Integrated Stochastic Supply-Chain Design Models

A supply chain is a network of facilities and activities that procure, produce, and distribute goods to customers. Many uncertainties within a supply chain can substantially affect its performance. The author discusses how to model these uncertainties in integrated supply-chain design models and studies their impacts on optimal supply-chain decisions.

Supply-chain management is a set of approaches that help efficiently integrate suppliers, manufacturers, warehouses, and retailers so that merchandise is produced and distributed at the right quantities, to the right locations, and at the right time. This minimizes system-wide costs (or maximizes profits) while satisfying service-level requirements.[1]

A supply chain's physical structure clearly impacts its performance, so major supply-chain decisions can vary widely: Which suppliers should we use? How many factories and warehouses should we have, and where should we locate them? How do we set the capacity at each facility? What products should each factory produce? Given locations and capacities, supply-chain decisions will try to answer other questions: What quantities should we produce and store at these locations? What quantities should we move from location to location, and at what time? (See the "Supply-Chain Decision Making" sidebar for a discussion on how previous research has attempted to answer these questions.)

1521-9615/07/$25.00 © 2007 IEEE
Copublished by the IEEE CS and the AIP

Zuo-Jun (Max) Shen
University of California, Berkeley

Most of the early supply-chain models are deterministic, assuming that decision makers know all variables of interest before implementing solutions. However, many decision parameters, such as demands and costs, can change dramatically from decision to implementation time. This calls for supply-chain design models that address the inherent uncertainties in facility location problems. Designing an efficient supply chain can provide a company with a tremendous competitive advantage in the marketplace.

This article focuses on integrating decisions at strategic, tactical, and operational levels. (See the "Treating Uncertainties in the Supply Chain" sidebar for a full discussion of related work in this area.) Specifically, I study strategic location decisions while taking into account the impact of inventory and shipment decisions. I also explicitly model stochastic demand and random supply in the supply chain. (For a more detailed review of joint optimization of decisions from different levels, see my related research.[2] Sunil Chopra and Peter Meindl[3] also discuss network design models that consider both demand and financial uncertainty.)

Basic Model Formulation

For this work, I considered a three-tiered supply-chain system consisting of one or more suppliers, distribution centers (DCs), and retailers. Each retailer has uncertain demand. The problem is de-

Supply-Chain Decision Making

We can roughly classify the necessary decisions in the supply chain into strategic, tactical, and operational levels. In the design phase, strategic decision such as facility location decisions play major roles. Once we've determined the supply chain's configuration, our focus shifts to tactical and operational decisions (such as inventory management decisions on raw materials, intermediate products, and end products) and production scheduling and product distribution decisions within the supply chain.

Researchers have typically treated decisions at different levels separately. The strategic location theory literature tends to focus on developing models for determining the number of distribution centers (DCs) and their locations as well as DC-retailer assignments. These decisions are evaluated based on resulting operational shipping costs and strategic location costs. With some notable exceptions, this work tends to ignore demand uncertainty. Furthermore, location models typically ignored inventory costs and their impact on location and shipping costs until recently.[1]

The inventory literature tends to ignore the strategic location decision and its associated costs. One reason for such a disconnection is that the decision maker doesn't possess detailed information at the nonstrategic level in the strategic design phase, thus facility location decisions are usually made without many inputs regarding inventory and distribution costs. However, failure to take the related inventory and shipment costs into consideration when determining the facility location can lead to suboptimal results.[2]

Thus, to achieve important cost savings, the supply chain should be optimized as a whole—that is, the major cost factors that can impact the supply chain's performance should be considered jointly in the decision model. This isn't only true for decisions at the same level (for instance, it's well known that the inventory-management scheme and the transportation strategy should be integrated), but it also applies to decisions at different levels.

References

1. Z.-J. Shen, C.R. Coullard, and M.S. Daskin, "A Joint Location-Inventory Model," *Transportation Science*, vol. 37, 2003, pp. 40–55.
2. Z.-J. Shen and L. Qi, "Incorporating Inventory and Routing Costs in Strategic Location Models," to be published in *European J. Operational Research*, vol. 179, no. 2, 2007.

termining how many DCs are necessary, where to locate them, which retailers to assign to each DC, how often to reorder at the DC, and what level of safety stock to maintain to minimize total location, shipment, and inventory costs while ensuring a specified level of service. (Safety stock is inventory held in case demand exceeds expectations.)

We can assume that location costs are incurred when DCs are established. Line-haul transportation costs are incurred for shipments from a supplier to the DCs. Local transportation costs are incurred when moving the goods from the DCs to the retailers. Inventory costs are incurred at each DC and consist of the carrying cost for the average inventory used over time as well as safety stock inventory carried to protect against uncertain retailer demand. We assume that the non-DC retailers maintain only a minimal amount of inventory, which I ignore in the model here. (I discuss how to deal with retailer inventory in the next section).

The inputs and parameters are

- I, a set of retailers;
- $\mathcal{J}$, a set of candidate DC locations;
- μ_i, the mean (daily) demand at retailer i for each $i \in I$;
- σ_i^2, the variance of (daily) demand at retailer I for each $i \in I$;
- f_j, the fixed (annual) cost of locating a DC at j for each $j \in \mathcal{J}$;
- d_{ij}, the cost of shipping a unit from DC j to retailer i for each $i \in I$ and $j \in \mathcal{J}$;
- α, the desired percentage of retailer orders satisfied (fill rate);
- γ, the weight factor associated with the routing;
- β, the weight factor associated with the shipment cost;
- θ, the weight factor associated with the inventory cost;
- z_α, the standard normal deviate such that $P(z \le z_\alpha) = \alpha$;
- h, the inventory holding cost per unit of product per year;
- F_j, the fixed administrative and handling cost of placing an order at DC j for each $j \in \mathcal{J}$;
- L, the DC order lead time in days;
- g_j, the fixed shipment cost per shipment from the supplier to DC j;
- $\bar{a}_j$, the cost per unit of a shipment from the supplier to DC j; and
- χ, a constant used to convert daily demand into annual demand (for example, 365 if demand occurs every day of the year).

