

Linear Programming Based Decomposition Methods for Inventory Distribution Systems

Sumit Kunnumkal

Indian School of Business, Gachibowli, Hyderabad, 500032, India
sumit_kunnumkal@isb.edu

Huseyin Topaloglu

School of Operations Research and Information Engineering,
Cornell University, Ithaca, New York 14853, USA
topaloglu@orie.cornell.edu
Tel: 1-607-255-0698 Fax: 1-607-255-9129

October 29, 2010

Abstract

We consider an inventory distribution system consisting of one warehouse and multiple retailers. The retailers face random demand and are supplied by the warehouse. The warehouse replenishes its stock from an external supplier. The objective is to minimize the total expected replenishment, holding and backlogging cost over a finite planning horizon. The problem can be formulated as a dynamic program, but this dynamic program is difficult to solve due to its high dimensional state variable. It has been observed in the earlier literature that if the warehouse is allowed to ship negative quantities to the retailers, then the problem decomposes by the locations. One way to exploit this observation is to relax the constraints that ensure the nonnegativity of the shipments to the retailers by associating Lagrange multipliers with them, which naturally raises the question of how to choose a good set of Lagrange multipliers. In this paper, we propose efficient methods that choose a good set of Lagrange multipliers by solving linear programming approximations to the inventory distribution problem. Computational experiments indicate that the inventory replenishment policies obtained by our approach can outperform several standard benchmarks by significant margins.

Keywords: Inventory distribution, approximate dynamic programming, inventory control.

1 INTRODUCTION

Many systems that involve storage and distribution of inventory operate in stages. A central warehouse distributing the inventory for the whole system may form the highest stage, several regional distribution centers supplied by the central warehouse may lie in the middle and the lowest stage of the system may be composed of a number of retailers that serve the end customer demand. Clearly, there are tradeoffs when deciding how much inventory should be stored at different stages. On one hand, there are lead times associated with shipping inventory from one stage to another and storing the inventory closer to the lower stages allows responding to the retailers more quickly. On the other hand, storing the inventory closer to the higher stages provides more flexibility. It may be difficult to ship inventory from one retailer to another, but the inventory stored at a regional distribution center may be used to serve the demand at a variety of retailers. Due to these tradeoffs, operating an inventory distribution system requires careful planning with a view of the overall system.

In this paper, we consider an inventory distribution system consisting of a single warehouse and multiple retailers. The retailers face random demand from the end customers and the unsatisfied demand at the retailers is backlogged. The warehouse supplies the retailers and replenishes its inventory from an external supplier. Inventory can be stocked both at the warehouse and at any of the retailers. We want to find a policy for the warehouse to supply the retailers and replenish itself so as to minimize the total expected cost over a finite planning horizon. We begin by formulating the problem as a dynamic program, which is computationally difficult to solve due to the high dimensionality of the state variable. However, this formulation allows us to observe that if the warehouse is allowed to ship negative quantities to the retailers, then the problem decomposes by the locations and can be solved by focusing on one location at a time. To exploit this observation, we relax the constraints that ensure the nonnegativity of the shipments to the retailers by associating Lagrange multipliers with them. In this case, we obtain an approximate solution to the dynamic programming formulation of the problem by solving a sequence of smaller dynamic programs, each involving a scalar state variable keeping track of the inventory position at one location.

The approach outlined above naturally brings up the question of what values to use for the Lagrange multipliers. Indeed, one of our main goals in this paper is to develop methods for choosing a good set of Lagrange multipliers in an efficient manner. The methods that we develop are based on linear programming approximations to the inventory distribution problem. In these linear programs, there are constraints that ensure the nonnegativity of the shipments to the retailers and we use the optimal values of the dual variables associated with these constraints as our Lagrange multipliers. We use two flavors of linear programming approximations. In the first flavor, we formulate a linear program under the assumption that the demand random variables take on their expected values. This way of choosing the Lagrange multipliers is quick to implement and solve, but it clearly ignores the inherent uncertainty in the demand random variables. In our second flavor, we formulate a linear program under the assumption that the realizations of the demand random variables are known a priori. The difficulty in this case is that the optimal values of the dual variables depend on what realizations we use for the demand random variables. We address this difficulty by solving the linear programming

approximation for many demand realizations. Each realization provides a set of optimal values for the dual variables and we use the average of the dual variables as our Lagrange multipliers. Since the set of dual variables that we obtain depend on the distributions of the demand random variables, the second method effectually uses the distribution information.

1.1 POSITIONING AND LITERATURE REVIEW

Our work here draws on two lines of research. The first line of research is on inventory distribution problems. There is a long line of work that is built on the observation that the dynamic programming formulations of inventory distribution problems decompose when we relax the constraints that ensure the nonnegativity of the shipments to the retailers. Clark and Scarf (1960) appear to be the first to make this observation and Eppen and Schrage (1981) and Federgruen and Zipkin (1984*b*) revive this observation two decades later. Gallego, Ozer and Zipkin (2007) give a recent and rigorous overview of the literature revolving around this observation. One common point made in these papers is that we obtain lower bounds on the value functions when we relax the nonnegativity requirement of the shipments to the retailers and we can obtain high quality inventory replenishment policies by using these lower bounds as approximations to the value functions. However, these papers either do not use Lagrange multipliers of any kind to penalize the violations of the relaxed constraints or do not give general purpose methods to compute Lagrange multipliers for any demand distribution.

Kunnumkal and Topaloglu (2008) build on the work described in the paragraph above by relaxing the constraints that ensure the nonnegativity of the shipments to the retailers, but they use Lagrange multipliers to explicitly penalize the violations of the relaxed constraints. Their contribution is to give a general purpose method that chooses a good set of Lagrange multipliers by solving a convex optimization problem. The computational experiments in Kunnumkal and Topaloglu (2008) indicate that the quality of the lower bounds and the performance of the inventory replenishment policies can be improved by associating Lagrange multipliers with the relaxed constraints. However, the unfortunate aspect of the work in Kunnumkal and Topaloglu (2008) is that solving the aforementioned convex optimization problem to find a good set of Lagrange multipliers is a computationally intensive and complex process. Solving their convex optimization problem through subgradient optimization requires calibrating the step size parameter and devising a stopping criterion, for which there are no hard and fast rules. Furthermore, computing a subgradient of their objective function requires a large amount of matrix algebra. The computational overhead associated with the large amount of matrix algebra, combined with a large number of iterations required for subgradient optimization, results in extensive computation times. Our interactions with supply chain practitioners indicate that the approach proposed by Kunnumkal and Topaloglu (2008) is theoretically interesting, but may not be practically attractive. In particular, extensive computation times may not only prevent dealing with industrial size systems, but may also prevent doing quick what if analyses.

One of our goals is to address the shortcomings of Kunnumkal and Topaloglu (2008). We propose efficient methods to obtain a good set of Lagrange multipliers. Our methods require solving linear

programs and can easily be implemented by using commercial linear programming software. It turns out that the methods that we propose can provide computation time improvements by about a factor of ten when compared with the work of Kunnumkal and Topaloglu (2008). However, our work is not limited to addressing the computational difficulties of Kunnumkal and Topaloglu (2008). We show that our linear programming approximations provide lower bounds on the optimal total expected cost. Furthermore, our computational experiments indicate the total expected costs incurred by our approach can be better than those incurred by Kunnumkal and Topaloglu (2008) by up to 7.62%. Therefore, our approach can provide improvements over Kunnumkal and Topaloglu (2008) in terms of both shorter computation times and better total expected costs.

The second line of research that we draw on is related to network revenue management. In the network revenue management literature, it is customary to formulate deterministic approximations by assuming that the demand random variables take on their expected values. This approach results in a linear program that is easy to solve and one can use the dual variables associated with the seat availability constraints to estimate the opportunity cost of a seat on a flight. Talluri and van Ryzin (2005) give a comprehensive overview of the literature revolving around this idea. Our linear programming approximations that work with expected values of the demand random variables are motivated by this literature. On the other hand, observing that a deterministic linear program does not capture the stochastic aspects of network revenue management problems, Talluri and van Ryzin (1999) use a randomized linear program that uses realizations of the demand random variables instead of the expected values. The authors solve the randomized linear program for many demand realizations and estimate the opportunity cost of a seat on a flight by averaging the optimal values of the dual variables obtained from different demand realizations. Our linear programming approximations that work with multiple samples of the demand random variables are motivated by the randomized linear programming approach of Talluri and van Ryzin (1999). To our knowledge, ours is the first application of the randomized linear programming idea outside network revenue management.

There is substantial literature on inventory distribution systems. In their seminal paper, Clark and Scarf (1960) initiate the study of inventory distribution problems. They point out that the optimal policy can be quite complex and proceed to developing an approximation method. Their approximation method hinges on the balance assumption, which is regarded as one of the pillars of the inventory control literature. In particular, the balance assumption amounts to assuming that it would never be desirable to redistribute the total amount of inventory available at the retailers even if it were allowed to do so. Clark and Scarf (1960) show that the optimal inventory replenishment policy under the balance assumption can be obtained in a tractable manner by focusing on one installation at a time, rather than the whole system. Federgruen and Zipkin (1984*b*) propose relaxing the constraints that ensure the nonnegativity of the shipments to the retailers. Under this relaxation, it is also possible to show that the optimal inventory replenishment policy can be obtained by focusing on one installation at a time. Furthermore, both the balance assumption of Clark and Scarf (1960) and the relaxation strategy of Federgruen and Zipkin (1984*b*) provide lower bounds on the value functions and these lower bounds can be useful to assess the optimality gap of any suboptimal policy. The papers by Federgruen and Zipkin (1984*a*) and Federgruen and Zipkin (1984*c*) follow up on the idea of relaxing the constraints

that ensure the nonnegativity of the shipments to the retailers and consider a variety of assumptions, including the finiteness of the planning horizon and the ability to hold stock at the warehouse.

The balance assumption has been employed and tested by many researchers. Eppen and Schrage (1981) use the balance assumption for a distribution system with a stockless warehouse and normally distributed end customer demands. Jackson (1988) extends this work to allow storing inventory at the warehouse. Axsater (2003) reviews the literature related to the balance assumption. Aviv and Federgruen (2001) consider a stockless warehouses with unknown demand distributions that need to be learned through observations. Gallego and Ozer (2003) and Ozer (2003) incorporate advance demand information into inventory systems with multiple echelons. Dogru (2005) and Dogru, de Kok and van Houtum (2005) provide extensive computational experiments that test the inventory replenishment policies obtained under the balance assumption. Federgruen (1993) gives a unified overview of the settings in which the relaxation strategy can be employed. Gallego et al. (2007) test the performances of different approximation strategies for inventory distribution systems, including the relaxation strategy of Federgruen and Zipkin (1984*b*). That paper also emphasizes the importance of computational tractability and performance speedup from the perspective of practical implementation.

1.2 SUMMARY OF CONTRIBUTIONS AND ORGANIZATION OF THE PAPER

In this paper, we make the following research contributions. 1) We develop tractable methods to make the replenishment decisions in inventory distribution systems. We exploit the dynamic programming formulation of the problem and use Lagrange multipliers to relax the constraints that ensure the nonnegativity of the shipments to the retailers. We develop new methods that choose a good set of Lagrange multipliers by solving linear programming approximations. Our methods are more tractable than the subgradient optimization idea used by Kunnumkal and Topaloglu (2008). Tractability is particularly crucial considering that the computational experiments in Kunnumkal and Topaloglu (2008) demonstrate that their relaxation strategy can provide significant improvements over several standard benchmarks, but the computation times are not short enough to allow what if analyses. 2) We prove that the methods that we develop provide lower bounds on the value functions. 3) Computational experiments on 483 test problems indicate that the inventory replenishment policies obtained by our methods can be significantly better than those obtained by standard benchmarks. We use the balance assumption of Clark and Scarf (1960) and the relaxation strategies of Federgruen and Zipkin (1984*b*) and Kunnumkal and Topaloglu (2008) as benchmarks and report significant improvements.

The rest of the paper is organized as follows. In Section 2, we formulate the inventory distribution problem as a dynamic program. In Section 3, we briefly review the Lagrangian relaxation strategy of Kunnumkal and Topaloglu (2008) and explain the portions of their results that are crucial to our work. In Section 4, we develop new methods that can be used to obtain a good set of Lagrange multipliers and show that these methods provide lower bounds on the value functions. In Section 5, we describe how we can apply the inventory replenishment policies obtained by our approach. In Section 6, we present our computational experiments. In Section 7, we conclude.

