

A Joint Replenishment Policy with Individual Control and Constant Size Orders

M. Mustafa Tanrikulu • Alper Şen • Osman Alp

Procter & Gamble, Kosifler İş Merkezi A Blok, Serin Sokak No:9, İçerenköy, 34752, İstanbul, Turkey

Department of Industrial Engineering, Bilkent University, Bilkent, 06800, Ankara, Turkey

Department of Industrial Engineering, Bilkent University, Bilkent, 06800, Ankara, Turkey

tanrikulu.mm@pg.com • alpersen@bilkent.edu.tr • osmanalp@bilkent.edu.tr

We consider inventory systems with multiple items under stochastic demand and jointly incurred order setup costs. The problem is to determine the replenishment policy that minimizes the total expected ordering, inventory holding, and backordering costs; the so-called stochastic joint replenishment problem. In particular, we study the settings in which order setup costs reflect the transportation costs and have a step-wise cost structure, each step corresponding to an additional transportation vehicle. For this setting, we propose a new policy which we call the (s, Q) policy, under which a replenishment order of constant size Q is triggered whenever the inventory position of one of the items drops to its reorder point s . The replenishment order is allocated to multiple items so that the inventory positions are equalized as much as possible. The policy is designed for settings in which backorder and setup costs are high, as it allows the items to independently trigger replenishment orders and fully exploits the economies of scale by consistently ordering the same quantity. A numerical study is conducted to show that the proposed (s, Q) policy outperforms the well-known (Q, S) policy when backorder costs are high and lead times are small.

1. Introduction

Many companies manage inventories of multiple items. The primary challenge in managing multi-item inventory systems is the fact that some of the costs are incurred jointly. In particular, the setup costs in production, purchasing and transportation are often incurred jointly for many items that are included in any given production batch, purchase order or shipment. Joint setups can be seen as an opportunity as well as a challenge, since scale economies can be exploited to reduce setup costs, cycle inventories or both, by carefully coordinating the replenishment of many items. The joint replenishment problem is to determine the inventory replenishment policy of multiple items that share a common setup.

A basic example of the joint replenishment problem occurs in a setting where multiple items are sourced from a common supplier. Setup costs in this setting may include transportation costs and purchase transaction costs. Since the 1980s, many manufacturing companies have been reducing their supplier bases. For example, Xerox reduced its supplier base in early 1980s from 5000 to 400 (Burt, 1989); Texas Instruments reduced its Maintenance, Repair and Operating (MRO) suppliers from 5000 to 750 between 1998 and 2000 (Pantumsinchai, 2000); Merck reduced its total global supplier base from 40,000 in 1992 to fewer than 10,000 in 1997 (Genna, 1997); IBM reduced the number of its suppliers to only 50 for 85% of its requirements (Carbone, 1999); and Sun Microsystems reduced its supplier base to 40 for 90% of its requirements (Carbone, 1996). Among other things, reduction of the supplier base helps companies decrease their inventory holding, transportation and purchasing costs by giving them the capability of jointly replenishing multiple items from common suppliers.

Supply base reduction is only one dimension of supply base optimization and is not always desired as it may lead to companies becoming highly dependent on few suppliers (Talluri and Narasimhan, 2005). However, sourcing from a common supplier is not a requirement for jointly replenishing multiple items. Companies are devising numerous strategies to leverage economies of scale of combining different items into a single delivery. Among these, the milk-run strategy allows the joint procurement of multiple SKUs from different suppliers located in close physical proximity, and helps companies consolidate smaller shipments into more efficient larger shipments (or move from infrequent independent shipments to more frequent joint shipments). This lets them reduce transportation costs and cycle stocks. For example, Toyota's Kentucky plant sources 80% of its parts from suppliers located within 200 miles of the plant. Milk-run vehicles serving these suppliers help Toyota receive deliveries on a JIT basis (Minahan, 1997). Another example is Eastman Kodak where the milk-run strategy significantly increased the frequency of inbound shipments to its plants (Cook et al., 2005). A final milk-run example is the commercial vehicle producer MAN. In 2005, MAN's Ankara plant successfully reduced its inbound transportation costs and component inventory by consolidating its shipments from various component manufacturers clustered in Northwestern Turkey (Bostancı et al., 2005).

Cross-docking is another strategy that allows companies to exploit economies of scale in inbound transportation; smaller shipments from multiple suppliers can be merged in a consolidation warehouse for a larger and more economical joint delivery. Cross-docking has

been a successful strategy, particularly in the famous Wal-Mart implementation (Simchi-Levi et al., 2003).

Joint replenishment is also relevant when replenishing a single item in multiple locations. As in the case of multi-item inventory systems, companies are developing strategies that will help them exploit economies of scale by combining shipments to multiple locations under their control. For example, a milk-run vehicle can depart from a supplier or a distribution center and replenish a group of production plants, reducing the transportation costs and cycle inventories. An example of this is again Eastman Kodak which also uses the milk-run strategy for shipments from its distribution center to multiple plants (Cook et al., 2005). Milk runs are widely used to replenish multiple retail store locations from retailer-owned distribution centers or from suppliers directly. Cross-docking also enables multiple facilities to consolidate their replenishment at least for a portion of the trip. Joint replenishment of multiple locations is possible when all these locations are centrally controlled or when these locations are in a coalition for joint replenishment. Under a Vendor Managed Inventory (VMI) contract between a supplier and multiple retailers (or other downstream players), the supplier takes control of the management of inventories at the retail locations. Among other benefits, VMI contracts allow the joint replenishment of multiple retail locations and help reduce transportation and inventory costs for the supply chain (Çetinkaya and Lee, 2000, and Cheung and Lee, 2002).

The main contribution of this paper is to propose a new continuous review policy for the SJRP under a step-wise truck cost structure. In particular, we develop a policy in which a replenishment order of fixed size (perhaps the capacity of the vehicle) is created whenever the inventory position for one of the N items (locations) drops to its reorder point. We call this policy (\mathbf{s}, Q) policy, where \mathbf{s} is the vector of reorder points for the N items, and Q is the constant reorder quantity. When items are identical, the contents of the replenishment order is decided in a way that equalizes the item inventory positions (to the extent that this is possible). For the case of non-identical items, we suggest a heuristic policy that allocates the incoming order such that inventory positions in excess of the re-order points are equalized. There are some advantages of Q being constant. First, since we are using capacitated vehicles for shipment, companies can maintain stable and acceptable utilization level on the trucks that they dispatch. Moreover, when Q is constant and there is a carefully selected allocation policy, it is possible to avoid ordering items which have higher inventory positions altogether (this is not true for policies with order-up-to levels).

A partial motivation for this study is our experience with a beverage producer in Turkey. This beverage producer manages the inventory of its distributors under a VMI like setting and dispatches trucks for the replenishment of about 100 SKUs at each distributor. The trucks that are used for the shipments are capacitated. Since the trucks travel large distances (up to 1000 kilometers), the transportation costs are substantial (as compared to inventory holding costs) and do not depend significantly on the load of the truck, the beverage producer almost always dispatches full trucks to its distributors. The company also wants to maintain a high service level at its distributors at which the demand for the SKUs can be highly uncertain. This rules out a policy that lacks the ability of each SKU individually triggering a replenishment order (one such policy is (Q, \mathbf{S}) policy, in which a replenishment order is triggered only when the total demand since last order reaches Q).