The decision variables are

Treating Uncertainties in the Supply Chain

Researchers have treated uncertainty in many different ways within the supply-chain literature, but the work typically focuses on the strategic, tactical, and operational levels. The models I review here typically assume full probabilistic characterization, but the research community is increasingly interested in applying learning techniques to problems without full probabilistic characterizations. When the data available for learning is limited, or the underlying uncertainty is nonstationary, these approaches can induce significant error, reducing the effectiveness of the policies derived. However, my colleagues and I discuss how to incorporate these errors in the model and describe different ideas in modeling model uncertainty,[1] finding the solution to this model using robust optimization and its implementation through learning. These models will be extremely helpful for tactical and operational decision making.

Strategic Level

At the strategic level, decisions can have a long-lasting impact on the firm's performance. For instance, facility location decisions, once implemented, won't change that often because of the huge setup costs involved.

With some notable exceptions, the location literature tends to ignore demand uncertainty, but doing so can result in bad facility locations and inefficiency and extra costs, even if the production, inventory, and shipment plans are well optimized. Mark Daskin and Susan Owen[2] provide an overview of facility location modeling as do other recent texts.[3,4] Additional related works[5,6] review facility location models in dynamic and uncertain environments.

Tactical Level

Tactical-level decisions are typically updated every few months, and they typically include production decisions, inventory, and transportation strategies (such as mode selection and frequency of visiting customers).

For inventory management problems, the objective is to decide the frequency and quantities of orders for supplying DCs and filling retailer orders. An inventory policy's performance is evaluated based on the resulting service levels (the percentage of retailer orders filled within the acceptable waiting period), shipping costs, inventory costs, and shortage costs (those incurred when an order can't be filled

- $X_j := 1$, if retailer j is selected as a DC location, but 0 otherwise for each $j \in \mathcal{J}$, and
- $Y_{ij} := 1$, if retailer i is served by a DC based at location j, but 0 otherwise for each $i \in I$ and each $j \in \mathcal{J}$.

To simplify notation, I assume all lead times are equal and the holding cost rates are the same at different DCs. I use the weight factors β, θ to adjust the relative proportion of different cost components.

Working Inventory Cost

For the inventory policy under this system, a DC orders inventory from the supplier using an (r, Q) policy with service-level constraints. The frequency of orders and the order quantity at each DC is determined by the mean demand served by the DC, which in turn is a function of the assignment of retailers to the DC.

Let S_j denote the set of retailers served by j, D_j denote the total annual (expected) demand going through DC

$$j(D_j = \Sigma_{i \in S_j} \mu_i = \Sigma_{i \in I} \mu_i Y_{ij}),$$

and n be the number of shipments per year from the supplier. The average shipment size in one shipment from the supplier to DC j is D_j/n, and the average working inventory cost at DC j is $\theta h D_j/(2n)$. Assuming we can calculate the delivery cost from the supplier to DC j as $g_j + \bar{a}_j D_j / n$, where g_j is the fixed cost of placing an order, then the total annual cost of ordering inventory from the supplier to DC j is given by

$$F_j n + \beta(g_j + \bar{a}_j D_j / n)n + \theta h D_j / (2n). \qquad (1)$$

The optimal value of n that minimizes this function is equal to

$$\sqrt{\theta h D_j / (2(F_j + \beta g_j))}\ .$$

We can express the corresponding total annual working inventory cost associated with DC j as

$$\sqrt{2\theta h D_j (F_j + \beta g_j)} + \beta \bar{a}_j D_j . \qquad (2)$$

Safety Stock Cost

Using Eppen's classical risk-pooling result, the amount of safety stock required to ensure that shortages occur with a probability of α or less is

$$z_\alpha \sqrt{L \Sigma_{i \in S_j} \sigma_i^2}\ .$$

within an acceptable waiting period). This line of research tends to incorporate demand uncertainty.[7,8]

Operational Level

Operational-level decisions deal with detailed short-term operations, including production scheduling and distribution. Although a vast amount of literature deals with production scheduling, most of the scheduling papers deal with deterministic problems. (Mike Pinedo[9] provides an excellent review of the subject.) In terms of product distribution, a heavily studied mathematical model is the vehicle routing problem (VRP), in which a set of customers must be served by a fleet of vehicles of limited capacity. The vehicles are initially located at a given depot, and the objective is to find a set of routes for the vehicles of minimal total length. Each route begins at the depot, visits a subset of the customers, and returns to the depot without violating the capacity constraint. Paolo Toth and Daniele Vigo[10] provide an excellent review of the model. (For reviews on stochastic VRP, see other related articles.[11,12])

