

Sustainable Inventory with Robust Periodic-Affine Policies and Application to Medical Supply Chains

Chaithanya Bandi,^a Eojin Han,^b Omid Nohadani^b

^a Kellogg School of Management, Northwestern University, Evanston, Illinois 60208; ^b Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208

Contact: c-bandi@kellogg.northwestern.edu (CB); eojinhan2020@u.northwestern.edu (EH); nohadani@northwestern.edu,  <http://orcid.org/0000-0001-6332-3403> (ON)

Received: June 23, 2017

Revised: December 27, 2017

Accepted: April 17, 2018

Published Online in Articles in Advance:
May 22, 2019

<https://doi.org/10.1287/mnsc.2018.3152>

Copyright: © 2019 INFORMS

Abstract. We introduce a new class of adaptive policies called *periodic-affine policies*, which allows a decision maker to optimally manage and control large-scale newsvendor networks in the presence of uncertain demand without distributional assumptions. These policies are data-driven and model many features of the demand such as correlation and remain robust to parameter misspecification. We present a model that can be generalized to multiproduct settings and extended to multiperiod problems. This is accomplished by modeling the uncertain demand via sets. In this way, it offers a natural framework to study competing policies such as base-stock, affine, and approximative approaches with respect to their profit, sensitivity to parameters and assumptions, and computational scalability. We show that the periodic-affine policies are sustainable—that is, time consistent—because they warrant optimality both within subperiods and over the entire planning horizon. This approach is tractable and free of distributional assumptions, and, hence, suited for real-world applications. We provide efficient algorithms to obtain the optimal periodic-affine policies and demonstrate their advantages on the sales data from one of India's largest pharmacy retailers.

History: Accepted by Yinyu Ye, optimization.

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/mnsc.2018.3152>.

Keywords: newsvendor network • robust optimization • demand uncertainty • correlation • affine policies • healthcare: pharmacy retailer

1. Introduction

Despite the physicians' diagnostic matching of patients to drugs, the heterogeneity in patients' illness, drug's efficacy, potential side effects, and varying length of treatment lead to sizable uncertainty in drug's demand (Crawford and Shum 2005). Retailers are mandated to service-level guarantees, and overstocking drugs is neither economical nor practical because they are perishable. Such healthcare problems affect a wide section of the population and have large societal implications. In this context, newsvendor models offer a natural framework and are used for decision making.

Practical solutions to such problems are critical to a broad range of industries. In particular, pharmaceutical companies with a large turnover are interested in optimal inventory management. GlaxoSmithKline spends over \$4.5 billion each year on manufacturing and supplying products. Johnson & Johnson spends approximately \$30 billion annually in leveraging its purchasing power to set sustainability expectations beyond its operations. Similarly, companies like Teva Pharmaceuticals, Pfizer, and Merck spend millions of dollars to ensure the safety and supply of their products, even though they have manufacturing units in multiple locations. Therefore, any variation in inventories can lead

to multiple disturbances in the system. A pharmacy's inventory represents its single largest investment. As discussed in Webman (2012), in a common pharmacy, cost of goods sold accounts for approximately 68% of total expenditures. For every 1% change in costs of goods, profits may increase or decrease by more than 20% (see Webman 2016). Thus, the sheer magnitude of dollars involved makes seemingly minor inefficiencies in purchasing and inventory control matters of great importance to both cash flow and profitability.

The challenges of such networks are multifold. Real-world settings are typically high-dimensional with multiple products and multiple stages of decision-making. These settings also suffer from substantial uncertainties in demand. Modeling such demand uncertainty is challenging because demand is often not stationary, or its uncertainty can depend on previous decisions (Nohadani and Sharma 2018).

In this work, we consider a newsvendor network with uncertain and correlated demand. Using the paradigm of robust optimization, we model such demand to reside in uncertainty sets and provide tractable formulations and associated algorithms for sustainable policies. To gain insight from a real-world setting, we apply the results to a major online pharmacy retailer in India, where a

prohibitively large penalty occurs when customers' demand is not satisfied. This company carries over 163 different brands, and the sales grow at about 23% per year. Their distribution network spans the entire country through fixed retail locations and online platforms. The decision makers of this company observe a sizable uncertainty in demand over the course of the year (in addition to seasonality) and significant correlations amongst various product categories. In close collaboration with this company's managers, we seek to design optimal implementable policies to control their inventory levels in their network.

Our contributions are as follows:

- *Modeling*: We provide a distribution-free description of uncertainty in demand using two types of sets. Independent demands are modeled via budget constraints. We also incorporate correlated demands using a factor model approach. The inventory-control problem is then cast as a multistage robust optimization problem. As a result, a novel solution concept of *periodic-affine policy* is provided for newsvendor networks with time-dependent and potentially correlated demand uncertainty.
- *Algorithms*: We provide a tractable algorithm that provides periodic-affine policies. These policies decompose the overall problem into a more tractable formulation than affine policies.
- *Application*: We analyze the sales data of a pharmacy retailer in India for the fourth quarter of 2016. This entails 1.5 million transactions for 228 different products. We construct the demand uncertainty set for the 20 most popular products, comprising 80% of all transactions. Our numerical experiments show that even for the single-station case, the computational burden for the optimal periodic-affine policies is significantly reduced over affine policies (by $100\times$ for a 15-period problem), making the proposed approach practical for real-world and large-sized problems. Moreover, the periodic-affine policy improved the cost-effectiveness of the operation by 19% over a base-stock policy for realistic penalty costs.

1.1. Literature Review

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 drops below s , an order is placed "up to S ." The (s, S) ordering policy has been proven optimal for simple stochastic inventory systems. Scarf (1960) proved that base-stock policies are optimal for a single installation model. Clark and Scarf (1960) extended the result to 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 nonstationary 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. For more work, refer to Langenhoff and Zijm (1990), Sethi and Cheng (1997), Muharremoglu and Tsitsiklis (2008), and Huh and Janakiraman (2008).

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 Kapuscinski and Tayur (1999) have developed various gradient-based algorithms to study inventory systems. These methods are practical whenever the input variables are continuous and their success depends on the quality of the gradient estimator.

On the other hand, Scarf (1958), Kasugai and Kasegai (1961), Gallego and Moon (1993), and Graves and Willems (2000) developed distribution-free approaches to inventory theory. Bertsimas and Thiele (2006) took a robust optimization approach to inventory theory and showed 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 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. Ben-Tal et al. (2004) extended the robust optimization framework to dynamic settings and explored the use of disturbance-affine policies by allowing the decision maker to adjust their strategy, leveraging the information revealed over time. Bertsimas and Thiele (2006) and Bienstock and Özbay (2008) studied the performance of base-stock policies, and Ben-Tal et al. (2005), Kuhn et al. (2011), and Bertsimas et al. (2010) investigated policies that are affine in prior demands under a robust optimization lens. Within the robust optimization framework, affine policies have gained much attention due to their tractability; depending on the class of the nominal problem, the optimal policy can be solved via linear, quadratic, conic, or semidefinite programs (see Löfberg 2003 and Kerrigan and Maciejowski 2004). Empirically, Ben-Tal et al. (2005) and Kuhn et al. (2011) have reported that affine policies have excellent performance and in many instances optimal.

Another approach is distributionally robust optimization, which assumes that the uncertainties follow a distribution within a prespecified set of distributions.

Such sets can be based on moment constraints (Delage and Ye 2010), phi-divergences (Ben-Tal et al. 2013), or Wasserstein metric (Mohajerin Esfahani and Kuhn 2018) to allow tractable reformulations. This approach typically yields less conservative solutions than deterministic robust optimization solutions. For multistage problems, Van Parys et al. (2016) proposed a tractable framework for distributionally robust linear feedback policy for discrete time linear control systems with quadratic objective functions.

In the context of pharmaceutical systems, Guerrero et al. (2013) provided a near-optimal base-stock policy for two-echelon distribution networks with multiple products, where every sink node is replenished by a single supplier. They provided a Markov chain formulation and a heuristic algorithm for Poisson distributed and independent demands. For a combined setting of a pharmaceutical company and a hospital, Uthayakumar and Priyan (2013) developed a two-echelon supply chain model to determine the optimal lot size, lead time, and total number of deliveries between the pharmaceutical company and a hospital. Using Lagrange multipliers, they provided decision tools for optimal costs while ensuring required service levels. In a two-level pharmaceutical supply chain, Baboli et al. (2011) studied a specific product with a constant demand rate and numerically showed that the overall cost is improved when pharmacies and hospitals are centralized.

Notation. Lowercase italic is used to denote scalars; lowercase bold is used to denote vectors, and uppercase bold is used to denote matrices. Sets are in calligraphic. Section-specific notation is introduced where needed. All proofs are relegated to the online appendix.

2. Model

We consider a newsvendor 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 adapted to systems with nonzero 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 occur and are filled or backlogged depending on the available inventory, and (3) the stock availability is updated for the next period.

- \mathcal{N} : Set of all installations where ordering decisions are made (source nodes) with $|\mathcal{N}| = m$
- \mathcal{S} : Set of all installations with external demand (sink nodes) with $|\mathcal{S}| = n$
- \mathcal{L} : Set of all links (edges) within the inventory network with $|\mathcal{L}| = p$
- \mathcal{N}_k : Set of source nodes supplying stock to a sink node $k \in \mathcal{S}$

- \mathcal{S}_v : Set of sink installations that are fed from a source node $v \in \mathcal{N}$
- s_t^v : Amount of order at the beginning of period t at a source $v \in \mathcal{N}$
- d_t^k : Demand observed at a sink $k \in \mathcal{S}$ throughout time period
- x_t^ℓ : Stock delivered along a link $\ell \in \mathcal{L}$ at time t
- $u_t^{s,v}$: Stock available after the period t at a source node $v \in \mathcal{N}$
- $u_t^{d,k}$: Backorders after the period t at a sink node $k \in \mathcal{S}$.


To track the system's operation, we capture information about the stock available and the stock ordered at source installations at the beginning of each time period as well as the demand at each sink installation throughout each time period (Figure 1). Specifically, assuming zero initial input and demands, we can express the dynamics of inventory levels and backlogged demands for $t = 1, \dots, T$ as

$$\begin{aligned} u_t^{s,v} &= u_{t-1}^{s,v} + s_t^v - \sum_{\ell=(v,k), k \in \mathcal{S}_v} x_t^\ell = \sum_{\tau=1}^t s_\tau^v - \sum_{\ell=(v,k), k \in \mathcal{S}_v} \sum_{\tau=1}^t x_\tau^\ell \\ &\quad \forall v \in \mathcal{N}, \\ u_t^{d,k} &= u_{t-1}^{d,k} + d_t^k - \sum_{\ell=(v,k), v \in \mathcal{N}_k} x_t^\ell = \sum_{\tau=1}^t d_\tau^k - \sum_{\ell=(v,k), v \in \mathcal{N}_k} \sum_{\tau=1}^t x_\tau^\ell \\ &\quad \forall k \in \mathcal{S}. \end{aligned} \quad (1)$$

Note that the ordering quantities $s_t^v = s_t^v(\pi, \mathbf{d})$, and therefore the amount of available stock $u_t^{s,v} = u_t^{s,v}(\pi, \mathbf{d})$ and backorders $u_t^{d,k} = u_t^{d,k}(\pi, \mathbf{d})$, are functions of the ordering policy π and the demand \mathbf{d} .

The high-dimensional nature of modeling demand uncertainty probabilistically and the complex dependence on random variables underscore the difficulty of analyzing and optimizing the expected total cost. Instead, we propose a framework that builds upon the robust optimization paradigm.

Figure 1. (Color online) Example of a Nine-Installation Network with $n = 4$ Sink Nodes and $m = 5$ Source Nodes



2.1. Robust Newsvendor Network Formulation

To describe our framework, we first introduce a robust approach to single-period models. Our models are based on the assumption that we have the following cost and revenue structure:

- c_S^v : Purchasing cost per unit at the source node $v \in \mathcal{N}$;
- c_H^v : Holding cost per unit for the leftover stock at the source node $v \in \mathcal{N}$;
- c_P^k : Penalty cost per unit for the unsatisfied demand at the sink node $k \in \mathcal{S}$;
- r^ℓ : Revenue by satisfying a unit demand occurred at the sink node k via $\ell = (v, k) \in \mathcal{L}$.

The goal of the decision maker is to order a proper amount of products $\{s_v : v \in \mathcal{N}\}$ and to process network activities $\{x_\ell : \ell \in \mathcal{L}\}$ to satisfy the customer demand at the sink nodes, so that the firm maximizes an overall profit. If we denote \mathcal{U} as a demand uncertainty set, then a single-period problem is formulated as a two-stage robust optimization problem

$$\begin{aligned} \max_{s_v \geq 0} & \left[- \sum_{v \in \mathcal{N}} c_S^v s_v + \min_{\mathbf{d} \in \mathcal{U}} \max_{x_\ell \geq 0} \left[\sum_{\ell \in \mathcal{L}} r_\ell x_\ell - \sum_{k \in \mathcal{S}} c_P^k \left(d_k - \sum_{\ell=(v,k), v \in \mathcal{N}_k} x_\ell \right) \right. \right. \\ & \left. \left. - \sum_{v \in \mathcal{N}} c_H^v \left(s_v - \sum_{\ell=(v,k), k \in \mathcal{S}_v} x_\ell \right) \right] \right] \\ \text{s.t.} & \sum_{\ell=(v,k), k \in \mathcal{S}_v} x_\ell \leq s_v \quad \forall v \in \mathcal{N}, \quad \sum_{\ell=(v,k), v \in \mathcal{N}_k} x_\ell \leq d_k \quad \forall k \in \mathcal{S}, \end{aligned} \quad (2)$$

where the constraints are network constraints and affect the inner maximization problem. Note that the order quantities $\{s_v : v \in \mathcal{N}\}$ are “here-and-now” decisions; it must be placed before demands are realized, whereas the network activities $\{x_\ell : \ell \in \mathcal{L}\}$ are “wait-and-see” solutions and assigned after demands are observed.

