Economic Order Quantity for Business-to-Consumer Fulfillment

Diego Klabjan
Department of Industrial Engineering and Management Sciences
Northwestern University
Evanston, IL
d-klabjan@northwestern.edu

August 5, 2009

Abstract

Order fulfillment is vital for successful business-to-consumer e-commerce firms. For an e-tailer, customers place orders from various geographically disperse locations, or, in the presence of multiple distribution channels, from various channels. We consider a constant demand rate for each one of these markets in the continuous time infinite time horizon setting. The distribution cost for shipping goods from fulfillment centers to markets, the variable and fixed procurement costs, and the holding costs at the fulfillment centers are captured. Demand from a market can be fulfilled from several fulfillment centers. We first study the case of two centers and a single market. We show that only one of the centers is used by an optimal policy. If, in addition, each facility has a market that must use an assigned center, this might no longer be the case. An optimal policy might use one center for a period of time and the other one for a different period of time. We develop simple cyclic policies that might outperform the strategy of using a single center. We also study the general case of several centers and markets. We give properties of optimal policies and we develop a lower bound. To this end, we first develop a Lagrangian type lower bound for general weakly coupled semi-Markov decision processes and then we apply this bound to our problem. Computational experiments show that the lower bound is often tight. We also numerically study the benefit of cyclic policies.

1 Introduction

In late nineties we witnessed a boom in e-commerce. The proliferation of internet led to new business concepts such as business-to-business exchanges, portals, e-procurement, online auctions, and business-to-consumer strategies. There are many firms that offer business-to-business services, e.g. ChemConnect for exchanges within the chemical industry, Dell uses its ValueChain.Dell.com system for rapid exchanges with its suppliers. Especially the latest paradigm of business-to-consumer had the most significant influence on everyday life of consumers since it involves them directly. In addition, online retailing is rising quickly; from $8 billion in revenues generated by the US retailers in 1998 to $90 billion in 2004, Grosso *et al.* (2005). Many firms use internet to sell directly to consumers, i.e. they use the direct channel. While direct sales are not confined to internet (catalog and mail sales exist for a long time), they definitely became the most widespread direct channel. The pioneers in this direction are the book retailer Amazon.com and the computer manufacturer Dell who sell only through the direct online channel.

On the other hand there are several firms that use several distribution channels. Barnes & Noble, another book retailer, uses its stores as a distribution channel and it offers a direct channel through their own web site. Several firms, e.g. Barnes & Noble, Gap, Best Buy, Levi Strauss & Co., used or are using both channels: the direct channel and the reseller channel. The direct channel also induces the so-called channel conflicts, which most often come in the form of undercutting the traditional reseller channel. Some companies use the direct channel as a mean of revenue while other companies are using it to induce sells in their reseller channel. The profile of online shoppers is typically different than the profile of traditional shoppers. As an example, online book buyers tend to buy many more books and are more time sensitive than the traditional book buyers.

To boost profitability, order fulfillment processes are critical. Amazon.com and Dell are two examples of efficiency machines when it comes to fulfillment. When an online order is placed, the firm faces the decision from which fulfillment center to ship. Fulfillment centers can be stores and warehouses owned by the company or outsourced facilities to carriers such as FedEx Express, DHL, and UPS. For example, Amazon.com has 19 fulfillment centers and the fulfillment decision factors in real-time order data and ship dates in order to develop optimal pick, pack, and ship processes. Which fulfillment center to use is an important decision especially when the shipping cost is covered by the company. This is usually the case when the customer selects the "standard" shipping option.

We study the inventory and fulfillment policies for the continuous time infinite time horizon problem with deterministic demand rates, i.e. the economic order quantity setting. A firm must make the procurement decision at any point in time during the time horizon and for each of the fulfillment centers. The demand rate is given for each sales location or market, which are geographically dispersed. In the case of multiple channel operations, each channel can have several sales locations. Pure e-tailers have only a single direct channel and therefore sales locations can be identified with geographic markets. In addition to the procurement cost, which includes the variable and the fixed component, the firm incurs the shipping cost, which depends on the sales location. Besides the procurement decision, at any point in time a decision is made for each center what fraction of a market's demand to fulfill. Each center incurs a linear inventory holding cost. In a manufacturing context, the procurement process is replaced by manufacturing, i.e. a decision is how many units to manufacture.

In the business-to-consumer fulfillment setting, consumers are very segmented. Our model is applicable at an aggregated level where consumers are aggregated based on demographics or channels. The presented model is general enough to handle also business-to-business situations. Consider, for example, a semiconductor manufacturer supplying a large computer manufacturer with several assembly plants. Each assembly plant can be served from various semiconductor plants. The problem faced by the semiconductor manufacturer is to select a set of its own plants to satisfy the demand from the assembly plants of the computer manufacturer.

In this paper we lay down the modeling framework for such infinite horizon planning. We start by studying the special case of two centers and a single market. Not surprisingly, in this case we show that it is optimal to fulfil all of the demand from a single center. Next we consider a similar case except that we assume that each center has also a market that must be satisfied only by this center (local markets). We show that using only a single center to fulfill all of the non local demand might not be optimal. We develop a cyclic policy that in some cases performs better. In such a policy for a certain period of time one of the two centers is used to fulfill the non local demand and in other times the other center is used. We also asses the gap between such policies and the policy of using only a single center for non local demand. We also consider the general case of several centers and markets. We first provide some structural properties of optimal policies. In order to develop a lower bound, we first model the problem as a semi-Markov decision process that is weakly coupled, i.e. only some constraints in the action space link various independent semi-Markov decision processes. Based on the Lagrangian relaxation principle, we develop a lower bound for such processes. This lower bound is then applied to our problem. While computing an optimal policy is very hard and we do not know of any efficient algorithms, the derived lower bound results in a relatively simple optimization problem that can easily be solved by standard optimization software tools. While computing an optimal policy is At the end we perform numerical experiments. They show that the lower bound is often tight. We also document the benefits of using the developed cyclic policies.

There are several important contributions of this work.

- We believe this is the first work on studying the problem with several fulfillment centers and markets in an economic order quantity setting. To this end, we provide the model and we analyze selected special cases, e.g., when is using a single center optimal.
- The developed cyclic policies are the second contribution. These policies are easy to implement.

- Another important contribution is in studying the weakly coupled semi-Markov decision processes. We give a lower bound in a very general setting. The approach is much more involved than the previously studied weakly coupled Markov decision processes. To the best of our knowledge, this is the first study of weakly coupled semi-Markov decision processes.

- Finally, it is a nontrivial and very technical task to use this semi-Markov bound in the context of business-to-consumers fulfillment.

We start the presentation with the problem statement in Section 2. The case with two facilities is considered in Section 3. The case with a single market is studied in Section 3.1. Two centers each one with a local market and a single non local market is discussed in Section 3.2. Properties of optimal policies are given in Section 4. The lower bound is developed in Section 5 and the computational experiments are given in Section 6. We finish the introduction with a literature review.

Literature Review

Most of the related research is related to the multi-channel studies of firms. Many authors studied decentralized systems with multiple channels. The main question is under what conditions is it beneficial to establish a direct channel. Another important question is the pricing strategies for the direct and the reseller channel, i.e. should the selling price be the same and if not, how to set it up. Tsay and Agrawal (2000), Chiang *et al.* (2003), and Tsay and Agrawal (2004) study cons and pros of using both channels in a decentralized system. They also compare three possible scenarios: the firm has only a direct channel, the firm has only the reseller channel, and the firm is using both channels. Cattani *et al.* (2003) consider also different price setting strategies between the two channels. Similarly, Boyaci and Gallego (2002) study pricing and channel profits in a single warehouse multiple store setting. Boyaci (2005) assumes that channels are differentiated based on the location and the channel related demand is substitutable. The distribution cost is not a factor. Bernstein *et al.* (2005) study the impact of setting up a channel to "taste" the product, which then hopefully induces additional reseller demand. Cattani *et al.* (2004) and Tsay and Agrawal (2004) provide recent surveys related to this line of research.

There is also limited inventory management literature in a multi-channel setting. Chiang and Monahan (2005) study the two echelon continuous review model with a single direct channel. A reseller channel consists of a warehouse that supplies a single retailer. The retailer faces exogenous stochastic demand, where costumers shop at the store. In addition, the direct demand is fulfilled from the warehouse. The authors study base stock policies. A similar system is studied by Allgor *et al.* (2004) where several heuristics are proposed for the multi-item version of the problem. Alptekinoglu and Tang (2005) study the stochastic problem with several cross-docking depots (not holding inventory) and several markets. Their model assumes stochastic demand but it is a finite horizon problem. The single period version where facilities carry inventory and the market demands are assumed to be stochastic is discussed in Klabjan (2009).

The business-to-consumer setting is also considered in Bagga *et al.* (2005). In their work a single warehouse supplies several stores, which fulfill the direct demand. They assume that a fixed order up-to-level replenishment policy is followed and they study day-to-day operations, i.e. execution planning. They do not allow demand from a location to be split among several stores, i.e. a single store must serve the entire demand from a location. They present an integer program that does this assignment.

There is vast literature on economic order quantity, i.e. single item continuous time infinite time horizon inventory problems. Many extensions to the basic model are given in Zipkin (2000). The more relevant to our model are those that embed transportation decisions. Note that linear distribution cost yields the same reorder quantity. Nonlinear distribution cost, such as those used by less than truckload carriers, is studied in Swensetha and Godfrey (2002) and Russell and Krajewski (1991). There are also several manuscripts that address production and distribution simultaneously. They are focused on operating a dedicated fleet, having a single manufacturing plant and potentially several customers, Blumenfeld *et al.* (1985), Blumenfeld *et al.* (1991), Hahm and Yano (1992), Burns *et al.* (1985). The integration of the economic order and production quantity is considered in Hall (1996).

There is similarity between the inventory routing problem and the problem studied herein. The literature on inventory routing is too vast to summarize but surveys and reviews can be found in Goyal and Satir (1989), Dror (2005), Campbell and Savelsbergh (2002), and Campbell and Savelsbergh (1998). In a single item version of the inventory routing problem, a dedicated fleet of trucks need to be dispatched from several depots to customers in an infinite time horizon setting. Consider the version of the inventory routing problem where truck routes are restricted to be single leg, there are several depots, trucks are uncapacitated, and customers do not incur any holding cost. Then this would be the problem studied here except for the following two complicating factors: (1) a customer can receive a replenishment from several depots simultaneously, and (2) there is holding cost at each depot. In addition, most of the inventory routing studies assume the periodic review setting and not the continuous time decisions. The problem studied herein is simultaneously a more restrictive version of the inventory routing model and also a generalization.

2 Problem Statement

We consider fulfillment centers or distribution facilities, for simplicity called facilities, $N = \{1, 2, \ldots, n\}$ operated by a single corporation, i.e. a centralized system, and sales locations or markets $M = \{1, 2, \ldots, m\}$ in an infinite time horizon and a single item. Replenishments and shipments can be carried out at any point in time (continuous time setting) and there is no leadtime. In a deterministic setting this is without loss of generality. The per item procurement cost of facility i is denoted by c_i. Whenever a replenishment order is placed by facility i, a fixed cost k_i is incurred. Each facility can carry inventory and let h_i be the per unit linear holding cost of facility i. Each market j has a constant deterministic demand rate L_j. At any point in time demand from market j can be simultaneously fulfilled from several facilities. The per unit distribution cost between facility i and market j is denoted by f_{ij}. This cost can, for example, be correlated to the distance between the facility and the market. No backlogging is allowed. Figure 1 depicts the materials flow. We assume that $k_i > 0, h_i > 0$ for every $i \in N$ and $L_j > 0$ for every $j \in M$. In addition, we impose $c_i \geq 0$ for every $i \in N$ and $f_{ij} \geq 0$ for every $i \in N, j \in M$.