References

1. A. Lim, J.G. Shanthikumar, and Z.-J. Shen, "Model Uncertainty, Robust Optimization and Learning," *Tutorials in Operations Research*, H.J. Greenberg, eds., INFORMS, 2006.
2. M.S. Daskin and S.H. Owen, "Location Models in Transportation," *Handbook of Transportation Science*, R. Hall, ed., Kluwer Academic, 1999, pp. 311–360.
3. P.B. Mirchandani and R.L. Francis, *Discrete Location Theory*, John Wiley & Sons, 1990.
4. M.S. Daskin, *Network and Discrete Location: Models, Algorithms, and Applications*, Wiley-Interscience, 1995.
5. S.H. Owen and M.S. Daskin, "Strategic Facility Location: A Review," *European J. Operational Research*, vol. 111, no. 3, 1998, pp. 423–447.
6. L.V. Snyder, "Facility Location under Uncertainty: A Review," *IIE Trans.*, vol. 38, no. 7, 2006, pp. 537–554.
7. E.L. Porteus, *Foundations of Stochastic Inventory Theory*, Stanford Univ. Press, 2002.
8. P.H. Zipkin, *Foundations of Inventory Management*, Irwin, 1997.
9. M. Pinedo, *Scheduling: Theory, Algorithms and Systems*, Prentice Hall, 1995.
10. P. Toth and D. Vigo, eds., *The Vehicle Routing Problem*, SIAM Press, 2002.
11. M. Dror, "Routing with Stochastic Demands: A Survey," *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*, M. Dror, P.L. Ecuyer, and F. Szydarovszky, eds., Kluwer Academic Publishers, 2002, pp. 629–653.
12. M. Dror, "Routing Propane Deliveries," *Logistics Systems: Design and Optimization*, A. Langevin and D. Riopel, eds., Springer, 2005, pp. 299–322.

The corresponding holding cost for the safety stock at DC j is

$$\theta h z_a \sqrt{L \sum_{i \in S_j} \sigma_i^2} = \theta h z_\alpha \sqrt{L \sum_{i \in I} \sigma_i^2 Y_{ij}} \ . \qquad (3)$$

Model Formulation

Using the cost items I described earlier, we can formulate the following supply-chain design model. First, minimize

$$\sum_{j \in \mathcal{J}} \left\{ \begin{array}{l} f_j X_j + \left[\sum_{i \in I} (\beta \mu_i d_{ij} + \beta a_j \mu_i) \chi Y_{ij} \right] + \\ + \sqrt{2\theta h (F_j + \beta g_j} \sqrt{\sum_{i \in I} \mu_i \chi Y_{ij}} \\ + \theta h z_\alpha \sqrt{\sum_{i \in I} \hat{\sigma}_i^2 Y_{ij}} \end{array} \right\}$$

$$= \sum_{j \in I} \left\{ \begin{array}{l} f_j X_j + \left(\sum_{i \in I} \hat{d}_{ij} Y_{ij} \right) \\ + K_j \sqrt{\sum_{i \in I} \mu_i Y_{ij}} + q \sqrt{\sum_{i \in I} \hat{\sigma}_i^2 Y_{ij}} \end{array} \right\}, \qquad (4)$$

which is subject to

$$\sum_{j \in \mathcal{J}} Y_{ij} = 1, \text{ for each } i \in I \qquad (5)$$

$$Y_{ij} - X_j \le 0, \text{ for each } i \in I, j \in \mathcal{J} \qquad (6)$$

$$Y_{ij} \in \{0, 1\}, \text{ for each } i \in I, j \in \mathcal{J} \qquad (7)$$

$$X_j \in \{0, 1\}, \text{ for each } j \in \mathcal{J}, \qquad (8)$$

where

$$\hat{d}_{ij} = \beta \chi \mu_i (d_{ij} + a_j)$$

$$K_j = \sqrt{2\theta h \chi (F_j + \beta g_j)}$$

$$q = \theta h z_\alpha$$

$$\hat{\sigma}_i^2 = L \sigma_i^2 .$$

The first two terms include the fixed cost of locating facilities and the delivery costs from the DCs to the retailers (represented by terms in d_{ij}) as well as the marginal cost of shipping a unit from a supplier to a DC (represented by terms in a_j). The last two terms are related to inventory costs, which are nonlinear in the assignment variables. K_j captures the inventory effects due to the fixed ordering costs at the DC and the fixed transport costs from a supplier to a DC. Finally, q captures the safety stock costs at the DCs.[4] The value of q depends on the desired service level. The model's constraints specify that every retailer must be served from an open DC. This problem is more difficult than the standard uncapacitated facility location problem, which is already a notorious NP-hard problem.

To solve this problem, we can either use the column generation or Lagrangian relaxation approaches. Both approaches utilize a low-order polynomial time algorithm for solving a subproblem of the following form. First, minimize

$$\sum_{i \in I} a_i Z_i + \sqrt{\sum_{i \in I} b_i Z_i} + \sqrt{\sum_{i \in I} c_i Z_i}\,, \tag{9}$$

which is subject to

$$Z_i \in \{0, 1\} \; \forall i \in I\,.$$

For each $j \in I$, define set function g_j on $E_j \equiv I \backslash \{j\}$ as follows: for each $S \subseteq E_j$,

$$g_j(S) \equiv$$

$$a_j + \sum_{i \in S} a_i + \sqrt{b_j + \sum_{i \in S} b_i} + \sqrt{c_j + \sum_{i \in S} c_i}\,. \tag{10}$$

Let y^* be an optimal solution to $\mathcal{P}_j$, with associated optimal objective value ω_j^*. The minimum reduced cost set $R_j^* \subset I$ is then the collection of retailers $i \in I$ with $z_i^* = 1$. If $\omega_j^* + f_j \geq 0$, then R_j^* has nonnegative reduced cost; moreover, we can conclude that there is no set $R \in \mathcal{R}$ having a designated DC j with negative reduced cost. Furthermore, if for each $j \in I$, we find that $\omega_j^* + f_j \geq 0$, then we can conclude that there is no set $R \in \mathcal{R}$ with negative reduced cost.

