Consolidation Strategies for the Delivery of Perishable Products

Christine Nguyen[a], Maged Dessouky[ac], and Alejandro Toriello[b]

[a]Daniel J. Epstein Department of Industrial and Systems Engineering
University of Southern California
3715 McClintock Ave
Los Angeles, CA 90089
{nguyen7, maged}@usc.edu

[b]H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
765 Ferst Drive, NW
Atlanta, Georgia 30332-0205
atoriello@isye.gatech.edu

[c]Corresponding author

Abstract

A set of agricultural suppliers with low demands can save on long-haul transportation costs by consolidating their product. We consider a system with stochastic demand and a single consolidation point near the suppliers. We propose a look-ahead heuristic that takes advantage of economies of scale by aiming to ship larger quantities. We experimentally compare the heuristic's performance against other simple policies, a rolling horizon algorithm, and a stochastic dynamic programming model. Our numerical results demonstrate that the heuristic provides solutions that are near the lower bound provided by the dynamic programming model, and that the benefits of consolidating depend on the size of the suppliers' demand. We also propose a proportional cost allocation rule that encourages the supplies to cooperate with each other instead of operating independently.

Keywords

1. Introduction

Collaborative strategies in the supply chain can greatly improve a company's performance. With the increase in competition, globalization and demanding customers, many firms believe they cannot continue to compete alone (Kumar and Banerjee 2012). Collaboration between firms offers opportunities to share risks, increase system efficiency, reduce costs, minimize unsatisfied customer demand, and increase their competitive advantage (Cao and Zhang 2010). However, collaboration will only work if *criticality* is present, where criticality is defined as "the notion of high recognized interdependence in which one supply chain member will not act in his own best interest to the detriment of the supply chain" (Spekman et al 1998).

In recent years, there is also an increasing interest in the potential savings of cooperation between multiple decision makers within or across supply chain levels. Specifically, joint strategies across multiple suppliers can decrease system-wide transportation costs through the consolidation of common products. Cooperating suppliers aim to minimize total joint costs but also need to determine an allocation that sustains continuous participation by each supplier. We focus primarily on terminal consolidation (defined by Hall 1987), where items from different origins are sorted at a single terminal to be shipped to different destinations on the same vehicle. The idea of consolidating to decrease costs is not a recent development. Early contributions to freight consolidation describing the opportunity for lower transportation costs and large shipment loads date back to at least the 1980s (Jackson 1985; Blumenfeld et al 1985; Closs and Cook 1987; Hall 1987). The terms *shipment consolidation* and *freight consolidation* are more popular in the current research literature than *terminal consolidation*.

In a shipment consolidation problem, we must determine how much to consolidate before shipping or how many time periods to consolidate before shipping the aggregate volume. Quantity-based polices determine a threshold weight or volume that must be accumulated before a shipment is released (Gupta and Bagchi 1987). Time-based policies dispatch after the first order in a consolidated load has waited for a predefined time (Mutlu and Çetinkaya 2010; Marklund 2011). Hybrid time-and-quantity polices release a shipment when either the quantity threshold is reached or the maximum waiting time has passed (Mutlu et al 2010). Çetinkaya and Bookbinder (2003) use renewal theory to investigate the quantity and time policy for private carriage and when transport is performed by a public, for-hire company. Ülkü (2009) also shows that quantity-based policies are the most cost-effective with unit-sized demands and Poisson arrivals in a shipment consolidation problem using private carriage. Higginson and Bookbinder (1995) use a discrete-time Markov decision process to study a sequential model where the shipper is required to reconsider the dispatch decision at the arrival of an order. Ülkü (2012) proposes a discrete-time based shipment consolidation policy that allows shipment release only at discrete times of the day while maintaining a certain customer service level.

Consolidating shipments allows shippers to take advantage of economies of scale and achieve decreased transportation costs. For example, Bausch et al (1995) showed that consolidation of Mobil Oil Corporation's heavy petroleum products could yield annual transportation savings of $1 million. Brown et al (2001) estimated approximately $35 million savings per year in inventory and distribution costs for Kellogg Company with the implementation of a new consolidation policy. Local and global third-party logistics companies also benefit from consolidation shipment-release policies (Lee et al 2003; Tyan et al 2003; Song et al 2008).

In this paper, we study a freight consolidation problem for perishable products where there is a hard time constraint for the product's stay in inventory at the consolidation center. This problem is motivated by the California cut flower industry, where growers currently do not consolidate their shipments (Nguyen et al 2013). That is, each grower sends shipments individually to its customers, primarily using a combination of less-than-truckload (LTL) rates and courier services instead of the more advantageous full truckload rates (FTL); high transportation costs are one of the major factors behind California's drop in U.S. cut flower market share from 64% to 20%. Growers from South America, who use a consolidation center in Miami, have concurrently seen their market share rise to 70%. Consolidation strategies are important to take advantage of economies of scale with perishable products and decrease system-wide transportation costs. For California flower growers, Nguyen et al (2013) showed that consolidation by 20 suppliers for all destinations in the United States could yield annual savings of at least $6 million. If 50 additional growers of the California Cut Flower Commission participated together, the savings could reach $17 million.

This paper is a generalization of the problem studied in Nguyen et al (2013). Instead of deterministic demand, we focus on stochastic demand. We also consider an environment in which the demand distribution changes between periods; in particular, we differentiate between peak and nonpeak periods, since some agricultural demand follows peak and nonpeak behavior and suppliers must plan accordingly in the weeks before harvesting and shipping. Nguyen et al (2013) considered the case where California cut flower growers would use a consolidation center in their transportation network. In order for the growers to maintain the claim that California cut flowers are fresher than imported cut flowers, the flowers could stay at the consolidation center at most one day. In this paper, we consider cases where the product can remain longer than a single day in inventory, but there is still a hard time constraint due to perishability. A time constraint greater than one day is realistic for other perishable products. For example, potatoes remain dormant 6 to 12 weeks after they are harvested and can be stored up to 2 to 3 months before they begin to sprout, depending on the variety and storage temperatures (Yanta and Tong 2013).

We solve the optimization problem using a stochastic dynamic programming approach, compare the performance of various heuristics, and study how changing demand distributions to the

system affect the benefit of consolidation. Furthermore, we propose a cost allocation policy and empirically show that the suppliers benefit from cooperating. We finally note that although this problem was motivated by the cut flower industry, the focus of this model is in developing strategies to lower long-haul transportation costs from a consolidation location to a break-bulk destination. We do not consider inventory costs and the transportation cost from the supplier to the consolidation center in this paper since it is assumed that these costs are significantly dominated by the long-haul transportation cost. These assumptions are reasonable when (1) products are harvested just before they are sent to the consolidation center, (2) perishability limits the amount of time the products can stay in the consolidation center, (3) the unit cost of the items is small, at least compared to the shipping cost, making the opportunity cost of capital insignificant, and (4) the suppliers are located close to the consolidation center. For agricultural products such as cut flowers where there is perishability and the growing cost is small relative to the shipping cost, the above assumptions tend to hold.

The remainder of the paper is organized in the following way. The next section gives a brief literature review. In Section 3, we present the dynamic programming formulation of the problem. We discuss heuristic approaches to the problem in Section 4. The numerical results are in Section 5, and we conclude the paper with possible extensions of the model and future considerations in Section 6.

2. Literature Review

The problem we consider has not been addressed to our knowledge in the research literature. However, some aspects are related to inventory control models such as the lot-sizing problem and joint replenishment problem. The following section is a brief review of the most recent developments.

2.1 Similar Inventory Control Problems

The classical *economic lot-sizing problem* (ELSP) is heavily studied in the inventory control and production planning research areas. Wagner and Whitin (1958) originally proposed the dynamic economic lot-sizing model and developed a dynamic programming algorithm to solve it. Since then, it continues to attract interest because it serves as the core problem to many applications. The literature is expansive, with many extensions and increasing complexity such as multiple suppliers, multiple retailers, multiple items, general cost functions, capacity constraints and cooperation.

Quadt and Kuhn (2008) provide a review of the capacitated lot-sizing problem with extensions such as backorders, linked lot sizes, sequence-dependent setups, and parallel machines. A review of exact and heuristic solution approaches for years 1988 to 2009 can be found in Robinson et al

(2009). Some exact algorithms like branch-and-cut (Guan et al 2006) have been applied to the stochastic economic lot-sizing problem, but a vast amount of proposed solution approaches are dynamic-programming based (Bai and Xu 2011; Kang and Lee 2013; Guan and Liu 2010).

Extensions and applications of the stochastic economic lot-sizing problem continue to increase in complexity and become computationally expensive to solve, making heuristics more appealing. For example, a fix-and-relax heuristic solves the stochastic lot-sizing problem in a reasonable amount of time (Beraldi et al 2006). The authors partition the original problem by time into a sequence of subproblems, where only a small number of variables must be integer. Alonso-Ayuso et al (2007) integrate the same heuristic into a fix-and-relax coordination framework that selectively explores the nodes of the search tree based on the characteristics of the non-anticipativity constraints.

Implementing a rolling horizon in the algorithmic framework is a common method used by managers to make lot-sizing decisions. Fisher et al (2001) propose an ending-inventory valuation algorithm that includes a valuation term in the objective function to offset end-effects. Their approach outperforms the Wagner-Whitin algorithm and the Silver-Meal heuristic (Silver 1976). Stadtler (2000) shows how an exact algorithm can perform at least as well as commonly used heuristics with an improved rolling horizon schedule. Zhang et al (2012) combine a dynamic programming algorithm with a rolling horizon heuristic to solve the stochastic uncapacitated ELSP with incremental quantity discounts.

