Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Multiperiod Stock Allocation via Robust Optimization

Peter L. Jackson, John A. Muckstadt, Yuexing Li

informs
http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Vol. 65, No. 2, February 2019, pp. 794–818
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Multiperiod Stock Allocation via Robust Optimization

Peter L. Jackson,[a] John A. Muckstadt,[a] Yuexing Li[a]

[a] School of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14853

Contact: pj16@cornell.edu, http://orcid.org/0000-0003-4031-3952 (PLJ); jmuckstadt@cornell.edu (JAM); yl2324@cornell.edu (YL)

Received: October 28, 2016
Revised: July 22, 2017
Accepted: September 21, 2017
Published Online in Articles in Advance: May 21, 2018

https://doi.org/10.1287/mnsc.2017.2969

Abstract. We consider a one-warehouse, N-retailer, multiperiod, stock allocation problem in which holding costs are identical at each location and no stock is received from outside suppliers for the duration of the planning horizon. No shipments are allowed between retailers. The only motive for holding inventory at the central warehouse for allocation in future periods is the so-called risk pooling motive. We apply robust optimization to this problem extending the inventory policy to allow for an adaptive, nonanticipatory shipment policy. We consider two alternatives for the uncertainty set, one in which risk pooling is implicit and another for which risk pooling is explicit. The explicit risk pooling uncertainty set grows by no more than the square of the number of retailers. The general problem can be solved using Benders' decomposition. A special case gives rise to closed-form solutions for both uncertainty set alternatives. The explicit risk pooling uncertainty set leads to a square root law in which the optimal stock to reserve at the central warehouse grows with the square root of the number of retailers. The experimental results confirm the value of the robust optimization approach and provide managerial insights into the operation of such systems.

History: Accepted by Yinyu Ye, optimization.

Keywords: multiechelon inventory optimization • robust optimization • inventory risk pooling

1. Introduction

In this paper, we focus on periodic-review, multiechelon inventory allocation problems. Surveys of this literature can be found in Axsäter (2003) and Doğru et al. (2009). In their seminal paper, Clark and Scarf (1960) use a dynamic programming approximation to solve a two-echelon central warehouse and multiple retailer distribution system problem under periodic review. Central to their approach is the so-called *balance assumption* under which it is never desirable to transfer inventory from one retailer to another. Jönsson and Silver (1987) make the observation that for items having relatively low coefficients of variation in their daily demand processes, inventory imbalance will occur among retailers, if at all, only at the end of a cycle, defined as the period between deliveries to the central warehouse from a supplier. This has led to many formulations of the problem with only two periods: a lengthy first period during which imbalance at the retailer locations can be ignored followed by a short period during which imbalance may be an issue at these locations. If imbalance is an issue, then it is necessary to include constraints requiring shipments to retailers to be nonnegative. The inclusion of these constraints leads to a dynamic programming problem formulation that is computationally intractable for realistic problem sizes. The balance assumption, on the other hand, permits the problem to be decomposed into single-location problems. Numerous papers have exploited the balance assumption, beginning with Clark and Scarf (1960), Eppen and Schrage (1981), Federgruen and Zipkin (1984a), and Federgruen and Zipkin (1984b). An analytical investigation of the imbalance assumption, including the impact of demand correlation, can be found in Zipkin (1984). A more recent review and application of the balance assumption can be found in Gallego et al. (2007).

Doğru et al. (2009) survey the papers in this area to highlight the widespread use of the balance assumption. To test the applicability of the balance assumption they compute a gap between an upper and lower bound on the optimal expected cost. The authors find that the gap is small for identical retailers if the coefficient of variation of demand is small or moderate (less than 0.5) or if the incremental holding cost at the retailers is high. The gap is small for nonidentical retailers in only a few of their parameters settings, even with moderate coefficients of variation. This suggests that either the lower bound or the heuristic upper bound policy, or both, are mediocre in settings characterized by high coefficients of variation or nonidentical retailers.

A number of authors have attempted to incorporate the nonnegative shipment constraints to improve the optimization results. Jackson and Muckstadt (1989) restrict attention to a two-period allocation cycle and show that when nonnegative shipments are enforced, the maximum post-allocation stockout probability across retailers converges to a constant as the number of retailers increases. Using this constant, they develop an approximate cost function that is separable by retailer and can be optimized to set their target inventory levels.

Axsäter et al. (2002) introduce a two-step allocation heuristic that is also based on a two-period allocation cycle. Their method uses convolutions of a three-point probability function to represent the joint probability distribution of period 1 demands. With this approximation, they find the optimal total stock to keep in reserve at the central warehouse.

Using a more general multiperiod approach, Kunnumkal and Topaloglu (2008) develop a lower bound on the cost of the optimal policy by associating Lagrange multipliers with the nonnegative shipment constraints. The resulting dynamic programming problem is separable into single-location dynamic programs that are solved easily. Subgradient optimization is then used to find a vector of Lagrange multipliers that maximizes the lower bound. Kunnumkal and Topaloglu (2011) explore approximations to speed the computations.

In this paper, we use a robust optimization approach to study the problem. Robust optimization has been proposed as a tractable optimization approach for stochastic planning problems that are too large to be solved by dynamic programming. There has been great interest in reformulating inventory planning problems as robust optimization problems beginning with Bertsimas and Thiele (2006). The problem with applying their original approach directly to stock allocation problems is that theirs is a static planning problem. To capture the role of risk pooling, we require an adaptive policy: one in which shipments will depend on the evolution of demand across time and retailers. Robust formulations of adaptive stochastic problems are intractable in general. However, if the adaptive policies can be expressed as affine functions of the uncertainty variables, then the problem becomes tractable (Ben-Tal et al. 2004). Mamani et al. (2017) provide a recent review of this literature while providing results on a static approach of their own, implemented in a rolling-horizon fashion.

We propose an adaptive policy but do not rely on an affine adaptive policy. Our approach follows Bredström et al. (2013) in modeling planning problems with right-hand-side uncertainty, specifically, demand uncertainty. We separate the decisions between inventory policy variables and actual stock allocation decisions. The former are chosen before demands are realized and the latter are consequences of the policy variables and the actual demands. As in Bredström et al., we use Benders' decomposition to solve the overall problem and face a bilinear program for the subproblem. In our case, the simplicity of our formulation allows us to recast the bilinear program as a mixed-integer linear program.

We also assume that demand distributions are characterized by their means and covariances and that demands are independent across time. We consider two alternative uncertainty sets in our formulation to model normalized demand. Both sets include box constraints to bound the uncertainty variables, but the phenomenon of risk pooling is captured in two different ways. In the first alternative, we apply an uncertainty budget to the weighted sum of normalized demands in each period. This is rather standard practice, and we refer to it as the implicit risk pooling alternative. The second alternative considers partial sums of uncertainty variables with explicit square root factors involving the cardinality of the sums in the bounds. We refer to this as the explicit risk pooling alternative. This latter approach appears in a number of robust optimization contexts such as Bertsimas and Sim (2004) and Bandi et al. (2015). What is unusual in our formulation is that we consider all possible partial sums. The explicit risk pooling approach leads to a novel insight in the case of symmetric demands. It also leads to generally good performance, empirically, in capturing the potential risk pooling benefit of centralized inventories. The implicit risk pooling approach also leads to good performance provided the uncertainty budget is chosen with care.

Our contribution to robust inventory modeling is thus to focus attention on the stock allocation problem, as distinct from the stock ordering problem that has been the focus of much of the literature, and to explore the use of nonaffine adaptive policies with alternative uncertainty sets.

In the course of evaluating the performance of the robust allocation approach, we consider a variety of situations in which risk pooling, and the balance assumption, may or may not play a role in system performance. This study reinforces some conclusions from earlier studies, but also points to new managerial insights. For example, Doğru et al. (2009) show the balance assumption to be violated to a significant extent when demand among retailers is unbalanced, even for moderate coefficients of variation. Our results point to a different conclusion, and this can be traced to our contention that coefficients of variation are inversely correlated with demand rates, as we have observed in several practical instances. The strength of the risk pooling phenomenon in this setting has been recently challenged in Bimpikis and Markakis (2016) when demand has a heavy-tailed distribution. Our numerical results are conducted using log-normally distributed demand and continue to show the importance of the risk pooling phenomenon in stock allocation.

2. Multiperiod Stock Allocation

In this section, we consider the general multiperiod stock allocation problem. There are several periods of uncertain demand at the retailers. Each location begins the cycle with an initial inventory. The central warehouse has the opportunity to make costless allocations to the retailers at the beginning of each period, to the extent that it has held stock in reserve for these periods. The goal is to minimize backorder costs over the cycle.

2.1. Expected Value Optimization Formulation

Let T denote the *stock allocation horizon*, the number of periods considered in the stock allocation problem. Period $T+1$ is assumed to be a period in which the system is replenished with sufficient inventory from the supplier to eliminate all backorders at all retailers. Let $t = 1, 2, \ldots, T$ index the periods. The periods are not necessarily of equal length: As discussed earlier, it is often desirable to allow the first period to cover a longer span of time than the other periods. The lead time to ship from the central warehouse to the retailers is assumed negligible.

Let the set of retailers be denoted by $\mathcal{N}$ and indexed by $i \in \mathcal{N}$. Let N denote the number of retailers: $N = |\mathcal{N}|$. For a generic period t, let v_i denote the beginning net inventory at retailer i, $i \in \mathcal{N}$, and let v_0, $v_0 \geq 0$, denote the stock held in reserve at the central warehouse. Let $v = (v_0, v_1, \ldots, v_N)$. Let x_i, $x_i \geq 0$, denote the allocation of stock from the central warehouse to retailer i at the beginning of the period. The total allocation must not exceed the reserve: $\sum_{i \in \mathcal{N}} x_i \leq v_0$.

Let $\tilde{d}_{it}$ denote the random demand occurring at retailer i during period t, and let $\tilde{d}_t = (\tilde{d}_{it})_{i \in \mathcal{N}}$ denote the vector of demands in period t. The stochastic process $\tilde{d} = \{\tilde{d}_t \colon t = 1, 2, \ldots, T\}$ is assumed to be independent from period-to-period.

We assume that the cost of holding inventory is the same at all locations and therefore irrelevant to the allocation problem. Our focus is on allocation policies that have risk pooling as the predominate consideration. Let w_{it} denote the per-unit backorder cost at retailer i in period t. Let $f_t(v)$ denote the minimal expected total backorders over periods $t, t+1, \ldots, T$, given that the system begins period t in state v. This function can be shown to satisfy the dynamic programming recursion given by the following: for $t = 1, 2, \ldots, T-1$,

$$f_T(v) = \min_{\substack{x_i \geq 0,\, i \in \mathcal{N} \\ \sum_{i \in \mathcal{N}} x_i \leq v_0}} \sum_{i \in \mathcal{N}} \mathrm{E}[w_{iT}(\tilde{d}_{iT} - v_i - x_i)^+];$$

$$f_t(v) = \min_{\substack{x_i \geq 0,\, i \in \mathcal{N} \\ \sum_{i \in \mathcal{N}} x_i \leq v_0}} \left\{ \begin{array}{l} \displaystyle\sum_{i \in \mathcal{N}} \mathrm{E}[w_{it}(\tilde{d}_{it} - v_i - x_i)^+] \\ \displaystyle + \mathrm{E}\left[f_{t+1}\left(\left(v_0 - \sum_{i \in \mathcal{N}} x_i, (v_i + x_i - \tilde{d}_{it})_{i \in \mathcal{N}} \right) \right) \right] \end{array} \right\}.$$

This formulation suffers from the curse of dimensionality: the state space over which each f_t must be evaluated is exponential in the number of retailers. As mentioned, past approaches in the literature have focused on different approximation techniques to reduce the computational burden.

2.2. Robust Optimization

2.2.1. The Risk Pooling Uncertainty Set. As motivation in designing an uncertainty set, we take the vector of demands, $\tilde{d}_t$, in any period t to be normally distributed with mean vector μ_t and variance–covariance matrix Σ_t. Since Σ_t is positive semidefinite, there exists a factorization, $\Sigma_t = C_t C_t^T$. Then, an equivalent representation of demand is given by

$$\tilde{d}_t = \mu_t + C_t \tilde{\varepsilon}_t, \tag{1}$$

where $\tilde{\varepsilon}_t$ is a vector of independent, $N(0,1)$ random variables.

Let c_{ijt} denote the entries in the matrix C_t. We model demand in period t for retailer i as

$$d_{it} = \mu_{it} + \sum_{j \in \mathcal{N}} c_{ijt} \varepsilon_{jt}, \tag{2}$$

where an adversary is restricted to choosing variables ε_{it} from some uncertainty set $U(\delta)$, to be defined. We drop the tilde (˜) decoration from d_{it} and ε_{it} because they are no longer interpreted as random variables. For much of the paper, we consider only cases in which the correlation entries, c_{ijt}, are nonnegative for all i, j, and t.

In general, we are concerned with restricting the adversary from choosing unrealistically large partial sums of demand, as this is what leads to inventory imbalance. We consider two forms of the risk pooling uncertainty set. The first form places upper bounds on the individual normalized demands and a single upper bound constraint per period on the aggregate scaled demand:

$$U(\delta) = \left\{ \varepsilon \colon \begin{array}{c} \varepsilon_{it} \leq \delta_0, \ \forall (i,t) \\ \displaystyle\sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} c_{ijt} \varepsilon_{jt} \leq \delta_1, \ t = 1, \ldots, T \end{array} \right\}, \tag{3}$$

for $\delta = \{\delta_0, \delta_1\}$. Observe that we do not place lower bounds on the normalized demands or the aggregate scaled demand, because we assume that the correlation entries are nonnegative. Nevertheless, we explore negative correlation cases in the numerical study, with suitable adjustment to the uncertainty set formulation. This form of the uncertainty set does not capture the risk pooling phenomenon explicitly. Care must be taken in choosing the parameter δ_1 to represent risk pooling adequately. We refer to this form as the *implicit risk pooling uncertainty set*.

Another approach is to consider partial sums of the normalized demand variables and impose explicit risk pooling constraints along the lines proposed in Bandi et al. (2015) and others. We define the multiperiod *explicit risk pooling uncertainty set* $U(\delta)$ as follows:

$$U(\delta) = \left\{ \varepsilon \colon \begin{array}{c} \varepsilon_{it} \leq \delta, \ \forall (i,t) \\ \displaystyle\sum_{i \in I} \sum_{t=1}^{t'} \varepsilon_{it} \leq \sqrt{|I| t'}\,\delta, \ \forall I \subseteq \mathcal{N}, \ t' = 1, \ldots, T \end{array} \right\}. \tag{4}$$

We treat δ as a scalar and use the same value in each constraint. This is for simplicity of presentation.

It could be replaced easily by a vector (one entry per constraint) that would provide greater management control over the resulting set. What is unusual here is that there are constraints limiting the partial sums of normalized demands for all possible subsets of retailers. This is motivated by the uncertainty over which retailers will participate in a stock allocation from the central warehouse.

There are two objections to this formulation of the uncertainty set. The first is that as the number of retailers increases, the probability that at least one of these constraints is violated increases. Hence, the set becomes less realistic. Secondly, the computational complexity will also grow in the number of retailers. In the interest of controlling both realism and computation time, we introduce another parameter, $\bar{n}$, to limit the cardinality of the subsets considered, $\bar{n} \in \{1, 2, \ldots, N\}$. We refer to $\bar{n}$ as the *depth* of the explicit risk pooling uncertainty set.

$$U(\delta, \bar{n}) = \left\{ \varepsilon: \begin{array}{c} \varepsilon_{it} \leq \delta, \forall (i,t) \\ \sum_{i \in I} \sum_{t=1}^{t'} \varepsilon_{it} \leq \sqrt{|I| t'} \delta, \\ \forall I \subseteq \mathcal{N}, |I| \leq \bar{n}, t' = 1, \ldots, T \end{array} \right\} \quad (5)$$

The number of sets satisfying $I \subseteq \mathcal{N}$ is 2^N, an exponential number of constraints. However, the problem of bounding the k largest entries of an n-vector is known to have an equivalent formulation with fewer constraints (Ben-Tal and Nemirovski 2001, Zakeri et al. 2014).

Lemma 1. *The set $Z = \{z = (z_i)_{i \in \mathcal{N}}: \sum_{i \in I} z_i \leq M_n, \forall I \subseteq \mathcal{N}, |I| = n\}$ for some M_n, $n \in \mathcal{N}$, is equivalent to the set Z' where*

$$Z' = \left\{ z = (z_i)_{i \in \mathcal{N}}: \begin{array}{c} \exists \alpha_n, \beta_n = (\beta_{ni})_{i \in \mathcal{N}} \text{ s.t.} \\ n\alpha_n + \sum_{i=1}^{|\mathcal{N}|} \beta_{ni} \leq M_n, \\ \alpha_n + \beta_{ni} \geq z_i, \, i \in \mathcal{N}, \\ \beta_{ni} \geq 0, \, i \in \mathcal{N}. \end{array} \right\}.$$

Proof. The proof of this lemma, and all subsequent results, can be found in Appendix B.

For a given t', z_i can play the role of $\sum_{t=1}^{t'} \varepsilon_{it}$, and M_n can play the role of $\sqrt{|I| t'} \delta$ in (5) where $|I| = n$. This yields an equivalent formulation of the explicit risk pooling uncertainty set, $U(\delta, \bar{n})$:

$$U(\delta, \bar{n}) = \left\{ \varepsilon = (\varepsilon_{it})_{i \in \mathcal{N}, t=1,\ldots,T}: \begin{array}{c} \varepsilon_{it} \leq \delta, \forall (i,t); \\ \text{for } \forall n \in \{1, 2, \ldots, \bar{n}\}, t = 1, \ldots, T: \\ \exists \alpha_{nt}, \beta_{nt} = (\beta_{nti})_{i \in \mathcal{N}} \text{ s.t.} \\ n\alpha_{nt} + \sum_{i \in \mathcal{N}} \beta_{nti} \leq \sqrt{nt}\delta, \\ \alpha_{nt} + \beta_{nti} \geq \sum_{t'=1}^{t} \varepsilon_{it'}, \forall i \in \mathcal{N}, \\ \beta_{nti} \geq 0, \forall i \in \mathcal{N}. \end{array} \right\}.$$

Proposition 1. *The number of constraints required to express the multiperiod uncertainty set is of order $TN\bar{n}$, where N is the number of retailers.*

2.2.2. Multiperiod Robust Optimization.

In this section, we formulate the stock allocation problem as a robust optimization problem. We assume that we are given an uncertainty set U of unspecified form (implicit or explicit risk pooling) and parameterization ($\{\delta_0, \delta_1\}$ or $\{\delta, \bar{n}\}$). We assume only that there exists a parameter δ such that $\varepsilon_{it} \leq \delta$ for all $i \in \mathcal{N}$ and $t = 1, \ldots, T$. The form of the uncertainty set becomes important in subsequent sections.

An Adaptive Shipment Policy. Suppose the initial net inventory at location i at the beginning of period 1 is given by v_i, $i \in \mathcal{N}$ and the initial reserve stock at the central warehouse at the beginning of period 1 is v_0. The variable y_{it} represents the target net inventory level at retailer i in period t just after allocation from the central warehouse. The variable x_{it} represents the corresponding allocation from the central warehouse. The symbols μ, σ, v, d, ε, y, and x with subscripts omitted represent vectors of parameters or variables ranging over the omitted subscripts.

