PRODUCTION AND OPERATIONS MANAGEMENT

PRODUCTION AND OPERATIONS MANAGEMENT
Vol. 22, No. 5, September–October 2013, pp. 1200–1213
ISSN 1059-1478 | EISSN 1937-5956 | 13 | 2205 | 1200
DOI 10.1111/j.1937-5956.2012.01424.x
© 2013 Production and Operations Management Society

The Retail Planning Problem Under Demand Uncertainty

George Georgiadis, Kumar Rajaram
UCLA Anderson School of Management, 110 Westwood Plaza, Los Angeles, California 90095-1481, USA
georgiadis@ucla.edu, krajaram@anderson.ucla.edu

We consider the retail planning problem in which the retailer chooses suppliers and determines the production, distribution, and inventory planning for products with uncertain demand to minimize total expected costs. This problem is often faced by large retail chains that carry private-label products. We formulate this problem as a convex-mixed integer program and show that it is strongly NP-hard. We determine a lower bound by applying a Lagrangian relaxation and show that this bound outperforms the standard convex programming relaxation while being computationally efficient. We also establish a worst-case error bound for the Lagrangian relaxation. We then develop heuristics to generate feasible solutions. Our computational results indicate that our convex programming heuristic yields feasible solutions that are close to optimal with an average suboptimality gap at 3.4%. We also develop managerial insights for practitioners who choose suppliers and make production, distribution, and inventory decisions in the supply chain.

Key words: retailing; facility location; inventory management; stochastic demand; nonlinear integer programming
History: Received: May 2011; Accepted: August 2012 by Vishal Gaur, after 2 revisions.

1. Introduction

Retail store chains typically carry private-label merchandise. For example, the department store chain Macy's carries several private-label brands such as Alfani, Club Room, Hotel Collection, and others. Similarly, Target, J. C. Penney, and others carry their own private-label brands. Other retail store chains such as GAP, H&M, and Zara carry private-label products exclusively. Private labels allow firms to differentiate their products from those of their competitors and enhance customer loyalty, and they typically provide higher profit margins. However, these benefits are accompanied by additional challenges. The retailer must plan the entire supply chain by selecting suppliers and by making decisions on production, distribution, and inventory at the retail (and possibly other) locations for each of these private-label products to minimize total costs. This problem can be complicated when there is a large number of products with uncertain demand that can be sourced from various suppliers, and they are distributed across various demand zones. An example of such a supply chain is illustrated in Figure 1.

Private-label products can be produced in-house, or production can be outsourced to third-party suppliers. Without loss of generality, we refer to these options as suppliers. Supplier choice entails fixed costs such as building and staffing a plant when producing in-house or negotiating, contracting, and control costs when outsourcing it. Each production facility can manufacture multiple products interchangeably, and there are economies of scale in manufacturing and distribution. Demand at each zone (i.e., store or city) is stochastic, and inventory is carried at every demand zone. Here, demand zones can be interpreted either as retail stores or as distribution centers (DCs).[1] The retailer incurs overstock and understock costs for leftover inventory and unmet demand, respectively. In this context, there are three types of decisions. First, the retailers need to decide which suppliers to choose. Second, they need to conduct production and logistics planning. Third, inventory management decisions on how much of each product to stock at each demand zone need to be made.

We develop the retail planning problem (RPP) under uncertainty to address these decisions. In this problem, we model the selection of suppliers, production, distribution, and inventory decisions faced by the retailer as a nonlinear mixed-integer program that minimizes total expected costs. We show that this problem is convex and strongly NP-hard. An interesting attribute of this problem is that it combines two well-known subproblems: a generalized multi-commodity facility location problem and a newsvendor problem. We exploit this structure to develop computationally efficient heuristics to generate feasible solutions. In addition, we apply a Lagrangian relaxation

Figure 1 The Retail Supply Chain for Private-label Products

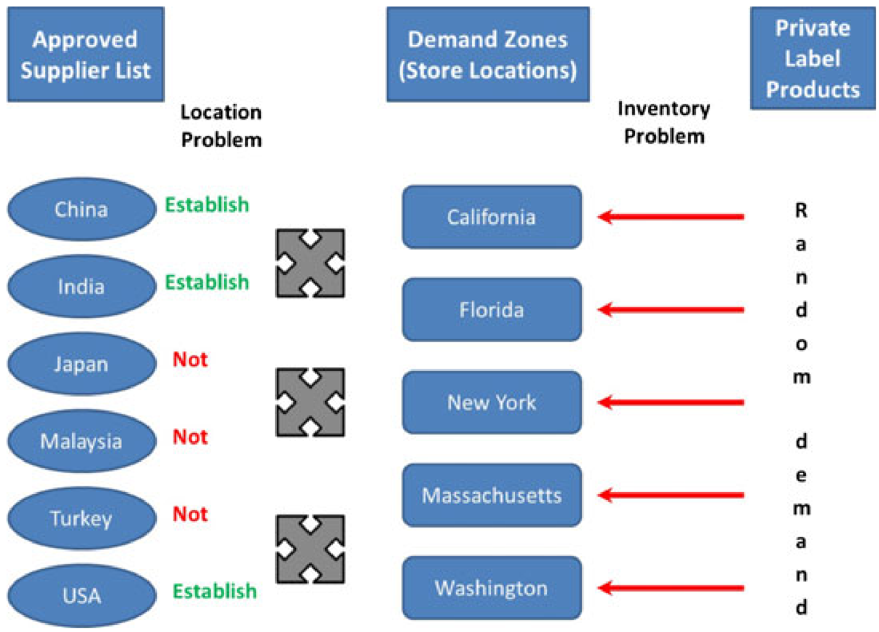

to obtain a lower bound, which we use to assess the quality of the feasible solutions provided by the heuristics. We show that the feasible solutions of a convex programming heuristic are close to optimal, on average within 3.4% of optimal, while in the majority of cases they are closer to optimal, as evidenced by the 2.8% median suboptimality gap. Furthermore, the performance gap of this heuristic improves with larger problem sizes, and the computational time of this heuristic scales up approximately linearly in the problem size. We also conduct robustness checks and find that the performance of this heuristic, as well as its advantage relative to the benchmark practitioner's heuristic, is not sensitive to changes in the problem parameters. All these are desirable attributes for potential implementation in large-sized, real applications.

Our analysis enables us to draw several managerial insights. First, the optimal inventory level when solving the joint supplier choice, production, distribution, and inventory problem is smaller than when the inventory subproblem is solved separately. This is because when solving the joint problem, the solution accounts for the fact that a larger downstream inventory level raises production quantities, which increases upstream production and distribution costs as well as the costs associated with establishing production capacity. In contrast, these costs are not considered when the inventory subproblem is solved separately, and hence result in a larger inventory level. Thus, to minimize total supply chain costs, one needs to adopt an integrated approach to solve the joint problem by considering the effect of downstream inventory decisions on upstream production and distribution costs. Our model provides a framework to analyze these decisions. Second, the two major costs that influence total (expected) supply chain costs are production costs and the understock costs associated with the variance in demand. Therefore, retailers should focus on reducing these costs first before considering the effects of supplier capacity and contracting costs. Third, it is important to consider establishment, production, distribution, and inventory costs together when choosing suppliers, because a supplier who is desirable in any one of these aspects may in fact not be the best overall choice. Our analysis provides a mechanism to integrate these aspects and pick the best set of suppliers.

As one of the decisions considered in the RPP under demand uncertainty is the establishment of production capacity by the explicit choice of suppliers, this problem can be placed in the broad category of facility location problems under uncertain demand. Aikens (1985), Drezner (1995), Melo et al. (2009), Owen and Daskin (1998), and Snyder (2006) provide extensive reviews. The problem with stochastic demand was first studied by Balachandran and Jain (1976) and Le Blanc (1977), who developed a branch and bound procedure and a Lagrangian heuristic, respectively. This study generalizes their models by considering multiple products, as well as incorporating economies of scale in production and distribution.

This study can also be placed in the general area of integrated supply chain models. Shen (2007) provided a comprehensive review of this area. In particular, this study is related to Daskin et al. (2002) and Shen et al. (2003), who studied a location-inventory problem in a supplier–DC–retailer network. Here, the planner's problem is to determine which DCs to establish, the inventory replenishment policy at each DC, and logistics between DCs and retailers. Daskin et al. (2002) and Shen et al. (2003) solved this problem by using a Lagrangian relaxation and a column-generation approach, respectively. Shen (2005) studied a multi-commodity extension of Daskin et al. (2002) with economies of scale but without explicitly modeling inventory decisions and without capacity constraints. Relative to these papers, we incorporate economies of scale in both production and distribution as well as capacity constraints at each supplier. Moreover, we explicitly model the inventory problem. Here, by using the newsvendor instead of a replenishment model to make inventory decisions, we capture features of the retail fashion industry, where lead times are long relative to product life cycles so that inventory cannot be replenished mid-season, and unmet demand is lost, resulting in underage costs.[2] A related problem was also studied by Oszen et al. (2008), who studied a capacitated extension of Shen et al. (2003). However, unlike these papers, we focus on the joint supplier choice, logistics, and inventory planning problem, as opposed to the risk pooling effects from strategically locating DCs. This is because manufacturing is often outsourced to third-party suppliers and contracts are volume based, and production, and inventory decisions are best made simultaneously (Fisher and Rajaram 2000).[3] Finally, in contrast to all these papers, we consider an important problem faced by retail chains carrying private-label products, propose an effective methodology to generate feasible solutions for this problem, test it on realistic data to assess its performance, and develop insights that practitioners can use for choosing suppliers and making production, distribution, and inventory decisions.