We conduct a numerical study to assess the performance of the proposed (\mathbf{s}, Q) policy against the (Q, \mathbf{S}) policy. The (Q, \mathbf{S}) policy is suggested by Cachon (2001) and is a natural candidate for comparison since it is a continuous review policy and employs constant size (capacity of the vehicle) replenishments. We did not use the other continuous review joint replenishment policies as these policies do not create constant-size replenishment orders. The results show that the (\mathbf{s}, Q) policy tends to outperform the (Q, \mathbf{S}) policy especially when the backorder costs are high and the lead times are small. This shows that the (\mathbf{s}, Q) policy is a proper alternative for replenishing multiple items (or locations) with joint setups in environments that exhibit these characteristics (e.g., the beverage producer mentioned above).

The rest of this paper is organized as follows. In Section 2, we review the relevant literature. In Section 3, we propose our new policy (\mathbf{s}, Q) along with a continuous time Markov chain to model the inventory in the system. In Section 4, we present numerical results comparing the two policies. Section 5 concludes the paper and suggests some avenues for future research.

2. Literature Review

Reflecting its importance in practice, the joint replenishment problem has been an extensively studied research topic since the pioneering works of Balintfy (1964) and Silver (1965). Formally, the problem is to determine the replenishment and inventory policy for N items (or locations) that minimizes the total setup, holding and shortage costs given joint setup costs.

In a more general setting, in addition to the setup costs that are common and incurred with each replenishment order regardless of which items are involved (major setups), item-specific setup costs may be incurred for each item in the order (minor setups). Research in this area has followed two separate paths depending on whether the demands are deterministic or stochastic; the latter is referred to as the stochastic joint replenishment problem (SJRP) and is the subject of this paper. There is a large body of literature on the deterministic joint replenishment problem; we refer the reader to two review articles by Aksoy and Erengüç (1988) and Khouja and Goyal (2007).

Balintfy (1964) is the first to study the SJRP. He proposes the so called can-order policy, denoted by $(\mathbf{s}, \mathbf{c}, \mathbf{S})$, where a joint order is triggered when the inventory position of an item drops to its reorder point s . The items with inventory positions below their corresponding can-order level, c , are also included in this order, and the inventory positions of all items included are raised to level S . Despite a simple structure, the exact analysis of the $(\mathbf{s}, \mathbf{c}, \mathbf{S})$ policy is extremely difficult. An approximate analysis of the system is provided in Silver (1974). Another continuous-review policy is the (Q, \mathbf{S}) policy first proposed by Renberg and Planche (1967). In this policy, when the total amount of demand observed since the previous order reaches Q , a joint order of Q units is placed to raise the inventory position of all items to \mathbf{S} . Pantumsinchai (1992) subsequently studies this policy under Poisson demand. He compares the (Q, \mathbf{S}) policy to the $(\mathbf{s}, \mathbf{c}, \mathbf{S})$ policy and shows that the (Q, \mathbf{S}) policy outperforms the other if the fixed ordering cost is high and the shortage cost is low. Atkins and Iyogun (1988) propose a class of periodic review replenishment policies, denoted by (\mathbf{R}, T) ; in this class, items' inventory positions are reviewed every T periods and are raised to level \mathbf{R} by placing a joint replenishment order. Their computational results indicate that as the fixed cost increases, (\mathbf{R}, T) policies outperform the can-order policy. Pantumsinchai (1992) shows that the performance of the (\mathbf{R}, T) policies is comparable to that of the (Q, \mathbf{S}) policy. Viswanathan (1997) suggests the periodic (\mathbf{s}, \mathbf{S}) policy where the inventory is reviewed periodically every T periods and all items with inventory positions below their s values are ordered up to S . Viswanathan's extensive numerical study shows that this policy generally outperforms the (\mathbf{R}, T) , (Q, \mathbf{S}) and can-order policies. Nielsen and Larsen (2005) study a stochastic review version of the policy of Viswanathan (1997) where the inventory is reviewed when the total demand since last order reaches Q . A computational study shows that this new policy $Q(\mathbf{s}, \mathbf{S})$ outperforms the periodic (\mathbf{s}, \mathbf{S}) policy in all cases and generally outperforms the standard (Q, \mathbf{S}) policy.

Recently, two new policies for the SJRP are proposed by Gürbüz et al. (2007) and Özkaya et al. (2006a). Gürbüz et al. (2007) propose the $(\mathbf{Q}, \mathbf{s}, \mathbf{S})$ policy for a supply chain environment where a cross-dock location serves multiple identical retailers. In this policy, a joint replenishment order is triggered when a total of Q demands are observed or when a retailer's inventory position drops to its reorder point s . Whenever a replenishment order is triggered, the inventory position of each retailer is raised to its order-up-to level S . The authors compare their policy to the (Q, \mathbf{S}) policy, the (\mathbf{R}, T) policy, and a special case of the can-order policy. The results show that the proposed policy performs better than the other policies under the settings considered. The authors also show that their policy outperforms other policies when additional transportation penalty costs are incurred when the number of units shipped exceeds a given truck capacity. Finally, Özkaya et al. (2006a) propose the (Q, \mathbf{S}, T) policy in a single location, N -items setting. According to this policy, a new joint replenishment order is triggered and inventory positions of all items are increased to their order-up-to points S , whenever a total of Q units are demanded or when T time units have elapsed since the previous replenishment order. It is shown that the (Q, \mathbf{S}, T) policy outperforms the (Q, \mathbf{S}) , $Q(\mathbf{s}, \mathbf{S})$, periodic (\mathbf{s}, \mathbf{S}) and can-order policies in most of the problem instances considered. This study is extended to multi-echelon setting by Özkaya et al. (2006b).

One major issue that has been overlooked in the existing literature is the fact that the setup costs (major or minor) may depend on the size of the order. It may be reasonable to assume that they are independent when the setup costs reflect the administrative costs that are related to a purchase order or the production setups that are incurred for a production batch. However, when the setup costs are due to transportation costs (perhaps the main motivation for joint replenishment), such an assumption is rather restrictive. In practice, vehicles with limited capacity are used for transportation. Thus, the setup cost structure is step-wise, each step corresponding to an additional vehicle. Such cost structures have been recently studied in the literature for single item stochastic inventory systems (see, for example, Alp et al., 2003) or for multiple items deterministic inventory systems (see, for example, Ben-Khedher and Yano, 1994). For multiple items stochastic inventory systems, the only related work that we are aware of is by Cachon (2001). This author considers a SJRP where multiple retailers are sourced by a single warehouse by using trucks of fixed capacity and a step-wise cost (truck cost) structure. He proposes three policy alternatives for this environment: the (Q, \mathbf{S}) , (\mathbf{R}, T) , and $(Q, \mathbf{S}|T)$ policies. The third policy is the minimum-quantity periodic-review policy where the retailer reviews its inventory position

every T time units and the maximum possible amount of the total demand observed during the last period is dispatched by full truckload shipments; the remaining quantity, if any, is shipped if it is greater than or equal to Q units.