In this approach, the target net inventory levels, y, are chosen without reference to a particular ε vector. The shipments in each period are allowed to be adaptive, but they must be nonanticipatory. In particular, we require

$$x_{it}(\varepsilon) = \left(y_{it} - v_i + \sum_{t'=1}^{t-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} - x_{it'}(\varepsilon) \right) \right)^+. \quad (6)$$

That is, in each period, the policy is to ship the minimum nonnegative quantity required to achieve the target net inventory level, subject to past demands, shipments, and initial net inventories.

Lemma 2. *Partial sums of required shipments satisfy the following relation:*

$$\sum_{t'=1}^{t} x_{it'}(\varepsilon) \geq y_{it} - v_i + \sum_{t'=1}^{t-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} \right), \quad (7)$$

for $t = 1, 2, \ldots, T$ with equality holding if $x_{it}(\varepsilon) > 0$.

Let $S(y, \varepsilon)$ denote the total shipments required over the stock allocation horizon following an adaptive, nonanticipatory policy for a given normalized demand vector, ε:

$$S(y, \varepsilon) = \sum_{i \in \mathcal{N}} \sum_{t=1}^{T} \left(y_{it} - v_i + \sum_{t'=1}^{t-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} - x_{it'} \right) \right)^+.$$

Proposition 2. *The total required shipment to support a target net inventory vector y given normalized demand* *vector ε can be determined using the following linear program*:

$$S(y,\varepsilon)=\min \sum_{i\in\mathcal{N}}\sum_{t=1}^{T} x_{it}$$

$$\text{subject to} \quad \sum_{t'=1}^{t} x_{it'}(\varepsilon) \geq y_{it} - v_i + \sum_{t'=1}^{t-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right),$$

$$x_{it} \geq 0,$$

for all $i \in \mathcal{N}$ and $t = 1, 2, \ldots, T$.

Corollary 1. *Alternatively,*

$$S(y,\varepsilon)=\max_{\pi} \sum_{i\in\mathcal{N}}\sum_{t=1}^{T} \pi_{it}\left(y_{it} - v_i + \sum_{t'=1}^{t-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right)\right) \tag{8}$$

$$\text{subject to} \quad \sum_{t'=t}^{T} \pi_{it'} \leq 1; \tag{9}$$

$$\pi_{it} \geq 0, \tag{10}$$

for all $i \in \mathcal{N}$ and $t = 1, 2, \ldots, T$.

The dual solution, π, has an interesting interpretation. For each retailer $i \in \mathcal{N}$, let $t_i(\varepsilon)$ denote the following maximizing period for retailer i:

$$t_i(\varepsilon) = \operatorname*{arg\,max}_{t=1,2,\ldots,T}\left(y_{it} - v_i + \sum_{t'=1}^{t-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right)\right).$$

In the case of ties, select the earliest period that achieves the optimum. The resulting maximum may not be positive. Consequently, we let $t_i^*(\varepsilon)$ equal zero in such cases:

$$t_i^*(\varepsilon) = \begin{cases} t_i(\varepsilon) & \text{if } y_{it_i(\varepsilon)} + \displaystyle\sum_{t'=1}^{t_i(\varepsilon)-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right) > v_i, \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 3. *The solution to the dual shipment problem,* (8)–(10), *is given by $\pi_{it}^*(\varepsilon) = 1_{\{t=t_i^*(\varepsilon)\}}$ for all $i \in \mathcal{N}$, $t = 1, 2, \ldots, T$.*

Let Π denote the set

$$\Pi = \left\{\pi\colon \sum_{t=1}^{T} \pi_{it} \leq 1; \pi_{it} \in \{0,1\}; i \in \mathcal{N}, t = 1, 2, \ldots, T\right\}.$$

Corollary 2. *There exists an optimal solution to* (8)–(10) *satisfying $\pi \in \Pi$.*

We can interpret the dual solution to the shipment problem as follows. Let $\tau_i(\varepsilon)$ denote the last period in which retailer i receives a shipment, or zero if retailer i receives no shipment:

$$\tau_i(\varepsilon) = \begin{cases} \operatorname*{arg\,max}_{t=1,2,\ldots,T} 1_{\{x_{it}(\varepsilon)>0\}} t & \text{if } \displaystyle\sum_{t=1}^{T} x_{it}(\varepsilon) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 4. *For a given target inventory vector y and normalized demand vector ε, the shipment vector $x(\varepsilon)$ satisfies $x_{it}(\varepsilon) = 0$, $\forall\, t > t_i^*(\varepsilon)$. Furthermore, $t_i^*(\varepsilon) = \tau_i(\varepsilon)$.*

Corollary 3. *Barring ties, the vector π optimizing* (8)–(10) *selects the last shipment period for each retailer. That is, $\pi_{it}^*(\varepsilon) = 1_{\{t=\tau_i(\varepsilon)\}}$.*

The Multiperiod Robust Stock Allocation Problem. We say that a target net inventory vector y is feasible if and only if its shipment requirements do not exceed the available reserve stock v_0 for any normalized demand vector in the uncertainty set. Let $S(y)$ be the worst-case shipping requirements: $S(y) = \max_{\varepsilon\in U} S(y,\varepsilon)$. Then, y is feasible *iff* $S(y) \leq v_0$.

Lemma 3. *$S(y)$ is nondecreasing in y.*

Recall that w_{it} denotes the weight to apply to backorders at retailer i in period t. Let B_t denote the maximum weighted backorders in period t across all retailers and all possible demand vectors. We assume that if $\varepsilon_{it} \in U$ there exists a parameter δ such that $\varepsilon_{it} \leq \delta$ for every $i \in \mathcal{N}$ and $t = 1, \ldots, T$. We also assume $c_{ijt} \geq 0$ for all i, j, and t. The problem is to minimize the sum of maximum weighted backorders across retailers in each period and across all possible normalized demands, as follows:

$$\min_{y,B} \sum_{t=1}^{T} B_t \tag{11}$$

$$\text{s.t.} \quad B_t \geq w_{it}(\bar{d}_{it}(\delta) - y_{it}), \quad \forall\, i \in \mathcal{N}, t = 1, \ldots, T; \tag{12}$$

$$B_t \geq 0, \quad t = 1, \ldots, T; \tag{13}$$

$$S(y) \leq v_0, \tag{14}$$

where $\bar{d}_{it}(\delta) = \mu_{it} + \sum_{j\in\mathcal{N}} |c_{ijt}|\delta$, $i \in \mathcal{N}$, $t = 1, \ldots, T$. The uncertainty set U has been removed from all but the subproblem, $S(y)$. Note that since we assume $c_{ijt} \geq 0$, the absolute value is unnecessary, but we do relax this assumption later and including it here adds no complexity.

Proposition 5. *There exists an optimal solution, (y^*, B^*), to the multiperiod stock allocation problem* (11)–(14) *satisfying*

$$y_{it}^* = \bar{d}_{it}(\delta) - w_{it}^{-1} B_t^*, \tag{15}$$

for all $i \in \mathcal{N}$, and $t = 1, 2, \ldots, T$.

By Corollary 2, we can restate the shipment requirement as follows:

$$S(y) = \max_{\varepsilon\in U}\max_{\pi\in\Pi} \sum_{i\in\mathcal{N}}\sum_{t=1}^{T} \pi_{it} \cdot \left(y_{it} - v_i + \sum_{t'=1}^{t-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right)\right). \tag{16}$$

For any $\pi \in \Pi$, let $S_\pi(y)$ be given by

$$\begin{aligned} S_\pi(y) &= \max_{\varepsilon \in U} \sum_{i\in\mathcal{N}} \sum_{t=1}^{T} \pi_{it}\left(y_{it} - v_i + \sum_{t'=1}^{t-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right)\right) \\ &= \sum_{i\in\mathcal{N}}\sum_{t=1}^{T}\pi_{it}y_{it} - \sum_{i\in\mathcal{N}}\sum_{t=1}^{T}\pi_{it}\left(v_i - \sum_{t'=1}^{t-1}\mu_{it'}\right) \\ &\quad + \max_{\varepsilon\in U}\sum_{i\in\mathcal{N}}\sum_{t=2}^{T}\pi_{it}\sum_{t'=1}^{t-1}\sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}. \end{aligned} \tag{17}$$

Observe that $S_\pi(y)$ is an affine function of y for each $\pi \in \Pi$. Furthermore, $S(y) = \max_{\pi\in\Pi} S_\pi(y)$. An equivalent formulation of the robust stock allocation problem (11)–(14) is as follows:

$$\min_{y,B} \sum_{t=1}^{T} B_t \tag{18}$$

$$\text{s.t. } B_t \ge w_{it}(\bar{d}_{it}(\delta) - y_{it}), \quad \forall\, i \in \mathcal{N},\ t = 1,\ldots,T; \tag{19}$$

$$B_t \ge 0, \quad t = 1,\ldots,T; \tag{20}$$

$$S_\pi(y) \le v_0, \quad \forall\, \pi \in \Pi. \tag{21}$$

Since for each $\pi \in \Pi$, $S_\pi(y)$ is affine in y and since Π is a finite set, it follows that the robust stock allocation problem can be expressed as a linear program. We refer to this formulation, (18)–(21), as the *Linear Form of the Robust Multiperiod Stock Allocation Problem*. In general, Π is a large set, and the coefficients of the linearized shipment requirement (17) can be expensive to compute. In the general case, we propose a decomposition approach. However, in a special case it is possible to solve the linear form directly. We treat the general case first.

3. Solution Techniques

3.1. The General Case: A Decomposition Algorithm

In this section, we apply Benders' decomposition approach to solve the general multiperiod robust stock allocation problem. We begin by reexpressing the worst-case shipment requirement (16), $\max_{\pi\in\Pi} S_\pi(y)$, which is a bilinear program, as a mixed-integer linear program (MILP) in the following manner. We know that for each retailer i there is a last period, $t_i^*(\varepsilon)$, in which it receives stock. For $t_i^*(\varepsilon) > 0$, let $u_{it} = 0$ if $t = t_i^*(\varepsilon)$ for retailer i and $u_{it} = 1$ otherwise. By Corollary 3, we will have $u_{it} = 1 - \pi_{it}$ at the optimum. Let $u_{i0} = 0$ if no shipment is made to retailer i during the time horizon and $u_{i0} = 1$ otherwise. This covers the situation when $t_i^*(\varepsilon) = 0$; that is, when v_i is sufficiently large that retailer i requires no shipment to meet the target inventory levels, y_{it}, for any $t = 1,2,\ldots,T$. Let S_i denote the total shipment required by retailer i under the adaptive shipment policy. By Lemma 2 and Proposition 4, S_i, for each retailer $i \in \mathcal{N}$, is given by

$$S_i = \begin{cases} y_{it_i^*(\varepsilon)} - v_i + \displaystyle\sum_{t'=1}^{t_i^*(\varepsilon)-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right) & \text{if } t_i^*(\varepsilon) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 6. *For a sufficiently large number, M, an equivalent MILP for determining the shipment requirement function is given by*

$$S(y) = \max_{\varepsilon,u,S} \sum_{i\in\mathcal{N}} S_i \tag{22}$$

$$\text{s.t. } S_i \le M u_{i0}, \quad i \in \mathcal{N}; \tag{23}$$

$$S_i \le y_{it} - v_i + \sum_{t'=1}^{t-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}\right) + M u_{it}, \quad i \in \mathcal{N},\ t = 1,2,\ldots,T; \tag{24}$$

$$\sum_{t=0}^{T} u_{it} = T, \quad i \in \mathcal{N}; \tag{25}$$

$$u_{it} \in \{0,1\}, \quad i \in \mathcal{N},\ t = 0,1,2,\ldots,T; \tag{26}$$

$$\varepsilon \in U. \tag{27}$$

Corollary 4. *At optimality for (16) and (22), there exist solutions satisfying $\pi_{it} = 1 - u_{it}$, for all $i \in \mathcal{N}$ and $t = 1,2,\ldots,T$.*

Applying Benders' approach, we state the master problem in linear form as

$$\min_{y,B} \sum_{t=1}^{T} B_t \tag{28}$$

$$\text{s.t. } B_t \ge w_{it}(\bar{d}_{it}(\delta) - y_{it}), \quad \forall\, i \in \mathcal{N},\ t = 1,\ldots,T; \tag{29}$$

$$B_t \ge 0, \quad t = 1,\ldots,T; \tag{30}$$

$$S_\pi(y) \le v_0, \quad \forall\, \pi \in \Pi^k, \tag{31}$$

where $S_\pi(y)$ is the affine function given by (17) and Π^k is some subset of the full set Π as of iteration k of the algorithm. The master problem can be solved as a linear program. Benders' decomposition algorithm as applied to this problem is described in Table 1.

Proposition 7. *Benders' decomposition algorithm (Table 1) solves the Linear Form of the Robust Multiperiod Stock Allocation Problem (18)–(21).*

Table 1. Benders' Decomposition Algorithm Applied to the Linear Form of the Robust Multiperiod Stock Allocation Problem

Input	$(\mu_{it}, \Sigma_t, w_{it})_{i\in\mathcal{N},t=1,2,\ldots,T}; v_0, (v_i)_{i\in\mathcal{N}}$; Assume $v_0 \ge 0$;
Output	$(y_{it}^*, B_t^*, \pi_{it}^*)_{i\in\mathcal{N},t=1,2,\ldots,T}$;
Step	
0	Set $k \leftarrow 0$; $\pi_{it}^0 \leftarrow 0, \forall\, i \in \mathcal{N}, t = 1,2,\ldots,T$; $\Pi^0 \leftarrow \{\pi^0\}$; $S_{\pi^0}(\cdot) \equiv 0$;
1	Solve the master problem (28)–(31) using Π^k; Extract solution y^k, B^k
2	Solve MILP (22)–(27) using y^k; Extract solution $S(y^k)$, u^k, and ε^k.
3	Set $\pi_{it}^k \leftarrow 1 - u_{it}^k, \forall\, i \in \mathcal{N}, t = 1,2,\ldots,T$;
4	If $S(y^k) \le v_0$ then stop: Set $(y^*, B^*, \pi^*) \leftarrow (y^k, B^k, \pi^k)$.
5	Else, using (17), express $S_{\pi^k}(y)$ in terms of y with ε^k and π^k;
6	Set $\Pi^k \leftarrow \Pi^k \cup \{\pi^k\}$;
7	Set $k \leftarrow k+1$; Return to step 1.

3.2. Two Periods, Identical Retailers, and Uncorrelated Demands

For the purpose of deriving insight, we consider the special case in which there are only two periods in the planning horizon, retailers are identical in their cost and demand characteristics, demands are uncorrelated between retailers, and there is no initial stock at any retailer. That is, $T=2$, $\mu_{it}=\mu_t$ and $w_{it}=w_t$ for all (i,t), $c_{ijt}=1_{\{i=j\}}\sigma_t$ for all (i,j,t), and $v_i=0$ for all $i\in\mathcal{N}$. In such cases, it is likely that each retailer will receive a shipment in the first period. The solution will therefore be indistinguishable from a situation in which initial retail inventories are nonzero but each retailer receives a shipment in the first period.

A consequence of (15) is that the optimal target stock levels are identical across retailers: $y^*_{it}=y^*_t$ for all $i\in\mathcal{N}$. Consequently, we restrict attention to policies with $y_{it}=y_t$ for all $i\in\mathcal{N}$. Let y_1 and y_2 denote the common target inventory levels for retailers in periods 1 and 2, respectively. Furthermore, we assume, for the balance of this section, $y\ge 0$ in the optimal solution to the robust stock allocation problem. This will likely be the case for sufficiently large values of the stock budget, v_0. By Corollary 3, the vector π selects the period in which retailers receive their last shipment. In the identical retailer case with two periods, let n denote the number of retailers that will receive their last shipment in period 2. Then $N-n$ is the number of retailers receiving their last shipment in period 1. The constraint on the total required shipment (21) can be simplified dramatically.

Proposition 8. *Under the assumptions of this section, (21) takes one of two forms, depending on the form of the uncertainty set. If $U=U(\delta,\bar{n})$ is given by (5), then (21) is given by*

$$\max_{n=0,1,\dots,N}\left\{(N-n)y_1+ny_2+n\mu_1+\frac{n}{\sqrt{\bar{n}\wedge n}}\sigma_1\delta\right\}\le v_0, \quad (32)$$

where $a\wedge b=\min(a,b)$ and the ratio $\frac{0}{0}$ is taken to be 0. Otherwise, if $U=U(\delta)$ is given by (3), then (21) is given by

$$\max_{n=0,1,\dots,N}\{(N-n)y_1+ny_2+n\mu_1+\min\{\delta_1,n\sigma_1\delta_0\}\}\le v_0. \quad (33)$$

Explicit Risk Pooling Uncertainty Set. We first exploit the proposition for the case of the explicit risk pooling uncertainty set (5). Following (15), we let $B_t=w_t(\bar{d}_t(\delta)-y_t)$ for $t=1,2$. We assume that δ is large enough to ensure $y\ge 0$. For the purpose of deriving analytical insight, we relax the integer constraint on n and ignore the bounds on y related to the maximum demand in each period. Substituting back into the linear form (18)–(21), and discarding constants from the objective function, results in a restatement of the robust stock allocation problem as

$$\max_{y_1,y_2}\ \{w_1y_1+w_2y_2\} \quad (34)$$

$$\text{s.t.}\ \ y_1\le\frac{v_0}{N}, \quad (35)$$

$$Ny_1+\max_{1\le n\le N}\left\{\frac{n}{\sqrt{\bar{n}\wedge n}}\sigma_1\delta-n(y_1-y_2-\mu_1)\right\}\le v_0. \quad (36)$$

We designate this as the *Relaxed Robust Two-Period Stock Allocation Problem with Identical Retailers* (*RR*2). The maximand of the left-hand side of (36) is a concave function of n, provided $\sigma_1\delta\ge 0$.

Before detailing a solution to RR2, we provide an intuitive sketch of the approach. For this purpose, we set $\bar{n}=N$. There are three degrees of freedom in RR2: the choice of target inventory levels, y_1 and y_2, and the choice of the number of retailers to participate in the allocation of stock in period 2, n. One of the degrees of freedom is removed by noting that the constraint (36) will be binding in an optimal solution. With $\bar{n}=N$ and (36) binding, we write this as an expression equating the stock held in reserve at the central warehouse, v_0-Ny_1, with the stock required to raise minimum period 2 inventory levels to y_2 among the n retailers receiving shipments:

$$v_0-Ny_1=\max_{1\le n\le N}\left\{n\left(y_2-y_1+\mu_1+\frac{\sigma_1\delta}{\sqrt{n}}\right)\right\}. \quad (37)$$

By Lemma 5(c) in Appendix B, $\mu_1+\sigma_1\delta/\sqrt{n}$ is the worst-case period 1 demand occurring at each of the retailers receiving a period 2 shipment. A second degree of freedom is removed by applying the first-order condition for optimality of the maximand on the right-hand side:

$$\sqrt{n}=\frac{1}{2}\frac{\sigma_1\delta}{y_1-y_2-\mu_1}. \quad (38)$$

Using this expression to eliminate $y_1-y_2-\mu_1$ from (37) and simplifying leads to

$$v_0-Ny_1=n\left\{-\frac{1}{2\sqrt{n}}\sigma_1\delta+\frac{\sigma_1\delta}{\sqrt{n}}\right\}=\frac{1}{2}\sqrt{n}\sigma_1\delta.$$

That is, the maximum stock required for allocation in period 2 is given by $\frac{1}{2}\sqrt{n}\sigma_1\delta$. Furthermore, we can view the remaining degree of freedom as either the choice of y_1 or n. Then,

$$y_1=\frac{1}{N}\left(v_0-\frac{1}{2}\sqrt{n}\sigma_1\delta\right).$$

In words, y_1 is the equal allocation, among N retailers, of initial inventory less the required reserve for period 2. Similarly, (38) can be written as

$$y_2=y_1-\mu_1-\frac{\sigma_1\delta}{\sqrt{n}}+\frac{(1/2)\sqrt{n}\sigma_1\delta}{n}.$$

That is, y_2 is equal to y_1 less the worst-case period 1 demand at a retailer receiving a period 2 shipment plus the equal allocation to the n retailers receiving shipments in period 2 of the stock held in reserve for period 2. The objective function of RR2 can be rewritten as

$$Z = (w_1 + w_2)\frac{1}{N}\left(v_0 - \frac{1}{2}\sqrt{n}\sigma_1\delta\right) - w_2\mu_1 - w_2\frac{\sigma_1\delta}{2\sqrt{n}}.$$

The remaining degree of freedom is removed by the first-order condition for a maximum of Z:

$$n = \frac{w_2}{w_1 + w_2}N.$$

That is, the optimized worst-case number of retailers to participate in the period 2 allocation is a simple fraction of the total number of retailers and that fraction expresses the optimal trade-off in costs between period 1 and period 2. This proof sketch has ignored complications that might arise from the constraints $1 \le n \le N$ and also the possibility that $\bar{n} < N$. We deal with these complications in the detailed results to follow.

