Multi-Item Two Echelon Distribution Systems with Random Demands: Bounds and Effective Strategies

Awi Federgruen
Graduate School of Business
Columbia University
af7@columbia.edu

Daniel Guetta
Graduate School of Business
Columbia University
guetta@cantab.net

Garud Iyengar
Industrial Engineering & Operations Research
Columbia University
garud@ieor.columbia.edu

October 11, 2015

Abstract

We address a general model for two-echelon, multi-item systems with an arbitrary number of retailers. The model considers a cost structure with the following four cost components: (i) costs associated with the orders placed with the suppliers, reflecting economies of scale and scope; (ii) shipment costs for the transfers from the depot to the retailers; (iii) for each item, carrying costs associated with the depot's and retailers' inventories at the end of each period, and (iv) for each item, backlogging costs at each of the retailers, a function of the backlog size there at the end of each period.

We develop an approach to compute a tractable lower bound dynamic program (DP) for the optimal system-wide costs, along with specific order, withdrawal and allocation policies that are derived from the strategy which is optimal in the lower bound DP. The lower bound DP is based on a combination of relaxation methods, Lagrangian and others.

An extensive numerical study, involving close to 1,200 instances and reported upon in our paper, compares the lower bound resulting from the Lagrangian dual with the upper bound associated with our novel system-wide order, withdrawal and allocation policy. We observe that almost across the entire parameter space considered, the lower bound is very accurate, and the proposed system-wide policy is close to optimal.

1 Introduction

Distribution systems for retailer organizations are complex, whether they sell via traditional brick-and-mortar stores, online systems, or combinations thereof (dual channels). The complexity of these systems arise from the fact that they deal with a large number of stock keeping units (SKUs) and that inventories need to be kept at many locations, often organized in two (or an even larger numbers of) echelons.

To appreciate the scope of the problem, consider the following: traditional brick-and-mortar retail chains have between scores and thousands of stores. At one end of the spectrum, Walmart has 5,000 stores and wholesale clubs in the US alone. As one of their major strategic advantages, internet retailers benefit from the fact that nationwide sales can be aggregated and fulfilled from national fulfilment centers, rather than a multitude of local stores, but even they find it desirable to operate from several such centers, each requiring its own inventory. Amazon, of course the largest such organization, has a network of over twenty fulfilment centers, for example.

In terms of the number of SKUs carried, the numbers are even more staggering. In April of 2010, Amazon sold some 140,000 items in the electronic category alone, and a similar number of tools and home improvement items. (See Table 1 in Jiang et. al. (2011) for these numbers, along with data from other categories). Progressive Grocer Magazine (as quoted by the Food Marketing Institute (2015)) reports that, in 2014, there were more than 37,000 supermarkets with annual sales of $2 million, or more, most of them *chains* with multiple sales locations. Based on similar data from the previous year, 2013, the same Food Marketing Institute report notes that the average number of items carried in a supermarket is 43,844! Net profits, after taxes, amount to no more than 1.3% of the sales revenues, which explains why efficient supply chain management, in particular efficient inventory and distribution planning is the key to survival and profitability.

The procurement planning process of these various SKUs, at the different facilities of the retail chain, need to be integrated into a system-wide planning model. For any given item, there are various interdependencies among the facilities. More specifically, consider the network structure discussed in this paper: the distribution system consists of a depot and multiple sales outlets, hereafter referred to as *retailers*, which face random demands for the various items sold there. The replenishment process consists of two stages: the depot procures the item from outside suppliers, employing a single supplier for any given item. The retailers' inventories are replenished by shipments from the depot. Both of the replenishment stages are associated with a given item- and facility- specific leadtime. The depot as well as the retailers face a limited inventory capacity, where the storage space is shared among all the items carried by the facility. Inventories are reviewed, orders are placed, and internal shipments are initiated on a periodic basis. When a retailer runs out of stock, unmet demand is backlogged. The interdependencies are then as follows (i) standard economies of scale in the order costs connect the procurement strategies of the various retail outlets; (ii) system-wide orders are to be determined on the basis of *all* the retailers' and the depot's information; moreover, the fact that the replenishment process is conducted in the above two phases allows us to exploit *statistical economics of scale*, a term coined by Eppen and Schrage (1981): when deciding on the *first stage* order quantity, only *aggregate* demands matter, the volatility of which is relatively lower than that at the individual retailers, thus requiring smaller safety stocks to reach a given service level; (iii) the existence of an inventory pool at the depot, from which all the retailers can draw, creates an additional interdepencency among the retailers. Similarly, the planning process for the different items interact with each other, at least because of (a) *shared storage capacity* at each of the facilities, (b) *economies of scope* in the order costs, for example, fixed cost components incurred irrespective of the composition of the order.

Arrow, Harris, and Marschak (1951) and Dvoretsky, Kiefer, and Wolfowitz (1953) initiated the field of stochastic inventory theory, by devising optimal replenishment strategies in systems with a *single* item sold from a *single* location. In the subsequent sixty years, many hundreds of papers have been published in this field, but, to our knowledge, none capable of handling the simultaneous challenge of dealing with a two echelon distribution system with a general number of retailers and a general number of items, under *random demands* and the above interdependencies.

In this paper, we address a general model for such two-echelon, multi-item systems with an arbitrary number of retailers. The model considers a cost structure with the following four cost components: (i) costs associated with the orders placed with the suppliers, reflecting economies of scale and scope; (ii) shipment costs for the transfers from the depot to the retailers; (iii) for each item, carrying costs associated with the depot's and retailers' inventories at the end of each period, and (iv) for each item, backlogging costs at each of the retailers, a function of the backlog size there at the end of each period. In terms of each period's demands for all items at all retailers, they are drawn from a general joint distribution. However, as in most inventory models, we assume that demands are independent across time. The objective is to minimize the expected discounted cost over the planning horizon.

In our base model, we make the following *standard* structural assumptions about the cost components: (i) for each item, the order costs, in each period, consist of a fixed and a variable component; (ii) internal shipment costs are *linear* with the shipment quantities, and item- and facility- specific cost rates; (iii) for each item, the carrying cost at each facility, in each period, is a convex function of its ending inventory level, (iv) in any given period, the backlogging cost, for each item, at each retailer, is a convex function of the end-of-period backlog size. Thus, in our base model, the interdependency among the different items is restricted to the joint storage capacity that prevails at each of the facilities. However, in section 7, we extend our results to settings with a *second* type of interdependency, i.e., *joint fixed costs* among the different items. More specifically, each supplier is selected to supply a complete line of items; a given joint fixed cost is incurred whenever orders are placed with this supplier, irrespective of which items are included in the order and how many of each item is ordered.

We develop an approach to compute a tractable lower bound dynamic program (DP) for the optimal system-wide costs, along with specific order, withdrawal and allocation policies that are derived from the strategy which is optimal in the lower bound DP. The lower bound DP is based on a combination of relaxation methods, Lagrangian and others. As such, the approach bears similarity to other relaxation approaches proposed for systems with a single item or a single sales outlet. In particular, Federgruen and Zipkin (1984a; 1984b), Kunnumkal and Topaloglu (2008; 2011), and Federgruen, Guetta, and Iyengar (2015). We exploit the fact that the planning problems are weakly coupled across the different items. This part of our Lagrangian relaxation approach is of the type discussed by Adelman and Mersereau (2008); see Section 2 for additional references. At the same time, the interdependencies among the different facilities are, as explained, of a fundamentally different type, and, therefore, require an altogether different relaxation approach to transform the exact DP into a lower bound DP that decomposes into independent DPs for each item, each with a *single-dimensional* state space.

A feasible system-wide strategy consists of three parts: (a) an order policy which determines, in each period, how much to order of the different items from each supplier, (b) a *withdrawal* policy, which specifies, in each period, how much inventory to withdraw from the depot for each of the items, and (c) an *allocation* policy, which determines how the withdrawal quantities are best allocated along the retailers. As to the first component, the order policy, we recommend implementing the order policy that is optimal in the Lagrangian dual. To make the withdrawal and allocation decisions, in each period, we design a new type of heuristic. This heuristic calls for the solution, on a rolling horizon basis, of a multi-period convex program, that determines optimal allocation quantities for all items for each of the periods in the look-ahead horizon considered. Clearly, the expected cost value of the combined order, withdrawal and allocation policy represents an *upper bound* on the optimal cost value of the full problem.

An extensive numerical study, involving close to 1,200 instances and reported upon in Section 8, compares the lower bound resulting from the Lagrangian dual with the upper bound associated with our novel system-wide order, withdrawal and allocation policy. The latter is assessed via Monte Carlo simulation. We observe that almost across the entire parameter space considered, the lower bound is very accurate, and the proposed system-wide policy is close to optimal. The computational complexity of the lower bound grows *linear* with the number of items I, the number of retailers J and the length of the planning horizon T.

In addition, we show that some of our fundamental results hold under fully general shapes of the various cost functions. Only the linearity of the internal shipping cost functions is essential to obtain a lower bound that decomposes into separate single-dimensional DPs for each item. The additional structural assumptions regarding cost components, (i), (iii), and (iv) are used to ensure that the order policies that are optimal in the Lagrangian relaxed DP are of a simple (s, S) type and that the mathematical programs to be solved to determine the depot withdrawal and allocation policies are convex.

The remainder of this paper is arranged as follows: Section 2 provides a review of the relevant literature. We introduce our model and associated notation in Section 3. Section 4 gives an exact DP formulation for the system-wide problem. Our Lagrangian lower-bound approach is developed in Section 5 while Section 6 covers our proposed heuristic policy. Section 7 exends our results to models with joint fixed order costs across products purchased from the same supplier. Our numerical study is described in Section 8, while section 9 concludes our paper.

2 Literature Review

We confine ourselves to a brief review of the literature on inventory models for systems with multiple storage locations or multiple items. Existing models capable of simultaneously handling product variety *along with* geographic dispersion of retailers appear to be confined to fully *deterministic* settings, thus ignoring the complexities of demand (or supply) risks.

These deterministic models include mathematical programming models for finite planning horizon problems; see Thomas and McClain (1993) and Shapiro (1993) for surveys. Incorporating fixed order or production costs into these models renders even these deterministic models complex, see e.g. Federgruen, Meissner, and Tzur (2007) and the references therein. For stationary infinite horizon models, under the long-run average cost objective, Maxwell and Muckstadt (1985), Queyranne (1985), Roundy (1986), Federgruen, Queyranne, and Zheng (1992), and Federgruen and Zheng (1992a) address the problem of determining replenishmen tstrategies for very general multi-item production and distribution networks. These authors show that policies of a simple so-called power-of-two structure are guaranteed to come within 2% of optimality, but only when demands are not only *deterministic*, but occur at *constant* rate. Moreover, their methods do not allow for joint capacity constraints.

Clark and Scarf (1960)'s seminal paper was the first to address a periodic review model for one warehouse multiple retailer systems with random demands, and a *single item* and no storage constraints. These authors proposed a heuristic approach based on the so-called *balance assumption*, according to which, at the start of each period, the distribution of inventories among the retailers is the one that minimizes future expected costs. This balance assumption results in a tractable *lower bound* for the actual system, and one that decomposes into J separate retailer problems, along with a DP for the depot, in which the cost functions are obtained from the value functions of the retailers' DPs. (The balance assumption is trivially satisfied when there is only *one* retailer; in this case, the proposed method is indeed exact). Clark and Scarf (1960) failed to report on a numerical study to assess the quality of their proposed lower bound and heuristic policy.

Half a century later, Kunnumkal and Topaloglu (2008) devised a lower bound approximation of a similar type: it also decomposes into J separate retailer problems, along with a DP for the depot whose cost functions are derived from value functions of the retailers' DPs. These authors arrive at *their* loiwer bound by applying a Lagrangian relaxation to the non-negativity constraints for the shipment values in the *exact* DP for the system-wide problem.

Like Clark and Scarf (1960), the Kunnumkal and Topaloglu (2008) model does not consider any capacity constraints, and, in addition, assumes that the order costs are linear, without any fixed components. The authors report on a numerical study that evaluates the accuracy gap of their lower bound and the optimality gap of their associated policy; it demonstrates superiority of their lower bound, with significant differences emerging when the problem instance is highly non-stationary and with significant heterogeneity among the retailers.

A *third* lower bound approach was proposed by Federgruen and Zipkin (1984a) for single-item uncapacitated systems *without* central inventories at the depot, as well as several other restrictions: (i) all single period demands are Normally distributed (or, slightly more formally, have a cdf which is unique up to a centralization and a scaling parameter), (ii) the backlogging and storage costs are *linear* functions, with cost rates that are *identical* across all retailers (under general cost rates, the authors' approach involves additional approximation steps, not necessarily resulting in a lower bound), (iii) the cost rates for shipments from the depot to the retailers are uniform across the retailers. The approach in Federgruen and Zipkin (1984a) relies on a *simple* relaxation of the above-mentioned non-negativity constraints for the shipment quantities. Its major advantage was that the resulting lower bound can be computed by solving a *single* DP with a one-dimensional state space, as opposed to a sequence of partially nested DPs, in the other two approaches.

Most recently, Federgruen, Guetta, and Iyengar (2015) developed a lower bound approach, and considered heuristic policies for *single* item systems. While dealing with a much more general model than Federgruen and Zipkin (1984a), the lower bound in Federgruen, Guetta, and Iyengar (2015) reduces, equally, to a *single* DP with a one-diemnsional state space. We refer the reader to that paper for a more detailed comparison of the four lower bound approaches in single-item systems, as well as a review of other papers that have built on these: Eppen and Schrage (1981), Federgruen and Zipkin (1984b; 1984c), Axsäter, Marklund, and Silver (2002), Axsäter (2003), Doğru (2005), Dogru, De Kok, and Van Houtum (2004) and Gallego, Özer, and Zipkin (2007).