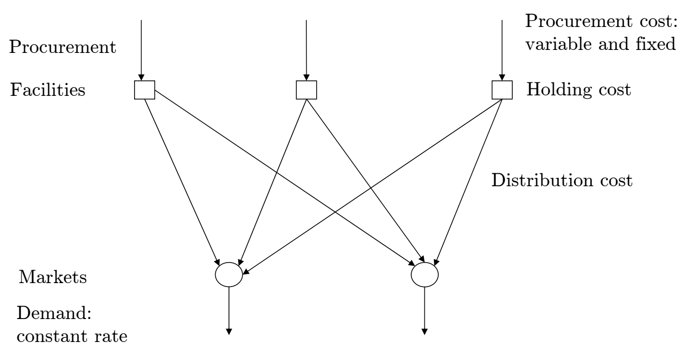

Figure 1: Configuration for $n = 3, m = 2$

Next we define policies. Let $I_i(t)$ be the inventory level at time t and facility i. Each decision epoch at time t consists of the following two types of actions at each facility.

Decision 1: Should facility i be replenished and if yes by how much?

Decision 2: What fraction of the demand of market j should each facility satisfy during the time period between now and the next decision epoch?

At each time t, let $D_{ij}(t)$ be the fraction of demand rate L_j of market j that is fulfilled from facility i between time t and the next decision epoch. The action space requirement Decision 2 from above imposes that at any point in time t we have $\sum_{i \in N} D_{ij}(t) = L_j$ for each $j \in M$. A possible trajectory for 2 facilities and a single market is depicted in Figure 2.

In addition to decisions D_{ij}, in each time period we decide upon the replenishment quantity y_i at facility i, Decision 1 above. We say that a facility has a *breakpoint* at a decision epoch if the trajectory has a breakpoint at this decision epoch, i.e. at the decision epoch at time t the total demand rate $\sum_{j \in M} D_{ij}(t^-)$ at facility i just before time t is different from the total demand rate $\sum_{j \in M} D_{ij}(t)$ at time t (or just after t).

Let us assume that the next decision epoch is at time $t_1 > t$. The total procurement cost is $\sum_{i \in N}(c_i y_i + k_i \delta(y_i))$, where $\delta(z)$ is 0 if $z = 0$ and 1 if $z > 0$. The distribution cost equals to $(t_1 - t)\sum_{j \in M, i \in N} f_{ij} D_{ij}(t)$ (on the link between facility i and market j we sell $D_{ij}(t) \cdot (t_1 - t)$ units for a per unit distribution cost f_{ij}). In addition each facility i incurs a linear holding cost with the per unit cost h_i. Note that holding cost accounting cannot be given by a simple formula since the trajectory at a facility does not have the nice saw-tooth structure.

The goal is to find a policy that minimizes the long-run average cost. It is easy to see that an optimal trajectory at each facility has the zero-inventory ordering property. Further properties of optimal policies are given in Section 4. We call this problem the *multi-market problem*.

3 The Two Facility Case

In this section we study the case with $n = 2$. In the first part we assume a single market and then we study the case when each facility has its own market whose demand must be fulfilled only from the corresponding facility.

3.1 A Single Market

Consider two facilities and a single market with demand rate L_1. The main result in this section states that it is optimal to use only a single facility.

Theorem 1. If there is a single market, there is an optimal policy where one facility serves all of the demand.

From Theorem 1, the value of the optimal policy is $\min\{(c_1+f_{11})L_1+\sqrt{2k_1L_1h_1}, (c_2+f_{21})L_1+\sqrt{2k_2L_1h_2}\}$ and the facility that attains this minimum serves all of the demand. The square root term comes from the standard economic order quantity long-run average cost and the first linear term accounts for the linear in time procurement and distribution costs. The technical proof of this result is given in Klabjan (2009). This result can be extended to the case of n facilities and a single market.

3.2 Three Markets

In this section we assume that each one of the two facilities has a market whose demand must be fulfilled from a particular facility. We can interpret the two facilities as being two retail stores. Each store serves its own demand from the geographical region corresponding to its location. In addition, the two stores must fulfill the demand from a remote market where extra distribution cost is incurred.

Formally, the remote market has a demand rate L and each facility i has its own local demand with rate L_i. (For ease of exposition we do not denote the market demand rates by L_1, L_2, L_3.) Thus the distribution cost equals to $f_{11} = f_{22} = 0, f_{12} = f_{21} = \infty$ and f_1, f_2 is the distribution cost between facility $1, 2$ and the remote market, respectively. At any time t we have $D_{11}(t) = L_1, D_{22}(t) = L_2, D_{12}(t) = D_{21}(t) = 0, D_1(t) + D_2(t) = L$. Here we denote by $D_1(t), D_2(t)$ the demand of facility 1,2 from the remote market, respectively. For ease of notation we define $\tilde{c} = c_1 + f_1 - c_2 - f_2$.

A $0/L$ policy is any policy such that in each time t either $D_1(t) = L, D_2(t) = 0$ or $D_1(t) = 0, D_2(t) = L$, and trajectory breakpoints occur only when one of the two facilities replenishes. A possible $0/L$ trajectory is depicted in Figure 3. Note that in a $0/L$ policy, at each point in time only one facility serves all of the remote demand. However, this facility might not be always the same one.

The following proposition can be proved along the same lines as the proof of Theorem 1, which is given in Klabjan (2009).

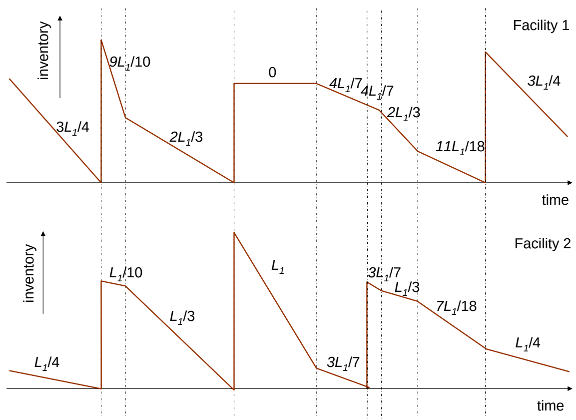

Figure 2: A possible trajectory for $n = 2, m = 1$

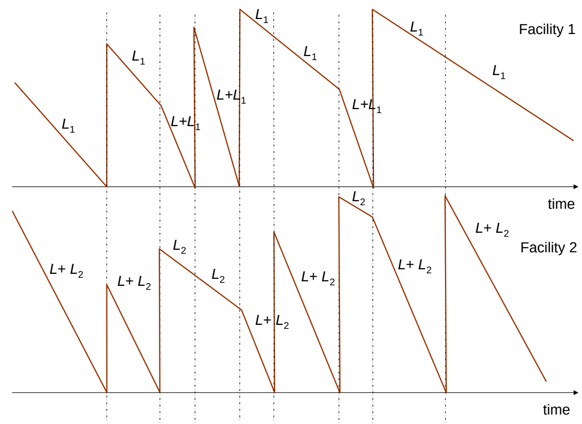

Figure 3: A $0/L$ trajectory

Theorem 2. There exists an optimal $0/L$ policy.

We show later that in this setting it might not be optimal to use a single facility to fulfill the remote demand. We next design a family of policies that sometimes perform better than the single facility policy. Consider the *2-slope cyclic policy*, Figure 4. Let $T = \alpha_1 + \alpha_2$ be the cycle length. The cost of such a policy is

$$g(\alpha_1, \alpha_2) = \frac{1}{T}\Bigg(k_1 + k_2 + c_1(\alpha_2 L_1 + \alpha_1(L + L_1)) + c_2(\alpha_1 L_2 + \alpha_2(L + L_2)) + f_1\alpha_1 L + f_2\alpha_2 L + \left[\frac{\alpha_1^2(L_1 + L)}{2} + \frac{\alpha_2^2 L_1}{2} + \alpha_1\alpha_2 L_1\right] h_1 + \left[\frac{\alpha_1^2 L_2}{2} + \frac{\alpha_2^2(L_2 + L)}{2} + \alpha_1\alpha_2 L_2\right] h_2\Bigg).$$

The best such policy is obtained by solving

$$\min_{\alpha_1 > 0, \alpha_2 > 0} g(\alpha_1, \alpha_2)\,.$$

A long but straightforward calculation from $\frac{\partial g}{\partial \alpha_1} = \frac{\partial g}{\partial \alpha_2} = 0$ shows that the best such policy is obtained by setting

$$\alpha_1 = \frac{h_2 T - \tilde{c}}{h_1 + h_2}$$
$$\alpha_2 = \frac{h_1 T + \tilde{c}}{h_1 + h_2}$$
$$T = \sqrt{\frac{2(k_1 + k_2)(h_1 + h_2) + L\tilde{c}^2}{L_1 h_1^2 + L_2 h_2^2 + h_1 h_2 (L + L_1 + L_2)}}\,.$$

In addition, we need to require that $\alpha_1 > 0, \alpha_2 > 0$. Note that one of them is always positive but in general it might not be the case that both of them are positive. If one of the two values is negative, then there does not exist a 2-slope cyclic policy.

The value of the policy that uses only a single facility to satisfy the remote demand is

$$Z^s = \min\{L_1 c_1 + \sqrt{2k_1 L_1 h_1} + c_2(L + L_2) + f_2 L + \sqrt{2k_2 h_2 (L_2 + L)}, L_2 c_2 + \sqrt{2k_2 L_2 h_2} + c_1(L + L_1) + f_1 L + \sqrt{2k_1 h_1 (L_1 + L)}\}\,.$$

We call such a policy the *trivial policy*. We now compare 2-slope cyclic and trivial policies.

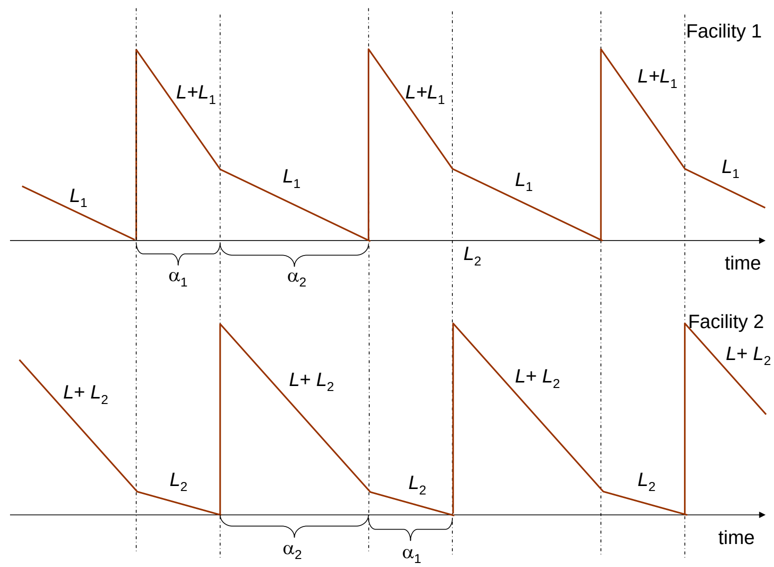

Figure 4: A 2-slope cyclic trajectory

Proposition 1. Let $c_1 = c_2 = f_1 = f_2 = 0$, i.e. we consider only the fixed replenishment cost and the holding cost. Then

1. the 2-slope cyclic policy can be better than the trivial policy,
2. the trivial policy can be arbitrarily better than the 2-slope cyclic policy,
3. the 2-slope cyclic policy can be at most by a factor of $\sqrt{3}$ better than the trivial policy.

Proof. Under the stated assumption of zero procurement and distribution cost, it is easy to see that the value of the optimal 2-slope cyclic policy is

$$Z^c = \sqrt{2(k_1 + k_2)(L_1h_1 + L_2h_2 + \frac{Lh_1h_2}{h_1 + h_2})}$$

and that it always exists.