Proposition 9. *If $\bar{n} \ge 4$, the maximand of (36) is optimized by*

$$n^* = \begin{cases} 1 & \text{if } y_1 - y_2 - \mu_1 > \frac{1}{2}\sigma_1\delta, \\ \hat{n} & \text{if } \zeta^+ < y_1 - y_2 - \mu_1 \le \frac{1}{2}\sigma_1\delta, \\ N & \text{if } y_1 - y_2 - \mu_1 \le \zeta^+, \end{cases} \tag{39}$$

where

$$\hat{n} = \left(\frac{\sigma_1\delta}{2(y_1 - y_2 - \mu_1)}\right)^2, \tag{40}$$

and

$$\zeta^+ = \frac{1}{2\sqrt{\bar{n}}}\sigma_1\delta\left\{1 + \sqrt{1 - \frac{\bar{n}}{N}}\right\}, \tag{41}$$

which is the unique solution in the interval $[(1/(2\sqrt{\bar{n}}))\sigma_1\delta, (1/\sqrt{\bar{n}})\sigma_1\delta]$ *satisfying*

$$\frac{N}{\sqrt{\bar{n}}}\sigma_1\delta = \frac{(\sigma_1\delta)^2}{4\zeta^+} + N\zeta^+.$$

Observe that for large values of $\bar{n}$ (i.e., values of $\bar{n}$ close to N), the critical point ζ^+ is well approximated by $\sigma_1\delta/(2\sqrt{\bar{n}})$. We are now in a position to characterize the optimal solution to RR2.

Theorem 1. *If $N > (w_1 + w_2)/w_2$, $\bar{n} \ge 4$, and $v_0 \ge \frac{1}{2}\sigma_1\delta \cdot \sqrt{w_2N/(w_1 + w_2)}$ then there exists a critical fraction $x_w \in [0, 1]$ such that the optimal solution to RR2 is given by the following*:

(1) *The worst-case number of retailers to receive shipments in period 2 is given by*

$$n^*(\bar{n}) = \begin{cases} \dfrac{w_2N}{w_1 + w_2} & \text{if } \bar{n} \ge x_wN, \\ \dfrac{\bar{n}}{(1 + \sqrt{1 - \bar{n}/N})^2} & \text{if } \bar{n} < x_wN. \end{cases} \tag{42}$$

(2) *The optimal amount of stock to hold at the central warehouse in period 1 is given by*

$$v_0 - Ny_1^* = \tfrac{1}{2}\sigma_1\delta\sqrt{n^*(\bar{n})}. \tag{43}$$

(3) *The optimal allocation to each retailer in period 1 is given by*

$$y_1^* = \frac{1}{N}\left(v_0 - \frac{1}{2}\sigma_1\delta\sqrt{n^*(\bar{n})}\right). \tag{44}$$

(4) *The optimal minimum inventory target level for period 2 is given by*

$$y_2^* = y_1^* - \mu_1 - \frac{1}{2}\sigma_1\delta\frac{1}{\sqrt{n^*(\bar{n})}}. \tag{45}$$

(5) *The critical fraction x_w is given by*

$$x_w = 1 - \left(\frac{w_1}{w_1 + 2w_2}\right)^2. \tag{46}$$

From (42), we see the centrality of the ratio $w_2/(w_1 + w_2)$. In practice, we would argue that this fraction is at least $\frac{1}{2}$ and more likely closer to 1 for the situations we envision. The reason is that backorders in period 1 will typically result in rapid shipments from the central warehouse, perhaps directly to the customer, to satisfy the backorder, and perhaps as soon as the backorders occur. Thus, the backorder is being satisfied out the reserve stock. The weight w_1, therefore, simply needs to capture the incremental shipping costs. Backorders in period 2 are different. During this period, the central warehouse is out of stock, having allocated all reserve stock to the retailers at the beginning of the period. A backorder in this period will not be satisfied until the supplier shipment is received and the next cycle begins. The weight w_2 in this situation must reflect the customer dissatisfaction and waiting time as well. For this reason, we are likely to see $n^*(\bar{n}) < w_2N/(w_1 + w_2)$ in the solution, except for large values of $\bar{n}$. That is, the value of $\bar{n}$ will directly restrict the optimal amount of stock held in reserve at the central warehouse through Equation (43). We use numerical experiments in the next section to explore the impact of changing $\bar{n}$.

If we allow $\bar{n} = N$, then we observe that the stock employed for pooling risk in this allocation problem, $v_0 - Ny_1^*$, grows with the square root of the number of retailers. Thus, we have arrived at a square root law without using expected value arguments. Another insight is that y_2^* approaches $y_1^* - \mu_1$ as $N \to \infty$. That is, due to risk pooling, the target minimum inventory level for each retailer in period 2 becomes less dependent on demand variability as the number of retailers increases.

In interpreting this theorem, keep in mind that it is the solution to a relaxed problem. In particular, we have ignored the bounds on y that limit target stock levels in the first period to the maximum possible demand in the period. We also assume that the initial inventory at the central warehouse is at least as large as the desired quantity of stock to withhold.

Implicit Risk Pooling Uncertainty Set. The approach for analyzing the impact of the implicit risk pooling uncertainty set parallels the approach for the explicit risk pooling set but the results are simpler. We now take $U = U(\delta)$ to be given by (3). Following the same analysis performed on the explicit risk pooling uncertainty set, we ignore the upper bound constraints on y_t and allow n, the number of retailers to receive shipments in period 2, to be continuous. By Proposition 8, the master problem can be relaxed as

$$\max_{y_1, y_2} \{w_1 y_1 + w_2 y_2\} \tag{47}$$

$$\text{subject to } \max_{0\le n\le N} \{(N-n)y_1 + n y_2 + n\mu_1 + \min\{\delta_1, n\sigma_1\delta_0\}\} \le v_0. \tag{48}$$

Then, the optimal solution to this problem can be characterized by the following theorem.

Theorem 2. *If $v_0 \ge N(\mu_1 + \sigma_1\delta_0)$, then*

(a) *the optimal solution (y_1^*, y_2^*) to the relaxed problem (47)–(48) is given by*

$$(y_1^*, y_2^*) = \begin{cases} \left(\dfrac{v_0}{N}, \dfrac{v_0}{N} - (\mu_1 + \sigma_1\delta_0)\right) & \text{if } \dfrac{w_2 N}{w_1 + w_2} \le \dfrac{\delta_1}{\sigma_1\delta_0}, \\ \left(\dfrac{v_0 - \delta_1}{N}, \dfrac{v_0 - \delta_1}{N} - \mu_1\right) & \text{if } \dfrac{w_2 N}{w_1 + w_2} > \dfrac{\delta_1}{\sigma_1\delta_0}; \end{cases} \tag{49}$$

(b) *the corresponding worst-case number of retailers to receive shipment in period 2 is degenerate, but one such solution is given by*

$$n^* = \frac{\delta_1}{\sigma_1\delta_0}; \tag{50}$$

(c) *the optimal amount of stock to hold in reserve at the central warehouse in period 1 is given by*

$$v_0 - N y_1^* = \begin{cases} 0 & \text{if } \dfrac{w_2 N}{w_1 + w_2} \le \dfrac{\delta_1}{\sigma_1\delta_0}, \\ \delta_1 & \text{if } \dfrac{w_2 N}{w_1 + w_2} > \dfrac{\delta_1}{\sigma_1\delta_0}. \end{cases} \tag{51}$$

This theorem is less satisfying than Theorem 1 because, instead of the insight of a square root law, we find only that the optimal stock to hold in reserve is given by the management parameter, δ_1, provided that the relative penalty $w_2/(w_1 + w_2)$ on period 2 backorders is sufficiently large. As mentioned earlier, care must be taken to choose δ_1 in a way that captures the risk pooling phenomenon. We provide such a method in Appendix A.

4. Experimental Results

In this section, we report on an empirical study conducted to assess the contribution that robust optimization can make to optimizing stock allocation problems. A general scheme for generating test cases is detailed in Appendix A. The relevant parameters are as follows:

- N, the number of retailers;
- T, the number of time periods (opportunities for allocation) in the cycle;
- $\bar{\mu}$, the average demand per retailer per day;
- $\bar{l}$, the average number of days per period;
- ψ, the coefficient of variation of daily demand of the smallest retailer (larger retailers experience lower coefficients of variation; hence, ψ is the maximum coefficient of variation across retailers);
- β_D, the Pareto demand shape parameter ($\beta_D = 0.2$ results in identically distributed demands across retailers, and $\beta_D = 0.8$ results in an 80-20 distribution: 80% of expected demand comes from only 20% of the retailers);
- β_L, the Pareto period-length shape parameter ($\beta_L = 0.2$ results in equal period lengths, and $\beta_L = 0.8$ results in an 80-20 distribution: 80% of total cycle days are concentrated in the first 20% of the number of periods);
- γ, the number of standard deviations of total system demand to add to expected demand for initial system inventory;
- ρ, the demand correlation coefficient, assumed to be the same for all pairs of retailers; ρ can be negative but not all negative correlations can be achieved, as explained in Appendix A;
- α, the probability of excess demand used to generate δ_1 in (3);
- θ, the growth factor for backorder costs over time.

The shape parameters allow us to generate test cases with nonidentical retailers and unequal period lengths. For most cases, we take $w_{it} = 1$ for all $i \in \mathcal{N}$ and all $t = 1, 2, \ldots, T$. In particular, we do not consider backorder costs that differ by location. We do consider different values of θ that allows different weights by period. The parameters used to generate test cases are summarized in Appendix A.

For the uncorrelated demand case, demands are simulated using a log-normal distribution. That is, we set $\tilde{d}_{it} = e^{(m_{it} + s_{it}\tilde{\varepsilon}_{it})}$, for each $i \in \mathcal{N}$ and all $t = 1, 2, \ldots, T$, where each $\tilde{\varepsilon}_{it}$ is drawn independently from a $N(0,1)$ distribution. The parameters m_{it} and s_{it} are chosen so that the mean and variance of $\tilde{d}_{it}$ are μ_{it} and σ_{it}^2, respectively. The log-normal is a heavy-tailed distribution and this by itself is more likely to give rise to situations violating the balance assumption than using a thin-tailed distribution such as the Poisson. While situations exist in practice in which the balance assumption is appropriate, our interest is in identifying stock allocation algorithms that perform well in situations in which the balance assumption fails. Further observe that the log-normal model of demand does not agree with the model of demand underlying the robust allocation approach (1): it is not possible to write demand as an affine function of a vector of mean-zero, unit-variance random variables. This puts the robust allocation approach at a potential disadvantage relative to heuristics that can exploit the true distribution. Note, also, that the allocation periods need not have identical lengths. Demand rates are scaled accordingly.

For the correlated demand case, we set

$$\tilde{d}_{it} = e^{(m_{it}+\sum_{j\in\mathcal{N}} p_{ijt}\tilde{\varepsilon}_{jt})} \tag{52}$$

for each $i \in \mathcal{N}$ and all $t = 1, 2, \ldots, T$. Similar to the uncorrelated demand case, the parameters m_{it} and p_{ijt} are chosen so that the mean and variance of $\tilde{d}_{it}$ are μ_{it}, σ_{it}^2 and the correlation between $\tilde{d}_{it}$ and $\tilde{d}_{jt}$ is ρ for all $i \neq j$, respectively. When $\rho = 0$, this expression of $\tilde{d}_{it}$ simplifies to one for the uncorrelated demand case.

4.1. Policies and Performance Metrics

Each simulation run corresponds to a sample vector $\tilde{\varepsilon}^k$, given by $\tilde{\varepsilon}^k = (\tilde{\varepsilon}_{it}^k)_{(i,t)\in\mathcal{N}\times\{1,2,\ldots,T\}}$ for sample index k. We generate a total of 10,000 sample vectors (or cycles) and reuse the same sample set with each experiment as a variance reduction technique. Each experiment is divided into 10 sample groups with $K = 1{,}000$ sample vectors in each group. Performance metrics are computed for each sample group and then averaged over all groups. A confidence interval based on a t-statistic with 95% confidence and nine degrees of freedom is computed for each performance metric.

We simulate the performance of different policies in a rolling-horizon setting and compare them under different metrics. For simplicity, we assume that initial inventories at all retailers are zero: $v_i = 0$, for all $i \in \mathcal{N}$. Let $\tilde{x}^{kP} = (\tilde{x}_{it}^{kP})$ denote the actual shipments made to retailers in each period under a generic policy P for a sample vector k. Two metrics considered for each sample group are average time-weighted backorders,

$$B^P = \frac{1}{K}\sum_{k=1}^{K}\sum_{i\in\mathcal{N}}\sum_{t=1}^{T} w_{it}\left(\sum_{t'=1}^{t}(\tilde{d}_{it'}^k - \tilde{x}_{it'}^{kP})\right)^+,$$

and average (unweighted) terminal backorders,

$$B_T^P = \frac{1}{K}\sum_{k=1}^{K}\sum_{i\in\mathcal{N}}\left(\sum_{t=1}^{T}(\tilde{d}_{it}^k - \tilde{x}_{it}^{kP})\right)^+.$$

Let D_T denote the average total demand in a sample group:

$$D_T = \frac{1}{K}\sum_{k=1}^{K}\sum_{i\in\mathcal{N}}\sum_{t=1}^{T}\tilde{d}_{it}^k.$$

The terminal fill rate of a policy is the fraction of demands that are satisfied by the end of the cycle:

$$\mathbf{F}_T^P = \left(1 - \frac{B_T^P}{D_T}\right) \times 100\%.$$

Robust Allocation in a Rolling Horizon. The *Linear Form of the Robust Multiperiod Stock Allocation Problem* is given by (18)–(21). Implicit in the definition of the shipment requirement function, $S_\pi(y)$, is the vector of initial inventories v. The robust allocation policy is formed by re-solving the robust multiperiod stock allocation problem at the beginning of each period, setting the problem data to match the remaining periods of allocation and the initial inventories to match the residual net inventories at the end of the previous period. The *Robust Allocation* policy is denoted as $\tilde{y}^{RA}$ and its computed average time-weighted backorders as B^{RA} and average terminal backorders as B_T^{RA}.

Lower Bound Policy. A lower bound on the minimal weighted number of expected backorders can be developed by allowing costless, instantaneous rebalancing of retail inventories in each period. Suppose, in each period, echelon inventory is instantaneously reallocated without restriction among retailers. We refer to this as the *Rebalance* policy and denote it with a superscript RB.

Upper Bound Policy. A natural upper bound on backorder metrics can be found by following the Eppen and Schrage *Ship All* policy in which all stock is allocated in the first period and no stock is held in reserve for rebalancing purposes. Under the Ship All policy, stock is allocated in the first period to minimize expected terminal unweighted backorders. We use a superscript SA to denote the Ship All policy.

Risk Pooling Capture. Since the Ship All policy involves no risk pooling and the Rebalance policy captures all possible risk pooling, the differences in metrics, $B^{SA} - B^{RB}$ and $B_T^{SA} - B_T^{RB}$, are upper bounds on the possible risk pooling benefit (average time-weighted backorders and average terminal backorders, resp.) from holding some stock in reserve at the central warehouse. For each experiment sample group, we score the Robust Allocation policy using percentages

$$\mathbf{C}^{RA} = \frac{B^{SA} - B^{RA}}{B^{SA} - B^{RB}} \times 100\%$$

and

$$\mathbf{C}_T^{RA} = \frac{B_T^{SA} - B_T^{RB}}{B_T^{SA} - B_T^{RB}} \times 100\%,$$

called the *capture percentage* and *terminal capture percentage*, respectively. Using the 10 sample groups for each experiment, we construct 95% confidence intervals around the mean capture percentages and mean terminal capture percentages.

Problem Difficulty. Since the test cases differ widely, we need a metric to compare the ease or difficulty of the underlying allocation problem. The terminal fill rate of the Rebalance policy, $\mathbf{F}_T^{RB}$, which we call the *Rebalance terminal fill rate*, serves this purpose. The lower this ideal fill rate is, the more difficult the underlying allocation problem must be. We also compute terminal fill rates for the Ship All policy, $\mathbf{F}_T^{SA}$, and for the Robust Allocation policy, $\mathbf{F}_T^{RA}$. As with the other metrics, we use the 10 sample groups to estimate means and confidence intervals for $\mathbf{F}_T^{SA}$, $\mathbf{F}_T^{RB}$, and $\mathbf{F}_T^{RA}$.

4.2. Competing Heuristics

The Robust Allocation (RA) heuristic has two forms, depending on the form of the risk pooling constraint set. If the implicit risk pooling uncertainty set (3) is used, then we refer to it as the Robust Allocation Implicit (RAI) heuristic. Otherwise, (RA) will refer to the heuristic with the explicit risk pooling uncertainty set (5). We select two competing heuristics from the literature for implementation and benchmarking. Both of the selected approaches use explicit distributional assumptions and so can be tailored for the log-normal demand model used in these experiments.

4.2.1. Infinite Retailers. Jackson and Muckstadt (1989) propose a heuristic for solving two-period allocation problems based on a limit theorem in which the number of retailers is allowed to go to infinity (with scaled demands). Their method can handle nonidentical retailers provided the demands are uncorrelated. It is relatively straightforward to extend their heuristic to more than two periods, though without the limit theorem justification. We refer to the resulting heuristic as the Infinite Retailers (IR) heuristic. Details are available on request.

4.2.2. Lagrangian Relaxation. Kunnumkal and Topaloglu (2008) extend the relaxation strategy of Federgruen and Zipkin (1984b) to a Lagrangian Relaxation strategy. An implicit assumption of their model as applied to the allocation problem of this paper is that the holding cost rate at the central warehouse is higher than the holding cost rate at the retailers. There is therefore always an economic motive for holding stock centrally. It can be shown that the solution to their model is unbounded in the case of identical holding cost rates. On the other hand, it is not difficult to extend their approach by adding a constraint to prevent an unbounded solution and to place a Lagrange multiplier on this constraint. We refer to this modified approach as the Lagrangian Relaxation (LR) heuristic. Details are available on request.

4.3. Test Results for Identical Retailers and Identical Period Lengths

In this section we report on all experiments in which the retailers are identical in their demand and cost characteristics and the period lengths are also identical. It happens that these are the conditions most favorable to the Robust Allocation policy. Unless otherwise noted in the test results, the number of retailers, N, is 4, the number of allocation periods, T, is 2, the average daily demand per retailer, $\bar{\mu}$, is 5, the average period length, $\bar{l}$, is 5, the safety stock factor, γ, is 2, the correlation, ρ, is 0, the explicit uncertainty set (4) is implemented, and its parameter, δ, is 2, the same value as the safety stock factor, γ.

