demand at the sink node following (a) normal distribution, (b) a lognormal distribution, (c) a gamma distribution, and (d) a uniform distribution, respectively.

where $\overline{C}$ and $\widetilde{C}$ are defined in Table 6. We report herein our results for simplified affine policies with $\tau = 2$, i.e., we assume the ordering amount at time t is an affine function of the demand realizations at times $t-1$ and $t-2$. Table 7 compares the performance of our approach and the traditional robust setting with respect to the optimal base-stock policy for Instances (2)-(5) for various demand distributions and time horizons. Note that we set the overall time limit to 7,200 seconds (2 hours) for the entire algorithm.

Affine policies obtained under our approach outperform the base-stock policies under the simplified parametrization with $\tau = 2$. Furthermore, our framework generates affine policies that allow to achieve lower costs compared to the traditional robust approach. Note that the values in Table 7 show that the difference between affine and base-stock policies tend to slim down for larger time horizons. The larger the time horizon (and the larger the size of the network), the larger are the

Table 8 Percent errors relative to the optimal base-stock solution[†].

		$\Gamma = 2$		Random Γ	
Instance	Demand[‡]	$T = 6$	$T = 9$	$T = 6$	$T = 9$
(2)	G	-8.39	-1.21	-14.7	-9.54
	U	-9.49	-2.56	-15.0	-9.76
(3)	G	-9.08	0.66	-14.1	-8.77
	N,L	-9.30	0.48	-14.2	-8.78
(4)	L	-5.26	1.22	-11.4	-7.05
	N,G,U	-6.50	0.02	-11.7	-7.34
(5)	U	-3.38	-2.53	-11.6	-5.64
	N,L,G,U	-4.30	-3.56	-12.8	-7.09

optimization problems we would need to solve. The incumbent solution to the MIO problems is more likely to be far from optimal for the larger problems by the time we reach a time limit of 300 seconds.

Computational Performance

Under the assumption that fixed costs are non-zero, the iterative algorithm takes longer to converge for problems optimizing affine policies compared to those optimizing base-stock policies. Figure 8 shows the rate of convergence for Instance (2) and $\tau = 2$ with time horizons ranging from $T = 6$ to $T = 12$. Figure 9 shows that the convergence of the algorithm is highly dependent on the size of the network. Consequently, for affine policies, the network size and length of the time horizon seem to have a direct effect on the rate of convergence. Runtimes in Table 8 reflect the tradeoff between the cost savings of implementing affine policies versus the associated computational challenge.

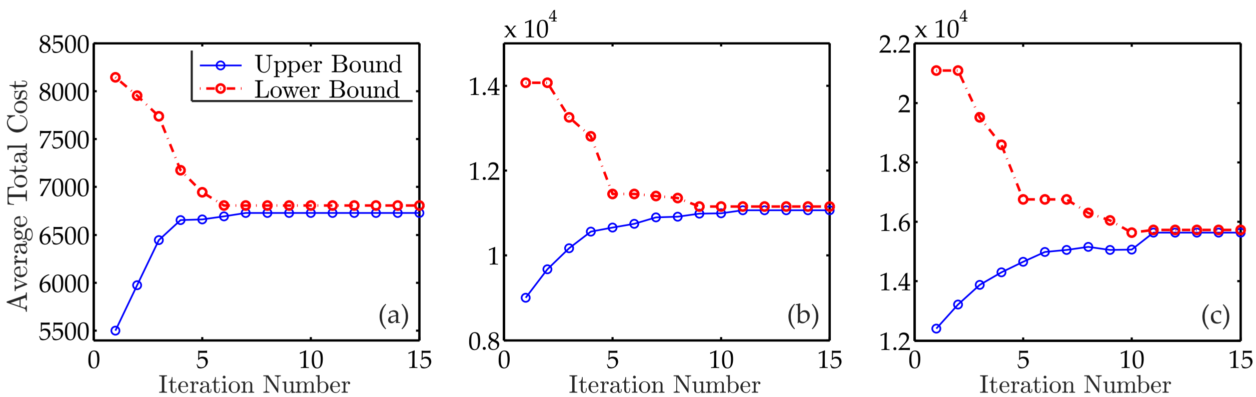

Figure 8 Evolution of the lower (solid line) and upper (dotted line) bounds through the iterative algorithm. Panels (a), (b) and (c) correspond to Instance (2) with three installations and a single sink nodes with an affine policy ($\tau = 2$) for $T = 6$, $T = 9$ and $T = 12$, respectively.

Note: For larger networks, the lower bound in Figure 8 may not increase monotonically. This is due to forcing a time limit of 300s to solve the decision maker's problem. As a result, the reported cost is associated with the incumbent solution retrieved at that time, and could be far from optimal.