To our knowledge, Li and Muckstadt (2014) and Li and Muckstadt (2015) are the first two papers to address systems that need to contend with multiple items as well as *geographic* dispersion across multiple retail outelts. However, in Li and Muckstadt (2014), there are no interdependencies across items, whether because of joint inventory constraints or because of economies of scope within the cost structure; their model therefore decomposes into separate planning problems for each item, as in the models reviewed above. Moreover, the authors, make an upfront restriction to order policies of a periodic base stock type: every m periods, an order is placed with the outside supplier to elevate the system-wide inventory position to a given base-stock level. For a given base-stock level, the planning problem then decomposes on a cycle-by-cycle basis. Based on the authors' work with an online retailer, a novel feature in the Li and Muckstadt (2014) model is that it distinguishes between *two* types of demands at each retailer, since customers are offered a choice between two possible delivery dates, at different shipment costs. Li and Muckstadt (2015) do consider interdependencies among items, in that customer orders consist of a combination of items. As well as its companion, this paper decouples the order policy part of the system-wide replenishment policy from the above withdrawal and allocation parts, by assuming that the system operates under a given periodic base-stock policy.

When comparing this paper with the single item system in Federgruen, Guetta, and Iyengar (2015), several important distinctions are to be noted: a different Lagrangian dual is needed to model and to handle joint storage capacity constraints, along with new challenges for an efficient evaluation of the required supergradients of the lower bound as a function of the Lagrange multipliers. This method bears similarity to that in Adelman and Mersereau (2008) for weakly coupled stochastic dynamic programs, itself based on Hawkins (2003); see there for other applications. When, as in Section 7, additional interdependencies are introduced through joint fixed order costs, an additional lower bounding approach needs to be added, based on a methodology initiated by Atkins and Iyogun (1988) for multi-item systems with a *single* sales location. In addition, to obtain high quality policies and to demonstrate the quality of the proposed lower bound, we develop a novel withdrawal and allocation policy. In the above prior approaches, withdrawal and allocations in any given period are determined to optimize expected costs in the *first* period in which these shipments are received (myopic allocation) or over a *time window* of several periods, but assuming that no other allocations are added over the course of the time window. Our new withdrawal and allocation policy solves a more intricate but still tractable program that allows future optimal shipments in all periods pertaining to the look-head time window.

3 Model and Notation

We consider a periodic review system with a finite planning horizon of $T < \infty$ periods. Our objective is to minimize expected discounted aggregate costs in the system. As in most standard inventory models, we assume that all stockouts at the retailers are fully backlogged, and that in any given period, the carrying and backlogging costs are assessed, for each product, as a retailer-, time-, and product-specific convex function of the end-of-period inventory and backlog size respectively.

We assume that at each of the system's facilities, aggregate space may be limited. However, when deciding on allocations to a given retailer, the resulting inventory levels upon receipt of these allocations a shipment leadtime later are hard to predict, since they depend on the random demand observed while the shipment is in transit. More specifically, the inventory level , a shipment leadtime hence, equals the current inventory position (i.e., the inventory level plus shipments in progress minus backlog), plus the new allocation, minus the (random) leadtime demand. For this reason, we model the capacity constraint by replacing the random leadtime demand for any given item by a user chosen low fractile thereof, say the 0.05 fractile. Thus, barring demand realisations that fall below these fractiles, the capacity constraint ensures that the actual inventory level a shipment leadtime later is guaranteed to fall below the prevailing capacity level.

This specification does not exclude the the possibility of an inventory overflow. However, as demonstrated in Section 8, by varying the above fractile, any desired overflow probability can be targeted. Of course, the possibility of an overflow can be completely prevented by replacing the random leadtime demand by the smallest value in its support. However, this approach is extremely conservative, and is likely to leave much of the capacity under-utilized at each retailer.

In a similar vein, we model any storage capacity constraint at the depot via a convex non-linear carrying cost function, the slope of which increases significantly at the 'normal' capacity bound, when applicable. (Specification via a simple hard capacity constraint may result in an infeasible DP).

The sequence of events in each period is as follows. (1) Deliveries to the depot and retailers are received. (2) Decisions are made as to how inventory at the depot should be apportioned, whether the depot should place a new order, and if so of what size. (3) Demand is observed. (4) Holding and backorder costs are assessed at the retailers, as well as carrying costs at the depot, given the remaining inventory levels at these various facilities.

Before we define our notation in detail, we note the following general conventions:

- Stores are indexed by $j = 1, \cdots, J$, positioned in the *subscript* of the variable in question.
- Products are indexed by $i = 1, \cdots, I$, positioned in the *superscript* of the variable of interest.
- Lowercase Latin letters denote costs and inventories at the stores, whereas upercase latin letters denote costs and inventories at the depot.
- We index time by $t = 1, \cdots, T$, positioned in parentheses in the *subscript* of the variable in question.

We use the following notation:

Infrastructure data :

- J: numbers of retailers, indexed by j.
- I: number of products our distribution system deals with, indexed by i.
- L^i: leadtime for orders of product i to the central depot.
- ℓ^i_j: leadtime for product i allocations to retailer j from the central depot.
- δ: the discount factor.

Capacity data :

- $\chi_{j,(t)}$: storage capacity of retailer j in period t.
- β_j: the maximum permitted probability of an overflow at retailer j at the start of any period.
- α_j: a specified fractile of the leadtime demand of any of the items coverd by retailer j.

Costs :

- $K^i_{(t)}$: the fixed cost in period t to place an order for product i from the supplier.
- $c^i_{(t)}$: the variable procurement cost rate in period t for product i, from the supplier.
- $p^i_{j,(t)}(\cdot)$: the backlogging costs for product i at retailer j at the end of period t, as a function of the total backlog size.
- $h_{j,(t)}(\cdot)$: the inventory carrying costs for product i at retailer j at the end of period t, as a function of the total inventory level.
- $H^i_{(t)}(\cdot)$: the product-i inventory carrying costs at the depot at the end of period t, as a function of the total inventory level.
- γ^i_j: the variable shipment cost rate for units of product i shipped to retailer j in period t.

Demand : The one-period demand for product i at retailer j in period t is represented by $u^i_{j,(t)}$. For any given period t, the random demand variables $\{u^i_{j,(t)} : i = 1, \cdots, I, j = 1, \cdots, J\}$ follow a general joint distribution. However, demands are independent across time.

We denote the marginal CDF of $u^i_{j,(t)}$ by $F^i_{j,(t)}$.

We are now ready to develop an exact DP formulation of the problem:

State of the system : The following variables fully determine the state of the system at the start of any given time period t:

- $x^i_{j,(t)}$: the total product-i inventory position at retailer j at the beginning of period t, before receipt of any shipment sent ℓ_j periods earlier. This includes the total inventory on hand at the retailer at the start of the period, as well as all allocations currently in the pipeline, on their way to the retailer.
- $X^i_{(t)}$: product-i inventory at the depot at the start of period t, before receipt of any order placed L^i periods earlier.

- $W^i_{(\tau)}$: the product-i order placed by the depot in period τ, for delivery in period $\tau + L^i$. At the start of period t, we shall only require these variables for time periods $\tau = t - L^i, \cdots, t-1$.

Thus, the $(J+L+1)I$-tuple $(x^1_{1,(t)}, \cdots, x^I_{J,(t)}, X^1_{(t)}, \cdots, X^I_{(t)}, W^1_{(t-L)}, \cdots, W^I_{(t-1)})$ serves as the state vector of the system at the beginning of period t.

Actions taken in period t : In period t, the following decisions need to be made. (We denote each decision by the same symbol as the corresponding state variable, but with the addition of a bar).

- $\bar{x}^i_{j,(t)}$: the product-i inventory position at retailer j at the start of period t, including any new allocation to retailer j in this period.
- $\bar{X}^i_{(t)}$: the product-i inventory at the depot, in period t, after inclusion of any newly arrived orders from the supplier, and after subtraction of any new allocations to the retailers.
- $\bar{W}^i_{(t)}$: the size of the order for product i to be placed with the external supplier in period t.

To formulate the problem, we use the following general conventions

- A + sign in the subscript (superscript) dnotes summation over all retailers (products). For example, $x^i_{+,(t)} = \sum_{j=1}^J x^i_{j,(t)}$.
- A bold symbol with a missing subscript (superscript) denotes a vector over all stores (products).

  (For example, $\boldsymbol{x}^i_{(t)} = \{x^i_{1,(t)}, \cdots, x^i_{J,(t)}\}$, and $\boldsymbol{x}_{j,(t)} = \{x^1_{j,(t)}, \cdots, x^I_{j,(t)}\}$.)

  $\boldsymbol{W}^i$ is used to represent the vector of product-i orders that are outstanding at the very start of period t, i.e. that were placed in periods $t - L^i$ to $t-1$, i.e.,

$$\boldsymbol{W}^i = \{W^i_{(t-L^i)}, \cdots, W^i_{(t-1)}\}$$

- A letter in a blackboard font with no retailer or product subscript or superscript denotes a matrix containing values of that variable for all retailers and products.

  For example, $\mathbb{x}_{(t)}$ is a matrix that contains all inventory positions at the start of period t.

$\mathbb{W}_{(t)}$ contains the vectors $\boldsymbol{W}^i_{(t)}$ for all i, i.e.

$$\mathbb{W}_{(t)} = \begin{bmatrix} W^1_{(t-L^1)} & \cdots & W^I_{(t-L^I)} \\ \vdots & \ddots & \vdots \\ W^1_{(t-1)} & \cdots & W^I_{(t-1)} \end{bmatrix}$$

Note that the $\boldsymbol{W}$ vectors may vary in size if L^i is not uniform across products, so this object is not necessarily a matrix.

- When considering demand variables, a hat indicates demand over the next L time periods, a dot demand over the next ℓ_j periods, and a tilde demand over the next $\ell_j + 1$ time periods. For example, $\hat{u}^i_{j,(t)} = \sum_{\tau=t}^{t+L-1} u^i_{j,(\tau)}$, $\dot{u}^i_{j,(t)} = \sum_{\tau=t}^{t+\ell_j-1} u^i_{j,(\tau)}$, and $\tilde{u}^i_{j,(t)} \sum_{\tau=t}^{t+\ell_j} u^i_{j,(\tau)}$.

4 An Exact Dynamic Programming Formulation

In this section, we derive an exact dynamic programming formulation, which uses the $(J+L+1)I$-tuple $(x^1_{1,(t)}, \cdots, x^I_{J,(t)}, X^1_{(t)}, \cdots, X^I_{(t)}, W^1_{(t-L)}, \cdots, W^I_{(t-1)})$ as the state of the system at the start of a given period t. This state description is maximally concise. As discussed in the introduction, the exact DP is, by itself, of very limited value, since in all but the most trivial applications, its $(J+L+1)I$-dimensional state space precludes computational tractability. However, the exact DP provides the starting point for the various relaxation approaches proposed in this paper.

To simplify our exposition, we assume all the shipment leadtimes are identical, i.e., $L^i = L$ for all i, and $\ell^i_j = \ell$ for all i and j. However, all results continue to hold when these leadtimes vary by product and retailer. We shall further assume that variable shipment rates to each retailer are zero, i.e., $\gamma^i_{j,(t)} = 0$ for all i, j, and t. Nonzero shipment rates can trivially be reintroduced using a simple transformation; see Federgruen, Guetta, and Iyengar (2015), for example.

Since any stockouts at a retailer are fully backlogged, a retailer's product-i inventory level at the end of period $(t + \ell_j)$ equals the inventory position $\bar{x}^i_{j,(t)}$ at the beginning of period t *after* inclusion of any allocations made by the depot, minus the cumulative demand $\tilde{u}^i_{j,(t)}$ that occurs in the interval $[t, t + \ell_j]$. This simple identity allows us to express the expected discounted carrying and backlogging costs, at the end of period $t + \ell_j$, as a function of $\bar{x}^i_{j,(t)}$ only; we call this function $Q^i_{j,(t)}$, where

$$Q^i_{j,(t)}(\bar{x}_{j,(t)}) \equiv \delta^{\ell_j} \mathbb{E} \left\{ h_{j,(t+\ell_j)}(\bar{x}_{j,(t)} - \tilde{u}_{j,(t)}) + p_{j,(t+\ell_j)}(\bar{x}_{j,(t)} - \tilde{u}_{j,(t)}) \right\} \tag{1}$$

The product-i inventory level at the start of period $t+\ell_j$ is given by $\left\{\bar{x}^i_{j,(t)} - \dot{u}^i_{j,(t)}\right\}$, so that the inventory on hand equals $\left[\bar{x}^i_{j,(t)} - \dot{u}^i_{j,(t)}\right]^+$.

Thus, to ensure that the aggregate inventory at the start of period t at retailer j meets capacity constraints there, the variables $\{\bar{x}^i_{j,(t)}, i = 1, \cdots, I\}$ must satisfy the chance constraint

$$\mathbb{P}\left(\sum_{i=1}^{I}\left[\bar{x}^i_{j,(t)} - \dot{u}^i_{j,(t)}\right]^+ \geq \chi_{j,(t+\ell_j)}\right) \leq \beta_j \tag{2}$$

In the single item case, this chance constraint is easily seen to be equivalent to a simple upper bound on the single $\bar{x}^i_{j,(t)}$ variables. However, in the multi-product case, the chance constraint confines the vector $\bar{\boldsymbol{x}}_{j,(t)}$ to a complex region in $\mathbb{R}^I$. Under general demand distributions, even the verification of whether a joint vector $\bar{\boldsymbol{x}}_{j,(t)}$ satisfies the chance constraint involves the evaluation of a complex I-dimensional integral. Closed form approximations can be devised, for example by replacing $\sum_{i=1}^{I}[\bar{x}^i_{j,(t)} - \dot{u}^i_{j,(t)}]^+$ by a Normal distribution. Such a replacement is asymptotically correct as the number of items becomes large, an application of the Central Limit Theorem. However, the resulting constraint has a highly complex non-linear dependence on the vector $\bar{\boldsymbol{x}}_{j,(t)}$.

We therefore adopt a different approach, as mentioned in the introduction. We replace each of the random variables in (2) by a low-fractile of its distribution, more specifically the α_jth fractile of the distribution, and specify that the aggregate inventory at the end of period $t+\ell_j$ remains bounded by the prevailing storage capacity $\chi_{j,(t+\ell_j)}$ as long as demand at the retailer for each product is at least as large as this α_jth fractile. This gives rise to the constraint

$$\sum_{i=1}^{I}\left[\bar{x}^i_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}\right]^+ \leq \chi_{j,(t+\ell_j)} \tag{3}$$

where

$$\dot{u}^{i,[\alpha_j]}_{j,(t)} = (\tilde{F}^i_{j,(t)})^{-1}(\alpha_j)$$

This constraint can be linearized by the addition of auxiliary variables $Y^i_{j,(t)}$. Replace (3) by

$$\sum_{i=1}^{I} Y^i_{j,(t)} \leq \chi_{j,(t+\ell_j)} \tag{4}$$

along with the constraints $Y^i_{j,(t)} \geq \bar{x}^i_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}$ and $Y^i_{j,(t)} \geq 0$ for all $i = 1, \cdots, I$ and $j = 1, \cdots, J$.