4.3.1. Varying COVs. In this set of experiments, we vary the coefficient of variation of daily retailer demand, ψ. Table 2 shows the results. We notice immediately that the Robust Allocation policy is capturing the risk pooling potential ranging from 33.57% to 65.11% in terms of time-weighted capture percentage and from 54.88% to 100% in terms of terminal capture percentage. These capture percentages decrease markedly with increases in the coefficient of variation. The Lagrangian Relaxation policy captures almost no risk pooling potential, indicating that it has no significant difference from the Ship All policy. With the exception of the lowest coefficient of variation case, the Infinite Retailers heuristic captures little of the potential risk pooling benefit as measured by C^{IR}. In fact, it

Table 2. Simulation Results for Four Identical Retailers, Two Equal-Length Periods, and Varying Coefficients of Variation, ψ, Showing Capture Percentages ($\mathbf{C}^{RA}$, $\mathbf{C}^{LR}$, and $\mathbf{C}^{IR}$), Terminal Capture Percentages ($\mathbf{C}_T^{RA}$, $\mathbf{C}_T^{LR}$, and $\mathbf{C}_T^{IR}$) and Terminal Fill Rates ($\mathbf{F}_T^{SA}$, $\mathbf{F}_T^{RB}$, $\mathbf{F}_T^{RA}$, $\mathbf{F}_T^{LR}$, and $\mathbf{F}_T^{IR}$)

ψ	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{C}^{LR}$	$\mathbf{C}_T^{LR}$	$\mathbf{C}^{IR}$	$\mathbf{C}_T^{IR}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$	$\mathbf{F}_T^{LR}$	$\mathbf{F}_T^{IR}$
0.5	65.11	100.00	0.17	0.17	88.35	100.00	98.44	99.18	99.18	98.45	98.18
	±1.71	±0.00	±0.01	±0.01	±0.90	±0.01	±0.06	±0.04	±0.04	±0.06	±0.04
1.0	53.95	99.19	0.00	0.00	4.40	99.78	96.46	98.01	98.00	96.46	98.01
	±1.80	±0.51	±0.00	±0.00	±2.64	±0.27	±0.13	±0.10	±0.10	±0.13	±0.10
1.5	53.19	89.82	0.00	0.00	−14.80	98.67	94.28	96.69	96.44	94.28	96.66
	±1.63	±1.19	±0.00	±0.00	±3.69	±0.43	±0.23	±0.16	±0.17	±0.23	±0.17
2.0	45.94	70.75	0.00	0.00	−9.03	95.79	92.12	95.36	94.41	92.12	95.23
	±1.48	±1.77	±0.00	±0.00	±4.00	±0.49	±0.32	±0.23	±0.24	±0.32	±0.23
2.5	37.24	56.96	0.00	0.00	−2.53	92.32	90.12	94.09	92.38	90.12	93.78
	±1.53	±1.83	±0.00	±0.00	±3.70	±0.90	±0.41	±0.30	±0.33	±0.41	±0.30
3.0	33.57	54.88	0.00	0.00	2.40	88.85	88.32	92.91	90.83	88.32	92.40
	±1.46	±1.89	±0.00	±0.00	±3.08	±1.19	±0.49	±0.37	±0.40	±0.49	±0.36

Table 3. Simulation Results for Four Identical Retailers, Three Equal-Length Periods, and Varying Coefficients of Variation, ψ, Showing Capture Percentages ($\mathbf{C}^{RA}$, $\mathbf{C}^{LR}$, and $\mathbf{C}^{IR}$), Terminal Capture Percentages ($\mathbf{C}_T^{RA}$, $\mathbf{C}_T^{LR}$, and $\mathbf{C}_T^{IR}$), and Terminal Fill Rates ($\mathbf{F}_T^{SA}$, $\mathbf{F}_T^{RB}$, $\mathbf{F}_T^{RA}$, $\mathbf{F}_T^{LR}$, and $\mathbf{F}_T^{IR}$)

ψ	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{C}^{LR}$	$\mathbf{C}_T^{LR}$	$\mathbf{C}^{IR}$	$\mathbf{C}_T^{IR}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$	$\mathbf{F}_T^{LR}$	$\mathbf{F}_T^{IR}$
0.5	57.41	100.00	0.00	0.00	92.99	100.00	98.74	99.55	99.55	98.74	99.55
	±2.09	±0.00	±0.00	±0.00	±0.96	±0.00	±0.03	±0.02	±0.02	±0.03	±0.02
1.0	42.42	99.11	0.00	0.00	23.69	99.98	97.16	98.87	98.85	97.16	98.87
	±2.46	±0.49	±0.00	±0.00	±2.82	±0.11	±0.08	±0.05	±0.05	±0.08	±0.05
1.5	42.81	94.13	0.00	0.00	−21.75	99.13	95.39	98.06	97.90	95.39	98.03
	±1.66	±1.15	±0.00	±0.00	±3.01	±0.45	±0.13	±0.09	±0.07	±0.13	±0.08
2.0	46.30	87.96	0.00	0.00	−23.90	97.56	93.58	97.19	96.75	93.58	97.10
	±1.45	±1.10	±0.00	±0.00	±3.25	±0.76	±0.20	±0.13	±0.11	±0.20	±0.11
2.5	47.67	80.97	0.00	0.00	−13.89	95.52	91.86	96.32	95.47	91.85	96.12
	±1.34	±1.12	±0.00	±0.00	±3.39	±1.02	±0.27	±0.17	±0.16	±0.27	±0.15
3.0	47.11	73.98	0.00	0.00	−5.54	93.44	90.27	95.48	94.12	90.26	95.14
	±1.22	±1.12	±0.00	±0.00	±3.39	±1.07	±0.33	±0.22	±0.21	±0.33	±0.20

is worse than the Ship All policy in many cases. On the other hand, it outperforms both Lagrangian Relaxation and Robust Allocation when measured by the terminal backorder capture ratio, C_T^{IR}. This suggests that the Infinite Retailers policy holds back an excessive amount of stock at the central warehouse. This pattern, in which the Lagrangian Relaxation heuristic holds back too little stock and the Infinite Retailers heuristic holds back too much, is common across all of the experimental runs. The Robust Allocation heuristic is the only one to consistently capture a share of the risk pooling benefit. Not surprisingly, we see that the ideal terminal fill rate, as measured by F_T^{RB}, decreases as the coefficient of variation increases, indicating that the difficulty of the underlying allocation problem is increasing in the coefficient of variation.

4.3.2. Three Equal-Length Periods. In this section, we consider varying coefficients of variation, but we extend the number of periods to $T = 3$. Table 3 shows the results, and we compare the results with Table 2. We note first that the rebalance terminal fill rate is generally higher for the three-period case than for the two-period case, suggesting that the opportunities for risk pooling are greater. The Robust Allocation policy terminal capture percentage is also generally higher for the three-period case than for the two-period case. It is therefore capturing a larger share of a larger pie. The story is mixed for the robust capture percentage: in the three-period test case, the capture percentage is smaller for lower coefficients of variation ($\psi = 0.5, 1, 1.5$) than in the two-period case, but the reverse is true for higher coefficients of variation ($\psi = 2, 2.5, 3$). The Lagrangian Relaxation policy still captures nothing and will be omitted from later comparison unless it yields interesting results. The Infinite Retailers heuristic fares worse than the Ship All policy in the capture percentage at high coefficients of variation but does well for very low coefficients of variation. It outperforms all other heuristics in terms of the terminal capture percentage. This continues to suggest that it withholds too much stock at the central warehouse, except for low coefficients of variation.

4.3.3. Varying Backorder Weight Growth Factors. In this subsection, we consider the effect of the backorder weight growth factor with $T = 2$. In this case, we set $w_{i1} = 1$ and $w_{i2} = \theta$ for all $i \in \mathcal{N}$. We would anticipate that the amount of stock held in reserve at the central warehouse is increasing in the backorder factor. To check this, we compute the average central stock level at the end of period 1 under each heuristic policy that exploits risk pooling: the Robust Allocation policy and the Infinite Retailer policy. These numbers are shown in Table 4, where R^P is used to denote the stock held in reserve at the central warehouse at the end of period 1 under generic policy P. Two observations are clear from this table. First, the Infinite Retailers policy holds almost three times as much central reserve stock as the Robust Allocation policy. This is consistent with observations in earlier sections that the Infinite Retailers policy is much more conservative than the Robust Allocation. Second, the amount of central reserve stock tends to increase with the backorder factor, particularly for the Infinite Retailers policy. The Robust Allocation central reserve stock is unchanged from $\theta = 1$ to $\theta = 2$. This might seem to contradict (42) and (43), which suggests a continuous change in the

Table 4. Simulation Results for Four Identical Retailers, Two Identical Periods, High Coefficient of Variation, $\psi = 3$, and Varying Backorder Weight Growth Factor, θ, Showing Average Stock Held in Reserve After Period 1 for RA and IR Policies, R^{RA} and R^{IR}

θ	R^{RA}	R^{IR}
1	52.23	149.74
2	52.23	169.74
4	62.27	185.74

Table 5. Simulation Results for Four Identical Retailers, Two Identical Periods, High Coefficient of Variation, $\psi = 3$, and Varying Backorder Weight Growth Factor, θ, Showing Capture Percentages ($\mathbf{C}^{RA}$ and $\mathbf{C}^{IR}$), Terminal Capture Percentages ($\mathbf{C}_T^{RA}$ and $\mathbf{C}_T^{IR}$), and Terminal Fill Rates ($\mathbf{F}_T^{SA}, \mathbf{F}_T^{RB}, \mathbf{F}_T^{RA}$, and $\mathbf{F}_T^{IR}$)

θ	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{C}^{IR}$	$\mathbf{C}_T^{IR}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$	$\mathbf{F}_T^{IR}$
1	33.57	54.88	2.40	88.85	88.32	92.91	90.83	92.40
	±1.46	±1.89	±3.08	±1.19	±0.49	±0.37	±0.40	±0.36
2	43.91	54.88	32.18	91.87	88.32	92.91	90.83	92.54
	±2.04	±1.89	±13.84	±0.99	±0.49	±0.37	±0.40	±0.36
4	54.21	60.78	62.43	93.74	88.32	92.91	91.10	92.62
	±1.72	±1.87	±1.42	±0.77	±0.49	±0.37	±0.40	±0.36

reserve stock. Recall, however, that we dropped the integrality restriction on n so it is not surprising to find the central reserve stock to be somewhat resistant to change. Table 5 shows the performance results. All of the capture ratios increase with the backorder growth rate factor. This is clearly because more stock is held in reserve to prevent period 2 backorders, and these backorders are weighted more heavily than in period 1.

4.3.4. Varying COVs and Correlations. In this subsection, we consider the impact of demand correlations. As shown in Zipkin (1984), the potential benefit of risk pooling decreases as demand correlation increases. To model correlated demand in (2), we must allow the c_{ijt} parameters to take on negative values. This change necessitates straightforward changes to the uncertainty sets. For example, (3) becomes

$$U(\delta) = \left\{ \varepsilon\colon \begin{array}{c} \varepsilon_{it} \le \delta_0, \forall (i,t); \\ \left| \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}} c_{ijt} \varepsilon_{jt} \right| \le \delta_1, t = 1, \dots, T \end{array} \right\},$$

and a similar change is required for (5). As noted in the introduction to this section, not all demand correlation coefficients are achievable. In particular, we restrict $\rho \ge -0.2$ for $\psi < 3.0$ and $\rho \ge -0.15$ otherwise. We consider two coefficients of variation ($\psi \in \{1, 3\}$), and a range of correlation coefficients from −0.2 to 1.0. Table 6 shows the results. The first thing to notice is that the ideal terminal fill rate, $\mathbf{F}_T^{RB}$, is generally decreasing in the correlation coefficient, ρ, for each coefficient of variation, ψ. This is consistent with Zipkin's result: there is less opportunity for risk pooling when demands are positively correlated. When demands are perfectly positively correlated, the ideal terminal fill rate is exactly the same as the terminal fill rate for the Ship All policy. As measured by terminal fill rates, the Robust Allocation policy is worse than the Ship All policy at high correlation coefficients and high coefficients of variation. It is competitive only when the correlation coefficient is at or below 0.2. The capture percentages tell a similar story. The values of capture percentages are not available when $\rho = 1.0$ because there ceases to

Table 6. Simulation Results for Four Identical Retailers, Two Identical Periods, Varying COVs, ψ, and Correlation ρ, Showing Robust Capture Percentage ($\mathbf{C}^{RA}$), Robust Terminal Capture Percentage ($\mathbf{C}_T^{RA}$), and Terminal Fill Rates ($\mathbf{F}_T^{SA}, \mathbf{F}_T^{RB}$, and $\mathbf{F}_T^{RA}$)

ψ	ρ	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$
1.0	−0.2	69.74 ± 1.52	94.57 ± 1.07	96.44 ± 0.11	98.24 ± 0.09	98.14 ± 0.09
	−0.1	67.91 ± 1.60	97.56 ± 0.74	96.46 ± 0.13	98.14 ± 0.09	98.09 ± 0.10
	0.0	53.95 ± 1.80	99.19 ± 0.51	96.46 ± 0.19	98.01 ± 0.10	98.00 ± 0.10
	0.2	63.94 ± 1.93	94.69 ± 1.02	96.46 ± 0.14	97.75 ± 0.11	97.68 ± 0.12
	0.5	47.99 ± 1.65	80.36 ± 1.22	96.45 ± 0.17	98.80 ± 0.17	96.75 ± 0.17
	0.8	7.93 ± 7.42	86.88 ± 2.74	96.45 ± 0.17	98.80 ± 0.17	96.75 ± 0.17
	1.0	NA	NA	92.01 ± 0.63	92.01 ± 0.63	92.01 ± 0.63
3.0	−0.15	13.89 ± 3.05	83.00 ± 1.90	88.34 ± 0.27	93.48 ± 0.36	92.61 ± 0.28
	−0.1	29.16 ± 2.15	76.64 ± 1.82	88.37 ± 0.39	93.39 ± 0.34	92.21 ± 0.32
	−0.05	35.06 ± 1.99	67.75 ± 2.09	88.34 ± 0.45	93.15 ± 0.36	91.60 ± 0.37
	0.0	33.57 ± 1.46	54.88 ± 1.89	88.32 ± 0.49	92.91 ± 0.37	90.83 ± 0.40
	0.2	18.31 ± 2.81	45.78 ± 1.51	88.33 ± 0.51	91.91 ± 0.44	89.97 ± 0.42
	0.5	−37.51 ± 8.01	0.70 ± 3.57	88.46 ± 0.59	90.58 ± 0.60	88.48 ± 0.56
	0.6	−69.11 ± 10.49	−28.92 ± 5.98	88.48 ± 0.65	90.14 ± 0.66	88.01 ± 0.63
	0.7	−105.05 ± 13.04	−60.23 ± 8.69	88.45 ± 0.71	89.69 ± 0.73	87.72 ± 0.71
	0.8	−128.91 ± 15.54	−66.45 ± 9.87	88.41 ± 0.80	89.24 ± 0.81	87.86 ± 0.81
	0.9	−171.83 ± 23.88	−90.71 ± 13.99	88.34 ± 0.90	88.76 ± 0.90	87.96 ± 0.92
	1.0	NA	NA	88.14 ± 1.03	88.14 ± 1.03	88.14 ± 1.03

Table 7. Simulation Results for Four Identical Retailers, Two Identical Periods, Varying COVs, ψ, and Probability of Excess Demand, α, Showing Robust Capture Percentage ($\mathbf{C}^{RAI}$), Robust Terminal Capture Percentage ($\mathbf{C}_T^{RAI}$), Robust Terminal Fill Rate ($\mathbf{F}_T^{RAI}$), and the Elapsed Time of Robust Allocation Policy (W_t^{RAI}), with the Implicit Uncertainty Set Implemented

ψ	α	$\mathbf{C}^{RAI}$	$\mathbf{C}_T^{RAI}$	$\mathbf{F}_T^{RAI}$	W_t^{RAI}
1.0	0.01	53.95 ± 1.80	99.19 ± 0.51	98.00 ± 0.10	0.39 ± 0.01
	0.02	53.95 ± 1.80	99.19 ± 0.51	98.00 ± 0.10	0.39 ± 0.02
	0.03	53.95 ± 1.80	99.19 ± 0.51	98.00 ± 0.10	0.40 ± 0.01
	0.05	53.95 ± 1.80	99.19 ± 0.51	98.00 ± 0.10	0.42 ± 0.02
	0.1	53.95 ± 1.80	99.19 ± 0.51	98.00 ± 0.10	0.40 ± 0.03
	0.2	53.95 ± 1.80	99.19 ± 0.51	98.00 ± 0.10	0.40 ± 0.02
	0.5	53.95 ± 1.80	99.19 ± 0.51	98.00 ± 0.10	0.40 ± 0.02
3.0	0.01	29.62 ± 0.81	76.02 ± 1.66	91.80 ± 0.38	0.35 ± 0.02
	0.02	18.14 ± 1.60	83.96 ± 1.35	92.17 ± 0.37	0.34 ± 0.02
	0.03	15.50 ± 1.85	85.04 ± 1.32	92.22 ± 0.37	0.33 ± 0.02
	0.05	21.43 ± 1.30	82.37 ± 1.40	92.10 ± 0.37	0.35 ± 0.02
	0.1	28.10 ± 0.84	77.60 ± 1.60	91.88 ± 0.38	0.32 ± 0.01
	0.2	33.18 ± 1.00	70.31 ± 1.81	91.54 ± 0.39	0.34 ± 0.01
	0.5	32.59 ± 1.49	51.62 ± 1.87	90.68 ± 0.41	0.28 ± 0.01

be a difference in performance between the Ship All and the Rebalance policies. The $\psi = 2.0$ case yields similar results and thus is omitted.

4.3.5. Varying COVs and Alphas.

The implicit risk pooling uncertainty set (3) is considered in this subsection. We change δ_1 in the implicit uncertainty set by varying α, the probability of excess demand, as described in Appendix A. The higher is the value of α, the lower will be the value of δ_1. Table 7 shows the results. When $\psi = 1.0$, it is easily seen that α has no impact on the risk pooling capture. Additionally, this result ($C^{RAI} = 53.95$) is exactly the same as the one in Table 2 for $\psi = 1.0$, where the result comes from the implementation of the explicit risk pooling uncertainty set (4) with $\delta = 2$. However, the computation time is shorter when using the implicit uncertainty set. This suggests that we can make use of the simpler implicit set to achieve the same risk pooling capture with higher efficiency when the coefficient of variation is low. When $\psi = 3.0$, α must be set at 0.5 to come close to matching the capture percentage of the explicit risk pooling set (Table 2). It is interesting that the terminal capture percentage peaks at $\alpha = 0.03$, whereas that value minimizes the capture percentage. We have checked the underlying solutions and see that throughout this range, the policy is to hold stock in reserve at the central warehouse and the quantity held in reserve varies with α. The $\psi = 2.0$ case yields similar pattern to the $\psi = 3.0$ case and is therefore omitted.

4.4. Test Results for Identical Retailers and Pareto Period Lengths

In this set of experiments, we consider identical retailers, but we make the first period longer than the second period by setting $\beta_L = 0.8$. As a result, the first period is eight days in length and the second period is two days in length. As motivated in Appendix A, we refer to this setting as *Pareto period lengths*. The general observation for this experiment set is that the Robust Allocation policy captures less of the potential risk pooling benefit in a relative sense. We argue that this is a consequence of Silver's insight that a longer first period reduces the risk of imbalance.