2 PROBLEM FORMULATION

Our problem formulation is standard and it is similar to the ones in Federgruen (1993) and Kunnumkal and Topaloglu (2008). We consider an inventory distribution system with a single warehouse supplying multiple retailers. We use ϕ to denote the warehouse and \mathcal{I} to denote the set of retailers. Throughout the paper, we use the term installation whenever we need to refer to the warehouse or a retailer without making a distinction. The problem takes place over the finite planning horizon $\mathcal{T} = \{1, \dots, \tau\}$. The demand at retailer i at time period t is given by the random variable d_{it} . We let $d_{\phi t} = \sum_{i \in \mathcal{I}} d_{it}$ so that we can also speak of the demand at the warehouse at time period t .

For notational clarity, we assume that the lead times for all replenishments are zero. In particular, the replenishment order shipped to a certain installation at a certain time period reaches its destination at the same time period. This assumption is only for notational brevity and all of our development in the paper extends in a straightforward and tractable manner to cover the case where the lead times are nonzero. The computational experiments that we present in Section 6 are indeed carried out under the assumption that the lead times are nonzero. Assuming that the lead times for all replenishments are zero, the following sequence of events take place at a particular time period. 1) The warehouse places its replenishment order to the external supplier. 2) Considering the inventory positions at the warehouse and at the retailers, the warehouse ships the replenishment orders to the retailers. 3) The warehouse receives its replenishment order from the external supplier and the retailers receive their replenishment orders from the warehouse. 4) The demand at the retailers is observed. The excess demand is backlogged by incurring a backlogging cost. The warehouse and the retailers incur holding costs for the inventory that they carry to the next time period.

We use x_{it} to denote the echelon inventory position at installation i at the beginning of time period t . For retailer i , the echelon inventory position is the difference between the inventory on hand and backlogs at this retailer. For the warehouse, the echelon inventory position is the inventory on hand at the warehouse, plus the inventory on hand at all of the retailers, minus the backlogs at all of the retailers. Noting that the echelon inventory position at the warehouse contains the inventory on hand and backlogs at all of the retailers, the warehouse has $x_{\phi t} - \sum_{i \in \mathcal{I}} x_{it}$ units of inventory on hand at the beginning of time period t . On the other hand, we use y_{it} to denote the echelon inventory position at installation i after receiving the replenishment quantity at time period t . Therefore, the difference $y_{it} - x_{it}$ gives the replenishment quantity shipped to installation i at time period t . Since the retailers receive their replenishment orders before the demand at the retailers is observed, retailer i has y_{it} units of inventory to serve the demand at time period t . We refer the reader to Clark and Scarf (1960) and Federgruen (1993) for further details of the echelon inventory concept.

We let h_{it} be the per unit holding cost at installation i at time period t and b_{it} be the per unit backlogging cost at retailer i at time period t . The expected holding and backlogging cost incurred by the whole inventory distribution system at time period t can be written as

$$h_{\phi t} \left[x_{\phi t} - \sum_{i \in \mathcal{I}} x_{it} + (y_{\phi t} - x_{\phi t}) - \sum_{i \in \mathcal{I}} (y_{it} - x_{it}) \right] + \sum_{i \in \mathcal{I}} h_{it} \mathbb{E} \{ [y_{it} - d_{it}]^+ \} + \sum_{i \in \mathcal{I}} b_{it} \mathbb{E} \{ [d_{it} - y_{it}]^+ \},$$

where we use $[\cdot]^+ = \max\{\cdot, 0\}$. We compute the expected holding cost at the warehouse by noting that $x_{\phi t} - \sum_{i \in \mathcal{I}} x_{it}$ gives the inventory on hand at the warehouse at the beginning of time period t , $y_{\phi t} - x_{\phi t}$ gives the replenishment quantity shipped to the warehouse at time period t and $\sum_{i \in \mathcal{I}} (y_{it} - x_{it})$ gives the total replenishment quantity shipped out of the warehouse at time period t . Since the warehouse does not face the customer demand, there is no backlogging cost associated with the warehouse. If we let $L_{\phi t}(y_{\phi t}) = h_{\phi t} y_{\phi t}$ and

$$L_{it}(y_{it}) = -h_{\phi t} y_{it} + h_{it} \mathbb{E}\{[y_{it} - d_{it}]^+\} + b_{it} \mathbb{E}\{[d_{it} - y_{it}]^+\},$$

then we can succinctly write the expected holding and backlogging cost incurred by the whole inventory distribution system at time period t as $\sum_{i \in \mathcal{I} \cup \{\phi\}} L_{it}(y_{it})$. This is a standard way of accounting for the expected holding and backlogging cost in inventory distribution systems.

Since the warehouse has $x_{\phi t} - \sum_{i \in \mathcal{I}} x_{it}$ units of inventory on hand at the beginning of time period t and we ship a total of $\sum_{i \in \mathcal{I}} (y_{it} - x_{it})$ units of inventory out of the warehouse at this time period, the replenishment quantities satisfy the constraint $\sum_{i \in \mathcal{I}} (y_{it} - x_{it}) \leq x_{\phi t} - \sum_{i \in \mathcal{I}} x_{it}$, which can be written as $\sum_{i \in \mathcal{I}} y_{it} \leq x_{\phi t}$. On the other hand, since $y_{it} - x_{it}$ is the replenishment quantity shipped to installation i at time period t , the nonnegativity of the replenishment quantities impose the constraint $y_{it} - x_{it} \geq 0$ for all $i \in \mathcal{I} \cup \{\phi\}$. Using $x_t = \{x_{it} : i \in \mathcal{I} \cup \{\phi\}\}$ as the state variable at time period t , $y_t = \{y_{it} : i \in \mathcal{I} \cup \{\phi\}\}$ as the decision variable at time period t and $d_t = \{d_{it} : i \in \mathcal{I} \cup \{\phi\}\}$ to denote the vector of demand random variables at time period t , the optimal inventory replenishment policy that minimizes the total expected cost over the planning horizon can be found by computing the value functions $\{V_t(\cdot) : t \in \mathcal{T}\}$ through the optimality equation

$$V_t(x_t) = \min \sum_{i \in \mathcal{I} \cup \{\phi\}} c_{it} [y_{it} - x_{it}] + \sum_{i \in \mathcal{I} \cup \{\phi\}} L_{it}(y_{it}) + \mathbb{E}\{V_{t+1}(y_t - d_t)\} \quad (1)$$

$$\text{subject to } \sum_{i \in \mathcal{I}} y_{it} \leq x_{\phi t} \quad (2)$$

$$y_{it} \geq x_{it} \quad i \in \mathcal{I} \cup \{\phi\}. \quad (3)$$

Noting that the echelon inventory position at the warehouse contains the inventory on hand and backlogs at all of the retailers, the echelon inventory position at the warehouse at the beginning of time period $t + 1$ is computed as $y_{\phi t} - d_{\phi t}$. Therefore, the term $y_t - d_t$ in the argument of the value function on the right side of the optimality equation above corresponds to the echelon inventory positions at the different installations. Constraints (2) impose the inventory availability at the warehouse, whereas constraints (3) ensure the nonnegativity of the replenishment quantities.

Due to the large number of dimensions of the state variable, it is computationally difficult to solve the optimality equation in (1)-(3) through standard dynamic programming methods. Kunnumkal and Topaloglu (2008) relax the constraints $y_{it} \geq x_{it}$ for all $i \in \mathcal{I}$ in problem (1)-(3) by associating Lagrange multipliers with them. In this case, the optimality equation in (1)-(3) decomposes by the installations and it can be solved by focusing on one installation at a time. In the next section, we briefly review this idea. After this, we develop new methods to choose a good set of Lagrange multipliers.

3 LAGRANGIAN RELAXATION STRATEGY

Associating the nonnegative Lagrange multipliers $\lambda = \{\lambda_{it} : i \in \mathcal{I}, t \in \mathcal{T}\}$ with the constraints $y_{it} \geq x_{it}$ for all $i \in \mathcal{I}$ in problem (1)-(3), we propose solving the optimality equation

$$V_t^\lambda(x_t) = \min_{y_{\phi t}} \quad c_{\phi t}[y_{\phi t} - x_{\phi t}] + L_{\phi t}(y_{\phi t}) + \sum_{i \in \mathcal{I}} [c_{it} - \lambda_{it}][y_{it} - x_{it}] + \sum_{i \in \mathcal{I}} L_{it}(y_{it}) + \mathbb{E}\{V_{t+1}^\lambda(y_t - d_t)\} \quad (4)$$

$$\text{subject to} \quad \sum_{i \in \mathcal{I}} y_{it} \leq x_{\phi t} \quad (5)$$

$$y_{\phi t} \geq x_{\phi t}, \quad (6)$$

where the superscript λ in the value functions emphasizes that the solution to the optimality equation above depends on the Lagrange multipliers. Since $y_{it} - x_{it}$ is the replenishment quantity of installation i at time period t , relaxing the constraints $y_{it} \geq x_{it}$ for all $i \in \mathcal{I}$ is equivalent to relaxing the constraints that ensure the nonnegativity of the shipments to the retailers.

Earlier work by Kunnumkal and Topaloglu (2008) gives two results that are useful for our work in this paper. First, the optimality equation in (4)-(6) decomposes by the installations and it can be solved by focusing on one installation at a time. Second, the value functions $\{V_t^\lambda(\cdot) : t \in \mathcal{T}\}$ computed through the optimality equation in (4)-(6) provide lower bounds on the value functions $\{V_t(\cdot) : t \in \mathcal{T}\}$ computed through the optimality equation in (1)-(3). In this case, we can choose the Lagrange multipliers to make the lower bounds provided by $\{V_t^\lambda(\cdot) : t \in \mathcal{T}\}$ as tight as possible. In the remainder of this section, we give short descriptions of these two results.

As Proposition 1 in Kunnumkal and Topaloglu (2008) indicates, it is possible to use an induction argument over the time periods to show that the optimality equation in (4)-(6) decomposes by the installations. In particular, the value functions computed under the Lagrangian relaxation strategy are separable functions of the form

$$V_t^\lambda(x_t) = \sum_{i \in \mathcal{I} \cup \{\phi\}} v_{it}^\lambda(x_{it}), \quad (7)$$

where the value functions $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ can be computed by solving dynamic programs with scalar state variables. In the expression above, the value functions $\{v_{it}^\lambda(\cdot) : t \in \mathcal{T}\}$ for retailer i are computed by solving the optimality equation

$$v_{it}^\lambda(x_{it}) = \min_{y_{it}} \left\{ [c_{it} - \lambda_{it}][y_{it} - x_{it}] + L_{it}(y_{it}) + \mathbb{E}\{v_{i,t+1}^\lambda(y_{it} - d_{it})\} \right\}. \quad (8)$$

This optimality equation computes the optimal inventory replenishment policy for retailer i under the assumption that the warehouse has infinite supply, the replenishment quantities of retailer i can be negative and the per unit replenishment cost at retailer i at time period t is deflated by λ_{it} . On the other hand, noting that the optimal solution to problem (8) does not depend on x_{it} and letting \hat{r}_{it}^λ be the optimal solution to this problem, the value functions $\{v_{\phi t}^\lambda(\cdot) : t \in \mathcal{T}\}$ for the warehouse are computed

by solving the optimality equation

$$v_{\phi t}^\lambda(x_{\phi t}) = \min_{y_{\phi t} \geq x_{\phi t}} \left\{ c_{\phi t} [y_{\phi t} - x_{\phi t}] + L_{\phi t}(y_{\phi t}) + \mathbb{E}\{v_{\phi, t+1}^\lambda(y_{\phi t} - d_{\phi t})\} \right\} + \Delta_t^\lambda(x_{\phi t}), \quad (9)$$

where the functions $\{\Delta_t^\lambda(\cdot) : t \in \mathcal{T}\}$ are given by

$$\Delta_t^\lambda(x_{\phi t}) = \min \sum_{i \in \mathcal{I}} [c_{it} - \lambda_{it}] [y_{it} - \hat{r}_{it}^\lambda] + \sum_{i \in \mathcal{I}} L_{it}(y_{it}) + \sum_{i \in \mathcal{I}} \mathbb{E}\{v_{i, t+1}^\lambda(y_{it} - d_{it})\} \\ - \sum_{i \in \mathcal{I}} L_{it}(\hat{r}_{it}^\lambda) - \sum_{i \in \mathcal{I}} \mathbb{E}\{v_{i, t+1}^\lambda(\hat{r}_{it}^\lambda - d_{it})\} \quad (10)$$

$$\text{subject to } \sum_{i \in \mathcal{I}} y_{it} \leq x_{\phi t}. \quad (11)$$

The optimality equation in (9) computes the optimal inventory replenishment policy for the warehouse under the assumption that the expected holding and backlogging cost at the warehouse at time period t is inflated by the function $\Delta_t^\lambda(\cdot)$. Without the function $\Delta_t^\lambda(\cdot)$, the optimality equation in (9) would compute the optimal policy for the warehouse with no regard to the retailers. It is possible to show that $\Delta_t^\lambda(x_{\phi t})$ is a convex function of $x_{\phi t}$ and gets larger as $x_{\phi t}$ gets smaller. Therefore, the role of the function $\Delta_t^\lambda(\cdot)$ is to penalize small echelon inventory levels at the warehouse. Similar functions appear in papers such as Clark and Scarf (1960), Chen and Song (2001) and Gallego and Ozer (2003) that analyze inventory problems with multiple echelons. Gallego and Ozer (2003) use the term implicit penalty function to refer to functions that play the role of $\Delta_t^\lambda(\cdot)$. Similar to $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I}, t \in \mathcal{T}\}$, the value functions that Gallego et al. (2007) compute for the retailers are based on an infinite supply assumption at the warehouse.