Notation. To simplify (2), we define $\mathbf{c}_S \in \mathbb{R}_+^m$, $\mathbf{c}_H \in \mathbb{R}_+^m$, $\mathbf{c}_P \in \mathbb{R}_+^n$, and $\mathbf{r} \in \mathbb{R}_+^p$ as cost and revenue vectors, and define $\mathbf{R}_S \in \mathbb{R}_+^{m \times p}$ and $\mathbf{R}_D \in \mathbb{R}_+^{n \times p}$ as matrices that describe the two constraints, respectively. Decision variables and uncertain demands are $\mathbf{s} \in \mathbb{R}_+^m$, $\mathbf{x} \in \mathbb{R}_+^p$, and $\mathbf{d} \in \mathbb{R}_+^n$. We obtain

$$\begin{aligned} \max_{\mathbf{s} \geq 0} & \left[- (\mathbf{c}_S + \mathbf{c}_H)^\top \mathbf{s} + \min_{\mathbf{d} \in \mathcal{U}} \max_{\mathbf{x} \geq 0} \left[\mathbf{v}^\top \mathbf{x} - \mathbf{c}_P^\top \mathbf{d} \right] \right] \\ \text{s.t.} & \mathbf{R}_S \mathbf{x} \leq \mathbf{s}, \quad \mathbf{R}_D \mathbf{x} \leq \mathbf{d}, \end{aligned}$$

with $\mathbf{v} = \mathbf{r} + \mathbf{R}_S^\top \mathbf{c}_H + \mathbf{R}_D^\top \mathbf{c}_P$.

2.2. Modeling 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). For each sink installation $k \in \mathcal{S}$, we denote the demand mean by μ_k and the demand standard deviation by σ_k . Our framework also captures correlation among the demand, where we denote $\Sigma \in \mathbb{R}^{n \times n}$ as the nominal

covariance matrix. Note that all these values can be inferred from historical data. Instead of describing the demand as a random variable, we describe the demand and its correlation by using budget uncertainty sets (Bertsimas and Sim 2004) and a factor-based approach (Bandi and Bertsimas 2012). Such sets do not require any distributional assumption other than first two moments, and, consequently, they are robust to the distribution choice.

We capture the correlations via the covariance matrix Σ with rank $l \leq n$. This means, there exist \mathbf{A} and $\lambda_1, \dots, \lambda_l > 0$ that satisfy $\Sigma = \mathbf{A} \cdot \text{diag}(\lambda_1^2, \dots, \lambda_l^2) \cdot \mathbf{A}^\top$.

Definition 1 (Single-Period Uncertainty Set). The uncertainty set for correlated demands at sink nodes $\mathbf{d} = (d_1, \dots, d_n)$ with variability parameters $\Gamma, \Gamma_B \geq 0$ is

$$\mathcal{U} = \left\{ \mathbf{d} \in \mathbb{R}_+^n \mid \mathbf{d} = \boldsymbol{\mu} + \mathbf{A} \cdot \mathbf{z}, \sum_{i=1}^l \left| \frac{z_i}{\lambda_i} \right| \leq \Gamma, \left| \frac{z_i}{\lambda_i} \right| \leq \Gamma_B \right. \\ \left. \forall i = 1, \dots, l \right\}. \quad (3)$$

Note that in this definition, Γ and Γ_B control the degree of conservatism. The first constraint in \mathcal{U} captures correlation, and the others are budget constraints that limit the absolute deviation from its nominal value. Although \mathcal{U} is data driven, it also captures previous results on the effect of mean and standard deviation on the profit in newsvendor networks. In particular, \mathcal{U} recovers the insightful properties in van Mieghem and Rudi (2002), as proposed in the following.

Proposition 1. For a single-period robust newsvendor network with the uncertainty set \mathcal{U} , the worst-case profit increases in μ_i and decreases in λ_i .

This proposition shows that our framework generalizes the structural properties from stochastic networks without distributional assumptions. We extend our model to multiperiod cases in the next section.

3. Multiperiod Robust Newsvendor Networks

To extend the single-period models into dynamic cases, we consider a decision maker who has multiple processing points of T periods. We assume that all parameters $\mathbf{c}_S, \mathbf{c}_H, \mathbf{c}_P, \mathbf{r}$ and matrices $\mathbf{R}_S, \mathbf{R}_D$ remain constant over the time horizon. As in Section 2, on-hand input stocks at source nodes and unsatisfied demand at sink nodes are backlogged to the next periods. We also assume that the demands are correlated over sink nodes, but independent over time, with nominal mean vector $\boldsymbol{\mu}_t$ and covariance matrix Σ_t for each time period $t = 1, \dots, T$.

Notation. Order quantities at time t are denoted by \mathbf{s}_t , customer demands by \mathbf{d}_t , and network activities by \mathbf{x}_t . Single-station quantities are denoted by s_t, d_t, x_t .

Aggregated amount of orders up to time t are denoted by \tilde{s}_t , customer demands by \tilde{d}_t , and network activities by \tilde{x}_t ($\tilde{s}_t, \tilde{d}_t, \tilde{x}_t$ for single-station). Inventory levels and backlogged demands after time t are denoted by u_t^s and u_t^d . Finally, $\mathbf{D}_{[t_1:t_2]} = (\mathbf{d}_{t_1}, \dots, \mathbf{d}_{t_2}) \in \mathbb{R}_+^{n \times (t_2 - t_1 + 1)}$ contains every realized demand from time t_1 to t_2 . Other quantities such as $\mathbf{S}_{[t_1:t_2]}$ and $\mathbf{X}_{[t_1:t_2]}$ are defined similarly. We define \mathbf{A}_t and $\lambda_{t,1}, \dots, \lambda_{t,l_t}$ for each t , with $\text{rank}(\Sigma_t) = l_t$ and $\Sigma_t = \mathbf{A}_t \cdot \text{diag}(\lambda_{t,1}^2, \dots, \lambda_{t,l_t}^2) \cdot \mathbf{A}_t^\top$.

In the following, we generalize Definition 1 for multiperiod demand.

Definition 2 (Multiperiod Uncertainty Set). The uncertainty set for the demand at sink nodes $(\mathbf{d}_1, \dots, \mathbf{d}_T) \in \mathbb{R}^{n \times T}$ over T periods is

$$\mathcal{U}^T = \left\{ (\mathbf{d}_1, \dots, \mathbf{d}_T) \mid \mathbf{d}_t = \boldsymbol{\mu}_t + \mathbf{A}_t \mathbf{z}_t \quad \forall t = 1, \dots, T \right. \\ \left. \sum_{t=1}^T \sum_{i=1}^{l_t} \left| \frac{z_{t,i}}{\lambda_{t,i}} \right| \leq \Gamma, \sum_{i=1}^{l_t} \left| \frac{z_{t,i}}{\lambda_{t,i}} \right| \leq \Gamma_t, \left| \frac{z_{t,i}}{\lambda_{t,i}} \right| \leq \Gamma_B \right. \\ \left. \forall i = 1, \dots, l_t, t = 1, \dots, T \right\}.$$

In this set, the additional constraint controls the absolute deviation over nodes and time periods. It prevents the demand to take extreme values in every period t , which reduces the conservatism over time. This definition can also describe seasonality of demands, which applies to many areas. When there is an explicit time dependence between the periods, \mathcal{U}^T can be expressed as a conic set (Nohadani and Roy 2017).

For the multiperiod newsvendor networks, we can express the dynamics of inventories and backlogged demands in (1) with vectors and matrices as

$$u_t^s = u_{t-1}^s + \mathbf{s}_t - \mathbf{R}_S \mathbf{x}_t = \sum_{\tau=1}^t (\mathbf{s}_\tau - \mathbf{R}_S \mathbf{x}_\tau) \\ u_t^d = u_{t-1}^d + \mathbf{d}_t - \mathbf{R}_D \mathbf{x}_t = \sum_{\tau=1}^t (\mathbf{d}_\tau - \mathbf{R}_D \mathbf{x}_\tau),$$

and model a multistage robust optimization problem as

$$\max_{\tilde{\mathbf{s}}_1(\mathbf{D}_{[1:t-1]})} \min_{\mathbf{D}_{[1:T]} \in \mathcal{U}^T} \max_{\tilde{\mathbf{x}}_t \in \mathcal{P}(\tilde{\mathbf{s}}_t, \tilde{\mathbf{d}}_t, \tilde{\mathbf{x}}_{t-1})} \\ \left[-\mathbf{c}_S^\top \tilde{\mathbf{s}}_T(\mathbf{D}_{[1:T-1]}) - \mathbf{c}_H^\top \sum_{t=1}^T [\tilde{\mathbf{s}}_t(\mathbf{D}_{[1:t-1]}) - \mathbf{R}_S \tilde{\mathbf{x}}_t(\mathbf{D}_{[1:t]})] \right. \\ \left. - \mathbf{c}_P^\top \sum_{t=1}^T [\tilde{\mathbf{d}}_t - \mathbf{R}_D \tilde{\mathbf{x}}_t(\mathbf{D}_{[1:t]})] + \mathbf{r}^\top \tilde{\mathbf{x}}_T(\mathbf{D}_{[1:T]}) \right], \quad (4)$$

where $\mathbf{D}_{[1:0]} = \mathbf{0}$, $\tilde{\mathbf{x}}_0 = \mathbf{0}$. Note that $\tilde{\mathbf{x}}_t$ is determined after $\tilde{\mathbf{s}}_t$ and $\tilde{\mathbf{d}}_t$, within a set

$$\mathcal{P}(\tilde{\mathbf{s}}_t, \tilde{\mathbf{d}}_t, \tilde{\mathbf{x}}_{t-1}) = \left\{ \tilde{\mathbf{x}}_t \in \mathbb{R}_+^n \mid \mathbf{R}_S \tilde{\mathbf{x}}_t \leq \tilde{\mathbf{s}}_t, \mathbf{R}_D \tilde{\mathbf{x}}_t \leq \tilde{\mathbf{d}}_t, \tilde{\mathbf{x}}_t \geq \tilde{\mathbf{x}}_{t-1} \right\},$$

which is defined for $\tilde{\mathbf{x}}_t$ to maximize profit, where the last constraint requires nonnegative network activities. The main difference between single-period and multiperiod models is that the order quantities are not *static*. That means, in order to obtain an optimal solution, one should find $\tilde{\mathbf{s}}_t$ as a function of $\mathbf{D}_{[1:t-1]}$ so that they are *fully adjustable* to all previous demands. Such policies also need to be *nonanticipative*—that is, adjustable decisions should only be based on realized uncertainties.

Even for $T = 1$, the problem (4) is a two-stage robust optimization problem and shown to be nondeterministic polynomial time (NP)-hard (Ben-Tal et al. 2004). For a multiperiod setting, the complexity only worsens, and, to our knowledge, no tractable algorithm has been proposed to exactly solve the general problem in (4). Because of this, restrictions to specific policies have been considered. In particular, *affine policies* have been proposed, where adaptive decisions are assumed to be affine functions of realized uncertainties.

Definition 3 (Affine Policy). A policy is called an affine policy if there exist $\{\mathbf{w}_t \in \mathbb{R}^m : 1 \leq t \leq T\}$ and $\{\mathbf{W}_{\tau,t} \in \mathbb{R}^{m \times n} : 1 \leq \tau \leq t-1, 1 \leq t \leq T\}$ such that

$$\mathbf{s}_1 = \mathbf{w}_1, \quad \mathbf{s}_t = \mathbf{w}_t + \sum_{\tau=1}^{t-1} \mathbf{W}_{\tau,t} \mathbf{d}_\tau \quad t = 2, \dots, T. \quad (5)$$

Affine policies have exhibited excellent performance in many real-world applications. With such policies, the multiperiod problem (4) converts to determining the affine weights. These policies force nonanticipativity of \mathbf{s}_t , and one can reformulate (4) as a two-stage adaptive linear optimization problem

$$\max_{\mathbf{w}_t, \mathbf{W}_{\tau,t}} \min_{\mathbf{d}_t} \max_{\tilde{\mathbf{x}}_t} \left[-\mathbf{c}_S^\top \tilde{\mathbf{s}}_T - \mathbf{c}_H^\top \sum_{t=1}^T (\tilde{\mathbf{s}}_t - \mathbf{R}_S \tilde{\mathbf{x}}_t) \right. \\ \left. - \mathbf{c}_P^\top \sum_{t=1}^T (\tilde{\mathbf{d}}_t - \mathbf{R}_D \tilde{\mathbf{x}}_t) + \mathbf{r}^\top \tilde{\mathbf{x}}_T \right] \quad (6)$$

$$\text{s.t. } \mathbf{w}_t + \sum_{\tau=1}^{t-1} \mathbf{W}_{\tau,t} \mathbf{d}_\tau \geq \mathbf{0}, \quad \mathbf{w}_1 \geq \mathbf{0} \quad \forall t = 2, \dots, T, \\ \forall (\mathbf{d}_1, \dots, \mathbf{d}_T) \in \mathcal{U}^T \quad (7)$$

$$\left. \begin{aligned} \tilde{\mathbf{s}}_1 &= \mathbf{w}_1 \\ \tilde{\mathbf{s}}_t &= \mathbf{w}_t + \sum_{j=2}^t (\mathbf{w}_j + \sum_{\tau=1}^{j-1} \mathbf{W}_{\tau,j} \mathbf{d}_\tau) \\ \mathbf{R}_S \tilde{\mathbf{x}}_t &\leq \tilde{\mathbf{s}}_t \\ \mathbf{R}_D \tilde{\mathbf{x}}_t &\leq \tilde{\mathbf{d}}_t \\ \tilde{\mathbf{x}}_T &\geq \tilde{\mathbf{x}}_{T-1} \geq \dots \geq \tilde{\mathbf{x}}_1 \geq \mathbf{0} \end{aligned} \right\} \quad \forall t = 1, \dots, T. \quad (8)$$

Constraint (7) implies that the order quantities are nonnegative for any realizations of past demands, and constraint (8) affects the inner maximization problem, which determines the processing activities.

Proposition 2. Finding an optimal affine policy for a multiperiod newsvendor network in (6–8) is a convex optimization problem.