4.4.1. Varying COVs.

Table 8 shows the results for varying coefficients of variation. They initially suggest that shifting the allocation point to later in the cycle generally makes things worse: the robust capture percentages have all decreased from their counterparts in Table 2, except when $\psi = 0.5$. The terminal capture percentages are uniformly worse for the Robust Allocation heuristic. However, these are relative comparisons. The actual terminal fill rates, F_T^{RA}, have increased, as have the ideal terminal fill rates, F_T^{RB}. As first noted by Silver, this suggests that shifting the allocation point to later in the cycle makes it easier to achieve high service levels because the risk of imbalance is lower. The Robust Allocation policy has less to offer in this setting.

Comparing the results with Table 3, we note that the risk pooling opportunities are greater for the two-period Pareto-length case than for the three-equal-periods case (the rebalance terminal fill rates are generally higher in the former case). However, the differences are not as pronounced as with Table 2. On the other hand, the capture percentages are much lower for the two-period Pareto-length case, except when $\psi = 0.5$. A tentative conclusion is that given a choice between two regimes for allocating stock, three equal periods versus two periods with Pareto-lengths, the Robust Allocation policy will fare better under the three-equal-periods regime.

4.4.2. Varying Safety Stock Factors and Uncertainty Set Parameters.

In this set of experiments, all of the

Table 8. Simulation Results for Four Identical Retailers, Two Pareto-Length Periods, and Varying Coefficients of Variation, ψ, Showing Capture Percentages ($\mathbf{C}^{RA}$, $\mathbf{C}^{LR}$, and $\mathbf{C}^{IR}$), Terminal Capture Percentages ($\mathbf{C}_T^{RA}$, $\mathbf{C}_T^{LR}$, and $\mathbf{C}_T^{IR}$), and Terminal Fill Rates ($\mathbf{F}_T^{SA}$, $\mathbf{F}_T^{RB}$, $\mathbf{F}_T^{RA}$, $\mathbf{F}_T^{LR}$, and $\mathbf{F}_T^{IR}$)

ψ	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{C}^{LR}$	$\mathbf{C}_T^{LR}$	$\mathbf{C}^{IR}$	$\mathbf{C}_T^{IR}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$	$\mathbf{F}_T^{LR}$	$\mathbf{F}_T^{IR}$
0.5	67.46	81.73	0.05	0.06	−14.51	99.26	98.44	99.62	99.41	98.44	99.61
	±1.49	±1.25	±0.00	±0.00	±1.71	±0.46	±0.07	±0.02	±0.03	±0.07	±0.03
1.0	39.82	68.45	0.00	0.00	−15.49	93.86	96.45	99.01	98.20	96.45	98.85
	±1.13	±1.25	±0.00	±0.00	±2.93	±0.78	±0.16	±0.06	±0.09	±0.16	±0.07
1.5	29.23	64.81	0.00	0.00	1.58	85.66	94.27	98.30	96.88	94.27	97.72
	±1.14	±1.21	±0.00	±0.00	±2.46	±1.11	±0.26	±0.11	±0.16	±0.26	±0.13
2.0	24.58	61.86	0.00	0.00	9.92	78.07	92.11	97.57	95.48	92.11	96.36
	±1.16	±1.25	±0.00	±0.00	±1.98	±1.25	±0.38	±0.16	±0.26	±0.38	±0.22
2.5	22.29	59.63	0.00	0.00	14.10	72.05	90.09	96.83	94.10	90.09	94.94
	±1.23	±1.35	±0.00	±0.00	±1.83	±1.38	±0.50	±0.20	±0.35	±0.50	±0.32
3.0	21.10	58.08	0.00	0.00	16.20	67.35	88.25	96.09	92.80	88.25	93.53
	±1.36	±1.43	±0.00	±0.00	±1.84	±1.52	±0.63	±0.25	±0.45	±0.63	±0.42

experiments are run under a high coefficient of variation ($\psi = 3$). From Table 8, we would expect this to result in relatively low capture percentages for the Robust policy. We consider safety stock factors ranging from $\gamma = 1$, a stressed setting, up to $\gamma = 2.5$, a relaxed setting. For each setting of the safety stock factor, we consider at least five possible settings for the uncertainty set parameter: $\delta \in \{\gamma, \gamma \pm 0.1, \gamma \pm 0.5\}$. This is by way of testing a conjecture that the best choice of δ is $\delta = \gamma$. Since this conjecture does not hold, we extend the range in certain cases to seek a local maximum of the capture ratio.

Table 9 presents the results. As anticipated, all of the robust capture percentages are relatively low (under or around 20%) because of the high coefficient of variation and the choice of period length. The uncertainty set parameter, δ, has a small effect on the terminal fill rate, F_T^{RA}, never changing it by more than two percentage points. Of more interest is the direction of change in capture percentages with respect to local changes in the uncertainty set parameter about the value of the safety stock factor. For the stressed setting ($\gamma = 1.0$), there seems to be little impact of the uncertainty set parameter over a broad range. In the relaxed settings ($\gamma = 2.5$), the robust capture percentage reaches its maximum value at $\gamma - 0.5$. It is outside of the scope of this paper to recommend a setting for the uncertainty set parameter; it appears to be a management parameter with interesting trade-offs. Note also that the terminal fill rates of the Ship All policy, the Rebalance policy, and the Infinite Retailers policy do not vary with the uncertainty set parameter δ because this parameter does not enter the implementation of these heuristics. The cases for $\gamma = 1.5$ and 2.0 are not shown since they yield similar observations.

4.4.3. Varying COVs and Depth Parameters. In this subsection, we examine the impact of the depth parameter, $\bar{n}$, in the explicit uncertainty set (5). For these experiments, the average daily demand, $\bar{\mu}$, is set at 2.5, and the uncertainty set parameter, δ, is 1.5. The computer elapsed time, W_t^{RA}, of the Robust Allocation policy implementation is also listed for comparison. Table 10 shows the result. For any given depth, both the capture percentages and the terminal fill rates decrease as the coefficient of variation increases. The depth parameter has no impact when $\psi = 0.5$, but the capture percentages increase with the depth parameter when $\psi = 3.0$. Meanwhile, the Robust Allocation policy experiences longer computation time when depth is higher. This suggests that limiting the depth of the explicit uncertainty set decreases the complexity of the set at the cost of risk pooling capture. The results for $\psi = 2.0$ are not shown, since they lead to similar conclusions to those for $\psi = 3.0$. We repeated this experiment with 10 identical retailers, but the conclusions are the same, so the results are omitted.

4.5. Test Results for Nonidentical Retailers

In this set of experiments, we consider nonidentical retailers with $\beta_D = 0.8$. As motivated in Appendix A, we refer to this case as having *Pareto retailers* since they follow an 80-20 rule with respect to demand rates. We consider a high coefficient of variation, $\psi = 3$, but this is only for the smallest retailer. We also consider two Pareto periods, with $\beta_L = 0.8$. The general observation from this experiment set is that the Robust Allocation policy has little to offer in capturing the risk pooling benefit of centralized inventories. We argue that the reason for this is the risk pooling inherent in aggregating demand into a small number of large retailers.

4.5.1. Eight Pareto Retailers, Pareto-Length Periods, Varying Safety Stock Factors. In this subsection, the daily demand rates that result from this setting are shown in Table 11. As the table shows, our test case generation approach results in the larger retailers having lower coefficients of variation. Because of the computational effort involved in the eight retailer scenario,

Table 9. Simulation Results for Four Identical Retailers, Two Pareto-Length Periods, a High Coefficient of Variation, $\psi = 3$, and Various Combinations of Safety Stock Factor, γ, and Explicit Uncertainty Set Parameter, δ, Showing Capture Percentages ($\mathbf{C}^{RA}$ and $\mathbf{C}^{IR}$), Terminal Capture Percentages ($\mathbf{C}_T^{RA}$ and $\mathbf{C}_T^{IR}$), and Terminal Fill Rates ($\mathbf{F}_T^{SA}$, $\mathbf{F}_T^{RB}$, $\mathbf{F}_T^{RA}$, and $\mathbf{F}_T^{IR}$)

γ	δ	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{C}^{IR}$	$\mathbf{C}_T^{IR}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$	$\mathbf{F}_T^{IR}$
1.0	0.5	17.42 ±0.43	53.99 ±0.85	12.70 ±0.93	64.71 ±0.85	81.73 ±0.69	91.99 ±0.41	87.26 ±0.59	88.36 ±0.55
	0.9	16.94 ±0.31	37.14 ±0.70	″	″	″	″	85.54 ±0.62	″
	1.0	17.54 ±0.31	40.25 ±0.75	″	″	″	″	85.86 ±0.62	″
	1.1	17.94 ±0.30	43.22 ±0.79	″	″	″	″	86.16 ±0.61	″
	1.5	17.49 ±0.42	53.65 ±0.85	″	″	″	″	87.23 ±0.59	″
2.5	1.0	11.78 ±2.62	72.50 ±1.51	17.39 ±2.16	67.60 ±1.60	90.40 ±0.58	97.12 ±0.24	95.26 ±0.33	94.93 ±0.36
	1.5	22.91 ±1.20	46.20 ±1.51	″	″	″	″	93.50 ±0.43	″
	2.0	23.14 ±1.57	55.79 ±1.67	″	″	″	″	94.14 ±0.40	″
	2.4	21.19 ±1.86	62.05 ±1.70	″	″	″	″	94.56 ±0.38	″
	2.5	20.43 ±1.92	63.45 ±1.67	″	″	″	″	94.66 ±0.37	″
	2.6	19.58 ±2.00	64.80 ±1.65	″	″	″	″	94.75 ±0.37	″
	3.0	15.27 ±2.33	69.72 ±1.56	″	″	″	″	95.08 ±0.35	″

we limit consideration to only two settings of the safety stock parameter: $\gamma = 1.5$ and 2.0. Table 12 shows the results. In this case, the robust capture percentages are negative: the Robust Allocation policy performs worse than the Ship All policy with respect to time-weighted capture ratios. On the other hand, it performs quite well with respect to terminal capture ratios (greater than 98.5% in both cases). The Lagrangian Relaxation approach yields results similar to the Ship All policy. The Infinite Retailers heuristic performs somewhat better than the Robust Allocation heuristic, but it is marginally better than the Ship All policy in only the stressed case ($\gamma = 1.5$). The main observation is that the allocation problem is quite easy when there are non-identical Pareto retailers: the terminal fill rates for the Ship All policy and the Rebalance policy are all quite high. This can be traced to our assumption that coefficients of demand variation decrease with retailer size. Consequently, all of the heuristics are performing well in absolute terms.

Table 10. Simulation Results for Six Identical Retailers, Two Pareto Periods, Varying COVs, ψ, and Depth Parameters, $\bar{n}$, Showing Robust Capture Percentage ($\mathbf{C}^{RA}$), Robust Terminal Capture Percentage ($\mathbf{C}_T^{RA}$), Robust Terminal Fill Rates ($\mathbf{F}_T^{RA}$), and the Elapsed Time of Robust Allocation Policy (W_t^{RA}), in Seconds

ψ	$\bar{n}$, depth	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{F}_T^{RA}$	W_t^{RA}
0.5	1	61.99 ± 1.25	69.68 ± 1.00	99.02 ± 0.04	1.94 ± 0.02
	2	61.99 ± 1.25	69.68 ± 1.00	99.02 ± 0.04	2.19 ± 0.01
	3	61.99 ± 1.25	69.68 ± 1.00	99.02 ± 0.04	2.33 ± 0.01
	4	61.99 ± 1.25	69.68 ± 1.00	99.02 ± 0.04	2.54 ± 0.01
	5	61.99 ± 1.25	69.68 ± 1.00	99.02 ± 0.04	2.81 ± 0.08
	6	61.99 ± 1.25	69.68 ± 1.00	99.02 ± 0.04	2.90 ± 0.03
3.0	1	0.00 ± 0.00	0.00 ± 0.00	86.04 ± 0.49	0.72 ± 0.01
	2	20.93 ± 0.63	30.35 ± 0.81	89.01 ± 0.44	1.26 ± 0.01
	3	26.25 ± 0.76	40.24 ± 1.00	89.98 ± 0.42	1.33 ± 0.01
	4	28.93 ± 0.83	46.44 ± 1.10	90.59 ± 0.42	1.48 ± 0.01
	5	31.29 ± 0.86	53.94 ± 1.21	91.33 ± 0.40	1.90 ± 0.01
	6	32.22 ± 0.89	59.86 ± 1.29	91.91 ± 0.39	2.57 ± 0.01

Table 11. Daily Demand Mean and Coefficient of Variation by Retailer When $N = 8$, $\bar{\mu} = 5$, $\beta_D = 0.8$, and $\psi = 3$

	1	2	3	4	5	6	7	8
Mean	22.08	9.91	4.45	2.00	0.90	0.40	0.18	0.08
C.V.	0.18	0.27	0.41	0.60	0.90	1.35	2.01	3.00

4.5.2. Varying Demand Shape Parameters.

To further explore the phenomenon discovered in the previous experiment, we consider a range of demand shape parameters varying from $\beta_D = 0.2$ (identical retailers)

Table 12. Simulation Results for Eight Pareto Retailers, Two Pareto Periods, High Coefficient of Variation, $\psi = 3$, and Varying Safety Stock Factors, γ, Showing Capture Percentages ($\mathbf{C}^{RA}$, $\mathbf{C}^{LR}$, and $\mathbf{C}^{IR}$), Terminal Capture Percentages ($\mathbf{C}_T^{RA}$, $\mathbf{C}_T^{LR}$, and $\mathbf{C}_T^{IR}$), and Terminal Fill Rates ($\mathbf{F}_T^{SA}$, $\mathbf{F}_T^{RB}$, $\mathbf{F}_T^{RA}$, $\mathbf{F}_T^{LR}$, and $\mathbf{F}_T^{IR}$)

γ	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{C}^{LR}$	$\mathbf{C}_T^{LR}$	$\mathbf{C}^{IR}$	$\mathbf{C}_T^{IR}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$	$\mathbf{F}_T^{LR}$	$\mathbf{F}_T^{IR}$
1.5	−51.76	99.20	−3.73	−8.67	7.01	98.92	98.58	99.58	99.57	98.49	99.57
	±3.25	±0.17	±0.39	±0.30	±1.75	±0.17	±0.04	±0.01	±0.01	±0.05	±0.01
2.0	−27.32	98.94	−4.31	−9.51	−20.60	99.53	99.00	99.78	99.77	98.92	99.77
	±3.25	±0.23	±0.73	±0.62	±3.02	±0.16	±0.03	±0.01	±0.01	±0.04	±0.01

Table 13. Simulation Results for Four Retailers, Two Pareto Periods, High Coefficient of Variation, $\psi = 3$, and Varying Demand Shape Parameters, β_D, Showing Capture Percentages ($\mathbf{C}^{RA}$ and $\mathbf{C}^{IR}$), Terminal Capture Percentages ($\mathbf{C}_T^{RA}$ and $\mathbf{C}_T^{IR}$), and Terminal Fill Rates ($\mathbf{F}_T^{SA}$, $\mathbf{F}_T^{RB}$, $\mathbf{F}_T^{RA}$, and $\mathbf{F}_T^{IR}$)

β_D	$\mathbf{C}^{RA}$	$\mathbf{C}_T^{RA}$	$\mathbf{C}^{IR}$	$\mathbf{C}_T^{IR}$	$\mathbf{F}_T^{SA}$	$\mathbf{F}_T^{RB}$	$\mathbf{F}_T^{RA}$	$\mathbf{F}_T^{IR}$
0.2	21.10	58.08	16.20	67.35	88.25	96.09	92.80	93.53
	±1.36	±1.43	±1.84	±1.52	±0.63	±0.25	±0.45	±0.42
0.5	2.25	82.44	1.80	83.24	94.15	98.19	97.47	97.51
	±2.46	±1.25	±2.67	±1.11	±0.28	±0.13	±0.14	±0.13
0.8	−74.46	95.42	−83.07	98.41	99.22	99.78	99.75	99.77
	±7.11	±0.77	±6.17	±0.41	±0.04	±0.02	±0.02	±0.02

to $\beta_D = 0.8$ (Pareto retailers), but limited to four retailers. Table 13 displays the results. The ideal fill rate, $\mathbf{F}_T^{RB}$, from the Rebalance policy, reveals that the allocation problem becomes progressively easier as demand among the retailers becomes more unbalanced. Unbalancing the retailers is, in itself, a form of risk pooling because concentrating demand in a smaller number of retailers reduces the safety stock requirements.

As an aside, recall that system service levels can be increased by supplying retailers more frequently. Consequently, when demand rates are very low and coefficients of variation are very high, as for the low demand rate retailers in these examples, one should consider resupplying the low demand rate retailers more frequently than high demand rate retailers, from an inventory perspective. In fact, in practical applications, it may be desirable not to stock anything at some low demand rate retailers, and simply fill backorders as they arise.

4.6. Validation and Computation Time

4.6.1. Lower Demand Rate. In these experiments, we consider identical retailers, two equal-length periods, and varying coefficients of variation, but we cut the demand rate in half, $\bar{\mu} = 2.5$, to see if the results are sensitive to scale. The results are not displayed as they are identical to Table 2. This is not surprising because the coefficient of variation is preserved through the scaling.

4.6.2. Computation Time. Finally, we consider the computation time required to compute various policies in a rolling horizon. For each test case and each demand sample, we capture the elapsed time to compute the policy. We then compute the average elapsed time over all demand samples for each test case. All of the experiments are implemented under the 1.2-GHz Intel Core m5 processor. Table 14 shows the result across test cases with similar dimensions. For the Robust Allocation policy, there is a dramatic increase in computation time when the number of retailers increases. For eight retailers, the Robust Allocation policy requires approximately 40 seconds of computation time. The approach is amenable to parallel processing if it must be applied to many part numbers independently. The computation time of the Robust Allocation policy also significantly increases when the number of periods goes up to three: it is almost the same as the computation time in a six-retailer, two-period case. The computation time of the Lagrangian Relaxation policy, however, is more sensitive to the number of periods than to the number of retailers. The Infinite Retailers policy is not sensitive to either N or T.

Table 14. Computation Time (in Seconds) to Compute Various Policies in a Rolling Horizon Fashion for a Single Demand Sample, Averaged Over Samples and Summarized by Test Case Dimension

Scenario	W_t^{RA}	W_t^{RAI}	W_t^{LR}	W_t^{IR}
$N = 4, T = 2$	0.57 ± 0.02	0.36 ± 0.01	1.97 ± 0.06	0.06 ± 0.01
$N = 4, T = 3$	2.27 ± 0.16	NA	7.72 ± 0.76	0.09 ± 0.01
$N = 6, T = 2$	2.02 ± 0.24	NA	2.38 ± 0.20	0.11 ± 0.02
$N = 8, T = 2$	42.22 ± 0.72	NA	5.54 ± 0.36	0.02 ± 0.00

5. Conclusion

In this paper, we have revisited a long-standing multiechelon inventory allocation problem and considered it afresh from a robust optimization perspective. A novel formulation of the uncertainty set leads to a tractable computation approach and analytical insights. The experimental results demonstrate that the method is more effective at capturing the potential risk pooling benefit of centralized allocation than previously proposed heuristics.

Acknowledgments

The authors gratefully acknowledge the inspiration of Dimitris Bertsimas in formulating this model and the observation of Michael Todd leading to a more elegant proof of Proposition 1. The contribution of the reviewers to a more thorough exploration of the topic is also appreciated.