The significance of the optimality equations in (8) and (9) is that we can compute the value functions $\{V_t^\lambda(\cdot) : t \in \mathcal{T}\}$ by first computing the value functions $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I}, t \in \mathcal{T}\}$ through the optimality equation in (8). This allows computing $\{\Delta_t^\lambda(\cdot) : t \in \mathcal{T}\}$ by solving problem (10)-(11). As a last step, we can compute the value functions $\{v_{\phi t}^\lambda(\cdot) : t \in \mathcal{T}\}$ through the optimality equation in (9). In this case, we can obtain $\{V_t^\lambda(\cdot) : t \in \mathcal{T}\}$ by noting (7). The fact that the optimality equation in (4)-(6) decomposes by the installations has close ties with the earlier literature. Federgruen and Zipkin (1984b) and Federgruen and Zipkin (1984c) show a similar decomposition result when one does not use Lagrange multipliers to penalize the violations of relaxed constraints. Ozer (2003) uses a Lagrangian relaxation strategy to compute optimal base stock levels after restricting attention to the class of policies that order up to a fixed level. Aviv and Federgruen (2001) propose a decomposition strategy that is applicable when no inventory is held at the warehouse, but they do not use Lagrange multipliers.

As Proposition 3 in Kunnumkal and Topaloglu (2008) indicates, it is possible to show that if the Lagrange multipliers are nonnegative, then the value functions computed through the optimality equation in (4)-(6) are lower bounds on the value functions computed through the optimality equation in (1)-(3). In other words, we have $V_t^\lambda(x_t) \leq V_t(x_t)$ for all $t \in \mathcal{T}$ whenever $\lambda \geq 0$. As a function of the initial echelon inventory positions x_1 , the optimal total expected cost over the planning horizon is given by $V_1(x_1)$. Since $V_1(x_1)$ is bounded from below by $V_1^\lambda(x_1)$, we can solve the problem

$$\max_{\lambda \geq 0} \left\{ V_1^\lambda(x_1) \right\} \quad (12)$$

to obtain the tightest possible lower bound on $V_1(x_1)$. Letting λ^* be the optimal solution to problem (12), Kunnumkal and Topaloglu (2008) choose their Lagrange multipliers as λ^* . In practice, however, solving problem (12) requires using subgradient optimization, which is computationally intensive and dependent on calibrating a step size parameter. To remedy these shortcomings, in the next section, we propose new methods that use linear programs to choose the Lagrange multipliers.

4 CHOOSING THE LAGRANGE MULTIPLIERS

In this section, we describe new methods that can be used to choose a good set of Lagrange multipliers in the Lagrangian relaxation strategy. These methods are based on solving linear programs and one should view them as computationally appealing alternatives to solving problem (12). If we use $d = \{d_{it} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ to denote the demand random variables over the whole planning horizon, then we can approximate the total cost incurred by the optimal policy under the demand realization d by using the optimal objective value of the problem

$$\min \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I} \cup \{\phi\}} c_{it} [y_{it} - x_{it}] + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I} \cup \{\phi\}} L_{it}(y_{it}) \quad (13)$$

$$\text{subject to } \sum_{i \in \mathcal{I}} y_{it} - x_{\phi t} \leq 0 \quad t \in \mathcal{T} \quad (14)$$

$$y_{it} - x_{it} \geq 0 \quad i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T} \quad (15)$$

$$y_{it} - x_{i,t+1} = d_{it} \quad i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T} \setminus \{\tau\} \quad (16)$$

with the understanding that the values of the decision variables $\{x_{i1} : i \in \mathcal{I} \cup \{\phi\}\}$ are fixed by the initial echelon inventory positions. In the problem above, the decision variable x_{it} is the echelon inventory position at installation i at the beginning of time period t , whereas the decision variable y_{it} is the echelon inventory position at installation i after receiving the replenishment order at time period t . Since d is a random variable, the optimal objective value of problem (13)-(16) is a random variable as well. Constraints (14) are similar to constraints (2) and they ensure that the total replenishment order of the retailers does not exceed the inventory on hand at the warehouse. Constraints (15) are similar to constraints (3) and they ensure the nonnegativity of the shipments to the installations. Constraints (16) compute the echelon inventory positions at the next time period. We note that the computation of $\{L_{it}(\cdot) : i \in \mathcal{I}, t \in \mathcal{T}\}$ in the objective function of problem (13)-(16) uses the distributions of the demand random variables. It is possible to replace $L_{it}(y_{it})$ in problem (13)-(16) with $-h_{\phi t} [y_{it}] + h_{it} [y_{it} - d_{it}]^+ + b_{it} [d_{it} - y_{it}]^+$, which only uses the demand sample d without taking expectations, but this does not make problem (13)-(16) much easier to solve. Finally, if the demand random variables take on discrete values, then $\{L_{it}(\cdot) : i \in \mathcal{I}, t \in \mathcal{T}\}$ are piecewise linear convex functions and problem (13)-(16) can be solved as a linear program.

One important feature of problem (13)-(16) is that it can be used to obtain lower bounds on the optimal total expected cost. In particular, we let $z_{LP}^*(d)$ be the optimal objective value of problem (13)-(16) as a function of the demand realizations d on the right side of constraints (16). In this case, the next proposition shows that both $\mathbb{E}\{z_{LP}^*(d)\}$ and $z_{LP}^*(\mathbb{E}\{d\})$ provide lower bounds on $V_1(x_1)$ and the lower

bound provided by $\mathbb{E}\{z_{LP}^*(d)\}$ is tighter than the one provided by $z_{LP}^*(\mathbb{E}\{d\})$. The lower bound provided by $z_{LP}^*(\mathbb{E}\{d\})$ is relatively simple to compute since it requires replacing the right side of constraints (16) with the expected values of the demand random variables and solving problem (13)-(16) once. In contrast, the lower bound provided by $\mathbb{E}\{z_{LP}^*(d)\}$ requires solving problem (13)-(16) for every possible realization of the demand random variables d , storing $z_{LP}^*(d)$ and computing the expectation $\mathbb{E}\{z_{LP}^*(d)\}$. Naturally, solving problem (13)-(16) for every possible realization of the demand random variables is not tractable in general, but we can solve problem (13)-(16) for N sampled demand realizations d^1, d^2, \dots, d^N and use the sample average $\frac{1}{N} \sum_{n=1}^N z_{LP}^*(d^n)$ to estimate the expectation $\mathbb{E}\{z_{LP}^*(d)\}$.

Proposition 1 *We have that $z_{LP}^*(\mathbb{E}\{d\}) \leq \mathbb{E}\{z_{LP}^*(d)\} \leq V_1(x_1)$.*

Proof By duality theory, $z_{LP}^*(d)$ is a convex function of d so that we obtain $z_{LP}^*(\mathbb{E}\{d\}) \leq \mathbb{E}\{z_{LP}^*(d)\}$ by Jensen's inequality. To show the second inequality, we let π^* be the optimal policy characterized by the solution to the optimality equation in (1)-(3). Given that the realizations of the demand random variables are d , we let $x^*(d) = \{x_{it}^*(d) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ be the sequence of states visited by policy π^* and $y^*(d) = \{y_{it}^*(d) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ be the sequence of actions taken by policy π^* . Naturally, both $x^*(d)$ and $y^*(d)$ are random processes. In this case, noting the cost components in the optimality equation in (1)-(3), the total expected cost incurred by policy π^* over the planning horizon is

$$\sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I} \cup \{\phi\}} c_{it} \mathbb{E}\{y_{it}^*(d) - x_{it}^*(d)\} + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I} \cup \{\phi\}} \mathbb{E}\{L_{it}(y_{it}^*(d))\}. \quad (17)$$

Since π^* is the optimal policy and the total expected cost incurred by the optimal policy over the planning horizon is given by $V_1(x_1)$, the expression in (17) is equal to $V_1(x_1)$.

Noting that the sequences of states visited and actions taken by policy π^* satisfy the inventory dynamics and constraints (2) and (3), we have $x_{i,t+1}^*(d) = y_{it}^*(d) - d_{it}$ for all $i \in \mathcal{I} \cup \{\phi\}$, $t \in \mathcal{T} \setminus \{\tau\}$, $\sum_{i \in \mathcal{I}} y_{it}^*(d) \leq x_{\phi t}^*(d)$ for all $t \in \mathcal{T}$ and $y_{it}^*(d) \geq x_{it}^*(d)$ for all $i \in \mathcal{I} \cup \{\phi\}$, $t \in \mathcal{T}$. Therefore, $x^*(d) = \{x_{it}^*(d) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ and $y^*(d) = \{y_{it}^*(d) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ form a feasible but not necessarily an optimal solution to problem (13)-(16), which implies that

$$z_{LP}^*(d) \leq \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I} \cup \{\phi\}} c_{it} [y_{it}^*(d) - x_{it}^*(d)] + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I} \cup \{\phi\}} L_{it}(y_{it}^*(d)).$$

In this case, the second inequality in the statement of the proposition follows by taking expectations in the expression above and noting that the expression in (17) is equal to $V_1(x_1)$. \square

Another important feature of problem (13)-(16) is that it can be used to choose the Lagrange multipliers in our Lagrangian relaxation strategy. As a function of the demand realizations d , we let $\lambda^*(d) = \{\lambda_{it}^*(d) : i \in \mathcal{I}, t \in \mathcal{T}\}$ be the optimal values of the dual variables associated with the constraints $y_{it} - x_{it} \geq 0$ for all $i \in \mathcal{I}$, $t \in \mathcal{T}$ in problem (13)-(16). Since these constraints are analogous to the constraints $y_{it} - x_{it} \geq 0$ for all $i \in \mathcal{I}$, $t \in \mathcal{T}$ in the optimality equation in (1)-(3), we can build on the dual variables $\{\lambda_{it}^*(d) : i \in \mathcal{I}, t \in \mathcal{T}\}$ to choose the Lagrange multipliers. One idea is to solve problem

(13)-(16) after replacing the right side of constraints (16) with the expected values of the demand random variables to obtain the Lagrange multipliers $\lambda^*(\mathbb{E}\{d\}) = \{\lambda_{it}^*(\mathbb{E}\{d\}) : i \in \mathcal{I}, t \in \mathcal{T}\}$. This approach requires solving problem (13)-(16) only once. Another idea is to solve problem (13)-(16) for every possible realization of the demand random variables d , store $\lambda^*(d)$ and compute the expectation $\mathbb{E}\{\lambda^*(d)\} = \{\mathbb{E}\{\lambda_{it}^*(d)\} : i \in \mathcal{I}, t \in \mathcal{T}\}$. Similar to $\mathbb{E}\{z_{LP}^*(d)\}$, solving problem (13)-(16) for every possible realization of the demand random variables is not tractable in general, but we can approximate the expectation $\mathbb{E}\{\lambda^*(d)\}$ by the sample average $\frac{1}{N} \sum_{n=1}^N \lambda^*(d^n)$, where d^1, d^2, \dots, d^N are N sampled demand realizations. The next proposition shows that if we choose the Lagrange multipliers as $\lambda^*(\mathbb{E}\{d\})$ or $\mathbb{E}\{\lambda^*(d)\}$ and compute $V_1^{\lambda^*(\mathbb{E}\{d\})}(x_1)$ or $V_1^{\mathbb{E}\{\lambda^*(d)\}}(x_1)$ by solving the optimality equation in (4)-(6), then we obtain lower bounds on the optimal total expected cost. Furthermore, the lower bound provided by $V_1^{\lambda^*(\mathbb{E}\{d\})}(x_1)$ is at least as tight as the one provided by $z_{LP}^*(\mathbb{E}\{d\})$.