Remark 1. The network activities \mathbf{x}_t maximize the net profit over the entire horizon, not just at time t —that is, we relax nonanticipativity of \mathbf{x}_t in the optimization problem (6–8). However, we claim that this relaxation will not be loose, because penalty cost and holding cost force \mathbf{x}_t to maximize profit in the corresponding period. As a special case, one can show that in single-station models, \mathbf{x}_t maximizes the overall profit if and only if it maximizes the profit at time t . This relaxation facilitates generality, as problem (6–8) is defined for any polyhedral uncertainty sets, whereas in the stochastic case, optimal strategies are only available for restricted cases (demands are independent and identically distributed over time as in van Mieghem and Rudi 2002).

Because the inner minimization problem over \mathbf{d}_t in (6–8) is nonconvex, the overall problem is solved with cut generation. If the uncertainty set \mathcal{U}^T is a polyhedron, then it is guaranteed to find an optimal solution within finite number of iterations (cuts). Therefore, our method only requires a polyhedral structure of \mathcal{U}^T . Note that the solution procedure does not exploit a specific structure of our uncertainty sets in Definition 2, and the main purpose of using the budgeted uncertainty sets is to reduce conservatism. Even though an optimal solution can be obtained within finite iterations, the problem is still NP-hard (Ben-Tal et al. 2004), and the computation grows significantly as T increases.

Although affine policies assume affine dependence of order quantities to realized demands, one can think of another class of policies, for which both ordering decisions and network activities are given as affine functions. We call these policies as *affine-approximation policies*, which are defined below.

Definition 4 (Affine-Approximation Policy). An adaptive policy is called an affine-approximation policy (Aff-approx) if there are $\{\mathbf{w}_t \in \mathbb{R}^m : 1 \leq t \leq T\}$ and $\{\mathbf{W}_{\tau,t} \in \mathbb{R}^{m \times n} : 1 \leq \tau \leq t-1, 1 \leq t \leq T\}$, $\{\mathbf{y}_t \in \mathbb{R}^p : 1 \leq t \leq T\}$ and $\{\mathbf{Y}_{\tau,t} \in \mathbb{R}^{p \times n} : 1 \leq \tau \leq t, 1 \leq t \leq T\}$ such that

$$\begin{aligned} \mathbf{s}_1 &= \mathbf{w}_1, \quad \mathbf{s}_t = \mathbf{w}_t + \sum_{\tau=1}^{t-1} \mathbf{W}_{\tau,t} \mathbf{d}_\tau \quad t = 2, \dots, T \\ \mathbf{x}_t &= \mathbf{y}_t + \sum_{\tau=1}^t \mathbf{Y}_{\tau,t} \mathbf{d}_\tau \quad t = 1, \dots, T. \end{aligned} \quad (9)$$

Note that \mathbf{x}_t depends on $\mathbf{D}_{[1:t]}$, whereas \mathbf{s}_t is a function of $\mathbf{D}_{[1:t-1]}$, as network activities are assigned after the demand is realized at each time. We provide two observations for affine-approximation policies.

(1) Aff-approx policies find affine parameters \mathbf{y}_t and $\mathbf{Y}_{\tau,t}$, so that \mathbf{x}_t is feasible to the innermost max operator. Thus, they give lower bounds to the affine policies.

(2) When Aff-approx policies are used, the problem (6–8) converts to a maximum-minimum (max-min) problem, where all the affine weights (both for \mathbf{s}_t and \mathbf{x}_t) are determined in the outer max operator. Using (9), both the objective function and the constraints can be expressed as functions of $\mathbf{w}_t, \mathbf{W}_{\tau,t}, \mathbf{y}_t, \mathbf{Y}_{\tau,t}$, and \mathbf{d}_t , which possibly include bilinear terms between the affine weights and \mathbf{d}_t . This type of problem is referred to as a (static) robust linear optimization problem, and they can be reformulated to a linear program, whenever \mathcal{U}^T is a polyhedron. Therefore, Aff-approx policies are tractable.

So far, we formulated multiperiod robust newsvendor network problems and introduced two policies. Affine policies convert the multiperiod problem into a two-stage problem, which is computationally intractable. On the contrary, Aff-approx policies solve a tractable linear program, but it only provides suboptimal solutions to affine policies (we will numerically study their suboptimality in Section 6). Our main contribution is motivated by taking an alternative approach to these policies, as presented in the next section.

4. Periodic-Affine Policies for Single-Station Models

As discussed, affine policies face computational difficulties when a decision maker has a larger number of resources and products over an extended period of time. We propose a new solution concept, denoted as *periodic-affine policies* (PA), where the overall time horizon is separated into subperiods, which are interconnected by the preceding surplus to become the proceeding demand. In this approach, the order quantities are determined as an affine function of past demands realized only within its subperiod, as opposed to affine and affine-approximation policies where all previous demands are considered. This scheme reduces the number of decision variables and consequently the computation time. Our framework constructs this policy by first formulating a dynamic programming (DP) problem, where each stage corresponds to a subperiod. We also propose an algorithm to compute such periodic-affine policies and show that they are computationally more tractable than affine policies. In addition, we present a sufficient condition that this algorithm provides the optimal solution to the DP problem. We first consider T -period single-station models in this section. However, our framework is naturally extended to multistation networks, which we discuss in the subsequent section.

Notation

We use same notations for all cost parameters c_S, c_H, c_P with revenue per item, r , and we may assume that $R_S = R_D = 1$ without loss of generality in single-station models. In this section, $\mathbf{d} = (d_1, \dots, d_T)$, $\mathbf{s} = (s_1, \dots, s_T)$, and $\mathbf{x} = (x_1, \dots, x_T)$. We denote $\pi(w_t, W_{\tau,t})$ as an affine

policy with affine parameters $\{w_t, W_{\tau,t} : 1 \leq \tau \leq t-1, 1 \leq t \leq T\}$. Furthermore, the problem of a T -period single-station newsvendor model is denoted as $\Phi(s_0, d_0)$ for an uncertainty set \mathcal{U}^T with initial input $s_0 \geq 0$ and demands $d_0 \geq 0$.

Analysis of Initial Input and Demand

We first study the role of initial input and demand for the optimal affine policy in the multiperiod model, given by

$$\Phi(s_0, d_0) := \max_{\pi} \min_{\mathbf{d} \in \mathcal{U}^T} \max_{\mathbf{x} \in \mathcal{X}(\pi, \mathbf{d}, s_0, d_0)} P(\pi(w_t, W_{\tau,t}), \mathbf{d}, \mathbf{x}; s_0, d_0),$$

where the profit during the period is

$$\begin{aligned} P(\pi(w_t, W_{\tau,t}), \mathbf{d}, \mathbf{x}; s_0, d_0) \\ = -c_S \left(\sum_{t=1}^T s_t \right) - c_H \sum_{t=1}^T \left(s_0 + \sum_{\tau=1}^t (s_\tau - x_\tau) \right) \\ - c_P \sum_{t=1}^T \left(d_0 + \sum_{\tau=1}^t (d_\tau - x_\tau) \right) + r \left(\sum_{t=1}^T x_t \right), \end{aligned}$$

and the feasible set $\mathcal{X}(\pi, \mathbf{d}, s_0, d_0)$ is given by

$$\mathcal{X}(\pi, \mathbf{d}, s_0, d_0) = \left\{ \mathbf{s}, \mathbf{x} \geq \mathbf{0} \left| \begin{array}{ll} s_1 = w_1, s_t = w_t + \sum_{\tau=1}^{t-1} W_{\tau,t} d_\tau & \forall t = 2, \dots, T \\ \sum_{\tau=1}^t x_\tau \leq s_0 + \sum_{\tau=1}^t s_\tau & \forall t = 1, \dots, T \\ \sum_{\tau=1}^t x_\tau \leq d_0 + \sum_{\tau=1}^t d_\tau & \forall t = 1, \dots, T \end{array} \right. \right\}. \quad (10)$$

The result is intuitive and plays a key role in establishing periodic-affine policies.

Proposition 3. For an optimal affine policy $\pi^*(w_t^*, W_{\tau,t}^*)$ of $\Phi(0, 0)$ with no initial input and demand, if $s_0 \leq w_1^*$, then:

(1) An optimal affine policy $\bar{\pi} = \bar{\pi}(\bar{w}_t^*, \bar{W}_{\tau,t}^*)$ of $\Phi(s_0, d_0)$ is characterized as

$$\begin{cases} \bar{w}_1^* = w_1^* - s_0 + d_0 \\ \bar{w}_t^* = w_t^* & \forall t = 2, \dots, T \\ \bar{W}_{\tau,t}^* = W_{\tau,t}^* & \forall \tau = 1, \dots, t-1, \forall t = 1, \dots, T. \end{cases} \quad (11)$$

(2) There exists a single worst-case demand $\mathbf{d}^* \in \mathcal{U}^T$ for both $\Phi(0, 0)$ and $\Phi(s_0, d_0)$.

This result implies that for small enough connecting inventories, the subperiods become effectively decoupled.

4.1. Model Formulation

We now introduce the DP formulation for a multiperiod newsvendor network.

Notation. For a T -period single-station model, we partition the time period into N subperiods sorted as $0 = t_0 < t_1 < \dots < t_{N-1} < t_N = T$. In interval $I_j = \{t_{j-1}+1, \dots, t_j\}$, the uncertainty set $\mathcal{U}^j \in \mathbb{R}_+^{|I_j|}$ for every $j = 1, \dots, N$. The amount of on-hand input stock and backlogged demands after time t are u_t^s and u_t^d . A class of affine policies for j th subperiod is denoted by $\Pi_{\text{aff}}(\mathcal{U}^j, \Xi_{j-1})$ on the uncertainty set \mathcal{U}^j , where the state Ξ_{j-1} contains all past information at the beginning of j th period with $\Xi_0 = 0$. In this section, $\mathbf{d}_j = (d_{t_{j-1}+1}, \dots, d_{t_j}) \in \mathbb{R}_+^{|I_j|}$ denotes the demand at the j th subperiod for $j = 1, \dots, N$. We proceed similarly for \mathbf{s}_j and \mathbf{x}_j .

DP Formulation. We consider an N -stage robust DP problem, where each stage corresponds to each subperiod. At the beginning of the j th subperiod, a decision maker obtains an affine policy $\pi_j \in \Pi_{\text{aff}}(\mathcal{U}^j, \Xi_{j-1})$ to make adaptive ordering decisions for the current subperiod. This can be formulated as

$$\begin{aligned} \max_{\pi_1 \in \Pi_{\text{aff}}(\mathcal{U}^1, \Xi_0)} & \left[\min_{\mathbf{d}_1 \in \mathcal{U}^1} \max_{\mathbf{x}_1 \in \mathcal{X}_1} \left[P_1(\pi_1, \mathbf{d}_1, \mathbf{x}_1; 0, 0) \right. \right. \\ & + \max_{\pi_2 \in \Pi_{\text{aff}}(\mathcal{U}^2, \Xi_1)} \left[\min_{\mathbf{d}_2 \in \mathcal{U}^2} \max_{\mathbf{x}_2 \in \mathcal{X}_2} \left[P_2(\pi_2, \mathbf{d}_2, \mathbf{x}_2; u_{t_1}^s, u_{t_1}^d) \dots \right. \right. \\ & + \max_{\pi_N \in \Pi_{\text{aff}}(\mathcal{U}^N, \Xi_{N-1})} \left[\min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}_N} P_N(\pi_N, \mathbf{d}_N, \mathbf{x}_N; \right. \\ & \left. \left. \left. u_{t_{N-1}}^s, u_{t_{N-1}}^d) \right] \dots \right] \right], \quad (12) \end{aligned}$$

where $P_j(\pi_j, \mathbf{d}_j, \mathbf{x}_j; u_{t_{j-1}}^s, u_{t_{j-1}}^d)$ is a profit generated during the j th subperiod with an initial input and demand

$$\begin{aligned} P_j(\pi_j, \mathbf{d}_j, \mathbf{x}_j; u_{t_{j-1}}^s, u_{t_{j-1}}^d) \\ = -c_S \left(\sum_{t \in I_j} s_t \right) - c_H \sum_{t \in I_j} \left(u_{t_{j-1}}^s + \sum_{\tau=t_{j-1}+1}^t (s_\tau - x_\tau) \right) \\ - c_P \sum_{t \in I_j} \left(u_{t_{j-1}}^d + \sum_{\tau=t_{j-1}+1}^t (d_\tau - x_\tau) \right) + r \left(\sum_{t \in I_j} x_t \right), \quad (13) \end{aligned}$$

and \mathbf{s}_j and \mathbf{x}_j are determined within a feasible set $\mathcal{X}_j = \mathcal{X}(\pi_j, \mathbf{d}_j, u_{t_{j-1}}^s, u_{t_{j-1}}^d)$ from (10). In (13), s_t denotes the order quantity at time t , if ordering decision is made by π_j and \mathbf{d}_j is realized. The only constraint for π_1, \dots, π_N is to ensure nonnegative order quantities for any demand realizations.

Recall that affine and Aff-approx policies are obtained before any demand realizations in (6–8), and the parameters are fixed over all time periods. In the DP formulation (12), the affine parameters are chosen dynamically at each subperiod, depending on the past information. Specifically, the affine parameters $w_t^{(j)}(\Xi_{j-1})$ and $W_{\tau,t}^{(j)}(\Xi_{j-1})$ of π_j can be any function of

the past realization Ξ_{j-1} . Hence, we can write the order quantities in (12) as

$$s_t(\mathbf{d}_j, \Xi_{j-1}) = \begin{cases} w_1^{(j)}(\Xi_{j-1}) & t = t_{j-1} + 1 \\ w_i^{(j)}(\Xi_{j-1}) + \sum_{\tau=1}^{i-1} W_{\tau,i}^{(j)}(\Xi_{j-1})d_{t_{j-1}+\tau} & t = t_{j-1} + i, i \geq 2, t \in I_j. \end{cases}$$

Every feasible affine and Aff-approx policy ensures nonnegative orders and, hence, is also feasible in (12). This implies that by solving (12), one can propose policies that have better worst-case profit than an optimal affine policy.