Appendix A. Experimental Design

In this section, we propose a model for generating test cases based on an economical set of generating parameters. Let μ_i denote the expected daily demand at retailer i, $i \in \mathcal{N}$. Let $N = |\mathcal{N}|$. Let $\bar{\mu}$ denote the average daily demand across retailers: $\bar{\mu} = N^{-1}\sum_{i=1}^{N}\mu_i$. We assume that the distribution of demand across retailers, in expected value, follows a Pareto law:

$$\mu_i = \alpha_D^{i-1}\mu_1, \tag{A.1}$$

for some $\alpha_D \le 1$, $i \in \mathcal{N}$. When $\alpha_D = 1$, the retailers are identical and $\mu_1 = \bar{\mu}$. When $\alpha_D < 1$, summing all μ_i and rearranging terms yield

$$\mu_1 = \frac{N\bar{\mu}(1-\alpha_D)}{1-\alpha_D^N}, \tag{A.2}$$

so the parameters $\bar{\mu}$ and α_D are sufficient to generate the complete vector of daily demands, $(\mu_1, \ldots, \mu_N)$. We seek to create demand distributions for which some prespecified fraction, β_D, of total demand is concentrated in the largest 20% of retailers. That is, we look to create distributions satisfying the relation

$$\frac{\sum_{i=1}^{m}\mu_i}{\sum_{i=1}^{N}\mu_i} = \beta_D,$$

where m is approximately 20% of N. This reduces to finding a value of α_D, called $\alpha_D(\beta_D)$, which satisfies

$$\lim_{\alpha\to\alpha_D(\beta_D)} \frac{1-\alpha^{\lceil 0.2N\rceil}}{1-\alpha^N} = \beta_D.$$

This equation does not have an analytical solution but can be solved easily by a search procedure. Observe that $\beta_D = 0.8$ will approximate the common 80-20 Pareto distribution and $\beta_D = 0.2$ will yield the equalized distribution, $\alpha_D(\beta_D) = 1$ (the limit of the left-hand side, when $\alpha_D(\beta_D) = 1$, is 0.2 by L'Hôpital's Rule). As a result, the two parameters $\bar{\mu}$ and β_D are sufficient to generate the complete vector of daily demand means. We refer to β_D as the Pareto demand-shaping parameter.

Let ψ denote the coefficient of variation of daily demand of the *smallest* retailer. The standard deviation of daily demand of the smallest retailer is therefore $\psi\mu_N$. The variance of demand of the smallest retailer is $(\psi\mu_N)^2$. Assume that the variance of daily demand of retailer i is given by $(\mu_i/\mu_N)(\psi\mu_N)^2$. For example, if a retailer has twice the rate of daily demand of retailer N (that is, $\mu_i = 2\mu_N$), then its demand behaves like the sum of two independent demand streams, each with the standard deviation of the smallest retailer. Let σ_i denote the standard deviation of daily demand for retailer i. According to the previous logic, we have

$$\sigma_i = \psi\sqrt{\mu_i\mu_N}, \tag{A.3}$$

$i \in \mathcal{N}$. Consequently, given the vector of demand rates, the parameter ψ is sufficient to generate reasonable choices for daily demand standard deviations.

Let T denote the number of allocation periods in the problem. Let l_t denote the length (number of days) of allocation period t. We assume that the periods are nonincreasing in length: $l_1 \ge l_2 \ge \cdots \ge l_T$. In parallel fashion to the way we generate the mean daily demand parameters, we assume we are given the average number of days per period, $\bar{l}$, and a Pareto period length shaping parameter, β_L. From these, we determine a value of α_L, called $\alpha_L(\beta_L)$, which solves

$$\lim_{\alpha\to\alpha_L(\beta_L)} \frac{1-\alpha^{\lceil 0.2T\rceil}}{1-\alpha^T} = \beta_L$$

and the length of the first period, l_1, using

$$l_1 = \frac{T\bar{l}(1-\alpha_L(\beta_L))}{1-\alpha_L(\beta_L)^T}.$$

The remaining period lengths are given by $l_t = \alpha_L(\beta_L)^{t-1}l_1$, $t = 2, 3, \ldots, T$. Setting $\beta_L = 0.2$ will result in period lengths being identical. We do not require period lengths to be integer because the period length is used only to scale demand rates. Combining the daily demand rates with the period lengths permits us to generate the period demand rates required by the optimization model: $\mu_{it} = l_t\mu_i, i \in \mathcal{N}, t = 1, \ldots, T$. Similarly, assuming that daily demands are independent and identically distributed over time, the standard deviations of period demands are given by $\sigma_{it} = \sqrt{l_t}\sigma_i$, $i \in \mathcal{N}$, $t = 1, \ldots, T$. In summary, using this test case model, the demand parameters for the optimization model can be generated from the following seven test parameters: N, $\bar{\mu}$, β_D, ψ, T, $\bar{l}$, and β_L.

Considering the initial conditions for the optimization model, we assume that initial inventories at all retailers are zero in all test cases: $v_i = 0$, $i \in \mathcal{N}$. This is equivalent to assuming that the delivery of stock that initiates the allocation cycle is sufficient to raise all retailers to their ideal target inventories and that no retailer has more than its ideal target level. The ideals are computed relative to the total amount of inventory in the system. So, for simplicity, we assume all inventory in the system is initially concentrated at the central warehouse. We express the initial system reserve stock, v_0, in terms of a parameter γ, the safety stock factor under the assumption that all demand is concentrated in a single location and time period:

$$v_0 = T\bar{l}N\bar{\mu} + \gamma\sqrt{\sum_{t=1}^{T} l_t \sum_{i\in\mathcal{N}} \sigma_i^2}.$$

That is, γ is the number of standard deviations of total system demand to hold as safety stock. Assuming that normal operation of the system would result in at least two standard deviations of total system demand, we assume that $\gamma \le 1$

would represent a system under stress and $\gamma \geq 3$ a system flush with stock. Substituting sequentially for σ_i, μ_N, and μ_1 from (A.3), (A.1), and (A.2), respectively, yields

$$v_0 = T\bar{l}N\bar{\mu} + \gamma\psi N\bar{\mu} \lim_{\alpha \to \alpha_D(\beta_D)} \sqrt{T\bar{l}\alpha^{N-1}\frac{1-\alpha}{1-\alpha^N}}.$$

For the special case of identical retailers, $\beta_D = 0.2$ and $\alpha_D(\beta_D) = 1$, this reduces to $v_0 = T\bar{l}N\bar{\mu} + \gamma\psi\bar{\mu}\sqrt{T\bar{l}N}$.

For the correlated demand case, the correlation between retailer demands within any given period is modeled by a single parameter ρ. We further assume that demands are independent across periods. Given a valid correlation parameter ρ, we are able to construct a covariance matrix Σ_t, one for each period t, whose diagonal entries are σ_{it}^2 and off-diagonal entries are $\rho\sigma_{it}\sigma_{jt}$. However, not all $\rho \in [-1, 1]$ are valid. In fact, if $\mu_{it} = \mu_t$ and $\sigma_{it} = \sigma_t$ for all $i \in \mathcal{N}$, then it can be proved that the lower bound for a valid ρ is $1/(N-1)$. For a valid covariance matrix Σ_t, there exists a decomposition $\Sigma_t = C_t C_t^T$, and the demand vector is expressed as $d_t = \mu_t + C_t\varepsilon_t$ in the modeling.

In the simulation, the random demand is assumed to follow the log-normal distribution and is given by (52). We choose parameters m_{it} and p_{ijt} so that

$$\mathrm{E}(\tilde{d}_{it}) = e^{m_{it} + (1/2)\sum_{j\in\mathcal{N}} p_{ijt}^2} = \mu_{it}, \quad \forall\, i \in \mathcal{N}; \tag{A.4}$$

$$\mathrm{Var}(\tilde{d}_{it}) = (e^{\sum_{j\in\mathcal{N}} p_{ijt}^2} - 1)e^{2m_{it} + \sum_{j\in\mathcal{N}} p_{ijt}^2} = \sigma_{it}^2, \quad \forall\, i \in \mathcal{N}; \tag{A.5}$$

$$\mathrm{Cov}(\tilde{d}_{it}, \tilde{d}_{kt}) = \mu_{it}\mu_{kt}(e^{\sum_{j\in\mathcal{N}} p_{ijt}p_{kjt}} - 1) = \rho\sigma_{it}\sigma_{kt}, \quad \forall\, i, k \in \mathcal{N}, i \neq k, \tag{A.6}$$

for any given period t. Squaring (A.4) and plugging it into (A.5) to get

$$\sum_{j\in\mathcal{N}} p_{ijt}^2 = \log\left(\frac{\sigma_{it}^2}{\mu_{it}^2} + 1\right), \quad \forall\, i \in \mathcal{N}, \tag{A.7}$$

and rearranging terms of (A.6) to obtain

$$\sum_{j\in\mathcal{N}} p_{ijt}p_{kjt} = \log\left(\frac{\rho\sigma_{it}\sigma_{kt}}{\mu_{it}\mu_{kt}} + 1\right), \quad \forall\, i, k \in \mathcal{N}, i \neq k. \tag{A.8}$$

Denote by P_t the matrix whose entries are p_{ijt} and let $Q_t = P_t P_t^T$. Then, (A.7) and (A.8) are sufficient to form the matrix Q_t as they are the diagonal and off-diagonal elements of Q_t, respectively. To apply the decomposition on Q_t to get P_t, we require that Q_t be positive semidefinite. This requirement, together with the positiveness of $\rho\sigma_{it}\sigma_{kt}/(\mu_{it}\mu_{kt}) + 1$ in (A.8), imposes another constraint on ρ. The lower bound on ρ is thus determined by the positive semidefiniteness of Σ_t and Q_t.

The optimization algorithm can be tuned by choosing different values of δ, the uncertainty set parameter. If the explicit uncertainty set (5) is implemented, then two natural values to consider would be $\delta = 2$ and $\delta = \gamma$. If the implicit uncertainty set (3) is implemented, then we slightly adjust it by adding the lower bound $-\delta_0$ on the normalized demand. Specifically, the *modified implicit risk pooling uncertainty set* is given by

$$U(\delta) = \left\{ \varepsilon = \varepsilon^+ - \varepsilon^- : \begin{array}{c} 0 \leq \varepsilon_{it}^+ \leq \delta_0, \ \forall\,(i, t) \\ 0 \leq \varepsilon_{it}^- \leq \delta_0, \ \forall\,(i, t) \\ \displaystyle\sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}} c_{ijt}\varepsilon_{jt}^+ \leq \delta_1, \ t = 1, \ldots, T \end{array} \right\}, \tag{A.9}$$

for $\delta = \{\delta_0, \delta_1\}$. Note that the coefficient c_{ijt} is allowed to be negative for all (i, j, t), so we are able to deal with negatively correlated demands. Denote by Z_t^+ the left-hand side of the constraint on the aggregate scaled demand in (A.9). We want the probability that the constraint is satisfied to be at least $1 - \alpha$, where $\alpha \in (0, 1)$ is called the *probability of excess demand*. Since we restrict attention to the two-period case, this relation can be expressed simply as $\mathbb{P}\{Z_1^+ \leq \delta_1\} \geq 1 - \alpha$. Determining the exact value of δ_1 requires the knowledge of the distribution of Z_1^+, which is difficult to be interpreted analytically. Therefore, we use the same set of 10,000 sample vectors $\tilde{\varepsilon}$ for the simulation to estimate an empirical distribution, $\tilde{F}(\cdot)$, of Z_1^+. We choose δ_1 to be the $100(1-\alpha)$-th percentile of the empirical distribution; that is, $\delta_1 = \tilde{F}^{-1}(1-\alpha)$. Thus, we can choose different δ_1 by varying α. This modified implicit uncertainty set captures risk pooling because the variance of Z_t^+ involves $N = |\mathcal{N}|$: $\mathrm{Var}(Z_t^+) = \sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{N}} c_{ijt}^2 \mathrm{Var}(\varepsilon_{jt}^+) = \sigma_+^2 \sum_{i\in\mathcal{N}} \sigma_{it}^2$, where $\sigma_+^2 = \mathrm{Var}(\varepsilon_{jt}^+)$. The last equation follows from the fact that $\sum_{j\in\mathcal{N}} c_{ijt}^2 = (C_t C_t^T)_{ii} = (\Sigma_t)_{ii} = \sigma_{it}^2$.

The optimization algorithm can be further tuned by choosing different values for the time period weights, w_t, $t = 1, \ldots, T$. We consider the following time period weighting scheme: $w_t = \theta^{t-1}$, $t = 1, \ldots, T$. Hence, $w_1 = 1$, $\theta = 1$ generates equal weights, $\theta > 1$ generates weights that increase over time, and $\theta < 1$ generates weights that decrease over time. We do not consider cases in which the weights differ by retailer: $w_{it} = w_t$ for all $i \in \mathcal{N}$. Table A.1 summarizes the parameters used to generate test cases.

Table A.1. Parameter Ranges for Generating Test Cases

Parameter name	Symbol	Value range
Number of retailers	N	4, 6, 8
Number of periods	T	2, 3
Mean daily demand	$\bar{\mu}$	2.5, 5
Mean days per period	$\bar{l}$	5
Coefficient of variation	ψ	0.5, 1, 1.5, 2, 2.5, 3
Pareto demand shape	β_D	0.2, 0.5, 0.8
Pareto period-length shape	β_L	0.2, 0.8
Safety stock factor	γ	1, 1.5, 2, 2.5
Demand correlation	ρ	−0.2, −0.15, −0.1, −0.05, 0.1, 0.2, ..., 1.0
Probability of excess demand	α	0.01, 0.02, 0.03, 0.05, 0.1, 0.2, 0.5
Backorder weight growth factor	θ	1, 2, 4

Appendix B. Proofs

The proofs for all results in the main body of the paper are presented here.

B.1. Proof of Lemma 1

If $z \in Z$, then order the elements of z from largest, $z_{[1]}$, to smallest, $z_{[|\mathcal{N}|]}$. Let $\alpha_n = z_{[n]}$ and $\beta_{ni} = (z_i - \alpha_n)^+$, $i \in \mathcal{N}$. Let I be the set of the n largest elements of z. It is clear that $\beta_{ni} = 0$ for all $i \notin I$. Since $z \in Z$ and $|I| = n$, we have

$$M_n \geq \sum_{i \in I} z_i = n\alpha_n + \sum_{i \in I} \beta_{ni} = n\alpha_n + \sum_{i \in \mathcal{N}} \beta_{ni},$$

establishing that $z \in Z'$. Now, suppose $z \in Z'$ and consider any subset I, $I \subseteq \mathcal{N}$ with $|I| = n$. Let (α_n, β_n) be any solution to the conditions in Z'. Then, since $\beta_{ni} \geq 0$ for all i:

$$M_n \geq n\alpha_n + \sum_{i \in \mathcal{N}} \beta_{ni} \geq \sum_{i \in I} (\alpha_n + \beta_{ni}) \geq \sum_{i \in I} z_i.$$

Hence, $z \in Z$. □

B.2. Proof of Proposition 1

There are constraints in $U(\delta, \bar{n})$ for all combinations of $i \in \mathcal{N}$, $n \in \{1, 2, \ldots, \bar{n}\}$, and $t \in T$. □

B.3. Proof of Lemma 2

By definition, $x_{it}(\varepsilon) \geq y_{it} - v_i + \sum_{t'=1}^{t-1} (\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} - x_{it'}(\varepsilon))$ with equality holding if $x_{it}(\varepsilon) > 0$. Rearranging terms yields the result. □

B.4. Proof of Proposition 2

Both the linear program and the expression for $S(y, \varepsilon)$ are separable by retailer. Let $S_i(y_i, \varepsilon) = \sum_{t=1}^{T} x_{it}(\varepsilon)$, where $x_{it}(\varepsilon)$ is given by (6). Note that the solution (6) is feasible for the linear program, so $S_i(y_i, \varepsilon)$ is an upper bound on the linear programming solution restricted to retailer i. Let x_i^* denote an optimal solution to the linear program. Let t be the first period for which $x_{it}^* \neq x_{it}(\varepsilon)$. Since $x_{it}(\varepsilon)$ is constructed to be the minimum nonnegative shipment quantity to achieve the target inventory level, we must have $x_{it}^* > x_{it}(\varepsilon)$. Define a new solution x_i' reducing the tth component of x_i^* by $x_{it}^* - x_{it}(\varepsilon)$ and increasing the $t + 1^{st}$ component of x_i^* by the same quantity. It is easily verified that this solution, x_i', is also optimal but now differs from $x_i(\varepsilon)$ starting from the $(t+1)^{st}$ component. Continuing in this manner, we find that $x_i(\varepsilon)$ is optimal for the linear program and so $S_i(y_i, \varepsilon)$ must equal the optimal objective value of the linear program restricted to retailer i. □

B.5. Proof of Corollary 1

The result follows easily by duality. □

B.6. Proof of Proposition 3

If $t_i^*(\varepsilon) = 0$, then no objective coefficient in (8) is positive for retailer i. In this case, the optimal solution will be to set $\pi_{it}^*(\varepsilon) = 0$ for all $t = 1, 2, \ldots, T$. If $t_i^*(\varepsilon) > 0$, then a marginal analysis argument can be used to show that, barring ties, any solution in which more than one $\pi_{it} > 0$ for any i can be improved by increasing the π_{it} for the period t with the largest objective coefficient, $t_i^*(\varepsilon)$, and decreasing the others. In the case of ties, the solution is unchanged by increasing one of the tied variables and decreasing the others. Hence, for each $i \in \mathcal{N}$, an optimal solution consists of setting one $\pi_{it} = 1$ and the rest equal to zero. □

B.7. Proof of Corollary 2

The optimal solution given in Proposition 3 satisfies $\pi \in \Pi$. □

B.8. Proof of Proposition 4

By (7)

$$\sum_{t'=1}^{t} x_{it'}(\varepsilon) \geq y_{it} - v_i + \sum_{t'=1}^{t-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} \right)$$

with equality holding if $x_{it}(\varepsilon) > 0$. Furthermore, $\sum_{t'=1}^{t} x_{it'}(\varepsilon)$ is nondecreasing in t. Suppose $\sum_{t=1}^{T} x_{it}(\varepsilon) > 0$. Let $\tau = \tau_i(\varepsilon)$ denote the last period in which $x_{it}(\varepsilon) > 0$. It follows that

$$\sum_{t'=1}^{\tau} x_{it'}(\varepsilon) = S_i(y_i, \varepsilon) = \max_{t=1,2,\ldots,T} \left(y_{it} - v_i + \sum_{t'=1}^{t-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} \right) \right),$$

where $S_i(y_i, \varepsilon)$ is the required shipment for location i. Note that we have dropped the positive part function under the assumption that $x_{it}(\varepsilon) > 0$. It follows that for this value of τ, we have

$$y_{i\tau} - v_i + \sum_{t'=1}^{\tau-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} \right) = \max_{t=1,2,\ldots,T} \left(y_{it} - v_i + \sum_{t'=1}^{t-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} \right) \right).$$