Proposition 2 *We have that $V_1^{\lambda^*(\mathbb{E}\{d\})}(x_1) \leq V_1(x_1)$ and $V_1^{\mathbb{E}\{\lambda^*(d)\}}(x_1) \leq V_1(x_1)$. Furthermore, the lower bound provided by $V_1^{\lambda^*(\mathbb{E}\{d\})}(x_1)$ is at least as tight as the one provided by $z_{LP}^*(\mathbb{E}\{d\})$ in the sense that $z_{LP}^*(\mathbb{E}\{d\}) \leq V_1^{\lambda^*(\mathbb{E}\{d\})}(x_1)$.*

Proof By dual feasibility to problem (13)-(16), we have $\lambda^*(d) \geq 0$ for any demand realization d . Thus, we have $\lambda^*(\mathbb{E}\{d\}) \geq 0$ and $\mathbb{E}\{\lambda^*(d)\} \geq 0$. On the other hand, Kunnumkal and Topaloglu (2008) show that $\max_{\lambda \geq 0} V_1^\lambda(x_1) \leq V_1(x_1)$. Therefore, the first two inequalities in the proposition follow from the fact that $V_1^{\lambda^*(\mathbb{E}\{d\})}(x_1) \leq \max_{\lambda \geq 0} V_1^\lambda(x_1) \leq V_1(x_1)$ and $V_1^{\mathbb{E}\{\lambda^*(d)\}}(x_1) \leq \max_{\lambda \geq 0} V_1^\lambda(x_1) \leq V_1(x_1)$.

To show the third inequality in the proposition, we let $\bar{d} = \mathbb{E}\{d\}$ and $\bar{\lambda}^* = \lambda^*(\bar{d})$ for notational brevity. Thus, if we solve problem (13)-(16) after replacing the right side of constraints (16) with $\bar{d} = \{\bar{d}_{it} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$, then the optimal objective value of the problem is $z_{LP}^*(\bar{d})$ and the optimal values of the dual variables associated with constraints $y_{it} - x_{it} \geq 0$ for all $i \in \mathcal{I}, t \in \mathcal{T}$ are $\bar{\lambda}^* = \{\bar{\lambda}_{it}^* : i \in \mathcal{I}, t \in \mathcal{T}\}$. In this case, by duality theory, the optimal objective value of the problem

$$\min \sum_{t \in \mathcal{T}} c_{\phi t} [y_{\phi t} - x_{\phi t}] + \sum_{t \in \mathcal{T}} L_{\phi t}(y_{\phi t}) + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} [c_{it} - \bar{\lambda}_{it}^*] [y_{it} - x_{it}] + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} L_{it}(y_{it}) \quad (18)$$

$$\text{subject to } \sum_{i \in \mathcal{I}} y_{it} - x_{\phi t} \leq 0 \quad t \in \mathcal{T} \quad (19)$$

$$y_{\phi t} - x_{\phi t} \geq 0 \quad t \in \mathcal{T} \quad (20)$$

$$y_{it} - x_{i,t+1} = \bar{d}_{it} \quad i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T} \setminus \{\tau\} \quad (21)$$

is equal to $z_{LP}^*(\bar{d})$. We now proceed to consider the optimality equation in (4)-(6) with the Lagrange multipliers $\bar{\lambda}^*$. We let $\pi^{\bar{\lambda}^*}$ be the policy characterized by the solution to this optimality equation, $x^{\bar{\lambda}^*} = \{x_{it}^{\bar{\lambda}^*} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ be the sequence of states visited by policy $\pi^{\bar{\lambda}^*}$ and $y^{\bar{\lambda}^*} = \{y_{it}^{\bar{\lambda}^*} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ be the sequence of actions taken by policy $\pi^{\bar{\lambda}^*}$. Clearly, both $x^{\bar{\lambda}^*}$ and $y^{\bar{\lambda}^*}$ are random processes. In this case, noting the cost components in the optimality equation in (4)-(6), the total expected cost incurred by policy $\pi^{\bar{\lambda}^*}$ over the planning horizon can be written as

$$\sum_{t \in \mathcal{T}} c_{\phi t} \mathbb{E}\{y_{\phi t}^{\bar{\lambda}^*} - x_{\phi t}^{\bar{\lambda}^*}\} + \sum_{t \in \mathcal{T}} \mathbb{E}\{L_{\phi t}(y_{\phi t}^{\bar{\lambda}^*})\} + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} [c_{it} - \lambda_{it}^*] \mathbb{E}\{y_{it}^{\bar{\lambda}^*} - x_{it}^{\bar{\lambda}^*}\} + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} \mathbb{E}\{L_{it}(y_{it}^{\bar{\lambda}^*})\}. \quad (22)$$

On the other hand, since $V_1^{\bar{\lambda}^*}(x_1)$ corresponds to the optimal total expected cost over the planning horizon in the optimality equation in (4)-(6), the total expected cost incurred by policy $\pi^{\bar{\lambda}^*}$ over the planning horizon is also equal to $V_1^{\bar{\lambda}^*}(x_1)$. Therefore, the expression in (22) is equal to $V_1^{\bar{\lambda}^*}(x_1)$.

Since the sequences of states visited and actions taken by policy $\pi^{\bar{\lambda}^*}$ satisfy the inventory dynamics and constraints (5) and (6), we have $x_{i,t+1}^{\bar{\lambda}^*} = y_{it}^{\bar{\lambda}^*} - d_{it}$ for all $i \in \mathcal{I} \cup \{\phi\}$, $t \in \mathcal{T} \setminus \{\tau\}$, $\sum_{i \in \mathcal{I}} y_{it}^{\bar{\lambda}^*} \leq x_{\phi t}^{\bar{\lambda}^*}$ for all $t \in \mathcal{T}$ and $y_{\phi t}^{\bar{\lambda}^*} \geq x_{\phi t}^{\bar{\lambda}^*}$ for all $t \in \mathcal{T}$ with probability 1. In this case, letting $\bar{x}_{it} = \mathbb{E}\{x_{it}^{\bar{\lambda}^*}\}$ and $\bar{y}_{it} = \mathbb{E}\{y_{it}^{\bar{\lambda}^*}\}$, we observe that $\{\bar{x}_{it} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ and $\{\bar{y}_{it} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ form a feasible solution to problem (18)-(21) and we obtain

$$\begin{aligned} z_{LP}^*(\bar{d}) &\leq \sum_{t \in \mathcal{T}} c_{\phi t} [\bar{y}_{\phi t} - \bar{x}_{\phi t}] + \sum_{t \in \mathcal{T}} L_{\phi t}(\bar{y}_{\phi t}) + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} [c_{it} - \lambda_{it}^*] [\bar{y}_{it} - \bar{x}_{it}] + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} L_{it}(\bar{y}_{it}) \\ &\leq \sum_{t \in \mathcal{T}} c_{\phi t} [\bar{y}_{\phi t} - \bar{x}_{\phi t}] + \sum_{t \in \mathcal{T}} \mathbb{E}\{L_{\phi t}(y_{\phi t}^{\bar{\lambda}^*})\} + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} [c_{it} - \lambda_{it}^*] [\bar{y}_{it} - \bar{x}_{it}] + \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} \mathbb{E}\{L_{it}(y_{it}^{\bar{\lambda}^*})\}. \end{aligned}$$

The first inequality follows by the fact that the optimal objective value of problem (18)-(21) is equal to $z_{LP}^*(\bar{d})$, and $\{\bar{x}_{it} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ and $\{\bar{y}_{it} : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ form a feasible but not necessarily an optimal solution to problem (18)-(21). The second inequality follows by the fact that $L_{it}(\mathbb{E}\{y_{it}^{\bar{\lambda}^*}\}) \leq \mathbb{E}\{L_{it}(y_{it}^{\bar{\lambda}^*})\}$ by Jensen's inequality. The result follows by noting that the last expression in the chain of inequalities above is the same as (22), which is, in turn, equal to $V_1^{\bar{\lambda}^*}(x_1)$. \square

5 APPLYING THE GREEDY POLICIES

In this section, we describe the inventory replenishment policies that we obtain by using the value functions $\{V_t^\lambda(\cdot) : t \in \mathcal{T}\}$ computed through the optimality equation in (4)-(6) as approximations to the value functions $\{V_t(\cdot) : t \in \mathcal{T}\}$ computed through the optimality equation in (1)-(3). We begin by noting that whether we choose the Lagrange multipliers by solving problem (12) as in Kunnumkal and Topaloglu (2008) or by solving problem (13)-(16) as in this paper, the value functions computed through the optimality equation in (4)-(6) are separable functions of the form $V_t^\lambda(x_t) = \sum_{i \in \mathcal{I} \cup \{\phi\}} v_{it}^\lambda(x_{it})$. Furthermore, it can be shown that $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ are convex function.

To make the inventory replenishment decisions, we replace $\{V_t(\cdot) : t \in \mathcal{T}\}$ in the right side of the objective function of problem (1)-(3) with $\{V_t^\lambda(\cdot) : t \in \mathcal{T}\}$ and solve this problem. In this case, since $\{V_t^\lambda(\cdot) : t \in \mathcal{T}\}$ are separable by the installations, problem (1)-(3) decomposes into two subproblems, one for the retailers and one for the warehouse. The subproblem for the retailers has the form

$$\min \sum_{i \in \mathcal{I}} c_{it} [y_{it} - x_{it}] + \sum_{i \in \mathcal{I}} L_{it}(y_{it}) + \sum_{i \in \mathcal{I}} \mathbb{E}\{v_{i,t+1}^\lambda(y_{it} - d_{it})\} \quad (23)$$

$$\text{subject to } \sum_{i \in \mathcal{I}} y_{it} \leq x_{\phi t} \quad (24)$$

$$y_{it} \geq x_{it} \quad i \in \mathcal{I}, \quad (25)$$

whereas the subproblem for the warehouse has the form

$$\min_{y_{\phi t} \geq x_{\phi t}} \left\{ c_{\phi t} [y_{\phi t} - x_{\phi t}] + L_{\phi t}(y_{\phi t}) + \mathbb{E}\{v_{\phi,t+1}^\lambda(y_{\phi t} - d_{\phi t})\} \right\}. \quad (26)$$

Problem (23)-(25) computes the replenishment orders of all of the retailers while ensuring that the total replenishment order of the retailers does not violate the availability at the warehouse and the replenishment order of each retailer is nonnegative. Problem (26) computes the replenishment order of the warehouse while ensuring the nonnegativity of the replenishment order. If the demand random variables take on discrete values, then it is possible to show that both $\{L_{it}(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ and $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ are piecewise linear convex functions, in which case problems (23)-(25) and (26) can be solved as linear programs.

The idea of constructing a policy by replacing the value function in the optimality equation with an approximation is frequently used in the literature. Bertsekas (2001) refers to the resulting policy as the greedy policy with respect to the value function approximation. Federgruen and Zipkin (1984b) and Gallego et al. (2007) use greedy approaches similar to ours in inventory distribution problems, but the computation of their value function approximations does not use Lagrange multipliers.

6 COMPUTATIONAL EXPERIMENTS

In this section, we compare the performances of the inventory replenishment policies obtained by our approach with those obtained by various benchmark methods.

6.1 BENCHMARK METHODS

We use the following five benchmark methods in our computational experiments.