The *joint replenishment problem* (JRP) focuses on finding replenishment policies for multiple products ordered by the same retailer, or for multiple locations ordering the same product. Khouja and Goyal (2008) provide a review of the classical JRP, stochastic JRP and dynamic JRP literature from 1989-2005. It summarizes heuristics and special approaches such as power-of-two policies and genetic algorithms. More recent developments include heuristics (Praharsi et al 2010; Zhang et al 2012), genetic algorithms (Hong and Kim 2009; Moon et al 2008; Cha et al 2008), and a linear programming-based heuristic framework (Amaya et al 2013).

The *one warehouse, multi-retailer* (OWMR) problem is a generalization both of the single-item lot-sizing problem and the joint replenishment problem. It focuses on determining the best lot-sizing policy of a warehouse and a set of retailers. A brief review can be found in Federgruen (1993). Extensions of the OWMR problem in recent years include dynamic allocation and retailer-reporting (Zhai et al 2011), and optimization of order-up-to policies (Wang 2013). Solution approaches to the problem include particle swarm optimization (Köchel and Thiem 2011), applying power-of-two policies (Chu and Shen 2010), and heuristics (Abdul-Jalbar et al 2010; Axsäter et al 2002).

2.2 Cost Allocation

The operations research literature for supply chain management and inventory decisions includes cost allocation strategies. While consolidation and cooperation can improve system-wide performance, the costs must be allocated in such a manner that individual decision makers have no incentive to leave the coalition.

Fiestras-Janeiro et al (2011) review the applications of cooperative game theory in centralized inventory management systems. Ben-Daya et al (2008) and Glock (2012) provide an extensive review of solution approaches to the joint economic lot-sizing problem.

Cooperation strategies have been applied to certain extensions of the economic lot sizing problem. Van den Heuvel et al (2007) show that economic lot sizing games with multiple retailers and deterministic demand are balanced, i.e. their core is non-empty. They also show that the ELS game with equal demand for each player and the 2-period ELS game are concave. Gopaladesikan et al (2012) use a primal-dual algorithm to find a cost allocation in the core of an economic lot sizing game. Xu and Yang (2009) propose a cost-sharing method and show that possessing properties such as cross-monotonicity, fairness and competitiveness makes it a viable option. Chen and Zhang (2006) use LP duality to show that an optimal dual solution defines an allocation in the core of an economic lot-sizing game with general concave ordering cost. Toriello and Uhan (2014) show how to calculate a dynamic cost allocation of a cooperative game where each player incurs demand in each period and coalitions can pool orders.

The cost allocation problem is of interest in the joint replenishment problem and similar inventory control problems as well. Anily and Haviv (2007) consider the cost allocation of the first-order interaction JRP under power-of-two policies. Elomri et al (2012) develop a procedure to form efficient coalition structures in a non-superadditive game of the JRP. Jokar and Sajadieh (2009) consider an integrated vendor-buyer production-inventory model. They demonstrate that the vendor and buyer benefit more from cooperating with each other than in competitive environments. Zhang (2009) applies a strong duality theorem to show that the joint replenishment game using power-of-two policies has a nonempty core. He et al (2012) apply general results from polymatroid optimization to the joint replenishment problem and show that the cooperative game is submodular.

3. Model

A set of suppliers ships product to a consolidation center, where products going to the same break-bulk destination leave in combined shipments. Transportation occurs from each supplier to the consolidation center, then from the consolidation center to the destination. We assume direct shipping methods and do not consider routing strategies. Since we assume that products going to different destinations do not share the same transportation resource, we consider only

the single destination problem. The products in this model are assumed to be perishable and can only stay at the consolidation center for a maximum amount of time. We define θ to represent the maximum number of days the product can stay at the consolidation center. We also assume the consolidation center is close to the suppliers and inventory is held at the consolidation center rather than at the suppliers. That is, the product is harvested immediately before shipment from the supplier to the consolidation center. The following list details all problem parameters:

S — Set of suppliers
T — Time horizon
θ — Maximum number of days inventory can stay at the consolidation center
D_{it} — Demand random variable with discrete distribution at time t for supplier i
κ_F — Capacity limit for a full truck
κ_L — Capacity limit for an LTL unit, e.g. a cubic foot or cube
c_F — FTL rate, per truck
c_L — LTL rate, per LTL cube
c_U — Courier service rate, per weight
α — Conversion factor, weight per cubic foot

Note that the demand distribution in each period t can vary for each supplier. In the cut flower application, there are two primary demand distributions: peak and nonpeak. Peak probability distributions occur during periods where demand is generally high. For example, higher demand for cut flowers occurs the week before Valentine's Day and the week before Mother's Day. All other days are considered to be nonpeak periods with a nonpeak demand probability distribution.

The product stays at the consolidation center for a relatively small amount of time, and we assume that the unit cost of the product is small relative to the shipping cost. We therefore do not consider inventory costs, because they are much lower than the transportation cost. Also, the model assumes that the consolidation center is close to the suppliers' sites, so the transportation cost to the consolidation center from the supplier is significantly smaller than the long-haul costs to each break-bulk destination. These assumptions hold for the cut flower industry.

3.1 Cost Function

The transportation cost includes three shipping options: full truckload (FTL), less-than-truckload (LTL) and courier services (i.e. UPS or FedEX). The FTL rate is a per truck rate. The LTL rate is a per volume unit rate (per cubic foot), and the courier services generally use a per weight rate such as per pound. All three rates are destination-dependent, but they follow a similar structure. Sending volume through a courier service is least expensive when the volume is extremely small but is the most expensive option for very large shipments. Shipping using FTL rates is more advantageous with very large volumes.
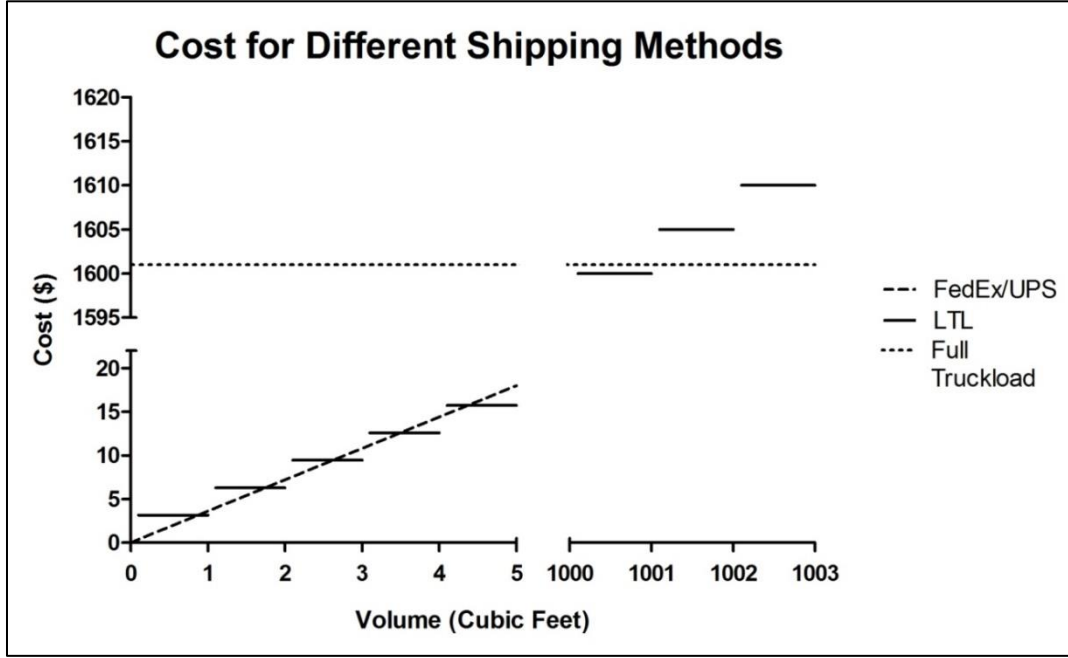
Figure 1. The cost of different shipping methods is destination-dependent.

Figure 1 illustrates the general cost structure the suppliers have. When a shipment is ready to leave, it can be shipped using a combination of the three methods. All shipments that can fit into a truckload are sent at the FTL rate, c_F. The remaining volume can be shipped at the FTL rate, LTL rate, combination of LTL and courier service or courier service only. The following function calculates the least-cost combination of the LTL and courier service shipping method.

$$\phi_{LU}(x) = \left\lfloor \frac{x}{\kappa_L} \right\rfloor c_L + \min\left(c_L, \left(x - \left\lfloor \frac{x}{\kappa_L} \right\rfloor \kappa_L\right) c_U\right) \tag{1}$$

Therefore, the cost function $\phi(x)$ for nonnegative values of volume x is the cost of shipping full truckloads, if $x > \kappa_F$, plus the least expensive shipping method for the remaining volume:

$$\phi(x) = \left\lfloor \frac{x}{\kappa_F} \right\rfloor c_F + \min\left(\left(x - \left\lfloor \frac{x}{\kappa_F} \right\rfloor \kappa_F\right) c_F, \phi_{LU}\left(x - \left\lfloor \frac{x}{\kappa_F} \right\rfloor \kappa_F\right)\right) \tag{2}$$

For volume up to a full truck, κ_F, we define two breakpoints: b_F and b_L. The breakpoint b_F is the volume above which the FTL rate is the least-cost shipping method. All volume less than b_F can be sent using a combination of LTL and courier service; we assume the volume can be separated. The breakpoint b_L is the volume above which sending using the LTL rate is more advantageous than using the courier service. For example, suppose κ_F is 2000 ft^3, κ_L is 1 ft^3, b_F is 1800 ft^3 and b_L is 0.7 ft^3. If the volume leaving today is 2800.9 ft^3, 2000 ft^3 is sent in a full truck. Since the remaining volume, 800.9 ft^3, is less than b_F, it is less expensive to send using LTL and courier

services. 800 ft^3 is best sent using the LTL rate because they are full units. The 0.9 ft^3 remaining is also best sent using a LTL rate because it is greater than b_L. If it were less than b_L, it would be sent at the courier rate.

$$b_F = \left\lfloor \frac{c_F}{c_L \kappa_L} \right\rfloor + \left(c_F - \left\lfloor \frac{c_F}{c_L \kappa_L} \right\rfloor c_L \kappa_L \right) \frac{1}{\alpha c_U} \tag{3}$$

$$b_L = \frac{c_L \kappa_L}{\alpha c_U} \tag{4}$$

We next present a stochastic dynamic programming formulation that determines the optimal movement of goods from the consolidation center to one destination.