The article is organized as follows: In section 2, we present the basic model formulation; in section 3, we discuss the corresponding Lagrangian relaxation; while in section 4, we propose heuristics. In section 5, we present results from our numerical study. In section 6, we summarize and provide future research directions.

2. Model Formulation

We formulate the RPP under uncertainty as a nonlinear mixed-integer program. To provide a precise statement of this problem, we define:

Indices:

I,J,K: The set of possible suppliers, demand zones, and products, respectively.
i,j,k: The subscripts for suppliers, demand zones, and products, respectively.

Parameters:

f_i: Fixed annualized cost associated with choosing supplier i.
d_{ik}: Setup cost associated with producing product k at supplier i.
e_{ij}: Setup cost associated with shipping from supplier i to demand zone j.
c_{ijk}: Marginal cost to produce and ship product k from supplier i to demand zone j.
L_i, U_i: Minimum acceptable throughput and capacity of supplier i, respectively.
α_{ijk}: Units of capacity consumed by a unit of product k at supplier i that is shipped to demand zone j.
h_{jk}/p_{jk}: Per unit overstock/understock cost associated with satisfying demand for product k at demand zone j.
$\Phi_{jk}(\xi)/\phi_{jk}(\xi)$: The cumulative/probability density function of the demand distribution for product k at demand zone j.

Decision variables:

z_i: 0–1 variable that equals 1 if supplier i is chosen to supply products and 0 otherwise.
w_{ik}: 0–1 variable that equals 1 if product k is produced in supplier i and 0 otherwise.
v_{ij}: 0–1 variable that equals 1 if supplier i ships to demand zone j and 0 otherwise.
x_{ijk}: Quantity of product k shipped from supplier i to demand zone j.
y_{jk}: Inventory level of product k carried at demand zone j.

To capture economies of scale so that per-unit production and shipping costs decrease in quantity, we approximate these costs by a setup cost d_{ik} that is incurred to initiate production for each product k at every supplier i, a setup cost e_{ij} that is incurred to ship from each supplier i to every demand zone j, and a constant marginal cost (c_{ijk}) that is incurred to produce and distribute each additional unit. While a more complex cost structure could be desirable in some applications, we employ this structure as it captures economies of scale and it permits structural analysis of the problem.

To model the inventory problem faced by the retailer, we employ the newsvendor model. In contrast to Daskin et al. (2002) and Shen et al. (2003), who use a (Q,r) replenishment model, this study is motivated by the fashion retail industry, where merchandise is often seasonal and lead times are long relative to the season length. Consequently, the retailer cannot replenish inventory mid-season, so unmet demand is lost, while leftover demand needs to be salvaged via mark-downs at the end of the season. Therefore, the standard single-period newsvendor model would seem most appropriate here. Under this model, let $S_{jk}(y)$ denote the expected overstock and understock cost associated with carrying y units of inventory for product k at demand zone j. This can be written as

$$\begin{aligned} S_{jk}(y) =& h_{jk}\int_0^y (y-\xi)\phi_{jk}(\xi)d\xi \\ &+ p_{jk}\int_y^\infty (\xi - y)\phi_{jk}(\xi)d\xi \\ \Longrightarrow S_{jk}(y) =& (h_{jk}+p_{jk})\int_0^y \Phi_{jk}(\xi)d\xi + p_{jk}[E(\xi)-y]. \end{aligned} \tag{1}$$

The problem of supplier selection, production, distribution, and inventory planning faced by the retailer can be expressed by the following nonlinear mixed-integer program, which we call the RPP:

$$\text{(RPP)}$$

$$\begin{aligned} Z_P = \min\Bigg\{ & \sum_{i\in I} f_i z_i + \sum_{i\in I}\sum_{k\in K} d_{ik}w_{ik} + \sum_{i\in I}\sum_{j\in J} e_{ij}v_{ij} \\ & + \sum_{i\in I}\sum_{j\in J}\sum_{k\in K} c_{ijk}x_{ijk} + \sum_{j\in J}\sum_{k\in K} S_{jk}(y_{jk}) \Bigg\} \end{aligned}$$

subject to

$$\sum_{i\in I} x_{ijk} = y_{jk}\, \forall j\in J, k\in K \tag{2}$$

$$L_i z_i \le \sum_j \sum_k \alpha_{ijk}x_{ijk} \le U_i z_i\, \forall i\in I \tag{3}$$

$$\sum_{j\in J} \alpha_{ijk}x_{ijk} \le U_i w_{ik}\, \forall i\in I, k\in K \tag{4}$$

$$\sum_{k\in K} \alpha_{ijk}x_{ijk} \le U_i v_{ij}\, \forall i\in I, j\in J \tag{5}$$

$$x_{ijk}\ge 0, y_{jk}\ge 0\, \forall i\in I, j\in J, k\in K \tag{6}$$

$$w_{ik}\in\{0,1\}, v_{ij}\in\{0,1\}, z_i\in\{0,1\}\, \forall i\in I, j\in J, k\in K. \tag{7}$$

The objective function of the RPP consists of four terms. The first term represents the annualized fixed cost associated with securing capacity at supplier i. The second term represents the setup cost associated with production, while the third term represents the setup cost associated with distribution. The fourth term represents the corresponding (constant) marginal production and distribution costs. The fifth term represents the total expected cost associated with carrying inventory at the demand zones.

Constraint (2) ensures that total inventory level for each product at every demand zone equals the total quantity produced and shipped to that zone. Note that it is also a coupling constraint. Were it not for Equation (2), the RPP would decompose by supplier i into a set of mixed-integer linear problems, and by demand zone j and product k into a set of newsvendor problems. This observation suggests that this may be a good candidate constraint to use in any eventual decomposition of the problem. The left-hand side inequality of Equation (3) imposes a lower bound on the minimum allowable throughput of a supplier, if the supplier is selected. A lower bound on a supplier's throughput may be desirable to achieve sufficient economies of scale. The right-hand side inequality of Equation (3) imposes the capacity constraint (i.e., U_i) for each supplier that is selected, and it enforces that no production will take place with suppliers that are not selected. Constraint (4) enforces the condition that $x_{ijk} > 0$ if and only if product k is produced at supplier i (i.e., iff $w_{ik} = 1$ for some $j \in J$), while Equation (5) enforces the condition that $x_{ijk} > 0$ if and only if some quantity is shipped from supplier i to demand zone j (i.e., iff $v_{ij} = 1$ for some $k \in K$). Finally, Equation (6) are nonnegativity constraints, while Equation (7) are binary constraints.

Observe that the RPP is a convex mixed-integer program, as it consists of a linear generalized facility location subproblem and a convex inventory planning subproblem. By noting that the capacitated plant location problem (CPLP) is strongly NP-hard (Cornuejols et al. 1991), it can be shown that the RPP is also strongly NP-hard.[4] Therefore, it is unlikely that real-sized problems can be solved to optimality. We verify this in our computational results. Consequently, it is desirable to develop heuristics to address this problem. The quality of these heuristics can be assessed by comparing them to a lower bound, which we establish in the next section.

3. Decomposition and Lower Bounds

To obtain a tight lower bound, we apply a Lagrangian relaxation to the RPP (see Fisher 1981, Geoffrion 1974). An important issue when designing a Lagrangian relaxation is deciding which constraints to relax. In making this choice, it is important to strike a suitable compromise between solving the relaxed problem efficiently and yielding a relatively tight bound. Observe that by relaxing (2), the problem can be decomposed into a mixed-integer linear program (MILP) containing the x_{ijk}, w_{ik}, v_{ij}, and z_i variables, and into a convex program containing the y_{jk} variables. Moreover, this relaxation enables us to further decompose the MILP by supplier (i.e., by i), and the convex program by demand zone and product (i.e., by j and k) into multiple subproblems. A key attribute of this decomposition is that all subproblems can be solved analytically. On the other hand, a potential concern is that this decomposition generates a relatively large number of dual multipliers: $J \times K$ of them, which we denote by λ_{jk}. Relaxing (2) for a given $J \times K-$ matrix λ of multipliers, the Lagrangian function takes the following form:

$$L(\lambda) = \min \sum_{i \in I} \left[f_i z_i + \sum_{k \in K} \left(d_{ik} w_{ik} + \sum_{j \in J} (c_{ijk} - \lambda_{jk}) x_{ijk} \right) + \sum_{j \in J} e_{ij} v_{ij} \right] + \sum_{j \in J} \sum_{k \in K} [\lambda_{jk} y_{jk} + S_{jk}(y_{jk})] \tag{8}$$

subject to Equations (3)–(7).