Lagrangian relaxation with deterministic linear program (DL) This is the solution method described in Section 4 and it is based on the idea that solves problem (13)-(16) only once. We solve problem (13)-(16) after replacing the right side of constraints (16) with the expected values of the demand random variables and let $\bar{\lambda}^* = \{\bar{\lambda}_{it}^* : i \in \mathcal{I}, t \in \mathcal{T}\}$ be the optimal values of the dual variables associated with the constraints $y_{it} \geq x_{it}$ for all $i \in \mathcal{I}, t \in \mathcal{T}$. We compute $\{v_{it}^{\bar{\lambda}^*}(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ by solving the optimality equations in (8) and (9) with the Lagrange multipliers $\bar{\lambda}^*$. In this case, DL makes the inventory replenishment decisions by solving problems (23)-(25) and (26) after replacing $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ with $\{v_{it}^{\bar{\lambda}^*}(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$.

Lagrangian relaxation with randomized linear program (RL) This is the solution method described in Section 4 and it is based on the idea that solves problem (13)-(16) with multiple demand realizations. We solve problem (13)-(16) for N demand realizations and let $\lambda^*(d^n) = \{\lambda_{it}^*(d^n) : i \in \mathcal{I}, t \in \mathcal{T}\}$ be the optimal values of the dual variables associated with the constraints $y_{it} \geq x_{it}$ for all $i \in \mathcal{I}, t \in \mathcal{T}$ when we use the demand realization d^n . Letting $\bar{\lambda}_{it} = \frac{1}{N} \sum_{n=1}^N \lambda_{it}^*(d^n)$ and $\bar{\lambda} = \{\bar{\lambda}_{it} : i \in \mathcal{I}, t \in \mathcal{T}\}$, we compute $\{v_{it}^{\bar{\lambda}}(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ by solving the optimality equations in (8) and (9) with the Lagrange multipliers $\bar{\lambda}$. In this case, RL makes the inventory replenishment decisions by solving problems (23)-(25) and (26) after replacing $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ with $\{v_{it}^{\bar{\lambda}}(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$. We use $N = 25$ in all of our computational experiments.

Lagrangian relaxation with subgradient optimization (SG) This is the solution method described in Kunnumkal and Topaloglu (2008). Except that it chooses the Lagrange multipliers by solving problem (12), the implementation of SG is the same as that of DL and RL. When solving problem (12) through subgradient optimization, we use the step size parameter $20/\sqrt{k}$ at iteration k and stop after 1,000 iterations. This is the same step size and stopping rule used in Kunnumkal and Topaloglu (2008).

Balance assumption of Clark and Scarf (BA) Clark and Scarf (1960) establish that if their widely known balance assumption is satisfied, then the value functions $\{V_t(\cdot) : t \in \mathcal{T}\}$ computed through the optimality equation in (1)-(3) are separable functions of the form $V_t(x_t) = \sum_{i \in \mathcal{I} \cup \{\phi\}} v_{it}(x_{it})$. BA makes the inventory replenishment decisions by using the separable functions $\{\sum_{i \in \mathcal{I} \cup \{\phi\}} v_{it}(\cdot) : t \in \mathcal{T}\}$ as value function approximations. In this case, one can see that BA makes the inventory replenishment decisions by solving problems that look very much like those in (23)-(25) and (26). Furthermore, it is possible to show that the value functions computed under the balance assumption provide lower bounds on the optimal total expected cost over the planning horizon. We do not go into the further details of the balance assumption here and refer the reader to Clark and Scarf (1960).

Relaxation strategy of Federgruen and Zipkin (RS) The relaxation strategy of Federgruen and Zipkin (1984b) relaxes the constraints $y_{it} \geq x_{it}$ for all $i \in \mathcal{I}$ in problem (1)-(3), but it does not associate Lagrange multipliers with them. Therefore, the implementation of RS is the same as that of DL, RL and SG, but RS sets all of the Lagrange multipliers to zero. This is to say that we compute $\{v_{it}^0(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ by solving the optimality equations in (8) and (9) with all of the Lagrange multipliers set to zero. In this case, RS makes the inventory replenishment decisions by solving problems (23)-(25) and (26) after replacing $\{v_{it}^\lambda(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$ with $\{v_{it}^0(\cdot) : i \in \mathcal{I} \cup \{\phi\}, t \in \mathcal{T}\}$.

6.2 EXPERIMENTAL SETUP

In our computational experiments, we work with distribution systems with identical as well as nonidentical retailers. For the setting with identical retailers, we start with a base case consisting of 21 test problems. In the base case, we have an inventory distribution system with three retailers and a planning horizon of 50 time periods. The holding and backlogging costs are $h_{\phi t} = 0.6$, $h_{it} = 1$ and $b_{it} = 19$ for all $i \in \mathcal{I}$, $t \in \mathcal{T}$. The replenishment order shipped to a certain installation at a certain time period reaches the installation at the next time period. This is in contrast with our earlier assumption that the lead times for all replenishments are zero. We assume that the demand at retailer i at time period t has a Poisson distribution with mean α_{it} and we generate three different demand profiles by generating $\{\alpha_{it} : i \in \mathcal{I}, t \in \mathcal{T}\}$ in three different ways. The demand profiles primarily differ in the degree of nonstationarity of the demand. For demand profile \mathcal{D}_0 , we simply generate α_{it} from the uniform distribution over $[5, 15]$. This demand profile represents a situation where the demand is relatively stationary and the goal of the other two demand profiles is to introduce additional nonstationarity. For demand profile \mathcal{D}_1 , we set $\alpha_{it} = 10$ with probability $1/2$ and have $\alpha_{it} = 0$ otherwise. For demand profile \mathcal{D}_2 , we set $\alpha_{it} = 0$ with probability $1/2$, but otherwise, generate α_{it} from the uniform distribution over $[5, 15]$. Both demand profiles \mathcal{D}_1 and \mathcal{D}_2 are examples of zero-inflated Poisson processes

described in Lambert (1992), but we note that demand profile \mathcal{D}_2 yields additional nonstationarity over demand profile \mathcal{D}_1 . We generate seven instances of each demand profile by using seven random seeds to generate $\{\alpha_{it} : i \in \mathcal{I}, t \in \mathcal{T}\}$. This approach provides the 21 test problems in the base case.

For each instance of the demand profiles, we vary one problem parameter at a time to be able to investigate how that problem parameter affects the performances of the benchmark methods. In particular, we vary the holding cost at the warehouse, the backloging cost at the retailers, the probability with which α_{it} is equal to 0 in demand profiles \mathcal{D}_1 and \mathcal{D}_2 , the number of retailers, the length of the lead times and the number of time periods in the planning horizon.

For the setting with nonidentical retailers, we consider an inventory distribution system with two retailers and a planning horizon of 50 time periods as the base case. The retailers are nonidentical in the sense that the mean demand at the second retailer is 90% of the mean demand at the first retailer. The remaining problem parameters are the same as for identical retailers. In particular, we have $h_{\phi t} = 0.6$, $h_{it} = 1$, $b_{it} = 19$ for all $i \in \mathcal{I}$, $t \in \mathcal{T}$ and it takes one time period for the replenishment order shipped to a certain installation to reach that installation. The demands at the retailers and the warehouse are generated according to demand profiles \mathcal{D}_1 and \mathcal{D}_2 with the understanding that we now have $\alpha_{2t} = 9$ with probability 1/2 for demand profile \mathcal{D}_1 and α_{2t} is generated from the uniform distribution over $[4.5, 13.5]$ with probability 1/2 for demand profile \mathcal{D}_2 .

Similar to our approach for the setting with identical retailers, we vary one problem parameter at a time to investigate how that problem parameter affects the performances of the benchmark methods. In particular, we vary the relative sizes of the two retailers as measured by the ratio of their mean demands and the relative holding costs and backloging costs at the two retailers.

6.3 COMPUTATIONAL RESULTS WITH IDENTICAL RETAILERS

We first describe the results of our computational experiments on the test problems with identical retailers. Table 1 shows the total expected costs incurred by the five benchmark methods for the 21 test problems in the base case. Each row corresponds to a particular instance of a demand profile. The first column shows the demand profile and the instance number. The second to sixth columns respectively show the total expected costs incurred by DL, RL, SG, BA and RS. We estimate these total expected costs by simulating the performances of the inventory replenishment policies obtained by the benchmark methods under multiple demand realizations. The seventh to tenth columns show the percent performance gaps between RL and the remaining four benchmark methods. We use RL as a reference point since it turns out to be one of the better benchmark methods. We simulate the performances of the inventory replenishment policies for enough number of demand realizations so that any performance gap above 1% is statistically significant at 5% level. It is not possible to guarantee that smaller performance gaps are statistically significant, since there are some test problems where the benchmark methods perform very similarly and making a statistically significant distinction between them requires a very large number of demand realizations.

Table 1 indicates all of the benchmark methods perform very similarly under demand profile \mathcal{D}_0 , which is the demand profile that yields relatively stationary demand. However, when we move to demand profiles \mathcal{D}_1 and \mathcal{D}_2 , we observe that RL can perform significantly better than DL, RS and BA. The performance gap between RL and DL can be as high as 21.20%, whereas the performance gap between RL and RS can be as high as 36.87%. This indicates that it is important to address the uncertainty in the demand random variables when choosing the Lagrange multipliers rather than just using the expected values of the demand random variables as in DL, and using the trivial value of zero for the Lagrange multipliers as in RS can result in poor performance. BA lags behind RL and RL can reduce the total costs by as much as 19.57% when compared to BA. Moreover, as we move from demand profile \mathcal{D}_0 to \mathcal{D}_1 and from \mathcal{D}_1 to \mathcal{D}_2 , we introduce additional nonstationarity and the performance gap between RL and BA increases. The performance of RL is comparable with that of SG. Although RL does not perform better than SG uniformly, it provides a small improvement over SG in many of the test problems. The results are quite encouraging, since implementing SG is much more complicated than implementing RL and the runtimes for SG are significantly longer. We compare the runtimes at the end of this section.

Table 2 shows the lower bounds on the optimal total expected cost provided by the five benchmark methods for the 21 test problems in the base case. The first column shows the demand profile and the instance number. Letting λ_{DL}^* , λ_{RL}^* and λ_{SG}^* respectively be the Lagrange multipliers chosen by DL, RL and SG, and noting that $z_{LP}^*(d)$ is the optimal objective value of problem (13)-(16), the second to fourth columns respectively show $V_1^{\lambda_{DL}^*}(x_1)$, $\max\{\mathbb{E}\{z_{LP}^*(d)\}, V_1^{\lambda_{RL}^*}(x_1)\}$ and $V_1^{\lambda_{SG}^*}(x_1)$. Propositions 1 and 2 show that $\mathbb{E}\{z_{LP}^*(d)\}$ and $V_1^{\lambda_{RL}^*}(x_1)$ are lower bounds on the optimal total expected cost. We use the tightest of these as the lower bound provided by RL. Proposition 2 shows that both $z_{LP}^*(\mathbb{E}\{d\})$ and $V_1^{\lambda_{DL}^*}(x_1)$ are lower bounds on the optimal total expected cost and the lower bound $V_1^{\lambda_{DL}^*}(x_1)$ is at least as tight as the lower bound $z_{LP}^*(\mathbb{E}\{d\})$. Therefore, we use $V_1^{\lambda_{DL}^*}(x_1)$ as the lower bound provided by DL. The fifth column in Table 2 shows $V_1^{BA}(x_1)$, where $\{V_t^{BA}(\cdot) : t \in \mathcal{T}\}$ are the value functions obtained under the balance assumption. The sixth column shows $V_1^0(x_1)$, which is the lower bound obtained by using the trivial value of zero for the Lagrange multipliers. The seventh to tenth columns simply show the percent gaps between the lower bounds provided by RL and the remaining four benchmark methods.