3.2 Dynamic Programming Model

The stochastic dynamic programming model reflects the expected cost for a time horizon of length T. At each time t, the state variable, $\underline{\delta}_t = (\delta_t^0, \delta_t^1 \ldots \delta_t^k \ldots \delta_t^\theta)$, is a vector of nonnegative values of length $\theta+1$ that represents the amount of demand that must leave by period $t + k$, for $k = 0$ to θ. In other words, the state variable is the inventory, subdivided by shipping deadline. The decision rule is: for a given state variable, δ_t^0 must leave now and the decision maker must decide how much of the remaining inventory, $\delta_t^1 \ldots \delta_t^k \ldots \delta_t^\theta$, to add to the outgoing shipment to minimize the cost.

We define the optimal value function, $g_t(\underline{\delta}_t)$ in equation 5, as the expected cost for periods t to T when the current inventory at the consolidation center is $\underline{\delta}_t$. The expected cost contains two parts. The first part is the shipping cost of the outgoing shipment, using the cost function $\phi(x)$ as defined in Section 3.1. The outgoing shipment includes the demand that must leave today because it has been in inventory for θ days. Some amount ρ from the inventory with later deadlines can be added to the outgoing shipment. The second part is the expected cost-to-go for time periods $t+1$ to T, a subproblem of this problem. We can write the recursion of the stochastic dynamic program as:

$$g_t(\underline{\delta}_t) = \min_\rho \left\{ \phi(\delta_t^0 + \rho) + E_D \left[g_{t+1} \left(f_t \left(\underline{\delta}_t, \underline{\rho}_t, D_t \right) \right) \right] \right\}, \qquad \forall t = 1 \ldots T - 1 \tag{5}$$

$$g_t(\underline{\delta}_t) = \min_\rho \left\{ \phi(\delta_t^0 + \rho) + g_{t+1} \left(f_t \left(\underline{\delta}_t, \underline{\rho}_t, d \right) \right) \left(\prod_{d_1 \in \widehat{D}_{1,t+1+\theta}, \ldots d_{|S|} \in \widehat{D}_{|S|,t+1+\theta}} p_{i,t+1+\theta}(d_i) \right) \right\}, \tag{6}$$
$$\forall t = 1 \ldots T - 1$$

The stochastic dynamic programming model begins with the calculation of $g_T(\underline{\delta}_T)$. At time T, the state variable $\underline{\delta}_t$ is empty except for the vector element δ_T^0, which represents the demand that must leave by time T. The other values of the state variable vector, δ_t^k, equal zero. This simplifies the boundary condition to equal the shipping cost of δ_T^0.

$$g_T(\underline{\delta}_T) = \phi(\delta_T^0) \tag{7}$$

Equation 6 demonstrates how to determine the expected cost with respect to the demand probability distributions. Let $\widehat{D}_{i,t}$ be a finite set of all possible demand values for supplier i at time t. The probability of a demand d_i for supplier i at time t is denoted as $p_{i,t}(d_i)$. Let $D_t = \sum_{i \in S} D_{it}$ be a random variable representing the aggregate demand across all suppliers for time period t, and let $d = \sum_{i \in S} d_i$ be the aggregate arriving demand. $\underline{\rho}_t = \left(\rho_t^1 \dots \rho_t^k \dots \rho_t^\theta \right)$ is a vector representing the inventory added to the outgoing shipment. The transition function $f_t\left(\underline{\delta}_t, \underline{\rho}_t, d \right)$ updates the inventory level, $\underline{\delta}_t$, and determines the state for the subproblem after the arrival of demand $d_i \in \widehat{D}_{i,t}$ for all suppliers i in S. The new inventory level is the current inventory level minus the outgoing shipment (shifted forward one period), plus the incoming demand d from the suppliers. ρ_t^k is the amount taken from the inventory that must leave by $t+k$ and added to the outgoing shipment at time t. $\rho = \sum_{k=1}^\theta \rho_t^k$ is chosen such that the overall cost is minimized.

$\rho_t^k \leq \delta_t^k$ must be true for $k = 1$ to θ at time t; i.e. we cannot add more than what is available in inventory. ρ is any value from zero to $\sum_{k=1}^\theta \delta_t^k$, which means the volume added to the current outgoing shipment can include a portion or all of the inventory. ρ cannot be greater than $\sum_{k=1}^\theta \delta_t^k$, because it implies that more product is added to today's shipment than what is available in inventory. Therefore, feasible values for the amount that can be added to the outgoing shipment are limited to what must still leave later, ie. $0 \leq \rho_t^k \leq \delta_t^k$. The principle of optimality with equations 5, 6 and 7 suggests how an optimal policy can be constructed: At any encountered state, choose ρ to minimize the immediate cost plus the expected cost-to-go (Bertsekas 2005).

3.2.1 Transition Function

The transition function $f_t\left(\underline{\delta}_t, \underline{\rho}_t, d \right)$ updates the state variable $\underline{\delta}_t$ from period t to $t+1$. We adjust each δ_t^k so that it references the correct time period after a shipment of size $\rho + \delta_t^0$ leaves. When $t+1+\theta \leq T$, the transition function shifts the demand values by one period and updates the demand values according to what was added to today's shipment (see equation 8). We generate new demand values for each supplier and add it to the end of the demand vector. When $t+1+\theta >$

T, the transition function simply shifts and updates the values of the vector and does not generate a new demand value because it would be outside the time horizon (see equation 9).

If $t + 1 + \theta \leq T$

$$\underline{\sigma}_{t+1} = f_t\left(\underline{\delta}_t, \underline{\rho}_t, d\right) = \begin{cases} \sigma_{t+1}^k = \delta_t^{k+1} - \rho_t^{k+1}, \ \forall \, k = 0 \dots \theta - 1 \\ \sigma_{t+1}^\theta = d \end{cases}$$ (8)

If $t + 1 + \theta > T$

$$\underline{\sigma}_{t+1} = f_t\left(\underline{\delta}_t, \underline{\rho}_t, d\right) = \begin{cases} \sigma_{t+1}^k = \delta_t^{k+1} - \rho_t^{k+1}, \ \forall \, k = 0 \dots T - 1 \\ \sigma_{t+1}^k = 0, \ \forall \, k = T - 1 \dots t + 1 + \theta \end{cases}$$ (9)

3.2.2 State Space

The state variable is a vector of length $\theta + 1$ that represents the demands that must leave today and in the next θ days. Let D_m be the maximum daily demand. The possible states for one day are based on the discretization of the demand. Let ω be the discretization value. For the $\theta = 1$ case, $Y = D_m / \omega + 1$ is the number of possible states in one day including zero. Suppose $\omega = 1$ and $D_m = 10$. The states for one day start at 0 and increase in increments of 1 to 10 so $Y = 11$. As ω decreases, the expected cost for the horizon T increases in accuracy. ω should be chosen at a value that includes all three shipping methods. If ω is too large, the dynamic programming model will not include all shipping methods. For example, if $\kappa_L = 1$ and $\omega = 1$, then the state demand values are in whole LTL units. The demand will never be less than 1, and then courier services will never be used.

For θ days of demand that can stay in inventory, the number of states in one day is $Y^{\theta+1}$ because the state vector is a length of $\theta+1$. There are a total of $TY^{\theta+1}$ possible states across the entire time horizon. As θ increases, the number of states increases exponentially and is the major factor in the increasing computational complexity of the dynamic programming model.

Various methods exist to decrease the number of states in the dynamic programming model. Here, we use the following pruning algorithm. For each demand that must leave at time t, the dynamic programming model cycles through all possible values of ρ that correspond to state variables at time $t+1$. However, there are some demand states that ρ does not need to consider. If the demand leaving today is greater than b_F but less than κ_F, then the least-cost shipment method is the FTL rate. That means for values of ρ corresponding to a total outgoing shipment volume between b_F and κ_F, the cost of outgoing demand does not increase. In other words, inventory is being added to the outgoing shipment at no additional cost. This is easily extended for demands that are larger κ_F. In Section 5, we show the differences in runtime for the dynamic programming model with and without pruning. Other pruning strategies may also improve the speed of the dynamic programming model, but they are outside the scope of this paper.

4. Heuristics

The stochastic dynamic programming model contains computational limitations as the time horizon, T, and the max inventory stay, θ, increase. The quality of the solution depends on the discretization of the demands for the state variables as well. In this section, we first present a computationally efficient, simple look-ahead heuristic that requires knowing only θ days of demand in advance and contains no discretization of demand values. We then discuss some alternative heuristic strategies.

4.1 θ-based Consolidation Heuristic

We develop a look-ahead heuristic that makes decisions based on the transportation savings through consolidation of up to θ periods. The cost function in Figure 1 shows that the cost per volume is lowest at the FTL rate. The next lowest cost per volume is the LTL rate, and the most expensive cost per volume is using a courier service. Therefore, it is advantageous to ship large quantities at the FTL rate. Consolidating demand across time can yield enough volume to ship economically at the cheaper FTL rate.