The actual overflow probability clearly increases with the choice of the fracticle α_j. This implies that any target overflow probability β_j may be targeted with a simple bisection search on the parameter α_j. See Section 8 for a demonstration of this approach.

The complexity of evaluating the functions $Q^i_{j,(t)}(\cdot)$ and the demand percentiles $\tilde{u}^{i,[\alpha_j]}_{j,(t)}$ depends on the ease with which the CDFs of the convolutions $\dot{u}_{j,(t)}$ and $\tilde{u}_{j,(t)}$ can be determined, as well as the complexity of the functions $h_{j,(t)}(\cdot)$ and $p_{j,(t)}(\cdot)$. Let

$$V_{(t)}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)}) \;=\; \text{The expected minimal present value of costs incurred in periods } t, t+1, \cdots, T \text{ when starting in state } (\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)})$$

It is easily verified that the value functions satisfy the following exact DP:

$$\begin{aligned}
&V_{(t)}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)}) \\
&\quad = \min_{\substack{\bar{\mathbb{x}}_{(t)}, \bar{\boldsymbol{X}}_{(t)}, \\ \bar{\boldsymbol{W}}_{(t)}}} \Big\{ \delta \mathbb{E}\left[V_{(t+1)}(\bar{\mathbb{x}}_{(t)} - \mathbb{u}_{(t)}, \bar{\boldsymbol{X}}_{(t)}, \mathbb{W}_{(t\to t+1)}) \right] \\
&\qquad\qquad + \sum_{i=1}^{I} K^i_{(t)} \mathbb{I}_{\bar{W}^i_{(t)} > 0} + c^i_{(t)} \bar{W}^i_{(t)} + H^i_{(t)}(\bar{X}^i_{(t)}) \\
&\qquad\qquad\qquad + \sum_{i=1}^{I} \sum_{j=1}^{J} Q^i_{j,(t)}(\bar{x}^i_{j,(t)}) \Big\} \qquad (5)
\end{aligned}$$

$$\begin{aligned}
\text{s.t.} \quad & \bar{x}^i_{+,(t)} + \bar{X}^i_{j,(t)} = x^i_{+,(t)} + X^i_{(t)} + W^i_{(t-L)} && \forall i && (6) \\
& \sum_{i=1}^{I} \max(\bar{x}^i_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}, 0) \le \chi_{j,(t+\ell_j)} && \forall j && (7) \\
& \bar{X}^i_{(t)} \ge 0 && \forall i && (8) \\
& \bar{x}^i_{j,(t)} \ge x^i_{j,(t)} && \forall i, j && (9)
\end{aligned}$$

where

$$\mathbb{W}_{(t\to t+1)} = \begin{bmatrix} W^1_{(t-L^1+1)} & \cdots & W^I_{(t-L^I+1)} \\ \vdots & \ddots & \vdots \\ W^1_{(t-1)} & \cdots & W^I_{(t-1)} \\ \bar{W}^1_{(t)} & \cdots & \bar{W}^I_{(t)} \end{bmatrix}$$

and with $V^{\lambda}_{(T-\ell+1)}(\cdot) = 0$.

Constraint (6) ensures, for each product, that the total inventory in our system at the start of period t plus any order to be received there at the start of that period must equal the sum of the retailers' inventory positions after new allocations, plus the amount of inventory kept at the depot. Each of the shipment quantities must be non-negative, as implied by constraint (9), which stipulates that the product i inventory allocated to retailer j must be greater than what was already there. Constraint (8) specifies that the inventory allocated to the depot must be non-negative. The remaining set of constraints (7) are identical to the inventory capacity constraints (see (3)).

5 A Lower Bound Approximation via Lagrangian Relaxation

In our base model, the interdependencies between the different items is confined to the joint storage capacity constraints of the different retailers. In section 7, we consider additional interdependencies among the items, in particular those arising from joint fixed order costs associated with groups of items procured from the same supplier.

Given the assumptions in our base model, the planning problems, complex as they may be, are nevertheless only *weakly* coupled among the different items, in the sense employed by Adelman and Mersereau (2008). Indeed, were the constraint set (7), which represents the joint storage capacity constraints, to be relaxed, the dynamic program would decompose into I independent DPs. Our first relaxation approach, therefore, consists of relaxing this constraint set in a Lagrangian manner, associating a non-negative Lagrange multiplier $\lambda_{j,(t)}$ to the constraint pertaining to retailer j's allocations in period t.

Under any given set of Lagrange multipliers $\{\lambda_{j,(t)} : j = 1, \cdots, J, t = 1, \cdots, T\}$, we obtain an *approximate* DP with value functions $V^{\lambda}_{(t)}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)})$ that satisfy the following recursions

$$
\begin{aligned}
& V^{\lambda}_{(t)}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)}) \\
&\quad = \min_{\substack{\bar{\mathbb{x}}_{(t)}, \bar{\boldsymbol{X}}_{(t)}, \\ \bar{\boldsymbol{W}}_{(t)}}} \Bigg\{ \delta \mathbb{E}\left[V^{\lambda}_{(t+1)}(\bar{\mathbb{x}}_{(t)} - \mathbb{u}_{(t)}, \bar{\boldsymbol{X}}_{(t)}, \mathbb{W}_{(t\to t+1)}) \right] \\
&\qquad\qquad + \sum_{i=1}^{I} K^i_{(t)} \mathbb{I}_{\bar{W}^i_{(t)} > 0} + c^i_{(t)} \bar{W}^i_{(t)} + H^i_{(t)}(\bar{X}^i_{(t)}) \\
&\qquad\qquad\quad + \sum_{j=1}^{J} \lambda_{j,(t)} \sum_{i=1}^{I} \max(\bar{x}^i_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}, 0) \\
&\qquad\qquad\qquad + \sum_{i=1}^{I}\sum_{j=1}^{J} Q^i_{j,(t)}(\bar{x}^i_{j,(t)}) + \Bigg\} - \sum_{j=1}^{J} \lambda_{j,(t)} \chi_{j,(t)}
\end{aligned} \tag{10}
$$

$$
\begin{aligned}
\text{s.t.} \quad & \bar{x}^i_{+,(t)} + \bar{X}^i_{j,(t)} = x^i_{+,(t)} + X^i_{(t)} + W^i_{(t-L)} && \forall i && (11) \\
& \bar{X}^i_{(t)} \geq 0 && \forall i && (12) \\
& \bar{x}^i_{j,(t)} \leq \chi_{j,(t+\ell_j)} + \dot{u}^{i,[\alpha_j]} && \forall i, j && (13) \\
& \bar{x}^i_{j,(t)} \geq x^i_{j,(t)} && \forall i, j && (14)
\end{aligned}
$$

with $V^{\lambda}_{(T-\ell+1)}(\cdot) = 0$. Via (13), we maintain simple upper bounds for the individual allocation variables $\bar{x}^i_{j,(t)}$, implied by the capacity constraints.

We now prove a Theorem that shows, for any vector of Lagrange multipliers $\boldsymbol{\lambda}$, that $V^{\lambda}_{(t)}(\cdot)$ is a lower bound for $V_{(t)}(\cdot)$, as well as a proposition that shows that the Lagrangian relaxation allows us to decompose this large DP into I smaller ones, one for each product.

Theorem 1. *For any $\lambda = (\lambda_{j,(t)}) \in \mathbb{R}^{JT}_{+}$ and any starting state $(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)})$,*

$$V^{\lambda}_{(t)}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)}) \leq V_{(t)}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)}) \tag{15}$$

Proof. We prove this lower bound by backward induction for $t = T-\ell, T-\ell-1, \cdots, 1$. Starting with the last period $(T-\ell)$, the bound follows from the fact that the minimization in the relaxed DP (10)-(14) for $V^{\lambda}_{(t)}$ is conducted over a larger feasible region than that in the exact DP, while the objective value in (10), for any feasible solution in the exact DP, is *lower* than that in the objective function (5), pertaining to the exact DP. Assume now that the bound applies pointwise for some period $(t+1)$. The proof that the bound applies to period t as well is analogous to the proof for the last period $(T-\ell)$. □

Proposition 1. *For each product i, define value functions $\theta^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})$ that satisfy the following recursions*

$$\begin{aligned}
&\theta^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)}) \\
&\quad= \min_{\substack{\bar{\boldsymbol{x}}^i_{(t)}, \bar{X}^i_{(t)}, \\ \bar{W}^i_{(t)}}} \Bigg\{ \delta \mathbb{E}\left[\theta^{i,\lambda}_{(t+1)}(\bar{\boldsymbol{x}}^i_{(t)} - \boldsymbol{u}^i_{(t)}, \bar{X}^i_{(t)}, \boldsymbol{W}^i_{(t\to t+1)})\right] \\
&\qquad\qquad + K^i_{(t)} \mathbb{I}_{\bar{W}^i_{(t)}>0} + c^i_{(t)} \bar{W}^i_{(t)} + H^i_{(t)}(\bar{X}^i_{(t)}) \\
&\qquad\qquad + \sum_{j=1}^{J} \lambda_{j,(t)} \max(\bar{x}^i_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}, 0) \\
&\qquad\qquad + \sum_{j=1}^{J} Q^i_{j,(t)}(\bar{x}^i_{j,(t)}) + \Bigg\} - \frac{1}{I} \sum_{j=1}^{J} \lambda_{j,(t)} \chi_{j,(t)}
\end{aligned} \tag{16}$$

$$s.t. \quad \bar{x}^i_{+,(t)} + \bar{X}^i_{j,(t)} = x^i_{+,(t)} + X^i_{(t)} + W^i_{(t-L)} \qquad \forall i \tag{17}$$

$$\bar{X}^i_{(t)} \geq 0 \qquad \forall i \tag{18}$$

$$\bar{x}^i_{j,(t)} \leq \chi_{j,(t+\ell_j)} + \dot{u}^{i,[\alpha_j]}_{j,(t)} \qquad \forall j \tag{19}$$

$$\bar{x}^i_{j,(t)} \geq x^i_{j,(t)} \qquad \forall j \tag{20}$$

with terminal condition $\theta^{i,\lambda}_{(T-L+1)}(\cdot) = 0$.

Then

$$V^{\lambda}_{(t)}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)}) = \sum_{i=1}^{I} \theta^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)}) \tag{21}$$

Proof. The result is trivially true for the terminal period $t = T - L + 1$.

Now, suppose (21) holds for time period t. Substituting equation (21) for $V^{\lambda}_{(t)}$ into the recursive equation (16)-(20) for $V^{\lambda}_{(t-1)}$, and noting that the resulting program decomposes across products, we conclude that (21) also holds for $t - 1$. □

Proposition 1 shows that $V^{\lambda}_{(t)}$ can be calculated as the solution of I independent $(J + L + 1)$-dimensional dynamic programs. Whilst this is a major improvement as compared to the $(J+L+1)I$-dimensional exact DP, this still leaves us with an intractable dynamic program.

However, the dynamic program (16)-(20) represents a version of the single item model discussed in Federgruen, Guetta, and Iyengar (2015). That paper proposed a close-to-accurate and highly tractable lower bound $\underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)}) \leq \theta^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})$ for the value functions of that DP. More specifically, the lower bound essentially amounts to the solution of a *single* dimensional DP, indeed one that can be interpreted as pertaining to a *single* item and *single* location inventory system in which a so-called (s, S) policy is optimal in each period. This means that the depot places an order whenever the *system-wide* inventory position is at or below s, to raise this inventory position to a level $S > s$. The system-wide inventory position consists of all inventory currently in the system, plus all outstanding orders with the supplier minus the retailers' backlog. In the lower bound DP with value functions $\underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})$, it is is optimal, in each period t, to determine the optimal withdrawal quantity and the new allocations to the retailers by solving a strictly convex mathematical program. In Section 5.1, we give a brief summary of the methodology required to develop the lower bounds $\underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})$.

While any vector of Lagrange multipliers $\boldsymbol{\lambda} = \{\lambda_{j,(t)}, j = 1, \cdots, J, t = 1, \cdots T\}$ generates a valid lower bound $\sum_{i=1}^{I} \underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)}) \leq V^{\lambda}_{t}(\mathbb{x}_{(t)}, \boldsymbol{X}_{(t)}, \mathbb{W}_{(t)})$, the most accurate such bound is obtained by maximizing this lower bound over the space of all non-negative such $\boldsymbol{\lambda}$-vectors.

$$\text{(D)} \qquad \max_{\lambda=\{\lambda_{j,(t)}\geq 0: j=1,\cdots,J, t=1,\cdots,T\}} \sum_{i=1}^{I} \underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)}) \tag{22}$$

The following theorem shows that the lower bound in (22) is a jointly concave function of $\boldsymbol{\lambda}$. This implies that the unconstrained dual optimization problem may be solved with a standard steepest ascent method, greatly facilitated by the fact that gradients to the lower bound function can be specified and computed, at any value $\boldsymbol{\lambda}^0$ of the multipliers at which the lower bound is evaluated.

Let $\bar{\boldsymbol{x}}^{\star,i,\lambda}_{(t)}$ denote the unique optimal vector of retailer allocations in the lower bound DP under the combined order, withdrawal and allocation policy that is optimal for the lower bound DP.