Note that (L_λ) decomposes by i into I independent production and distribution subproblems, and by j and k into $J \times K$ independent inventory subproblems. More specifically, Equation (8) can be rewritten as follows:

$$L(\lambda) = \sum_{i \in I} L_i^{milp}(\lambda) + \sum_{j \in J} \sum_{k \in K} L_{jk}^{cvx}(\lambda)$$

where

$$L_i^{milp}(\lambda) = \min \left\{ f_i z_i + \sum_{k \in K} \left[d_{ik} w_{ik} + \sum_{j \in J} (c_{ijk} - \lambda_{jk}) x_{ijk} \right] + \sum_{j \in J} e_{ij} v_{ij} \right\}$$

and

$$L_{jk}^{cvx}(\lambda) = \min\{\lambda_{jk} y_{jk} + S_{jk}(y_{jk})\}.$$

Note that the Lagrangian multipliers in the production and distribution subproblems (i.e., $L_i^{milp}(\lambda)$) can be interpreted as the cost saved (or cost incurred if $\lambda_{jk} < 0$) from producing and distributing an additional unit of product k to demand zone j. On the other hand, the Lagrangian multipliers in the inventory subproblems (i.e., $L_{jk}^{cvx}(\lambda)$) can be interpreted as the change in holding cost associated with carrying an additional unit of inventory of product k at demand zone j.

For any given set of multipliers λ, Proposition 1 determines the optimal solution for (L_λ), thus providing a lower bound for the RPP.

Proposition 1. *For given set of multipliers $\lambda \in \mathbb{R}^{J \times K}$, a lower bound for the RPP is given by*

$$L(\lambda) = \sum_{i \in I} \min \left\{ \min_{j \in J} \left\{ e_{ij} + \min_{k \in K} \left\{ d_{ik} + (c_{ijk} - \lambda_{jk}) \frac{U_i}{\alpha_{ijk}} \right\} \right\} + f_i, 0 \right\} + \sum_{j \in J} \sum_{k \in K} \left[p_{jk} \mathbb{E}_{jk}(\xi) - (p_{jk} + h_{jk}) \int_0^{y_{jk}(\lambda_{jk})} \xi \phi_{jk}(\xi) d\xi \right], \tag{9}$$

where

$$y_{jk}(\lambda) = \begin{cases} \Phi_{jk}^{-1}(1) & \text{if } \lambda_{jk} \le -h_{jk} \\ \Phi_{jk}^{-1}\left(\frac{p_{jk} - \lambda_{jk}}{p_{jk} + h_{jk}}\right) & \text{if } -h_{jk} \le \lambda_{jk} \le p_{jk} \\ \Phi_{jk}^{-1}(0) & \text{if } \lambda_{jk} \ge p_{jk}. \end{cases} \tag{10}$$

Proof. To begin, fix $\lambda \in \mathbb{R}^{J \times K}$. Let us first consider each production and distribution subproblem. To solve each subproblem, we apply the integer linearization principle by Geoffrion (1974). First, observe that if $z_i = 0$, then $L_i^{milp} = 0$. Hence, the optimal solution must satisfy $L_i^{milp}(\lambda) \le 0$. As a result, we fix $z_i = 1$ and solve

$$L_i^{milp}(\lambda, z_i = 1) \triangleq \min \sum_{k \in K} \left[d_{ik} w_{ik} + \sum_{j \in J} (c_{ijk} - \lambda_{jk}) x_{ijk} \right] + \sum_{j \in J} e_{ij} v_{ij} + f_i$$

subject to Equations (3)–(7).

Because the problem is linear, using Equation (3) it can easily be shown that $x_{ijk}(\lambda) \in \{0, \frac{U_i}{\alpha_{ijk}}\}$. Using that $L_i^{milp}(\lambda) \le 0$, Equations (4), (5), and (7), it follows that

$$L_i^{milp}(\lambda) = \min \left\{ \min_{j \in J} \left\{ e_{ij} + \min_{k \in K} \left\{ d_{ik} + (c_{ijk} - \lambda_{jk}) \frac{U_i}{\alpha_{ijk}} \right\} \right\} + f_i, 0 \right\}.$$

Next, consider each inventory subproblem. It is easy to show that this problem is convex in y_{jk} and by solving the first-order condition with respect to y_{jk}, we obtain Equation (10), where $\Phi_{jk}^{-1}(\bullet)$ denotes the inverse of $\Phi_{jk}(\bullet)$. Finally, by using Equation (1) and $y_{jk}(\lambda_{jk})$, it is easy to show that for each $j \in J$ and $k \in K$, $L_{jk}^{cvx}(\lambda_{jk})$ can be written as

$$L_{jk}^{cvx}(\lambda_{jk}) = p_{jk} \mathbb{E}_{jk}(\xi) - (p_{jk} + h_{jk}) \int_0^{y_{jk}(\lambda_{jk})} \xi \phi_{jk}(\xi) d\xi.$$

By noting that a lower bound can be obtained by $L(\lambda) = \sum_{i \in I} L_i^{milp}(\lambda) + \sum_{j \in J} \sum_{k \in K} L_{jk}^{cvx}(\lambda)$, the proof is complete. □

Note that the Lagrangian solution will chose a supplier (i.e., set $z_i(\lambda) = 1$) if and only if the cost savings associated with producing and distributing an additional unit of product k to demand zone j exceed the fixed cost associated with choosing this supplier for at least some j and k (i.e., if and only if $-\min_{j \in J} \left\{ e_{ij} + \min_{k \in K} \left\{ d_{ik} + (c_{ijk} - \lambda_{jk}) \frac{U_i}{\alpha_{ijk}} \right\} \right\} \ge f_i$.

However, this solution may not be feasible. Thus, the purpose of this solution is more to establish the value of the objective function of (L_λ), which is a lower bound on the value of the optimal solution of the RPP. This lower bound can then be used to evaluate the quality of any feasible solution generated by heuristics for this problem. In the unlikely event that the corresponding solution is feasible for the original problem, it then solves the RPP optimally.

In Lemma 1, we show that the Lagrangian problem (L_λ) does not possess the integrality property (see Geoffrion 1974). Therefore, the Lagrangian bound is likely to be strictly better than that of a convex programming relaxation (i.e., the relaxation that is obtained by replacing the binary constraints in Equation (7) by the continuous interval [0,1] for the RPP). We confirm this in our computational results in section 5.

LEMMA 1. *The Lagrangian problem (L_λ) does not possess the integrality property.*

PROOF. It suffices to show that a convex programming relaxation of the RPP where Equation (7) is replaced by

$$0 \le w_{ik} \le 1, 0 \le v_{ij} \le 1, \text{ and } 0 \le z_i \le 1\, \forall i \in I, j \in J, k \in K$$

does not yield a solution such that the w, v, and z variables are integral. We prove this by constructing a counterexample as follows: Let $|J| = |K| = 1$, $e_{i1} = d_{i1} = L_i = 0\, \forall i \in I, \alpha_{i11} = 1\, \forall i \in I$, and $\Phi_{11}(\xi) = \xi$. To simplify exposition, in the remainder of this proof we drop the subscripts j and k. Observe that by cost minimization, $\forall i$ we will have that $z_i = \frac{x_i}{U_i}$. As a result, it suffices to show that there exists an instance of the convex programming relaxation of the RPP with optimal solution $x_i^* \notin \{0, U_i\}$ for some $i \in I$ (and hence $z_i^* \notin \{0, 1\}$. To proceed, by noting that Slater's condition is satisfied for the primal problem, we dualize (2) and write the Lagrangian

$$L(v) = \min_{0 \le x_i \le U_i}\left\{\sum_{i\in I}\left(c_i + \frac{f_i}{U_i} + v\right)x_i + (h+p)\int_0^y \xi d\xi + \frac{p}{2} - (v+p)y\right\}.$$

It is straightforward to check that for any given dual multiplier v, the Lagrangian program assumes the following optimal solution:

$$x_i(v) = \begin{cases} U_i & \text{if } c_i + \frac{f_i}{U_i} + v < 0 \\ \in [0, U_i] & \text{if } c_i + \frac{f_i}{U_i} + v = 0, \text{ and } y(v) = \frac{v+p}{h+p} \\ 0 & \text{otherwise}. \end{cases}$$

Observe that a solution of the form $x_i \in \{0, U_i\}$ will be optimal (and hence $z_i \in \{0, 1\}$) if and only if there exists a dual multiplier v such that

$$\sum_{i\in I} U_i \mathbf{1}_{\left\{c_i + \frac{f_i}{U_i} + v \le 0\right\}} = \frac{v+p}{h+p}.$$