$t + \theta$ is the deadline to ship product that arrives at the consolidation center at time t. From time t to $t + \theta$, this product can either wait in inventory until its deadline or be shipped earlier with an outgoing shipment to take advantage of a less expensive rate; i.e. from LTL rate to FTL rate. At time $t + \theta$, the outgoing shipment must include all the demand that arrived at time t and has not been shipped.

We defined the breakpoints, b_F and b_L, in Section 3.1. The heuristic uses the breakpoints to determine the best shipping option, and the amount necessary to ship at the FTL, LTL or courier service rate. The flowchart in Figure 2 provides an overview of the heuristic for $\theta = 1$. If there is an outgoing shipment in period t, then we determine the number of full trucks. Any trucks filled to capacity are shipped at the FTL rate. The remaining amount that still needs to be shipped is the volume that partially fills a truckload. If the remaining volume is zero, then the heuristic is done for today and updates to the next day. If the remaining amount plus tomorrow's demand is greater than b_F, then we add tomorrow's demand to fill the excess capacity of the partial truck and ship it at the FTL rate. Otherwise, we must determine if it should be shipped at the LTL rate. The heuristic calculates the number of full LTL units and ships those at the LTL rate. The remaining amount does not fill an LTL unit. If this remaining amount plus tomorrow's demand is greater than b_L, then we add enough of tomorrow's demand to fill the LTL unit and ship at the LTL rate. Otherwise, we ship the remaining volume at the courier rate and do not add tomorrow's demand.
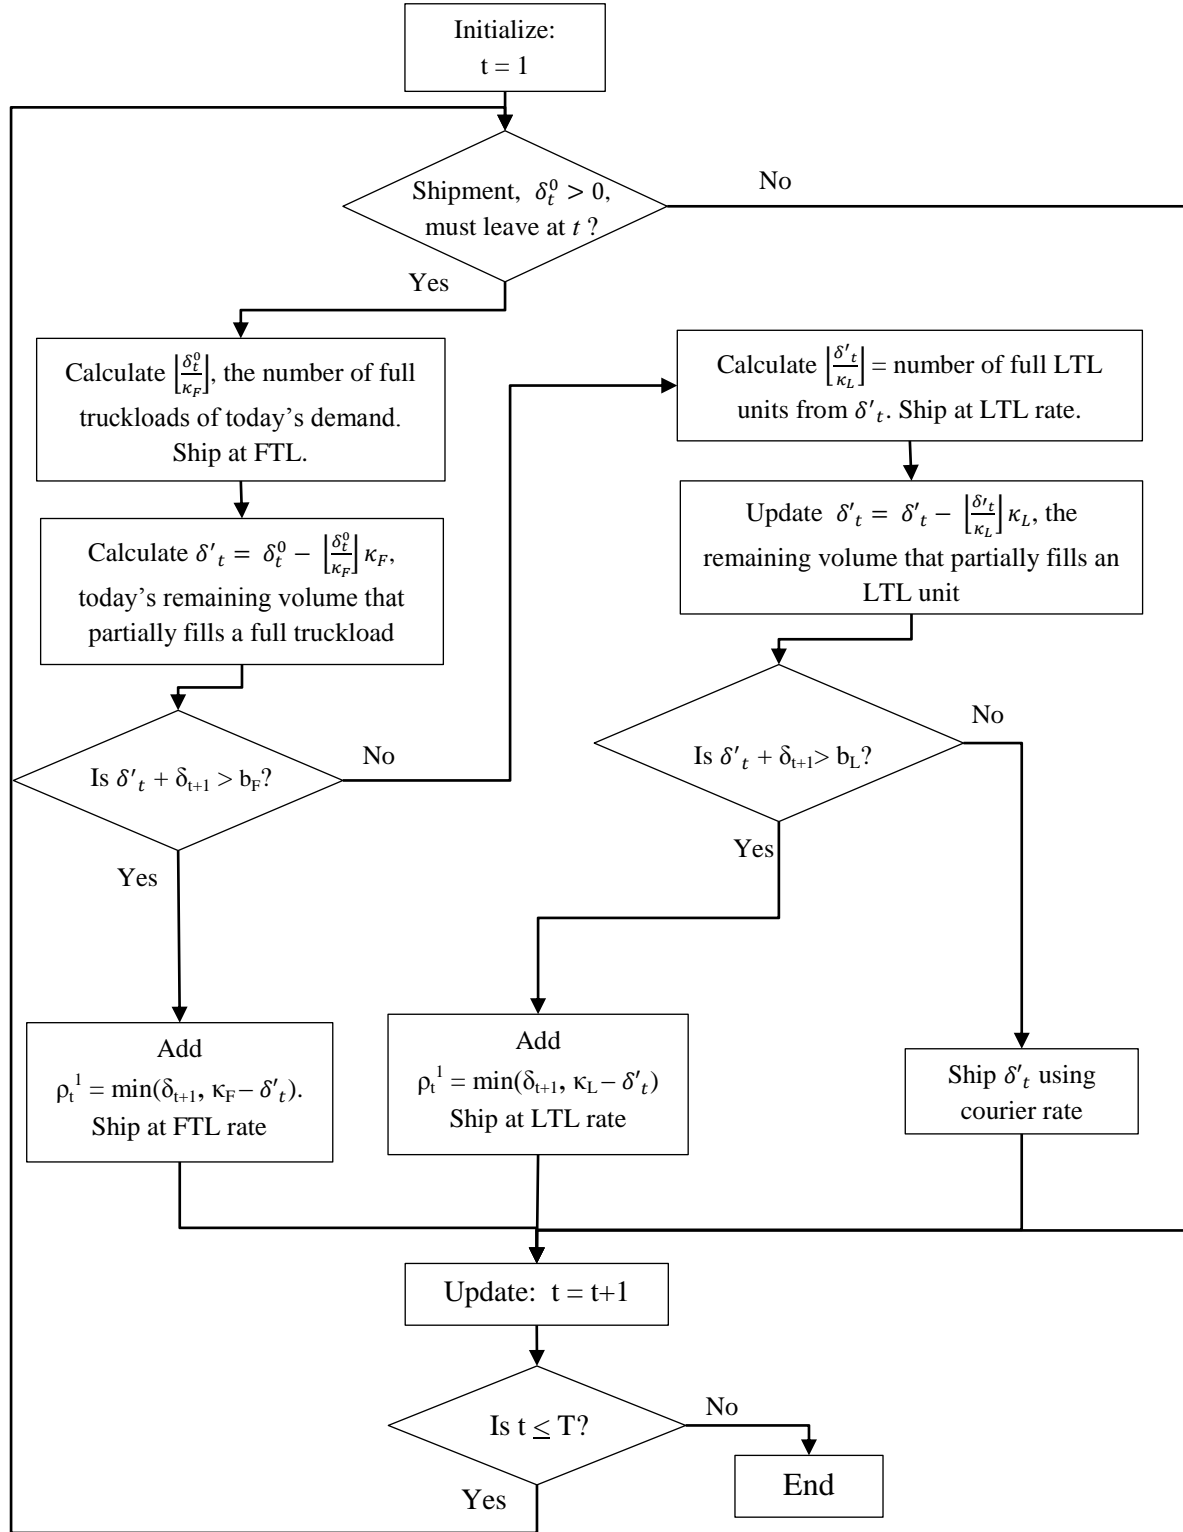
Initialize:
t = 1

Shipment, $\delta_t^0 > 0$, must leave at t ?

No

Yes

Calculate $\left\lfloor \frac{\delta_t^0}{\kappa_F} \right\rfloor$, the number of full truckloads of today's demand. Ship at FTL.

Calculate $\delta'_t = \delta_t^0 - \left\lfloor \frac{\delta_t^0}{\kappa_F} \right\rfloor \kappa_F$, today's remaining volume that partially fills a full truckload

Calculate $\left\lfloor \frac{\delta'_t}{\kappa_L} \right\rfloor$ = number of full LTL units from δ'_t. Ship at LTL rate.

Update $\delta'_t = \delta'_t - \left\lfloor \frac{\delta'_t}{\kappa_L} \right\rfloor \kappa_L$, the remaining volume that partially fills an LTL unit

Is $\delta'_t + \delta_{t+1} > b_F$?

No

Is $\delta'_t + \delta_{t+1} > b_L$?

No

Yes

Yes

Add $\rho_t^1 = \min(\delta_{t+1}, \kappa_F - \delta'_t)$. Ship at FTL rate

Add $\rho_t^1 = \min(\delta_{t+1}, \kappa_L - \delta'_t)$ Ship at LTL rate

Ship δ'_t using courier rate

Update: t = t+1

Is t ≤ T?

No

Yes

End

Figure 2. A flowchart of the θ-based Consolidation Heuristic for $\theta=1$.

The steps in the heuristic are very similar for greater values of θ. The only section of the heuristic that changes with θ is the product that is added to a partial truck or partial LTL unit. For the $\theta=1$ case, we consider only adding tomorrow's demand to the outgoing shipment. When θ increases, we consider the θ days of demand stored in inventory at the consolidation center when we compare with the breakpoints b_F and b_L to determine the shipping method. For example, if today's remaining demand plus all θ days of demand in inventory is greater than b_F, then we ship at the FTL rate. If there is excess capacity to ship at the FTL rate, the heuristic adds demands up to θ days from today to be added until the truck is filled to κ_F. The same occurs if the outgoing shipment should be shipped at the LTL rate. We consider demands that must leave up to θ days from today to be added to the partial LTL unit until it is filled to capacity, κ_L. We next present the detailed steps for the *θ-based consolidation heuristic* with a general θ value.