Theorem 2. *For any given set of values* $(\mathbb{x}_{(1)}, \boldsymbol{X}_{(1)}, \mathbb{W}_{(1)})$*,*

(a) *The lower bound* $\underline{V}^{\lambda}_{(t)}(\mathbb{x}_{(1)}, \boldsymbol{X}_{(1)}, \mathbb{W}_{(1)}) = \sum_{i=1}^{I} \underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(1)}, X^i_{(1)}, \boldsymbol{W}^i_{(1)})$ *is a jointly concave function of* $\boldsymbol{\lambda}$*.*

(b) *The lower bound* $\underline{V}^{\lambda}_{(1)}(\mathbb{x}_{(1)}, \boldsymbol{X}_{(1)}, \mathbb{W}_{(1)})$ *is a differentiable function of the vector of Lagrange multipliers* $\boldsymbol{\lambda}$*, with*

$$\frac{\partial \underline{V}^{\lambda}_{(1)}(\mathbb{x}_{(1)}, \boldsymbol{X}_{(1)}, \mathbb{W}_{(1)})}{\partial \lambda_{j,(t)}} = \sum_{i=1}^{I} \mathbb{E}\left[\bar{x}^{\star,i,\lambda}_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}\right]^+ - \chi_{j,(t+\ell_j)}, \tag{23}$$

where the expectation is taken under the combined order-, withdrawal-, and allocation policy that is optimal in the lower bound DP with value functions $\underline{\theta}^{i,\lambda}_{(t)}(\cdot)$*.*

Proof. To prove part (a), note that the value functions $\underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})$ solve a DP of the following form:

$$\begin{aligned}
\theta^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)}) \quad &= \min_{\bar{\boldsymbol{x}}^i_{(t)}, \bar{X}^i_{(t)}, \bar{W}^i_{(t)}} \quad \Big\{ \delta \mathbb{E}[\theta^{i,\lambda}_{(t+1)}(\bar{\boldsymbol{x}}^i_{(t)} - \boldsymbol{u}^i_{(t)}, \bar{X}^i_{(t)}, \boldsymbol{W}^i_{(t\to t+1)})] + K^i_{(t)} \mathbb{I}_{\bar{W}^i_{(t)}>0} \\
&\qquad + c^i_{(t)} \bar{W}^i_{(t)} + H^i_{(t)}(\bar{X}^i_{(t)}) + \sum_{j=1}^{J} \lambda_{j,(t)} \left[\bar{x}^i_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}\right]^+ \\
&\qquad + \sum_{j=1}^{J} Q^i_{j,(t)}(\bar{x}^i_{j,(t)}) \Big\} - \frac{1}{I} \sum_{j=1}^{J} \lambda_{j,(t)} \chi_{j,(t)} \qquad (24) \\
\text{s.t.} \qquad & [\bar{\boldsymbol{x}}^i_{(t)}, \bar{X}^i_{(t)}, \bar{W}^i_{(t)}] \in \mathcal{F}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})
\end{aligned}$$

As described in Section 5.1, the *lower bound* value functions $\underline{\theta}^{i,\lambda}_{(t)}$ are obtained by relaxing the feasible action spaces $\mathcal{F}$. Depending on the specific implementation, the relaxation may simply eliminate some of the constraints, or relax them in a Lagrangian manner, in which case additional terms are added to the minimand in (24).

Note first that the feasible action sets $\mathcal{F}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})$ are compact and remain so under the above described relaxations. The concavity proof of the value functions $\underline{\theta}^{i,\lambda}_{(t)}(\cdot)$ can be demonstrated by induction, for $t = T-L, \cdots, 1$. In period $t = T-L$, concavity follows from Danskin's Theorem, since the minimand (24) in the recursive equations is an affine, and hence concave function of λ while the feasible action sets are compact. Assume, therefore, that the concavity property has been demonstrated for the value functions $\underline{\theta}^{i,\lambda}_{(t+1)}(\cdot)$ for some $t = 1, \cdots, T-L-1$. This implies that the minimand in the recursive equation for period t is a concave function of λ. The above argument can be repeated to obtain concavity of the value function $\underline{\theta}^{i,\lambda}_{(t)}(\cdot)$, thus completing the induction proof.

To prove part (b), note that differentiability of each of the value functions $\underline{\theta}^{i,\lambda}_{(t)}(\cdot)$ can once again be shown by induction, again employing Danskin's theorem as well as the fact that in every possible state of the system and in every period, the optimal order, withdrawal and allocation quantities are unique. Uniqueness of the withdrawal and allocation quantities follows from the fact that they are part of the optimal solution of a *strictly convex* mathematical program.

This leaves us with the verification of the derivatives in (23). It suffices to show that for any given $j = 1, \cdots, J$ and $t = 1, \cdots, T-L$

$$\frac{\partial \underline{\theta}^{i,\lambda}_{(1)}(\mathbb{x}_{(1)}, \boldsymbol{X}_{(1)}, \mathbb{W}_{(1)})}{\partial \lambda_{j,(t)}} = \mathbb{E}\left[\bar{x}^{\star,i,\lambda}_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}\right]^+ - \frac{\chi_{j,(t+\ell_j)}}{I}$$

We show this by noting that for $\tau = 1, \cdots, t$,

$$\frac{\partial \underline{\theta}^{i,\lambda}_{(\tau)}(\mathbb{x}_{(\tau)}, \boldsymbol{X}_{(\tau)}, \mathbb{W}_{(\tau)})}{\partial \lambda_{j,(t)}} = \mathbb{E}\left[\bar{x}^{\star,i,\lambda}_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}\right]^+ - \frac{\chi_{j,(t+\ell_j)}}{I} \tag{25}$$

and for any $\tau = t+1, \cdots, T-L$,

$$\frac{\partial \underline{\theta}^{i,\lambda}_{(\tau)}(\mathbb{x}_{(\tau)}, \boldsymbol{X}_{(\tau)}, \mathbb{W}_{(\tau)})}{\partial \lambda_{j,(t)}} = 0 \tag{26}$$

The proof of (26) is immediate. The proof of (25) follows by backwards induction for $\tau = t, t-1, \cdots, 1$, employing Danskin's Theorem, and justifying interchanges of the differentiation and expectation operators using the fact that $\left[\bar{x}^{\star,i,\lambda}_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}\right]^+ \leq \left[\chi_{j,(t+\ell_j)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}\right]^+ < \infty$. □

5.1 A Tractable Lower Bound for the Single-Product Value Function

As mentioned, for any item $i = 1, \cdots, I$, the lower-bound DP (16)-(20) with value function $\theta^{i,\lambda}_{(t)}(\cdot)$ is easily interpreted as an instance of the *single item* model addressed in Federgruen, Guetta, and Iyengar (2015). This paper develops a lower bound DP for this model, by relaxing the constraints (20), either completely or in a Lagrangian manner. Recall that these constraints ensure that the shipment quantities from the depot to the retailers are non-negative. The (Lagrangian) relaxation of this single set of constraints allows us to reduce the $(J+L+1)$-dimensional DP (16)-(20) to one that has a one-dimensional state-space, and is structurally identical to a single item, single location inventory model for which an (s, S) policy is optimal.

For notational simplicity, we will derive this lower bound when the constraint set (20) is completely eliminated. We direct the reader to Federgruen, Guetta, and Iyengar (2015) for details of this relaxation when a Lagrangian relaxation is used instead.

First note that, after eliminating constraint set (20), the dynamic program for $\underline{\theta}^{i,\lambda}_{(t)}$ becomes

$$
\begin{aligned}
&\underline{\theta}^{i,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)}) \\
&= \min_{\substack{\bar{\boldsymbol{x}}^i_{(t)}, \bar{X}^i_{(t)}, \\ \bar{W}^i_{(t)}}} \Bigg\{ \delta\mathbb{E}\left[\underline{\theta}^{i,\lambda}_{(t+1)}(\bar{\boldsymbol{x}}^i_{(t)} - \boldsymbol{u}^i_{(t)}, \bar{X}^i_{(t)}, \boldsymbol{W}^i_{(t\to t+1)})\right] \\
&\qquad + K^i_{(t)}\mathbb{I}_{\bar{W}^i_{(t)}>0} + c^i_{(t)}\bar{W}^i_{(t)} + H^i_{(t)}(\bar{X}^i_{(t)}) \\
&\qquad + \sum_{j=1}^{J} \lambda_{j,(t)} \max(\bar{x}^i_{j,(t)} - \dot{u}^{i,[\alpha_j]}_{j,(t)}, 0) \\
&\qquad + \sum_{j=1}^{J} Q^i_{j,(t)}(\bar{x}^i_{j,(t)}) + \Bigg\} - \frac{1}{I}\sum_{j=1}^{J}\lambda_{j,(t)}\chi_{j,(t)} \\
\text{s.t.}\quad & \bar{x}^i_{+,(t)} + \bar{X}^i_{j,(t)} = x^i_{+,(t)} + X^i_{(t)} + W^i_{(t-L)} \qquad \forall i \\
& \bar{X}^i_{(t)} \geq 0 \qquad \forall i \\
& \bar{x}^i_{j,(t)} \leq \chi^i_{j,(t+\ell_j)} \qquad \forall j
\end{aligned}
$$

Now, define a new quantity $A^i_{(t)} \equiv x^i_{+,(t)} + X^i_{(t)}$ to denote the total system-wide inventory level of product i at the beginning of period t, which resides somewhere in the distribution system, either at the depot or the retailers, or in transit between them. Observe that in the DP above, $\underline{\theta}^{i,\lambda}_{(t)}$ depends on the geographically disaggregated inventory information $(\boldsymbol{x}^i_{(t)}, X_{(t)})$ only via its aggregation $A^i_{(t)}$. The original set of constraints (20) did require knowledge of each individual inventory position $x^i_{j,(t)}$. However, after our relaxation of these constraints, only the aggregate system-wide inventory level $A^i_{(t)}$ matters, rather than its disaggregation across the different retailers and the depot. Furthermore, note that the minimization in this relaxed DP decomposes into two separate minimizations, one depending only on $\bar{\boldsymbol{x}}^i_{(t)}$ and $\bar{X}^i_{(t)}$, and one depending only on $\bar{W}^i_{(t)}$.

We can therefore define a new set of value functions $\hat{\underline{\theta}}^{i,\lambda}_{(t)}(A^i_{(t)}, \boldsymbol{W}^i_{(t)})$ that satisfy the following recurions

$$
\begin{aligned}
\hat{\underline{\theta}}^{i,\lambda}_{(t)}(A^i_{(t)}, \boldsymbol{W}^i_{(t)}) = R^{i,\lambda}_{(t)}(A^i_{(t)} + W^i_{(t-L)}) + \min_{\bar{W}^i_{(t)}\geq 0}\Big\{ K^i_{(t)}\mathbb{I}_{\bar{W}^i_{(t)}>0} + c^i_{(t)}\bar{W}^i_{(t)} \\
+\beta\mathbb{E}\left[\hat{\underline{\theta}}^{i,\lambda}_{(t+1)}(A^i_{(t)} - u^i_{+,(t)} + W^i_{(t-L)}, \boldsymbol{W}^i_{(t\to t+1)})\right]\Big\} - \frac{1}{I}\sum_{j=1}^{J}\lambda_{j,(t)}\chi_{j,(t)} \quad (27)
\end{aligned}
$$

where

$$R^{i,\lambda}_{(t)}(A) \quad = \min_{\bar{\boldsymbol{x}}^i_{(t)}, \bar{X}^i_{(t)}} \quad H^i_{(t)}(\bar{X}^i_{(t)}) + \sum_{j=1}^{J} Q^i_{j,(t)}(\bar{x}^i_{j,(t)}) + \lambda_{j,(t)} \max(\bar{x}^i_{j,(t)} - \tilde{u}^{i,[\alpha_j]}_{j,(t)}, 0) \tag{28}$$

$$\text{s.t.} \qquad \bar{x}^i_{+,(t)} + \bar{X}^i_{(t)} = \tilde{A} \tag{29}$$

$$\bar{X}_{j,(t)} \geq 0 \tag{30}$$

We have therefore successfully reduced the dimension of our dynamic program from $(J+L+1)$ to $(L+1)$. Furthermore, since the functions $Q^i_{j,(t)}(\cdot)$ and $H_{(t)}(\cdot)$ are strictly convex, it is easily verified that each $R^{i,\lambda}_{(t)}(\cdot)$ is strictly convex as well. This dynamic program therefore has a form of a simple single location, single item inventory problem, in which the cost term that depends on the system-wide inventory position is *convex*.

Finally, we can reduce the $(L+1)$-dimensional program to a 1-dimensional program by using a similar accounting scheme to the one we used the construct the functions $Q^i_{j,(t)}(\cdot)$. Instead of assigning at the beginning of period t the cost $R^{i,\lambda}_{(t)}(A^i_{(t)} + W^i_{(t-L)})$ as a function of the system-wide inventory known at that time, we assign to period t the discounted expectation of this cost value a lead time later, which is a function of the system-wide inventory position at time t. We refer the reader to Federgruen, Guetta, and Iyengar (2015) for details of this final reduction in the state space.

6 An Upper Bound: a Heuristic Policy

In this section, we propose a specific feasible strategy to govern the distribution system. The expected cost, under this policy, is of course an *upper bound* on the optimal cost value, and this for any starting state and time horizon. While relatively simple, its cost performance is still too difficult to be assessed analytically. Instead, we evaluate this with Monte Carlo simulations. Extensive numerical studies reported in section 8 compare the *lower bound*, resulting from the approximate DP in the previous section, with this *upper* bound.

A replenishment strategy consists of three components. (a) An *order policy* which dynamically prescribes when a system-wide order is to be placed with the external supplier for each product, and of what size; (b) an *allocation policy* which prescribes how much inventory, if any, is to be withdrawn from the depot, and how this inventory should be allocated among the different retailers. As to the order policy, it is natural to adopt the $\left\{(s^i_t, S^i_t) : t = 1, \cdots, T\right\}$ policies which are optimal in the approximate DP for each item $i = 1, \cdots, I$.

Given the finite-horizon setting in which we are operating, a different policy parameter pair (s_t^i, S_t^i) is to be determined for each period. These can be obtained by solving a regular one-dimensional DP for each product, albeit with various possible simplifications. For example, if for a given inventory level Z^0, it is optimal not to place an order, the same is true for any $Z \geq Z^0$. If, on the other hand, it is optimal to place an order which elevates the inventory position to some level S, then the same is true for all $Z \leq Z^0$.

It is close-to-optimal to adopt the order policies that are optimal in the lower bound DPs as part of the combined order, withdrawal and allocation strategy. After all, these order policies remain *feasible* in the *actual* DP, and optimal in a relaxation thereof. The same cannot be said for the remainder of the strategy. More specifically, the allocation strategies that are optimal in the lower bound DPs *ignore* the joint storage constraints (7), and allow for relaxation of the constraint (9), which ensure that new shipments to the retailers are non-negative.