3. Model

We consider a supply chain environment where inventories of multiple items at a single location or a single item at multiple locations are controlled in a continuous and coordinated fashion by a centralized authority. For consistency, we use the “multiple items single location” setting throughout the text without loss of generality. Items are sourced from one or more suppliers, all of which we refer to as the “warehouse”. The location where the items observe customer demand is referred to as the “retailer”. Joint orders are delivered by trucks of fixed and identical sizes. We assume that there is an ample supply of items at the warehouse and the fleet size is sufficiently large so that shipment of a joint order can be initiated without any delay whenever it is placed. The transit time between the warehouse and the retailer constitutes the replenishment lead time, which is denoted by L_i for item i . We assume that the demand for each item follows an independent Poisson process at the retailer with a rate of λ_i and that unsatisfied demands are fully backordered. The holding cost, h_i , is incurred at the retailer per item per unit time basis. There is a backorder cost of π_i per unit per occasion and p_i per unit per unit of time. The fixed ordering cost, K , is associated with the use of trucks, i.e., for each truck of capacity C utilized for shipment, a fixed cost of K is incurred independently of the quantity loaded. The retailer aims to find a coordinated replenishment policy for the inventory management of her N items to minimize the total expected ordering, inventory holding, and backordering costs. As explained in Section 1, the (Q, \mathbf{S}) policy and the proposed (\mathbf{s}, Q) policy are two policies that can be used in such environments.

3.1 The (Q, \mathbf{S}) policy

This policy is first suggested by Renberg and Planche (1967) for the general joint replenishment problem. The policy under Poisson demand is studied by Pantumsinchai (1992). The policy under a capacitated vehicle is studied by Cachon (2001). In this policy, when the total amount of demand since the previous order reaches to Q units, an order in the amount of Q is placed so that the retailer raises the inventory positions of all items up to

the vector $\mathbf{S} = (S_1, S_2, \dots, S_N)$, where S_i is the order-up-to level of item i . That is, when the total inventory position $IP(t) = \sum_{i=1}^N IP_i(t)$ drops to $S_T - Q$, where $S_T = \sum_{i=1}^N S_i$, an order amount of Q is placed to raise the inventory positions of all items up to their order-up-to level S_i . $S_T - Q$ can be assumed as the system reorder point.

Pantumsinchai (1992) presents the derivation of the total expected cost function of this policy. Since all of the inventory positions are raised to their order-up-to points, S_i whenever an order is triggered, inventory positions become a regenerative process. Therefore, the inventory positions of items reach to a steady state and their limiting probability can be computed. The cumulative demand for an item since the last order is binomially distributed for given cumulative demand for all items. If we let X_i be the random variable for the cumulative demand since the last order for item i and X_0 be the $\sum_{i=1}^N X_i$, $P(X_i|X_0)$ becomes binomial with parameters x_0 and θ_i , where $\theta_i = \lambda_i/\lambda_0$, where $\lambda_0 = \sum_{i=1}^N \lambda_i$ and x_0 is uniformly distributed between 0 and $Q - 1$. Therefore, the marginal distribution of X_i , referred to as $u(x_i)$ becomes:

$$u_i(x_i) = \frac{1}{Q} \sum_{x_0=x_i}^{Q-1} \binom{x_0}{x_i} \theta_i^{x_i} (1 - \theta_i)^{x_0 - x_i}, \quad x_i = 0, 1, \dots, Q - 1,$$

as shown in Pantumsinchai (1992).

It can be shown using recursive calculations that the marginal distribution of X_i is given by:

$$u_i(x_i) = \frac{1}{\theta_i Q} (1 - B(x_i; Q, \theta_i)), \quad x_i = 0, 1, \dots, Q - 1,$$

where $B(x_i; Q, \theta_i)$ is the cumulative binomial probability. The expected value of X_i is $\theta_i(Q - 1)/2$ and the variance is $\theta_i(1 - \theta_i)(Q - 1)/2 + \theta_i^2(Q^2 - 1)/12$.

In order to calculate different cost components, we need to calculate the stockout probabilities and the expected backorder size at any time which require us to derive the probability distributions of inventory positions and the net inventories of the items. For this, for now we drop the index i . It is known that the inventory position at any time t depends on the fixed lead time L . Assuming that the inventory position of an item is z at time $t - L$ and the demand for the item between $t - L$ and t is d , the net inventory at any time t becomes $z - d = S - v$ where $v = x + d$. This is because, items ordered before time $t - L$ will be on hand by time t but the items ordered after time $t - L$ will not be on hand by time t . If we

let $m(v)$ be the probability distribution of v , $m(v)$ becomes:

$$m(v) = \sum_{x=0}^{\min(v, Q-1)} u(x)r(v-x), \quad v = 0, 1, 2, \dots,$$

where $r(\cdot)$ is the probability of demand during lead time.

We now can calculate the stockout probability and expected size of backorder at any time. If we let $P(S, Q)$ be the stockout probability and $B(S, Q)$ be the expected size of backorder at any time then the equations become:

$$P(S, Q) = Pr(v \geq S) = \sum_{v=S}^{\infty} m(v),$$

and

$$B(S, Q) = \sum_{v=S+1}^{\infty} (v-S)m(v).$$

With all this information and re-introducing the index for the items, we write the expected total cost equation per unit of time as follows:


$$\begin{aligned} ETC(Q, \mathbf{S}) &= \frac{K\lambda_0}{Q} + \sum_{i=1}^N h_i(S_i - \theta_i(Q-1)/2 - \lambda_i L_i) \\ &+ \sum_{i=1}^N (p_i + h_i)B_i(S_i, Q) + \sum_{i=1}^N \pi_i \lambda_i P_i(S_i, Q). \end{aligned} \quad (1)$$

3.2 (\mathbf{s}, Q) policy

In the (\mathbf{s}, Q) policy, a joint replenishment order of size Q is triggered when the inventory position of an item falls to its reorder point s_i , where $\mathbf{s} = (s_1, s_2, \dots, s_N)$ denotes the vector of the reorder points of items. The total order size Q is then allocated to the items accordingly. Unlike the (Q, \mathbf{S}) policy, there is no fixed order-up-to point in the (\mathbf{s}, Q) policy. Hence, inventory positions of items at each replenishment epoch are not necessarily equal to each other even if all items are identical. Figure 1 illustrates a snapshot of the inventory levels and positions of two identical items that are controlled by the (\mathbf{s}, Q) policy. In what follows, we present a modeling framework which differs from that of (Q, \mathbf{S}) policy since the inventory positions of items do not form a regenerative process in the (\mathbf{s}, Q) policy.