The results in Table 2 indicate that all of the benchmark methods provide very similar lower bounds under demand profile \mathcal{D}_0 . It is also important to emphasize that under this demand profile, the lower bounds in Table 2 are essentially equal to the total expected costs in Table 1. Therefore, all of the benchmark methods obtain essentially the optimal policy for the test problems under demand profile \mathcal{D}_0 . We see differences between the benchmark methods when we move to demand profiles \mathcal{D}_1 and \mathcal{D}_2 . SG provides the tightest lower bounds and is followed by RL, DL and BA. RS consistently provides the loosest lower bounds. SG can improve upon the lower bounds provided by RL by up to 1.59%. However, we also observe test problems where RL provides a slightly tighter lower bound than SG. Since $\lambda_{RL}^* \geq 0$, we trivially have $V_1^{\lambda_{RL}^*}(x_1) \leq \max_{\lambda \geq 0} \{V_1^\lambda(x_1)\} = V_1^{\lambda_{SG}^*}(x_1)$. Therefore, for RL to provide a tighter lower bound than SG, we need to have $\mathbb{E}\{z_{LP}^*(d)\} \geq V_1^{\lambda_{SG}^*}(x_1)$ so that we can get $\max\{\mathbb{E}\{z_{LP}^*(d)\}, V_1^{\lambda_{RL}^*}(x_1)\} \geq V_1^{\lambda_{SG}^*}(x_1)$. This indeed happens to be the case for some

test problems. RL can improve upon the lower bounds provided by DL, BA and RS by as much as 2.51%, 3.93% and 14.83%, respectively. We also note that under demand profiles \mathcal{D}_1 and \mathcal{D}_2 , there are significant gaps between the lower bounds in Table 2 and the total expected costs in Table 1. Therefore, there is still some progress to be made in terms of either tightening the lower bounds or improving the total expected costs. Throughout the rest of this section, we omit the results for demand profile \mathcal{D}_0 and focus on demand profiles \mathcal{D}_1 and \mathcal{D}_2 . Under demand profile \mathcal{D}_0 , all of the benchmark methods continue to perform very similarly and the results are not discriminative.

In Tables 3 to 9, we show how the five benchmark methods perform when we vary different problem parameters. We begin by varying the holding cost at the warehouse in Table 3. For economy of space in Tables 3 to 9, we only show the average results over seven instances of a particular demand profile. More specifically, the two portions of Table 3 respectively focus on demand profiles \mathcal{D}_1 and \mathcal{D}_2 . In each portion, the first column shows the holding cost at the warehouse. The second to fifth columns show the percent gaps between the total expected costs incurred by RL and the remaining four benchmark methods, averaged over seven instances of a particular demand profile. The sixth to ninth columns do essentially the same thing as the previous four columns, but they compare the lower bounds on the optimal total expected cost provided by RL with those provided by the remaining four benchmark methods. The percent gaps are calculated in such a manner that positive values always favor RL and negative values favor the other benchmark methods. We also note that averaging does not change the interpretation of the results significantly. This is to say that if the average performance of a benchmark method is noticeably better than that of another benchmark method, then the performance is generally better over all of the seven instances.

Table 3 indicates that RL performs better than BA and the average performance gap between RL and BA increases as the holding cost at the warehouse increases. We interpret this trend shortly in this section. The performance of RL remains comparable with that of SG throughout. DL can perform well when the holding cost at the warehouse is very low, but it performs worse than RL otherwise. Overall, the quality of the lower bounds provided by RL is reasonable when compared with the other benchmark methods and RL can provide the tightest lower bounds in certain cases.

Table 4 varies the backlogging cost at the retailers. BA incurs the lowest total expected costs among all of the benchmark methods when the backlogging cost at the retailers is 1. Nevertheless, a backlogging cost of 1 is quite low noting that the holding cost at the retailers is also 1. As the backlogging cost at the retailers increases, the performance of RL starts dominating the performance of BA. We shortly give an interpretation for this trend as well.

Table 5 varies the probability with which α_{it} is equal to zero. Using $p = \mathbb{P}\{\alpha_{it} = 0\}$ to denote this probability, under demand profiles \mathcal{D}_1 and \mathcal{D}_2 , we have $\mathbb{E}\{d_{it}\} = (1 - p)10$ and $\text{Var}(d_{it}) = (1 - p)(1 + 10p)10$. Therefore, the relative variation as measured by the coefficient of variation is an increasing function of p . Table 5 indicates that the performance gap between RL and the other solution methods increases with p . We note that setting $p = 0$ for demand profile \mathcal{D}_1 corresponds to the case of stationary demand. In that case, we see that all the benchmark methods perform equally well.

Table 6 varies the number of retailers in the system. The average performance gap between RL and BA increases as the number of retailers increases. When there are five retailers, the average performance gap between RL and BA can be as high as 16.66%. Clark and Scarf (1960) show that BA obtains the optimal policy when there is only one retailer in the system. RL does not have such a guarantee and it performs slightly worse than BA for the test problems with one retailer.

Table 7 varies the lead time for the replenishments of the warehouse and Table 8 varies the lead time for the replenishments of the retailers. As the lead time for the warehouse increases, the performances of RL, SG and BA become comparable. RL continues to provide substantial improvements over DL and RS. RL maintains its advantage over BA under a variety of lead times for the retailers. The performances of RL and SG become comparable as the lead times for the retailers increase.

Table 9 varies the number of time periods in the planning horizon. There does not appear to be a systematic change in the relative performances of the benchmark methods as we vary the length of the planning horizon. The overall performance of RL is quite satisfactory and RL performs significantly better than BA, DL and RS. Under demand profile \mathcal{D}_2 , RL has a small advantage over SG.

To sum up, RL provides advantages compared with the remaining benchmark methods when the demand distributions involve significant nonstationarity or have a high coefficient of variation. In addition to these factors, high holding cost at the warehouse and high backlogging cost at the retailers appear as two other factors that may improve the relative performance of RL. When the holding cost at the warehouse and the backlogging cost at the retailers are high, a naive policy has the tendency to push the inventory to the retailers so as not to incur the high holding cost at the warehouse and the high backlogging cost at the retailers. This results in keeping a relatively small amount of inventory at the warehouse and a large amount of inventory at the retailers. Since the inventory at a particular retailer can only be used at that retailer, whereas the inventory at the warehouse can be shipped to any of the retailers, this situation ultimately damages the flexibility of the system. RL appears to overcome the tendency of a naive policy to push the inventory to the retailers.

6.4 COMPUTATIONAL RESULTS WITH NONIDENTICAL RETAILERS

For the test problems with nonidentical retailers, the base case has two retailers and the mean demand at the second retailer is 90% of the mean demand at the first retailer. We use the mean demand as a proxy for the size of the retailer so that the second retailer is the smaller retailer. Table 10 shows the total expected costs incurred by the five benchmark methods for the 14 test problems in the base case. The rows and columns have the same interpretation as in Table 1. Table 10 indicates that RL provides the best performance, followed by SG. The performances of BA and DL are comparable. RS consistently performs the worst among the benchmark methods.

Table 11 shows the lower bounds on the optimal total expected costs provided by the five benchmark methods for the 14 problems in the base case with nonidentical retailers. SG provides the tightest lower bounds followed by BA. RL generates tighter lower bounds than DL and RS.

Table 12 varies the relative sizes of the two retailers as measured by the ratio of the mean demand at second retailer with respect to the first retailer. RL continues to provide substantial improvements over the other benchmark methods. We note that we have identical retailers when the ratio of the mean demands at the two retailers is equal to one. The performance gap between RL and DL gets larger as the relative sizes of the retailers become more different, but there does not appear to be a systematic change in the other performance gaps. Overall, RL continues to provide significant improvements over all of the other benchmark strategies.

Table 13 varies the holding cost at the second, smaller retailer. We note that when the holding cost is 0.6, it costs the same to hold a unit of inventory at the second retailer as at the warehouse. On the other hand, when the holding cost is 1.2, the second retailer incurs the highest holding cost in the distribution system. The performance gap between RL and the other benchmark methods increases as the holding cost at the second retailer decreases. The performances of BA and RS seem to deteriorate significantly as the holding cost at the second retailer decreases. This phenomenon is also observed by Dogru et al. (2005), who refer to it as the forwarding to the small retailer effect. Table 14 varies the backlogging cost at the second retailer. RL performs significantly better than the other benchmark methods and the performance gap increases as the backlogging cost at the second retailer gets larger. The average performance gap between RL and SG can be as high as 4.64% while that between RL and BA can be as high as 14.67%.

To sum up, RL provides advantages over the other benchmark methods when the smaller retailer has a low holding cost or a high backlogging cost. When the smaller retailer has a low holding cost and a high backlogging cost, a naive policy has the tendency to push the inventory from the warehouse to the smaller retailer. Since the smaller retailer experiences a lower mean demand, this results in build up of inventory at the smaller retailer. RL appears to overcome this tendency of a naive policy.

To give an overall statistic, RL incurs the lowest total expected cost in 268 of our 483 test problems. Of the 215 test problems where RL does not incur the lowest total expected cost, there are 115 test problems where there is a statistically significant difference between the performances of RL and the best performing benchmark method. Table 15 compares RL with the remaining four benchmark methods over all of our test problems. The second to fifth columns respectively correspond to DL, SG, BA and RS. The first row shows the number of test problems where RL performs better than a particular benchmark method and the performance gap is statistically significant. The second row shows the number of test problems where there is no statistically significant performance gap. The third row shows the number of test problems where RL performs worse than a particular benchmark method and the performance gap is statistically significant. When we compare RL with SG, which is its closest competitor, there are only 74 test problems where RL performs worse than SG and the performance gap is statistically significant. In the remaining 409 test problems, either RL performs better than SG or there is no statistically significant performance gap between the two benchmark methods. These observations indicate that RL is a very strong competitor of SG. Similar observations apply when we compare RL with DL, BA and RS.

All of our computational experiments were carried out on an Intel Xeon server with 3 GHz CPU and 4 GB RAM running Windows Server 2003. Table 16 shows the CPU seconds required by the different benchmark methods to compute one set of value function approximations. The top and bottom portions of this table respectively show the CPU seconds as a function of the number of time periods in the planning horizon and the number of retailers in the system. RS and BA have the shortest CPU seconds and they are followed by DL. SG has the longest CPU seconds. The CPU seconds for RL lie somewhere between those for SG and DL. RL essentially solves the same linear programming approximation solved by DL for 25 demand samples, but the CPU seconds for RL are about five times the CPU seconds for DL. This is due to the fact that when solving the linear programming approximation for one demand sample, RL starts from the optimal basis obtained for the previous demand sample. This warm start significantly boosts the CPU seconds for RL. The CPU seconds for RL are noticeably shorter than those for SG and the CPU seconds for RL and SG can differ by almost a factor of ten. Furthermore, solving problem (12) for SG requires solving a nonsmooth constrained convex optimization problem, which, in turn, requires calibrating the step size parameter and devising a stopping criterion in a subgradient search procedure. On the other hand, RL solves the linear program in (13)-(16) for multiple demand samples and this can be done by using commercial linear programming software. Given that the performance of RL is comparable or superior to the performance of SG for a large majority of our test problems, and RL can perform significantly better than the other benchmark methods, RL appears to be a viable approach to solve inventory distribution problems. Furthermore, its ease of implementation and relatively short CPU seconds enable RL to carry out quick what if analyses, and hence, make RL quite appealing to the practitioners.

7 CONCLUSIONS

In this paper, we developed a tractable method to make the inventory replenishment decisions in a distribution system. Our approach is based on the dynamic programming formulation of the inventory distribution problem, but due to the high dimensionality of the state variable, the dynamic programming formulation ends up being difficult to solve. We relax the constraints that ensure the nonnegativity of the shipments to the retailers by associating Lagrange multipliers with them. As a result of this relaxation, the dynamic programming formulation of the inventory distribution problem decomposes by the installations and can be solved by focusing on one installation at a time. This approach naturally brings up the question of how to choose the Lagrange multipliers. To answer this question, Kunnumkal and Topaloglu (2008) solve a nonsmooth constrained convex optimization problem, but their method can be difficult to implement in practice and can require excessive computation time. In this paper, we proposed alternative methods for choosing a good set of values for the Lagrange multipliers. One method solves a linear programming approximation to the inventory distribution problem that is formulated under the assumption that the demand random variables take on their expected values, whereas the other method solves a linear programming approximation that is formulated under the assumption that the realizations of the demand random variables are known a priori. Computational experiments demonstrated that the Lagrange multipliers provided by especially the second method can perform significantly better than several standard benchmarks. In particular, the total expected costs incurred

by our second method can be better than those incurred by Kunnumkal and Topaloglu (2008) by up to 7.62%. Furthermore, both of our methods can easily be implemented by using commercial linear programming software and their computation times can be shorter than those for the method proposed by Kunnumkal and Topaloglu (2008) by about a factor of ten. Thus, given their ease of implementation, short computation time and superior solution quality, the methods that we propose in this paper form viable alternatives for making the inventory replenishment decisions in distribution systems.

ACKNOWLEDGEMENTS

The authors thank two anonymous referees for their useful comments that especially strengthened the exposition and computational results. This work was supported in part by National Science Foundation grant CMMI-0758441. The first author acknowledges the support of Indian School of Business.