Indeed, in the lower bound DP, it is optimal to determine all withdrawal and allocation quantities for each of the items *separately*, and in a manner that minimizes a *myopic* allocation problem, i.e., one that minimizes the expected costs in the very period period in which the shipments are received by the retailers, and, implicitely allowing for negative shipments. In other words, the allocation strategies that are optimal in the lower bound DPs suffer from three deficiencies. They ignore (a) the interdependencies among the items, (b) the cost consequences in later periods beyond the very first period in which the shipments are received, and (c) the shipment quantities' non-negativity constraints.

We therefore design a new type of allocation policy which addresses all three of the above differences; (i) it simultaneously determines allocations across *all* items, explicitly reinstating the joint storage constraints, for each retailer and each time period, therefore addressing the interdependencies among the various items. All the constraints in (9) are also reinstated, ensuring that shipment quantities remain non-negative. (ii) it determines the current period's allocations so as to minimize total expected discounted costs over an appropriate chosen planning window of κ periods, rather than only focusing myopically on the *first* period in which the current allocations have an impact. (iii) When minimizing total expected discounted costs over the full planning window, current allocations are made bearing in mind the possibility that any inventory remaining at the depot may be re-allocated in later periods in the window.

We therefore consider an allocation problem in which we make simultaneous allocations across all items in period t considering the cost impact of our decisions not only in period $t + \ell$, but also in periods $t+\ell+1,\cdots,t+\kappa$ (i.e. over $\kappa+1$ periods).

Specifically, we assume that once we have made our allocations in period t, any inventory remaining in the depot (or any inventory in our pipeline $\mathbb{W}_{(t)}$ that will arrive later at the depot) might later be assigned to some or all of the retailers in periods $t+1,\cdots,t+\kappa$.

Define the following decision variables, using a dot instead of a bar to distinguish these variables from those in the full DP:

- $\dot{x}^i_{j,(\tau)}$: the inventory position of product i at retailer j, at the beginning of period t, plus all shipments of this product, to this retailer, introduced in periods $t,t+1,\cdots,\tau$.
- $\dot{X}^i_{(\tau)}$: the inventory of product i assigned to the depot, at the start of period τ.

Given a specific set of allocations, our program minimizes expected discounted costs in periods $[t+\ell,t+\ell+\kappa]$. To express these costs succintly, let $u^i_{j,(t\to\tau)}$ denote the sum of all demands from time t through time $\tau+\ell_j$, i.e. $u^i_{j,(t\to\tau)}=\sum_{\tau'=t}^{\tau+\ell_j}u^i_{j,(\tau')}$. Note that $u^i_{j,(t\to t)}=\tilde{u}^i_{j,(t)}$. The product i inventory level at retailer j at the *end* of period τ is then given by $\dot{x}^i_{j,(\tau)}-u^i_{j,(t\to\tau)}$. We can therefore express the holding and backlogging costs at the end of that period as a function of $\dot{x}^i_{j,(\tau)}$ only; we call this function $Q^i_{j,(t\to\tau)}$, where

$$Q^i_{j,(t\to\tau)}(\dot{x}^i_{j,(\tau)})\equiv p^i_{j,(\tau+\ell)}\mathbb{E}\left(\dot{x}^i_{j,(\tau)}-u^i_{j,(t\to\tau)}\right)^-+h^i_{j,(\tau+\ell)}\mathbb{E}\left(u^i_{j,(t\to\tau)}-\dot{x}^i_{j,(\tau)}\right)^+ \tag{31}$$

It is, again, easy to verify that these functions are strictly convex.

In the true system, as reflected by the exact DP, the future values of the quantities $\{\dot{x}^i_{j,(tau)},\dot{X}^i_{(\tau)},\tau=t+1,\cdots,t+\kappa\}$ are determined based on the most updated information available, which includes all demands observed between time t and period τ. In the κ-allocation problem, thes quantiites are determined statically, upfront, as a function of the information available up to period t, i.e., as a function of the state of the system at the beginning of period t: $(\mathbb{x}_{(t)},\boldsymbol{X}_{(t)},\mathbb{W}_{(t)})$.

We are now ready to formulate a so called κ-allocation mathematical program that minimizes expected discounted costs in the next κ periods, assuming we are currently at the start of period t and in state $(\mathbb{x}_{(t)},\boldsymbol{X}_{(t)},\mathbb{W}_{(t)})$:

(κA) - κ-*Allocation*

$$\min_{\substack{\dot{\mathbb{x}}_{(t)},\cdots,\dot{\mathbb{x}}_{(t+\kappa)}\\ \dot{\boldsymbol{X}}_{(t)},\cdots\dot{\boldsymbol{X}}_{(t+\kappa)}}} \left\{\sum_{\tau=t}^{t+\kappa}\sum_{i=1}^{I}\left(H^i_{(\tau)}(\dot{X}^i_{(\tau)}) + \sum_{j=1}^{J} Q^i_{j,(t\to\tau)}(\dot{x}^i_{j,(\tau)})\right)\right\} \tag{32}$$

s.t.

$$\dot{x}^i_{+,(t)} + \dot{X}^i_{(t)} = x^i_{+,(t)} + X^i_{(t)} + W^i_{(t-L)} \qquad \forall i \tag{33}$$

$$\dot{x}^i_{+,(\tau)} + \dot{X}^i_{(\tau)} = \dot{x}^i_{+,(\tau-1)} + \dot{X}^i_{(\tau-1)} + W^i_{(\tau-L)} \qquad \forall i, \tau > t \tag{34}$$

$$\sum_{i=1}^{I} \max(\dot{x}^i_{j,(\tau)} - \tilde{u}^{i,[\alpha_j]}_{j,(t\to\tau)}, 0) \le \chi_{j,(\tau+\ell_j)} \qquad \forall j, \tau \tag{35}$$

$$\dot{x}^i_{j,(t)} \ge x^i_{j,(t)} \qquad \forall i, j \tag{36}$$

$$\dot{x}^i_{j,(\tau)} \ge \dot{x}^i_{j,(\tau-1)} \qquad \forall i, j, \tau > t \tag{37}$$

$$\dot{X}^i_{(\tau)} \ge 0 \qquad \forall i, \tau \tag{38}$$

where constraints over $\forall i$ are taken for $i = 1, \cdots I$, over $\forall j$ are taken for $j = 1, \cdots, j$ and over $\forall \tau$ are taken for $\tau = t, \cdots, t + \kappa$.

Constraints (33) are identical to (6) in the exact DP and ensure that, in period t, we only assign inventory that is available in that period. Constraints (34) reflect the fact that, for any item i, only $W^i_{(\tau-L)}$ (i.e., the order that is received at the beginning of period τ) can be used to augment $\dot{x}^i_{j,(\tau)} + \dot{X}^i_{j,(\tau)}$ from its value in the period period $(\tau - 1)$. Constraints (35) are identical to the original capacity constraints (7). Constraints (36) and (37) ensure that all shipments to the retailers are positive. Finally, constraint (38) ensures that no backlogs are allowed at the depot in any period, again identical to (8) in the original DP.

Note that even though this optimization problem determines allocations $(\dot{\mathbb{x}}_{(\tau)}, \dot{\boldsymbol{X}}_{(\tau)})$ for all periods $\tau = t, \cdots, t + \kappa$, the only part of these allocations to be implemented is $(\dot{\mathbb{x}}_{(t)}, \dot{\boldsymbol{X}}_{(t)})$, so as to make allocations in the current period t. In other words, the (κA)-problem is solved on a rolling horizon basis.

There are several possible choices for the length of the time window κ.

We first note that in period t, our pipeline of orders to the depot $\mathbb{W}_{(t)}$ contains orders that will be delivered there in periods $t, \cdots, t+L$. Therefore, one possible choice for κ is $\kappa = \kappa_1 \equiv L$, which optimizes over the next $L-1$ periods, and ensures that our non-myopic problem always accurately reflects the total inventory in the system. After period $t+L$, future orders placed after t may arrive at the depot, and our non-myopic problem would no longer correctly reflect the total inventory in our system in these periods.

Note that in (34), at time t, the magnitude of the incoming order $W^i_{(\tau-L)}$ is known for $\tau = t, \cdots, T+L$. Thus, any value of κ with $1 \leq \kappa \leq L$ could be chosen. In additional, it is possible to choose $\kappa > L$ and subsitute $W^i_{(t+L+1)} = \cdots = W^i_{(t+\kappa-L)} = 0$, if it is reasonble to expect that the next oder for item i will not be placed before time $(t+\kappa-L+1)$. Indeed, the exact *expected* number of periods with the next order can be computer as follows. Let $Z^i_{(t)}$ denote the system-wide inventory position for item i, in period t, and define

$$\begin{aligned}
\epsilon^i_{(t)}(Z^i_{(t)}) &= \text{The expected number of periods until the next order is placed} \\
&\quad\ \text{when the current initial system-wide inventory position is } Z^i_{(t)} \\
&= \mathbb{E} \sum_{\tau=(t+1)}^{T-L} \mathbb{I}_{\{\text{Next order does not occur in period } \tau\}} \qquad (39) \\
&= \sum_{\tau=(t+1)}^{T-L} \mathbb{P}\left(Z^i_{(t)} - \sum_{\tau'=t}^{\tau-1} u^i_{+,(\tau')} \geq s^i_{(\tau)} \right) \qquad (40)
\end{aligned}$$

When the demand distributions are discrete, $\epsilon^i_{(t)}(Z^i_{(t)})$ may be computed from a standard (non-stationary) renewal equation recursion. Whether we choose $\kappa^i = L^i$ (in the general case where leadtimes are item-dependent) or $\kappa^i = L + \epsilon^i_{(t)}(Z^i_{(t)})$, both choices would result in an *item dependent* planning window, which is entirely feasible.

The allocation problem (κA) is a convex optimization problem. It has a non-linear objective function and one set of non-linear constraints (35). However, as explained in Section 3, this set of constraints can be linearized with the help of auxiliary variables. Similarly, when all cost functions $H^i_{(t)}(\cdot)$, $h^i_{j,(t)}(\cdot)$, and $p^i_{j,(t)}(\cdot)$ are piecewise linear or approximately piecewise linear functions at every retailer, the objection function of (κA) is piecewise linear and convex as well. This implies that the entire optimization problem (κA) can be formulated as a linear program, and solved with any number of standard LP solvers.

The allocation problem (κA) is essentially different and superior compared to the allocation mathematical programs employed in the literature, e.g. Federgruen and Zipkin (1984a) and Federgruen, Guetta, and Iyengar (2015). The latter, for example, determines the allocation quantities, in each period, by minizming the sum of discounted costs over a given planning window, but in doing so, assumes that no additional allocations with be made after the initial period t.

7 Cost Interdependencies Among Items

In our base model, we assumed that the interdependencies among the items is confined to the joint storage constraints at the various facilities. In practice, other types of interdependencies prevail, most specifically *joint order costs*. Here is the simplest such setting: assume there are S suppliers, each selected to procure a distinct subset of the I items. Let $I(s)$ constain the set of items purchased from supplier s, $s = 1, \cdots, S$. The item collection $\{I(1), \cdots, I(S)\}$ is assumed to be a *partition* of the full item set $1, \cdots, I$. In particular, we assume that each item is procured from a *single* supplier.

The simplest, but frequently prevalent, join cost structure assumes that, in every period t, there exists a fixed order cost $K^{0,s}_{(t)}$ that is incurred whenever some order is placed with supplier s, irrespective of its composition, and *in addition* to the item-specific fixed order costs $\{K^i_{(t)}\}$ already incorporated in our base model. The above cost structure has been studied, extensively, in the single-location case, where the planning problem is referred to as the "joint replenishment problem" (JRP). Scores of papers have been devoted to this specific model, both under deterministic and stochastic demands. See Khouja and Goyal (2008) for a review of the literature up to 2005. Even under deterministically known demands, the problem is NP-complete, a result shown by Arkin, Joneja, and Roundy (1989).

The mere Lagrangian relaxation of the storage constraints (7) in the exact DP is no longer sufficient to decompose the problem on an item-by-item basis. To arrive at a *decoupled lower bound* DP, a different relaxation needs to be added. We propose adopting the approach of Atkins and Iyogun (1988) for the single location JRP. The idea is to allocate the joint order cost $K^{0,s}_{(t)}$ among the various items supplied by supplier s, in given proportions $\{\alpha^i_{(t)} : \sum_{i \in I(s)} \alpha^i_{(t)} = 1\}$, thereby replacing the true cost structure with an approximate one in which the item-dependent fixed costs are adjusted as follows

$$\underline{K}^i_{(t)} = K^i_{(t)} + \alpha^i_{(t)} K^{0,s}_{(t)} \qquad\qquad i \in I(s), t = 1, \cdots, T \tag{41}$$

This relaxation reduces the problem back to an instance of the model in Section 3, the solution of which is a *lower bound*, since at any given order epoch, the total order cost in the approximate model is bounded from above by the true costs. Together with the Lagrangian relaxation of constraints (7), this gives rise to a set of lower bound value functions $\underline{\theta}^{i,\alpha,\lambda}_{(t)}(\boldsymbol{x}^i_{(t)}, X^i_{(t)}, \boldsymbol{W}^i_{(t)})$. The best lower bound of this type is now obtained by maximizing over the vectors $\boldsymbol{\lambda}$ and $\boldsymbol{\alpha}$:

$$
\begin{array}{llll}
\text{(D)} & \max_{\boldsymbol{\lambda},\boldsymbol{\alpha}} & \sum_{i=1}^{I} \underline{\theta}_{(t)}^{i,\alpha,\lambda}(\boldsymbol{x}_{(t)}^{i}, X_{(t)}^{i}, \boldsymbol{W}_{(t)}^{i}) & \\
 & \text{s.t.} & \lambda_{j,(t)} \geq 0 & \forall j, t \\
 & & \sum_{i \in I(s)} \alpha_{(t)}^{i} = 1 & \forall s, t \\
 & & \alpha_{(t)}^{i} \geq 0 & \forall i, t
\end{array}
$$

Analogous to the proof of Theorem 2, one can show that the lower bound value functions $\underline{\theta}_{(t)}^{i,\alpha,\lambda}(\cdot)$ are *jointly concave* in the vector $\boldsymbol{\alpha}$, and that they are differentiable as well. As in the base model, let $\underline{V}_{(1)}^{\alpha,\lambda}(\cdot) = \sum_{i=1}^{I} \underline{\theta}_{(1)}^{i,\alpha,\lambda}(\cdot)$. Then, for each supplier s,

$$
\frac{\partial \underline{V}_{(1)}^{\alpha,\lambda}(\mathbb{x}_{(1)}, \boldsymbol{X}_{(1)}, \mathbb{W}_{(1)})}{\partial \alpha_{(t)}^{i}} = \mathbb{P}(\text{An order from supplier } s \text{ is placed in period } t \text{ under the strategy which is optimal in the lower bound problems}) \quad (42)
$$

We propose to solve the dual problem (D) in an iterative way: for a given allocation vector α, find the corresponding vector of optimal Lagrange multipliers $\boldsymbol{\lambda}(\boldsymbol{\alpha})$ with the steepest ascent method of Section 5. Then, using this set of Lagrange multipliers, carry out a search for the optimal α-vector, using a steepest ascent method for constrained problems, once again using the fact that the gradient (with respect to $\boldsymbol{\alpha}$) is easily computed via (42).