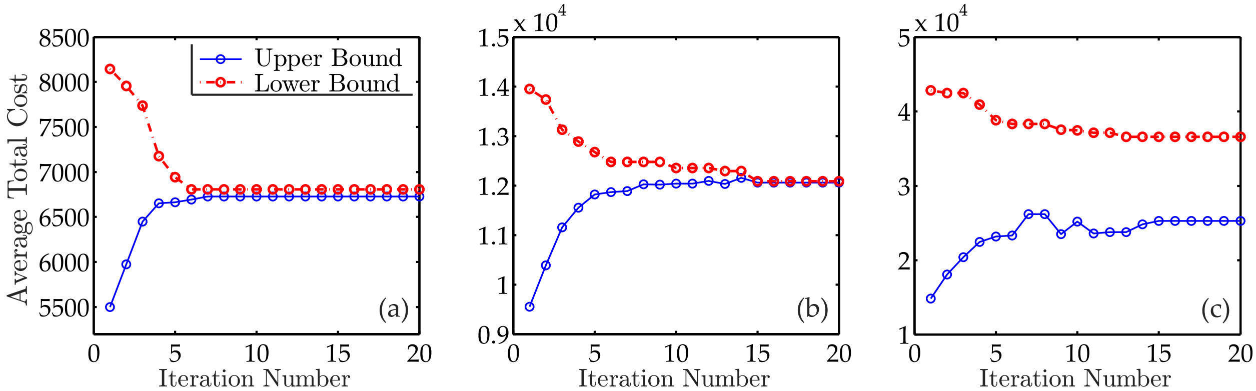

Figure 9 **Evolution of the lower (solid line) and upper (dotted line) bounds through the iterative algorithm. Panels (a), (b) and (c) correspond to an inventory network with a horizon $T = 6$, an affine policy with $\tau = 2$ for Instances (2), (4) and (5), respectively.**

Table 9 **Number of iterations and runtime (in seconds)†.**

	$T = 6$		$T = 9$	
Instance	Iterations	Runtime	Iterations	Runtime
(2)	5	20.7	7	796.7
(3)	6	328.9	13	2,589.1
(4)	7	547.0	13	5,574.5
(5)	>20	>7,200	>7	>7,200

† Convergence within 5% between the lower and upper bound

5. 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 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 and affine policies for complex supply chain networks. Our computations suggest that our approach (a) obtained base-stock levels whose expected performance matches that of optimal base-stock levels obtained via stochastic optimization, (b) provided optimal affine policies which often times yield better results compared with optimal base-stock policies, and (c) generated policies that consistently outperform the solutions obtained via the traditional robust optimization approach.

In summary, our proposed framework (a) avoids the challenges of fitting probability distributions to the uncertain variables, (b) eliminates the need to generate 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. While randomness is viewed probabilistically in stochastic optimization and deterministically in robust optimization, our approach bridges the strength of the limit laws of probability and the tractability of the deterministic robust setting in view of understanding and optimizing systems under uncertainty.

References

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, 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, S. Boyd, and A. Nemirovski. Control of uncertainty-affected discrete time linear systems via convex programming. *Mathematical Programming*, 99(2):351–376, 2004a.

A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004b.

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 I. Dunning. Multistage robust mixed integer optimization with adaptive partitions. *Operations Research*, 64(4):980–998, 2016.

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. A robust optimization approach to inventory theory. *Operations Research*, 54: 150–168, 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. *Mathematical Programming, to appear*, 2017.

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. Tayur. 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.* Springer, 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. Tayur. *Quantitative Models for Supply Chain Management*, 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.

E. Kerrigan and J. Maciejowski. Properties of a new parametrization for the control of constrained systems with disturbances. *Proceedings of the 2004 American Control Conference*, 5:4669–4674, 2004.

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.

J. Löfberg. Approximations of closed-loop minimax mpc. *Proceedings of the 42nd IEEE Conference on Decision Control*, 2:1438–1442, 2003.

T. Morton. The non-stationary infinite horizon inventory problem. *Management Science*, 24:1474–1482, 1978.

A. Muharremoglu and J. Tsitsiklis. A singe-unit decomposition approach to multi-echelon inventory systems. *Operations Research*, 56:1089–1103, 2008.

K. Postek and D. D. Hertog. Multi-stage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set (in revision). *INFORMS Journal of Computing*, 28(3):553–574, 2016.

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.

A. Thorsen and Y. Tao. Robust inventory control under demand and lead time uncertainty. *Annals of Operations Research*, pages 1–30, Dec 2015.