The inventory positions of items under the (\mathbf{s}, Q) policy can be modeled by a continuous-time Markov chain. Let $\mathbf{IP} = (IP_1, IP_2, \dots, IP_N)$ be the vector of inventory positions at some

Figure 1: How the proposed (\mathbf{s}, Q) policy functions



given time. A continuous-time Markov chain can be constructed by setting the \mathbf{IP} vectors as the state variables. The system leaves its current state whenever one of the items observes one unit of demand. Since the demand of each item is Poisson with rate λ_i , the inter-arrival times of demand realizations for each item is exponentially distributed with the same rate. Hence, the time that the system stays at any given state is given by the minimum of these inter-arrival times, and thus, is exponentially distributed with rate $\lambda_0 = \sum_{i=1}^N \lambda_i$. Moreover, the probability that the minimum is due to item j (or that the system leaves the current state due to item j) is λ_j/λ_0 . Thus, the rate of leaving any state \mathbf{IP} due to item j is simply given by


$$\frac{\lambda_j}{\lambda_0} \times \lambda_0 = \lambda_j.$$

First note that $s_i < IP_i$ for all $i \in \{1, \dots, N\}$ since the policy authorizes to place a joint order whenever the IP of an item hits its reorder level, s_i . There are two types of transitions in this Markov chain. First, in a regular transition, if the current state is $\mathbf{IP} = (IP_1, IP_2, \dots, IP_N)$, the system moves to state $(IP_1, IP_2, \dots, IP_j - 1, \dots, IP_N)$ if an item j with $IP_j > s_j + 1$ observes a demand. The second type of transition, a replenishment transition, happens when an item j with $IP_j = s_j + 1$ observes a demand. In this case, a joint order of size Q is placed and the system moves to a state that is fully characterized by the allocation policy that allocates the joint order Q to multiple items in the system.

Hence, for any given allocation policy, one can define a state transition function that will yield possible states that can be visited next and their corresponding transition rates. In this study, we propose a policy that allocates the joint order to the items in a way that their inventory positions in excess of their reorder points (i.e. $IP_i - s_i$) are equalized to the extent possible. An example state transition diagram for $N = 2$ and $Q = 4$ is given in Figure 2. Note that if all items are identical in terms of demand rate and cost parameters, then this policy is the optimal allocation policy. However, if items are non-identical, then this policy is merely a heuristic allocation policy. Better allocation mechanisms can be applied in this context (see, for example, Federgruen and Zipkin (1984)), however we do not attempt to explore other allocation mechanisms in this paper, as our major objective is to study the impact of the individual control provided by this new replenishment policy and the heuristic has a sufficiently good performance to highlight this impact as shown in Section 3.3.

We now make two crucial observations about this Markov chain.

Figure 2: State transition diagram when $N = 2, Q = 4$.



Observation 1 Under the allocation policy proposed, $s_i < IP_i \leq s_i + Q$ for all $i \in \{1, \dots, N\}$ and the resulting Markov chain has Q^N states.

Observation 2 *The resulting Markov chain is ergodic and the steady state probabilities are independent of the reorder levels.*

Let $\pi(s_1 + k_1, s_2 + k_2, \dots, s_N + k_N)$ be the steady state probability that the system is in state $(IP_1, IP_2, \dots, IP_N)$ where $IP_i = s_i + k_i$ for all i . Due to Observation 2, $\pi(s_1 + k_1, s_2 + k_2, \dots, s_N + k_N) = \pi(k_1, k_2, \dots, k_N)$. Hence, it is sufficient to solve the Markov chain once with states (k_1, k_2, \dots, k_N) for $k_i \in \{1, \dots, Q\}$ for all i in order to calculate the steady state probabilities for the inventory positions. Let $\pi(k_i, \mathbf{k}_{-i}) = \pi(k_1, \dots, k_{i-1}, k_i, k_{i+1}, \dots, k_N)$ where $\mathbf{k}_{-i} = (k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_N)$. Then, for any given vector $\mathbf{s} = (s_1, s_2, \dots, s_N)$ we have:

$$P\{IP_i = s_i + k_i\} = \sum_{\mathbf{k}_{-i}} \pi(k_i, \mathbf{k}_{-i}).$$

Expected Total Cost Function

Knowing the steady state probabilities of the inventory positions, the steady state probabilities of the inventory levels of each item can be found by

$$P\{IL_i = x\} = \sum_{y=\max\{x, s_i+1\}}^{s_i+Q} P\{IP_i = y\} P\{D(L_i) = y - x\}$$

where $D(L_i)$ is the demand during the lead time, which is Poisson distributed with rate $\lambda_i L_i$. Then expected total cost of ordering, inventory carrying, and backordering can be written as follows:

$$\begin{aligned} ETC(\mathbf{s}, Q) &= \frac{K\lambda_0}{Q} + \sum_{i=1}^N \sum_{x=1}^{s_i+Q} h_i x P(IL_i = x) \\ &+ \sum_{i=1}^N \sum_{x=1}^{\infty} p_i x P(IL_i = -x) + \sum_{i=1}^N \sum_{x=1}^{\infty} \pi_i \lambda_i P(IL_i = -x). \end{aligned} \quad (2)$$

After finding the steady state probabilities, we search for the optimal \mathbf{s} and Q values to minimize the expected total cost. For the case of capacitated trucks, we only search the set $\{1, 2, \dots, C\}$ for the optimal value of Q , since delaying the shipment of a fully loaded truck cannot be optimal under the (\mathbf{s}, Q) policy (this is true for any continuous review policy). All of our numerical studies show that the expected total cost function is convex in \mathbf{s} for a given Q . Similarly, the expected total cost function is observed to be convex in Q for a given \mathbf{s} .

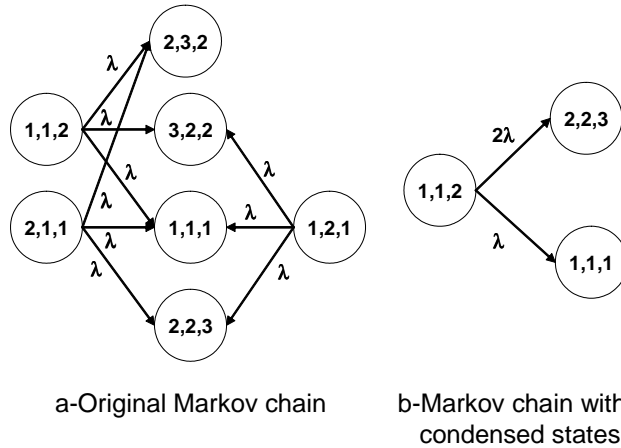
Reduction of the State Space

In the above representation of the system, the Markov chain has Q^N states. In this part, we propose an equivalent Markov chain with considerably reduced number of states with a redefinition of the state variables when the items have identical demand rates. We first make the following observation.

Observation 3 *Let $\mathbf{IP}^1, \mathbf{IP}^2, \dots, \mathbf{IP}^K$ be the states of the original system whose elements are exactly the same but sequenced in a different order. Then $\pi(\mathbf{IP}^1) = \pi(\mathbf{IP}^2) = \dots = \pi(\mathbf{IP}^K)$ if the demand rate of all items are equal.*

We call such states a ‘group’. Each group can be represented by (or condensed to) a single state whose elements are sorted in the ascending order. In particular, a state $\overline{IP} = (IP_1, IP_2, \dots, IP_N)$ is a condensed state if $IP_1 \leq IP_2 \leq \dots \leq IP_N$. In the original state transition diagram, from each state of a group, there are connections to the states of some other groups with the same transition rates. Hence we can reduce the state space of the original system by constructing a continuous time Markov chain based on only the condensed states. As an example, let $N = 3, Q = 4$, and consider states $(1, 1, 2), (1, 2, 1)$ and $(2, 1, 1)$. Apparently, these three states form a group and the condensed state is $(1, 1, 2)$. Figure 3a depicts the states emancipating from these three states in the original system.