In the base model, the order policy component of the overall replenishment strategy could be chosen, straightforwardly, as the (s, S) order policies which are optimal in the best lower bound DPs; this in contrast to the withdrawal and allocation policy components for which we developed a new approach based on a mathematical program. Under joint order costs, this is no longer effective. After all, under the $(s_{(t)}, S_{(t)})$ policies in the lower-bound DPs, no incentive is provided to combine orders pertaining to different items. Instead, we propose adopting a $(T^i, v_{(t)}^i)$ policy for each item i, under which the system-wide inventory position of the item is replenished every T^i periods. If period t is a replenishment epoch, an order is placed to raise the system-wide inventory position to a target level $v_{(t)}^i$.

The replenishment cycles $\{T^i : i \in I(s)\}$ are to be selected as the optimal power-of-two policy in the setting where all demands are deterministic and *constant* over time, as well as all cost parameters. This optimal vector of power-of-two intervals $\{T^i : i \in I(s)\}$, $s = 1, \cdots, S$ can be found with a simple $O(T \log T)$ procedure, see e.g. Jackson, Maxwell, and Muckstadt (1985) and Roundy (1985). The latter proved that in a stationary, deterministic setting, the optimal power-of-two policy is guaranteed to come within 6% of optimality. Note that when all replenishment cycles are power-of-two multiples of the same base period, maximal coordination of the order process is achieved, and hence maximum advantage is taken of the economies of scope associated with the joint order costs. As to the withdrawal and allocation part of the replenishment strategy, the same heuristic can be employed as the one that is described in Section 6.

Finally, the joint order costs may have a more complex structure than those described in this section. Most generally, the fixed order costs associated with an order may depend on the precise composition of the order, according to a general *submodular* set function $K^s_{(t)}(I)$, where $I \in I(s)$ is the specific set of items included in the order. It is still possible to adopt the above relaxation approach, replacing the joint cost cost structure by a separable one, that results in a *lower bound*. See Federgruen and Zheng (1992b) or Viswanathan (2007) for details.

8 Numerical Study

In this section, we describe an extensive numerical study, consisting of 1,162 problem instances, designed to assess the accuracy of our lower bound DP derived in section 5, and the optimality gap of the proposed heuristic policy, described in section 6.

All instances consider a horizon of $T = 20$ periods, and a system with $J = 8$ retailers, each carrying $I = 7$ items. All instances use linear holding and backlogging costs at the depot and the retailers, and we assume, throughout, that the storage space at the depot is ample. All demand distributions are Normal; however, at various parts of the lower and upper bound calculations, we employ a fine discretization of the relevant Normal distribution, employing a 49 point approximation. The cost performance of the proposed heuristic strategy is evaluated via Monte Carlo simulation, using a sample of 40 paths cosisting of 20 periods each, and solving an instance of the κ-allocation mathematical program in each period for each of the generated scenarios.

The median computational time required to calculate any given instance was 167 CPU minutes. This number relates to a current implementation of our algorithm, with much of the work done in R, and, as mentioned, using a very fine 49 point discretization of the underlying demand distribution. The current computational times explain the size of our numerical study. We are confident that a significantly faster implementation can be achieved by implementing our entire system in a lower-level language like C++, along with various other speedups.

Table 1 summarizes the parameters and meta-parameters that are used to generate the problem instances in our numerical study.

Parameter	Values	
T	20	The number of periods in the horizon.
J	8	The number of retailers. A value of 8 gives a distribution network with 9 locations.
I	7	The number of items.
L	3, 4	The leadtime from the supplier to the depot.
ℓ	2, 3	The leadtime from the depot to the retailers.
$c_{i,(t)} = c$	5	The variable procurement cost from the supplier.
`overPenalty`	0.5, 1	A 'meta-parameter' specifying the ratio of the holding cost rate at the depot to the *smallest* of the holding rates at the retailers. The same ratio is used for all items.
`chiBase`	-1, 5 1000	A 'meta-parameter' used to specify the capacity at each retailer.
`targetEpochs`	3, 7	A 'meta-parameter' used to calculate the fixed cost of procuring each item from the supplier. See the text of the paper for details.
`baseCostRatio`	4, 10	A parameter used to specify the ratio of backorder to holding costs. See the text of the paper for details.
`CVBase`	0.15, 0.3, 0.4	A 'meta-parameter' used the specify the coefficient of variation of the single-period demand at each retailer.

Table 1: Parameters used in our numerical study. See the body of the paper for more details.

The remaining parameters are generated as follows.

Cost Ratios

We investigate two configurations for the holding and backlogging cost rates for each item at the various retailers.

Proportional costs : All of the holding cost rates $h^i_{j,(t)}$ are chosen uniformly and independently from the interval $[8, 12]$. All of the backlogging cost rates $p^i_{j,(t)}$ are specified as $p^i_{j,(t)} = \texttt{baseCostRatio} \cdot h^i_{j,(t)}$.

Random costs : All of the retailers' holding costs are chosen uniformly and independently from the interval $[8, 12]$. Likewise, the backlogging costs are chosen independently and uniformly from the interval $[8, 12] \cdot$ `baseCostRatio`.

For each item $i = 1, \cdots, I$, holding cost rates at the depot are set to `overPenalty` times the smallest of the holding cost rates at the retailers:

$$H^i_{(t)} = \texttt{overPenalty} \cdot \min_j h^i_{j,(t)}$$

It is well-known that holding cost rates *increase* as we progress from one echelon in a supply network to the next. A large, often dominant part of the holding costs consists of capital costs. Similarly, large depots are typically located where rents and other real-estate costs are lower; unit costs at the depot are additionally reduced by virtue of economies of scale. It is therefore without loss of generality that our study limits the values of the `overPenalty` parameter to 1.

Retailers' Capacities

The retailers' capacities are specified to add a given multiple (`chiBase`) of the average standard deviation of aggregate demands, across all items, to the average mean aggregate demand in a single period. Here, the average is taken over all T periods in the planning horizon. Thus

$$\chi_{j,(t)} \equiv \chi_j \equiv \frac{1}{T}\sum_{\tau=1}^{T}\mu^{+}_{j,(\tau)} + \texttt{chiBase}\cdot\left[\frac{1}{T}\sum_{\tau=1}^{T}\sqrt{\sum_{i=1}^{I}(\sigma^{i}_{j,(\tau)})^2}\right]$$

Demand Distributions

As mentioned, all one-period demand distributions are Normal. The distribution pertaining to item i and retailer j, in preiod t, is $N(\mu^i_{j,(t)}, \sigma^i_{j,(t)})$. We investigated two ways to generate the means and standard deviations, assuming all demand distributions are independent of each other

Constant CV : All means $\mu^i_{j,(t)}$ are chosen, uniformly and independently, from the interval $[80, 120]$. The standard deviations $\sigma^i_{j,(t)}$ are specified to maintain *uniform* coefficients of variation across all demand distributions, i.e.,

$$\sigma^i_{j,(t)} = \texttt{CVBase}\cdot\mu^i_{j,(t)}$$

for all i, j, and t.

Random CV : All means $\mu^i_{j,(t)}$ are again chosen uniformly and independently from the interval $[80, 120]$. All coefficients of variation $\sigma^i_{j,(t)}/\mu^i_{j,(t)}$ are chosen independently and uniformly from the interval $[0.05, \texttt{CVBase}]$.

Fixed Order Costs

The fixed order costs are chosen to be stationary, i.e., $K^i_{(t)} = K^i$ for all t and i. They are selected so as to target a specific frequency of orders with the items' supplier, or, equivalently, a specific average cycle length between consecutive orders, determined by the parameter `targetEpochs`. In doing so, we assume that this average cycle length is approximately the optimal cycle length in the Economic Order Quantity model that arises when all demands occur at constant deterministic rates given by the average mean of aggregate demands across all periods

$$\mu^i_{+,(\mathrm{av})} = \frac{1}{T}\sum_{\tau=1}^{T}\sum_{j=1}^{J}\mu^i_{j,(\tau)} \qquad i = 1, \cdots, I \tag{43}$$

and a constant holding cost rate, given by the average of the items' holding cost rates across all retailers and periods, i.e.,

$$h^i_{(\mathrm{av}),(\mathrm{av})} = \frac{1}{TJ} \sum_{\tau=1}^{T} \sum_{j=1}^{J} h^i_{j,(\tau)} \tag{44}$$

The optimal cycle length is then given by the well known formula as

$$\texttt{targetEpochs} = \sqrt{\frac{2K^i}{\mu^i_{+,(\mathrm{av})} h^i_{(\mathrm{av}),(\mathrm{av})}}}$$

this results in the specification

$$K^i = \frac{\texttt{targetEpochs}^2 \cdot \mu^i_{+,(\mathrm{av})} h^i_{(\mathrm{av}),(\mathrm{av})}}{2}$$

Numerical Results

Our numerical results are arranged in 16 tables, each reporting on a set of 72 instances. Each table lists instances with a specific combination of `targetEpochs`, `baseCostRatio`, L, and ℓ parameters, and in each cell, displays the ratio $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$, where UB denotes the expected cost of the proposed system-wide heuristic strategy and LB the value of the Lagrangian dual. This ratio is, of course, an upper bound for the strategy's optimality gap, and the accuracy gap of the lower bound.

In each table, there are three row sections corresponding to the three values of `chiBase`, which specify how much storage capacity is available to each retailer. Within each row section, we display two lines, corresponding to the two procedures used to generate the one-period demand distributions.

The twelve columns in our table are arranged in three column sections, corresponding to the three `CVBase` values used in this study. Within each of these sections, we differentiate between the Proportional and Random cost methods discussed above to generate the retailers' holding and backlogging costs used in each period, as well as the two `overPenalty` values that determine the holding costs at the depot.

Tables 2 and 3 are two of the tables from our numerical studies. The remaining 14 tables are provided in an appendix.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase = −1	Const. CV	0.18	3.02	0.19	3.1	1.09	1.61	1.13	1.58	1.25	1.58	1.28	1.68
	Rand. CV	0.08	4.49	0.1	4.65	0.1	4.14	0.17	3.16	0.12	3.62	0.4	2.31
chiBase = 5	Const. CV	0.18	2.05	0.19	2.87	0.63	1.19	0.58	1.26	0.97	1.43	0.98	1.44
	Rand. CV	0.08	3.72	0.1	3.48	0.1	2.64	0.17	2.44	0.12	2.2	0.25	1.83
chiBase = 1000	Const. CV	0.18	0.65	0.19	0.65	0.56	1.2	0.5	1.1	0.84	1.63	0.8	1.69
	Rand. CV	0.08	0.63	0.1	0.63	0.1	0.6	0.17	0.64	0.12	0.6	0.25	0.72

Table 2: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 4$, $\ell = 3$, targetEpochs $= 7$, and baseCostRatio $= 10$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase = −1	Const. CV	0.16	0.94	0.15	0.94	0.84	1.49	0.86	2.19	2.01	2.46	1.24	4.51
	Rand. CV	0.11	2.72	0.1	1.71	0.16	1.66	0.18	0.89	0.2	1.73	0.25	0.6
chiBase = 5	Const. CV	0.16	1.04	0.15	1.57	0.52	1.47	0.52	1.46	0.93	2.06	0.98	1.84
	Rand. CV	0.11	2.01	0.1	2.19	0.16	1.44	0.18	0.92	0.2	1.5	0.25	1.14
chiBase = 1000	Const. CV	0.16	0.67	0.15	0.6	0.52	1.46	0.52	1.46	0.93	2.21	1	2.32
	Rand. CV	0.11	0.74	0.1	0.59	0.16	0.7	0.18	0.6	0.2	0.67	0.25	0.78

Table 3: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 4$, $\ell = 2$, targetEpochs $= 3$, and baseCostRatio $= 4$.

Our overall conclusion is that the lower bound is remarkably accurate, and the proposed heuristic is close-to-optimal over all investigated problem instances. When the `overPenalty` is 0.5 (i.e., when it is significantly cheaper to carry inventories at the depot) the above mentioned gaps $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ are 2% or lower, across the board! When the `overPenalty` is 1, so that storage at the depot is as expensive at storage at the retailers, the gaps are slightly larger, but still remain below 8%, and are usually considerably smaler.

Comparing consecutive row sections, one discovers that almost always, gaps improve as more storage capacity is available. As can be expected, the gaps are larger when the order cycles are larger; this can be observed by changing the `targetEpochs` meta-parameter. This phenomenon was observed in the earlier models reviewed in the Introduction and Literature Review: as the order cycles get larger, there are less frequent opportunities to rectify emerging imbalances across the inventories. Finally, the gaps decrease as the backlogging costs decrease relative to the holding costs; this phenomenon was observed by varying the `baseCostRatio` meta-parameter. There is no apparent pattern for the dependence of the gaps on the leadtime parameters.

In addition to calculating the gaps $100 \cdot (\text{UB} - \text{LB})/\text{LB}$, we also estimated the overflow probabilities at the retailers for each problem instance. Here, the overflow probability is defined as the percentage of periods, over all T periods observed across all simulations, in which at least one retailer experiences an inventory overflow. The largest observed overflow probability over all our instances was 0.77%. The value of α used for each of these instances was 0.1.