We show the first statement by means of an example. Let $L = 2, L_1 = L_2 = h_1 = h_2 = k_1 = k_2 = 1$. For this choice we have $Z^s = \sqrt{2} + \sqrt{6}$. On the other hand $Z^c = \sqrt{12} < Z^s$. This shows that the optimal 2-slope cyclic policy outperforms the trivial policy.

Next we show that the trivial policy can be arbitrarily better than the 2-slope cyclic policy. Let A_1 denote the first term in the definition of Z^s and A_2 the second term. Consider L arbitrary small, i.e. $L \to 0$. Then

$$\frac{Z^c}{A_1} \to \frac{\sqrt{2(k_1 + k_2)(L_1h_1 + L_2h_2)}}{\sqrt{2k_1L_1h_1} + \sqrt{2k_2L_2h_2}}.$$

As $k_1 \to \infty, k_2 \to 0$, $L_2h_2 \to \infty, L_1h_1 \to 0$, it is easy to see that $\frac{Z^c}{A_1} \to \infty$. We conclude that $\frac{Z^c}{Z^s} \to \infty$.

We now show that the trivial policy cannot be worse than a factor of $\sqrt{3}$ from the optimal 2-slope cyclic policy. Let us assume that $h_1 > h_2$. We have

$$\begin{aligned}\frac{A_1}{Z^c} &= \frac{1+\sqrt{\frac{k_2}{k_1}\cdot\frac{L_2+L}{L_1}\cdot\frac{h_2}{h_1}}}{\sqrt{(1+\frac{k_2}{k_1})(1+\frac{L_2 h_2}{L_1 h_1}+\frac{L h_2}{L_1(h_1+h_2)})}} \\ &\le \frac{1+\sqrt{\frac{k_2}{k_1}\frac{L_2+L}{L_1}\frac{h_2}{h_1}}}{\sqrt{1+\frac{k_2}{k_1}\frac{L_2+L}{2L_1}\frac{h_2}{h_1}}} \\ &= \frac{1+\sqrt{2x}}{\sqrt{1+x}} \le \sqrt{3}.\end{aligned}$$

Here we denote $x = \frac{k_2}{k_1}\cdot\frac{L_2+L}{2L_1}\cdot\frac{h_2}{h_1}$. The first inequality is a long but straightforward calculation and it uses the fact $h_1 - h_2 \ge 0$. The last inequality is easy to check as well. Clearly now it follows $\frac{Z^s}{Z^c} \le \frac{A_1}{Z^c} \le \sqrt{3}$. □

The notion of cyclic policies can easily be extended to the general multi facility and multi market setting. Details about some of these more general cyclic policies are given in Section 6 and Klabjan (2009). Note also that there is a symmetric 2-slope cyclic policy. This policy for facility 1 during the time period corresponding to α_1, α_2 the slope is $L_1, L + L_1$, respectively. For facility 2, during the time period corresponding to α_1, α_2 the slope is $L + L_2, L_2$, respectively. The best of the two policies is from now on called the 2-slope cyclic policy.

4 Properties of Optimal Policies

In this section we state results regarding the structure of optimal policies in the general setting. Most of these results will be used later in Section 5.

Based on our definition of a decision epoch, we allow that at a decision epoch none of the facilities replenishes and therefore all of the trajectories have a breakpoint at such a decision epoch. The next theorem states that there is an optimal policy where in each decision epoch at least one facility replenishes. Since we have the zero-inventory ordering property, this also implies that at least one facility stocks out.

Theorem 3. There is an optimal policy where in each decision epoch at least one facility replenishes.

The technical proof of this result is given in Klabjan (2009). In our problem statement we do not explicitly require that the inventory level be below a certain upper bound at any point in time. It is easy to explicitly add such a requirement. Next we show that even if arbitrarily large inventory levels are allowed, in an optimal policy the inventory level is always bounded. This is not surprising due to the holding cost.

Proposition 2. There is an optimal policy such that the inventory level in each facility at any point in time is always less than or equal to

$$\max_{i\in N}\sqrt{\frac{56\cdot k_i\cdot\sum_{j\in M} L_j}{h_i}}\,. \tag{1}$$

Proof. We only sketch the proof here. The details are given in Klabjan (2009). The main idea is that if a large replenishment is made, then it is less costly to make several smaller replenishments, see Figure 5. By using such a strategy, we clearly decrease the holding cost, however the fixed cost increases since the new trajectory has more replenishments. It can be shown that for a replenishment larger than the quantity given in (1), the holding cost savings outweigh the increase in the fixed cost. It is easy to see that the variable procurement and the distribution costs do not change. □

We have already introduced trivial policies in Section 3.2 in the case of $n = 2$. A more general definition follows next. A *trivial policy* is any policy where the demand $D_{ij}(t)$ does not depend on t for every $i \in N, j \in M$. In a trivial policy the demand rate on any distribution link is constant throughout the planning

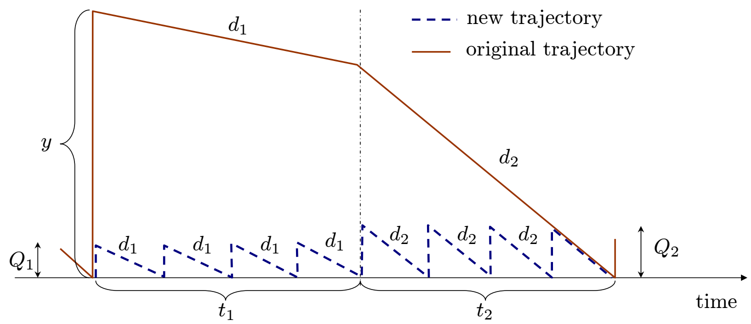

Figure 5: The trajectories for the inventory upper bound proof

horizon, i.e. it is a stationary policy. The optimal trivial policy can be computed by solving the following nonlinear optimization problem:

$$
\begin{aligned}
\min \sum_{i\in N}(c_i \sum_{j\in M} d_{ij} + \sum_{j\in M} f_{ij} d_{ij} + \sqrt{2k_i h_i \sum_{j\in M} d_{ij}})\\
\sum_{i\in N} d_{ij} = L_j \qquad j \in M\\
d \geq 0\,.
\end{aligned}
$$

Here d_{ij} is the stationary fraction of jth market demand satisfied by facility i. Note that this is a concave program and therefore an extreme point optimal solution exists, i.e. for every $j \in M$ there exists $i_j \in N$ such that $d^*_{i_j j} = L_j$ in an optimal solution d^*. We study the effectiveness of trivial policies by means of numerical experiments in Section 6.

5 Lagrangian Lower Bound

In this section we model our problem as a semi-Markov decision process. For a more general discussion of semi-Markov decision processes see for example Bhattacharya and Majumdar (1989), Vega-Amaya (1993), and Luque-Vásquez and Hernández-Lerma (1999). The resulting semi-Markov model is weakly coupled. We develop a Lagrangian relaxation theory for weakly coupled semi-Markov decision processes. Weakly coupled Markov decision processes have been recently studied, Hawkins (2003), Meuleau *et al.* (1998), Adelman and Mersereau (2004), however we are not aware of any such study for semi-Markov decision processes. The treatment of such processes is more involved. At the end of this section we apply this theory to the multi-market problem in order to obtain a lower bound on the optimal value.

5.1 Semi-Markov Decision Processes

We first define semi-Markov decision processes and stationary deterministic policies. We do not define them in a stochastic setting but only in a deterministic setting. Such processes are called semi-Markov decision processes with Dirac's transition kernels, Klabjan and Adelman (2006). It turns out that in many respects such processes are even harder to study and analyze than their stochastic counterparts. Along the same line, we do not consider randomized policies but only deterministic stationary policies.

We denote a semi-Markov decision process as sMDP(c, τ, s, S, K), where c is the cost function, τ the transit time function, s the transition function, S the state space, and K the state-action space. For simplicity we consider only state and action spaces in Euclidian spaces. Thus $S \subseteq \mathbb{R}^n$ and $K \subseteq \mathbb{R}^{n+n}$ for some n. Given a state x, we can take any action $a \in A(x)$. As a result we next move to the state $s(x, a) \in S$.

The move takes $\tau(x, a)$ units of time and we incur cost $c(x, a)$. The goal is to minimize the long-run average cost.

Formally, we have $K = \{(x, a)|x \in S, a \in A(x)\}$ and $c : K \to \mathbb{R}, \tau : K \to \mathbb{R}_+, s : K \to S$. In addition we define the action space $A = \{A(x)|x \in S\} \subseteq \mathbb{R}^n$. A *policy* is a function $f : S \to A$ such that for every $x \in S$ we have $f(x) \in A(x)$. For every policy f we define the *long-run average cost* starting in state x as

$$J(x, f) = \limsup_{T\to\infty} \frac{\sum_{j=0}^{T-1} c(x^j, f(x^j))}{\sum_{j=0}^{T-1} \tau(x^j, f(x^j))} ,$$

where $x^0 = x$ and $x^j = s(x^{j-1}, f(x^{j-1}))$ for $j = 1, 2, \dots$.

We want to find a policy with as low long-run average cost as possible. For every $x \in S$ let the optimal average cost function be defined as $J(x) = \inf_f J(x, f)$. Note that we deliberately write infimum instead of minimum since a single optimal policy might not exist. Let also $\rho^* = \inf_{x\in S} J(x)$. The *average cost problem* is to find ρ^* and the corresponding x^*, f^*, if they exist, such that $J(x^*, f^*) = \rho^*$. We always assume that $\rho^* < \infty$, i.e. there exists at least one policy with a finite long-run average cost.

A solution (u, ρ^*) to the *optimality equation*

$$u(x) = \min\{c(x, a) - \rho^*\tau(x, a) + u(s(x, a))\}$$

implies the existence of an optimal stationary deterministic policy, i.e. we can replace the infimums above with minimums. The existence of such a pair is hard to show, Klabjan and Adelman (2006).

5.2 The Multi-Market Problem as a Semi-Markov Decision Process

Here we model our problem as a semi-Markov decision process. We give a formulation that is suitable for applying the Lagrangian relaxation. We use the convention that $0/0 = \infty$ and also $0/\infty = 0, \infty/0 = \infty$.

We clearly need to track the inventory level at each facility. Let x_i be the current inventory level at facility i. In addition, for cost accounting, we need to know the demand rate just before each decision epoch. Let D_i be the demand rate at facility i just before the current decision epoch, i.e. between the previous decision epoch and the current one. The need for D_i will be clear later. Similarly, let F_i be the distribution cost incurred at facility i just before the decision epoch. We argue later about the role of F_i. For ease of notation we write $x = (x_i)_{i\in N}, D = (D_i)_{i\in N}, F = (F_i)_{i\in N}$.

We have two sets of action variables. Let y_i be the replenishment quantity of facility i and let d_{ij} be the demand rate between facility i and market j during the time until the next decision epoch. Let $d = (d_{ij})_{i\in N, j\in M}$. In our case the action space does not depend on the state. Formally,

$$A = A(x, D, F) = \{(y, d)| \sum_{i\in N} d_{ij} = L_j \text{ for all } j \in M, y \geq 0, d \geq 0\} ,$$

and $S = \mathbb{R}^{3n}_+, K = S \times A$.