That is, τ is a candidate for $t_i^*(\varepsilon)$. In the case of ties, we defined $t_i^*(\varepsilon)$ to be the earliest period t optimizing the right-hand side. Consequently, $t_i^*(\varepsilon) \leq \tau$. Suppose, by way of contradiction, that $t_i^*(\varepsilon) < \tau$. In that case, we would have, by (7),

$$\sum_{t'=1}^{t_i^*(\varepsilon)} x_{it'}(\varepsilon) \geq y_{it_i^*(\varepsilon)} - v_i + \sum_{t'=1}^{t_i^*(\varepsilon)-1} \left(\mu_{it'} + \sum_{j \in \mathcal{N}} c_{ijt'} \varepsilon_{jt'} \right) = \sum_{t'=1}^{\tau} x_{it'}(\varepsilon),$$

leading to the conclusion

$$\sum_{t'=t_i^*(\varepsilon)+1}^{\tau} x_{it'}(\varepsilon) = 0,$$

which contradicts the assumption that $x_{i\tau}(\varepsilon) > 0$. Consequently, $t_i^*(\varepsilon) = \tau$. On the other hand, if $\sum_{t=1}^{T} x_{it}(\varepsilon) = 0$, then $\tau_i(\varepsilon) = 0$, by definition, and the result holds trivially. □

B.9. Proof of Corollary 3

This follows directly from Propositions 3 and 4. □

B.10. Proof of Lemma 3

For every $\varepsilon \in U$, if x is feasible for y, then it is also feasible for $y' \leq y$. Hence, $S(y', \varepsilon) \leq S(y, \varepsilon)$ for all $\varepsilon \in U$. □

B.11. Proof of Proposition 5

This is a straightforward consequence of the fact that $S(y)$ is nondecreasing in y. That is, if (y^0, B^*) is any optimal solution to (11) with $y_{it}^0 > \bar{d}_{it}(\delta) - w_{it}^{-1} B_t^*$, then it can be shown that y^* as defined by (15) satisfies $y^* < y^0$ and, hence, $S(y^*) \leq S(y^0) \leq v_0$. The combination (y^*, B^*) is therefore feasible and yields the same objective value as the optimal solution (y^0, B^*). □

B.12. Proof of Proposition 6

Denote an optimal solution to the MILP using the superscript "$*$." The binary variables u_{it}^* act as selectors. For each $i \in \mathcal{N}$, there will be exactly one $u_{it}^* = 0$, and hence if $t > 0$, then $S_i^* = y_{it} - v_i + \sum_{t'=1}^{t-1}(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}^*)$ for that value of t. If $u_{i0} = 0$, then $S_i^* = 0$. It follows that

$$S_i^* = \max_{t=1,2,\ldots,T}\left(y_{it} - v_i + \sum_{t'=1}^{t-1}\left(\mu_{it'} + \sum_{j\in\mathcal{N}} c_{ijt'}\varepsilon_{jt'}^*\right)\right)^+$$

and that $S(y) = \sum_{i\in\mathcal{N}} S_i^*$. □

B.13. Proof of Corollary 4

This is a direct application of Corollary 3. □

B.14. Proof of Proposition 7

Initially, i.e., on iteration $k = 0$, we set $S_{\pi^0}(\cdot) \equiv 0$. That is, we ignore the shipment constraint on the first iteration. Let (y^k, B^k) denote the optimal solution to the master problem on iteration k. The subproblem at the kth iteration is to solve $\max_{\pi\in\Pi} S_\pi(y^k)$. The optimal selector solution to the kth subproblem is denoted by π^k, given by $\pi^k = 1 - u^k$, where $(\varepsilon^k, u^k, S(y^k))$ is the optimal solution to (22)–(27). It follows that if $S_{\pi^k}(y^k) \le v_0$, then (y^k, B^k) is an optimal solution to the original problem. In this case, the algorithm is terminated. Otherwise, $\pi^k \notin \Pi^k$. We then express $S_{\pi^k}(y)$ in terms of y with ε^k and π^k using (17) and add the constraint $S_{\pi^k}(y) \le v_0$ to the master problem. We replace Π^k with $\Pi^k \cup \{\pi^k\}$ and repeat the process. Since Π is finite and a new element, π^k, is identified on each iteration that fails to find an optimal solution, the algorithm must terminate with an optimal solution in a finite number of iterations. □

B.15. Proof of Proposition 8

We prove this proposition by first showing some lemmas. Let us first consider the special case presented at the beginning of Section 3.2 but drop temporarily the assumption of identical retailers. Let $I_1 = \{i: \pi_{i1} = 1\}$ and $I_2 = \{i: \pi_{i2} = 1\}$. By Corollary 3, I_1 is the set of retailers receiving their last shipment in the first period, and I_2 is the set of retailers receiving a shipment in the second period. In this case, the shipment requirement (17) simplifies to the following:

$$S_\pi(y) = \sum_{i\in I_1} y_{i1} + \sum_{i\in I_2} y_{i2} + \sum_{i\in I_2}\mu_{i1} + \max_{\varepsilon\in U}\sum_{i\in I_2}\sigma_{i1}\varepsilon_{i1}.$$

Consider the last term:

$$\max_{\varepsilon\in U}\sum_{i\in I_2}\sigma_{i1}\varepsilon_{i1}. \tag{B.1}$$

The solution to this optimization problem depends on the form of the uncertainty set. If the implicit risk pooling uncertainty set is used, the solution is quite simple.

Lemma 4. *If $U = U(\delta)$ is given by (3), then*

$$\max_{\varepsilon\in U(\delta)}\sum_{i\in I_2}\sigma_{i1}\varepsilon_{i1} = \min\left\{\delta_1, \left(\sum_{i\in I_2}\sigma_{i1}\right)\delta_0\right\}. \tag{B.2}$$

Proof. Since demands are uncorrelated, this optimization problem can be simplified as

$$\begin{aligned}\max_{\varepsilon_{i1}} \quad & \sum_{i\in I_2}\sigma_{i1}\varepsilon_{i1}\\ \text{subject to} \quad & \varepsilon_{i1} \le \delta_0,\\ & \sum_{i\in\mathcal{N}}\sigma_{i1}\varepsilon_{i1} \le \delta_1.\end{aligned}$$

If $(\sum_{i\in I_2}\sigma_{i1})\delta_0 \le \delta_1$, then $\varepsilon_{i1} = \delta_0$ for all $i \in I_2$ is clearly an optimal solution. In this case, the optimal objective value is $(\sum_{i\in I_2}\sigma_{i1})\delta_0$. Otherwise, the constraint $\sum_{i\in\mathcal{N}}\sigma_{i1}\varepsilon_{i1} \le \delta_1$ must be binding, and thus the objective value is δ_1, because otherwise we can always optimize the objective by increasing some ε_{i1} for $i \in I_2$ and decreasing some ε_{i1} for $i \notin I_2$ by the same amount. Therefore, the optimal objective value is given by $\min\{\delta_1, (\sum_{i\in I_2}\sigma_{i1})\delta_0\}$. □

On the other hand, if the explicit risk pooling uncertainty set is used, the situation is more complicated. Let $[j]$ index the retailer with the jth largest value of σ_{i1} for $i \in I_2$, and let $n_2 = |I_2|$. Then, assuming $n_2 > 0$, $\sigma_{[1]1} \ge \sigma_{[2]1} \ge \cdots \ge \sigma_{[n_2]1}$. The optimization problem (B.1) then can be written as

$$\max \quad \sum_{j=1}^{n_2}\sigma_{[j]1}\varepsilon_{[j]1} \tag{B.3}$$

$$\text{subject to} \quad \sum_{i\in I}\varepsilon_{i1} \le \sqrt{|I|}\delta, \quad \forall\, I \subseteq I_2, |I| \le \bar{n}. \tag{B.4}$$

There are multiple possible optimal extreme points to this problem, depending on the values of $\bar{n}$ and σ_{i1}, $i \in I$. Let $\bar{n}_2 = \min(\bar{n}, n_2)$, the largest value of $|I|$ in (B.4). The following partial characterization of an optimal solution is possible.

Lemma 5. *Assume $U = U(\delta, \bar{n})$ is given by (5).*

(a) *If $n_2 > 0$, then there is an optimal solution to (B.3)–(B.4) such that*

$$\varepsilon_{[1]1} \ge \varepsilon_{[2]1} \ge \cdots \ge \varepsilon_{[n_2]1} \ge \delta\left(\sqrt{\bar{n}_2} - \sqrt{\bar{n}_2 - 1}\right) > 0.$$

(b) *If $\bar{n} \ge n_2$, an optimal solution is to set*

$$\varepsilon_{[j]1} = \delta\left(\sqrt{j} - \sqrt{j-1}\right) \tag{B.5}$$

for $j = 1, 2, \ldots, n_2$.

(c) *If $\sigma_{i1} = \sigma_1$ for all $i \in \mathcal{N}$, then an optimal solution is given by*

$$\varepsilon_{i1} = \frac{\delta}{\sqrt{\bar{n}_2}} \tag{B.6}$$

for $i = 1, 2, \ldots, n_2$.

Proof. (a) It is easily seen that if $\varepsilon \in U(\delta, \bar{n})$ then ε with any permutation of retailer indices is also feasible. It follows from the ordering of objective coefficients that an optimal solution will satisfy $\varepsilon_{[1]1} \ge \varepsilon_{[2]1} \ge \cdots \ge \varepsilon_{[n_2]1}$. Let I with $|I| = j$ identify any binding constraint involving $\varepsilon_{[n_2]1}$. $\sum_{i\in I}\varepsilon_{[i]1} = \sqrt{j}\delta$. There must be at least one such constraint in an optimal solution, given that $\sigma_{[n_2]1} > 0$. Then,

$$\begin{aligned}\varepsilon_{[n_2]1} &= \sqrt{j}\delta - \sum_{i\in I\setminus\{n_2\}}\varepsilon_{[i]1} \ge \sqrt{j}\delta - \sqrt{j-1}\delta\\ &\ge \delta\left(\sqrt{\bar{n}_2} - \sqrt{\bar{n}_2 - 1}\right) > 0,\end{aligned}$$

where the first inequality results from the constraint on sets of size $j-1$, and the remaining inequalities result from the strict concavity of $\sqrt{j}$.

(b) The proposed solution satisfies the following constraints of (B.4) with equality:

$$\begin{aligned}\varepsilon_{[1]1} &= \delta\\ \varepsilon_{[1]1} + \varepsilon_{[2]1} &= \sqrt{2}\delta\\ &\vdots\\ \varepsilon_{[1]1} + \varepsilon_{[2]1} + \cdots + \varepsilon_{[n_2]1} &= \sqrt{n_2}\delta.\end{aligned}$$

These correspond to subsets $I = \{[1]\}, \{[1],[2]\}, \ldots, \{[1],[2], \ldots,[n_2]\}$ of I_2, respectively. Denote the dual variables for these constraints by $\eta_{\{[1]\}}, \eta_{\{[1],[2]\}}, \ldots, \eta_{\{[1],[2],\ldots,[n_2]\}}$, respectively. Set all other dual variables, η_I, $I \subseteq I_2$, to zero and solve the dual equations to find

$$\begin{aligned} \eta_{\{[1],[2],\ldots,[n_2]\}} &= \sigma_{[n_2]1} \\ \eta_{\{[1],[2],\ldots,[n_2-1]\}} &= \sigma_{[n_2-1]1} - \sigma_{[n_2]1} \\ &\vdots \\ \eta_{\{[1]\}} &= \sigma_{[1]1} - \sigma_{[2]1}. \end{aligned}$$

Since the sequence $\sigma_{[1]1}, \sigma_{[2]1}, \ldots, \sigma_{[n_2]1}$ is nonincreasing, $\eta \geq 0$ and so η is dual feasible. Thus, the primal objective is given by

$$\begin{aligned} \sum_{j=1}^{n_2} \sigma_{[j]1}\varepsilon_{[j]1} &= \sum_{j=1}^{n_2} \sigma_{[j]1}\delta(\sqrt{j} - \sqrt{j-1}) \\ &= \sum_{j=1}^{n_2-1} (\sigma_{[j]1} - \sigma_{[j+1]1})\sqrt{j}\delta + \sigma_{[n_2]1}\sqrt{n_2}\delta. \end{aligned}$$

The dual objective is given by

$$\begin{aligned} \sum_{I \subseteq I_2} \eta_I \sqrt{|I|}\delta &= \sum_{n=1}^{n_2} \eta_{\{[1],[2],\ldots,[n]\}}\sqrt{n}\delta \\ &= \sum_{n=1}^{n_2-1} (\sigma_{[n]1} - \sigma_{[n+1]1})\sqrt{n}\delta + \sigma_{[n_2]1}\sqrt{n_2}\delta. \end{aligned}$$

By duality, we have found an optimal solution.

(c) It is easily verified that $\varepsilon_{i1} = \delta/\sqrt{\bar{n}_2}$ is primal feasible. Consider the set of equations $\sum_{i \in I} \varepsilon_{i1} = \sqrt{\bar{n}_2}\delta$ for all $I \subseteq \mathcal{N}$, $|I| = \bar{n}_2$. The number of equations in this set is the number of ways to choose $\bar{n}_2$ items from a collection of n_2 items; i.e., $\binom{n_2}{\bar{n}_2}$. Let $\kappa = \binom{n_2}{\bar{n}_2}$ and let $k = 1, 2, \ldots, \kappa$ number these equations. Let $e_{ik} = 1$ if retailer i is present in the set I represented by the kth equation, and zero otherwise. Then the equations can be written as $\sum_{i=1}^{n_2} e_{ik}\varepsilon_{i1} = \sqrt{\bar{n}_2}\delta$, for $k = 1, 2, \ldots, \kappa$. Let η_k denote the dual multiplier associated with constraint k, $k = 1, 2, \ldots, \kappa$, and consider the corresponding dual constraints of the form: $\sum_{k=1}^{\kappa} \eta_k e_{ik} = \sigma_1$, for $i = 1, 2, \ldots, n_2$. By the definition of e_{ik}, the sum $\sum_{k=1}^{\kappa} e_{ik}$ is the number of primal equations with $|I| = \bar{n}_2$ containing ε_{i1} and is therefore the number of ways to choose $\bar{n}_2 - 1$ items from a collection of $n_2 - 1$ items; i.e., $\sum_{k=1}^{\kappa} e_{ik} = \binom{n_2-1}{\bar{n}_2-1}$. It follows that the dual solution

$$\eta_k = \frac{\sigma_1}{\binom{n_2-1}{\bar{n}_2-1}}$$

is feasible. We extend the vector (η_k) to a vector, η, over all constraints in the primal problem by setting the dual multipliers equal to zero outside of the set of equations considered in the solution for (η_k). This vector is dual feasible. By strong duality, we conclude that this solution pair is optimal for the primal and dual problems, since

$$\sum_{k=1}^{\kappa} \eta_k \sqrt{\bar{n}_2}\delta = \binom{n_2}{\bar{n}_2} \frac{\sigma_1\sqrt{\bar{n}_2}\delta}{\binom{n_2-1}{\bar{n}_2-1}} = \sum_{i=1}^{n_2} \sigma_1\varepsilon_{i1}. \quad \square$$

Lemma 6. *If $\sigma_{i1} = \sigma_1$ for all $i \in \mathcal{N}$, and $n_2 > 0$, then*

$$\max_{\varepsilon \in U(\delta,\bar{n})} \sum_{i \in I_2} \sigma_{i1}\varepsilon_{i1} = \frac{n_2}{\sqrt{\bar{n}_2}}\sigma_1\delta.$$

Proof. Plugging (B.6) into (B.1) yields the result. $\square$

One implication of this lemma is simply that if $\sigma_{i1} = \sigma_1$ for all $i \in \mathcal{N}$, then

$$\max_{\varepsilon \in U(\delta,\bar{n})} \sum_{i \in I_2} \sigma_{i1}\varepsilon_{i1} > \max_{\varepsilon \in U(\delta)} \sum_{i \in I_2} \sigma_{i1}\varepsilon_{i1}$$

whenever $\bar{n} < n_2$. That is, the relaxation in the uncertainty set leads to strictly larger partial sums of worst-case demand whenever the number of retailers in the allocation exceeds $\bar{n}$.

Lemma 7. *Under either risk pooling uncertainty set,* (3) *or* (5), *if $y \geq 0$, $\mu_{i1} > 0$, and π is an optimal solution to Problem* (16) *for the given y, then $\pi_{i1} + \pi_{i2} = 1$ for each $i \in \mathcal{N}$. Furthermore, $I_1 \cup I_2 = \mathcal{N}$ and $I_1 \cap I_2 = \varnothing$.*

Proof. By (7), $\sum_{t'=1}^{2} x_{it'}(\varepsilon) \geq y_{i2} + \mu_{i1} + \sigma_{i1}\varepsilon_{i1}$ for each $i \in \mathcal{N}$. By Proposition 5, the optimal ε is positive. Under the conditions of the corollary, therefore, for each $i \in \mathcal{N}$, we have $\sum_{t'=1}^{2} x_{it'}(\varepsilon) > 0$. That is, the optimal ε requires that the shipments to each retailer be strictly positive in at least one of the two periods. Consequently, by Corollary 3, exactly one of π_{i1} and π_{i2} must equal one, for each $i \in \mathcal{N}$. It follows directly, by the definition of I_1 and I_2, that $I_1 \cup I_2 = \mathcal{N}$ and $I_1 \cap I_2 = \varnothing$. $\square$

We now proceed to prove the proposition by imposing the assumption of identical retailers. Recall that we have assumed $y \geq 0$ in the optimal solution, so by Lemma 7, we need consider only $\pi \in \Pi$ with $\pi_{i1} + \pi_{i2} = 1$, $\forall\, i$. If $U = U(\delta)$ is given by (5), then by Lemma 6, the shipment requirement is equivalent to

$$\max_{I_2 \subseteq \mathcal{N}} \left\{ \sum_{i \in \mathcal{N}\setminus I_2} y_{i1} + \sum_{i \in I_2} y_{i2} + |I_2|\mu_1 + \frac{|I_2|}{\sqrt{\bar{n} \wedge |I_2|}}\sigma_1\delta \right\} \leq v_0, \qquad \text{(B.7)}$$

where $a \wedge b = \min(a, b)$ and the ratio $\frac{0}{0}$ is taken to be 0. Let $n = |I_2|$, the number of retailers receiving a shipment in the second period. Let y_1 and y_2 denote the common target inventory levels for retailers in periods 1 and 2, respectively. The two-period shipment constraint (B.7) simplifies to (32). In particular, knowledge of the set of retailers receiving shipments in period 2 collapses to knowledge of just n, the number of retailers receiving a shipment in period 2. If $U = U(\delta)$ is given by (3), then by Lemma 4 and the assumption of identical retailers, (33) can be easily obtained. This completes the proof of the proposition. $\square$

B.16. Proof of Proposition 9

Let $M_1(n,u) = \sqrt{n}\sigma_1\delta - nu$ and $M_2(n,u) = (n/\sqrt{\bar{n}})\sigma_1\delta - nu$. Let $n' = \operatorname{argmax}_{1 \leq n \leq \bar{n}} M_1(n,u)$ and $n'' = \operatorname{argmax}_{\bar{n} \leq n \leq N} M_2(n,u)$. Let $M_1^*(u) = M_1(n',u)$ and $M_2^*(u) = M_2(n'',u)$. Then, the unconstrained optimizer $\hat{n}$ for n in $M_1(n,u)$ is $\hat{n} = (\sigma_1\delta/(2u))^2$. By concavity,

$$n' = \begin{cases} 1 & \text{if } u > \dfrac{1}{2}\sigma_1\delta, \\ \bar{n} & \text{if } u < \dfrac{1}{2\sqrt{\bar{n}}}\sigma_1\delta, \\ \hat{n} & \text{if otherwise.} \end{cases}$$

The second maximand, $M_2(n,u)$, is linear in n, so

$$n'' = \begin{cases} \bar{n} & \text{if } u > \dfrac{1}{\sqrt{\bar{n}}}\sigma_1\delta, \\ N & \text{if } u \leq \dfrac{1}{\sqrt{\bar{n}}}\sigma_1\delta. \end{cases}$$

Suppose $\bar{n} \geq 4$, then $(1/(2\sqrt{\bar{n}}))\sigma_1\delta \leq (1/\sqrt{\bar{n}})\sigma_1\delta \leq \frac{1}{2}\sigma_1\delta$. There are four regions of interest when comparing $M_1^*(u)$ and $M_2^*(u)$.