REFERENCES

- Aviv, Y. and Federgruen, A. (2001), ‘Design for postponement: A comprehensive characterization of its benefits under unknown demand distributions’, *Operations Research* **49**(4), 578–598.
- Axsater, S. (2003), Supply chain operations: Serial and distribution inventory systems, in S. C. Graves and A. G. de Kok, eds, ‘*Handbook of Operations Research, Supply Chain Management: Design, Coordination and Operation*’, Elsevier Science Publishers B. V., Amsterdam, North Holland, pp. 525–559.
- Bertsekas, D. P. (2001), *Dynamic Programming and Optimal Control*, Athena Scientific, Belmont, MA.
- Chen, F. and Song, J.-S. (2001), ‘Optimal policies for multiechelon inventory problems with markov modulated demand’, *Operations Research* **49**(2), 226–234.
- Clark, A. J. and Scarf, H. (1960), ‘Optimal policies for a multi-echelon inventory problem’, *Management Science* **64**, 475–490.
- Dogru, M. K. (2005), Optimal Control of One-Warehouse Multi-Retailer Systems: An Assessment of the Balance Assumption, PhD thesis, Technische Universiteit Eindhoven, Eindhoven, Netherlands.
- Dogru, M. K., de Kok, A. G. and van Houtum, G. J. (2005), A numerical study on the effect of the balance assumption in one-warehouse multi-retailer inventory systems, Beta working paper, WP 135, Department of Technology Management, Technische Universiteit Eindhoven.
- Eppen, G. L. and Schrage, L. (1981), Centralized ordering policies in a multi-warehouse system with leadtimes and random demand, in L. Schwarz, ed., ‘*Multi-Level Production/Inventory Control Systems: Theory and Practice*’, North Holland, Amsterdam, The Netherlands, pp. 51–69.
- Federgruen, A. (1993), Centralized planning models for multi-echelon inventory systems under uncertainty, in S. C. Graves, A. H. G. R. Kan and P. H. Zipkin, eds, ‘*Handbook in Operations Research and Management Science, Volume on Logistics of Production and Inventory*’, North Holland, Amsterdam, pp. 3–50.
- Federgruen, A. and Zipkin, P. (1984a), ‘Allocation policies and cost approximations for multilocation inventory systems’, *Naval Research Logistics* **31**, 97–129.
- Federgruen, A. and Zipkin, P. (1984b), ‘Approximations of dynamic, multilocation production and inventory problems’, *Management Science* **30**(1), 69–84.
- Federgruen, A. and Zipkin, P. (1984c), ‘Computational issues in an infinite horizon multi-echelon inventory model’, *Operations Research* **32**, 818–836.

Test prob.	Total exp. cost					Percent gap with RL			
	DL	RL	SG	BA	RS	DL	SG	BA	RS
($\mathcal{D}_0, 1$)	2,660	2,660	2,660	2,660	2,660	0.00	0.00	0.00	0.00
($\mathcal{D}_0, 2$)	2,562	2,563	2,562	2,562	2,562	-0.04	-0.04	-0.04	-0.04
($\mathcal{D}_0, 3$)	2,680	2,680	2,680	2,680	2,680	0.00	0.00	0.00	0.00
($\mathcal{D}_0, 4$)	2,659	2,659	2,659	2,659	2,659	0.00	0.00	0.00	0.00
($\mathcal{D}_0, 5$)	2,714	2,714	2,714	2,714	2,714	0.00	0.00	0.00	0.00
($\mathcal{D}_0, 6$)	2,675	2,675	2,675	2,675	2,675	0.00	0.00	0.00	0.00
($\mathcal{D}_0, 7$)	2,695	2,695	2,695	2,695	2,695	0.00	0.00	0.00	0.00
Average						0.00	0.00	0.00	0.00
($\mathcal{D}_1, 1$)	1,878	1,773	1,732	1,959	2,108	5.91	-2.33	10.50	18.88
($\mathcal{D}_1, 2$)	2,053	1,758	1,755	1,985	2,264	16.79	-0.16	12.95	28.78
($\mathcal{D}_1, 3$)	1,807	1,764	1,781	2,022	2,318	2.42	0.94	14.60	31.36
($\mathcal{D}_1, 4$)	1,590	1,500	1,507	1,741	1,991	5.99	0.45	16.08	32.76
($\mathcal{D}_1, 5$)	1,796	1,655	1,659	1,945	2,230	8.50	0.25	17.53	34.76
($\mathcal{D}_1, 6$)	1,887	1,699	1,756	2,002	2,290	11.04	3.36	17.79	34.77
($\mathcal{D}_1, 7$)	1,898	1,770	1,785	2,013	2,250	7.23	0.87	13.73	27.12
Average						8.27	0.48	14.74	29.78
($\mathcal{D}_2, 1$)	1,728	1,659	1,643	1,935	2,113	4.13	-0.96	16.60	27.36
($\mathcal{D}_2, 2$)	1,876	1,806	1,840	1,992	2,166	3.88	1.90	10.30	19.90
($\mathcal{D}_2, 3$)	1,716	1,613	1,639	1,794	1,930	6.39	1.60	11.21	19.70
($\mathcal{D}_2, 4$)	1,956	1,614	1,672	1,905	2,143	21.20	3.59	18.05	32.76
($\mathcal{D}_2, 5$)	1,773	1,699	1,752	2,032	2,226	4.33	3.14	19.57	30.99
($\mathcal{D}_2, 6$)	1,746	1,707	1,718	2,025	2,337	2.26	0.62	18.58	36.87
($\mathcal{D}_2, 7$)	1,813	1,718	1,745	1,977	2,146	5.56	1.58	15.09	24.93
Average						6.82	1.64	15.63	27.50

Table 1: Total expected costs incurred by the benchmark methods for the 21 test problems in the base case with identical retailers.

- Gallego, G. and Ozer, O. (2003), ‘Optimal replenishment policies for multiechelon inventory problems under advance demand information’, *Manufacturing & Service Operations Management* **5**(2), 157–175.
- Gallego, G., Ozer, O. and Zipkin, P. (2007), ‘Bounds, heuristics, and approximations for distribution systems’, *Operations Research* **55**(3), 503–517.
- Jackson, P. L. (1988), ‘Stock allocation in a two-echelon distribution system or “what to do until your ship comes in”’, *Management Science* **34**(7), 880–895.
- Kunnumkal, S. and Topaloglu, H. (2008), ‘A duality-based relaxation and decomposition approach for inventory distribution systems’, *Naval Research Logistics Quarterly* **55**(7), 612–631.
- Lambert, D. (1992), ‘Zero-inflated poisson regression models with an application to defects in manufacturing’, *Technometrics* **34**, 1–14.
- Ozer, O. (2003), ‘Replenishment strategies for distribution systems under advance demand information’, *Management Science* **49**(3), 255–272.
- Talluri, K. T. and van Ryzin, G. J. (2005), *The Theory and Practice of Revenue Management*, Springer, New York, NY.
- Talluri, K. and van Ryzin, G. (1999), ‘A randomized linear programming method for computing network bid prices’, *Transportation Science* **33**(2), 207–216.

Test prob.	Lower bound					Percent gap with RL			
	DL	RL	SG	BA	RS	DL	SG	BA	RS
($\mathcal{D}_0, 1$)	2,651	2,650	2,651	2,651	2,651	-0.04	-0.04	-0.04	-0.04
($\mathcal{D}_0, 2$)	2,560	2,558	2,560	2,560	2,560	-0.08	-0.08	-0.08	-0.08
($\mathcal{D}_0, 3$)	2,680	2,679	2,680	2,680	2,680	-0.04	-0.04	-0.04	-0.04
($\mathcal{D}_0, 4$)	2,659	2,657	2,659	2,659	2,659	-0.08	-0.08	-0.08	-0.08
($\mathcal{D}_0, 5$)	2,714	2,712	2,714	2,714	2,714	-0.07	-0.07	-0.07	-0.07
($\mathcal{D}_0, 6$)	2,675	2,673	2,675	2,675	2,675	-0.07	-0.07	-0.07	-0.07
($\mathcal{D}_0, 7$)	2,694	2,693	2,694	2,694	2,694	-0.04	-0.04	-0.04	-0.04
Average						0.00	0.00	0.00	0.00
($\mathcal{D}_1, 1$)	1,406	1,425	1,428	1,389	1,234	1.31	-0.20	2.52	13.39
($\mathcal{D}_1, 2$)	1,536	1,539	1,561	1,533	1,407	0.21	-1.45	0.36	8.60
($\mathcal{D}_1, 3$)	1,476	1,490	1,496	1,465	1,328	0.95	-0.44	1.67	10.84
($\mathcal{D}_1, 4$)	1,162	1,192	1,195	1,163	1,015	2.51	-0.28	2.44	14.83
($\mathcal{D}_1, 5$)	1,374	1,393	1,401	1,359	1,210	1.40	-0.59	2.44	13.18
($\mathcal{D}_1, 6$)	1,397	1,424	1,422	1,392	1,246	1.86	0.09	2.22	12.51
($\mathcal{D}_1, 7$)	1,486	1,510	1,513	1,486	1,356	1.56	-0.23	1.56	10.20
Average						1.40	-0.44	1.89	11.94
($\mathcal{D}_2, 1$)	1,393	1,399	1,411	1,378	1,237	0.39	-0.89	1.49	11.56
($\mathcal{D}_2, 2$)	1,511	1,510	1,534	1,503	1,376	-0.07	-1.59	0.46	8.83
($\mathcal{D}_2, 3$)	1,372	1,402	1,395	1,368	1,236	2.17	0.55	2.42	11.86
($\mathcal{D}_2, 4$)	1,292	1,323	1,319	1,286	1,132	2.35	0.28	2.76	14.42
($\mathcal{D}_2, 5$)	1,366	1,386	1,378	1,332	1,181	1.47	0.58	3.93	14.81
($\mathcal{D}_2, 6$)	1,351	1,376	1,371	1,335	1,186	1.83	0.33	2.96	13.81
($\mathcal{D}_2, 7$)	1,447	1,458	1,463	1,425	1,288	0.81	-0.34	2.26	11.72
Average						1.28	-0.15	2.33	12.43

Table 2: Lower bounds on the optimal total expected cost provided by the benchmark methods for the 21 test problems in the base case with identical retailers.

\mathcal{D}_1

Hld. cst.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.1	-4.94	-1.16	3.67	30.42	2.50	1.52	-1.03	25.52
0.3	0.87	0.96	8.14	29.08	1.94	0.06	-0.62	18.41
0.6	8.27	0.48	14.74	29.78	1.40	-0.44	1.89	11.94
0.9	13.83	-0.14	21.88	28.13	1.11	-0.53	5.78	8.26

\mathcal{D}_2

Hld. cst.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.1	-5.41	0.56	4.53	25.93	2.21	1.58	-0.85	26.09
0.3	0.81	1.01	9.59	26.15	1.72	0.31	-0.29	18.94
0.6	6.82	1.64	15.63	27.50	1.28	-0.15	2.33	12.43
0.9	12.86	-0.48	21.55	25.83	1.04	-0.24	6.14	8.62

Table 3: Effect of the holding cost at the warehouse on the performance of the benchmark methods.

\mathcal{D}_1

Bck. cst.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	24.30	2.35	-2.19	2.62	3.47	-1.72	-5.56	-1.71
9	18.38	0.36	5.55	17.27	3.56	-0.99	-0.87	8.40
19	8.27	0.48	14.74	29.78	1.40	-0.44	1.89	11.94
29	0.78	0.22	20.23	36.70	0.73	-0.27	2.89	13.07
39	-1.16	1.13	26.95	46.55	0.50	0.09	3.74	13.81

\mathcal{D}_2

Bck. cst.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	8.75	1.97	-2.29	2.15	0.55	-1.61	-5.73	-1.59
9	14.96	0.64	5.94	15.83	3.32	-0.61	-0.32	9.26
19	6.82	1.64	15.63	27.50	1.28	-0.15	2.33	12.43
29	2.12	1.95	22.44	36.11	0.66	-0.03	3.24	13.24
39	-2.13	2.15	30.71	43.93	0.41	0.10	3.71	13.46

Table 4: Effect of the backlogging cost at the retailers on the performance of the benchmark methods.