Figure 3: State transitions out of states $(1, 1, 2), (1, 2, 1)$ and $(2, 1, 1)$



Note that the system moves to state $(1, 1, 1)$ with rate λ from each of these states and moves to any two states of the group $((2, 2, 3), (2, 3, 2), (3, 2, 2))$ with rate 2λ . Hence, we can represent all these transitions with a fewer number of transitions as depicted in Figure 3b,

by using only the condensed states. Note that the transition diagram of Figures 3a and 3b are equivalent to each other.

The following proposition characterizes the number of states in the “reduced system” which is obtained by the condensed states only.

Proposition 1 *The total number of states in the “reduced system” is given by*

$$\sum_{k=1}^N \binom{Q}{k} \binom{N-1}{k-1}. \quad (3)$$

Proof: Any given state of the reduced system may have k distinct IP values where $k \in \{1, 2, \dots, \min(N, Q)\}$. First, note that there are $\binom{Q}{k}$ possible groups with k distinct IP values. For each of these groups, let $f(N, k)$ be the total number of states with k distinct IP values. When the IP of the first item is different than the others, i.e. $i_1 \neq i_j$ for all $j \in \{2, \dots, N\}$, then there are $f(N-1, k-1)$ state combinations due to the definition of the function f . Similarly when the IP of the first two items are the same but different than the other IP values, i.e. $i_1 = i_2$ and $i_1 \neq i_j$ for all $j \in \{3, \dots, N\}$, then there are $f(N-2, k-1)$ state combinations. In a similar fashion, we can obtain the following recursion:

$$f(N, k) = f(N-1, k-1) + f(N-2, k-1) + \dots + f(k-1, k-1). \quad (4)$$

Since this recursion is valid for any N and k value, we can rewrite equation (4) as

$$f(N, k) = f(N-1, k-1) + f(N-1, k). \quad (5)$$

Noting that $f(1, 1) = 1$, $f(2, 1) = 1$ and $f(2, 2) = 1$, we observe that $f(N, k)$ is equal to the k th element of the N th row of Pascal’s Triangle. Hence $f(N, k) = \binom{N-1}{k-1}$. Multiplying this number with total number of combinations with k distinct values and adding them up from 1 to N , we obtain equation (3). \square

4. Performance of the (s, Q) policy

In this section, we compare performance of the proposed (s, Q) policy to the (Q, \mathbf{S}) policy under a variety of settings through a computational study. The comparison is based on the optimal expected total cost rates of the two policies for several problem instances with

different parameters. For each of the two policies and for every problem instance considered, we find the optimal solutions by enumerating a sufficiently wide range of parameter values and evaluating the expected total cost rates for these values. The expected total cost rates for the (Q, \mathbf{S}) and (\mathbf{s}, Q) policies are given in (1) and (2), respectively. We use a standard non-iterative sparse solver (Super LU solver version 3.0, Demmel et al. (1999)) to solve the linear equations system and derive the steady state probabilities that are used in the expected total cost rate expression given in (2). The demand rates are identical throughout the analysis in this section. Hence, solving the linear system for the reduced Markov chain once for each Q and N in consideration was sufficient.

Let $TC_{(\mathbf{s}, Q)}^*$ and $TC_{(Q, \mathbf{S})}^*$ denote the optimal expected total cost rates of the (\mathbf{s}, Q) and (Q, \mathbf{S}) policies, respectively. We define the following measure to evaluate the relative performance of the (\mathbf{s}, Q) policy over (Q, \mathbf{S}) :

$$\Delta = \frac{TC_{(Q, \mathbf{S})}^* - TC_{(\mathbf{s}, Q)}^*}{TC_{(Q, \mathbf{S})}^*} \times 100.$$

Hence, a positive value of Δ for a given problem instance indicates that the (\mathbf{s}, Q) policy performs better than the (Q, \mathbf{S}) policy; a negative value indicates a superior performance of the (Q, \mathbf{S}) policy.

4.1 Comparison with exogenous Q

We first compare the performance of the policies assuming a constant Q . This reflects settings where the batch sizes are exogenously set due to a technological constraint and/or significantly large fixed costs lead to fully utilizing fixed capacity vehicles. For now, we assume that the items are identical. The total demand rate λ is 10 and it is equally split to N items. We study three values for N : 2, 4 or 6. Throughout the section, the inventory holding cost per unit per time h is taken as 6. We use the backorder cost per unit per unit of time p to be zero throughout the analysis. The remaining parameters take one of the following values: $\pi \in \{50, 100, 200\}$ and $L \in \{0.25, 0.50\}$. We study four different values for Q : 5, 10, 15 and 20. With these parameters, a total of 72 instances are solved for each policy.

Table 1 presents the results for $N = 2$. The fourth column is the optimal S value for the (Q, \mathbf{S}) policy for the given Q and the fifth and sixth columns are the corresponding inventory holding and backorder cost rates. The seventh column is the optimal s value for the (\mathbf{s}, Q)

policy for the given Q and the eighth and ninth columns are the corresponding inventory holding and backorder cost rates. Since, Q is exogenous, the setup costs are excluded from the analysis and Δ in the tenth column reports the performance of the (s, Q) policy with respect to the sum of inventory holding and backorder costs.

Table 1: Comparison of (Q, \mathbf{S}) and (s, Q) for $N = 2$ with exogenous Q

Q	L	π	(Q, \mathbf{S}) policy			(s, Q) policy			Δ
			S^*	Holding	Backorder	s^*	Holding	Backorder	
5	0.25	50	6	45.11	12.82	3	51.69	5.37	1.48
		100	7	57.02	7.06	3	51.69	10.75	2.56
		200	7	57.02	14.12	4	63.67	4.86	3.69
	0.50	50	8	54.16	12.93	5	60.73	6.90	-0.82
		100	9	66.05	9.17	5	60.73	13.80	0.91
		200	10	78.01	5.83	6	72.68	8.79	2.83
10	0.25	50	9	66.08	7.27	2	59.27	9.79	5.84
		100	9	66.08	14.53	3	71.20	5.28	5.11
		200	10	78.02	9.28	3	71.20	10.57	6.33
	0.50	50	10	63.25	15.95	4	68.30	9.56	1.70
		100	11	75.09	13.43	5	80.22	6.83	1.67
		200	12	87.03	10.20	5	80.22	13.65	3.45
15	0.25	50	11	75.18	12.33	2	77.54	6.44	4.03
		100	12	87.06	10.06	2	77.54	12.89	6.89
		200	13	99.02	7.18	3	89.50	6.97	9.16
	0.50	50	12	72.41	20.41	3	74.72	15.42	2.89
		100	13	84.17	19.87	4	86.57	12.52	4.76
		200	14	96.07	17.61	5	98.51	8.96	5.45
20	0.25	50	13	84.35	17.73	1	83.56	14.43	4.01
		100	14	96.14	17.20	2	95.44	9.59	7.33
		200	15	108.05	14.95	3	107.41	5.19	8.45
	0.50	50	14	81.62	25.14	3	92.57	11.41	2.61
		100	16	105.13	13.74	4	104.46	9.28	4.32
		200	17	117.05	12.69	4	104.46	18.56	5.18

Note that the (s, Q) policy outperforms the (Q, s) policy in all but one problem instances for $N = 2$. The backorder cost π has a significant impact on the performance gap between the two policies. Recall that the (Q, \mathbf{S}) policy does not provide an individual control over the items, but controls the system only in aggregate terms. Therefore, even if the inventory of a particular item is dangerously low, this policy does not replenish that item if the total demand since the last order has not reached Q . However with the (s, Q) policy, each item can trigger a replenishment order individually. This difference becomes more important as the unit backorder cost π is higher. Thus, Δ generally increases as π increases (except when

$Q = 10$).