As stated in Section 3, T represents the entire time horizon and θ represents the maximum number of days demand can stay in inventory at the consolidation center. We use the breakpoints, b_F and b_L, in our heuristic to determine the least-cost shipping method for a given volume.

Step 0.   Set $t = 1$. Set $\underline{\delta_t} = \left(\delta_t^0, \delta_t^1 \dots \delta_t^k \dots \delta_t^\theta \right)$

Step 1.   Set $\underline{\rho_t} = \left(\rho_t^1 \dots \rho_t^k \dots \rho_t^\theta \right)$ to zero.

   If $\delta_t^0 > 0$, then a shipment must leave today, go to Step 2a.
   Otherwise, go to Step 4.

Step 2a. $\left\lfloor \frac{\delta_t^0}{\kappa_F} \right\rfloor$ is the number of full trucks to ship. Ship at the FTL rate.

   Let $\delta'_t = \delta_t^0 - \left\lfloor \frac{\delta_t^0}{\kappa_F} \right\rfloor \kappa_F$ be the remainder.

   If $\delta'_t + \delta_t^1 \geq b_F$, then ship at FTL rate, go to Step 2b.
   Otherwise, ship using LTL rates, go to Step 3a.

Step 2b.   We can fill the remaining FTL space $(\kappa_F - \delta'_t)$ at no additional cost.
   Set $k = 1$.

Step 2c.   Add $\rho_t^k = \min \left(\kappa_F - \delta'_t - \sum_{m=1}^{k-1} \rho_t^m, \ \delta_t^k \right)$ to the truck.

Step 2d.   If $\kappa_F - \delta'_t - \sum_{m=1}^{k-1} \rho_t^m - \delta_t^k > 0$, then go to Step 2e.
   Otherwise, go to Step 4.

Step 2e.   Update $k = k + 1$.
   If $k > \theta$, go to Step 4. Otherwise, go to Step 2c.

Step 3a. $\left\lfloor \frac{\delta'_t}{\kappa_L} \right\rfloor$ is the number of full LTL units to be shipped.

Update $\delta'_t = \delta'_t - \left\lfloor \frac{\delta'_t}{\kappa_L} \right\rfloor \kappa_L$

If $\delta'_t + \delta_t^1 \geq b_L$, then go to Step 3b.

Otherwise, go to Step 3f.

Step 3b. We can fill the remaining LTL space $(\kappa_L - \delta'_t)$ at no additional cost.

Set $k = 1$.

Step 3c. Add $\rho_t^k = \min \ (\kappa_L - \delta'_t - \sum_{m=1}^{k-1} \rho_t^m , \ \delta_t^k)$ to the truck.

Step 3d. If $\kappa_L - \delta'_t - \sum_{m=1}^{k-1} \rho_t^m - \delta_t^k > 0$, then go to Step 3e.

Otherwise, go to Step 4.

Step 3e. Update $k = k+1$.

If $k > \theta$, go to Step 4. Otherwise, go to Step 3c.

Step 3f. Send the remaining amount, δ'_t, using a courier service. Go to Step 4.

Step 4. If $t > T$, then end.

Otherwise, update the following:

Generate demand for each supplier at time $t+1+\theta$: $d_i \in \widehat{D}_{i,t+1+\theta}, \forall \ i \ \in S$

Update the demand vector for the next time period, $\underline{\delta}_{t+1} = f_t \left(\underline{\delta}_t, \underline{\rho}_t, d \right)$ using

aggregate demand $d = \sum_{i \in S} d_i$

Update to the next time period, $t = t+1$.

Go to Step 1.

4.2 Other Algorithms and Heuristics

To determine how well our heuristic performs, we compare it to a set of other algorithms and heuristics.

The *every $(\theta + 1)$ days* heuristic consolidates the demand for θ days before shipping. For example, if $\theta = 1$, then the heuristic ships every two days; i.e. one day's demand waits until the next day, when it is shipped with the second day's demand.

The *cost-to-go policy* creates a cost-to-go table for the entire horizon. This policy performs a deterministic dynamic programming algorithm where the demand for each day is the expected demand of the corresponding probability distribution for that period. The table contains the

expected cost-to-go for each demand discretization and each day of the year. The decisions are based on the cost-to-go tables.

The *rolling horizon* algorithm performs the deterministic dynamic programming algorithm for a number of predetermined days, R, and the first M days are implemented. The first M days includes the decisions that affect the first $M + \theta$ days. However, this also assumes that $M + \theta$ is less than or equal to the number of days in the run length ($M + \theta \leq R$). The deterministic dynamic programming model begins its run after the time horizon shifts forward M days, and it assumes the demand for R days is known; that is, this heuristic requires additional knowledge of future demand.

5. Results and Discussion

The data we used was taken from the California cut flower industry for the year 2010. We chose five suppliers shipping to the same destination who had average daily demand under a full truckload. Because of the low volumes, the suppliers would ship their product using a combination of courier services and less than truckload rates if working individually.
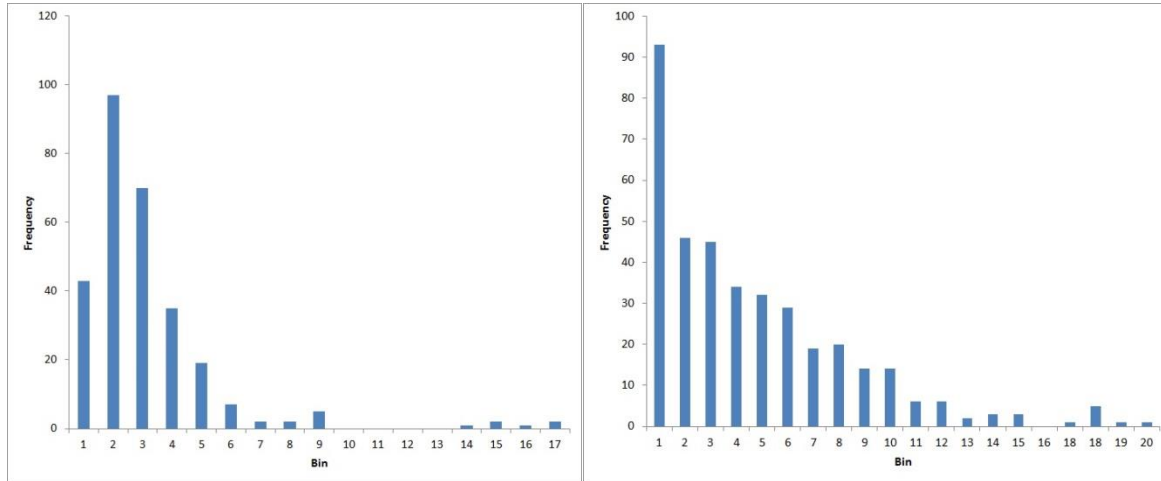

Figure 3. Demand distribution for two suppliers shipping to the same destination.

Figure 3 illustrates the empirical distributions of demand for two different suppliers shipping to the same destination. In Figure 3, the size of the bin was 100 cubic feet of demand. For each supplier, we estimated two empirical distributions - one for the nonpeak days of demand and the other for the peak days. Peak days in the data were days 34 to 43 for Valentine's Day shipments and days 116 to 127 for Mother's Day shipments. All other days were nonpeak demand days. The demand distributions for all supplier-destination pairs followed a different pattern, so we used empirical distributions in our simulation.

5.1 Comparison of Heuristics

We generated 100 sets of demand each with a time horizon T of 365 days from the empirical distributions for the five suppliers. A lower bound was determined by solving the deterministic dynamic program of the problem on each realized set of sampled demand. The *θ-based consolidation heuristic* and *every (θ+1) heuristic* assume that θ days of demand were known ahead of time.

Table 1 below shows the averages of the 100 runs, where each run was one sample of realized demand for the year. The stochastic and deterministic dynamic programming formulations were based on a demand discretization value of 0.5. We ran all the dynamic programing simulations with pruning, yet it was only computationally possible to find the optimal solution from the stochastic dynamic programming formulation for $\theta=1$. It took approximately 13.1 hours of CPU time on a machine with an Intel Core i3 CPU 530 processor @ 2.93 GHz to find this optimal solution, whereas the lower bound for 100 samples of annual demand could be computed in approximately 12 minutes of CPU time. The last three columns of the table show the ratio of each heuristic solution's cost over the best possible bound. For $\theta=1$, it is the ratio of the corresponding heuristic divided by the stochastic DP optimal cost. For $\theta >1$, it is the ratio with the lower bound value.

Table 1. Comparison of algorithms and heuristics

θ	Stochastic DP	Lower Bound	θ-based consolidation heuristic	Every (θ + 1)	Cost-To-Go
1	$202,903	$195,303	1.01	1.01	1.24
2	-	$141,594	1.04	1.09	1.62
3	-	$119,408	1.04	1.19	1.75
4	-	$110,475	1.04	1.26	1.80
5	-	$107,309	1.04	1.25	1.79

The *θ-based consolidation heuristic* provided a reasonable schedule that costs only about 4% more than the lower bound and in the $\theta=1$ case is within 1% of the optimal solution. As θ increased, the annual cost decreased due to the increased availability of product to be consolidated. A larger θ value means demand can be held at the consolidation center for a longer period of time, and more periods of demand can be consolidated. The *θ-based consolidation heuristic* considered the demand that must leave within the next θ days to determine if the FTL rate could be achieved. If there was excess capacity, then it continued to add demand and take advantage of adding product at no additional cost, since the suppliers paid for the entire truck regardless of whether it was filled or not. If it was sent at the LTL rate instead, the heuristic added inventory to fill an LTL unit since the LTL rate is a per cubic foot rate. Shipping one full truck on two separate days or shipping out two full trucks on one day costs exactly the same in this model; therefore, the heuristic took advantage of θ and allowed product to stay in the

consolidation center as long as possible to increase the opportunity for shipping at the cheaper FTL rate.