Given a finite set E, a real-valued function $h(\cdot)$ that's defined on the subsets of E is called submodular if, for every pair $S, T \subseteq E$, we have

$$h(S) + h(T) \geq h(S \cap T) + h(S \cup T).$$

Theorem 1: The $g_j(S)$ that arises from the pricing problem is a submodular function.[5]

The pricing problem in the column generation phase is thus a submodular function minimization problem (one for each problem). The result I've presented implies the pricing subproblem is polynomially solvable, but the algorithms ($O(n^7 \log n)$) still aren't very efficient.

Solving the Pricing Problem

Let's look at a much faster algorithm ($O(n^2 \log n)$) to solve the pricing problem.

Lemma 1: Given a retailer $j \in I$, and associated minimum-reduced-cost set $R_j^* \subset I$, for every $i \in R_j^* \setminus \{j\}$, $a_i < 0$.

Proof: Let $i \in R_j^* \setminus \{j\}$. Because $b_i, c_i > 0$, if $a_i \geq 0$, then for any solution $\bar{z}$ with $\bar{z}_i = 1$, the objective function value is strictly greater than that of the solution obtained from $\bar{y}$ by setting $\bar{z}_i = 0$.

Hence we can restrict our search for R_j^* to retailers in I^-, where $I^- \equiv \{i \in I \backslash \{j\} : a_i < 0\}$. We next identify a nice structural property of the set R_j^* by extending an argument in a related work.[6]

Let $a_S = \sum_{i \in S} a_i$, $b_S = \sum_{i \in S} b_i$, and $c_S = \sum_{i \in S} c_i$. Define a new function

$$h_j(x, y, z) := (a_j + x) + \sqrt{b_j + y} + \sqrt{c_j + z}\,. \tag{11}$$

Note that $h_j(x, y, z)$ is a separable concave function, and

$$\min_{S \subseteq I} - g_j(S) = \min_{S \subseteq I} - h_j(a_S, b_S, c_S). \tag{12}$$

Because the set of ordered pairs $\{(a_S, b_S, c_S) : S \subseteq I^-\}$ is finite, its convex hull, H, is a convex polyhedron. It now follows from Equation 10 that

$$\begin{aligned} \min_{S \subseteq I} - g_j(S) &= \min_{S \subseteq I} - h_j(a_S, b_S, c_S) \\ &\geq \min_{(A,B,C) \in H} h_j(A, B, C). \end{aligned}$$

Because the function $h_j(A, B, C)$ is concave in the variables (A, B, C), the latter minimization problem attains a minimum at an extreme point of H.

Let (A^*, B^*, C^*) be an extreme point of H. Because H is a polyhedron, a linear function f on H exists that attains its unique minimum over H at (A^*, B^*, C^*). Because f is linear, it has a representation $f(A, B, C) = \alpha A + \beta B + \gamma C$ defined by real numbers α, β, and γ. The uniqueness of (A^*, B^*, C^*) as the minimizer of f over H assures that we don't have $\alpha = \beta = \gamma = 0$.

Because H is the convex hull of $\{(a_S, b_S, c_S) : S \subseteq I^-\}$,

$$\begin{aligned} \alpha A^* + \beta B^* + \gamma C^* &= \min_{(A,B,C) \in H} \alpha A + \beta B + \gamma C \\ &= \min_{S \subseteq I} - \alpha a_S + \beta b_S + \gamma c_S \\ &= \min_{S \subseteq I} - \textstyle\sum_{i \in S} (\alpha a_i + \beta b_i + \gamma c_i). \end{aligned}$$

The set $S^* = \{i : \alpha a_i + \beta b_i + \gamma c_i < 0\}$ is clearly optimal for the last optimization problem. Hence, we conclude from the uniqueness of (A^*, B^*, C^*) as the minimizer of f over H that $(A^*, B^*, C^*) = (a_{S^*}, b_{S^*}, c_{S^*})$. Furthermore, $R_j^* = S^*$.

Note that $S^* = \{i : \alpha a_i + \beta b_i + \gamma c_i < 0\} = \{i : \alpha + \beta(b_i/a_i) + \gamma(c_i/a_i)\} > 0\} = \{i : \beta x_i + \gamma y_i < \alpha\}$, where $x_i = -b_i/a_i$ and $y_i = -c_i/a_i$. Note also that $x_i, y_i \geq 0$ for all i.

Although there are infinitely many choices for the parameters α, β, and γ, it turns out that the number of distinct partitions obtained by varying the parameters is limited. This follows from a general result in the theory of VC dimension. To describe this result, I need to first introduce some notation.

The VC dimension is defined for any set system $S \subset 2^X$ on an arbitrary set X. It's the supremum of the sizes of all shattered subsets $\mathcal{A} \subset X$; here, $\mathcal{A}$ is shattered if $S|_{\mathcal{A}} = 2^{\mathcal{A}}$—that is, for any $\mathcal{B} \subset \mathcal{A}$, a set $S \in S$ exists such that $\mathcal{B} = \mathcal{A} \cap S$. So if $\mathcal{H}$ denotes the system of all closed half-planes in the plane, it isn't difficult to check that the VC dimension of the set system $\mathcal{H}$ is three, because no four points in the plane can be shattered by using only half-planes.

The following well-known result shows that the number of possible candidates for S^* are essentially small.

Lemma 2:[7,8] For any set system S of VC dimension at most d, we have $|S|_X| \leq \Phi_d(|X|)$, where

$$\Phi_d(m) = \binom{m}{0} + \binom{m}{1} + \ldots + \binom{m}{d}.$$

This lemma suggests that we need to search among at most $O(n^3)$ possible subsets to determine S^*.