Table 1 also shows that the Δ values for $L = 0.25$ are always higher than Δ values for $L = 0.50$. The reason behind this result is again the individual control that is possible with the (\mathbf{s}, Q) policy. Since replenishment occurs quickly due to short lead time, in both of the policies less inventory is kept and less backorder occurs. Therefore, the sum of the backorder and holding costs of both policies decreases as the lead times decrease. However, the reduction in total backorder and holding costs of the (\mathbf{s}, Q) policy is more than that of the (Q, \mathbf{S}) policy because of the effective control at the individual level.

Table 2 shows the results for $N = 4$ and $N = 6$. The fourth column represents the optimal S value, the fifth column represents the corresponding total cost for the (Q, \mathbf{S}) policy for $N = 4$. The sixth column represents the optimal s value and the seventh column represents the performance of the (\mathbf{s}, Q) policy for $N = 4$. Columns 8–12 are similarly defined for $N = 6$.

We see in Table 2 along with the results in Table 1 that the relative performance of the (\mathbf{s}, Q) policy compared to the (Q, \mathbf{S}) deteriorates as N increases for all instances. This is due to the fact that the ability to trigger constant size orders individually (the major advantage of the (\mathbf{s}, Q) policy) is less important when more items are involved. However we also see that for sufficiently large backorder costs and sufficiently small lead times, the (\mathbf{s}, Q) policy still provides significant improvements over the (Q, \mathbf{S}) policy even when N is large. The impact of lead times and the backorder costs remain unchanged for $N = 4$ and $N = 6$. For all N values, the batch size Q does not seem to have an impact on the relative performance of the two policies.

4.2 Comparison with endogenous Q

The analysis so far assumes that Q is exogenous. We now study the case where each batch incurs a setup cost and order size Q also is a decision variable. The setup cost corresponds to the fixed cost of a trip that a truck makes to replenish the multiple items. We also assume that there is an upper bound C on the value that Q can take. This reflects the finite capacity of the transportation vehicle (e.g., truck). For a given capacity, C , one must search the set $\{1, 2, \dots, C\}$ for the optimal value of Q in both policies. In Table 3, we provide a comparison of the (\mathbf{s}, Q) and (Q, \mathbf{S}) policies for $N = 2$ and $K \in \{100, 200\}$. The capacity is set at $C = 20$. In Table 3, the fourth and the fifth column shows the optimal Q and S values in the (Q, \mathbf{S}) policy. The sixth column is the corresponding optimal total cost rate including

Table 2: Comparison of (Q, \mathbf{S}) and (\mathbf{s}, Q) policies for $N = 4$ and $N = 6$ with exogenous Q

Q	L	π	$N = 4$				$N = 6$			
			S^*	Total Cost	s^*	Δ	S^*	Total Cost	s^*	Δ
5	0.25	50	4	82.80	1	0.80	3	101.59	1	-5.07
		100	4	96.43	2	2.40	3	121.89	1	5.17
		200	5	104.11	2	3.08	4	131.17	1	-1.60
	0.50	50	5	94.63	2	-0.24	4	117.74	1	-0.12
		100	6	110.95	3	2.47	4	133.18	2	-1.27
		200	6	119.85	3	-0.16	5	151.64	2	0.55
10	0.25	50	5	99.34	1	0.56	4	121.04	0	1.41
		100	6	114.61	1	5.83	4	139.67	1	0.85
		200	6	127.12	2	5.43	5	156.22	1	5.79
	0.50	50	6	107.88	2	-1.81	4	129.99	1	-1.92
		100	7	125.07	2	3.31	5	149.84	1	-0.01
		200	7	139.00	3	3.37	6	173.30	2	3.81
15	0.25	50	6	115.21	0	1.12	4	137.22	0	-0.87
		100	7	132.18	1	5.25	5	157.98	0	2.07
		200	8	148.85	1	7.85	6	179.81	1	5.89
	0.50	50	7	121.59	1	0.46	5	142.64	0	-2.11
		100	8	140.45	2	2.65	6	168.47	1	1.79
		200	9	158.66	2	4.77	6	192.57	1	2.47
20	0.25	50	7	130.09	0	0.25	5	149.91	0	-6.81
		100	8	149.94	1	3.63	6	176.49	0	2.49
		200	9	167.55	1	8.50	7	202.17	1	4.85
	0.50	50	8	135.11	1	-1.58	6	158.39	0	-1.68
		100	9	156.27	1	1.75	6	186.58	1	0.60
		200	10	175.55	2	5.56	7	210.09	1	3.97

the setup costs. The seventh and eighth column shows to the optimal s and Q values in the (\mathbf{s}, Q) policy. The ninth column is the corresponding optimal total cost rate including the setup costs.

Table 3: Comparison of the (Q, \mathbf{S}) and (\mathbf{s}, Q) policies for $N = 2$ with endogenous Q

K	L	π	(Q, \mathbf{S}) policy			(\mathbf{s}, Q) policy			Δ
			Q^*	S^*	Total Cost	s^*	Q^*	Total Cost	
100	0.25	50	18	12	151.86	1	19	147.83	2.65
		100	18	13	162.45	2	18	154.59	4.84
		200	16	13	171.33	3	17	161.65	5.65
	0.50	50	19	14	156.07	3	19	153.66	1.54
		100	18	15	168.26	4	18	163.28	2.96
		200	18	16	178.73	5	18	172.30	3.60
200	0.25	50	20	13	202.08	1	20	197.99	2.02
		100	20	14	213.34	2	20	205.03	3.90
		200	20	15	223.00	3	20	212.60	4.66
	0.50	50	20	14	206.76	3	20	203.98	1.35
		100	20	16	218.87	4	20	213.74	2.34
		200	20	17	229.74	4	20	223.02	2.92

In all 12 instances, the (\mathbf{s}, Q) policy performs better than the (Q, \mathbf{S}) policy. For $K = 100$, all Q^* values are below the capacity C . For $K = 200$, both policies set the batch size at the capacity. For both setup costs, the (\mathbf{s}, Q) policy again performs better with higher backorder cost and smaller lead time; this time without any exceptions. The relative performance of the (\mathbf{s}, Q) policy is smaller when $K = 200$. This effect is due to the fact that higher setup costs tend to produce larger ordering and inventory holding costs, and this significantly diminishes the impact of the differences in backorder costs for which the policy choice is important.