For ease of exposition we define the following quantities.

$$\begin{aligned}
\bar{d}_i &= \sum_{j\in M} d_{ij} \\
\bar{F} &= (\sum_{j\in M} f_{1j}d_{1j}, \sum_{j\in M} f_{2j}d_{2j}, \dots, \sum_{j\in M} f_{nj}d_{nj}) \\
\bar{d} &= (\bar{d}_1, \bar{d}_2, \dots, \bar{d}_n) \\
\tilde{d}_i &= (d_{i1}, d_{i2}, \dots, d_{im})
\end{aligned}$$

In view of Theorem 3 and since backlogging is not allowed, the next decision epoch is going to be at the time when the first facility stocks out. Therefore

$$\tau(x, D, F; y, d) = \min_{i\in N} \frac{x_i + y_i}{\bar{d}_i}$$

and the inventory levels after this amount of time are $x + y - \tau(x, D, F; y, d) \cdot \bar{d}$. By taking into account the definition of D and F we conclude that $F = \bar{F}$ and $D = \bar{d}$ in the next decision epoch. To summarize, the transition function is

$$s(x, D, F; y, d) = (x + y - \tau(x, D, F; y, d) \cdot \bar{d}, \bar{d}, \bar{F}) \, .$$

Cost accounting is very technical. The total procurement cost is $\sum_{i\in N} c_i y_i$ and the total fixed cost is $\sum_{i\in N} k_i \delta(y_i)$. Holding cost accounting is more involved and we explain it by an example, Figure 6. Consider five decision epochs A_1, A_2, A_3, A_4, C_2 spanning two consecutive replenishments at facility i. The holding cost corresponds to the shaded area multiplied by the per unit holding cost. It is formed by four separated regions. Consider the following relationship, where $\triangle XYZ$ denotes the area of the triangle spanned by vertices X, Y, Z.

$$\begin{aligned}
&① + ② + ③ + ④ \\
&\quad = (\triangle A_1B_1C_4 - \triangle A_2B_2C_4) + (\triangle A_2B_2C_1 - \triangle A_3B_3C_1) + (\triangle A_3B_3C_3 - \triangle A_4B_4C_3) + \triangle A_4B_4C_2 \\
&\quad = \triangle A_1B_1C_4 + (\triangle A_2B_2C_1 - \triangle A_2B_2C_4) + (\triangle A_3B_3C_3 - \triangle A_3B_3C_1) + (\triangle A_4B_4C_2 - \triangle A_4B_4C_3)
\end{aligned} \tag{2}$$

The last three expressions in brackets correspond to subtracting the area of the triangle defined by the current inventory level and the previous demand rate from the area of the triangle defined by the current inventory level and the next demand rate. Considering that the previous demand rate is encoded in the state space, this can be calculated. The former demand rate is $\bar{d}_i$ and the latter one is D_i. Hence each one of these three terms can be expressed as

$$\frac{x_i^2}{2\bar{d}_i} - \frac{x_i^2}{2D_i} = \frac{x_i^2}{2}\left(\frac{1}{\bar{d}_i} - \frac{1}{D_i}\right) = \frac{(x_i + y_i)^2}{2}\left(\frac{1}{\bar{d}_i} - \frac{1}{D_i}\right) ,$$

since in these decision epochs we do not replenish, i.e. $y_i = 0$.

The first term in (2) equals to $y_i^2/(2\bar{d}_i)$, which is equivalent to $(x_i + y_i)^2/(2\bar{d}_i)$, since at this decision epoch $x_i = 0$. The two expressions can be written in a compact form as

$$\frac{(x_i + y_i)^2}{2}\left(\frac{1}{\bar{d}_i} - \frac{1 - \delta(y_i)}{D_i}\right) .$$

We conclude that the total holding cost is

$$\sum_{i\in N}\left(\frac{h_i(x_i + y_i)^2}{2}\left(\frac{1}{\bar{d}_i} - \frac{1 - \delta(y_i)}{D_i}\right)\right) .$$

The same trick can be used for the distribution cost. We obtain that the total distribution cost equals to

$$\sum_{i\in N}(x_i + y_i)\left(\frac{\sum_{j\in M} f_{ij} d_{ij}}{\bar{d}_i} - \frac{F_i(1 - \delta(y_i))}{D_i}\right) .$$

To summarize, the total cost equals to

$$\begin{aligned}
c(x, D, F; y, d) = \sum_{i\in N}\Bigg[& c_i y_i + k_i\delta(y_i) \\
& + \sum_{i\in N}(x_i + y_i)\left(\frac{\sum_{j\in M} f_{ij} d_{ij}}{\bar{d}_i} - \frac{F_i(1 - \delta(y_i))}{D_i}\right) \\
& + \sum_{i\in N}\left(\frac{h_i(x_i + y_i)^2}{2}\left(\frac{1}{\bar{d}_i} - \frac{1 - \delta(y_i)}{D_i}\right)\right)\Bigg] .
\end{aligned}$$

There are a few important observations about our model. The cost function is separable with respect to facilities. It is possible to model the problem only with the inventory variables x, however in this case the

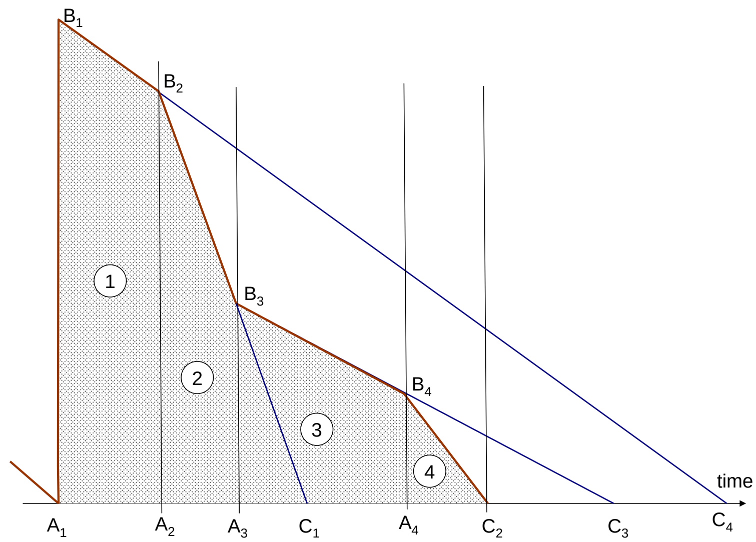

Figure 6: Holding cost accounting

cost function depends on τ and therefore it is no longer separable. Same cannot be said about the transit time function, however τ is the minimum of several single facility functions. The real linking part among the several facilities are the action space constraints $\sum_{i\in N} d_{ij} = L_j$. If these constraints were not present, then it seems plausible that the problem would decompose with respect to facilities. At this point we will use the idea from the Lagrangian relaxation, which is standard in integer programming and it has been extended to Markov decision processes. We neglect these constraints but we add a penalty λ_j for each violation. In the next section we present a general framework for weakly coupled semi-Markov decision processes.

5.3 Lagrangian Lower Bound for Weakly Coupled Semi-Markov Decision Processes

In this section we consider semi-Markov problems that are weakly coupled. Such problems consist of several independent processes that are linked only by constraints in the action space and in a specific form by the transition function and the transit cost function.

We consider a semi-Markov decision process sMDP(c, τ, s, S, K) with the following properties and optimal value ρ^*.

$$c(x,a) = \sum_{i\in N} c_i(x_i, a_i) \tag{3}$$

$$\tau(x,a) = \min_{i\in N} \tau_i(x_i, a_i) \tag{4}$$

$$s(x,a) = \alpha(x,a) - \tau(x,a)\cdot\beta(x,a) \tag{5}$$

$$\alpha(x,a) = (\alpha_1(x_1,a_1), \alpha_2(x_2,a_2), \ldots, \alpha_n(x_n,a_n)) \tag{6}$$

$$\beta(x,a) = (\beta_1(x_1,a_1), \beta_2(x_2,a_2), \ldots, \beta_n(x_n,a_n)) \tag{7}$$

$$H(x,a) = \sum_{i\in N} H^i(x_i, a_i) \tag{8}$$

$$A(x) = \{a \in \mathbb{R}^n | a_i \in A_i(x_i), H(x,a) \geq b\} \tag{9}$$

Here $c_i, \tau_i, \alpha_i, \beta_i$ are functions from $\mathbb{R}^2$ to $\mathbb{R}$ for each i. For each $i \in N$ the function H^i is a function from $\mathbb{R}^2$ to $\mathbb{R}^m$ for a given m. We denote by H^i_k the kth component. In addition, b is a vector in $\mathbb{R}^m$ and A_i are any one dimensional action spaces. Condition (3) states that the cost is separable, while (4) states that the joint transit time is the minimum of individual transit times. Requirements (5)-(7) impose a weak separability of the transition function. If $\beta(x,a)=0$, then the transition function is indeed separable. Finally, note that (8) and (9) require that the action space is separable except for constraints $H(x,a)\ge b$.

Let us fix an arbitrary $\lambda\in\mathbb{R}^m_+$. For each $i\in N$ we consider the semi-Markov decision processes defined by sMDP$^\lambda_i(c_i-\lambda H^i\tau_i,\tau_i,\alpha_i-\tau_i\beta_i,S_i,K_i)$, where $S=S_1\times S_2\times\cdots\times S_n$ and $K_i=\{(x_i,a_i)|a_i\in A_i(x_i),x_i\in S_i\}$. We require that for each $i\in N$ there exists a pair $(v^\lambda_i,\rho^\lambda_i)$ that solves the optimality equation

$$v^\lambda_i(x_i)=\min_{a_i\in A_i(x_i)}\{c_i(x_i,a_i)-\lambda H^i(x_i,a_i)\tau_i(x_i,a_i)-\rho^\lambda_i\tau_i(x_i,a_i)+v^\lambda_i(\alpha_i(x_i,a_i)-\tau_i(x_i,a_i)\beta(x_i,a_i))\}\,. \tag{10}$$

This implies that $sMPD^\lambda_i$ has an optimal stationary deterministic policy. Here we implicitly assume that the minimum in the right-hand side of (10) is attained. The following is the main result of this section.

Theorem 4. Let us assume that

1. $H\ge 0$,
2. $\rho^\lambda_i\ge 0$ for every $i\in N$,
3. $\beta_i(x_i,a_i)\ge 0$ for every $i\in N$ and $(x_i,a_i)\in K_i$,
4. v^λ_i is nondecreasing for every $i\in N$,
5. $||v^\lambda_i||_\infty<M<\infty$ for every $i\in N$, and
6. there exists a constant $\zeta>0$ such that $\tau_i(x_i,a_i)\ge\zeta$ for every $(x_i,a_i)\in K_i$ and $i\in N$.

Then

$$\lambda b+\sum_{i\in N}\rho^\lambda_i\le\rho^*\,.$$

If $\beta=0$, then condition 4 is not needed.

Proof. Since $H(x,a)\ge b$ for any $(x,a)\in K$ and $\lambda\ge 0$, we obtain $\lambda H(x,a)\ge\lambda b$. In addition, since $\tau\ge 0$ we in turn get $(\lambda\cdot H(x,a))\tau(x,a)\ge(\lambda\cdot b)\tau(x,a)$ for every $(x,a)\in K$.

We consider sMDP$(c+\lambda b-\lambda H\tau,\tau,s,S,K)$ and let $\tilde\rho^\lambda$ be the optimal value. It is clear from the above that $\tilde\rho^\lambda\le\rho^*$. We next show that $\lambda b+\sum_{i\in N}\rho^\lambda_i\le\tilde\rho^\lambda$.

Let $\bar v^\lambda(x)=\sum_{i\in N}v^\lambda_i(x_i)$ for every $x\in S$ and $\bar\rho^\lambda=\lambda b+\sum_{i\in N}\rho^\lambda_i$. We first show that

$$\bar v^\lambda(x)\le\inf_{\substack{a_i\in A(x_i)\\ \text{for every } i\in N}}\{c(x,a)-\lambda H(x,a)\tau(x,a)-(\bar\rho^\lambda-\lambda b)\tau(x,a)+\bar v^\lambda(s(x,a))\} \tag{11}$$

for every $x\in S$.