In fact, we can enumerate the candidate solution more efficiently. We first observe that at the optimal set S^*, the parameters α, β, and γ satisfy the additional properties:

$$\alpha = 1,$$

$$1/(2\sqrt{b_j + \Sigma_{i \in S} b_i}) \leq \beta \leq 1/(2\sqrt{b_j}), \text{ and}$$

$$1/(2\sqrt{c_j + \Sigma_{i \in S} c_i}) \leq \gamma \leq 1/(2\sqrt{c_j}).$$

These equations follow from gradient conditions at the optimal solution because the concave objective function is of the form $h(x, y, z) = x + \sqrt{y} + \sqrt{z}$ and $0 \leq z_i \leq 1$ for all i.

We know $S^* = \{i : \alpha a_i + \beta b_i + \gamma c_i < 0\} = \{i : \alpha + \beta(b_i/a_i) + \gamma(c_i/a_i) > 0\} = \{i : \beta x_i + \gamma y_i < \alpha\}$ for some choice of α, β, and γ. Furthermore, the additional properties let us restrict our search to finding β and γ such that $S^* = \{i : \beta x_i + \gamma y_i < 1, \beta > 0, \gamma > 0\}$. The possible choices for β and γ now lie in the positive orthant. Furthermore, the inequality $\beta x_i + \gamma y_i < 1$ denotes a half-space in this region.

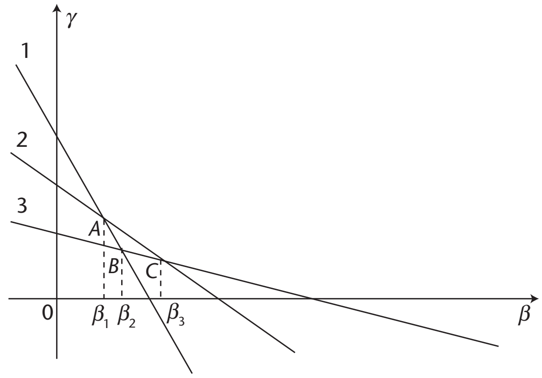

Figure 1. Illustration of the dual algorithm. The diagram efficiently enumerates the candidate solutions.

For each pair of i, j, we solve the equation

$$\begin{cases} \beta x_i + \gamma y_i = 1 \\ \beta x_j + \gamma y_j = 1. \end{cases}$$

This gives rise to a solution $(\beta_{ij}, \gamma_{ij})$. We can discard the solution if any of the β_{ij}, γ_{ij} is nonnegative.

Figure 1 uses a three retailer example to illustrate how to get the optimal set S^*. We first sort all the intersection points according to the value of the β coordinates. For ease of exposition, I relabeled the points $(\beta_{ij}, \gamma_{ij})$ as (β_k, γ_k) so that $\beta_k \leq \beta_{k+1}$ for all k, $k = 1, 2, \ldots, m$, and $m \leq n^2$.

When $\beta \in [0, \beta_1)$, the changes of possible candidates for S^* as γ varies follow an obvious pattern. In Figure 1, the possible candidates are {1, 2, 3}, {1, 2}, and {1} (as γ increases). Similarly, the possible candidates for S^* are {2, 1, 3}, {2, 1}, and {2} when $\beta \in [\beta_1, \beta_2)$; the possible candidates for S^* are {2, 3, 1}, {2, 3}, and {2} when $\beta \in [\beta_2, \beta_3)$; and the possible candidates for S^* when $\beta \in [\beta_3, \infty)$ are {3, 2} and {3}, respectively.

More formally, I can describe the algorithm as follows:

0. Given points (β_k, γ_k), $k = 1, 2, \ldots, m$, with $\beta_1 \leq \beta_2 \leq \ldots \leq \beta_m$.
1. For each k in 0, 1, 2, …, m (define $\beta_0 = 0$),
   a. For each line i, let $\Gamma_i \equiv (1 - \beta_k x_i)/y_i$ —that is, when the β value is set at β_k.
   b. Sort the lines in nondecreasing value of Γ_i. Let $k_1 \leq \ldots \leq k_n$ denote the ordering of the lines when the β value is set at β_k.
   c. The candidate solutions are $\{k_j, k_{j+1}, \ldots, k_n\}$, for each j in 1, 2, …, n, provided $\Gamma_{k_j} \geq 0$.

Because $m \le n^2$, we can execute the sorting procedure in step 0 in $O(n^2 \log n)$ time. To execute step 1 more efficiently, we use the observation that we can obtain the ordering obtained for $k = b + 1$ in step 1b from the ordering for $k = b$ by an inversion (that is, by inverting the order of two adjacent points in the ordering provided by $k = b$). This is true because the β_j coordinates are already sorted in nondecreasing order. Hence, there's no need to explicitly execute steps 1a and 1b for $k > 0$. In this way, we can execute step 1 in $O(n^2)$ time. The bulk of the computational complexity thus lies in step 0, which takes $O(n^2 \log(n))$.

(More details on the model and solution algorithm appear elsewhere.[9–11])

Model with Unreliable Supply

Except for random demands from customers, other uncertain factors exist in a supply-chain network, such as supplier yields and delivery reliability. Ignoring these uncertain factors when designing supply chains can result in inefficient systems. In a recent work, my student and I discussed a supply-chain design model that considers uncertain factors other than demand.[12]

Specifically, we considered the following multiperiod problem: in each period, multiple retailers order a specific product from a supplier, and the supplier ships that product to intermediate facilities selected from a set of candidate locations. Distributors can perform some assembly and packaging activities to satisfy orders from different retailers. Due to a certain service requirement, some amount of final product inventory must be kept in these facilities and be ready for delivery to retailers at the beginning of each period. (We assume that more than one facility can serve any retailer.) However, the amount of final product delivered on time to a retailer might not equal the amount this retailer requests from the supplier because of the quality issues resulting from different production and assembly capabilities in different facilities, mistakes made during assembly and packaging operations, the weather, or other factors that can impact on-time delivery from facilities to retailers. Decision makers must consider all these unreliabilities when designing a supply chain.