A similar computational experiment is conducted for $N = 4$ and $N = 6$ and the results are shown in Table 4. As for the case of $N = 2$, the (\mathbf{s}, Q) policy performs relatively better for larger backorder costs and smaller lead times. In parallel to the results in Section 4.1, as N grows the performance of the (\mathbf{s}, Q) policy deteriorates. However, as before, for sufficiently large backorder costs and sufficiently small lead times, the (\mathbf{s}, Q) policy still offers significant savings over the (Q, \mathbf{S}) policy. We also note that the optimal Q value obtained under one policy is not necessarily smaller or larger than the optimal Q value in the other policy. This observation (together with the fact that Δ is independent of Q found in Section 4.1) tells us that the vehicle capacities do not necessarily favor one policy over the other.

Table 4: Comparison of the (Q, \mathbf{S}) and (\mathbf{s}, Q) policies for $N = 2$ and $N = 4$ with endogenous Q

K	L	π	$N = 4$						$N = 6$					
			Q^*	S^*	Total Cost	s^*	Q^*	Δ	Q^*	S^*	Total Cost	s^*	Q^*	Δ
100	0.25	50	19	7	179.81	0	18	0.64	19	5	199.88	0	15	-2.60
		100	16	7	198.45	1	16	3.49	15	5	224.65	0	17	2.05
		200	16	8	214.11	1	18	5.61	16	6	245.62	1	15	3.96
	0.50	50	20	8	185.11	1	17	-0.49	17	5	206.93	0	18	-1.42
		100	16	8	205.71	2	17	1.41	17	6	232.55	1	16	0.33
		200	17	9	223.08	2	18	3.59	18	7	257.01	1	18	2.31
200	0.25	50	20	7	230.09	0	20	0.14	20	5	249.91	0	20	-4.09
		100	20	8	249.94	1	20	2.18	20	6	276.49	0	20	1.59
		200	20	9	267.55	1	20	5.32	20	7	302.17	1	20	3.25
	0.50	50	20	8	235.11	1	20	-0.91	20	6	258.39	0	20	-1.03
		100	20	9	256.27	1	20	1.07	20	6	286.58	1	20	0.39
		200	20	10	275.55	2	20	3.54	20	7	310.09	1	20	2.69

4.3 Non-identical items

We now investigate the case of non-identical items. For simplicity of the computational analysis, we assume that the items have identical demand rates, but differ in their backorder costs. When the items are non-identical, a policy is necessary to allocate the incoming replenishment order (of size Q) to multiple items based on their inventory positions. As discussed in Section 3, we use a simple allocation policy for this purpose. The incoming order is allocated in such a way that the inventory positions above the re-order points (s_i) are equalized to the extent possible. This is obviously a heuristic for this problem since the differences in item characteristics are taken into account only when setting different re-order points.

The results of the computational study for the non-identical case is given in Table 5 for $N = 2$ and $N = 4$. We use a single value for the lead time $L = 0.50$. For each N , we study a number of backorder cost vectors $\boldsymbol{\pi}$ that reflect the differences in backorder costs among items. For example, for $N = 2$, $\boldsymbol{\pi}$ takes three values $(25, 175)$, $(50, 150)$ and $(100, 100)$; the last one corresponding to the identical case.

The results in Table 5 shows that in 21 of 24 instances, the (\mathbf{s}, Q) policy outperforms the (Q, \mathbf{S}) . This is despite the fact that the (\mathbf{s}, Q) policy uses a heuristic allocation policy. It is expected that more advanced allocation mechanisms will yield to larger Δ values. We

Table 5: Comparison of (Q, \mathbf{S}) and (\mathbf{s}, Q) for non-identical items with exogenous Q

N	Q	π	(Q, \mathbf{S}) policy		(\mathbf{s}, Q) policy		Δ
			\mathbf{S}^*	Total Cost	\mathbf{s}^*	Total Cost	
2	5	25,175	7,9	70.38	4,6	69.44	1.34
		50,150	8,9	73.45	5,6	73.45	-0.01
		100,100	9,9	75.22	5,5	74.53	0.91
	10	25,175	9,12	82.32	3,5	80.25	2.52
		50,150	10,12	86.94	4,5	84.16	3.20
		100,100	11,11	88.52	5,5	87.05	1.67
	15	25,175	11,14	95.75	2,5	93.03	2.84
		50,150	12,14	101.05	3,5	97.69	3.32
		100,100	13,13	104.04	4,4	99.09	4.76
	20	25,175	13,17	109.86	1,4	106.54	3.02
		50,150	14,16	116.25	3,4	111.18	4.36
		100,100	16,16	118.87	4,4	113.74	4.32
4	5	25,25,175,175	4,4,6,6	98.98	1,1,3,3	100.39	-1.42
		25,75,125,175	4,5,6,6	103.49	1,3,3,3	104.30	-0.79
		100,100,100,100	6,6,6,6	110.95	3,3,3,3	108.22	2.47
	10	25,25,175,175	5,5,7,7	112.93	1,1,3,3	112.71	0.20
		25,75,125,175	5,6,7,7	118.12	1,2,2,3	116.82	1.11
		100,100,100,100	7,7,7,7	125.07	2,2,2,2	120.93	3.31
	15	25,25,175,175	6,6,8,8	128.32	0,0,2,2	125.61	2.11
		25,75,125,175	6,7,8,8	134.04	0,1,2,2	130.93	2.32
		100,100,100,100	8,8,8,8	140.45	2,2,2,2	136.73	2.65
	20	25,25,175,175	7,7,10,10	142.29	0,0,2,2	140.58	1.20
		25,75,125,175	7,9,9,10	149.28	0,1,2,2	146.11	2.12
		100,100,100,100	9,9,9,9	156.27	1,1,1,1	153.54	1.75

Table 6: Comparison of (Q, \mathbf{S}) and (\mathbf{s}, Q) for non-identical items with endogenous Q

N	K	π	(Q, \mathbf{S}) policy			(\mathbf{s}, Q) policy			Δ
			Q^*	\mathbf{S}^*	Total Cost	\mathbf{s}^*	Q^*	Total Cost	
2	100	25,175	19	12,16	159.64	1,4	20	156.54	1.94
		50,150	19	14,16	165.37	3,4	19	160.93	2.69
		100,100	18	15,15	168.26	4,4	18	163.28	2.96
	200	25,175	20	13,17	209.86	1,4	20	206.54	1.58
		50,150	20	14,16	216.25	3,4	20	211.18	2.35
		100,100	20	16,16	218.87	4,4	20	213.74	2.34
4	100	25,25,175,175	18	6,6,9,9	191.55	0,0,2,2	18	189.76	0.94
		25,75,125,175	18	6,8,9,9	198.06	0,1,2,2	18	195.22	1.43
		100,100,100,100	16	8,8,8,8	205.71	2,2,2,2	17	202.81	1.41
	200	25,25,175,175	20	7,7,10,10	242.29	0,0,2,2	20	240.58	0.71
		25,75,125,175	20	7,9,9,10	249.28	0,1,2,2	20	246.11	1.27
		100,100,100,100	20	9,9,9,9	256.27	1,1,1,1	20	253.54	1.07

see that the asymmetry in backorder costs does not have a pronounced effect on the relative performance of the (\mathbf{s}, Q) policy.

In Table 6, we report the performance of the (\mathbf{s}, Q) policy for the case of non-identical items and when Q is endogenously determined based on a setup cost of K for each replenishment. Again a capacity C of 20 is considered. In this case, the (\mathbf{s}, Q) policy performs better than the (Q, \mathbf{S}) policy in all 12 instances.