1. $u > \frac{1}{2}\sigma_1\delta$. In this region, $n' = 1$. Furthermore, $\frac{1}{2}\sigma_1\delta \geq (1/\sqrt{\bar{n}})\sigma_1\delta$, so $n'' = \bar{n}$. Observe that $M_1(\bar{n}, u) = M_2(\bar{n}, u)$ for all u, so the optimal value n^* of n is equal to 1.

2. $(1/\sqrt{\bar{n}})\sigma_1\delta \leq u \leq \frac{1}{2}\sigma_1\delta$. In this region, $n'' = \bar{n}$. Because $(1/(2\sqrt{\bar{n}}))\sigma_1\delta \leq (1/\sqrt{\bar{n}})\sigma_1\delta$, $n' = \hat{n}$. The two maximands equal each other at $n = \bar{n}$. Combining that with the concavity of the first maximand, so $n^* = \hat{n}$.

3. $u < (1/(2\sqrt{\bar{n}}))\sigma_1\delta$. In this region, $n' = \bar{n}$ and $n'' = N$. By continuity at $\bar{n}$, and since $M_2(n, u)$ is increasing in n for $n \geq \bar{n}$, we conclude that $n^* = N$ in this case.

4. $(1/(2\sqrt{\bar{n}}))\sigma_1\delta < u < (1/\sqrt{\bar{n}})\sigma_1\delta$. In this region, $n' = \hat{n}$ and $n'' = N$. Furthermore,

$$M_1^*(u) = M_1(\hat{n}, u) = \frac{(\sigma_1\delta)^2}{2u} - \left(\frac{\sigma_1\delta}{2u}\right)^2 u = \frac{(\sigma_1\delta)^2}{4u},$$
$$M_2^*(u) = M_2(N, u) = \frac{N}{\sqrt{\bar{n}}}\sigma_1\delta - Nu.$$

Consider the difference $M_1^*(u) - M_2^*(u)$ as a function of u. Noting that $(1/(2\sqrt{\bar{n}}))\sigma_1\delta < u$ in this region, we have

$$\frac{d(M_1^*(u) - M_2^*(u))}{du} > \frac{-(\sigma_1\delta)^2}{4((1/(2\sqrt{\bar{n}}))\sigma_1\delta)^2} + N = N - \bar{n} \geq 0.$$

Consequently, $M_1^*(u) - M_2^*(u)$ is strictly increasing over this region. Evaluated at $u = (1/(2\sqrt{\bar{n}}))\sigma_1\delta$, $M_2^*(u) = (N/(2\sqrt{\bar{n}}))\cdot\sigma_1\delta \geq (\bar{n}/(2\sqrt{\bar{n}}))\sigma_1\delta = M_1^*(u)$, while at $u = (1/\sqrt{\bar{n}})\sigma_1\delta$, $M_2^*(u) = 0 < \sqrt{\bar{n}}(\sigma_1\delta)/4 = M_1^*(u)$, so there must be a unique value ζ at which $M_1^*(\zeta) = M_2^*(\zeta)$. Rewriting this as a quadratic equation and solving it yield two real roots:

$$\zeta = \frac{\sigma_1\delta}{2N}\left\{\frac{N}{\sqrt{\bar{n}}} \pm \sqrt{\frac{N^2}{\bar{n}} - N}\right\}.$$

Denote them as ζ^+ and ζ^-, respectively. It follows that $\zeta^- < (\sigma_1\delta/(2N))(N/\sqrt{\bar{n}}) = \sigma_1\delta/(2\sqrt{\bar{n}})$, which falls outside of the region of interest. Therefore, the unique value of ζ must be

$$\zeta^+ = \frac{\sigma_1\delta}{2N}\left\{\frac{N}{\sqrt{\bar{n}}} + \sqrt{\left(\frac{N}{\sqrt{\bar{n}}}\right)^2 - N}\right\} = \frac{1}{2\sqrt{\bar{n}}}\sigma_1\delta\left\{1 + \sqrt{1 - \frac{\bar{n}}{N}}\right\}.$$

In summary, $n^* = N$ for $(1/(2\sqrt{\bar{n}}))\sigma_1\delta < u < \zeta^+$ and $n^* = \hat{n}$ for $\zeta^+ \leq u < (1/\sqrt{\bar{n}})\sigma_1\delta$ in this region.

Assembling the results leads to the proposition. □

B.17. Proof of Theorem 1

Let $y_2^*(y_1)$ denote the optimal target minimum inventory level in period 2 as a function of the target inventory level in period 1. The constraint (36) will be binding in any optimal solution to RR2. Let $Z = w_1y_1 + w_2y_2$, the objective function for RR2, and let $Z(y_1) = w_1y_1 + w_2y_2^*(y_1)$. We first prove the following lemma:

Lemma 8. *If $\bar{n} \geq 4$, the optimal target inventory level of period 2 is given by*

$$y_2^*(y_1) = \begin{cases} y_1 - \mu_1 + v_0 - Ny_1 - \sigma_1\delta \\ \qquad \text{if } v_0 - Ny_1 < \frac{1}{2}\sigma_1\delta, \\ y_1 - \mu_1 - (\sigma_1\delta)^2/(4(v_0 - Ny_1)) \\ \qquad \text{if } \frac{1}{2}\sigma_1\delta \leq v_0 - Ny_1 < (\sigma_1\delta)^2/(4\zeta^+), \\ v_0/N - \mu_1 - \sigma_1\delta/\sqrt{\bar{n}} \\ \qquad \text{if } (\sigma_1\delta)^2/(4\zeta^+) \leq v_0 - Ny_1, \end{cases} \tag{B.8}$$

and $Z(y_1)$ is a continuous function given by

$$Z(y_1) = \begin{cases} (w_1 + w_2 - w_2N)y_1 - w_2(\mu_1 - v_0 + \sigma_1\delta) \\ \qquad \text{if } v_0 - Ny_1 < \frac{1}{2}\sigma_1\delta, \\ (w_1 + w_2)y_1 - w_2\mu_1 - w_2(\sigma_1\delta)^2/(4(v_0 - Ny_1)) \\ \qquad \text{if } \frac{1}{2}\sigma_1\delta \leq v_0 - Ny_1 < (\sigma_1\delta)^2/(4\zeta^+), \\ w_1y_1 + w_2(v_0/N - \mu_1 - \sigma_1\delta/\sqrt{\bar{n}}) \\ \qquad \text{if } (\sigma_1\delta)^2/(4\zeta^+) \leq v_0 - Ny_1. \end{cases} \tag{B.9}$$

Proof. By applying (39) in Proposition 9 and noting that (36) holds with equality, (B.8) is easily obtained. Plugging (B.8) into the objective function of RR2 yields (B.9). The continuity of $Z(y_1)$ can be verified by checking at the two endpoints of the middle condition. □

We begin to solve RR2 by considering the derivative of $Z(y_1)$:

$$\frac{dZ}{dy_1}(y_1) = \begin{cases} w_1 + w_2(1 - N) \quad \text{if } v_0 - Ny_1 < \dfrac{1}{2}\sigma_1\delta, \\ (w_1 + w_2) - \dfrac{w_2(\sigma_1\delta)^2N}{4(v_0 - Ny_1)^2} \\ \qquad \text{if } \dfrac{1}{2}\sigma_1\delta \leq v_0 - Ny_1 < \dfrac{(\sigma_1\delta)^2}{4\zeta^+}, \\ w_1 \quad \text{if } \dfrac{(\sigma_1\delta)^2}{4\zeta^+} \leq v_0 - Ny_1. \end{cases}$$

Consequently, the optimal value of y_1 will satisfy $v_0 - Ny_1 = \frac{1}{2}\sigma_1\delta$ in the first region and $v_0 - Ny_1 = (\sigma_1\delta)^2/(4\zeta^+)$ in the third region. By the continuity of $Z(y_1)$ (Lemma 8), the optimal solution must lie in the middle region. In this interval, we note that the function $Z(y_1)$ is concave in y_1 since $(d^2Z/dy_1^2)(y_1) = -w_2(\sigma_1\delta N)^2/(2(v_0 - Ny_1)^3) \leq 0$, provided $v_0 \geq Ny_1$. Therefore, setting $dZ/dy_1 = 0$ yields the unconstrained maximizer $\hat{y}_1 = v_0/N - (\sigma_1\delta/(2\sqrt{N}))\sqrt{w_2/(w_1+w_2)}$. The condition $v_0 \geq \frac{1}{2}\sigma_1\delta\sqrt{w_2N/(w_1+w_2)}$ ensures that $\hat{y}_1 \geq 0$. Considering the bounds in this region, we observe that if $N > (w_1+w_2)/w_2$, then the lower bound is naturally satisfied. On the other hand, it is possible that $v_0 - N\hat{y}_1 \geq (\sigma_1\delta)^2/(4\zeta^+)$ for sufficiently small $\bar{n}$. In that event, $\hat{y}_1$ should be chosen to satisfy $v_0 - N\hat{y}_1 = \frac{1}{2}\sigma_1\delta\sqrt{w_2N/(w_1+w_2)} = (\sigma_1\delta)^2/(4\zeta^+)$. Let $f(x) = x(1 + \sqrt{1-x})^{-2}$, then the critical event can be rewritten as $w_2/(w_1+w_2) = f(\bar{n}/N)$. It is easily checked that $f(x)$ is continuous and increasing on $[0,1]$. Furthermore, $f(0) = 0$ and $f(1) = 1$, so by the intermediate value theorem there is a unique solution, x_w, in $(0,1)$, such that $x_w = \bar{n}/N$. Assembling these results, we have

$$v_0 - Ny_1^* = \begin{cases} \dfrac{1}{2}\sigma_1\delta\sqrt{\dfrac{w_2 N}{w_1+w_2}} & \text{if } \bar{n} \geq x_w N, \\ \dfrac{1}{2}\sigma_1\delta\sqrt{\bar{n}}\left(1+\sqrt{1-\dfrac{\bar{n}}{N}}\right)^{-1} & \text{if } \bar{n} < x_w N, \end{cases}$$

where y_1^* is the constrained optimizer of $Z(y_1)$ in the middle interval. Let $n^*(\bar{n})$ be defined as (42), then (43) follows by a simply substitution in the expression above. Rearranging terms yields (44). By (B.8), (45) follows by substitution for y_1^* in (44).

Next, we establish the interpretation of $n^*(\bar{n})$. In the middle region,

$$y_1^* - y_2^* - \mu_1 = \begin{cases} \dfrac{1}{2}\sigma_1\delta\sqrt{\dfrac{w_1+w_2}{w_2 N}} & \text{if } \bar{n} \geq x_w N, \\ \zeta^+ & \text{if } \bar{n} < x_w N. \end{cases}$$

Thus, by Proposition 9, the worst-case number of retailers to receive shipments in period 2 is given by

$$n^* = \left(\frac{\sigma_1\delta}{2(y_1^* - y_2^* - \mu_1)}\right)^2 = n^*(\bar{n}),$$

which confirms the interpretation of (42).

Finally, solving $x(1+\sqrt{1-x})^{-2} = w_2/(w_1+w_2)$ for x gives (46). □

B.18. Proof of Theorem 2

Let $f(y_1, y_2)$ denote the optimal value of the maximization problem on the left-hand side of (48). For any given y_1 and y_2, there are two cases of interest:

1. If $n < \delta_1/(\sigma_1\delta_0)$, then $f(y_1,y_2) = Ny_1 + \max_n\{(y_2+\mu_1 - y_1 + \sigma_1\delta_0)n\}$. Clearly, if $y_2+\mu_1-y_1 > -\sigma_1\delta_0$, the optimizer $n^* = \delta_1/(\sigma_1\delta_0)$; if $y_2+\mu_1-y_1 < -\sigma_1\delta_0$, $n^*=0$. n^* can be any number in $[0, \delta_1/(\sigma_1\delta_0))$ if $y_2+\mu_1-y_1 = -\sigma_1\delta_0$.

2. If $n \geq \delta_1/(\sigma_1\delta_0)$, then $f(y_1,y_2) = Ny_1 + \delta_1 + \max_n\{(y_2 + \mu_1 - y_1)n\}$. Similarly, if $y_2+\mu_1-y_1 > 0$, $n^* = N$; if $y_2+\mu_1-y_1 < 0$, $n^* = \delta_1/(\sigma_1\delta_0)$. n^* can be any number in $[\delta_1/(\sigma_1\delta_0), N]$ if $y_2+\mu_1-y_1 = 0$.

Thus, after plugging the value of n^* into $f(y_1,y_2)$ and collecting terms, we have

$$f(y_1,y_2) = \begin{cases} Ny_1 \quad \text{if } y_2+\mu_1-y_1 \leq -\sigma_1\delta_0, \\ \left(N - \dfrac{\delta_1}{\sigma_1\delta_0}\right)y_1 + \dfrac{\delta_1}{\sigma_1\delta_0}y_2 + (\mu_1+\sigma_1\delta_0)\dfrac{\delta_1}{\sigma_1\delta_0} \\ \qquad \text{if } -\sigma_1\delta_0 < y_2+\mu_1-y_1 < 0, \\ Ny_2 + N\mu_1 + \delta_1 \quad \text{if } y_2+\mu_1-y_1 \geq 0. \end{cases}$$

By checking $y_2+\mu_1-y_1$ on the boundaries, it is easily verified that $f(y_1,y_2)$ is continuous. Since both coefficients of y_1 and y_2 are nonnegative in (47) and in $f(y_1,y_2)$, (48) will be binding at optimality; that is, $f(y_1^*, y_2^*) = v_0$. Based on the analysis above, we examine the optimality of the relaxed problem (47)–(48) in terms of y_1 and y_2 in three regions, respectively:

1. When $y_2+\mu_1-y_1 \leq -\sigma_1\delta_0$, $y_1^* = v_0/N$ and $y_2^* = y_1^* - (\mu_1 + \sigma_1\delta_0) = v_0/N - (\mu_1+\sigma_1\delta_0)$. The optimal objective value equals $M_1^* = (w_1+w_2)(v_0/N) - w_2(\mu_1+\sigma_1\delta_0)$.

2. When $y_2+\mu_1-y_1 \geq 0$, $y_2^* = (v_0-\delta_1)/N - \mu_1$ and $y_1^* = y_2^* + \mu_1 = (v_0-\delta_1)/N$. The optimal objective value equals $M_2^* = (w_1+w_2)((v_0-\delta_1)/N) - w_2\mu_1$.

3. When $-\sigma_1\delta_0 < y_2+\mu_1-y_1 < 0$, writing y_2^* in terms of y_1^* and plugging it into (47) gives

$$\left(w_1+w_2-w_2N\frac{\sigma_1\delta_0}{\delta_1}\right)y_1^* + w_2 v_0\frac{\sigma_1\delta_0}{\delta_1} - w_2(\sigma_1+\sigma_1\delta_0).$$

(a) If $w_2N/(w_1+w_2) \leq \delta_1/(\sigma_1\delta_0)$, then y_1^* will take the upper bound $y_2^*+\mu_1+\sigma_1\delta_0$ prescribed by the constraint $-\sigma_1\delta_0 < y_2+\mu_1-y_1$. This is equivalent to the first case above.

(b) If $w_2N/(w_1+w_2) > \delta_1/(\sigma_1\delta_0)$, then y_1^* will take the lower bound $y_2^*+\mu_1$ determined by the constraint $y_2+\mu_1-y_1 < 0$. This is equivalent to the second case above.

Assembling the above analysis leads to (49) and (51). Note that $v_0 \geq N(\mu_1+\sigma_1\delta_0)$ ensures that $y_1^*, y_2^* \geq 0$. The optimal solution for n^* is degenerate, and we use $\delta_1/(\sigma_1\delta_0)$ as a representative. □

References

Axsäter S (2003) Supply chain operations: Serial and distribution systems. Graves SC, de Kok AG, eds. *Handbooks in OR and MS, Supply Chain Management: Design, Coordination and Operation*, Vol. 11, Chap. 10 (North-Holland, Amsterdam), 525–559.

Axsäter S, Marklund J, Silver EA (2002) Heuristic methods for centralized control of one-warehouse, N-retailer inventory systems. *Manufacturing Service Oper. Management* 4(1):75–97.

Bandi C, Bertsimas D, Youssef N (2015) Robust queueing theory. *Oper. Res.* 63(3):676–700.

Ben-Tal A, Nemirovski A (2001) *Lectures on Modern Convex Optimization—Analysis, Algorithms, and Engineering Applications*, MOS-SIAM Series on Optimization (Society for Industrial and Applied Mathematics, Philadelphia).

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

Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52(1):35–53.

Bertsimas D, Thiele A (2006) A robust optimization approach to inventory theory. *Oper. Res.* 54(1):150–168.

Bimpikis K, Markakis MG (2016) Inventory pooling under heavy-tailed demand. *Management Sci.* 62(6):1800–1813.

Bredström D, Flisberg P, Rönnqvist M (2013) A new method for robustness in rolling horizon planning. *International J. Production Econom.* 143(1):41–52.

Clark AJ, Scarf H (1960) Optimal policies for a multi-echelon inventory problem. *Management Sci.* 6(4):475–490.

Doğru MK, de Kok AG, van Houtum GJ (2009) A numerical study on the effect of the balance assumption in one-warehouse multi-retailer inventory systems. *Flexible Services Manufacturing J.* 21(3–4):114–147.

Eppen G, Schrage L (1981) Centralized ordering policies in a multi-warehouse system with lead times and random demand. Schwarz L, ed. *Multi-Level Production/Inventory Systems: Theory and Practice*, TIMS Studies in the Management Sciences, Vol. 16 (North-Holland, Amsterdam), 51–67.

Federgruen A, Zipkin P (1984a) Allocation policies and cost approximations for multilocation inventory systems. *Naval Res. Logist.* 31(1):97–129.

Federgruen A, Zipkin P (1984b) Approximations of dynamic, multilocation production and inventory problems. *Management Sci.* 30(1):69–84.

Gallego G, Özer Ö, Zipkin PH (2007) Bounds, heuristics, and approximations for distribution systems. *Oper. Res.* 55(3):503–517.

Jackson PL, Muckstadt JA (1989) Risk pooling in a two-period, two-echelon inventory stocking and allocation problem. *Naval Res. Logist.* 36:1–26.

Jönsson H, Silver EA (1987) Analysis of a two-echelon inventory control system with complete redistribution. *Management Sci.* 33(2):215–227.

Kunnumkal S, Topaloglu H (2008) A duality-based relaxation and decomposition approach for inventory distribution systems. *Naval Res. Logist.* 55(7):612–631.

Kunnumkal S, Topaloglu H (2011) Linear programming based decomposition methods for inventory distribution systems. *Eur. J. Oper. Res.* 211(2):282–297.

Mamani H, Nassiri S, Wagner MR (2017) Closed-form solutions for robust inventory management. *Management Sci.* 63(5): 1625–1643.

Zakeri G, Craigie D, Philpott A, Todd M (2014) Optimization of demand response through peak shaving. *Oper. Res. Lett.* 42(1): 97–101.

Zipkin P (1984) On the imbalance of inventories in multi-echelon systems. *Math. Oper. Res.* 9(3):402–423.