We modeled the amount of goods delivered from a facility to a retailer by the product of this retailer's order quantity and a random variable associated with this facility, which is called the *reliability coefficient*. This method is prevalent in the random yield literature.[12,13]

Model Formulation

In our previous work, we also considered facility location, working inventory, and safety stock costs at facilities as well as the penalty and transportation costs associated with retailers.[12] Given the retail price of the product at each retailer, our model's objective is to maximize the expected annual profit. We define the following additional notation for this problem:

- c, the purchasing price from the supplier per unit of product;
- R_j, the reliability coefficient associated with facility $j, j \in \mathcal{J}$, which is a random variable between 0 and 1 (let $\theta_j = E(R_j)$ and $\tau_j^2 = Var(R_j)$);
- p_i, the retail price at retailer i per unit of product, $i \in I$;
- π_i, the penalty cost of lost good will at retailer i per unit of product, $i \in I$; and
- v_i, the salvage value at retailer i per unit of product, $i \in I$.

The decision variables include

- X_j, $\begin{cases} 1 & \text{facility } j \in \mathcal{J} \text{ is open} \\ 0 & \text{otherwise} \end{cases}$

  and
- Q_{ij}, the order quantity at facility $j \in \mathcal{J}$ from retailer $i \in I$ in each period. (We assume that the order quantity from retailer i to facility j is the same in each period.)

Let $\boldsymbol{X}$ denote the $1 \times m$ matrix $(X_j, j \in \mathcal{J})$ and $\boldsymbol{Q}$ denote the $n \times m$ matrix $(Q_{ij}, i \in I, j \in \mathcal{J})$. $R_j Q_{ij}$ represents the actual quantity retailer i receives from facility j in each period if retailer i orders Q_{ij} from facility j, and the reliability coefficient associated with facility j is R_j.

We assume that open facilities only deliver products to retailers at the beginning of every period and that each retailer acts like a newsboy in the newsboy problem and maintains only a minimal amount of inventory. (In the newsboy problem, which is a famous stochastic inventory control problem, we can determine the optimal order quantity if only one order can be placed before actual demand is realized, given a known stochastic distribution for a product's demand.) We can therefore ignore the holding cost of the retailers' inventory in the integrated model.

We formulate retailer i's $(i \in I)$ inventory problem as a classic newsboy problem.[14]

Maximize

$$T_i(Q) \equiv E\left\{ p_i D_i + v_i[\sum_{j\in\mathcal{J}} R_j Q_{ij} - D_i]^+ - (p_i+\pi_i)[D_i - \sum_{j\in\mathcal{J}} R_j Q_{ij}]^+ - \sum_{j\in\mathcal{J}} d_{ij} R_j Q_{ij} \right\}$$

subject to $Q_{ij} \geq 0 \qquad j \in \mathcal{J}$.

Integrated Model

We assume that the per-unit purchase and transportation costs are based on the quantity ordered, not the quantity actually received by the retailers.[12] Based on this assumption, we formulate problem P.

Maximize

$$-\sum_{j\in\mathcal{J}} \left\{ \begin{array}{l} f_j X_j \\ +c\chi \sum_{i\in I} Q_{ij} + \left(a_j \chi \sum_{i\in I} Q_{ij} + K_j \sqrt{\sum_{i\in I} Q_{ij}} \right) \end{array} \right\} + \chi \sum_{i\in I} T_i(Q) = \phi(Q) - \sum_{j\in\mathcal{J}} f_j X_j$$

subject to $1 - e^{-\beta Q_{ij}} \leq X_j \; i \in I, j \in \mathcal{J}$

$Q_{ij} \geq 0 \qquad i \in I, j \in \mathcal{J}$

$X_j \in \{0, 1\} \qquad j \in \mathcal{J}$,

where

$$\phi(Q) \equiv -\sum_{j\in\mathcal{J}} \left\{ (c + a_j)\chi \sum_{i\in I} Q_{ij} + K_j \sqrt{\Sigma_{i\in I} Q_{ij}} \right\} + \chi \Sigma_{i\in I} T_i(Q).$$

The objective here is to maximize the entire system's expected annual profit including all facilities and retailers. In the objective function, the first term represents the facility location cost for opening facilities, and the second term is the annual purchasing cost from the supplier. The third and forth terms represent the working inventory and safety stock costs associated with each facility, respectively. The last term is the profit earned for the retailers.

The first constraint stipulates that retailers can only order from open facilities. We use an exponential function to formulate this restriction instead of other commonly used methods such as the big-M method because of the exponential function's quick convergence property. (Using the big-M method, we would model the constraint as $Q_{ij} <= M.X_j$, where M is a big constant.) We use the positive constant β in the first constraint to expedite the convergence.