By noting that the RHS is a smooth function strictly increasing in v while the LHS is a step function decreasing in v, it follows that there may exist at most one v such that the above equality is satisfied. We now construct an example in which there exists no v such that the above equality is satisfied. Letting $h = p = 1$, $|I| = 2$, $U_i = \frac{i}{2}$, and $c_i + \frac{f_i}{U_i} = \frac{i}{3}$, observe that if

$$\begin{array}{rl} -\frac{1}{3} < v & (\text{LHS}) = 0 < \frac{v+1}{2} = (\text{RHS}) \\ -1 < v \le -\frac{1}{3} \text{ then} & (\text{LHS}) = \frac{1}{2} > \frac{v+1}{2} = (\text{RHS}) \\ v < -1 & (\text{LHS}) = \frac{3}{2} > \frac{v+1}{2} = (\text{RHS}). \end{array}$$

We have thus constructed an instance for which the convex programming relaxation does not yield an optimal solution that is integral and hence proven that the Integrality Property does not hold. □

We next consider the problem of choosing the matrix of Lagrangian multipliers λ to tighten the bound $L(\lambda)$ as much as possible. Specifically, we are interested in the tightest possible lower bound, which can be obtained by solving

$$LB_{LR} = \max_{\lambda \in \mathbb{R}^{J\times K}} L(\lambda).$$

One way to maximize $L(\lambda)$ is by using a traditional subgradient algorithm (see Fisher 1985 for details). However this technique may be computationally intensive in our problem, as we have $J \times K$ Lagrangian multipliers.

To overcome this difficulty, we exploit the structure of the dual problem to demonstrate how the optimal set of Lagrangian multipliers λ can in some cases be fully or partially determined analytically. In preparation, we establish Lemma 2.

LEMMA 2. *The optimal set of Lagrangian multipliers $\lambda^* \in J \times K$ satisfy*

$$\min\left\{\min_{i\in I}\left\{c_{ijk} + \frac{\alpha_{ijk}}{U_i}(d_{ik} + e_{ij} + f_i)\right\}, p_{jk}\right\} \le \lambda^*_{jk} \le p_{jk}$$

$\forall j \in J$ and $k \in K$.

PROOF. First, it is easy to check from the first line of Equation (9) that $L^{milp}(\lambda)$ decreases in λ, and $L^{milp}(\lambda) = 0$ if $d_{ik} + e_{ij} + (c_{ijk} - \lambda_{jk})\frac{U_i}{\alpha_{ijk}} + f_i \ge 0\, \forall i, j, k$. By rearranging terms, one can show that $L^{milp}(\lambda) = 0$ if $\lambda_{jk} \le c_{ijk} + \frac{\alpha_{ijk}}{U_i}(d_{ik} + e_{ij} + f_i)\quad \forall i, j, k$. It is also easy to verify from the second line of (8) that $L_{jk}^{cvx}(\lambda_{jk})$ increases in λ_{jk}, and $L_{jk}^{cvx}(\lambda_{jk}) = p_{jk}\mathbb{E}_{jk}(\xi)$ if $\lambda_{jk} \geq p_{jk} \forall j,k$.

To show that $\min\left\{\min_{i\in I}\left\{c_{ijk}+\frac{\alpha_{ijk}}{U_i}(d_{ik}+e_{ij}+f_i)\right\},\right.$ $\left.p_{jk}\right\} \leq \lambda^*_{jk} \leq p_{jk}$, first suppose that the LHS inequality is not satisfied for some j,k. Then, $L_i^{milp}(\lambda^*) = L_i^{milp}(\hat{\lambda})$ and $L_{jk}^{cvx}(\lambda^*_{jk}) \leq L_{jk}^{cvx}(\hat{\lambda}_{jk})$, where $\hat{\lambda} = \max\left\{\lambda^*, \min\left\{\min_{i\in I}\left\{c_{ijk} + \frac{\alpha_{ijk}}{U_i}(d_{ik} + e_{ij} + f_i)\right\}, p_{jk}\right\}\right\}$. As a result, $L(\lambda^*) \leq L(\hat{\lambda})$ and hence λ^* cannot be optimal. Now suppose that $\lambda^*_{jk} > p_{jk}$ for some j, k. Then $L_{jk}^{cvx}(\lambda^*_{jk}) = L_{jk}^{cvx}(p_{jk})$ and $L_i^{milp}(\lambda^*) \leq L_i^{milp}(\lambda^*_{jk})$, where λ^*_{jk} denotes the set of Lagrangian multipliers λ^*, in which the $j-k^{th}$ element has been replaced by p_{jk}. As a result $L(\lambda^*) \leq L(\lambda^*_{jk})$, and hence λ^* cannot be optimal. This completes the proof. $\square$

Lemma 2 states that the optimal set of Lagrangian multipliers λ^* lies in a well-defined compact set. Observe from the left-hand side expression in Lemma 2 that $\lambda^*_{jk} > 0 \forall j,k$. From Equation (10), observe that the optimal inventory level $y_{jk}(\lambda^*_{jk})$ is strictly smaller than the optimal inventory level that would be determined from solving the inventory subproblem separately from the supplier choice and production planning subproblem. This is a direct consequence of performing production, distribution, and inventory planning in an integrated manner. The second implication of Lemma 2 is that the optimal solution of the Lagrangian relaxation will always satisfy $\sum_{i\in I} x_{ijk}(\lambda^*) \geq y_{jk}(\lambda^*)$, and if the set defined in Lemma 2 is a singleton for some j and k, then it is possible to partially characterize the optimal set of Lagrangian multipliers *ex ante*. When these sets are singletons for all j and k, then it is possible to completely characterize λ^* *ex ante*. This is established by Proposition 2.

Proposition 2. *If* $\min_{i\in I}\left\{c_{ijk} + \frac{\alpha_{ijk}}{U_i}(d_{ik} + e_{ij} + f_i)\right\} \geq p_{jk}$, *then the optimal Lagrangian multiplier* $\lambda^*_{jk} = p_{jk}$. *If this inequality holds* $\forall j \in J$ *and* $k \in K$, *then* $\lambda^* = p$ *and* $Z_P = LB_{LR}$ *(i.e., the Lagrangian relaxation solves the RPP).*

Proof. For any j and k, if $\min_{i\in I}\left\{c_{ijk} + \frac{\alpha_{ijk}}{U_i}(d_{ik} + e_{ij} + f_i)\right\} \geq p_{jk}$, then by Lemma 2 $\lambda^*_{jk} = p_{jk}$. If this condition holds for all j and k, then it follows that $\lambda^*_{jk} = p_{jk}$, and by substituting $\lambda^*_{jk} = p_{jk}$ into Equation (8), it is easy to check that (L_λ) is feasible for RPP. This completes the proof. $\square$

Observe that $\min_{i\in I}\left\{c_{ijk} + \frac{\alpha_{ijk}}{U_i}(d_{ik} + e_{ij} + f_i)\right\}$ can be interpreted as the lowest marginal cost associated with establishing capacity at some supplier, producing product k, and distributing it to demand zone j. As a result, when this marginal cost exceeds the marginal underage cost, it is optimal not to produce any quantity of product k for demand zone j and incur the expected underage cost; that is, set $\lambda_{jk} = p_{jk}$, which yields $y_{jk}(p_{jk}) = 0$ by applying (10).

By using Lemma 2 and Proposition 2 we now establish a worst-case error bound for the Lagrangian relaxation studied in this section.

Corollary 1. *The worst-case error bound for this Lagrangian relaxation satisfies*

$$\epsilon_{LR} \geq 1 + \max\left\{\frac{-\sum_{j,k}(p_{jk}+h_{jk})\int_0^{y(\lambda^1_{jk})}\xi\phi_{jk}(\xi)d\xi}{\sum_{j,k}p_{jk}\mathbb{E}_{jk}(\xi)}, \frac{\sum_i \min\left\{\min_{j\in J}\left\{e_{ij} + \min_{k\in K}\left\{d_{ik} + (c_{ijk} - p_{jk})\frac{U_i}{\alpha_{ijk}}\right\}\right\}, 0\right\}}{\sum_{j,k}p_{jk}\mathbb{E}_{jk}(\xi)}\right\}$$

where $\epsilon_{LR} = \frac{LB_{LR}}{Z_P}$ *and* $\lambda^1_{jk} = \min\left\{\min_{i\in I}\left\{c_{ijk} + \frac{\alpha_{ijk}}{U_i}\cdot (d_{ik} + e_{ij} + f_i)\right\}, p_{jk}\right\} \forall j$ *and* k. *Moreover, there exists a problem instance of the RPP such that the bound is tight (i.e.,* $\epsilon_{LR} = 1$*).*