The *every (θ +1)* and *cost-to-go* heuristics do not check for the same consolidation opportunities as the *θ-based consolidation heuristic*. The performance for both policies decreased as θ increased from 1 to 5. The *every (θ +1)* policy is a time-based policy that will ship θ days of demand even if the most recently arrived demand is better off being shipped later. This heuristic consolidates even at additional cost; as θ increased, more demand was shipped out before its deadline. The *cost-to-go* heuristic's table recommended an amount to consolidate with the outgoing shipment. This amount was based on deterministic values equal to the expected value of the empirical distributions. When realized demand was greater than the expected value, the amount added to the outgoing shipment exceeded any excess capacity of a truck or LTL unit. This increased the total annual cost compared to the heuristic. When realized demand was less than the expected value, the inventory added was less than the recommendation, and the outgoing shipment was not shipped at the cheaper rate. The transportation cost was higher since that demand could have stayed in inventory to be consolidated with another shipment. As θ increased, the *cost-to-go* performed poorly since the amount recommended to consolidate was based on an increasing number of days of expected demand. Increasing θ resulted in a higher difference in the amount the *cost-to-go* table recommended to be consolidated and the amount of realized demand that was consolidated. This difference increased, and the transportation costs increased.

We also tested the pruning algorithm on data sets generated to find the Lower Bound in Table 1. Table 2 shows the difference in runtime before and after the pruning is added.

Table 2. The average runtime for a dynamic programming model with and without pruning.

Theta	Average Runtime (seconds)	
	Without Pruning	With Pruning
1	37.9	13.0
2	74.7	17.7
3	153.0	28.3
4	158.3	43.5
5	202.7	57.5

The dynamic programming model with pruning significantly reduces the runtime. For $\theta = 5$, the runtime decreases from approximately 3 minutes to 1 minute. Additional pruning techniques may further reduce the runtime.

Table 3. The comparison between the θ-based consolidation heuristic and the rolling horizon algorithm.

θ	θ-based consolidation heuristic	Ratio by Run Length					
		5 Days	10 Days	20 Days	30 Days	40 Days	50 Days
1	$206,160	1.03	1.01	0.95	0.95	0.95	0.95
2	$146,483	1.11	1.07	0.98	0.97	0.97	0.97
3	$123,218	1.23	1.16	1.03	1.00	0.98	0.97
4	$114,676	1.21	1.16	1.08	1.04	1.01	0.99
5	$111,348	1.25	1.17	1.14	1.11	1.06	1.03

For the *rolling horizon* algorithm, we chose an implementation period of five days to correspond with a typical work week. We tested the algorithm with 5, 10, 20, 30 40 and 50 days of known demand. In this experiment, we compared the annual cost of the rolling horizon algorithm with the *θ-based consolidation heuristic*. Table 3 shows the comparison between the *θ-based consolidation heuristic* and the *rolling horizon* algorithm. The figures in the second column represent the annual cost of the *θ-based consolidation heuristic* and the other columns show the ratio of the *rolling horizon* algorithm's cost over the heuristic.

The *rolling horizon* algorithm performed poorly against the *θ-based consolidation heuristic* until the number of days of known demand was greater than or equal to 20 days, and even then it did not beat the heuristic when θ was large. This result was most likely due to end-period effects of the *rolling horizon* algorithm since it forced the schedule to consider an ending inventory of zero. The decisions in the last θ days of the run length included more outgoing shipments than the consolidation decision because there were no incoming demands. By defining an ending inventory of zero, we restricted the consolidation opportunities for the last θ days. As the number of days of known demand increased, the *rolling horizon* performed better. However, it is highly unlikely to know demand exactly that far in advance.

5.2 Sensitivity of Consolidation

We next compare the benefits of consolidating versus each supplier operating independently as a function of the demand size. In this analysis, we used the *θ-based consolidation heuristic* to generate the solutions for the consolidation and independent strategies. Figure 4 illustrates the benefits of consolidation as expected demand varies.
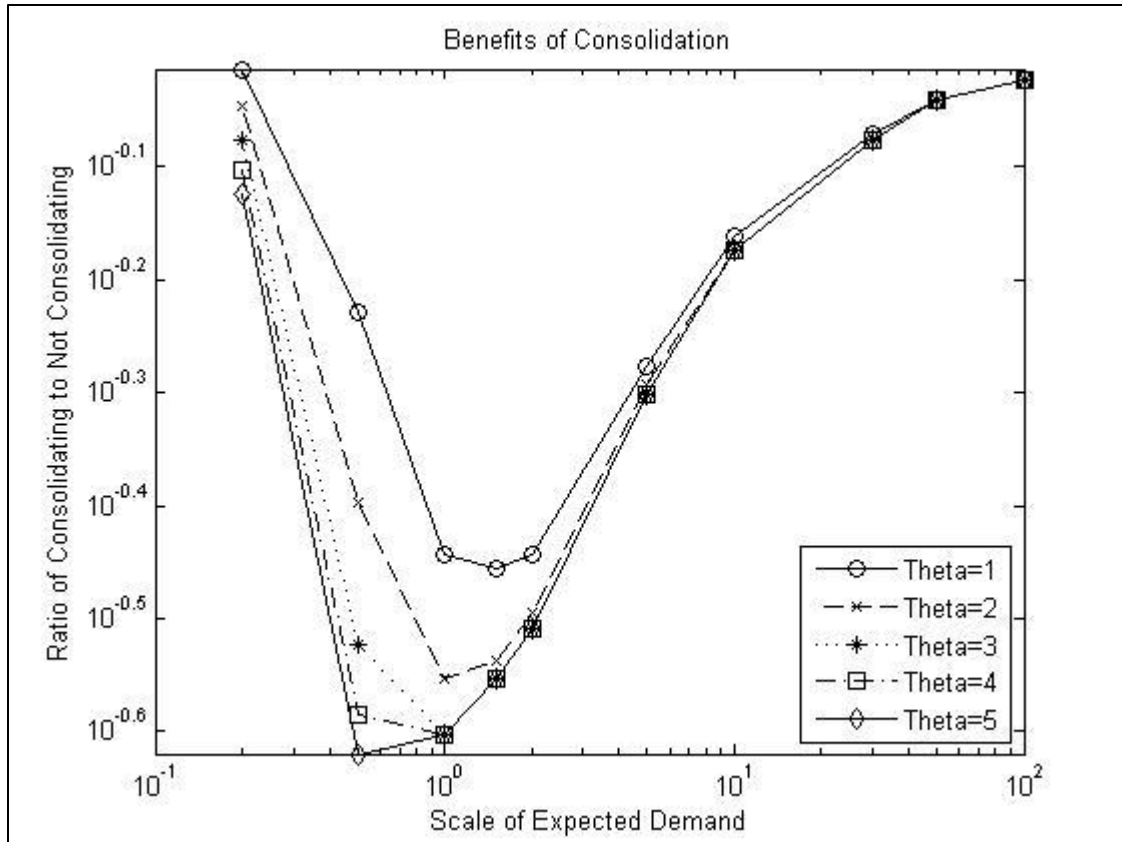
Figure 4. Ratio of consolidating using the θ-based consolidation heuristic versus operating independently.

We graphed the annual cost of the consolidation strategy divided by the annual cost of shipping independently. In other words, the values graphed are ratio values comparing a consolidation scenario versus a scenario with no consolidation. We scaled the actual demand by 0.2, 0.5, 1, 1.5, 2, 5, 10, 30, 50 and 100. The pattern illustrated in this graph shows that at very low demand distributions, the benefits of consolidating are small compared to not consolidating. This is a result of the cumulative demand being so low that the least-cost shipping option is still LTL. As demand increases, the aggregate demand over all five suppliers is sent at the more advantageous FTL rate (ratio values are closer to 0). However, as demand increases to extremely large values (50x and 100x), the benefits disappear because each supplier's individual demand is large enough to ship at the FTL rate without additional demand (ratio values are closer to 1). Consolidating at high volumes results in approximately the same number of full truckloads as operating independently.