Periodic-Affine Policy Formulation. With initial input u^s bounded above with $w_1^{(j)}$, we define *affine-Initial Base-Stock (affine-IBS)* policies by modifying the initial period of an affine policy $\pi_j(w_t^{(j)}, W_{\tau,t}^{(j)})$.

Definition 5 (Affine-IBS). For j th subperiod, the *affine Initial Base-Stock policy* $\pi_j(w_t^{(j)}, W_{\tau,t}^{(j)})$ associated with an affine policy $\pi_j(w_t^{(j)}, W_{\tau,t}^{(j)})$ determines order quantity by

$$s_t(u^s, u^d, \mathbf{d}_j) = \begin{cases} w_1^{(j)} - u^s + u^d & t = t_{j-1} + 1 \\ w_i^{(j)} + \sum_{\tau=1}^{i-1} W_{\tau,i}^{(j)}d_{t_{j-1}+\tau} & t = t_{j-1} + i, i \geq 2, t \in I_j. \end{cases}$$

Note that at each subperiod, affine-IBS policies adapt to initial input and demand by adjusting the order quantity at the first period. From the second period, affine-IBS and its associated affine policies are equivalent.

We now consider a sequence of affine-IBS policies $\bar{\pi} = (\bar{\pi}_1, \dots, \bar{\pi}_N)$, where each $\bar{\pi}_j = \bar{\pi}_j(w_t^{(j)}, W_{\tau,t}^{(j)})$ is for j th subperiod. Note that this policy may not be well defined for each subperiod because it does not guarantee that every order quantity is nonnegative. That means, if an input stock after t_j is greater than $w_1^{(j+1)}$, then the policy would not be feasible. To account for this, we impose

$$w_1^{(j+1)} \geq u_{t_j}^s = u_{t_j}^s(\bar{\pi}_j, \mathbf{d}_j) \quad \forall \mathbf{d}_j \in \mathcal{U}^j \quad \forall j = 1, \dots, N-1.$$

By Definition 5, the right-hand side is equivalent to

$$\begin{aligned} u_{t_j}^s(\bar{\pi}_j, \mathbf{d}_j) &= \max \left(0, \left(u_{t_{j-1}}^s + \sum_{t \in I_j} s_t(u_{t_{j-1}}^s, u_{t_{j-1}}^d, \mathbf{d}_j) \right) \right. \\ &\quad \left. - \left(u_{t_{j-1}}^d + \sum_{t \in I_j} d_t \right) \right) \\ &= \max \left(0, \left(\sum_{t=1}^{t_j-t_{j-1}} w_t^{(j)} + \sum_{t=2}^{t_j-t_{j-1}} \sum_{\tau=1}^{t-1} W_{\tau,t}^{(j)}d_{t_{j-1}+\tau} \right) \right. \\ &\quad \left. - \sum_{t=1}^{t_j-t_{j-1}} d_{t_{j-1}+t} \right). \end{aligned}$$

Because $w_1^{(j+1)}$ has to be nonnegative for any demand realization, the periodic-affine policy is well-defined if

$$w_1^{(j+1)} \geq \theta_j^* := \max_{\mathbf{d}_j \in \mathcal{U}^j} \left[\sum_{t=1}^{t_j-t_{j-1}} w_t^{(j)} + \sum_{t=2}^{t_j-t_{j-1}} \sum_{\tau=1}^{t-1} W_{\tau,t}^{(j)}d_{t_{j-1}+\tau} - \sum_{t=1}^{t_j-t_{j-1}} d_{t_{j-1}+t} \right], \quad (14)$$

for every $j = 0, \dots, N-1$, where θ_j^* denotes the maximum leftover input after j th subperiod with $\theta_0^* = 0$. Now we can define periodic-affine policies.

Definition 6 (Periodic-Affine Policy). A *periodic-affine policy* $\bar{\pi}_{\text{PA}} := (\bar{\pi}_1, \dots, \bar{\pi}_N)$ is an affine-IBS policy $\bar{\pi}_i$ satisfying (14) for affine policies π_i .

For a periodic-affine policy $\bar{\pi}_{\text{PA}} = (\bar{\pi}_1, \dots, \bar{\pi}_N)$, where $\bar{\pi}_j = \bar{\pi}_j(w_t^{(j)}, W_{\tau,t}^{(j)})$, order quantities at time t are determined as follows:

$$\begin{aligned} s_t &= s_t(\bar{\pi}_{\text{PA}}) \\ &= \begin{cases} w_1^{(j)} - u_{t_{j-1}}^s + u_{t_{j-1}}^d & \forall t = t_{j-1} + 1 \\ w_i^{(j)} + \sum_{\tau=1}^{i-1} W_{\tau,i}^{(j)}d_{t_{j-1}+\tau} & \forall t = t_{j-1} + i, i \geq 2, t \in I_j, \end{cases} \end{aligned} \quad (\text{PA})$$

for every $j = 1, \dots, N$. Because $\bar{\pi}_j \in \Pi_{\text{aff}}(\mathcal{U}^j, \Xi_{j-1})$ and $\bar{\pi}_{\text{PA}}$ satisfies (14), every periodic-affine policy is a feasible solution to the DP in (12).

In the next section, we present our algorithm to compute periodic-affine policies.

4.2. Periodic-Affine Algorithm

Our algorithm obtains affine-IBS policies for each subperiod by solving smaller subproblems. However, because affine-IBS policies take initial input and demands into account, we construct the objective function to account for leftover resources and demands. We identify such objective functions from the DP problem (12). We first show that if initial input is small, an affine-IBS policy will be optimal among $\Pi_{\text{aff}}(\mathcal{U}^N, \Xi_{N-1})$. The proof is similar to the Proposition 3 and is omitted.

Corollary 1. Let $\pi_N(w_t^{(N)}, W_{\tau,t}^{(N)})$ be an optimal affine policy with zero initial input and demands. If $u_{t_{N-1}}^s \leq w_1^{(N)}$ for any realization of $u_{t_{N-1}}^s$, then its associated affine-IBS policy $\bar{\pi}_N$ is an optimal solution among $\Pi_{\text{aff}}(\mathcal{U}^N, \Xi_{N-1})$. Moreover,

$$\begin{aligned} &\max_{\pi \in \Pi_{\text{aff}}(\mathcal{U}^N, \Xi_{N-1})} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}^N} [P_N(\pi, \mathbf{d}_N, \mathbf{x}_N; u_{t_{N-1}}^s, u_{t_{N-1}}^d)] \\ &= c_S u_{t_{N-1}}^s + (r - c_S) u_{t_{N-1}}^d + \max_{\pi \in \Pi_{\text{aff}}(\mathcal{U}^N, 0)} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}^N} [P_N(\pi, \mathbf{d}_N, \mathbf{x}_N; 0, 0)]. \end{aligned} \quad (15)$$

Using this corollary, we reformulate an optimality condition for the last stage as

$$\begin{aligned} & \max_{\pi_N \in \Pi_{\text{aff}}(\mathcal{U}^N, \Xi_{N-1})} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}_N} \left[P_N(\pi_N, \mathbf{d}_N, \mathbf{x}_N; u_{t_{N-1}}^s, u_{t_{N-1}}^d) \right] \\ &= c_S u_{t_{N-1}}^s + (r - c_S) u_{t_{N-1}}^d \\ &+ \max_{\pi_N \in \Pi_{\text{aff}}(\mathcal{U}^N, 0)} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}_N} \left[P_N(\pi_N, \mathbf{d}_N, \mathbf{x}_N; 0, 0) \right]. \end{aligned}$$

In single-station cases, $u_{t_{N-1}}^s$ and $u_{t_{N-1}}^d$ can be rewritten as

$$\begin{aligned} u_{t_{N-1}}^s &= u_{t_{N-2}}^s + \sum_{t \in I_{N-1}} (s_t - x_t), \\ u_{t_{N-1}}^d &= u_{t_{N-2}}^d + \sum_{t \in I_{N-1}} (d_t - x_t). \end{aligned}$$

This can be incorporated with $P_{N-1}(\pi_{N-1}, \mathbf{d}_{N-1}, \mathbf{x}_{N-1}; u_{t_{N-2}}^s, u_{t_{N-2}}^d)$ as

$$\begin{aligned} & P_{N-1}(\pi_{N-1}, \mathbf{d}_{N-1}, \mathbf{x}_{N-1}; u_{t_{N-2}}^s, u_{t_{N-2}}^d) + c_S u_{t_{N-1}}^s + (r - c_S) u_{t_{N-1}}^d \\ &= -c_S \left(\sum_{t \in I_{N-1}} s_t \right) - c_H \sum_{t \in I_{N-1}} \left(u_{t_{N-2}}^s + \sum_{\tau=t_{N-2}+1}^t (s_\tau - x_\tau) \right) \\ &\quad - c_P \sum_{t \in I_{N-1}} \left(u_{t_{N-2}}^d + \sum_{\tau=t_{N-2}+1}^t (d_\tau - x_\tau) \right) \\ &\quad + r \left(\sum_{t \in I_{N-1}} x_t \right) + c_S \left(u_{t_{N-2}}^s + \sum_{t \in I_{N-1}} (s_t - x_t) \right) \\ &\quad + (r - c_S) \left(u_{t_{N-2}}^d + \sum_{t \in I_{N-1}} (d_t - x_t) \right) \\ &= c_S u_{t_{N-2}}^s + (r - c_S) u_{t_{N-2}}^d - c_H \sum_{t \in I_{N-1}} \left(u_{t_{N-2}}^s + \sum_{\tau=t_{N-2}+1}^t (s_\tau - x_\tau) \right) \\ &\quad - c_P \sum_{t \in I_{N-1}} \left(u_{t_{N-2}}^d + \sum_{\tau=t_{N-2}+1}^t (d_\tau - x_\tau) \right) + (r - c_S) \left(\sum_{t \in I_{N-1}} d_t \right). \end{aligned} \quad (16)$$

We define a modified objective function \tilde{P}_{N-1} as

$$\begin{aligned} & \tilde{P}_{N-1}(\pi_{N-1}, \mathbf{d}_{N-1}, \mathbf{x}_{N-1}; u_{t_{N-2}}^s, u_{t_{N-2}}^d) \\ &= -c_H \sum_{t \in I_{N-1}} \left(u_{t_{N-2}}^s + \sum_{\tau=t_{N-2}+1}^t (s_\tau - x_\tau) \right) \\ &\quad - c_P \sum_{t \in I_{N-1}} \left(u_{t_{N-2}}^d + \sum_{\tau=t_{N-2}+1}^t (d_\tau - x_\tau) \right) + (r - c_S) \left(\sum_{t \in I_{N-1}} d_t \right). \end{aligned} \quad (17)$$

If we assume that both $u_{t_{N-2}}^s$ and $u_{t_{N-1}}^s$ are small, we can rewrite

$$\begin{aligned} & \max_{\pi_{N-1} \in \Pi_{\text{aff}}(\mathcal{U}^{N-1}, \Xi_{N-2})} \min_{\mathbf{d}_{N-1} \in \mathcal{U}^{N-1}} \max_{\mathbf{x}_{N-1} \in \mathcal{X}_{N-1}} \\ & \left[P_{N-1}(\pi_{N-1}, \mathbf{d}_{N-1}, \mathbf{x}_{N-1}; u_{t_{N-2}}^s, u_{t_{N-2}}^d) \right. \\ & \left. + \max_{\pi_N \in \Pi_{\text{aff}}(\mathcal{U}^N, \Xi_{N-1})} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}_N} \left[P_N(\pi_N, \mathbf{d}_N, \mathbf{x}_N; u_{t_{N-1}}^s, u_{t_{N-1}}^d) \right] \right] \end{aligned}$$

$$\begin{aligned} &= \max_{\pi_{N-1} \in \Pi_{\text{aff}}(\mathcal{U}^{N-1}, \Xi_{N-2})} \min_{\mathbf{d}_{N-1} \in \mathcal{U}^{N-1}} \max_{\mathbf{x}_{N-1} \in \mathcal{X}_{N-1}} \\ & \left[P_{N-1}(\pi_{N-1}, \mathbf{d}_{N-1}, \mathbf{x}_{N-1}; u_{t_{N-2}}^s, u_{t_{N-2}}^d) + c_S u_{t_{N-1}}^s \right. \\ & \left. + (r - c_S) u_{t_{N-1}}^d + \max_{\pi_N \in \Pi_{\text{aff}}(\mathcal{U}^N, 0)} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}_N} \right. \\ & \left. \left[P_N(\pi_N, \mathbf{d}_N, \mathbf{x}_N; 0, 0) \right] \right] \\ &= \max_{\pi_{N-1} \in \Pi_{\text{aff}}(\mathcal{U}^{N-1}, \Xi_{N-2})} \min_{\mathbf{d}_{N-1} \in \mathcal{U}^{N-1}} \max_{\mathbf{x}_{N-1} \in \mathcal{X}_{N-1}} \left[c_S u_{t_{N-2}}^s + (r - c_S) u_{t_{N-2}}^d \right. \\ & \left. + \tilde{P}_{N-1}(\pi_{N-1}, \mathbf{d}_{N-1}, \mathbf{x}_{N-1}; u_{t_{N-2}}^s, u_{t_{N-2}}^d) \right] \\ & \quad + \max_{\pi_N \in \Pi_{\text{aff}}(\mathcal{U}^N, 0)} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}_N} \left[P_N(\pi_N, \mathbf{d}_N, \mathbf{x}_N; 0, 0) \right] \\ &= c_S u_{t_{N-2}}^s + (r - c_S) u_{t_{N-2}}^d + \max_{\pi_{N-1} \in \Pi_{\text{aff}}(\mathcal{U}^{N-1}, 0)} \min_{\mathbf{d}_{N-1} \in \mathcal{U}^{N-1}} \max_{\mathbf{x}_{N-1} \in \mathcal{X}_{N-1}} \\ & \left[\tilde{P}_{N-1}(\pi_{N-1}, \mathbf{d}_{N-1}, \mathbf{x}_{N-1}; 0, 0) \right] \\ & \quad + \max_{\pi_N \in \Pi_{\text{aff}}(\mathcal{U}^N, 0)} \min_{\mathbf{d}_N \in \mathcal{U}^N} \max_{\mathbf{x}_N \in \mathcal{X}_N} \left[P_N(\pi_N, \mathbf{d}_N, \mathbf{x}_N; 0, 0) \right]. \end{aligned} \quad (18)$$

Note that the second equality comes from (16), and one can verify similarly from Corollary 1 that the last equality holds. This reformulation shows that if leftover input after every subperiod is small enough, we can solve the DP problem in (12) by solving smaller subproblems. These subproblems are defined with modified objective function \tilde{P}_j with no backlogged input and demand; hence, we can solve them independently. Proceeding iteratively, we define an objective $P_j^{\text{PA}}(\pi_j, \mathbf{d}_j, \mathbf{x}_j)$ as

$$\begin{aligned} & P_j^{\text{PA}}(\pi_j, \mathbf{d}_j, \mathbf{x}_j) \\ &= \begin{cases} -c_H \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (s_\tau - x_\tau) \right) \\ -c_P \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (d_\tau - x_\tau) \right) \\ + (r - c_S) \sum_{t \in I_j} d_t, & j = 1, \dots, N-1 \\ -c_S \sum_{t \in I_j} s_t - c_H \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (s_\tau - x_\tau) \right) \\ -c_P \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (d_\tau - x_\tau) \right) + r \sum_{t \in I_j} x_t, & j = N. \end{cases} \end{aligned} \quad (19)$$

We now propose the periodic-affine algorithm. This algorithm (1) ensures that the solution is well-defined and (2) exploits the modified objective functions P_j^{PA} .