Proof. First note that the Lagrangian dual is a concave maximization problem, and recall from Lemma 2 that $\lambda^*_{jk} \geq \min\{\min_{i\in I}\{c_{ijk} + \frac{\alpha_{ijk}}{U_i}(d_{ik} + e_{ij} + f_i)\}, p_{jk}\} = \lambda^1_{jk}$. Moreover, it is easy to check that a trivial feasible solution can be obtained by setting $z_i = w_{ik} = x_{ijk} = y_{jk} = 0 \ \forall i,j,k$, in which case the objective function is equal to $\sum_{j,k} p_{jk}\mathbb{E}_{jk}(\xi)$. As a result, the following inequalities hold:

$$\max\{L(\lambda^1), L(p)\} \leq LB_{LR} \leq Z_P \leq \sum_{j,k} p_{jk}\mathbb{E}_{jk}(\xi).$$

Hence, $\epsilon_{LP} = \frac{LB_{LR}}{Z_P} \geq \frac{\max\{L(\lambda^1),L(p)\}}{\sum_{j,k}p_{jk}\mathbb{E}_{jk}(\xi)}$, and the result follows by substituting $L(\lambda^1)$ and $L(p)$ from Equation (9). To show that there exists an instance such that this bound is tight, for every $i \in I$, pick f_i such that $\min_{j,k}\left\{d_{ik} + e_{ij} + (c_{ijk} - p_{jk})\frac{U_i}{\alpha_{ijk}}\right\} + f_i \geq 0$. Then, it is easy to check that $\epsilon_{LR} \geq 1$. Because $\epsilon_{LR} \leq 1$ by definition, we conclude that $\epsilon_{LR} = 1$ in this instance. This completes the proof. $\square$

4. Heuristics and Upper Bounds

In this section, we develop heuristics, which can be used to obtain feasible solutions for the RPP. These heuristics can be used in conjunction with the lower bound developed in section 3 to provide upper bounds for a branch and bound algorithm or to generate a feasible solution for the RPP. We initially propose two intuitive heuristics. The first is a practitioner's heuristic developed based on observed practice at a large retail chain. The second is a sequential heuristic, which solves the inventory management subproblem first, and then it solves the remaining standard facility location problem by applying the well-known *Drop* procedure (Klincewicz and Luss 1986).

These two heuristics can be used to benchmark the performance of the analytically more rigorous heuristics we develop. The first is a convex programming-based heuristic, which generates a feasible solution by solving a sequence of convex programs. We also propose a simpler LP-based heuristic, which is computationally more efficient. This heuristic uses the inventory levels from the Lagrangian problem (i.e., $y(\lambda^*)$), and it generates a feasible solution by solving a sequence of linear programs. We next present these heuristics, and we evaluate their performance in section 5.

4.1. Practitioner's Heuristic

This heuristic first chooses the inventory level for every product at each demand zone to equal the respective expected demand; that is, $y_{jk} = \mu_{jk}$ $\forall j \in J, k \in K$. Second, suppliers are sorted according to the ratio $R_i = \frac{f_i}{U_i}$, which captures the fixed cost per-unit of capacity associated with choosing supplier i. Third, the algorithm establishes sufficient capacity to satisfy the total inventory by choosing suppliers that have the lowest R_i. For example, if $R_1 \le R_2 \ldots \le R_I$, then the algorithm will set $z_i = 1$ $\forall i \in \{1, \ldots, n\}$ and $z_i = 0$ otherwise, where $n = \min\left\{n \le I : \sum_{i=1}^{n} U_i \ge \sum_{j \in J} \sum_{k \in K} y_{jk}\right\}$. Finally, production and transportation decisions are made by solving a relaxed version of the RPP, where the fixed cost variables w_{ik} and v_{ij} are relaxed to lie in [0,1]. Here, a feasible solution is obtained by rounding to 1 the fractional w_{ik} and v_{ij} variables and by re-solving the linear program with respect to $x_{ijk} \ge 0$. Note that this heuristic does not take into account the underage and overage costs due to the variation in demand as inventory levels are set to simply equal the mean demand. We denote the objective function of this heuristic by UB_{Pr}. This procedure is formalized in Algorithm 1.

A more sophisticated version of this heuristic can be obtained by choosing the inventory levels according to the newsvendor model, and then using the same approach as described in Algorithm 1 to choose suppliers and conduct logistics planning. We call this the newsvendor-based practitioner's heuristic, and we denote its objective function by UB_{Pr-NV}.

Algorithm 1. Practitioner's Heuristic

1: Let $R_i = \frac{f_i}{U_i}$, and sort candidate facilities such that $R_1 \le R_2 \ldots \le R_I$.
2: Fix $y_{jk} = \mu_{jk}$ $\forall j$ and k.
3: Let $n = \min\left\{n \le I : \sum_{i=1}^{n} U_i \ge \sum_{j \in J} \sum_{k \in K} y_{jk}\right\}$.
4: Fix $z_i = 1$ $\forall i = 1, \ldots, n$ and $z_i = 0$ otherwise.
5: Solve the RPP with relaxed variables $v_{ij}, w_{ik} \in [0, 1]$.
6: Fix to 1 any $v_{ij} > 0$ and $w_{ik} > 0$, re-solve LP, and compute objective function UB_{Pr}.

4.2. Sequential Heuristic

This heuristic obtains a feasible solution for the RPP in two stages: in the first stage, it fixes the inventory level for each product at every demand zone by solving $J \times K$ newsvendor problems. This reduces the problem to a standard capacitated facility location problem with piece-wise linear costs. Then, in the second stage, it uses a *Drop* heuristic–a well-known construction heuristic for facility location problems to determine which suppliers to choose. The general idea of the Drop heuristic is to start with a solution in which all candidate suppliers are chosen (i.e., $z_i = 1\,\forall i$), iteratively deselect one supplier at a time, and solve the remaining subproblem in which the fixed cost variables w_{ik} and v_{ij} are relaxed to lie in [0,1]. Then, any fractional w_{ik} and v_{ij} variables are rounded to 1, and the problem is resolved with respect to the x_{ijk} variables. In each loop, the heuristic permanently deselects the supplier who provides the greatest reduction in total expected costs, and it terminates if no further cost reduction is possible. Since exactly one z_i is dropped in each loop, and at least one supplier must be selected in any feasible solution, the algorithm needs at most $I(I-1)$ iterations in total, and two convex programs are solved in each iteration. We denote the objective function of this heuristic by UB_{Seq}. This procedure is formalized in Algorithm 2.

For completeness, we also consider a variant of the sequential heuristic that fixes the inventory level for each product at every demand zone to equal the respective expected demand. We call this the *simplified* sequential heuristic, and we denote its objective function by $UB_{Seq-Simple}$.

Algorithm 2. Sequential Heuristic

1: Fix $y_{jk} = y(\text{newsvendor})$ $\forall j$ and k
2: Fix $z_i = 1$ $\forall i$ and $UB_{Seq} = +\infty$.
3: **for** $n = 1$ **to/do**
4:   **for** $m = 1$ **to/do**
5:     **if** $z_m = 1$ **do**
6:       Fix $z_i^m = z_i$ $\forall i \ne m$ and $z_m^m = 0$.
7:       Solve the RPP with z_i^m and relaxed variables $v_{ij}, w_{ik} \in [0, 1]$.
8:       Fix to 1 any $v_{ij} > 0$ and $w_{ik} > 0$, and re-solve RPP to find x_{ijk} variables.
9:       Compute objective function UB_{Seq}^m.
10:     **end if**
11:   **end for**
12:   **if** $\min_m UB_{Seq}^m < UB_{Seq}$ **do**
13:     $UB_{Seq} = \min_m UB_{Seq}^m$ and $z_{m^*} = 0$, where $m^* = \arg\min_m UB_{Seq}^m$.
14:     *terminate*
15:   **end if**
16: **end for**

4.3 Convex Programming-Based Heuristic

One disadvantage of the practitioner's and the sequential heuristics is that inventory decisions are made independent of supplier selection and logistics decisions. Moreover, the Drop approach used in the sequential heuristic can be computationally intensive. Therefore, we construct a convex programming-based heuristic as an alternative way to obtain a feasible solution for the RPP.

The heuristic begins by solving a relaxed RPP where the fixed cost variables z_i, w_{ik}, and v_{ij} have been relaxed to lie in [0,1]. First, it temporarily fixes the largest fractional z_i to 1, solves the remaining (relaxed) problem, and rounds to 1 any fractional w_{ik} and v_{ij} variables. Second, it temporarily fixes the smallest fractional z_i to 0, and again it solves the remaining (relaxed) problem and rounds to 1 any fractional w_{ik} and v_{ij} variables. The algorithm then permanently fixes the z_i that yielded the lowest total expected costs, and it continues to iterate until all z_i variables have been fixed to 0 or 1. The assumption behind this approach is that the fractional value of z_i is a good indicator of the "worthiness" of choosing supplier i. Since at least one z_i is fixed in each loop, the algorithm needs at most i iterations in total, and two convex programs are solved in each iteration. We denote the objective function of this heuristic by UB_{Cvx}. This procedure is formalized in Algorithm 3.