Figure 5 re-emphasizes how consolidating benefits those supplier with lower daily volumes. The number of suppliers shipping together decreases as each supplier's demand increases. As the supplier's daily demand increases to κ_F, the need to consolidate decreases because individual suppliers can send out full truckloads. The excess capacity in a partial truckload decreases, and

fewer suppliers are able to add product to a partially loaded truck. With small daily demand values, more suppliers consolidate to achieve the FTL rate.
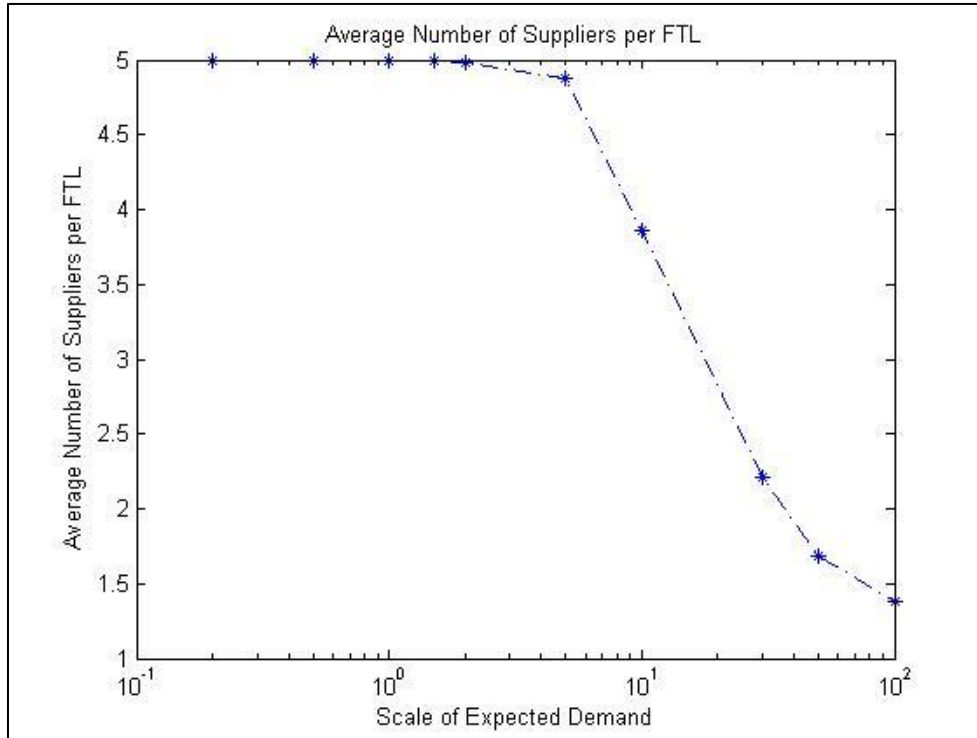

Figure 5. The average number of suppliers per full truckload decreased as demand increased.

5.3 Cost Allocation

While consolidating across suppliers with low values of demand provides system-wide benefits, independent suppliers will not cooperate unless there are individual advantages as well. The question of how to allocate the costs among the suppliers still remains. In this section, we briefly look at a potential cost allocation policy and determine whether it would support the cooperation of suppliers at the consolidation center.

We next propose a policy that charges each supplier a proportion of the total cost equal to the proportion of demand in the current shipment that belongs to the supplier. The cost is divided every time there is a shipment.   That is, let c be the total cost for the outgoing shipment to a destination, and let β_i be the proportion of volume that originates from supplier i. Then the supplier will pay $c\beta_i$.

We performed this allocation on 100 samples of annual demand using only nonpeak distributions. Full cooperation, known as the *grand coalition*, includes all 5 suppliers. The average daily cost for suppliers 1 through 5 when cooperating fully was $147.37, $100.38, $150.43, $73.70 and $153.57, respectively (see Table 4). We compared these costs with

scenarios where a subset of the 5 suppliers, or *coalition*, formed. This cost allocation rule encourages participation in the grand coalition only if there is no other coalition with greater savings. To do this, we considered other coalition possibilities of the 5 suppliers.

We considered all possibilities to create a coalition of size 4, 3, 2 and 1 using the original 5 suppliers. For example, the coalition of size 2 was all cooperation possibilities between 2 suppliers, which included the following: {1, 2}, {1, 3}, {1, 4}, {1, 5}, {2, 3}, {2,4}, {2,5}, {3,4}, {3,5}, and {4,5}. We applied the proportional allocation policy for each possible coalition and calculated the average daily cost across all the coalitions of the same size. Table 4 contains the average daily cost for the grand coalition. Each value in Table 5 is a ratio of the average daily cost for the corresponding coalition size divided by the average daily cost of fully cooperating (values in Table 4).

Table 4. The average daily cost under the proportional allocation policy.

Coalition size	Supplier 1	Supplier 2	Supplier 3	Supplier 4	Supplier 5
5	$147.37	$100.38	$150.43	$73.70	$153.57

Table 5. Ratio comparison against the grand coalition of 5 suppliers

Coalition size	Supplier 1	Supplier 2	Supplier 3	Supplier 4	Supplier 5
4	1.17	1.21	1.15	1.21	1.17
3	1.49	1.62	1.43	1.69	1.47
2	2.04	2.33	1.92	2.52	1.99
1	2.77	3.03	2.57	3.00	2.72

Because all ratio values are greater than 1, the corresponding supplier would pay more on average if he were in a smaller coalition or operating independently versus the grand coalition. Every ratio value was greater than 1 for subset sizes 1 through 4, and from subset size 1 to 4 the ratio for each supplier decreases. These results imply that under the proposed cost allocation policy, every supplier benefits from a full cooperation versus operating alone or in a smaller coalition. However, the benefit for each supplier is not necessarily the same. One supplier might benefit much more than the other suppliers if his demand is relatively low and therefore easily consolidated with the remaining suppliers' partial LTL units or partial truckloads.

6. Conclusion

We studied a consolidation problem for long-haul shipments with multiple suppliers and one consolidation center delivering perishable products to a single destination. We formulated a stochastic dynamic programming model that solved the optimization problem. In addition, we

developed a look-ahead heuristic and compared it with various other heuristics. The *θ-based consolidation heuristic* yields a good solution compared to the optimal solution and other policies. For example, the *θ-based consolidation heuristic* requires knowledge of fewer days of deterministic demand versus the *rolling horizon* algorithm; the *rolling horizon* algorithm performs better only if the length of the horizon is long and the demand during this horizon is known. In general, the *θ-based consolidation heuristic* performs the best of all the heuristics tested. We showed how the benefits of consolidation vary as the demand sizes change. Consolidation gives the highest benefit when demand levels are moderate. At the very low demand levels, consolidation does not take advantage of the cheaper full-truckload rates. At the very high demand levels, there is no need to consolidate to achieve economies of scale.

There are several limitations to the *θ-based consolidation heuristic*. The delay at the consolidation center influences the shelf life of the product, and this model does not measure the quality of the product. It uses one parameter to determine how long a product stays in inventory. This would be effective if the product were allowed to stay in inventory a very short period of time, such as a few days. However, the quality of the product might decrease with each passing day, which could influence the decision to ship or not.

Future work in this area includes the development of a quality measure for the perishable products. Inventory costs at the consolidation center were not included in this model, and if they are, the *θ-based consolidation heuristic* may not perform as well. The trade-off between inventory costs and transportation costs will need to be considered to develop a new heuristic. Vehicle routing strategies offer opportunities for more transportation savings as well. Multiple pick-ups from suppliers would utilize more space in the vehicle and decrease costs compared to direct shipping. Incorporating routing into the existing model would also consider the travel time between the suppliers and the consolidation center.

References

Abdul-Jalbar, B., A. Segerstedt, J. Sicilia, and A. Nilsson. 2010. "A New Heuristic to Solve the One-Warehouse N-Retailer Problem." *Computers & Operations Research* 37 (2): 265–272.

Alonso-Ayuso, A., L.F. Escudero, M.T. Ortuño, and C. Pizarro. 2007. "On a Stochastic Sequencing and Scheduling Problem." *Computers & Operations Research* 34 (9): 2604–2624.

Amaya, C.A., J. Carvajal, and F. Castaño. 2013. "A Heuristic Framework Based on Linear Programming to Solve the Constrained Joint Replenishment Problem." *International Journal of Production Economics* 144 (1): 243–247.

Anily, S. and M. Haviv. 2007. "The Cost Allocation Problem for the First Order Interaction Joint Replenishment Model." *Operations Research* 55 (2) : 292–302.

Axsäter, S., J. Marklund, and E.A. Silver. 2002. "Heuristic Methods for Centralized Control of One-Warehouse, N -Retailer Inventory Systems." *Manufacturing & Service Operations Management* 4 (1): 75–97.

Bai, Q.G. and J.T. Xu. 2011. "Optimal Solutions for the Economic Lot-Sizing Problem with Multiple Suppliers and Cost Structures." *Journal of Applied Mathematics and Computing* 37 (1-2): 331-345

Bausch, D.O., G.G. Brown, and D. Ronen. 1995. "Consolidating and Dispatching Truck Shipments of Mobil Heavy Petroleum Products." *Interfaces* 25 (2): 1–17.

Ben-Daya, M., M. Darwish, and K. Ertogral. 2008. "The Joint Economic Lot Sizing Problem: Review and Extensions." *European Journal of Operational Research* 185 (2): 726–742.

Beraldi, P., G. Ghiani, A. Grieco, and E. Guerriero. 2006. "Fix and Relax Heuristic for a Stochastic Lot-Sizing Problem." *Computational Optimization and Applications* 33 (2-3): 303–318.

Bertsekas, D.P. *Dynamic Programming and Optimal Control*. 3rd ed. Boston, Massachusetts: Athena Scientific, 2005.

Blumenfeld, D.E., L.D.Burns, D.J. Diltz, and C.F. Daganzo. 1985. "Analyzing Trade-Offs Between Transportation, Inventory and Production Costs on Freight Networks." *Transportation Research Part B* 19 (5): 361–380.

Brown, G., J. Keegan, B. Vigus, and K. Wood. 2001. "The Kellogg Company Optimizes Production, Inventory, and Distribution." *Interfaces* 31 (6): 1–15.

Cao, M. and Q. Zhang. 2010. "Supply Chain Collaborative Advantage: A Firm's Perspective." *International Journal of Production Economics* 128 (1): 358–367.

Çetinkaya, S. and J.H. Bookbinder. 2003. "Stochastic Models for the Dispatch of Consolidated Shipments." *Transportation Research Part B:Methodological* 37 (8): 747-768

Cha, B.C., I.K. Moon, and J.H. Park. 2008. "The Joint Replenishment and Delivery Scheduling of the One-Warehouse, n-Retailer System." *Transportation Research Part E: Logistics and Transportation Review* 44 (5): 720–730.

Chen, X. and J. Zhang. "Duality Approaches to Economic Lot-Sizing Games." IOMS: Operations Management Working Paper OM-2006-01, Stern School of Business, New York University, 2006.