Let us consider an arbitrary (x,a) such that $(x_i,a_i)\in K_i$ for every $i\in N$. We have

$$s(x,a)_i=\alpha_i(x_i,a_i)-\tau(x,a)\beta(x_i,a_i)\ge\alpha_i(x_i,a_i)-\tau_i(x_i,a_i)\beta(x_i,a_i)\,. \tag{12}$$

Here we used condition 3 and (4). Since v^λ_i is nondecreasing (condition 4), it follows

$$v^\lambda_i(s(x,a)_i)\ge v^\lambda_i(\alpha_i(x_i,a_i)-\tau_i(x_i,a_i)\beta(x_i,a_i))$$

and therefore

$$\bar v^\lambda(s(x,a))=\sum_{i\in N}v^\lambda_i(s(x,a)_i)\ge\sum_{i\in N}v^\lambda_i(\alpha_i(x_i,a_i)-\tau_i(x_i,a_i)\beta(x_i,a_i))\,. \tag{13}$$

By definition we have $c(x,a)=\sum_{i\in N}c_i(x_i,a_i)$. Since $H\ge 0$ (condition 1), $\lambda\ge 0$, and $\tau=\min\{\tau_1,\dots,\tau_n\}$, we get

$$-\lambda H(x,a)\tau(x,a)\ge-\sum_{i\in N}\sum_{k=1}^m\tau_i(x_i,a_i)\lambda_kH^i_k(x_i,a_i)=-\sum_{i\in N}\lambda H^i(x_i,a_i)\tau(x_i,a_i)\,. \tag{14}$$

From condition 2 and (4) we obtain

$$\rho_i^\lambda \tau(x,a) \le \rho_i^\lambda \tau_i(x_i,a_i)\,. \tag{15}$$

From (13), (14), and (15) we obtain

$$\begin{aligned} &c(x,a) - \lambda H(x,a)\tau(x,a) - (\bar\rho^\lambda - \lambda b)\tau(x,a) + \bar v^\lambda(s(x,a)) \\ &\quad\ge \sum_{i\in N}\left[c_i(x_i,a_i) - \lambda H^i(x_i,a_i)\tau_i(x_i,a_i) - \rho_i^\lambda\tau_i(x_i,a_i) + v_i^\lambda(\alpha_i(x_i,a_i) - \tau_i(x_i,a_i)\beta(x_i,a_i))\right] \\ &\quad\ge \sum_{i\in N} v_i^\lambda(x_i) = \bar v^\lambda(x)\,. \end{aligned}$$

We have also used (10). Since x, a are arbitrary, we conclude that (11) holds.

To finish the proof, we argue that $\bar\rho^\lambda - \lambda b \le \tilde\rho^\lambda$. For ease of notation let $\tilde c = c - \lambda H\tau$. Let $\epsilon > 0$ be arbitrary and let us fix $x \in S$. By definition of $\tilde\rho^\lambda$ there exists a policy $f = f_\epsilon$ such that

$$\limsup_{T\to\infty} \frac{\sum_{j=0}^{T-1}\tilde c(x^j, f(x^j))}{\sum_{j=0}^{T-1}\tau(x^j, f(x^j))} \le \tilde\rho^\lambda + \epsilon\,.$$

From (11) we obtain

$$\bar v^\lambda(x^j) \le \tilde c(x^j, f(x^j)) - (\bar\rho^\lambda - \lambda b)\tau(x^j, f(x^j)) + \bar v^\lambda(s(x^j, f(x^j)))$$

for $j = 0, 1, \ldots, T-1$, where $x = x^0$. By summing these inequalities we obtain

$$\bar\rho^\lambda - \lambda b \le \frac{\sum_{j=0}^{T-1}\tilde c(x^j, f(x^j))}{\sum_{j=0}^{T-1}\tau(x^j, f(x^j))} + \frac{\bar v^\lambda(x_T) - \bar v^\lambda(x)}{\sum_{j=0}^{T-1}\tau(x^j, f(x^j))}\,. \tag{16}$$

From condition 5 we obtain $|\bar v^\lambda(x_T) - \bar v^\lambda(x)| \le 2Mn$ and from condition 6 it follows $\sum_{j=0}^{T-1}\tau(x^j, f(x^j)) \ge T\zeta$. As T goes to infinity, the last term in (16) tends to 0. Therefore we obtain

$$\bar\rho^\lambda - \lambda b \le \frac{\sum_{j=0}^{T-1}\tilde c(x^j, f(x^j))}{\sum_{j=0}^{T-1}\tau(x^j, f(x^j))} \le \tilde\rho^\lambda + \epsilon\,.$$

Since ϵ is arbitrary, it follows that $\bar\rho^\lambda - \lambda b \le \tilde\rho^\lambda$. This completes the proof.

If $\beta = 0$, then in (12) we have equality and subsequently the nondecreasing property is not required. □

It is interesting to note that we are not able to prove that $\bar v^\lambda(x)$ satisfies the optimality equation for sMDP$(c + \lambda b - \lambda H\tau, \tau, s, S, K)$. Inequality (11) shows one direction. We point out that the theorem can be easily extended to a stochastic setting by appropriately adapting conditions (5)-(7).

5.4 A Lower Bound for the Multi-Market Problem

In this section we apply the Lagrangian lower bound to the formulation presented in Section 5.2. In order to carry out the analysis, we make a slight change to the actions space. Let $\hat L = \sum_{j\in M} L_j$. We redefine

$$A = A(x, D, F) = \{(y,d) | \sum_{i\in N} d_{ij} \ge L_j \text{ for all } j \in M, \sum_{j\in M} d_{ij} \le \hat L \text{ for all } i \in N, y \ge 0, d \ge 0\}\,.$$

We can justify this in two steps. If we retain $\sum_{i\in N} d_{ij} = L_j, j \in M$, then it is obvious that constraints $\sum_{j\in M} d_{ij} \le \hat L$ are redundant and therefore they can be added to the actions space. In the second step, since the cost function is nonnegative, it is easy to see that $\sum_{i\in N} d_{ij} \ge L_j, j \in M$ always implies equality. To summarize, the dynamic program with the new actions space is identical to the original one. As in Section 5.3, we denote by ρ^* the optimal value of the dynamic program for the multi-market problem. We first summarize some easy facts about the standard economic order quantity model and then we develop a lower bound for the multi-market problem.

5.4.1 Economic Order Quantity as Semi-Markov Decision Process

Consider the basic economic order quantity model with fixed procurement cost k and holding cost h. Let the demand rate be denoted by r. Then the total cost in a cycle of length T is $k + hTQ/2$, where Q is the order quantity. Since $Q = Tr$ we obtain that the cost per unit time is $kr/Q + hQ/2$. The optimal order quantity is $Q = \sqrt{2kr/h}$.

Now let us also assume that the per unit procurement cost is c and that for every demanded item we pay a distribution cost of f per unit. Then the total cost of a cycle of length T is $fTr + cQ + k + hTQ/2$. After using $Q = Tr$ we obtain that the cost per unit time is $(f + c)r + kr/Q + hQ/2$, which gives the same order quantity. In this case the long-run average cost is

$$\rho = (c + f)r + \sqrt{2khr}\,. \tag{17}$$

We can also model this simple problem as a semi-Markov decision process with the profit-to-go or value function v and the long-run average cost of ρ. The state space x corresponds to the current inventory and the replenishment quantity is y. The transition time between two decision epochs is $(x + y)/r$. The optimality equation reads

$$v(x) = \min_{y\geq 0}\{cy + f(x + y) + k\delta(y) + \frac{h(x + y)^2}{2r} - \rho\frac{x + y}{r} + v(0)\}\,. \tag{18}$$

It is clear that $y = 0$ if $x > 0$ and $y = Q$ if $x = 0$. The profit-to-go function v is given by

$$v(x) = \begin{cases} (c + f)x & 0 \leq x < Q \\ hx^2/(2r) - \rho x/r & x \geq Q\,. \end{cases} \tag{19}$$

This function is depicted in Figure 7. Note that v is an increasing function.

5.4.2 Application to The Multi-Market Problem

In order to apply the result of Theorem 4, we need to identify the structure required by (3)-(9). We are relaxing constraints $\sum_{i\in N} d_{ij} \geq L_j, j \in M$. We have

$$\begin{aligned}
c_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= c_iy_i + k_i\delta(y_i) + (x_i + y_i)\left(\frac{\sum_{j\in M} f_{ij}d_{ij}}{\bar{d}_i} - \frac{F_i(1 - \delta(y_i))}{D_i}\right) \\
&\quad + \frac{h_i(x_i + y_i)^2}{2}\left(\frac{1}{\bar{d}_i} - \frac{1 - \delta(y_i)}{D_i}\right) && i \in N \\
\tau_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= \frac{x_i + y_i}{\bar{d}_i} && i \in N \\
\alpha_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= x_i + y_i && i \in N \\
\alpha_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= \bar{d}_i && i = n + 1, \ldots, 2n \\
\alpha_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= \sum_{j\in M} f_{ij}d_{ij} && i = 2n + 1, \ldots, 3n \\
\beta_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= \bar{d}_i && i \in N \\
\beta_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= 0 && i = n + 1, \ldots, 3n \\
A_i(x_i, D_i, F_i) &= \{(y_i, \tilde{d}_i) : \sum_{j\in M} d_{ij} \leq \hat{L}, y_i \geq 0, \tilde{d}_i \geq 0\} && i \in N \\
D_i(x_i, D_i, F_i; y_i, \tilde{d}_i) &= \tilde{d}_i && i \in N\,.
\end{aligned}$$

We also have $b_j = L_j$ for every $j \in M$.

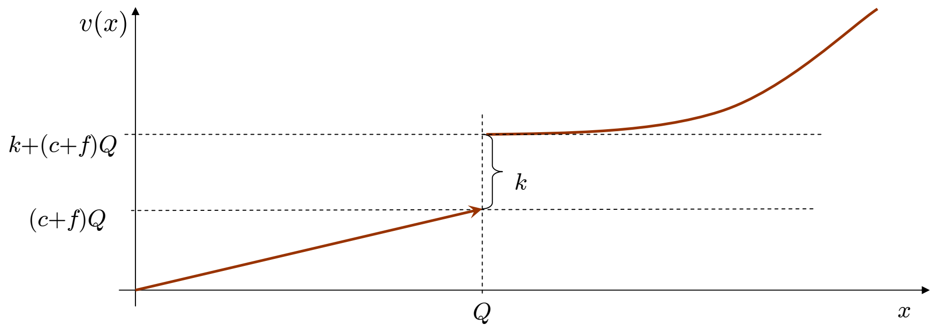

Figure 7: The profit-to-go function of standard economic order quantity

The resulting optimality equation (10) for the decoupled semi-Markov decision process reads

$$\begin{aligned} v_i^\lambda(x_i, D_i, F_i) = \min_{y_i \ge 0, \tilde{d}_i \ge 0, \bar{d}_i \le \hat{L}} \Big[c_i(x_i, D_i, F_i; y_i, \tilde{d}_i) - \frac{x_i + y_i}{\bar{d}_i} \sum_{j \in M} \lambda_j d_{ij} - \rho_i^\lambda \frac{x_i + y_i}{\bar{d}_i} \\ + v_i^\lambda(0, \tilde{d}_i, \sum_{j \in M} f_{ij} d_{ij}) \Big] . \end{aligned} \tag{20}$$

Using the same arguments for cost accounting to those in Section 5.2, except that we reverse them, we obtain that (20) is equivalent to

$$u_i^\lambda(x_i) = \min_{y_i \ge 0, \tilde{d}_i \ge 0, \bar{d}_i \le \hat{L}} \left[x_i \delta(y_i) + c_i y_i + \frac{h_i (x_i + y_i)^2}{2 \bar{d}_i} + \frac{x_i + y_i}{\bar{d}_i} \sum_{j \in M} (f_{ij} - \lambda_j) d_{ij} - \rho_i^\lambda \frac{x_i + y_i}{\bar{d}_i} + u_i^\lambda(0) \right] .$$

A rigorous argument will be given later.