The j th subproblem can be solved by the following optimization problem

$$\begin{aligned} \max_{\pi_j \in \Pi_{\text{aff}}(\mathcal{U}, 0)} \quad & \min_{\mathbf{d}_j \in \mathcal{U}} \max_{\mathbf{x}_j, \mathbf{s}_j} P_j^{\text{PA}}(\pi_j, \mathbf{d}_j, \mathbf{x}_j) \\ \text{s.t.} \quad & (\mathbf{x}_j, \mathbf{s}_j) \in \mathcal{X}(\pi_j, \mathbf{d}_j, 0, 0) \\ & w_1^{(j)} \geq \theta_{j-1}^* \end{aligned} \quad (20)$$

The last constraint ensures that periodic-affine policy is well-defined, where θ_{j-1}^* is the maximum amount of on-hand input after $(j-1)$ th subperiod, computed by (14). The overall procedure solves (20) and (14) iteratively, as summarized in Algorithm 1.

Algorithm 1 (Periodic-Affine Algorithm for Single-Station Problems)

Given. Time indices $0 = t_0 < t_1 < \dots < t_N = T$, uncertainty set $\mathcal{U} = \mathcal{U}^1 \times \dots \times \mathcal{U}^N$, $j = 1$, and $\theta_0^* = 0$.

Step 1. Solve (20) to obtain π_j for the j th subperiod.

Step 2. Using π_j , compute the maximum leftover input θ_j^* by (14).

Step 3. If $j = N$, return $\bar{\pi}_{\text{PA}} = (\bar{\pi}_1, \dots, \bar{\pi}_N)$ and STOP. Otherwise, $j \leftarrow j + 1$ and go to *Step 1*.

We present an additional useful property of the periodic-affine algorithm, namely, that the worst-case scenario can be obtained from each iteration.

Proposition 4. Let $\bar{\pi}_{\text{PA}}$ be a solution of periodic-affine algorithm and \mathbf{d}_j^* be a worst-case scenario from the j th subproblem. Then $\bar{\pi}_{\text{PA}}$ has a worst-case scenario $(\mathbf{d}_1^*, \dots, \mathbf{d}_N^*)$.

So far, we discussed a single-station, multiperiod robust newsvendor model, where the uncertainty set over time periods is defined as a Cartesian product of prespecified uncertainty sets for each subperiod. We formulated a dynamic programming problem, where each stage corresponds to each subperiod. Motivated from this formulation, we developed an algorithm to find a periodic-affine policy, by defining the modified objective functions P_j^{PA} 's.

Computational Advantages. We provide two theoretical evidences that periodic-affine policies are computationally advantageous over affine policies. First, the number of affine parameters in PA grows linearly in T , whereas it increases quadratically in affine policies. Second, solving subproblems in PA requires significantly less computational effort than solving one large problem in affine policies. This is because an optimal affine policy solves (6–8) by cut generation. The number of extreme points in multiperiod uncertainty sets (Definition 2) grows exponentially in T , and this results in a very large number of iterations for affine policies. In PA, however, one can arbitrarily choose the length of subperiods. Hence, the number of iterations grows at most linearly in T . Strong empirical evidence in Section 6 demonstrate these advantages.

Also note that PA does not suffer from the curse of dimensionality. Even though PA is motivated from the

DP formulation (12), an optimal solution is not obtained by solving the Bellman equation; our algorithm decomposes the overall problem into smaller subproblems in order to achieve tractable solutions.

In the next section, we present theoretical properties of periodic-affine policies, where we provide a sufficient condition for the algorithm to have an optimal solution to the DP (12).

4.3. Optimality of Periodic-Affine Policies

In this section, we present theoretical properties of PA by analyzing the effect of base-stock levels on the worst-case performance. Specifically, we provide a sufficient condition under which the periodic-affine algorithm solves the DP problem (12). To compare the worst-case performance of PA with affine policies, we consider affine policies under the rectangular uncertainty set $\mathcal{U} = \mathcal{U}^1 \times \dots \times \mathcal{U}^N$ so that both policies are defined equivalently. Moreover, we present an analytical approximation for the suboptimality of PA. Note that this is a *posterior* approximation—that is, it is computed during the algorithm.

Let the worst-case profit of the two policies be V_{PA}^* and V_{Aff}^* , which are evaluated by (12), and V_{DP}^* as an optimal value of the DP problem (12).

The following assumption guarantees the optimality of PA policies.

Assumption 1. For a solution of the PA algorithm $\bar{\pi}_{\text{PA}}$, assume that the maximum leftover input level after each subperiod θ_j^* satisfies the last constraint in (20).

In other words, for a solution of the PA algorithm, the last constraint in (20) is not active at every iteration. Because an optimal periodic-affine policy maximizes the overall profit, it tends to have lower leftovers at each time period and, hence, is likely to satisfy Assumption 1. This assumption is not too restrictive as evidenced in the numerical experiments in Section 6.

Remark 2. We also suggest that it is possible to enforce Assumption 1 to hold, or reduce the suboptimality of PA, by a proper choice of partitioning the time horizon. For example, consider a problem with $T = 12$, where the nominal mean is 10 for the first 6 periods and decreases to 2 for the last 6 periods. Then, Assumption 1 may not hold if PA is obtained by partitioning 12 periods into 2 (of each 6 periods). However, one can obtain a better result by partitioning the periods into 3 subperiods of each 4 periods.

Theorem 1. For a single-station network, if Assumption 1 holds, then $V_{\text{Aff}}^* \leq V_{\text{PA}}^* = V_{\text{DP}}^*$.

Because affine policies are defined on a subset of every feasible solution of the DP, the worst-case profit of PA is guaranteed to be greater than or equal to that of affine policies under Assumption 1. We show in the following proposition that in single-station problems, their worst-case performance is indeed equal.

Proposition 5. For single-station models under Assumption 1, $V_{\text{Aff}}^* = V_{\text{PA}}^* = V_{\text{DP}}^*$.

Remark. All previous theoretical properties of PA rely on Assumption 1. In other words, the worst-case performance of PA may not match that of affine policies without the assumption.

We now relax Assumption 1 and provide a suboptimality bound for PA. For this, we compare the solutions of the following optimization problems

$$\begin{aligned} \mathcal{P}_j &:= \max_{\pi_j \in \Pi_{\text{aff}}(\mathcal{W}, 0)} \min_{\mathbf{d}_j \in \mathcal{W}} \max_{\mathbf{x}_j, \mathbf{s}_j} P_j^{\text{PA}}(\pi_j, \mathbf{d}_j, \mathbf{x}_j) \\ \text{s.t. } &(\mathbf{x}_j, \mathbf{s}_j) \in \mathcal{X}(\pi_j, \mathbf{d}_j, 0, 0) \\ &w_1^{(j)} \geq \theta_{j-1}^*, \\ \tilde{\mathcal{P}}_j &:= \max_{\pi_j \in \Pi_{\text{aff}}(\mathcal{W}, 0)} \min_{\mathbf{d}_j \in \mathcal{W}} \max_{\mathbf{x}_j, \mathbf{s}_j} P_j^{\text{PA}}(\pi_j, \mathbf{d}_j, \mathbf{x}_j) \\ \text{s.t. } &(\mathbf{x}_j, \mathbf{s}_j) \in \mathcal{X}(\pi_j, \mathbf{d}_j, 0, 0), \end{aligned}$$

where \mathcal{P}_j is the j th subproblem during the PA algorithm and $\tilde{\mathcal{P}}_j$ solves a subproblem with assuming no leftover input from the previous subperiods.

Theorem 2. For any single-station newsvendor networks with an objective value f_j^* of $\tilde{\mathcal{P}}_j$,

$$V_{\text{PA}}^* \leq V_{\text{Aff}}^* \leq V_{\text{DP}}^* \leq \tilde{V}_{\text{PA}}^* := \sum_{j=1}^N \tilde{f}_j^*.$$

Theorem 2 provides a tight bound. All the inequalities hold with equalities if Assumption 1 holds. Note that one may not need to resolve a problem $\tilde{\mathcal{P}}_j$; recall that every subproblem in the periodic-affine algorithm is solved by generating cuts. Once a solution of \mathcal{P}_j is obtained, one can relax the last constraint in \mathcal{P}_j and continue cut generation in order to solve $\tilde{\mathcal{P}}_j$. This requires fewer iterations, because (1) an optimal solution of \mathcal{P}_j serves as a warm-start initial point for additional cuts, and (2) the previously generated cuts are still valid without any modifications to $\tilde{\mathcal{P}}_j$.

So far, we introduced theoretical properties of periodic-affine policies. We showed that under mild condition, the algorithm finds an optimal solution to the DP problem, and thus achieve a worst-case performance equal to that of an optimal affine policy for single-station problems. If this assumption does not hold, we provided a tight bound that can measure the gap between periodic and affine policies. Moreover, this gap can be computed by minimally modifying the algorithm with similar computational requirements. In the proceeding section, we extend this framework to general multistation newsvendor networks and infinite-horizon problems.

5. Extensions of Periodic-Affine Policies

We extend our approach in Section 4 to general multistation networks, where a decision maker intends to satisfy customers' demand at multiple locations. First, we extend

Algorithm 1 for multistation networks. Then, we develop periodic-affine policies for infinite horizon problems.

5.1. Multistation Networks

Here, we follow the flow of Section 4. To set this up, we define a matrix that plays a key role in implementing periodic-affine policies.

Definition 7 (Basic Matrix). Let $\ell_k \in \mathcal{L}$ be an optimal solution to satisfy a unit demand at a sink node $k \in \mathcal{S}$. A basic matrix $\mathbf{R}_B \in \mathbb{R}^{p \times n}$ is given by

$$\mathbf{R}_B(\ell, k) = 1 \text{ if } \ell = \ell_k \forall k, \text{ and } 0 \text{ otherwise.}$$

Using such a basic matrix, we obtain closed-form expressions of ordering quantities and network activities if the demand is deterministic. In particular, for any $\mathbf{d} \in \mathbb{R}_+^n$, an optimal decision is given by $\mathbf{s} = \mathbf{R}_S \mathbf{R}_B \mathbf{d}$ and $\mathbf{x} = \mathbf{R}_D \mathbf{d}$.

Recall that we have defined a modified objective functions for each stage of the DP problem (12) to separate the overall problem into subproblems. In single-station models, the values of on-hand products and backlogged demands at the beginning of the $(j+1)$ th subperiod are expressed as

$$\begin{aligned} c_S u_{t_j}^s &= c_S \cdot \sum_{t \in I_j} (s_t - x_t), \\ (r - c_S) u_{t_j}^d &= (r - c_S) \cdot \sum_{t \in I_j} (d_t - x_t), \end{aligned}$$

which are taken into j th subproblem. After the j th subperiod, $\mathbf{u}_{t_j}^s$ and $\mathbf{u}_{t_j}^d$ are deterministic, and, hence, their values can be expressed by using the basic matrix \mathbf{R}_B as

$$\begin{aligned} \mathbf{c}_S^\top \mathbf{u}_{t_j}^s &= \mathbf{c}_S^\top \sum_{t \in I_j} (\mathbf{s}_t - \mathbf{R}_S \mathbf{x}_t), \\ (\mathbf{R}_B^\top \mathbf{r} - \mathbf{R}_B^\top \mathbf{R}_S^\top \mathbf{c}_S)^\top \mathbf{u}_{t_j}^d &= (\mathbf{R}_B^\top \mathbf{r} - \mathbf{R}_B^\top \mathbf{R}_S^\top \mathbf{c}_S)^\top \\ &\quad \cdot \sum_{t \in I_j} (\mathbf{d}_t - \mathbf{R}_D \mathbf{x}_t). \end{aligned}$$

Note that the value of $\mathbf{u}_{t_j}^d$ is determined by ordering $\mathbf{R}_B \mathbf{u}_{t_j}^d$ and processing $\mathbf{R}_S \mathbf{R}_B \mathbf{u}_{t_j}^d$. This allows us to extend the definition of affine-IBS policies as follows.