Chu, L. and Z.J.M. Shen. 2010. "A Power-of-Two Ordering Policy for One-Warehouse Multiretailer Systems with Stochastic Demand." *Operations Research* 58 (2): 492–502.

Closs, D.J. and R.L. Cook. 1987. "Multi-Stage Transportation Consolidation Analysis Using Dynamic Simulation." *International Journal of Physical Distribution and Logistics Management* 17 (3): 28–45.

Elomri, A., A. Ghaffari, Z. Jemai, and Y. Dallery. 2012. "Coalition Formation and Cost Allocation for Joint Replenishment Systems." *Production and Operations Management* 21 (6): 1015–1027.

Federgruen, A. 1993. "Centralized Planning Models for Multi-Echelon Inventory Systems Under Uncertainty." In *Handbook in Operations Research and Management Science*, edited by S Graves, A Rinnooy Kan, and P Zipkin, 133–173. North Holland, Amsterdam, The Netherlands: Logistics of Production and Inventory.

Fiestras-Janeiro, M.G., I. García-Jurado, A. Meca, and M.A. Mosquera. 2011. "Cooperative Game Theory and Inventory Management." *European Journal of Operational Research* 210 (3): 459–466.

Fisher, M., K. Ramdas, and Y.S. Zheng. 2001. "Ending Inventory Valuation in Multiperiod Production Scheduling." *Management Science* 47 (5): 679–692.

Glock, C.H. 2012. "The Joint Economic Lot Size Problem: A Review." *International Journal of Production Economics* 135 (2): 671–686.

Gopaladesikan, M., N.A. Uhan, and J. Zou. 2012. "A Primal–dual Algorithm for Computing a Cost Allocation in the Core of Economic Lot-Sizing Games." *Operations Research Letters* 40 (6): 453–458.

Guan, Y., S. Ahmed, G.L. Nemhauser, and A.J. Miller. 2006. "A Branch-and-Cut Algorithm for the Stochastic Uncapacitated Lot-Sizing Problem." *Mathematical Programming Series A* 105 (1): 55–84.

Guan, Y. and T. Liu. 2010. "Stochastic Lot-Sizing Problem with Inventory-Bounds and Constant Order-Capacities." *European Journal of Operational Research* 207 (3): 1398–1409.

Gupta, Y.P. and P.K. Bagchi. 1987. "Inbound Freight Consolidation Under Just-In-Time Procurement." *Journal of Business Logistics* 8 (2): 74–94.

Hall, R. W. 1987. "Consolidation Strategy : Inventory, Vehicles And Terminals." *Journal of Business Logistics* 8 (2): 57–73.

He, S., J. Zhang, and S. Zhang. 2012. "Polymatroid Optimization, Submodularity, and Joint Replenishment Games." *Operations Research* 60 (1): 128–137.

Higginson, J.K. and J.H. Bookbinder. 1995. "Markovian Decision Processes in Shipment Consolidation." Transportation Science 29 (3): 242-255.

Hong, S.P. and Y.H. Kim. 2009. "A Genetic Algorithm for Joint Replenishment Based on the Exact Inventory Cost." *Computers & Operations Research* 36 (1): 167–175.

Jackson, G.C. 1985. "A Survey of Freight Consolidation Practices." *Journal of Business Logistics* 6 (1): 13–34.

Jokar, M.R.A. and M.S. Sajadieh. 2009. "Optimizing a Joint Economic Lot Sizing Problem with Price-Sensitive Demand." *Transaction E: Industrial Engineering* 16 (2): 159–164.

Kang, H.Y. and A.H.I. Lee. 2013. "A Stochastic Lot-Sizing Model with Multi-Supplier and Quantity Discounts." *International Journal of Production Research* 51 (1): 37–41.

Khouja, M. and S. Goyal. 2008. "A Review of the Joint Replenishment Problem Literature: 1989–2005." *European Journal of Operational Research* 186 (1): 1–16.

Köchel, P. and S. Thiem. 2011. "Search for Good Policies in a Single-Warehouse, Multi-Retailer System by Particle Swarm Optimisation." *International Journal of Production Economics* 133 (1): 319–325.

Kumar, G. and R.N. Banerjee. 2012. "Collaboration in Supply Chain: an Assessment of Hierarchical Model using Partial Least Squares." *International Journal of Productivity and Performance Management* 61 (8): 897-918.

Lee, C.Y., S. Çetinkaya, and W. Jaruphongsa. 2003. "A Dynamic Model for Inventory Lot Sizing and Outbound Shipment Scheduling at a Third-Party Warehouse." *Operations Research* 51 (5): 735–747.

Marklund, J. 2011. "Inventory Control in Divergent Supply Chains with Time-Based Dispatching and Shipment Consolidation." *Naval Research Logistics* 58 (1): 59-70.

Moon, I.K., S.K. Goyal, and B.C. Cha. 2008. "The Joint Replenishment Problem Involving Multiple Suppliers Offering Quantity Discounts." *International Journal of Systems Science* 39 (6): 629–637.

Mutlu, F., S. Çetinkaya, and J.H. Bookbinder. 2010. "An Analytical Model for Computing the Optimal Time-and-Quantity-Based Policy for Consolidated Shipments." *IIE Transactions* 42 (5): 367–377.

Mutlu, F. and S. Çetinkaya. 2010. "An Integrated Model for Stock Replenishment and Shipment Scheduling Under Common Carrier Dispatch Costs." *Transportation Research Part E: Logistics and Transportation Review* 46 (6): 844–854.

Nguyen, C., A. Toriello, M. Dessouky and J.E. Moore II. 2013. "Evaluation of Transportation Practices in the California Cut Flower Industry." *Interfaces* 43 (2): 182-193.

Praharsi, Y., H.D. Purnomo, and H.M. Wee. 2010. "An Innovative Heuristic for Joint Replenishment Problem with Deterministic and Stochastic Demand." *International Journal of Electronic Business Management* 8 (3): 223–230.

Quadt, D. and H. Kuhn. 2008. "Capacitated Lot-Sizing with Extensions: a Review." *4or* 6 (1): 61–83.

Robinson, P., A. Narayanan, and F. Sahin. 2009. "Coordinated Deterministic Dynamic Demand Lot-Sizing Problem: A Review of Models and Algorithms." *Omega* 37 (1): 3–15.

Silver, E.A. 1976. "A Simple Method of Determining Ordr Quantities in Joint Replenishments Under Deterministic Demand." *Management Science* 22 (12): 1351–1361.

Song, H., V.N. Hsu, and R.K. Cheung. 2008. "Distribution Coordination Between Suppliers and Customers with a Consolidation Center." *Operations Research* 56 (5): 1264–1277.

Spekman, R.E., J.W. Kamauff Jr, and N. Myhr. 1998. "An Empirical Investigation into Supply Chain Management: a Perspective on Partnerships." *International Journal of Physical Distribution and Logistics Management* 48 (8): 630-650

Stadtler, H. 2000. "Improved Rolling Schedules for the Dynamic Single-Level Lot-Sizing Problem." *Management Science* 46 (2): 318–326.

Toriello, A. and N.A. Uhan. 2014. "Dynamic Cost Allocation for Economic Lost Sizing Games." *Operations Research Letters* 42 (1): 82-84.

Tyan, J.C., F.K. Wang, and T.C. Du. 2003. "An Evaluation of Freight Consolidation Policies in Global Third Party Logistics." *Omega* 31 (1): 55–62.

Ülkü, M.A. 2009. "Comparison of Typical Shipment Consolidation Programs: Structural Results." *Management Science and Engineering* 3 (4): 27–34.

Ülkü, M.A. 2012. "Dare to Care: Shipment Consolidation Reduces not only Costs, but also Environmental Damage." *International Journal of Production Economics* 139(2): 438-446.

Van den Heuvel, W., P. Borm, and H. Hamers. 2007. "Economic Lot-Sizing Games." *European Journal of Operational Research* 176 (2): 1117–1130.

Wagner, H.M. and T.M. Whitin. 1958. "Dynamic Version of the Economic Lot Size Model." *Management Science* 5 (1): 89–96.

Wang, Q. 2013. "A Periodic-Review Inventory Control Policy for a Two-Level Supply Chain with Multiple Retailers and Stochastic Demand." *European Journal of Operational Research* 230 (1): 53–62.

Xu, D. and R. Yang. 2009. "A Cost-Sharing Method for an Economic Lot-Sizing Game." *Operations Research Letters* 37 (2): 107–110.

Yanta, J.P. and Tong, C.B. 2013. "Commercial Postharvest Handling of Potatoes (Solanum tuberosum)." Regents of the University of Minnesota. Accessed May 10, 2014. http://www.extension.umn.edu/garden/fruit-vegetable/commercial-postharvest-handling-of-potatoes/

Zhai, X., J.E. Ward, and L.B. Schwarz. 2011. "Coordinating a One-Warehouse N-Retailer Distribution System Under Retailer-Reporting." *International Journal of Production Economics* 134 (1): 204–211.

Zhang, J. 2009. "Cost Allocation for Joint Replenishment Models." *Operations Research* 57 (1): 146–156.

Zhang, R, I. Kaku, and Y. Xiao. 2012. "Model and Heuristic Algorithm of the Joint Replenishment Problem with Complete Backordering and Correlated Demand." *International Journal of Production Economics* 139 (1): 1–9.

Zhang, Y., S. Song, C. Wu, and W. Yin. 2012. "Dynamic Programming and Heuristic for Stochastic Uncapacitated Lot-Sizing Problems with Incremental Quantity Discount." *Mathematical Problems in Engineering* 2012: 1–21.