Let ($\boldsymbol{X}^*$, $\boldsymbol{Q}^*$) denote the optimal solution to problem P. Because problem P is a highly nonlinear and mixed-integer optimization problem, with neither a convex nor a concave objective function, it's difficult to solve directly using any standard algorithm. We first study the relationship between problem P and its Lagrangian dual problem by relaxing the first constraint.[12] Here's problem LR:

$$v(\lambda) = \text{Maximize } \phi(Q) - \sum_{j\in\mathcal{J}} f_j X_j + \sum_{j\in\mathcal{J}} \sum_{i\in I} \lambda_{ij} \{X_j - 1 + e^{-\beta Q_{ij}}\}$$

subject to

$Q_{ij} \geq 0 \qquad i \in I, j \in \mathcal{J}$

$X_j \in \{0, 1\} \qquad j \in \mathcal{J}$,

Let $\boldsymbol{\lambda}$ denote the $n \times m$ matrix (λ_{ij}, $i \in I, j \in \mathcal{J}$), and rewrite the objective function of problem LR as

$$-\sum_{j\in\mathcal{J}} (f_j - \Sigma_{i\in I} \lambda_{ij}) X_j + \phi(Q) + \Sigma_{j\in\mathcal{J}} \Sigma_{i\in I} \lambda_{ij} (e^{-\beta Q_{ij}} - 1).$$

Because $\boldsymbol{X}$ and $\boldsymbol{Q}$ are independent in problem LR, we can determine the optimal solutions for $\boldsymbol{X}$ and $\boldsymbol{Q}$ separately. Because $X_j \in \{0, 1\}$, $\forall j \in \mathcal{J}$, we can directly determine its optimal solution using its corresponding coefficient, $-f_j + \Sigma_{i\in I}\lambda_{ij}$. If the coefficient of X_j is positive, then $X_j = 1$; if it's negative, $X_j = 0$. If the coefficient of X_j equals 0, we use the following rule to get a feasible solution: $X_j = 0$ when $Q_{ij} = 0$, and $X_j = 1$; otherwise, $\forall i \in I$. In a previous work, we proposed an algorithm[12] based on the bisection search and the outer approximation algorithm[15,16] and showed that it's more efficient than the outer approximation algorithm when applied to this problem.

Parameter Uncertainty

All the models I've just discussed assume that the decision maker knows the demand parameters—that is, the mean μ_i and the standard deviation σ_i of retailer i's demand. Larry Snyder, Mark Daskin, and C.P. Teo[17] presented a stochastic version of the model in the "Basic Model Formulation" section that explicitly handles parameter uncertainty by describing parameters via discrete scenarios, each with a specified probability of occurrence. The goal is to choose DC locations, assign retailers to DCs, and set inventory levels at DCs to minimize the total system-wide cost. To model this problem, we need to define the following additional notation: S is the set of scenarios indexed by s.

The parameters include

- μ_{is}, the mean daily demand at retailer i in scenario s for $i \in I, s \in S$;
- σ_{is}^2 , the variance of daily demand at retailer i in scenario s for $i \in I, s \in S$;
- d_{ijs}, the per-unit cost to ship from DC j to retailer i in scenario s for $i \in I, j \in \mathcal{J}, s \in S$; and
- q_s, the probability that scenario s occurs for $s \in S$.

The decision variables include

- x_j := 1, if j is selected as a facility location, but 0 otherwise for each $j \in \mathcal{J}$, and
- y_{ijs} := 1, if retailer $i \in I$ is served by DC $j \in \mathcal{J}$ in scenario $s \in S$ but 0 otherwise.

Now we can model this problem as follows. Minimize

$$\sum_{s \in S} \sum_{j \in \mathcal{J}} \left\{ \begin{array}{l} f_j X_j + \left(\sum_{i \in I} \hat{d}_{ijs} Y_{ijs} \right) \\ + K_j \sqrt{\sum_{i \in I} \mu_{is} Y_{ijs}} + q \sqrt{\sum_{i \in I} L_j \sigma_{is}^2 Y_{ijs}} \end{array} \right\},$$

which is subject to

$$\sum_{j \in \mathcal{J}} Y_{ijs} = 1 \text{ for each } i \in I, s \in S$$

$$Y_{ijs} - X_j \leq 0 \text{ for each } i \in I, j \in \mathcal{J}, s \in S$$

$$Y_{ij} \in \{0, 1\} \text{ for each } i \in I, j \in \mathcal{J}, s \in S$$

$$X_j \in \{0, 1\} \text{ for each } j \in \mathcal{J},$$

where

$$\hat{d}_{ijs} = \beta\chi\mu_{is}(d_{ijs} + a_j)$$

$$q = \theta h z_\alpha .$$

The authors also presented a Lagrangian-relaxation-based solution algorithm for this model.[18] They showed that the Lagrangian subproblem is a nonlinear integer program, but we can solve it by using a low-order polynomial algorithm. The authors also presented quantitative and qualitative computational results on problems with up to 150 nodes and nine scenarios, and described both algorithms' performance and solution behavior as key parameters changed. One of their major findings was that the stochastic (or min-expected-cost) solutions and the individual scenario solutions differed substantially in their choices of DC locations. This suggests that each of the single-scenario solutions would perform poorly in long-run expected cost. Furthermore, they observed that implementing the stochastic solution would entail roughly 8 percent regret on average and nearly 25 percent regret in the worst case. (The *regret* of a scenario is the difference [absolute or percentage] between the cost of the chosen solution under a given scenario and the cost of the optimal solution for that scenario.) Finally, they noted that, on average, half the retailers were assigned to different DCs in different scenarios, indicating the value of letting retailer assignments was scenario-dependent.

The algorithms I've reviewed here can also apply to a range of other concave cost-minimization problems. I believe the field offers many interesting future research problems. The "Basic Model Formulation" section in this article, for example, uses direct shipping to estimate the distribution cost, but a more realistic modeling of the vehicle routing cost could provide more valuable information on the benefit of integrating decisions from different levels.[19] Decision makers must also consider other important factors when designing a supply chain, such as exchange rates, production scheduling requirements, sale prices, sourcing flexibility, and import tariffs. In particular, future research related to robustness and risk management is important. Other disciplines such as financial engineering offer some risk-management tools, but the question is how to apply them to supply-chain design problems.