Definition 8 (Affine-IBS for Multistation). For j th subperiod, the affine-IBS policy $\bar{\pi}_j(\mathbf{w}_t^{(j)}, \mathbf{W}_{\tau,t}^{(j)})$ associated with an affine policy $\pi_j(\mathbf{w}_t^{(j)}, \mathbf{W}_{\tau,t}^{(j)})$ determines order quantity by

$$\mathbf{s}_t(\mathbf{u}^s, \mathbf{u}^d, \mathbf{D}_j) = \begin{cases} \mathbf{w}_1^{(j)} - \mathbf{u}^s + \mathbf{R}_S \mathbf{R}_B \mathbf{u}^d & t = t_{j-1} + 1 \\ \mathbf{w}_i^{(j)} + \sum_{\tau=1}^{i-1} \mathbf{W}_{\tau,i}^{(j)} \mathbf{d}_{t_{j-1}+\tau} & t = t_{j-1} + i, i \geq 2, t \in I_j. \end{cases}$$

Periodic-Affine Policy for Multistation Networks. As in Definition 6, periodic-affine policies are defined as a sequence of affine-IBS policies. Equation (14), which is required for periodic-affine policies to be well defined, is readily extended by replacing with a vector inequality. With this generalization, all the arguments in Section 4.2 can be repeated in multistation network setting.

As a result, the objective function for each subproblem is given by

$$\mathbf{P}_j^{\text{PA}}(\pi_j, \mathbf{D}_j, \mathbf{X}_j) = \begin{cases} -\mathbf{c}_H^\top \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (\mathbf{s}_\tau - \mathbf{R}_S \mathbf{x}_\tau) \right) \\ -\mathbf{c}_P^\top \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (\mathbf{d}_\tau - \mathbf{R}_D \mathbf{x}_\tau) \right) \\ + \mathbf{v}_d^\top \sum_{t \in I_j} \mathbf{d}_t + \mathbf{v}_x^\top \sum_{t \in I_j} \mathbf{x}_t, & j \leq N-1 \\ -\mathbf{c}_S^\top \sum_{t \in I_j} \mathbf{s}_t - \mathbf{c}_H^\top \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (\mathbf{s}_\tau - \mathbf{R}_S \mathbf{x}_\tau) \right) \\ -\mathbf{c}_P^\top \sum_{t \in I_j} \left(\sum_{\tau=t_{j-1}+1}^t (\mathbf{d}_\tau - \mathbf{R}_D \mathbf{x}_\tau) \right) + \mathbf{r}^\top \sum_{t \in I_j} \mathbf{x}_t, & j = N, \end{cases} \quad (21)$$

where $\mathbf{v}_d = \mathbf{R}_B^\top \mathbf{r} - \mathbf{R}_B^\top \mathbf{R}_S^\top \mathbf{c}_S$ and $\mathbf{v}_x = \mathbf{r} - \mathbf{R}_S^\top \mathbf{c}_S - \mathbf{R}_D^\top \mathbf{R}_B^\top \mathbf{r} + \mathbf{R}_D^\top \mathbf{R}_B^\top \mathbf{R}_S^\top \mathbf{c}_S$.

Periodic-Affine Algorithm for Multistation Networks. As in Section 4, Problem (20) can be readily converted into multidimensional form by replacing the last inequality with a vector inequality. However, it is challenging to obtain the multistation version of (14), which computes the maximum amount of leftover resources. Therefore, we incorporate these into a single robust two-stage optimization problem, as follows:

$$\begin{aligned} \max_{\theta_j, \pi_j \in \Pi_{\text{aff}}(\mathcal{U}, 0)} \quad & \min_{\mathbf{D}_j \in \mathcal{U}} \max_{\mathbf{S}_j, \mathbf{X}_j} \mathbf{P}_j^{\text{PA}}(\pi_j, \mathbf{D}_j, \mathbf{X}_j) - \delta \cdot \mathbf{1}^\top \theta_j \\ \text{s.t.} \quad & (\mathbf{S}_j, \mathbf{X}_j) \in \mathcal{X}(\pi_j, \mathbf{D}_j, 0, 0) \\ & \mathbf{w}_1^{(j)} \geq \theta_{j-1}^* \\ & \sum_{t \in I_j} (\mathbf{s}_t - \mathbf{R}_S \mathbf{x}_t) \leq \theta_j, \end{aligned}$$

where $\delta > 0$ is a small real number. In this way, the periodic-affine algorithm for multistation networks proceeds by iteratively solving subproblems, similar to Algorithm 1.

Properties of Periodic-Affine Policies for Multistation Networks. We now generalize the theoretical properties for the multistation networks. We consider the DP problem (12) by replacing every single-dimensional quantity by multidimensional quantities. Assumption 1 is extended with vector inequalities, each of which is for

each source node. We use V_{Aff}^* , V_{PA}^* , and V_{DP}^* for the worst-case objective values for affine, PA, and the DP problem, and define \tilde{V}_{PA}^* similar to single-station problems.

Theorem 3. For multistation networks, if Assumption 1 holds, then $V_{\text{Aff}}^* \leq V_{\text{PA}}^* = V_{\text{DP}}^*$. Otherwise, $V_{\text{PA}}^* \leq V_{\text{DP}}^* \leq \tilde{V}_{\text{PA}}^*$.

Note that Proposition 5 cannot be extended to multistation networks. In other words, PA policies for multistation networks are not necessarily affine policies. Theorem 3 implies that an optimal PA policy has a worst-case performance not less than an optimal affine policy. However, for multistation networks, we cannot compare the two policies without Assumption 1.

5.2. Infinite Horizon Problems

So far, PA policies are based on multiperiod problems of finite horizon. In this section, we extend these PA policies to infinite horizon problems with a discount factor of $\beta < 1$. For this, we assume that nominal means and covariances of demands have periodicity with the period $T_0 \geq 1$. We then define an uncertainty set $\mathcal{U}^{T_0} \in \mathbb{R}^{n \times T_0}$ to describe demand uncertainties for each period. This framework models settings where demand has stationary mean and covariance along the periods.

Definition 9 (Infinite-Horizon Uncertainty Set). The infinite-horizon demand uncertainty set is a Cartesian product of \mathcal{U}^{T_0} via

$$\mathcal{U}^\infty = \mathcal{U}^{T_0} \times \mathcal{U}^{T_0} \times \mathcal{U}^{T_0} \times \dots$$

We implement PA policies for infinite horizon problems by replicating policies over the periods. We construct a PA policy of period T_0 by solving a single problem of duration T_0 . As in previous sections, we define the objective function $(\pi, \mathbf{D}_{[1:T_0]}, \mathbf{X}_{[1:T_0]})$ by taking leftover inventories, unsatisfied demands, and the discount factor into account as

$$\begin{aligned} \mathbf{P}_\infty^{\text{PA}}(\pi, \mathbf{D}_{[1:T_0]}, \mathbf{X}_{[1:T_0]}) &= -\mathbf{c}_H^\top \sum_{t=1}^{T_0} \beta^t \left(\sum_{\tau=1}^t (\mathbf{s}_\tau - \mathbf{R}_S \mathbf{x}_\tau) \right) \\ &\quad - \mathbf{c}_P^\top \sum_{t=1}^{T_0} \beta^t \left(\sum_{\tau=1}^t (\mathbf{d}_\tau - \mathbf{R}_D \mathbf{x}_\tau) \right) + \mathbf{r}^\top \sum_{t=1}^{T_0} \beta^t \mathbf{x}_t \\ &\quad - \mathbf{c}_S^\top \sum_{t=1}^{T_0} \beta^t \mathbf{s}_t + \beta^{T_0+1} \mathbf{c}_S^\top \sum_{t=1}^{T_0} (\mathbf{s}_t - \mathbf{R}_S \mathbf{x}_t) \\ &\quad + \beta^{T_0+1} \left(\mathbf{R}_B^\top \mathbf{r} - \mathbf{R}_B^\top \mathbf{R}_S^\top \mathbf{c}_S \right)^\top \sum_{t=1}^{T_0} (\mathbf{d}_t - \mathbf{R}_D \mathbf{x}_t) \\ &= - \sum_{t=1}^{T_0} \beta^t \mathbf{c}_{S,t}^* \mathbf{s}_t - \sum_{t=1}^{T_0} \beta^t \mathbf{c}_{H,t}^* \left(\sum_{\tau=1}^t (\mathbf{s}_\tau - \mathbf{R}_S \mathbf{x}_\tau) \right) \\ &\quad - \sum_{t=1}^{T_0} \beta^t \mathbf{c}_{P,t}^* \left(\sum_{\tau=1}^t (\mathbf{d}_\tau - \mathbf{R}_D \mathbf{x}_\tau) \right) + \sum_{t=1}^{T_0} \beta^t \mathbf{r}_t^\top \mathbf{x}_t, \end{aligned}$$

where $\mathbf{c}_{S,t}^* = \mathbf{c}_S$, $\mathbf{r}_t^* = \mathbf{r}$ for $t = 1, \dots, T_0$, and

$$\mathbf{c}_{H,t}^* = \begin{cases} \mathbf{c}_H & 1 \leq t \leq T_0 - 1 \\ \mathbf{c}_H - \beta \mathbf{c}_S & t = T_0 \end{cases}$$

$$\mathbf{c}_{P,t}^* = \begin{cases} \mathbf{c}_P & 1 \leq t \leq T_0 - 1 \\ \mathbf{c}_P - \beta \mathbf{R}_B^\top \mathbf{r} + \beta \mathbf{R}_B^\top \mathbf{R}_S^\top \mathbf{c}_S & t = T_0. \end{cases}$$

As a result, an optimal PA policy is obtained by solving a single optimization problem

$$\begin{aligned} \max_{\theta, \pi} \quad & \min_{\mathbf{D}_{[1:T_0]} \in \mathcal{U}^{T_0}} \max_{\mathbf{S}_{[1:T_0]}, \mathbf{X}_{[1:T_0]}} \mathbf{P}_\infty^{\text{PA}}(\pi, \mathbf{D}_{[1:T_0]}, \mathbf{X}_{[1:T_0]}) \\ \text{s.t.} \quad & (\mathbf{S}_{[1:T_0]}, \mathbf{X}_{[1:T_0]}) \in \mathcal{X}(\pi, \mathbf{D}_{[1:T_0]}, \mathbf{0}, \mathbf{0}) \\ & \mathbf{w}_1 \geq \theta, \end{aligned} \quad (22)$$

where the last constraint ensures that the solution is replicable over time periods. Based on the solution of (22), an infinite PA policy determines order quantity as

$$\mathbf{s}_t = \begin{cases} \mathbf{w}_1 - \mathbf{u}_{t-1}^s + \mathbf{R}_S \mathbf{R}_B \mathbf{u}_{t-1}^d & t = nT_0 + 1, n = 0, 1, 2, \dots \\ \mathbf{w}_i + \sum_{\tau=1}^{i-1} \mathbf{W}_{\tau,i} \mathbf{d}_{nT_0+\tau} & t = nT_0 + i, 1 < i < T_0, n = 0, 1, 2, \dots \end{cases} \quad (23)$$

We next present our main result for infinite-horizon cases.

Theorem 4. *For an infinite-horizon multistation network and the uncertainty set \mathcal{U}^∞ , if Assumption 1 holds, then the infinite periodic-affine policy (23) is optimal to the DP in (12).*

In summary, we generalized the periodic-affine policies into multistation networks and infinite horizon problems. In both these cases, we presented periodic-affine algorithms and showed that the theoretical properties hold. We next discuss a numerical case study to demonstrate the practical applicability of these findings and the performance of the proposed policies.

6. Discussion: Insights and Implications

In this section, we present various implications of our modeling and solution approach and demonstrate the following advantages of our modeling approach and solution algorithm:

- **Practical Relevance:** The relevance of an approach and corresponding algorithms hinges on the ability to model features in a real-world setting and provide implementable solutions. We demonstrate that our approach is able to achieve this. In particular, we show that PA policies perform well in large-scale and data-driven environments by studying the case of a major pharmacy retailer in India. We also demonstrate that we are able to model the service level guarantees by using the robust optimization approach. This ensures that the demands are satisfied for all scenarios in an

uncertainty set. We also demonstrate the robustness to misspecification and study the performance for a spectrum of various cost parameters.

- **Generalizability and Extendability:** It is also important for the approach to be generalizable and extendable in order to accommodate higher-dimensional versions of the problem and newer types of demand information. We achieve this by modifying the uncertainty set based on the available demand information and by showing that our approach naturally extends to multidimensional settings. In particular, we incorporate correlation information in computing the optimal PA policy and demonstrate our algorithm on the high-dimensional real-world case study.

- **Computational Tractability:** An algorithm suited for real applications needs to be tractable and implementable. We demonstrate tractability of the PA policies by presenting empirical evidence on the computational times on simulated data.

Next, we elaborate on each of these advantages.

6.1. Practical Relevance: Case Study of a Pharmacy Retailer

We analyze the sales data of a leading pharmacy retailer in India to probe the performance of the policies in a real-world setting. A common problem in forecasting demands is that sales records do not necessarily imply customers' demands, because product shortage is not reflected in sales data. However, because pharmacy retailers in India face a prohibitively large penalty for unmet demand, we can interpret the sales records as demands for this numerical study.

The data consists of more than 1.5 million transactions over 40 days (end of September to early November of 2016) for 228 different products. To reduce the problem size, we analyzed the 20 most popular products, comprising nearly 80% of all transactions. Hierarchical clustering (Maechler et al. 2016) is used to bundle the products into groups, within which demands are highly correlated. Moving averages and residuals are extracted from the sales records and used as nominal means and variances to define data-driven uncertainty sets for each group. Penalty and holding costs are not revealed in the sales records. Therefore, we fix penalty and holding cost rates, and compute penalty and holding costs as a product of the corresponding rates and net profit per unit.

Uncertainty Sets. We defined the multiperiod uncertainty sets following Definition 2, where the variability parameters are defined as $\Gamma = 2\sqrt{nT}$, $\Gamma_t = 2\sqrt{n}$, and $\Gamma_B = 2$. The factor 2 is inspired from the fact that $\mathbb{P}(|Z| \leq 2) \approx 0.95$ for a standard normal random variable Z . Such a setting is typical in the robust optimization literature for describing random variables bounded by uncertainty sets (Bertsimas et al. 2017). The parameters Γ_t and Γ are chosen to be proportional to \sqrt{n} and \sqrt{nT} ,

motivated by the central limit theorem (Bertsimas and Sim 2004).