\mathcal{D}_1

$\Pr(\alpha_{it} = 0)$	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
0.25	6.00	-0.81	3.83	7.61	0.24	-0.75	-0.04	3.54
0.50	8.27	0.48	14.74	29.78	1.40	-0.44	1.89	11.94
0.75	26.58	7.62	19.45	48.85	6.24	0.40	1.18	20.67

\mathcal{D}_2

$\Pr(\alpha_{it} = 0)$	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.00	0.00	0.00	0.00	0.00	-0.06	-0.06	-0.06	-0.06
0.25	7.36	-0.43	4.37	8.28	0.17	-0.89	-0.18	3.46
0.50	6.82	1.64	15.63	27.50	1.28	-0.15	2.33	12.43
0.75	13.54	6.29	19.24	44.61	4.99	0.95	2.05	20.25

Table 5: Effect of the probability with which α_{it} is equal to zero on the performance of the benchmark methods.

\mathcal{D}_1

No. ret.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	1.71	-1.02	-1.00	-0.10	-0.84	-2.81	-3.43	5.93
2	6.44	3.32	10.02	26.50	0.85	-2.75	-1.55	8.76
3	8.27	0.48	14.74	29.78	1.40	-0.44	1.89	11.94
4	9.23	-1.39	17.21	30.18	1.53	0.69	3.91	13.49
5	8.74	-2.41	16.08	27.30	1.69	1.31	5.09	14.29

\mathcal{D}_2

No. ret.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	5.59	-1.44	-1.47	-0.72	1.18	-3.36	-4.01	6.69
2	4.45	3.03	9.48	20.31	0.83	-2.32	-1.12	9.13
3	6.82	1.64	15.63	27.50	1.28	-0.15	2.33	12.43
4	7.30	-1.02	16.08	27.09	1.30	0.72	4.13	14.13
5	7.54	-1.12	16.66	26.26	1.40	1.30	5.55	15.51

Table 6: Effect of the number of retailers on the performance of the benchmark methods.

\mathcal{D}_1

Wh. l.t.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	8.27	0.48	14.74	29.78	1.40	-0.44	1.89	11.94
2	27.21	0.95	6.41	19.68	1.45	-3.21	-2.99	6.85
3	31.94	-0.24	1.84	12.95	1.86	-5.53	-6.04	3.30

\mathcal{D}_2

Wh. l.t.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	6.82	1.64	15.63	27.50	1.28	-0.15	2.33	12.43
2	25.42	-0.15	6.88	17.31	1.33	-2.95	-2.57	7.38
3	30.53	-0.66	1.92	11.62	1.48	-5.45	-5.90	3.57

Table 7: Effect of the lead time for the replenishments of the warehouse on the performance of the benchmark methods.

\mathcal{D}_1

Ret. l.t.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	8.27	0.48	14.74	29.78	1.40	-0.44	1.89	11.94
2	1.99	-0.56	10.58	16.44	0.39	-0.02	1.79	6.40
3	0.11	0.31	7.65	10.50	0.18	-0.12	0.93	3.18

\mathcal{D}_2

Ret. l.t.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	6.82	1.64	15.63	27.50	1.28	-0.15	2.33	12.43
2	0.35	0.05	9.98	14.84	0.61	0.09	1.78	6.37
3	0.97	0.30	6.33	8.78	0.35	0.00	0.84	3.12

Table 8: Effect of the lead time for the replenishments of the retailers on the performance of the benchmark methods.

\mathcal{D}_1

Pln. hor.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
25	8.12	-0.38	12.89	28.77	1.30	-0.16	1.64	9.87
50	8.27	0.48	14.74	29.78	1.40	-0.44	1.89	11.94
75	6.83	-0.18	14.52	28.78	1.32	-0.57	1.73	11.79
100	6.45	0.53	14.96	30.82	1.19	-0.70	1.88	12.78

\mathcal{D}_2

Pln. hor.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
25	12.39	1.16	13.66	25.80	1.42	0.16	2.31	11.36
50	6.82	1.64	15.63	27.50	1.28	-0.15	2.33	12.43
75	6.49	0.20	15.64	28.75	0.99	-0.80	1.62	12.34
100	5.51	0.83	15.10	29.57	1.26	-0.47	2.36	13.81

Table 9: Effect of the length of the planning horizon on the performance of the benchmark methods.

Test prob.	Total exp. cost					Percent gap with RL			
	DL	RL	SG	BA	RS	DL	SG	BA	RS
($\mathcal{D}_1, 1$)	1,240	1,182	1,190	1,267	1,444	4.89	0.66	7.18	22.16
($\mathcal{D}_1, 2$)	1,260	1,212	1,259	1,309	1,490	3.96	3.88	8.02	22.96
($\mathcal{D}_1, 3$)	1,098	1,046	1,060	1,123	1,355	5.01	1.42	7.42	29.56
($\mathcal{D}_1, 4$)	1,076	1,019	1,021	1,097	1,233	5.63	0.25	7.63	21.05
($\mathcal{D}_1, 5$)	1,077	963	1,000	1,082	1,265	11.85	3.83	12.41	31.35
($\mathcal{D}_1, 6$)	1,355	1,139	1,198	1,252	1,575	18.94	5.15	9.93	38.21
($\mathcal{D}_1, 7$)	1,207	1,152	1,199	1,290	1,484	4.80	4.08	11.99	28.85
Average						7.87	2.75	9.22	27.73
($\mathcal{D}_2, 1$)	1,089	1,082	1,080	1,168	1,330	0.72	-0.14	7.99	22.96
($\mathcal{D}_2, 2$)	1,296	1,240	1,280	1,376	1,532	4.53	3.22	10.92	23.50
($\mathcal{D}_2, 3$)	1,229	1,193	1,199	1,237	1,323	3.02	0.53	3.75	10.88
($\mathcal{D}_2, 4$)	1,071	928	971	1,031	1,163	15.35	4.64	11.06	25.29
($\mathcal{D}_2, 5$)	1,397	1,200	1,243	1,318	1,434	16.47	3.60	9.84	19.55
($\mathcal{D}_2, 6$)	1,001	952	1,012	1,092	1,248	5.16	6.36	14.72	31.12
($\mathcal{D}_2, 7$)	1,205	1,144	1,198	1,268	1,378	5.33	4.72	10.85	20.47
Average						7.23	3.27	9.88	21.97

Table 10: Total expected costs incurred by the benchmark methods for the 14 test problems in the base case with nonidentical retailers.

Test prob.	Lower bound					Percent gap with RL			
	DL	RL	SG	BA	RS	DL	SG	BA	RS
($\mathcal{D}_1, 1$)	971	973	1,002	988	887	0.17	-2.97	-1.50	8.81
($\mathcal{D}_1, 2$)	1,040	1,045	1,069	1,058	973	0.51	-2.28	-1.23	6.95
($\mathcal{D}_1, 3$)	868	879	899	895	813	1.29	-2.28	-1.80	7.50
($\mathcal{D}_1, 4$)	828	841	864	857	768	1.54	-2.70	-1.92	8.66
($\mathcal{D}_1, 5$)	771	778	804	790	696	0.89	-3.39	-1.58	10.54
($\mathcal{D}_1, 6$)	874	888	913	896	785	1.48	-2.91	-0.97	11.54
($\mathcal{D}_1, 7$)	963	973	1,001	993	899	0.94	-2.94	-2.13	7.53
Average						0.97	-2.78	-1.59	8.79
($\mathcal{D}_2, 1$)	852	855	881	867	771	0.40	-2.96	-1.36	9.86
($\mathcal{D}_2, 2$)	1089	1101	1120	1111	1030	1.16	-1.68	-0.85	6.48
($\mathcal{D}_2, 3$)	1025	1032	1054	1046	959	0.68	-2.14	-1.37	7.07
($\mathcal{D}_2, 4$)	712	724	749	740	637	1.74	-3.37	-2.09	12.07
($\mathcal{D}_2, 5$)	968	971	993	977	876	0.29	-2.33	-0.61	9.80
($\mathcal{D}_2, 6$)	742	760	773	763	664	2.37	-1.78	-0.45	12.56
($\mathcal{D}_2, 7$)	1003	1008	1032	1026	936	0.51	-2.39	-1.83	7.11
Average						1.02	-2.38	-1.22	9.28

Table 11: Lower bounds on the optimal total expected cost provided by the benchmark methods for the 14 test problems in the base case with nonidentical retailers.

\mathcal{D}_1

Rel. size.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.5	12.96	3.38	9.37	23.21	1.04	-3.13	-2.09	8.18
0.6	12.24	3.73	10.31	26.84	1.22	-2.99	-1.89	8.64
0.7	9.69	3.17	9.58	25.37	0.66	-2.86	-1.72	8.74
0.8	8.08	2.75	9.64	25.53	0.87	-2.83	-1.69	8.77
0.9	7.87	2.75	9.22	27.73	0.97	-2.78	-1.59	8.79
1.0	6.44	3.32	10.02	26.50	0.85	-2.75	-1.55	8.76

\mathcal{D}_2

Rel. size.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.5	9.99	2.09	9.79	21.74	1.15	-2.72	-1.70	9.26
0.6	8.79	2.08	9.34	21.05	1.04	-2.53	-1.46	9.29
0.7	9.83	2.36	9.24	21.59	0.99	-2.45	-1.28	9.47
0.8	7.83	2.28	9.14	22.06	0.90	-2.38	-1.14	9.49
0.9	7.23	3.27	9.88	21.97	1.02	-2.38	-1.22	9.28
1.0	4.45	3.03	9.48	20.31	0.83	-2.32	-1.12	9.13

Table 12: Effect of the relative sizes of the retailers on the performance of the benchmark methods.

\mathcal{D}_1

Hld. cst.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.6	5.00	3.25	37.64	39.28	0.45	-2.52	0.95	6.64
0.8	7.60	3.05	14.43	31.58	0.68	-2.53	-0.53	7.97
1.0	7.87	2.75	9.22	27.73	0.97	-2.78	-1.59	8.79
1.2	7.60	2.53	7.23	24.65	1.18	-2.93	-2.36	9.41

\mathcal{D}_2

Hld. cst.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
0.6	6.73	1.74	25.97	26.26	0.46	-2.07	1.07	7.14
0.8	6.52	2.56	12.81	23.21	0.79	-2.19	-0.31	8.39
1.0	7.23	3.27	9.88	21.97	1.02	-2.38	-1.22	9.28
1.2	7.26	1.94	8.15	19.75	1.29	-2.51	-1.90	9.94

Table 13: Effect of the holding cost at the smaller retailer on the performance of the benchmark methods.

\mathcal{D}_1

Bck. cst	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	27.49	2.63	-0.80	5.24	1.29	-1.28	0.26	8.68
9	10.61	0.80	5.98	21.04	1.83	-3.05	-2.49	7.70
19	7.87	2.75	9.22	27.73	0.97	-2.78	-1.59	8.79
29	4.84	3.36	12.15	28.96	0.58	-2.34	2.27	12.52
39	3.39	4.64	14.67	34.28	0.46	-2.04	6.08	16.10

\mathcal{D}_2

Bck. cst.	Total exp. cost percent gap with RL				Lower bound percent gap with RL			
	DL	SG	BA	RS	DL	SG	BA	RS
1	29.28	1.09	1.29	7.03	0.95	-1.02	0.70	9.57
9	8.78	1.19	6.45	17.00	1.77	-2.80	-2.09	8.30
19	7.23	3.27	9.88	21.97	1.02	-2.38	-1.22	9.28
29	4.48	2.81	11.67	24.34	0.67	-2.07	2.74	13.07
39	3.00	3.45	13.13	25.66	0.43	-1.80	6.62	16.74

Table 14: Effect of the backloging cost at the smaller retailer on the performance of the benchmark methods.

No. test prob. with	DL	SG	BA	RS
RL better	381	230	421	447
No difference	66	179	36	34
RL worse	36	74	26	2

Table 15: Comparison of RL with the remaining four benchmark methods over all of our test problems.

Pln. hor.	CPU seconds				
	DL	RL	SG	BA	RS
25	3	14	182	0.2	0.2
50	10	55	422	0.8	0.5
75	23	94	850	0.9	0.8
100	31	150	944	1.1	1.0

No. ret.	CPU seconds				
	DL	RL	SG	BA	RS
1	1	7	109	0.2	0.2
2	5	31	351	0.3	0.4
3	10	55	422	0.8	0.5
4	18	89	722	0.9	0.6
5	30	137	985	1.0	1.0

Table 16: CPU seconds required by the benchmark methods to compute one set of value function approximations as a function of the number of time periods and the number of retailers.