As discussed, the relationship between the α parameter and the overflow probability is unfortunately difficult to characterize, and depends on the demand distributions, the capacities at the retailers, and the other parameters of the problem. As mentioned, the probability of an overflow clearly increases monotonically with the parameter α, so any given probability of overflow can be achieved or closely approximated with a simple bi-section method with respect to the α parameter. (When the demand distributions are continuous, any targeted overflow probability can be achieved exactly. When they are discrete, the overflow probability is an increasing step function of α and may overshoot the demand overflow probability threshold without achieving it at any value of α.)

To illustrate this point, we focus on one specific instance, with `CVBase` $= 0.4$, `depotCost` $= 1$, `chiBase` $= 5$, `baseCostRatio` $= 10$, `targetEpochs` $= 3$, random coefficients of variation, random costs, $L = 4$, and $\ell = 3$. We run this instance 100 times, with values of α ranging from 0.01 to 0.6 on an equally spaced grid. The same seed was used the generate the sample paths required to calculate the upper bound in all cases, to simplify comparison between instances. In Figure 1, we

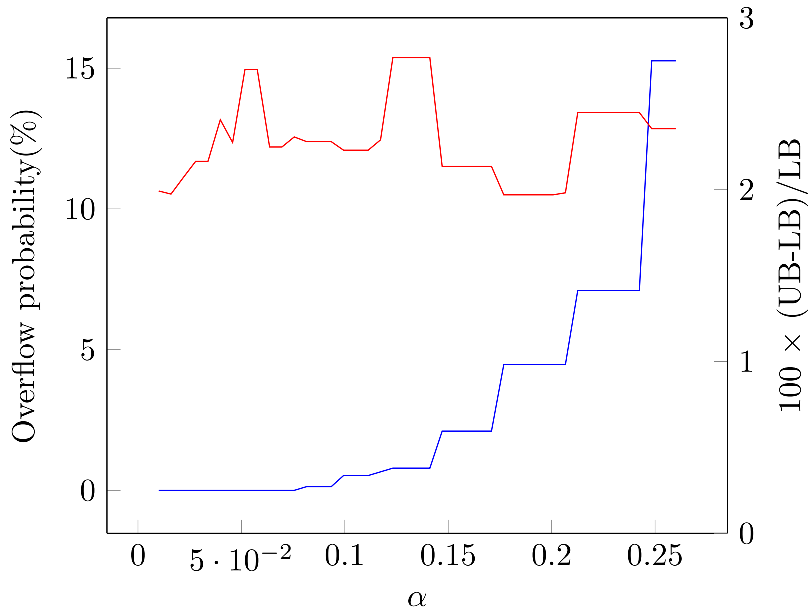

Figure 1: Investigating the effect of α on overflow probability. The blue line (left axis) plots the overflow probability for each of the instances in our study. The red line (right axis) plots the performance gap between the upper and lower bound for each instance. We restrict our plot to instances with $\alpha \leq 0.26$. Overflow probabilities above this fractile were too large to be of practical interest without resulting in significant cost improvements.

plot the oveflow probability and the gap $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ against α. In Figure 2, we plot the value of the upper bound against α.

We also observe that there is no obvious pattern for the dependence of the optimality gaps on the parameter α. This is reassuring – regardless of the overflow probability required, the gap remains roughly constant and small. Finally, it is informative to consider the value of the upper bound as a function of α. Clearly, the smaller the value of α (ie: the more stringently we apply our capacity constraint) the higher the cost of our heuristic policy.

9 Conclusion

In this paper, we have developed a methodology to determine effective replenishment strategies for two-echelon distribution systems, consisting of a depot and an arbitrary number of retailers, each selling a variety of items. The interdependencies among the facilities are as follows (i) standard economies of scale in the order costs connect the procurement strategies of the various retail outlets; (ii) system-wide orders are to be determined on the basis of *all* the retailers' and the depot's information; moreover, the fact that the replenishment process is conducted in the above two

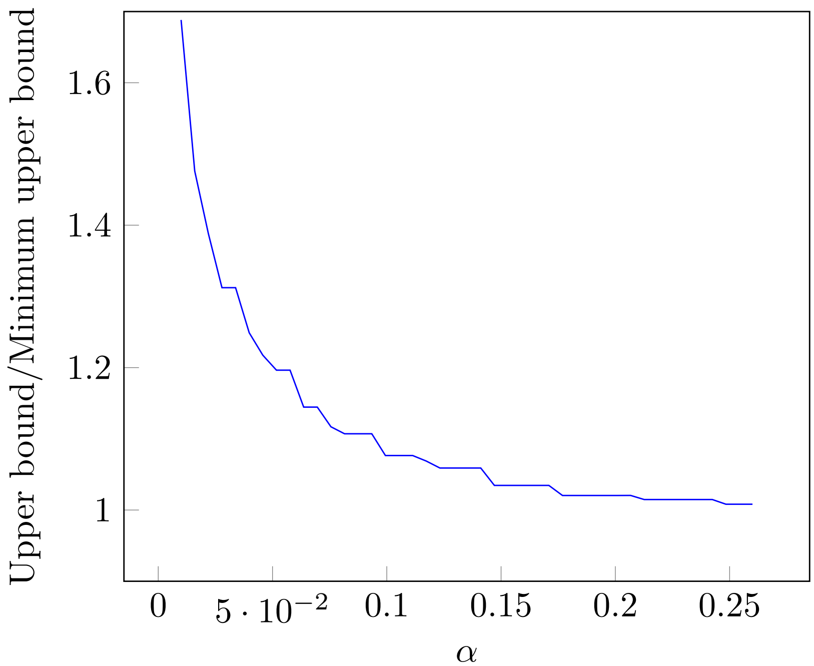

Figure 2: Investigating the effect of α on system costs. This graph plots the cost of the upper bound for each of the instances in our study, as a multiple of the smallest observed cost. We restrict our plot to instances with $\alpha \leq 0.26$. Overflow probabilities above this fractile were too large to be of practical interest without resulting in significant cost improvements.

phases allows us to exploit *statistical economics of scale*: when deciding on the *first stage* order quantity, only *aggregate* demands matter, the volatility of which is relatively lower than that at the individual retailers, thus requiring smaller safety stocks to reach a given service level; (iii) the existence of an inventory pool at the depot, from which all the retailers can draw, creates an additional interdepecency among the retailers. Similarly, the planning process for the different items interact with each other, at least because of (a) *shared storage capacity* at each of the facilities, (b) *economies of scope* in the order costs, for example, fixed cost components incurred irrespective of the composition of the order.

We have developed a new methodology to replace the fully intractable exact DP by a tractable series of lower bound DPs. This methodology combines several types of relaxations and finds the best bound within the resulting class of lower bounds. Of equal importance is the development of a new approach to determine withdrawal and allocation quantities, in each period, by solving a multi-period convex allocation program on a rolling horizon basis. An extensive numerical study shows that the lower bound is close to accurate, and the proposed replenishment strategy is close to optimal, throughout the explored parameter space. The maximum gap between the upper bound and the lower bound in our study of 1,162 instances is 8%, but the average gap is 1.2%, a significantly lower number.

The applications of the above results go beyond the derivation of effective replenishment strategies for multi-item two-echelon distribution systems. The availability of close-to-accurate lower bounds allows one to explore a series of strategic questions, for example the impact of leadtime reductions, changes in the cost functions, investment in storage capacity, the addition of fulfilment centers, etc...

References

Adelman, D. and A. J. Mersereau (2008). "Relaxations of weakly coupled stochastic dynamic programs". In: *Operations Research* 56.3, pp. 712–727.

Arkin, E., D. Joneja, and R. Roundy (1989). "Computational complexity of uncapacitated multi-echelon production planning problems". In: *Operations Research Letters* 8.2, pp. 61–66.

Arrow, K. J., T. Harris, and J. Marschak (1951). "Optimal inventory policy". In: *Econometrica: Journal of the Econometric Society*, pp. 250–272.

Atkins, D. R. and P. O. Iyogun (1988). “Periodic versus “can-order” policies for coordinated multi-item inventory systems”. In: *Management Science* 34.6, pp. 791–796.

Axsäter, S. (2003). “Supply Chain Operations: Serial and Distribution Inventory Systems”. eng. In: ed. by S C Graves and T de Kok. Supply chain management : design, coordination and operation. Elsevier, pp. 525–559. ISBN: 0-444-51328-0.

Axsäter, S., J. Marklund, and E. A. Silver (2002). “Heuristic Methods for Centralized Control of One-Warehouse, N-Retailer Inventory Systems”. In: *Manufacturing & Service Operations Management* 4.1, pp. 75–97. DOI: 10.1287/msom.4.1.75.291. URL: http://pubsonline.informs.org/doi/abs/10.1287/msom.4.1.75.291.

Clark, A. J. and H. Scarf (1960). “Optimal policies for a multi-echelon inventory problem”. In: *Management science* 6.4, pp. 475–490.

Doğru, M. K. (2005). “Optimal control of one-warehouse multi-retailer systems: An assessment of the balance assumption.” PhD thesis. Technische Universiteit Eindhoven, Eindhoven, The Netherlands.

Dogru, M. K., A. G. De Kok, and G. J. Van Houtum (2004). “Optimal control of one-warehouse multi-retailer systems with discrete demand”. In: *Dept. of Technology Management, Technische Universiteit Eindhoven*.

Dvoretsky, A., J. Kiefer, and J. Wolfowitz (1953). “On the optimal character of the (s, S) policy in inventory policy”. In: *Econometrica, October*.

Eppen, G. L. and L. Schrage (1981). “Centralized Ordering Policies in a Multi-Warehouse System with Lead Times and Random Demand”. In: *TIMS Studies in the Management Sciences* 16, pp. 51–67.

Federgruen, A., D. Guetta, and G. Iyengar (2015). “Two Echelon Distribution Systems with Random Demands and Storage Constraints”. Columbia University Working Paper.

Federgruen, A., J. Meissner, and M. Tzur (2007). “Progressive interval heuristics for multi-item capacitated lot-sizing problems”. In: *Operations Research* 55.3, pp. 490–502.

Federgruen, A., M. Queyranne, and Y-S. Zheng (1992). “Simple power-of-two policies are close to optimal in a general class of production/distribution networks with general joint setup costs”. In: *Mathematics of Operations Research* 17.4, pp. 951–963.

Federgruen, A. and Y-S. Zheng (1992a). “An efficient algorithm for computing an optimal (r, Q) policy in continuous review stochastic inventory systems”. In: *Operations research* 40.4, pp. 808–813.

Federgruen, A. and Y-S. Zheng (1992b). "The joint replenishment problem with general joint cost structures". In: *Operations Research* 40.2, pp. 384–403.

Federgruen, A. and P. Zipkin (1984a). "Allocation policies and cost approximations for multilocation inventory systems". In: *Naval Research Logistics Quarterly* 31.1, pp. 97–129.

— (1984b). "Approximations of dynamic, multilocation production and inventory problems". In: *Management Science* 30.1, pp. 69–84.

— (1984c). "Computational issues in an infinite-horizon, multiechelon inventory model". In: *Operations Research* 32.4, pp. 818–836.

Food Marketing Institute (2015). `http://www.fmi.org/research-resources/supermarket-facts`. Accessed: 2015-08-11.

Gallego, G., Ö. Özer, and P. Zipkin (2007). "Bounds, heuristics, and approximations for distribution systems". In: *Operations Research* 55.3, pp. 503–517.

Hawkins, J. T. (2003). "A Langrangian decomposition approach to weakly coupled dynamic optimization problems and its applications". PhD thesis. Massachusetts Institute of Technology.

Jackson, P., W. Maxwell, and J. Muckstadt (1985). "The joint replenishment problem with a powers-of-two restriction". In: *IIE transactions* 17.1, pp. 25–32.

Khouja, M. and S. Goyal (2008). "A review of the joint replenishment problem literature: 1989–2005". In: *European Journal of Operational Research* 186.1, pp. 1–16.

Kunnumkal, S. and H. Topaloglu (2008). "A duality-based relaxation and decomposition approach for inventory distribution systems". In: *Naval Research Logistics (NRL)* 55.7, pp. 612–631. ISSN: 1520-6750. DOI: `10.1002/nav.20306`. URL: `http://dx.doi.org/10.1002/nav.20306`.

— (2011). "Linear programming based decomposition methods for inventory distribution systems". In: *European Journal of Operational Research* 211.2, pp. 282–297.

Li, J. and J. Muckstadt (2014). "Managing inventories in multi-echelon on-line retail fulfillment system with different response lead time demands". In: *Computational Intelligence in Production and Logistics Systems (CIPLS), 2014 IEEE Symposium on*. IEEE, pp. 75–82.

— (2015). "Fulfilling Orders in a Multi-Echelon Capacitated On-line Retail System: PART TWO, real-time purchasing and ful llment decision making". In:

Maxwell, W. L. and J. A. Muckstadt (1985). "Establishing consistent and realistic reorder intervals in production-distribution systems". In: *Operations Research* 33.6, pp. 1316–1341.

Queyranne, M. (1985). *A polynomial-time sub-modular extension to Roundy's 98% effective heuristic for production*. Tech. rep. inventory systems. Working paper.

Roundy, R. (1985). “98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems”. In: *Management science* 31.11, pp. 1416–1430.

— (1986). “A 98%-effective lot-sizing rule for a multi-product, multi-stage production/inventory system”. In: *Mathematics of operations research* 11.4, pp. 699–727.

Shapiro, J. F. (1993). “Mathematical programming models and methods for production planning and scheduling”. In: *Handbooks in operations research and management science* 4, pp. 371–443.

Thomas, L. J. and J. O. McClain (1993). “An overview of production planning”. In: *Handbooks in operations research and management science* 4, pp. 333–370.

Viswanathan, S. (2007). “An algorithm for determining the best lower bound for the stochastic joint replenishment problem”. In: *Operations research* 55.5, pp. 992–996.