Performance of PA Policies. We compute ordering policies for each product group. We begin by assuming that the sales of different product groups are independent of each other, and later consider the more realistic case of correlated sales. We compare the performance of three policies: PA, Aff-approx, and base-stock policies for these product groups. We implement two different base-stock policies. Myopic base-stock policies are implemented, where the order-up-to level at each period is determined myopically by using the nominal means and variances, assuming normally distributed demand. We also compute the base-stock levels with sample average approximation (SAA), using the approach discussed in Bertsimas et al. (2018).

For precomputed policies, we generate random demand samples to evaluate the two policies. The samples are generated independently over time with normal distributions, in which nominal mean and variances are used. For given cost and revenue parameters, profit is calculated for each sample. We compare 5%, 25% quantiles, and median for the policies.

Table 1 displays performance of various policies for different values of the penalty cost rate. We observe that PA performs better than Aff-approx in terms of worst-case performance. In Section 6.3, we will show that this is also consistent with synthetic data. On the other hand, if the penalty cost is low, both PA and Aff-approx are not effective, and the base-stock policy outperforms them. However, PA yields better lower-percentile performance than the other policies for increased penalty cost. This is because PA maximizes

the worst-case profit. We also observe that under significant penalty costs, PA not only protects the worst-case performance and lower percentiles (improves by 19% over base-stock at 5th percentile for cost of 10) but also leads to better average profit and historical backtesting than the other policies. We also notice that although SAA improves over the myopic policy, both Aff-approx and PA outperform it.

Sensitivity to Model Misspecification. Given that all these policies are implemented by using the nominal mean and variance inferred from past records, it is important to measure their robustness to errors in model calibration. For this, we consider demand realizations to have mean greater than (Figure 2(a)), same as (Figure 2(b)), or less than (Figure 2(c)) their nominal values for varying holding and penalty costs. We observe that when the realized demand distribution differs from the assumed one, the region of parameters (*phase*), for which PA outperforms base-stock policy, changes. For example, for a holding cost of 0.1 and penalty of 1, the PA policy outperforms base-stock, if the mean of the assumed demand coincides with the realized one (see Figure 2(b)). However, only 5% increase of the means is sufficient for the base-stock to prevail (see Figure 2(a)).

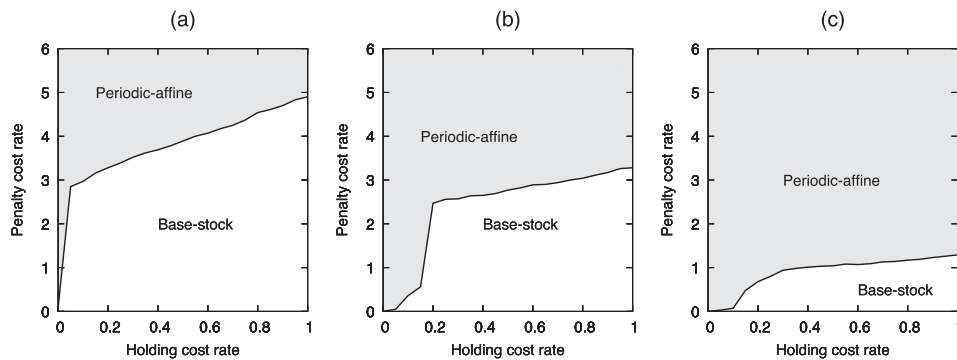
Performance Dependence on Holding and Penalty Cost. As the costs vary, we observe a *phase transition* between a phase where PA outperforms the base-stock policies and a phase in which the base-stock policy outperforms. The phase diagrams in Figure 2 allow the decision maker to select the policy based on the given cost and demand structure. In fact, for pharmacy retailers, who face a

Table 1. Performance of Policies for Different Penalty Cost Rates

Penalty cost rate	Policy	Worst	5% quantile	25% quantile	Median	Historical
0.2	Affine approximation	2.68	3.50	3.64	3.74	3.81
	Periodic-affine	2.72	3.64	3.77	3.85	3.97
	Sample average approximation	N/A	3.73	4.02	4.13	4.41
	Myopic	N/A	3.80	4.16	4.37	4.67
1.0	Affine approximation	1.96	2.87	3.27	3.51	3.86
	Periodic-affine	2.49	3.30	3.41	3.48	3.56
	Sample average approximation	N/A	3.43	3.53	3.81	3.96
	Myopic	N/A	3.27	3.66	3.96	4.25
5.0	Affine approximation	1.49	3.31	3.60	3.79	4.04
	Periodic-affine	1.78	3.37	3.64	3.82	3.83
	Sample average approximation	N/A	3.21	3.58	3.61	3.52
	Myopic	N/A	2.85	3.27	3.58	3.68
10.0	Affine approximation	1.26	3.10	3.45	3.68	4.15
	Periodic-affine	1.39	3.28	3.62	3.84	3.77
	Sample average approximation	N/A	3.12	3.39	3.64	3.61
	Myopic	N/A	2.76	3.19	3.49	3.40

Notes. All values are multiplied by 10^6 . Percentiles and medians are calculated from 1,000 samples. The last column is from the historical sales data. N/A, not applicable.

Figure 2. Phase Diagram of Periodic-Affine and the Base-Stock Policies



Note. The realized demand means are increased by 5% (a), not changed (b), and decreased by 5% (c) from the nominal values.

substantial penalty with unsatisfied demand, it shows that our proposed PA policy is preferable. On the other hand, if the decision maker is committed to a certain ordering policy (e.g., contractually), the phase diagrams in Figure 2 can suggest suitable changes to the cost structure in order to make the policy superior.

Impact of High Penalty Cost. When analyzing backlogged demands and inputs for different values of the penalty cost, Figure 3 shows that the three policies react differently for high penalty costs. First, the base-stock policy does not effectively control the backlogged demands. Although the penalty cost is accounted for in the newsvendor quantile to avoid high backlogs, increasing it slightly decreases the amount of backlogged demands. On the other hand, Aff-approx determines order quantities more conservatively than PA. Under high penalty cost, Aff-approx satisfies nearly all customer demands by ordering an excessive amount of input. This causes a significantly larger holding cost and leads to less profit than PA. Finally, PA controls both leftover input and backlogged demands. As the penalty cost increases, PA not only reduces backlogged demands (same as Aff-approx), but also maintains much lower

input levels than Aff-approx. This leads to a higher profit than the other policies.

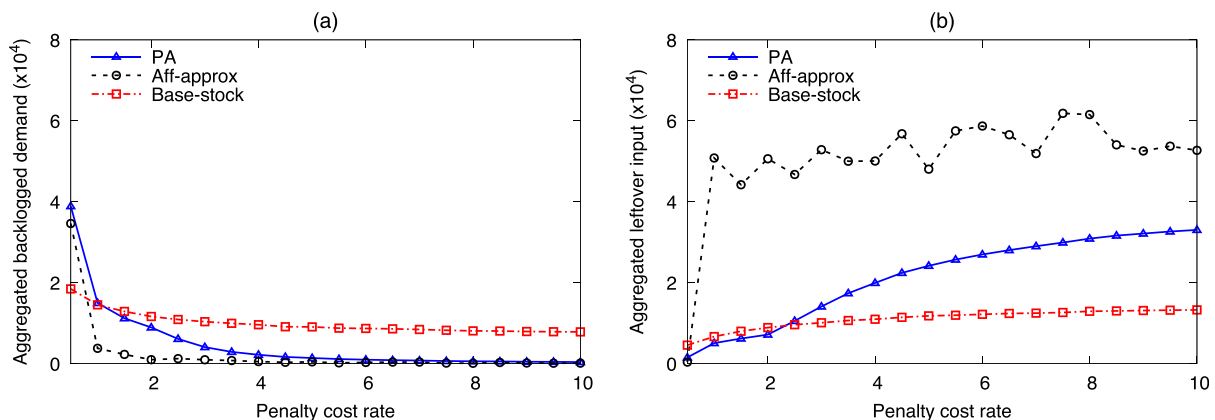
Summary. In summary, the case study with pharmacy retailer's data demonstrates sizable increase in performance for the periodic-affine policies. In particular, in the presence of high penalty cost, PA improves the lower-quantile performance by more than 10% than the base-stock policies. It also allows decision makers to identify the optimal policy based on their respective cost and demand structures.

6.2. Generalizability and Extentability: Modeling Correlation and Solving Multistation Newsvendor Problems


To adequately discuss the performance of PA under a multistation setting, we take demand correlation information into account. We model correlations in demand using the *correlated uncertainty sets* presented in (3) by using a factor model approach. For the numerical analysis, we consider two products over 15 time periods with subperiods of length 5.

The benefits of modeling correlation become apparent when comparing the following two policies: single-station

Figure 3. (Color online) Impact of Penalty Cost on Backlogged Demands (a) and Leftover Resources (b) on Random Samples



Note. Aff-approx, affine approximation; PA, periodic-affine.

Figure 4. (Color online) Relative Performance for Different Correlations ρ 

Note. Insets are the corresponding uncertainty sets.

PA (PA-single) for each product using marginal mean and variance and multistation PA (PA-multi) using the correlated uncertainty set. For comparison, we compute the *relative performance* (RP) of PA-multi over PA-single as

$$RP = \frac{\text{profit of PA-multi} - \text{profit of PA-single}}{\text{profit of PA-single}}.$$

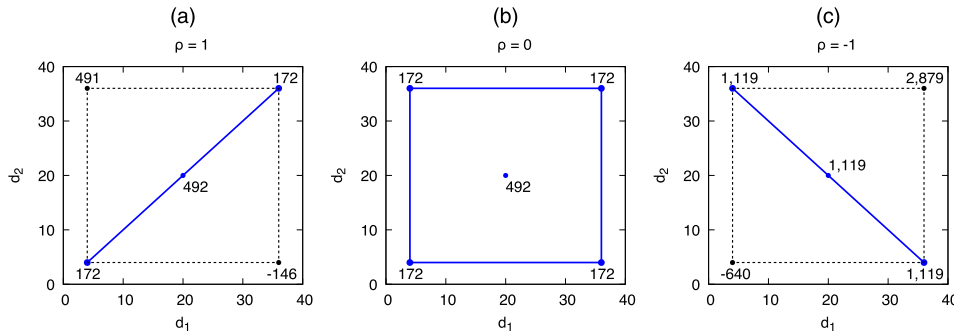
After the two policies are implemented, we generate random demand with nominal mean and covariance and evaluate the relative performance for each sample. For each correlation ρ , correlated uncertainty sets are defined by substituting

$$A = \begin{bmatrix} \rho & \sqrt{1-\rho^2} \\ 0 & 1 \end{bmatrix},$$

into (3), and the same matrix is used to generate multivariate normal random demands. Figure 4 displays

this relative performance for different correlation coefficient ρ . We observe that the median RP for every ρ is positive. However, the behavior differs for positively or negatively correlated products. Although for highly correlated products, the RP slightly decreases with growing ρ , significant improvements are made for negatively correlated products. In fact, RP increases by more than 17%, when the products have a strong negative correlation ($\rho = -0.9$).

This observation can be interpreted by the structure of the uncertainty sets, as shown in the insets of Figure 4. Positively correlated products lead to sets that allow both uncertain demands d_1 and d_2 to be concurrently at their maximum or minimum value. As ρ decreases, the area of the polyhedron shrinks; however, the extreme points are unaffected. However, when ρ becomes negative—that is, the products are negatively correlated—if one of the uncertain demands can take its maximum value, the

Figure 5. (Color online) An Illustrative Example for Correlation and Multistation Modeling

Notes. Two products are perfectly correlated (a), uncorrelated (b), and perfectly anticorrelated (c). Uncertainty sets are defined as within the contours along with the profit of optimal policies for extreme and nominal points.

Table 2. Performance of Policies for Two Correlated Product Groups for Different Penalty Cost Rates

Penalty cost rate	Policy	Worst	5% quantile	25% quantile	Median	Historical
0.2	Single-station affine approximation	6.20	9.09	9.65	10.03	10.60
	Multistation affine approximation	7.03	8.65	9.72	10.31	9.51
	Single-station periodic-affine	6.32	8.59	8.97	9.21	9.66
	Multistation periodic-affine	7.17	9.24	10.24	10.71	11.12
	Sample average approximation	N/A	9.12	10.04	10.51	10.93
	Myopic	N/A	8.50	9.97	10.74	10.81
1.0	Single-station affine approximation	4.07	6.44	7.59	8.15	8.39
	Multistation affine approximation	5.29	8.22	8.83	9.34	10.02
	Single-station periodic-affine	5.69	7.76	8.07	8.31	8.97
	Multistation periodic-affine	6.64	8.65	9.30	9.70	10.54
	Sample average approximation	N/A	7.49	7.82	9.35	9.13
	Myopic	N/A	6.70	8.10	9.21	8.98
5.0	Single-station affine approximation	2.63	5.19	6.82	7.90	7.21
	Multistation affine approximation	4.20	7.44	8.44	9.01	9.37
	Single-station periodic-affine	3.44	7.57	8.33	8.77	9.75
	Multistation periodic-affine	4.95	7.82	8.73	9.43	9.25
	Sample average approximation	N/A	6.41	7.23	8.61	8.92
	Myopic	N/A	5.04	6.61	7.76	7.03
10.0	Single-station affine approximation	1.69	5.06	6.67	7.64	6.88
	Multistation affine approximation	3.46	7.07	8.17	8.85	9.88
	Single-station periodic-affine	2.11	7.26	8.24	8.81	9.01
	Multistation periodic-affine	3.98	7.04	8.25	8.98	8.87
	Sample average approximation	N/A	6.01	7.12	7.83	7.61
	Myopic	N/A	4.61	6.27	7.40	6.42

Notes. All values are multiplied by 10^6 . Percentiles and medians are calculated from 1,000 samples. The last column is from the historical sales data. N/A, not applicable.

other is forced to its lowest, and vice versa. This effect forces an increase in RP as $\rho \rightarrow -1$. The extreme cases are illustrated in the example of Figure 5. For perfectly correlated products, Figure 5(a) shows that even though the uncertainty set is dramatically shrunk, the worst-case profit cannot be improved over the uncorrelated demand setting (Figure 5(b)), because the worst-case is often captured when both demands are high or low. However, for negatively correlated demands, the uncertainty set does not contain this region (high/high or low/low), allowing for substantial improvement in worst-case profit, as shown in Figure 5(c).