The numerical study presented in this section shows that the (\mathbf{s}, Q) policy is viable alternative for jointly replenishing multiple items (or locations) when the setup costs are due to capacitated vehicles or manufacturing processes. Numerical results show that the (\mathbf{s}, Q) policy outperforms the (Q, \mathbf{S}) policy for sufficiently large backorder costs and sufficiently small lead times. These settings coincides with the settings that the beverage producer discussed as the primary motivation of this study in Section 1. The backorder penalties are prohibitive for this company and therefore the company maintains a very high level of service. In addition, there is ample supply in bottling facilities and most transportation lead times are less than a day which leads to small replenishment lead times. Since the transportation is carried out with fixed capacity trucks and the fuel charges are high, the company also wants to maintain full truck loads. The numerical study presented here shows that the (\mathbf{s}, Q) policy is a good candidate for this company and others that exhibit similar characteristics. The numerical study also shows that the (\mathbf{s}, Q) policy yields better results under similar settings when the items are non-identical even when a heuristic allocation

policy is employed.

5. Conclusion

In this paper, we consider the stochastic joint replenishment problem with cost structures determined by transporting the items in fixed-capacity vehicles. As far as we know, there is only one other study in the literature in which vehicle capacities are explicitly considered; the study by Cachon (2001) suggests the (Q, \mathbf{S}) policy for the continuous review case. We propose a new policy called the (\mathbf{s}, Q) policy, which creates a joint replenishment order of size Q whenever one of the items falls to its reorder point and allocates this order to multiple items based on their inventory positions. We show that the inventory positions of the items can be modeled using a continuous Markov chain.

We conducted an extensive numerical study to assess the performance of these two policies. The results show that the (\mathbf{s}, Q) policy usually outperforms the (Q, \mathbf{S}) policy especially when the backorder penalties are high and the replenishment lead times are small. Thus, we conclude that our policy would be an appropriate choice for jointly replenishing multiple items (or locations) when backorder costs or the service levels are high, lead times are small and when there are important economies of scale or capacities in transportation vehicles or the manufacturing processes that replenish these items.

The study here can be extended in several directions. First, one can consider a case where the warehouse does not have ample supply and hence also manages inventories and operates a replenishment policy. Another extension could be the incorporation of minor setup costs in addition to the major setup costs here. Another direction could be to devise different allocation rules for the non-identical case and determine how these perform against the existing policies.

References

- Aksoy, Y., S. Erengüç. 1988. Multi-item models with coordinated replenishments: A survey. *International Journal of Production Management* **8** 63–73.
- Alp, O., N. K. Erkip, R. Güllü. 2003. Optimal lot-sizing/vehicle-dispatching policies under stochastic lead times and stepwise fixed costs. *Operations Research* **51** 160–166.
- Atkins, D. R., P. O. Iyogun. 1988. Periodic versus can-order policies for coordinated multi-item inventory systems. *Management Science* **34** 791–796.
- Balintfy, J. L. 1964. On a basic class of inventory problems. *Management Science* **10** 287–297.

- Ben-Khedher, N., C. A. Yano. 1994. Multi-item joint replenishment problem with transportation and container effects. *Transportation Science* **28** 37–54.
- Bostancı, E., Z. A. Cengar, B. Çağlar E. Güler, M. Subaşı, E. Çiftçi, A. Şen. 2005. Computer based supply chain optimization: MAN Turkey. *Industrial Projects*. Bilkent University, Department of Industrial Engineering, Ankara, Turkey, 79–90.
- Burt, D. N. 1989. Managing suppliers up to speed. *Harvard Business Review* July–August, 127–135.
- Cachon, G. 2001. Managing a retailer’s shelf space, inventory and transportation. *Manufacturing and Service Operations Management* **3** 211–229.
- Carbone, J. 1996. Sun shines by taking out time. *Purchasing* September 19.
- Carbone, J. 1999. Reinventing purchasing wins the medal for Big Blue. *Purchasing* September 16, 38–62.
- Çetinkaya, S., C. Y. Lee. 2000. Stock replenishment and shipment scheduling for vendor managed inventory systems. *Management Science* **16** 217–232.
- Cheung, K. L., H. L. Lee. 2002. The inventory benefit of shipment coordination and stock rebalancing in a supply chain. *Management Science* **48** 300–306.
- Cook, R. L., B. Gibson, D. MacCurdy. 2005. A lean approach to cross docking. *Supply Chain Management Review* **9** 54.
- Demmel, J. W., S. C. Eisenstat, J. R. Gilbert, X. S. Li, J. W. H. Liu. 1999. A supernodal approach to sparse partial pivoting. *SIAM J. Matrix Analysis and Applications* **20**(3) 720–755.
- Federgruen, A., P. Zipkin. 1984. Allocation policies and cost approximations for multi-location inventory systems. *Naval Research Logistics* **37** 97–130.
- Genna, A. 1997. How Merck leverages supply for profit. *Purchasing* September 4.
- Gürbüz, M. C., K. Moinzadeh, Y. P. Zhou. 2007. Coordinated replenishment and shipping strategies in inventory/distribution systems. *Management Science* **53** 293–307.
- Khouja, M., S. Goyal. 2007. A review of the joint replenishment problem literature: 1989-2005. *European Journal of Operational Research* In press.
- Minahan, T. 1997. JIT – a process with many faces. *Purchasing* **123,3** 42–46.
- Nielsen, C., C. Larsen. 2005. An analytical study of the $Q(s, S)$ policy applied to the joint replenishment problem. *European Journal of Operational Research* **163** 721–732.
- Özkaya, B. Y., Ü. Gürler, E. Berk. 2006a. The stochastic joint replenishment problem: a new policy, analysis, and insights. *Naval Research Logistics* **53** 525–546.
- Özkaya, B. Y., Ü. Gürler, E. Berk. 2006b. Stochastic joint replenishment problem in a two-echelon inventory system: analysis and insights. Tech. rep., Department of Industrial Engineering, Bilkent University.
- Pantumsinchai, P. 1992. A comparison of three joint ordering policies. *Decision Sciences* **23** 111–127.
- Pantumsinchai, P. 2000. E-procurement takes on the untamed supply chain. *iSource* November, 108.
- Renberg, B., R. Planche. 1967. Un modle pour la gestion simultane des n articles d’un stock. *Revue Francaise d’Informatique et de Recherche Oprationelle* **6** 47–59.
- Silver, E. A. 1965. Some characteristics of a special joint-order inventory model. *Operations Research* **13** 319–322.
- Silver, E. A. 1974. A control system for coordinated inventory replenishment. *International Journal of Production Research* **12** 647–671.

- Simchi-Levi, D., P. Kaminsky, E. Simchi-Levi. 2003. *Designing & managing the supply chain*. McGraw-Hill.
- Talluri, S., R. Narasimhan. 2005. A note on “A methodology for supply base reduction”. *IEEE Transactions on Engineering Management* **52** 130–139.
- Viswanathan, S. 1997. Periodic review (s, S) policies for joint replenishment inventory systems. *Management Science* **43** 1447–1454.