To gauge the value of joint logistics and inventory planning, we also consider a simplified version of the convex programming heuristic in which inventory levels are selected in advance using the solution corresponding to the lower bound from the Lagrangian relaxation (i.e., $y_{jk}(\lambda^*)$ $\forall j$ and k). Then, the problem of finding a feasible solution reduces to solving a sequence of linear programs, which are easier to solve than convex programs. We denote the objective function associated with this LP-based heuristic by UB_{Lp}.

Algorithm 3. Convex Programming-Based Heuristic

Initiate $z_i^{min} = 0$ and $z_i^{max} = 1\ \forall i$
while $z_i^{max} > z_i^{min}$ for some i **do**
Solve the RPP with relaxed variables $v_{ij}, w_{ik} \in [0,1]$ and $z_i^{min} \le z_i \le z_i^{max}$
  if $z_i \in \{0,1\}$ **do**
    Set $z_i^{min} = z_i^{max} = z_i$
  end if
  Let $i_{max} = \arg\max\{z_i : z_i \in (0,1)\}$ and $i_{min} = \arg\min\{z_i : z_i \in (0,1)\}$.
  Solve the RPP with relaxed variables $v_{ij}^+, w_{ik}^+ \in [0,1]$, $z_i^{min} \le z_i^+ \le z_i^{max}$, and $z_{i_{max}}^+ = 1$.
  Fix to 1 any $v_{ij}^+ > 0$ and $w_{ik}^+ > 0$, and compute objective function UB_{CVX}^+.
  Solve the RPP with relaxed variables $v_{ij}^-, w_{ik}^- \in [0,1]$, $z_i^{min} \le z_i^- \le z_i^{max}$, and $z_{i_{min}}^- = 0$.
  Fix to 1 any $v_{ij}^- > 0$ and $w_{ik}^- > 0$, and compute objective function UB_{CVX}^-.
  if $Z^+ > Z^-$ do
    $z_{i_{max}}^{min} = 1$
  else
    $z_{i_{min}}^{max} = 0$
  end if
end while
Fix to 1 any $v_{ij} > 0$ and $w_{ik} > 0$, re-solve the convex program, and compute UB_{Cvx}.

5. Computational Results

In this section, we present a computational study to evaluate the performance of the heuristics. In addition, we investigate the key factors that drive their performance and also examine their robustness. In addition, we investigate the key factors that drive their performance and we examine their robustness.

To test our methods across a broad range of data, we randomly generated the parameter values using a realistic set of data made available to us by a large retailer. We generated 500 random problem instances, each comprising between 5 and 20 candidate suppliers, 10 and 40 demand zones, and 1 and 25 products (i.e., $I \sim U\{5,\ldots,20\}$, $J \sim U\{10,\ldots,40\}$ and $K \sim U\{1,\ldots,25\}$). The parameters we used in our computational study are summarized in Table 1. To solve the optimization problems associated with the bounding techniques we propose, we used the CVX solver for Matlab (CVX 2011) running on a computer with an Intel Core i7-2670QM 2.2GHz processor and 6 GB of RAM memory.

To evaluate the performance of the Lagrangian lower bound, we benchmark it against a standard convex programming relaxation, in which the integrality constraints are relaxed so that Equation (7) is replaced by

$$0 \le w_{ik} \le 1, \quad 0 \le v_{ij} \le 1, \quad 0 \le z_i \le 1\ \forall i \in I, j \in J, k \in K.$$

In every one of the problem instances tested, the Lagrangian relaxation generated a better lower bound than the convex programming relaxation, on average by 2.34%. This is consistent with Lemma 1, which asserts that the Lagrangian problem L_λ does not possess the Integrality Property.

To test the performance of the heuristics developed in section 4, we evaluate the suboptimality gaps relative to the Lagrangian lower bound. Let $\Delta_x = 100\% \frac{UB_x - LB_{LR}}{LB_{LR}}$, where $x \in \{Cvx, Lp, Seq, Seq-Simple, Pr, Pr-NV\}$ denote the convex programming based, the LP-based, the sequential, the simplified sequential, the practitioner's, and the newsvendor-based practitioner's heuristics, respectively. The average and median values, as well as the range of these metrics, are illustrated in Figure 2.

Table 1. Summary of Parameters Used in Our Computational Study

Parameters	Distribution of values	Parameters	Distribution of values
Fixed cost of choosing a supplier	$\bar{f} \sim N(50, 10)$ $f_i \sim N(\frac{2JK}{3I}\bar{f}, \frac{3JK}{2I}\bar{f})$	Overstock cost	$\bar{h} \sim N(5, 1)$ $h_{jk} \sim U[\frac{2}{3}\bar{h}, \frac{3}{2}\bar{h}]$
Setup cost associated with production	$\bar{d} \sim N(200, 40)$ $d_{ik} \sim U[0, \bar{d}]$	Understock cost	$\bar{p} \sim N(50, 10)$ $p_{jk} \sim U[\frac{2}{3}\bar{p}, \frac{3}{2}\bar{p}]$
Setup cost associated with distribution	$\bar{e} \sim N(200, 40)$ $e_{ij} \sim U[0, \bar{e}]$	Mean demand	$\bar{\mu} \sim N(20, 4)$ $\mu_{jk} \sim U[\frac{2}{3}\bar{\mu}, \frac{3}{2}\bar{\mu}]$
Marginal production and distribution cost	$\bar{c} \sim N(10, 2)$ $c_{ijk} \sim U[\frac{2}{3}\bar{c}, \frac{3}{2}\bar{c}]$	Demand variance	$\bar{\sigma} \sim N(5, 1)$ $\sigma_{jk} \sim U[\frac{2}{3}\bar{\sigma}, \frac{3}{2}\bar{\sigma}]$
Supplier capacity	$\bar{U} \sim N(100, 20)$ $U_i \sim N(\frac{40JK}{I}\bar{U}, \frac{90JK}{I}\bar{U})$	Weights min. throughput	$\alpha_{ijk} = 1$ $L_i = 0$

In addition, to get an idea of the computational complexity of these heuristics, Table 2 reports the mean, median, and maximum computational times for the problem instances tested.

First, observe from Figure 2 that the convex programming-based heuristic unambiguously outperforms the other heuristics. In particular, it provides feasible solutions that are on average within 3.44% of optimal, and range from 0.41% to 18.76%. While the gap of the LP-based heuristic is higher than the convex programming-based heuristic on average, in the majority of cases it generates a feasible solution that is quite close to optimal as evidenced by the median gap of 4.32%. The practitioner's heuristics generate feasible solutions that are on average 19.95% and 36.32% from optimal for the standard and the newsvendor-based versions, respectively. On the other hand, the suboptimality gap for the sequential heuristics is on average 36.72% and 23.7% for the standard and the simplified versions, respectively.

Interestingly, with both the practitioner's and the sequential heuristics, the version in which inventory levels are set equal to the mean demand outperforms the version in which inventory levels are chosen according to the newsvendor solution. This is because the understock costs are generally larger than the overstock costs, and, hence, the newsvendor model leads to a larger stocking quantity than the average demand. This in turn increases production and distribution costs as well as the fixed costs associated with establishing capacity in excess of the benefit of reducing underage costs.

In addition, the inventory levels corresponding to the solution of the convex programming-based heuristic are always lower than those determined by the newsvendor solution, and they are often lower than those chosen by the LP-based heuristic. This is because the convex programming heuristic solves the joint problem in contrast to the LP-based heuristic as

Figure 2 Suboptimality Gap

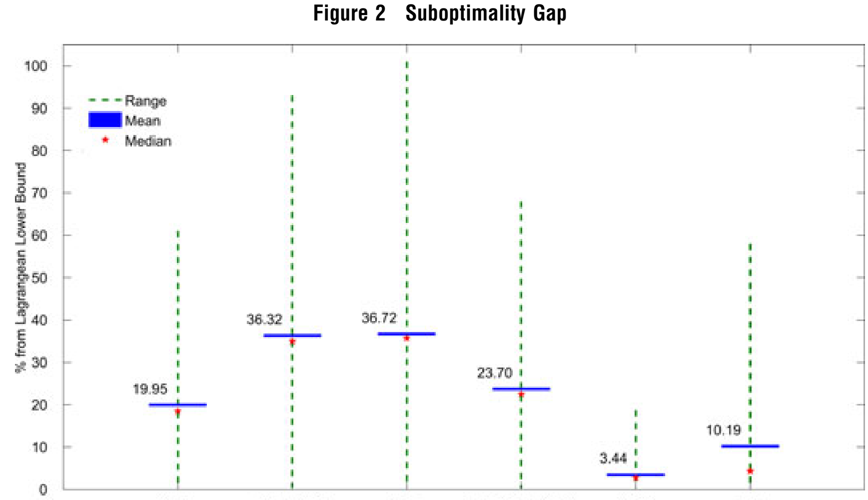

Table 2. Computational Times (sec)

	UB_{Pr}	UB_{Pr-NV}	UB_{Seq}	$UB_{Seq-Simple}$	UB_{Cvx}	UB_{Lp}
Mean	2.02	2.03	80.85	187.41	175.76	105.04
Median	1.68	1.72	59.48	91.00	96.78	52.48
Max	9.44	10.23	463.98	1974.96	1472.18	949.64

well as other heuristics in which the inventory levels are chosen separately from the joint problem. When one solves for the joint problem, the solution accounts for the fact that a larger downstream inventory level raises production quantities, which increases upstream production and distribution costs as well as the costs associated with establishing production capacity. In contrast, these costs are not considered when the inventory subproblem is solved separately, and hence result in a larger inventory level. The takeaway from this is that when planning the entire supply chain, it is important to consider the effect of downstream inventory decisions to the upstream production and distribution costs. Retailers often underestimate the impact of upstream costs in their urge to have a higher market share associated with higher fill rates. When such costs are adequately represented, a lower fill rate may actually be preferable to lower total costs. Finally, note that the cost reduction resulting from the convex programming-based heuristic relative to the other heuristics is important, because retailers operate in a highly competitive environment with very low margins and even a small cost reduction can lead to a large profit increase.