It is easy to see that only one market will be used, i.e. $\sum_{j \in M} (f_{ij} - \lambda_j) d_{ij} = (f_{ik} - \lambda_k) d_{ik}$ for a $k \in M$. If we fix the demand rate d_{ik} (note that $d_{ik} = \bar{d}_i$), then by comparing this optimality equation with (18), we see that the resulting problem is the standard economic order quantity problem with procurement cost c_i, the distribution cost $f_{ik} - \lambda_k$, fixed cost k_i, and demand rate d_{ik}. Thus based on (17) the long-run average cost is

$$\bar{g}(d_{ik}) = (c_i + f_{ik} - \lambda_k) d_{ik} + \sqrt{2 k_i h_i d_{ik}} \,.$$

Let $\mu = \min_{0 \le d \le \hat{L}} \bar{g}(d)$. Clearly $\mu = 0$ if $c_i + f_{ik} - \lambda_k \ge 0$ since the optimal demand rate d is 0. Let now $c_i + f_{ik} - \lambda_k < 0$. An elementary calculation shows that $\bar{g}$ is increasing in $[0, \frac{k_i h_i}{2(c_i + f_{ik} - \lambda_k)^2}]$ and decreasing on $[\frac{k_i h_i}{2(c_i + f_{ik} - \lambda_k)^2}, \infty]$. In addition $\lim_{d \to \infty} \bar{g}(d) = -\infty$ and $\bar{g}(\hat{d}) = 0$ for $\hat{d} = \frac{2 k_i h_i}{\lambda_k - c_i - f_{ik}}$. Thus if $\hat{L} < \hat{d}$, then $\mu = 0$, and if $\hat{L} \ge \hat{d}$, then the minimum is attained at $d = \hat{L}$. We conclude that if $\lambda_k \le \sqrt{\frac{2 k_i h_i}{\hat{L}}} + c_i + f_{ik}$, then $\mu = 0$. If $\lambda_k > \sqrt{\frac{2 k_i h_i}{\hat{L}}} + c_i + f_{ik}$, then $\mu = \hat{L}(c_i + f_{ik} - \lambda_k) + \sqrt{2 k_i h_i \hat{L}}$. From this expression it also follows that $k = \mathrm{argmax}_{j \in M} (\lambda_j - f_{ij})$. We conclude that

$$\rho_i^\lambda = \begin{cases} 0 & \lambda_k \le c_i + f_{ik} + \sqrt{\frac{2 k_i h_i}{\hat{L}}} \\ \hat{L}(c_i + f_{ik} - \lambda_k) + \sqrt{2 k_i h_i \hat{L}} & \text{otherwise.} \end{cases}$$

If the assumptions in Theorem 4 are satisfied, then

$$\sum_{j \in M} \lambda_j L_j + \sum_{i \in N} \rho_i^\lambda$$

provides a lower bound on ρ^*. Clearly we want to find λ that maximizes this lower bound. This can be written as the following optimization problem.

$$\max \sum_{j\in M} \lambda_j L_j + \sum_{i\in N} \left[(\hat{L}(c_i - t_i) + \sqrt{2\hat{L}k_i h_i})\chi(t_i - c_i - \sqrt{2k_i h_i/\hat{L}}) \right] \tag{21a}$$

$$\lambda_j - t_i \leq f_{ij} \qquad j \in M, i \in N \tag{21b}$$

$$\lambda \geq 0, t \text{ unrestricted.} \tag{21c}$$

Here $\chi(x)$ is 1 if $x > 0$ and 0 if $x \leq 0$. It is easy to see that in an optimal solution $t_i = \max_{j\in M}(\lambda_j - f_{ij})$. The following proposition is straightforward to prove and it shows that (21) is a concave optimization problem.

Proposition 3. The objective function (21a) is separable and concave.

To validate that (21) gives a lower bound, we need to check the requirements stated in Theorem 4. Conditions 1 to 3 are obvious. From (19) and the aforementioned analysis it is easy to see that

$$u_i^\lambda(x_i) = \begin{cases} (c_i + f_{ik} - \lambda_k)\hat{L} & x_i < Q_i \\ \frac{h_i x_i^2}{2\hat{L}} - \frac{\rho_i^\lambda x_i}{\hat{L}} & x_i \geq Q_i\,, \end{cases}$$

where

$$\rho_i^\lambda = \sqrt{2k_i h_i \hat{L}} + (c_i + f_{ik} - \lambda_k)\hat{L} \tag{22}$$

$$Q_i = \sqrt{2k_i \hat{L}/h_i}\,. \tag{23}$$

This holds only if $\lambda_k - f_{ik} > Q_i - c_i$. If this condition is violated, then $u_i^\lambda(x_i) = 0$ and $\rho_i^\lambda = 0$. In addition we have $v_i^\lambda(x_i, D_i, F_i) = u_i^\lambda(x_i)$. From Figure 7 it follows that v_i^λ is increasing and therefore condition 4 holds. By Proposition 2 we can assume from the very beginning that x is in a compact space, which implies that it suffices to consider v_i^λ only on a compact subset. This implies that these functions are bounded and therefore condition 5 is satisfied. From the above analysis it also follows that $\tau_i(x_i, a_i) \geq Q_i/\hat{L}$ for every $i \in N$ and $(x_i, a_i) \in K_i$, which justifies condition 6.

6 Numerical Experiments

In this section we evaluate the proposed policies and the lower bound by means of computational experiments. All of the experiments were performed on a Toshiba Portégé M200 Tablet PC with an Intel Pentium Mobile 1.70 GHz central processing unit and 512 MB of main memory. The optimization routines were coded in Microsoft Visual Basic 6.3 and What's Best 7.0 was used as the optimization solver.

The main purposes of the computational study are: (1) to establish how often and by how much the alternative policies outperform the trivial policy, and (2) to evaluate the quality of the Lagrangian lower bound, which is computed by solving (21). We perform an extensive sensitivity analysis study. The input data were randomly generated as described later. We start by studying the two facility setting with local markets and a single remote market.

6.1 Two Facilities with Local Markets and a Single Remote Market

In this section we use the notation introduced in Section 3.2. The input data are generated randomly with the underlying distributions being uniform over a given interval. The holding cost at facility i is always a random fraction r_i of the procurement cost, i.e. $h_i = r_i \cdot c_i$. We always express r_i as percents of c_i. The default values are: $k_i \in [100, 400], c_i \in [5, 15], r_i \in [50\%, 80\%], f_j \in [1, 5], L_j \in [3000, 6000], L \in [3000, 6000]$ for $i = 1, 2$ and $j = 1, 2$. For example, the fixed cost at both facilities is a random number between 100 and 400. For each experiment we generated 1,000 random instances and for each instance we computed the lower bound, the optimal 2-slope cyclic policy, and the optimal trivial policy. Every number in the figures and tables that follow is an average over all randomly generated instances. To compute the optimal trivial policy, the optimal 2-slope cyclic policy, and the lower bound it takes less than a second of computational time.

We first study the effects of various cost components and demand values with respect to the trivial and the 2-slope cyclic policy. In addition we also report the optimality gaps. For each experiment we document the effectiveness of the 2-slope cyclic policies with respect to the optimal trivial policy. We record the *frequency*, which is defined as the number of times the 2-slope cyclic policy outperforms the trivial policy. The relative improvement, which we also call the *improvement gap*, is defined as 100 times the ratio of the value of the trivial policy minus the value of the 2-slope cyclic policy over the value of the 2-slope cyclic policy. We also measure the average of the improvement gaps and the maximum improvement gap of the 2-slope cyclic policy. The left charts in Figures 8-11 always show these values. In these charts the right axis shows the frequency values and the other axis the values corresponding to the average and the minimum. We also show the relative gap called the *optimality gap* between the two policies and the lower bound. The values for "Maximum" correspond to the maximum optimality gap in all 1,000 instances. The average values correspond to the average optimality gap among all of the instances. The minimum value shows the minimum optimality gap among all positive optimality gaps. These three statistics are shown on right charts in Figures 8-11. The two values pertaining to the minimum and the average are shown on the left axis and the maximum scale is on the right axis. The frequency of zero optimality gaps is shown later.

In the first experiment we vary the fixed cost from the average value of a 100 to 1,500, Figure 8. As the fixed cost increases, the effectiveness of the 2-slope cyclic policy keeps improving. The frequency varies from 2% for the fixed cost in the range $[0, 200]$ to 15% for the fixed cost in the range $[1000, 2000]$. The value 15% means that for 150 instances out of 1,000, the 2-slope cyclic policy outperforms the trivial policy. Similar trend is observed for the maximum and the average improvement gaps. The maximum observed improvement is approximately 1.3% while the corresponding average improvement gap of the two policies is 0.4%. The improvements and their frequencies are definitely low. The same increasing trend is seen with respect to the optimality gap. While the gap can be relatively large (up to 40%), the average optimality gap is very low. It varies from 0.50% to 2.00%. Whenever the gap is positive, we see that it is always larger than approximately 0.04%.

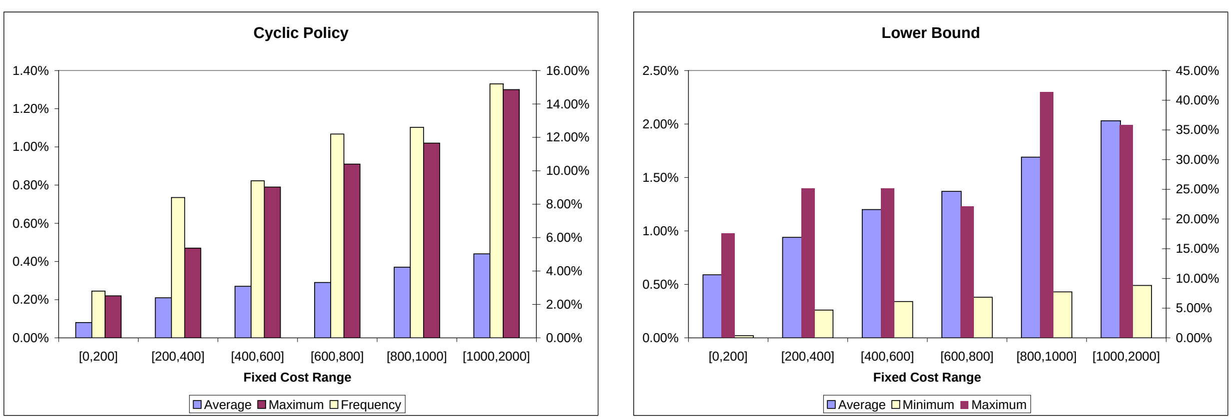

Figure 8: The effect of the fixed cost

The sensitivity analysis with respect to the distribution cost is shown in Figure 9. All of the remaining cost values correspond to the default ones, e.g. the fixed costs are random numbers in the range $[100, 400]$. The improvement and the optimality gap ranges are of the same order. We observe that with the increased distribution cost, the 2-slope cyclic policy is less effective and the optimality gap decreases as the distribution cost increases. The trends are less pronounced as when varying the fixed cost.