Case Study of a Pharmacy Retailer—Revisited. Here, we account for correlation amongst the product groups and compare the following five policies in Table 2: PA-multi, PA-single, Aff-approx with correlated uncertainty set (Aff-approx-multi), Aff-approx for each group (Aff-approx-single), SAA base-stock, and Myopic policies. For lower quantiles, we observe that the base-stock policy performs poorly compared with the single-station models. Both PA-single and PA-multi yield significantly greater profit in lower quantiles than the base-stock policy. The multistation framework and correlated demand uncertainty sets offer better performance than the single-station framework. In particular, PA-multi achieves at least 7% more profit than

PA-single for moderate choice of the penalty cost. However, for extremely high penalty cost rates (e.g., 10), PA-single performs slightly better than PA-multi for lower quantiles, even though the worst-case objective value is lower. This is due to samples that are generated outside of the correlated uncertainty sets. Note that SAA and the Myopic policies do not take correlations into account and hence underperform.

Summary. In summary, Table 2 demonstrates that capturing correlation using the periodic-affine policies in multistation setting outperforms the base-stock policies for moderate penalties. Because the demand of the studied pharmaceutical product groups is positively correlated, variability of the base-stock policy increases, causing a sizable degradation of the profit when compared with periodic-affine policies.

6.3. Computational Tractability

In order to focus on computational performance in a sterile environment, we consider the following simulation environment. We simulate three cases with duration $T \in \{10, 15, 20\}$ with a subperiod consisting of five time periods. We randomly generate 100 instances of single-station newsvendor problems for each T . Nominal means are generated by autoregression processes AR(1) and nominal coefficient of variations are

Table 3. Average Performance of Three Policies on Randomly Generated Instances

Policy	Affine			Periodic-affine			Affine approximation		
Time periods	10	15	20	10	15	20	10	15	20
Computation time (seconds)	3.8	88.5	1388.3	0.18	0.24	0.29	0.02	0.09	0.22
Worst objective value	11,035	16,833	22,440	11,033	16,831	22,438	9,523	14,601	19,479
Difference to affine (%)	—	—	—	−0.015	−0.011	−0.006	−13.39	−13.03	−12.96

chosen uniformly in $(0.3, 0.5)$. Unless modified, cost parameters are $c_S = 20$, $r = 120$, and $c_H = c_P = 20$, and all variability parameters are set similarly as in Section 6.1. We then compare the PA policies with affine and Aff-approx policies.

Although we evaluate the worst-case performances of PA using (12) with Proposition 4, the other policies are calculated using (6–8) to overcome tractability issues. Note that for any nonanticipative policies, (6–8) is an upper bound for the DP (12). Hence, the optimality gap between PA and the other policies, evaluated in (12), is closer than the gaps presented in Table 3 and Figure 6.


For the three policies, Table 3 displays the computation times and worst objective values on the same uncertainty set. We observe that the computation of PA is significantly faster than affine policies, because PA is tractable. However, the worst objective values of PA are very close (within 0.1%) to affine policies, whereas Aff-approx consistently deviates by $\geq 10\%$ from the others. Indeed, only 13 out of the 300 artificial instances have greater worst objective values in affine policies. Moreover, there is only one instance in which PA loses more than 1% of optimality. This implies that Assumption 1, which is a sufficient optimality condition of PA, holds for fairly general settings. This means that PA is as competitive as affine policies in worst-case values.

Comparison of PA and Aff-Approx Policies on Synthetic Data. We next compare the performance of PA and Aff-approx policies for a spectrum of parameters.

Figure 6 shows that the gap between the two policies are different for holding and for penalty cost. Figure 6(a) shows that the gap between PA and Aff-approx decreases as holding cost increases. However, PA protects the worst-case profit significantly better than Aff-approx as penalty cost increases, shown in Figure 6(b). These results suggest that depending of the holding and penalty costs, PA orders the proper amount of input to meet demands—that is, both on-hand input and unsatisfied demands are well-controlled. Therefore, its worst-case performance is far less affected by larger cost parameters than Aff-approx. On the other hand, Figure 6(b) indicates that Aff-approx may not manage the left-over resources and backlogged demands as well as PA, which yields poor performance for higher penalty cost.

The significant relative decrease of worst objective values between the two policies of up to 35%, as shown in Figure 6 and also observed in Table 3, implies that the degree of suboptimality of Aff-approx may render it inferior to PA. Although theoretical bounds of affine approximation have been proposed (Bertsimas and Goyal 2012, Bertsimas and Bidkhori 2015), these results indicate that the affine approximation cannot be successfully applied to general settings, despite its computational advantage of tractability.

Summary. In summary, the experiment on artificial data reveals the optimality and tractability of periodic-affine policies, whereas the other methods did not

Figure 6. (Color online) Impact of Holding Cost (a) and Penalty Cost (b) on Worst-Case Performance

Note. PA, periodic-affine.

achieve both properties. In addition, we observe that periodic-affine policies perform substantially better for larger T than Aff-approx in terms of protecting its worst-case performance at high penalty cost. This implies that PA is particularly useful when a decision maker faces a massive penalty for unsatisfied demands, such as in the case of the pharmacy retailer in India.

7. Conclusions

In this paper, we consider the problem of optimal control in multiperiod and multistage newsvendor networks. To this end, we introduce a new class of adaptive policies called *periodic-affine policies*. These policies are data-driven and incorporate the correlation amongst products, which is an instrumental feature of real-world settings. These policies also remain robust to parameter misspecification. For this, we model the uncertain demand via sets, which can incorporate correlations and can be generalized to multiproduct settings and extended to multiperiod problems. This approach offers a natural framework to study current competing policies of base-stock, affine, and approximative approaches with respect to their profit, sensitivity to parameters and assumptions, and computational scalability. We show that the periodic-affine policies are sustainable—that is, time consistent—because they warrant optimality both within subperiods and over the entire planning horizon.

We presented empirical evidence of tractability and robustness, which makes our approach well suited for real-world applications. We demonstrate the advantages of our approach by considering the problem of one of India's largest pharmacy retailers using their sales data. We show that the periodic-affine policies are capable to increase the profits by up to 17% over base-stock policies. This study reveals that capturing the demand correlation can sizably affect the performance. Furthermore, we offer a phase diagram for managers to select the optimal policy based on their cost and demand structures.

In the future, we intend to incorporate time-dependent uncertainty sets (Nohadani and Roy 2017) to more accurately model seasonal demand. This step forward will lend itself well to incorporate returns—that is, feedback from satisfied demand that can guide the next period's decisions.

Acknowledgments

The authors thank Jan A. van Mieghem for insightful comments and two anonymous referees for their thorough and suggestive reviews.

References

- Arrow KJ, Harris T, Marschak J (1951) Optimal inventory policy. *Econometrica* 19(3):250–272.
- Baboli A, Fondrevelle J, Tavakkoli-Moghaddam R, Mehrabi A (2011) A replenishment policy based on joint optimization in a

downstream pharmaceutical supply chain: Centralized vs. decentralized replenishment. *Internat. J. Adv. Manufacturing Tech.* 57(1–4):367–378.

- Bandi C, Bertsimas D (2012) Tractable stochastic analysis in high dimensions via robust optimization. *Math. Programming* 134(1):23–70.
- Ben-Tal A, Golany B, Nemirovski A, Vial J (2005) Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing Service Oper. Management* 7(3):248–271.
- Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Programming* 99(2):351–376.
- Ben-Tal A, Den Hertog D, De Waegenaere A, Melenberg B, Rennen G (2013) Robust solutions of optimization problems affected by uncertain probabilities. *Management Sci.* 59(2):341–357.
- Bertsimas D, Biddkhor H (2015) On the performance of affine policies for two-stage adaptive optimization: A geometric perspective. *Math. Programming* 153(2):577–594.
- Bertsimas D, Goyal V (2012) On the power and limitations of affine policies in two-stage adaptive optimization. *Math. Programming* 134(2):491–531.
- Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52(1):35–53.
- Bertsimas D, Thiele A (2006) Robust and data-driven optimization: Modern decision making under uncertainty. *Models, Methods, and Applications for Innovative Decision Making*, TutORials in Operations Research (INFORMS, Catonsville, MD), 95–122.
- Bertsimas D, Gupta V, Kallus N (2018) Robust sample average approximation. *Math. Programming* 171(1–2):217–282.
- Bertsimas D, Iancu D, Parillo P (2010) Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.* 35(2):363–394.
- Bienstock D, Özbay N (2008) Computing robust base-stock levels. *Discrete Optim.* 5(2):389–414.
- Clark AJ, Scarf H (1960) Optimal policies for a multi-echelon inventory problem. *Management Sci.* 6(4):475–490.
- Crawford GS, Shum M (2005) Uncertainty and learning in pharmaceutical demand. *Econometrica* 73(4):1137–1173.
- Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* 58(3):595–612.
- Federgruen A, Zipkin P (1986) An inventory model with limited production capacity and uncertain demands. I. The average-cost criterion. *Math. Oper. Res.* 11(2):193–207.
- Fu M, Healy K (1997) Techniques for simulation optimization: An experimental study on an (s, S) inventory system. *IIE Trans.* 29(3):191–199.
- Fu M, Glover F, April J (2005) Simulation optimization: A review, new developments, and applications. Kuhl ME, Steiger NM, Armstrong FB, Joines JA, eds. *Proc. 2005 Winter Simulation Conf.* (Association for Computing Machinery, New York), 83–95.
- Fu MC (1994) Sample path derivatives for (s, S) inventory systems. *Oper. Res.* 42(2):351–364.
- Gallego G, Moon I (1993) The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.* 44(8):825–834.
- Glasserman P, Tayur S (1995) Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Sci.* 41(2):263–281.
- Graves S, Willems S (2000) Optimizing strategic safety-stock placement in supply chains. *Manufacturing Service Oper. Management* 2(1):68–83.
- Guerrero WJ, Yeung T, Guéret C (2013) Joint-optimization of inventory policies on a multi-product multi-echelon pharmaceutical system with batching and ordering constraints. *Eur. J. Oper. Res.* 231(1):98–108.
- Huh W, Janakiraman G (2008) A sample-path approach to the optimality of echelon order-up-to policies in serial inventory systems. *Oper. Res. Lett.* 36(5):547–550.

- Kapuscinski R, Tayur S (1999) Optimal policies and simulation-based optimization for capacitated production inventory systems. *Quantitative Models for Supply Chain Management* (Springer, New York), 7–40.
- Karlin S (1960) Dynamic inventory policy with varying stochastic demands. *Management Sci.* 6(3):231–258.
- Kasugai H, Kasegai T (1961) Note on minimax regret ordering policy-static and dynamic solutions and a comparison to maximin policy. *J. Oper. Res. Soc. Japan* 3(4):155–169.
- Kerrigan E, Maciejowski J (2004) Properties of a new parametrization for the control of constrained systems with disturbances. *Proc. 2004 Amer. Control Conf.*, vol. 5 (Institute of Electrical and Electronic Engineers, Piscataway, NJ), 4669–4674.
- Kuhn D, Wiesmann W, Georghiou A (2011) Primal and dual linear decision rules in stochastic and robust optimization. *Math. Programming* 130(1):177–209.
- Langenhoff L, Zijm W (1990) An analytical theory of multi-echelon production/distribution systems. *Statist. Neerlandica* 44(3):149–174.
- Löfberg J (2003) Approximations of closed-loop minimax mpc. *Proc. 42nd IEEE Conf. Decision Control*, vol. 2 (Institute of Electrical and Electronic Engineers, Piscataway, NJ), 1438–1442.
- Maechler M, Rousseeuw P, Struyf A, Hubert M, Hornik K (2016) cluster: Cluster analysis basics and extensions. R package version 2.0.5.
- Mohajerin Esfahani P, Kuhn D (2018) Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Math. Programming* 171(1/2):115–166.
- Morton TE (1978) The nonstationary infinite horizon inventory problem. *Management Sci.* 24(14):1474–1482.
- Muharremoglu A, Tsitsiklis JN (2008) A single-unit decomposition approach to multiechelon inventory systems. *Oper. Res.* 56(5):1089–1103.
- Nohadani O, Sharma K (2018) Optimization under decision-dependent uncertainty. *SIAM J. Optim.* 28(2):1773–1795.
- Nohadani O, Roy A (2017) Robust optimization with time-dependent uncertainty in radiation therapy. *IISE Trans. Healthcare Systems Engrg.* 7(2):81–92.
- Rikun A (2011) Applications of robust optimization to queueing and inventory systems. Unpublished doctoral thesis, Massachusetts Institute of Technology, Cambridge.
- Rosling K (1989) Optimal inventory policies for assembly systems under random demands. *Oper. Res.* 37(4):565–579.
- Scarf H (1958) A min-max solution of an inventory problems. Arrow KJ, Karlin S, Scarf H, eds. *Studies in the Mathematical Theory of Inventory and Production* (Stanford University Press, Stanford, CA).
- Scarf H (1960) The optimality of (s, S) policies in the dynamic inventory problem. Arrow KJ, Karlin S, Suppes P, eds. *Mathematical Methods in the Social Sciences* (Stanford University Press, Stanford, CA).
- Sethi S, Cheng F (1997) Optimality of (s, S) policies in inventory models with Markovian demand. *Oper. Res.* 45(6):931–939.
- Uthayakumar R, Priyan S (2013) Pharmaceutical supply chain and inventory management strategies: Optimization for a pharmaceutical company and a hospital. *Oper. Res. Health Care* 2(3):52–64.
- van Mieghem JA, Rudi N (2002) Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing Service Oper. Management* 4(4):313–335.
- van Parys BP, Kuhn D, Goulart PJ, Morari M (2016) Distributionally robust control of constrained stochastic systems. *IEEE Trans. Automatic Control* 61(2):430–442.
- Webman E (2012) How much is it worth? Accessed December 19, 2018, <http://www.ncpa.co/pdf/APNOV12-HowMuchWorth.pdf>.
- Webman E (2016) Value added. Accessed December 19, 2018, <http://www.ncpa.co/pdf/webman-pharmacy-valuation-article.pdf>.