Next, consider the computation time for each heuristic. Observe that both practitioner's heuristics are computationally very fast, while both sequential heuristics are quite slow. Also note that the standard sequential heuristic is computationally less intensive than its simplified counterpart. Because the standard sequential heuristic chooses the stocking quantities according to the newsvendor model, which in general are higher than the expected demand, the Drop procedure needs fewer iterations in the standard sequential heuristic. Finally, observe that the convex programming-based heuristic is about as computationally intensive as the simplified sequential heuristic, but leads to much lower average gaps. Thus, it clearly dominates both versions of the sequential heuristic. However, as expected, it is computationally more intensive than the LP-based heuristic.

As the convex programming heuristic dominates the other heuristics in terms of the gap from the lower bound, we focus on this heuristic to examine (i) how the computational time scales up with the size of the problem, (ii) how the suboptimality gap and its performance advantage relative to the practitioner's heuristic depend on the parameters of the problem, and (iii) which parameters have the greatest impact on the total expected costs.

To conduct this analysis, we regress the computational times, the suboptimality gap of the convex programming heuristic (i.e., Δ_{Cvx}), the gap between the convex programming and the practitioner's heuristic (i.e., $100\% \frac{UB_{Cvx} - UB_{Pr}}{UB_{Pr}}$), and the total expected cost associated with the convex programming heuristic (i.e., UB_{Cvx}) of the 500 problem instances tested earlier on the size (i.e., I,J,K), and the parameters of the problem (i.e., $\bar{\mu}, \bar{\sigma}, \bar{h}, \bar{p}, \bar{c}, \bar{d}, \bar{e}, \bar{f}, \bar{U}$). Table 3 summarizes the results.

First note that the computational time of the convex programming heuristic is strongly dependent on the problem size (i.e., I, J, and K), while it is insensitive to the other parameters of the problem. More interestingly, the relatively large R^2 ratio implies that the computational time of the convex programming heuristic is explained by a linear model well, which in turn suggests that the computational time scales up approximately linearly in the problem size.

From the second column, observe that the suboptimality gap decreases in the size of the problem (I, j, and K), and this effect is significant at the 1% level. This finding is encouraging: it predicts that the convex programming heuristic will perform even better in larger problem instances that could be expected in some applications. The suboptimality gap increases in the capacity of the candidate suppliers ($\bar{U}$), while it decreases in the mean demand (μ) and the fixed costs associated with choosing a supplier ($\bar{f}$). The suboptimality gap also increases in the demand variance ($\bar{\sigma}$), the underage and overage costs ($\bar{p}$ and $\bar{h}$), and the production costs ($\bar{c}$, $\bar{d}$, and $\bar{e}$), but this effect is not significant at the 10% level. Finally, note that a 95% confidence interval for each regression coefficient can be obtained from (regression coeff.) ± 1.9648(SE).[5] Therefore, as seen in Table 3, because the values of all regression coefficients and their respective standard errors are close to zero, |(regression coeff.) ± 1.9648(SE)| is close to zero for all parameters. This shows that the performance of the convex programming heuristic is robust to changes in the parameters of the RPP.

The third column examines how the performance advantage of the convex programming heuristic relative to the practitioner's heuristic depends on the parameters of the problem. Observe that the performance advantage of the convex programming heuristic becomes larger in the size of the problem, while it is insensitive to the cost parameters, as evidenced by the small regression coefficients and the small (respective) standard errors. This, together with the finding that the value of the intercept is negative at the 1% significance level, reinforces the benefits from using the convex programming heuristic, as one could expect even larger problems with different cost parameters in certain applications.

The fourth column considers the relationship between the total expected cost of the feasible solutions generated by the convex programming heuristic and the parameters of the problem. Predictably, the expected cost increases in the size of the problem (I, J, and K), in the mean demand ($\bar{\mu}$), in the production costs ($\bar{c}$, $\bar{d}$, and $\bar{e}$), in the fixed costs associated with

Table 3. Suboptimality Gap vs. Problem Parameters (Convex Programming Heuristic)

	Computational time	Suboptimality gap	Cvx. vs. Pract. H.	Expected cost
I	23.93*** (1.301)	−0.13*** (0.034)	0.32*** (0.089)	5247.77*** (1094.73)
J	7.66*** (0.651)	−0.09*** (0.016)	0.06 (0.042)	9798.92*** (517.59)
K	18.33*** (0.803)	−0.16*** (0.026)	0.15** (0.067)	19,590.85*** (822.44)
$\bar{\mu}$	−0.35 (0.893)	−0.12*** (0.017)	0.25*** (0.044)	11,042.93*** (543.84)
$\bar{\sigma}$	0.04 (3.732)	0.101 (0.071)	−0.12 (0.185)	48.77 (2278.22)
$\bar{h}$	7.79 (5.706)	0.12 (0.109)	0.12 (0.287)	7702.49** (3521.77)
$\bar{p}$	0.023 (0.548)	0.012 (0.0104)	0.25*** (0.027)	2048.43*** (334.45)
$\bar{c}$	6.13 (2.69)	0.055 (0.051)	−0.038 (0.134)	8629.04*** (1647.61)
$\bar{d}$	0.02 (1.108)	0.026 (0.0204)	−0.053 (0.054)	261.74 (657.6)
$\bar{e}$	−1.59 (1.099)	0.011 (0.021)	−0.007 (0.055)	1269.01* (672.77)
$\bar{f}$	−1.40** (0.551)	−0.00003*** (0.000003)	−0.00008*** (0.00001)	2.29*** (0.12)
$\bar{U}$	1.38*** (0.266)	0.0015*** (0.0001)	0.002*** (0.0003)	−37.82*** (3.35)
Intercept	−689.2*** (88.8)	8.52*** (1.469)	−36.62*** (3.851)	−830, 954.6*** (47,333.18)
R^2	0.641	0.40	0.315	0.891

The values before the parentheses denote the regression coefficients corresponding to the parameter in the left column. The values in parentheses denote standard errors.
*Significance at 10% level.
**Significance at 5% level.
***Significance at 1% level.

choosing a supplier ($\bar{f}$), as well as in the underage and overage costs ($\bar{p}$ and $\bar{h}$). On the other hand, the expected cost decreases in the capacity of the candidate suppliers ($\bar{U}$), while the effect of the demand variance ($\bar{\sigma}$) is insignificant. Therefore, our findings suggest that besides the problem size (i.e., I, J, K, and $\bar{\mu}$), the two most important factors that affect the expected cost of a feasible solution are (i) the marginal production cost and (ii) the inventory underage and overage costs. The latter observation emphasizes the value of an improved demand forecast. On the other hand, the capacity of a supplier and the fixed contracting costs appear to have a secondary effect. This is consistent with the initiatives undertaken at several retailers to reduce the impact of production, inventory underage, and overage costs (Fisher and Raman 2010).

As the gaps of the convex programming-based heuristic are the smallest, we analyze the solutions to develop some insights about how it chooses suppliers. This could be useful for practitioners who make such decisions. We find that suppliers are chosen in increasing order of the ratio r_i, where

$$r_i = \frac{1}{U_i}\left[f_i + \frac{1}{|J||K|}\sum_{j,k}\left(d_{ik} + v_{ij} + c_{ijk}\frac{\alpha_{ijk}}{U_i}\right)\right].$$

The term in brackets represents the sum of fixed establishment costs and the average production and distribution costs across products and demand zones when a supplier is fully utilized. Therefore, the ratio r_i can be interpreted as the average total cost per unit of capacity at supplier i. This suggests that it is important to consider establishment, production and distribution costs together when choosing suppliers, and it is beneficial to choose suppliers with the lowest total average cost per unit of capacity.