Next we vary the holding cost, Figure 10, while all other values correspond to the default ones. Recall

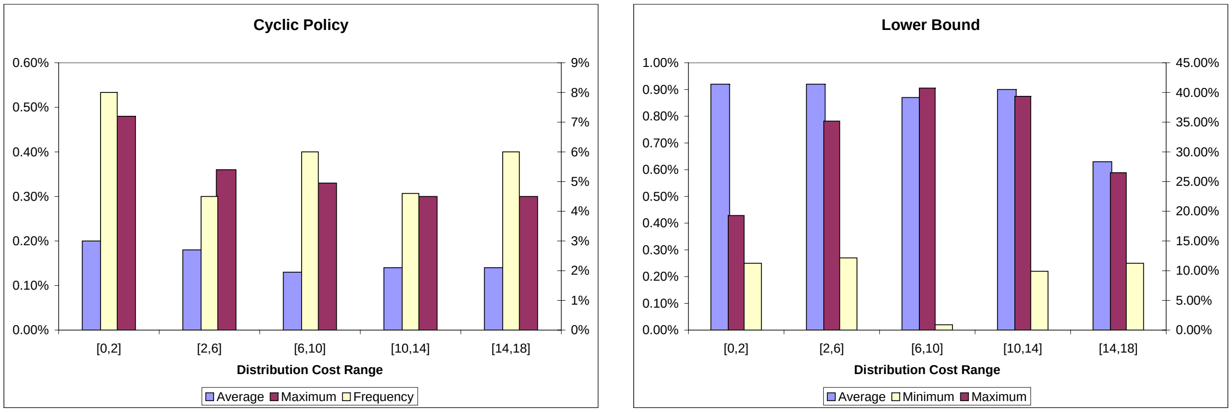

Figure 9: The effect of the distribution cost

that the holding cost is expressed relative to the procurement cost, which is constant in this experiment. Thus in this experiments we vary only r_i. It is clear that with the increased holding cost, the 2-slope cyclic policy becomes better. Likewise, the optimality gap keeps increasing.

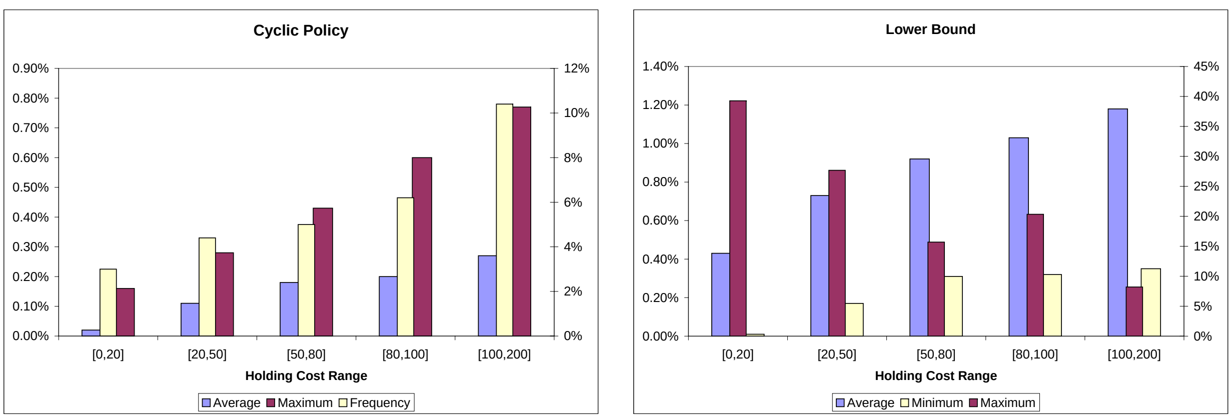

Figure 10: The effect of holding cost

To wrap up the sensitivity analysis we study the effect of different demand ranges. We are interested in the ratio of the remote demand and the local demands. The demand ratio is defined as the average demand of the two local markets over the average demand of the remote market. Note that the two local markets have always the same range. For example, if $L_1, L_2 \in [3000, 7000], L \in [1000, 3000]$, then the demand ratio is $5,000/2,000 = 2.5$. The demand ratio of the default data is 1.0. The results are shown in Figure 11. All of the other parameters are kept at their default values. With the increased demand ratio, we observe a decreasing trend in the frequency when the 2-slope cyclic policy outperforms the trivial policy. On the other hand, the maximum and the average improvement gaps keep increases (with the exception of the demand ratio 1.0). It seems that the 2-slope cyclic policy less frequently makes an improvement, however, whenever it does, the improvement is larger. This has not been observed in the previous charts. The optimality gap shows no trend.

We do not show the sensitivity analysis with respect to the procurement cost since the computational experiments revealed that the effectiveness of the 2-slope cyclic policy and the optimality gaps are indifferent to this cost.

In Figure 12 we show the corresponding histograms with respect to the default ranges. For the 2-slope cyclic policy chart, we consider only those instances when an improvement is made. Based on the left chart

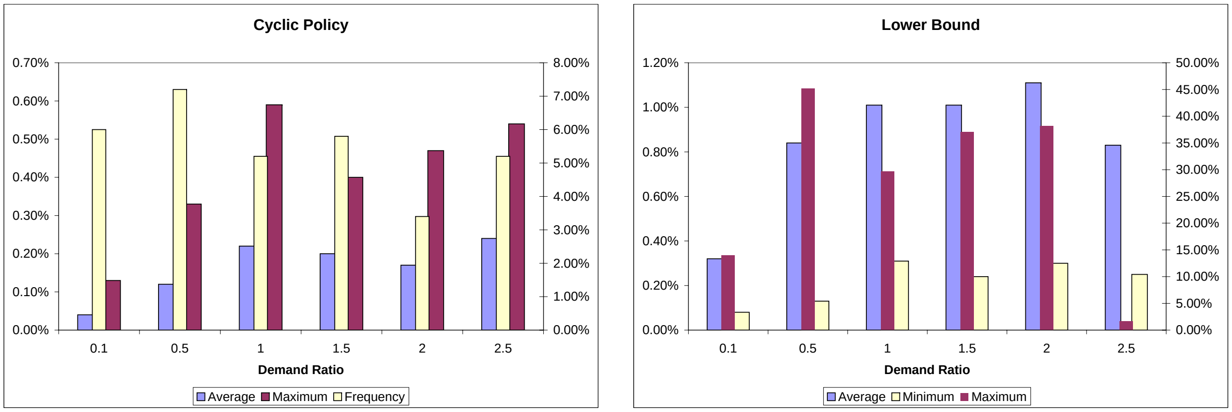

Figure 11: The effect of the demand ratio

we see that most of the time only a slight improvement is made. Very seldom the improvements larger than 0.50% occur. This clearly explains why is the average improvement gap low. The right figure shows the optimality gap, which exhibits the usual normal distribution like shape. Large optimality gaps of 40% happen rarely. Unfortunately zero gap instances are non existing.

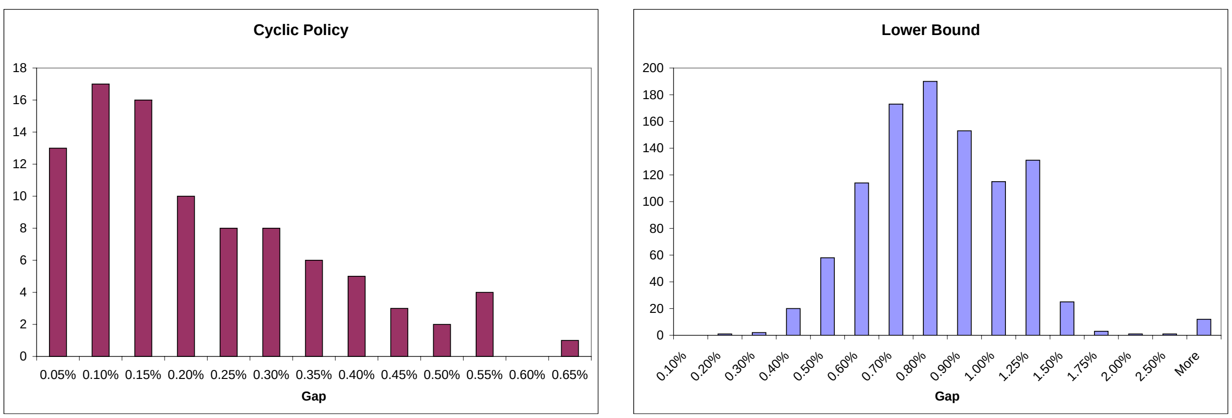

Figure 12: Histograms

6.2 3-slope Cyclic Policy and the General Case with 2 Facilities and 2 Markets

In this section we assume that $n = 2, m = 2$ and thus there is a distribution cost from both facilities to either of the two markets. Due to higher computational times (approximately a second for each one of the policy and the lower bound), in this section for each experiment we generated 500 instances.

We also study a new set of policies. A *3-slope cyclic policy* is depicted in Figure 13. In this policy during a period or cycle facility 1 replenishes once and facility 2 replenishes twice. A similar policy can be obtained by imposing that facility 1 replenishes twice and facility 2 once during the period. We call the 3-slope cyclic policy the best of the two policies. We can compute the best 3-slope cyclic policy by solving two nonlinear optimization problems that are given in Klabjan (2009). We call the 2-slope cyclic and the 3-slope cyclic policies simply the *slope policies*.

We consider 3 cases, which are summarized in Table 1. We vary all of the parameters except the distribution costs. The effect of the distribution costs is studied later. The values are selected in such a way that we have a variety of different cost parameters for all cost components except the distribution cost.

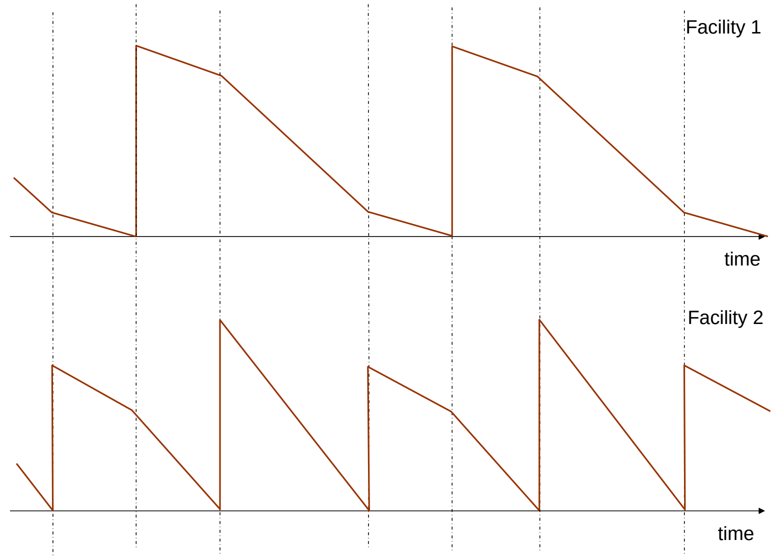

Figure 13: The 3-slope cyclic policy

	k_i	c_i	f_{ij}	r_i	L_j
Case 1	[100,300]	[5,10]	[0.1,2.1]	[80%,100%]	[2000,5000]
Case 2	[1000,1500]	[5,10]	[0.1,2.1]	[100%,300%]	[5000,10000]
Case 3	[500,1000]	[2,5]	[0.1,2.1]	[80%,100%]	[3000,6000]

Table 1: The input data for the three cases

The results are shown in Table 2. The table shows the effectiveness of the 3-slope cyclic policy, the performance of slope policies, and the optimality gap under this more general setting. Frequency and maximum have the same meaning as in Section 6.1. The first row "slopes vs. trivial" shows the relevant values for the best value between the slope policies and the trivial policy. The average and the maximum values correspond to the average and the maximum improvement of the slope policies with respect to the trivial policy. The second row "gap" shows similar statistics but we compare the best of the policies and the lower bound. The frequency value shows how often the optimality gap equals zero. The next row compares the 3-slope cyclic policy versus the 2-slope cyclic policy. The frequency shows how often the 3-slope cyclic policy outperforms the 2-slope cyclic policy. The average and the maximum columns show the average and the maximum improvement of the 3-slope cyclic policy with respect to the 2-slope cyclic policy. The last row compares the 3-slope cyclic policy versus the best value among all 3 policies. The frequency column shows how often is the 3-slope cyclic policy the best one. The remaining two values show the improvement whenever there is one.