One of the first papers along this line[20] applied the conditional value-at-risk (CVaR) idea[21] to a location model. CVaR approximately equals the average of the worst-case α percent scenarios. The model minimizes the expected regret with respect to an endogenously selected subset of worst-case scenarios with a collective probability of occurrence that is at most $1 - \alpha$. This model, the α-reliable mean-excess regret model, demonstrated significant improvements over the α-reliable p-median minimax regret model in numerical tests. In addition, the paper's authors presented a heuristic that efficiently solves the α-reliable p-median minimax regret model by solving a series of mean-excess subproblems. My colleagues and I plan to apply the CVaR concept to the supply-chain design models I describe here in this article.

Unlike the classical facility location models that implicitly assume facilities will never fail, Snyder and Daskin[22] proposed location models in which facilities do fail from time to time due to poor weather, disasters, changes of ownership, or other factors. They applied a reliability concept to the median and uncapacitated facility location problems and presented an optimal Lagrangian relaxation algorithm. My colleagues and I[23] have proposed an approximation algorithm for the Snyder and Daskin models[22] and a worst-case bound for the algorithm. We also derived a model that deals with the more general case in which each facility has a different probability of failing. It would be interesting to apply the same reliability concept to the supply-chain design models I review in this article.

References

1. D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi, *Designing and Managing the Supply Chain*, McGraw-Hill, 2000.
2. Z.-J. Shen, "Integrated Supply Chain Design Models: A Survey and Future Research Directions," to be published in *J. Industrial and Management Optimization*, vol. 3, no. 1, 2007.
3. S. Chopra and P. Meindl, *Supply Chain Management: Strategy, Planning, and Operation*, 3rd ed., Pearson-Prentice Hall, 2007.
4. G. Eppen, "Effects of Centralization on Expected Costs in a Multi-Location Newsboy Problem," *Management Science*, vol. 25, no. 5, 1979, pp. 498–501.
5. Z.-J. Shen, *Approximation Algorithms for Various Supply Chain Problems*, doctoral thesis, Dept. of Industrial Eng. and Management Sciences, Northwestern Univ., 2000.
6. A.K. Chakravarty, J.B. Orlin, and U.G. Rothblum, "Consecutive Optimizers for a Partitioning Problem with Applications to Optimal Inventory Groupings for Joint Replenishment," *Operations Research*, vol. 33, no. 4, 1985, pp. 820–834.
7. N. Sauer, "On the Density of Families of Sets," *J. Combinatorial Theory Series A*, vol. 1, 1972, pp. 145–172.
8. V.N. Vapnik and A.Ya. Chervonenkis, "On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities," *Theory Probab. Appl.*, vol. 16, no. 2, 1971, pp. 264–280.
9. M. Daskin, C. Coullard, and Z.-J. Shen, "An Inventory-Location Model: Formulation, Solution Algorithm and Computational Results," *Annals of Operations Research*, vol. 110, 2001, pp. 83–106.
10. Z.-J. Shen, C.R. Coullard, and M.S. Daskin, "A Joint Location-Inventory Model," *Transportation Science*, vol. 37, no. 1, 2003, pp. 40–55.
11. J. Shu, C.P. Teo, and Z.-J. Shen, "Stochastic Transportation-Inventory Network Design Problem," *Operations Research*, vol. 53, no. 1, 2004, pp. 48–60.
12. L. Qi and Z.-J. Shen, "A Supply Chain Design Model with Unreliable Supply," working paper, Univ. of California, Berkeley; for a copy, please contact shen@ieor.berkeley.edu.
13. C.A. Yano and H.L. Lee, "Lot Sizing with Random Yields: A Review," *Operations Research*, vol. 43, no. 2, 1995, pp. 311–334.
14. M. Dada, N. Petruzzi, and L.B. Schwarz, "A Newsvendor Model with Unreliable Suppliers," in review with *Manufacturing & Service Operations Management*.
15. R. Horst, P. Pardalos, and N.V. Thoai, *Introduction to Global Optimization*, Kluwer Academic, 2000.
16. R. Horst and H. Tuy, *Global Optimization*, Springer-Verlag, 1990.
17. L.V. Snyder, M.S. Daskin, and C.-P. Teo, "The Stochastic Location Model with Risk Pooling," to be published in *European J. Operational Research*, 2007.
18. M. Parlar and D. Perry, "Analysis of a (q, r, t) Inventory Policy with Deterministic and Random Yields When Future Supply Is Uncertain," *European J. Operational Research*, vol. 84, 1995, pp. 431–443.
19. Z.-J. Shen and L. Qi, "Incorporating Inventory and Routing Costs in Strategic Location Models," to be published in *European J. Operational Research*, vol. 179, no. 2, 2007.
20. G. Chen et al., "The α-Reliable Mean-Excess Regret Model for Stochastic Facility Location Modeling," *Naval Research Logistics*, vol. 53, no. 7, 2006, pp. 617–626.
21. R.T. Rockafellar and S. Uryasev, "Optimization of Conditional Value-at-Risk," *J. Risk*, vol. 2, no. 3, 2000, pp. 21–41.
22. L.V. Snyder and M.S. Daskin, "Reliability Models for Facility Location: The Expected Failure Cost Case," *Transportation Science*, vol. 39, no. 3, 2005, pp. 400–416.
23. Z.-J. Shen, R. Zhan, and J. Zhang, "Reliability Models for Facility Location: Formulations and Algorithms," working paper, Univ. of California, Berkeley; for a copy, please contact shen@ieor.berkeley.edu.

***Zuo-Jun (Max) Shen** is an assistant professor in the Department of Industrial Engineering and Operations Research at the University of California, Berkeley. His research areas include integrated supply-chain design and management, market mechanism design, and applied optimization. Shen has a PhD in industrial engineering and management sciences from Northwestern University. Contact him at shen@ieor.berkeley.edu.*