6. Conclusions

We analyze a multi-product RPP under demand uncertainty in which the retailer jointly chooses suppliers, plans production and distribution, and selects inventory levels to minimize total expected costs. This problem typically arises in retail store chains carrying private-label products, who need to plan the entire supply chain by making decisions with respect to (i) supplier selection for their private-label products, (ii) distribution of products from suppliers to demand zones (i.e., stores or DCs), and (iii) the inventory levels for every product at each demand zone. This problem is formulated as a mixed-integer convex program.

As the RPP is strongly NP-hard, we use a Lagrangian relaxation to obtain a lower bound, and we develop heuristics to generate feasible solutions. First, we develop an analytic solution for the Lagrangian problem (Proposition 1), and we establish conditions under which the Lagrangian dual can be solved analytically (see Proposition 2). We first develop a practitioner's and a sequential heuristic. We then propose two heuristics, which reduce the problem of generating a feasible solution to solving a sequence of convex or linear programs. To test the performance and the robustness of our methods, we conduct an extensive computational study. The convex programming-based heuristic and its LP-based counterpart yield feasible solutions that are on average within 3.4% and 10.2% from optimal, respectively. Sensitivity analysis suggests that the computational time of the convex programming heuristic scales up approximately linearly in the problem size, while it is stable to changes in problem parameters. Finally, these heuristics outperform both the sequential and the practitioner's heuristics, and the performance advantage of the convex programming-based heuristic relative to the practitioner's heuristic is robust to the parameters of the problem. All these are desirable features for any eventual implementation in large-sized real applications.

Several managerial insights can be drawn from this work. First, solving the more complicated joint supplier choice, production, distribution, and inventory problem leads to a leaner supply chain with lower inventory levels than solving the inventory subproblem separately from the supplier choice and logistics subproblem. This highlights the importance of considering the effect of inventory decisions on upstream production and distribution costs. Our methodology provides an effective approach to solve this joint problem. Second, the major costs that influence supply chain costs across the retailer are production costs, as well as the understock and overstock costs associated with carrying inventory at the demand zones. Therefore, retailers should focus on reducing these costs first before considering the effects of supplier capacity and contracting costs. Third, it is important to consider establishment, production, distribution, and inventory costs together when choosing suppliers, because a supplier who is desirable in any one of these aspects may in fact *not* be the best overall choice. Our analysis provides a mechanism to integrate these aspects and pick the best set of suppliers.

This study opens up several opportunities for future research. First, this problem could be extended to explicitly model nonlinear production and shipping costs, which is of particular interest for applications that exhibit significant economies of scale. In that case, the problem formulation is a mixed-integer nonlinear program that is neither convex nor concave (see Caro et al. 2012 for details about addressing a related problem in the process industry with uncertain yields). Second, our model could be extended to incorporate multiple echelons in the supply chain (i.e., wholesalers and DCs) and allow multiple echelons to carry inventory. Third, it may be desirable to incorporate side constraints pertaining to facilities, production, and distribution (i.e., v, x, w, and z variables) as in Geoffrion and McBride (1978). Undoubtedly, all these extensions would require significant, nontrivial modifications to our model. Finally, further work could be done to improve the heuristics to further reduce the suboptimality gap.

In conclusion, we believe the methods described in this article provide an effective methodology to address the RPP under demand uncertainty.

Acknowledgments

We thank two anonymous referees, the associate editor, and the department editor Vishal Gaur for suggestions that improved this article. We also thank Professors Reza Ahmadi, Arthur Geoffrion, and Christopher Tang for useful comments.

Notes

[1]With the latter interpretation we implicitly assume (i) that the locations of DCs and the assignment of stores to DCs are predetermined and (ii) that stores maintain only a minimal amount of inventory so that inventory costs at individual stores are negligible. This latter assumption is consistent with the existing literature (e.g., Shen et al. 2003). While it is plausible that management must also determine the location of DCs and allocate stores to DCs, we leave this important problem for future research.

[2]Specific lead times faced by manufacturers are reported to be 7 months for Oxford shirts ordered by J. C. Penney, and 5 months for Benetton apparel (Iyer and Bergen 1997).

[3]For example, leading retailers H&M and GAP outsource 100% of their manufacturing, while Zara outsources approximately 40% of its manufacturing to the third-party suppliers (Tokatli 2008). Anecdotal evidence suggests that Macy's outsources all of its manufacturing.

[4]This result can be shown by reducing an instance of the RPP to the CPLP. Specifically, in this reduction, let (i) demand assume a degenerate probability distribution, (ii) the overage and underage costs to be arbitrarily large (i.e., h_{jk} and $p_{jk} \to \infty \, \forall j,k$), (iii) $d_{ik} = 0$ and $e_{ij} = 0 \; \forall i,j,k$, (iv) $L_i = 0 \; \forall i \in I$, and (v) U_i's take values from the set $\{1,\ldots, p\}$ for any fixed $p \geq 3 \; \forall i \in I$.

[5]± 1.9648 corresponds to the 2.5 and 97.5 percentile of a t-distribution with $500 - 13$ degrees of freedom.

References

Aikens, C. H. 1985. Facility location models for distribution planning. *Eur. J. Oper. Res.* **22**: 263–279.

Balachandran, V., S. Jain. 1976. Optimal facility location under random demand with general cost structure. *Nav. Res. Logistics Q.* **23**: 421–436.

Caro, F., K. Rajaram, J. Wollenweber. 2012. Process location and product distribution with uncertain yields. *Oper. Res.* Forthcoming.

Cornuejols, G., R. Sridharan, J. M. Thizy. 1991. A comparison of heuristics and relaxations for the capacitated plant location problem. *Eur. J. Oper. Res.* **50**: 280–297.

Daskin, M. S., C. R. Coullard, Z. J. M. Shen. 2002. An inventory location model: Formulation, solution algorithm and computational results. *Ann. Oper. Res.* **110**: 83–106.

Drezner, Z. 1995. *Facility Location: A Survey of Applications and Methods*. Springer, New York.

Fisher, M. L. 1981. The Lagrangian relaxation method for solving integer programming problems. *Manage. Sci.* **27**(1): 1–18.

Fisher, M. L. 1985. An applications oriented guide to Lagrangian relaxation. *Interfaces* **15**(2): 10–21.

Fisher, M. L., K. Rajaram. 2000. Accurate retail testing of fashion merchandise. *Mark. Sci.* **19**(3): 266–278.

Fisher, M. L., A. Raman. 2010. *The New Science of Retailing*. Harvard Business Press, Cambridge, MA.

Geoffrion, A. M. 1974. Lagrangian relaxation for integer programming. *Math. Programming Study* **2**: 82–114.

Geoffrion, A. M., R. McBride. 1978. Lagrangian relaxation applied to the capacitated facility location problem. *AIIE Trans.* **10**: 40–47.

Grant, M., S. Boyd. 2011. CVX: Matlab Software for Disciplined Convex Programming, version 1.21. http://www.cvxr.com/cvx (accessed date November 16, 2012).

Iyer, A., M. Bergen. 1997. Quick response in manufacturer-retailer channels. *Manage. Sci.* **43**(4): 559–570.

Klincewicz, J. G., H. Luss. 1986. A Lagrangian relaxation heuristic for capacitated facility location with single-source constraints. *J. Oper. Res. Soc.* **37**(5): 495–500.

Krarup, J., P. M. Pruzan. 1989. Ingredients of locational analysis. P. B. Mirchandani, R. L. Francis, eds. *Discrete Location Theory*. John Wiley & Sons, New York, 1–54.

Le Blanc, L. J. 1977. A heuristic approach for large-scale discrete stochastic transportation-location problems. *Comput. Math. Appl.* **3**: 87–94.

Melo, M. T., S. Nickel, F. Saldanha-da-Gama. 2009. Facility location and supply chain management—A review. *Eur. J. Oper. Res.* **1962**: 401–412.

Oszen, L., R. C. Collette, M. S. Daskin. 2008. Capacitated warehouse location model with risk pooling. *Nav. Res. Logistics* **55**(4): 295–312.

Owen, S. H., M. S. Daskin. 1998. Stratergic facility location: A review. *Eur. J. Oper. Res.* **111**: 423–447.

Shen, Z.-J. M. 2005. A multi-commodity supply chain design problem. *IIE Trans.* **37**: 753–762.

Shen, Z.-J. M. 2007. Integrated supply chain models: A survey and future research direction. *J. Ind. Manage. Optimiz.* **3**(1): 1–27.

Shen, Z.-J. M., C. Coullard, M. S. Daskin. 2003. A joint location-inventory model. *Transp. Sci.* **37**(1): 40–55.

Snyder, L. V. 2006. Facility location under uncertainty: A review. *IIE Trans.* **38**: 537–554.

Tokatli, N. 2008. Global sourcing: Insights from the global clothing industry—The case of Zara, a fast fashion retailer. *J. Econ. Geogr.* **8**(1): 21–38.