The improvements of the slope policies are at the same level as those already observed. The improvement happens rarely (less than 4%) and the average improvement is low. The results with respect to the optimality gap are improved. First, note that the optimality gap is often zero (in more than half of the instances), which means that the best of the trivial and the slope policy is optimal in many cases. The average optimality gap is relatively low and the maximum gap is low as well (less than 8% in two cases and 29% in Case 2). The maximum optimality gap is substantially lower than in Section 6.1, where it was observed to be as high as 45%. The last two rows reveal an interesting fact about the 3-slope cyclic policy. The next to the last row shows that the 3-slope cyclic policy almost always outperforms the 2-slope cyclic policy and

	Frequency			Average			Maximum		
	Case 1	Case 2	Case 3	Case 1	Case 2	Case 3	Case 1	Case 2	Case 3
Slopes vs. trivial	3.40%	0.40%	3.20%	0.17%	0.42%	0.26%	0.44%	0.55%	0.75%
Gap	61.00%	66.00%	53.40%	0.30%	0.32%	0.66%	5.34%	28.20%	7.62%
3-slope vs. 2-slope	95.40%	98.80%	94.40%	1.01%	2.08%	1.88%	2.31%	4.69%	4.15%
3-slope best	0.04%	0.00%	0.01%	0.09%	0.00%	0.17%	0.10%	0.00%	0.24%

Table 2: Effectiveness of using the 3-slope cyclic policy

the improvements can be quite high (up to 5%). On the other hand, as the last row shows, the 3-slope cyclic policy is rarely the best policy among those considered. From these findings we observe an interesting anomaly. In almost all of the cases, whenever the 2-slope cyclic policy outperforms the trivial policy, the 3-slope cyclic policy is worse than the 2-slope cyclic policy. On the other hand, whenever the 2-slope cyclic policy is worse than the trivial policy, the 3-slope cyclic policy outperforms the 2-slope cyclic policy but not the trivial policy.

In Table 3 we study the effect of the distribution costs. Case 1 instance has the distribution cost in the range $[0.1, 2]$ and the next one has $f_{ij} \in [1, 5]$ for every $i = 1, 2, j = 1, 2$. In the third case we pick a random arc among the four of them and we assign a random value in $[0, 2]$. Next we randomly pick an arc among the remaining three arcs and give value $[10, 12]$. We next select a random arc between the two remaining arcs and assign the distribution cost in the range $[20, 22]$. The last arc gets the cost in the range $[30, 32]$. Case 4 is generated in the same way except that the interval length 2 is replaced by 5. All of the remaining values equal to the default values in Section 6.1. Average, maximum, and frequency columns correspond to the improvement of the slope policies with respect to the trivial policy. The interpretation of the fifth and the sixth column are identical to the corresponding columns in Table 2. The last four columns correspond to the comparison of the best policy value and the lower bound. The last column shows the average optimality gap among all instances that produce a positive optimality gap.

	Policy					Lower bound			
	Average	Maximum	Frequency	3-slope best	3-slope vs. 2-slope	gap=0	Average	Maximum	Average Positive
Case 1	0.00%	0.00%	0.00%	0.00%	71.00%	24%	0.66%	1.93%	0.87%
Case 2	0.00%	0.00%	0.00%	0.00%	77.00%	33%	0.57%	3.51%	0.75%
Case 3	0.09%	0.28%	0.28%	0.80%	94.20%	45%	0.34%	25.47%	0.61%
Case 4	0.06%	0.19%	2.60%	0.60%	95.80%	64%	0.21%	37.52%	0.59%

Table 3: The effect of distribution costs in the case $n = 2, m = 2$

From Table 3 we observe that if all of the distribution costs are in the same range, then the slope policy is ineffective. In this case it is better to simply run the trivial policy. As the distribution cost ranges vary between facilities and markets (Case 3 and Case 4), the slope policies become more effective. In Case 4, in more than 2% of the instances the slope policies outperform the trivial policy. In these two cases, the 3-slope cyclic policy is the best policy on average on 0.7% of the instances, which is better than what was observed in Table 2. Unfortunately the improvements are not large. The lower bound as well becomes more effective in Cases 3 and 4. In Case 4 it is zero in more than half of the cases. In all of the cases the average optimality gap is low and it is the lowest in Case 4.

6.3 Several Markets

In this section we study the case of 2 facilities and several markets. In particular, we consider $m = 2, 6, 13, 26$. The 2-slope cyclic policy can easily be extended to the case of a general number of markets m. The underlying optimization problem for computing an optimal 2-slope cyclic policy in this general case is given in Klabjan (2009). Each reported value is averaged over 500 randomly generated instances.

The results are summarized in Table 4. The input data correspond to the default data given in Section 6.1. This table shows the benefits of using the 2-slope cyclic policy versus the trivial policy. As the number of markets m increases, the 2-slope cyclic policy more frequently outperforms the trivial policy. For $m = 26$, this frequency is 11%, which is much better than all of these values observed so far. On the downside, the relative improvements stay low. They actually decrease with an increasing number of markets.

	Average	Maximum	Frequency
2x2	0.06%	0.19%	2.60%
2x6	0.01%	0.41%	6.00%
2x13	0.03%	0.11%	7.80%
2x26	0.02%	0.04%	11.00%

Table 4: Increasing number of markets

6.4 Concluding Remarks

We next summarize our main findings with respect to policies and the lower bound based on the computational experiments. The most important conclusion is that the trivial policy is very efficient and in many cases very close to the optimal policy (less than 1% optimality gap). The slope policies become more effective when the fixed and the holding costs are high. In addition, if the distribution costs vary among facilities and markets, then slope policies more frequently outperform the trivial policy. Their benefits also increase with an increased number of markets. With 26 markets they outperform the trivial policy 11% of the times. However, the improvements are not large, on average approximately 0.02%.

The lower bound usually produces very good bounds. The average optimality gap is typically around 0.5% and approximately half of the times there is no gap at all. However, the gap does not appear to be bounded by a constant since we have encountered cases with a 50% gap.

References

Adelman, D. and Mersereau, A. J. (2004). Relaxations of weakly coupled stochastic dynamic programs. Technical report, Graduate School of Business, The University of Chicago, Chicago, IL.

Allgor, R., Graves, S., and Xu, P. (2004). Traditional inventory models in an E-retailing setting: A two-stage serial system with space constraints. In *Proceedings of 2004 SMA Conference*, St. Pete Beach, FL.

Alptekinoglu, A. and Tang, C. (2005). A model for analyzing multi-channel distribution systems. *European Journal of Operational Research*, **163**, 802–824.

Bagga, J., Jain, A., and Palekar, U. (2005). An evolution of fulfillment strategies for business-to-consumer electronic commerce. Technical report, Department of Mechanical and Industrial Engineering, University of Illinois at Urbana-Champaign, Urbana, IL.

Bernstein, F., Song, J.-S., and Zheng, X. (2005). Free riding in a multi-channel supply chain. Technical report, The Fuque School of Business, Duke University, Durham, NC.

Bhattacharya, R. and Majumdar, M. (1989). Controlled semi-Markov models under long-run average rewards. *Journal of Statistical Planning and Inference*, **22**, 223–242.

Blumenfeld, D., Burns, L., Diltz, L., and Daganzo, C. (1985). Analyzing trade-offs between transportation, inventory and production costs on freight networks. *Transportation Research Part B: Methodological*, **19**, 361–380.

Blumenfeld, D., Burns, L., and Daganzo, C. (1991). Synchronizing production and transportation schedules. *Transportation Research Part B: Methodological*, **25**, 23–27.

Boyaci, T. (2005). Competitive stocking and coordination in a multiple-channel distribution system. *IIE Transactions*, **37**, 407–427.

Boyaci, T. and Gallego, G. (2002). Coordinating pricing and inventory replenishment policies for one wholesaler and one or more geographically dispersed retailers. *International Journal of Production Economics*, **77**, 95–111.

Burns, L., Blumenfeld, R. H. D., and Daganzo, C. (1985). Distribution strategies that minimize transportation and inventory costs. *Operations Research*, **33**, 469–490.

Campbell, A. and Savelsbergh, M. (1998). The inventory routing problem. In T. G. Crainic and G. Laporte, editors, *Fleet Management and Logistics*. Kluwer Academic Publishers.

Campbell, A. and Savelsbergh, M. (2002). Inventory routing in practice. In P. Toth and D. Vigo, editors, *SIAM Monographs on Discrete Mathematics and Applications*. Society for Industrial and Applied Mathematics.

Cattani, K., Gilland, W., and Swaminathan, J. (2003). Adding a direct channel? How autonomy of the direct channel affects prices and profits. Technical report, Kenan-Flagler Business School, The University of North Carolina at Chapel Hill, Chapel Hill, NC.

Cattani, K., Gilland, W., and Swaminathan, J. (2004). Coordinating traditional and internet supply chains. In D. Simchi-Levi, D. Wu, and M. Shen, editors, *In Supply Chain Analysis in the eBusiness Era*. Kluwer Academic Publishers.

Chiang, W. and Monahan, G. (2005). Managing inventories in a two-echelon dual-channel supply chain. *European Journal of Operational Research*, **162**, 325–341.

Chiang, W., Chhajed, D., and Hess, J. (2003). Direct marketing, indirect profits: A strategic analysis of dual-channel supply-chain design. *Management Science*, **29**, 1–20.

Dror, M. (2005). Routing propane deliveries. In A. Langevin and D. Riopel, editors, *Logistics Systems: Design and Optimization*. Springer Verlag.

Goyal, S. K. and Satir, A. T. (1989). Joint replenishment inventory control: Deterministic and stochastic models. *European Journal of Operational Research*, **38**, 2–13.

Grosso, C., McPherson, J., and Shi, C. (2005). Retailing: What's working online. *The McKinsey Quarterly*, **3**.

Hahm, J. and Yano, C. (1992). The economic lot and delivery scheduling problem: the single item case. *International Journal of Production Economics*, **28**, 235–252.

Hall, R. (1996). On the integration of production and distribution: Economic order and production quantity implications. *Transportation Research Part B: Methodological*, **30**, 387–403.

Hawkins, J. (2003). *A Lagrangian Decomposition Approach to Weakly Coupled Dynamic Optimization Problems and its Applications*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Klabjan, D. (2009). Business-to-consumer single period fulfillment strategies. Technical report, Department of Industrial Engineering and Management Sciences, Northwestern University, IL.

Klabjan, D. (2009). Appendix: Economic order quantity for business-to-consumer fulfillment. Technical report, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL. Available from http://www.klabjan.dynresmanagement.com/.

Klabjan, D. and Adelman, D. (2006). Existence of optimal policies for semi-Markov decision processes using duality for infinite linear programming. *SIAM Journal on Control and Optimization*, **44**, 2104–4122.

Luque-Vásquez, F. and Hernández-Lerma, O. (1999). Semi-Markov control models with average costs. *Applicationes Mathematicae*, **26**, 315–331.

Meuleau, N., Hauskrecht, M., Kim, K., Peshkin, L., Kaelbling, L., Dean, T., and Boutilier, C. (1998). Solving very large weakly coupled Markov decision processes. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 165–172, Madison, WI.

Russell, R. and Krajewski, L. (1991). Optimal purchase and transportation cost lot sizing for a single item. *Decision Scienses*, **22**, 940–952.

Swensetha, S. and Godfrey, M. (2002). Incorporating transportation costs into inventory replenishment decisions. *Internation Journal of Production Economics*, **77**, 113–130.

Tsay, A. and Agrawal, N. (2000). Channel dynamics under price and service competition. *Manufacturing & Service Operations Management*, **2**, 372–391.

Tsay, A. and Agrawal, N. (2004). Channel conflict and coordination in the E-commerce age. *Production and Operations Management*, **13**, 93–110.

Vega-Amaya, O. (1993). Average optimality in semi-Markov control models on Borel spaces: unbounded cost and controls. *Boletín de la Sociedad Matemática Mexicana*, **38**, 47–60.

Zipkin, P. (2000). *Foundations of Inventory Management*. McGraw-Hill.