A Results for the Multi-Product Numerical Study

In this appendix, we list all numerical results in our study.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase = −1	Const. CV	0.18	0.64	0.2	0.52	1.07	2.04	0.86	2.79	2.09	3.49	1.39	2.88
	Rand. CV	0.13	1.09	0.15	0.65	0.19	2.6	0.26	0.52	0.22	0.66	0.35	1.02
chiBase = 5	Const. CV	0.18	1.85	0.2	0.91	0.56	1.35	0.56	1.24	1	2.08	1.09	2
	Rand. CV	0.13	1.22	0.15	0.91	0.19	0.98	0.26	0.84	0.22	0.87	0.34	1.15
chiBase = 1000	Const. CV	0.18	0.6	0.2	0.54	0.56	1.53	0.56	1.47	1.01	2.38	1.02	2.35
	Rand. CV	0.13	0.58	0.15	0.52	0.19	0.56	0.26	0.57	0.22	0.63	0.34	0.9

Table 4: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 4$, $\ell = 3$, targetEpochs $= 3$, and baseCostRatio $= 4$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.14	0.83	0.21	0.71	1.73	1.78	1.81	1.96	1.89	2.12	1.83	2.05
	Rand. CV	0.12	1.12	0.17	0.86	0.17	0.79	0.28	0.83	0.18	0.74	0.6	0.93
chiBase $= 5$	Const. CV	0.14	0.83	0.21	0.64	0.54	1.28	0.58	1.15	1.22	2.39	1.39	2.28
	Rand. CV	0.12	1.31	0.17	1.32	0.17	0.86	0.28	1.42	0.18	0.79	0.35	0.84
chiBase $= 1000$	Const. CV	0.14	0.6	0.21	0.58	0.45	1.73	0.49	1.57	0.88	2.69	0.96	2.57
	Rand. CV	0.12	0.61	0.17	0.56	0.17	0.6	0.28	0.59	0.18	0.59	0.35	0.74

Table 5: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 4$, $\ell = 3$, targetEpochs $= 3$, and baseCostRatio $= 10$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase = −1	Const. CV	0.68	1.84	0.69	2.12	0.65	1.41	0.69	1.78	0.72	1.64	0.76	1.65
	Rand. CV	0.79	2.48	0.78	2.95	0.73	2.15	0.67	2.17	0.68	1.96	0.64	1.85
chiBase = 5	Const. CV	0.68	2.04	0.69	1.58	0.67	1.25	0.66	0.79	0.8	1.03	0.82	1.03
	Rand. CV	0.79	2.43	0.78	2.33	0.73	1.98	0.67	2.2	0.68	1.56	0.63	1.15
chiBase = 1000	Const. CV	0.68	0.96	0.69	1.12	0.67	0.95	0.66	1.07	0.81	1.33	0.82	1.41
	Rand. CV	0.79	1.36	0.78	1.44	0.73	1.14	0.67	1.1	0.68	1	0.63	0.98

Table 6: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 4$, $\ell = 3$, targetEpochs = 7, and baseCostRatio = 4.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty $\rightarrow$	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.16	1.72	0.17	1.46	1.53	1.56	1.5	1.61	1.98	1.99	1.97	2.19
	Rand. CV	0.1	3.13	0.12	2.15	0.16	2.29	0.2	1.4	0.18	1.97	0.28	1.05
chiBase $= 5$	Const. CV	0.16	1.57	0.17	1.39	0.39	1.05	0.35	0.94	1.03	1.71	0.92	1.97
	Rand. CV	0.1	3.09	0.12	1.88	0.16	2.4	0.2	1.79	0.18	1.59	0.25	1.14
chiBase $= 1000$	Const. CV	0.16	0.74	0.17	0.61	0.39	1.57	0.35	1.62	0.83	2.67	0.98	2.65
	Rand. CV	0.1	0.83	0.12	0.63	0.16	0.78	0.2	0.61	0.18	0.76	0.25	0.7

Table 7: $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ results for $L = 4$, $\ell = 2$, targetEpochs $= 3$, and baseCostRatio $= 10$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty $\rightarrow$	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.76	4.23	0.78	4.41	0.83	3.05	0.84	3.17	0.88	2.67	0.91	2.85
	Rand. CV	0.79	5.28	0.81	5.03	0.76	4.81	0.76	4.4	0.73	4.54	0.74	3.81
chiBase $= 5$	Const. CV	0.76	3.4	0.78	3.28	0.77	1.29	0.77	1.43	0.89	1.09	0.89	1.38
	Rand. CV	0.79	4.78	0.81	4.52	0.76	4.24	0.76	3.62	0.73	3.72	0.74	2.48
chiBase $= 1000$	Const. CV	0.76	1.25	0.78	1.31	0.77	1.16	0.77	1.21	0.89	1.5	0.9	1.57
	Rand. CV	0.79	1.56	0.81	1.56	0.76	1.39	0.76	1.28	0.73	1.27	0.74	1.16

Table 8: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 4$, $\ell = 2$, targetEpochs $= 7$, and baseCostRatio $= 4$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase = −1	Const. CV	0.17	6.43	0.17	6.04	0.96	3.14	0.96	2.85	1.32	2.32	1.35	2.3
	Rand. CV	0.07	8.16	0.07	7.78	0.08	7.26	0.14	5.89	0.1	6.67	0.23	5.29
chiBase = 5	Const. CV	0.17	4.85	0.17	4.7	0.53	2.27	0.48	2.35	0.85	1.79	0.8	1.89
	Rand. CV	0.07	6.87	0.07	6.49	0.08	5.73	0.14	4.74	0.1	5.04	0.21	3.77
chiBase = 1000	Const. CV	0.17	0.69	0.17	0.71	0.53	1.23	0.48	1.23	0.8	1.78	0.76	1.8
	Rand. CV	0.07	0.71	0.07	0.68	0.08	0.69	0.14	0.7	0.1	0.68	0.21	0.79

Table 9: $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ results for $L = 4$, $\ell = 2$, targetEpochs $= 7$, and baseCostRatio $= 10$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty $\rightarrow$	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.2	0.54	0.19	0.69	0.94	3.53	0.89	3.7	1.77	3.83	1.56	3.43
	Rand. CV	0.14	1.47	0.15	2.98	0.21	1.1	0.26	0.57	0.26	0.73	0.36	0.92
chiBase $= 5$	Const. CV	0.2	2.33	0.19	0.73	0.63	1.22	0.59	1.26	1.03	3.61	1.02	2.54
	Rand. CV	0.14	1.05	0.15	0.9	0.21	1.08	0.26	0.87	0.26	1.07	0.35	0.83
chiBase $= 1000$	Const. CV	0.2	0.63	0.19	0.62	0.63	1.37	0.59	1.38	1.04	2.12	1	2.32
	Rand. CV	0.14	1.36	0.15	1.11	0.21	0.98	0.26	0.67	0.26	0.77	0.35	0.78

Table 10: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 3$, $\ell = 3$, targetEpochs $= 3$, and baseCostRatio $= 4$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase = −1	Const. CV	0.24	1.46	0.26	1.02	1.82	1.91	1.82	1.74	2.02	2.14	1.91	2.3
	Rand. CV	0.17	6.75	0.18	1.71	0.23	1.85	0.32	0.97	0.25	1.39	0.64	1.14
chiBase = 5	Const. CV	0.24	0.82	0.26	2.33	0.65	0.98	0.67	1.4	1.34	1.87	1.43	2.02
	Rand. CV	0.17	1.48	0.18	2.56	0.23	1.14	0.32	3.2	0.25	1.3	0.42	0.84
chiBase = 1000	Const. CV	0.24	1.24	0.26	1.24	0.55	1.53	0.54	1.51	1.06	2.71	1.12	2.54
	Rand. CV	0.17	2.58	0.18	2.13	0.23	1.88	0.32	1.26	0.25	1.53	0.42	1.03

Table 11: $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ results for $L = 3$, $\ell = 3$, targetEpochs $= 3$, and baseCostRatio $= 10$.

		$\texttt{CVBase} = 0.15$				$\texttt{CVBase} = 0.3$				$\texttt{CVBase} = 0.4$			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	$\texttt{overPenalty} \rightarrow$	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
$\texttt{chiBase} = -1$	Const. CV	0.72	2.04	0.73	2.48	0.76	1.82	0.78	2.02	0.78	1.81	0.81	1.69
	Rand. CV	0.78	3.4	0.78	3.46	0.75	2.92	0.72	2.46	0.72	2.7	0.71	1.82
$\texttt{chiBase} = 5$	Const. CV	0.72	1.83	0.73	1.8	0.73	1.67	0.72	0.92	0.86	1.67	0.87	1.28
	Rand. CV	0.78	3	0.78	2.62	0.75	2.42	0.72	1.75	0.72	1.95	0.7	1.34
$\texttt{chiBase} = 1000$	Const. CV	0.72	1.11	0.73	1.18	0.73	1.06	0.72	1.11	0.86	1.41	0.86	1.44
	Rand. CV	0.78	1.46	0.78	1.45	0.75	1.27	0.72	1.16	0.72	1.15	0.7	1.05

Table 12: $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ results for $L = 3$, $\ell = 3$, $\texttt{targetEpochs} = 7$, and $\texttt{baseCostRatio} = 4$.

		CVBase $= 0.15$				CVBase $= 0.3$				CVBase $= 0.4$			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty $\rightarrow$	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.21	3.53	0.21	3.42	1.18	1.59	1.19	1.62	1.3	1.54	1.3	1.69
	Rand. CV	0.09	5.23	0.11	5.2	0.12	4.79	0.2	3.38	0.15	3.95	0.43	2.29
chiBase $= 5$	Const. CV	0.21	2.5	0.21	2.81	0.67	1.43	0.61	1.59	1.04	1.3	1.02	1.68
	Rand. CV	0.09	4.22	0.11	4.12	0.12	3.32	0.2	2.82	0.15	2.9	0.28	1.95
chiBase $= 1000$	Const. CV	0.21	0.84	0.21	0.84	0.6	1.16	0.53	1.15	0.87	1.63	0.83	1.7
	Rand. CV	0.09	1.47	0.11	1.27	0.12	1.11	0.2	0.82	0.15	0.92	0.28	0.8

Table 13: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 3$, $\ell = 3$, targetEpochs $= 7$, and baseCostRatio $= 10$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty $\rightarrow$	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.16	0.87	0.15	0.82	1.29	1.82	1.19	1.48	1.64	2.59	1.18	3.61
	Rand. CV	0.1	2.57	0.11	2.49	0.16	1.56	0.19	0.84	0.19	1.19	0.27	0.68
chiBase $= 5$	Const. CV	0.16	2.57	0.15	1.52	0.5	1.28	0.47	1.17	0.86	1.85	0.86	1.81
	Rand. CV	0.1	2.64	0.11	1.46	0.16	3.48	0.19	1.76	0.19	1.96	0.27	0.97
chiBase $= 1000$	Const. CV	0.16	0.8	0.15	0.79	0.5	1.49	0.47	1.51	0.86	2.26	0.87	2.22
	Rand. CV	0.1	1.58	0.11	1.23	0.16	1.18	0.19	0.83	0.19	0.96	0.27	0.9

Table 14: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 3$, $\ell = 2$, targetEpochs $= 3$, and baseCostRatio $= 4$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty $\rightarrow$	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.19	1.67	0.17	1.16	1.63	1.45	1.51	1.61	1.96	1.93	1.99	2.2
	Rand. CV	0.12	2.99	0.13	1.94	0.19	2.83	0.23	1.11	0.22	2.3	0.32	1.14
chiBase $= 5$	Const. CV	0.19	1.54	0.17	1.18	0.41	1.82	0.37	1.07	1.11	1.71	0.78	1.96
	Rand. CV	0.12	2.48	0.13	2.22	0.19	3.06	0.23	1.26	0.22	1.47	0.29	0.94
chiBase $= 1000$	Const. CV	0.19	1.36	0.17	1.31	0.41	1.64	0.38	1.68	0.85	2.77	0.81	2.64
	Rand. CV	0.12	0.74	0.13	2.21	0.19	0.86	0.23	1.36	0.22	1.64	0.29	1.07

Table 15: $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ results for $L = 3$, $\ell = 2$, targetEpochs $= 3$, and baseCostRatio $= 10$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase $= -1$	Const. CV	0.81	4.05	0.82	3.8	0.9	2.72	0.9	2.47	0.94	2.37	0.95	2.42
	Rand. CV	0.78	5.25	0.8	4.94	0.78	4.4	0.81	4.03	0.78	4.22	0.81	3.16
chiBase $= 5$	Const. CV	0.81	2.99	0.82	2.86	0.84	1.44	0.82	1.52	0.93	1.31	0.93	1.23
	Rand. CV	0.78	4.61	0.8	4.03	0.78	3.49	0.81	2.88	0.78	3.17	0.81	2.16
chiBase $= 1000$	Const. CV	0.81	1.23	0.82	1.22	0.84	1.17	0.82	1.17	0.93	1.51	0.93	1.55
	Rand. CV	0.78	1.49	0.8	1.45	0.78	1.35	0.81	1.25	0.78	1.26	0.81	1.12

Table 16: $100 \cdot (\text{UB} - \text{LB})/\text{LB}$ results for $L = 3$, $\ell = 2$, targetEpochs $= 7$, and baseCostRatio $= 4$.

		CVBase = 0.15				CVBase = 0.3				CVBase = 0.4			
		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs		Prop. Costs		Rand. Costs	
	overPenalty →	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1	0.5	1
chiBase = −1	Const. CV	0.19	6.31	0.19	6.06	0.96	2.77	0.98	2.56	1.28	2.18	1.2	2.17
	Rand. CV	0.07	8.12	0.09	7.72	0.1	7.35	0.17	6.03	0.13	6.63	0.26	4.24
chiBase = 5	Const. CV	0.19	4.67	0.19	4.45	0.54	2.22	0.5	2.17	0.86	1.76	0.81	1.83
	Rand. CV	0.07	6.78	0.09	6.09	0.1	5.45	0.17	4.37	0.13	4.74	0.24	3.51
chiBase = 1000	Const. CV	0.19	0.82	0.19	0.85	0.54	1.2	0.5	1.23	0.82	1.77	0.79	1.73
	Rand. CV	0.07	1.43	0.09	1.3	0.1	1.09	0.17	0.84	0.13	0.88	0.24	0.81

Table 17: $100 \cdot (\mathrm{UB} - \mathrm{LB})/\mathrm{LB}$ results for $L = 3$, $\ell = 2$, targetEpochs $= 7$, and baseCostRatio $= 10$.