

A Constant Approximation Algorithm for the One-Warehouse Multi-Retailer Problem

Retsef Levi* Robin Roundy† David B. Shmoys‡ Maxim Sviridenko §

Submitted January 2004, Revised May 2006, Revised March 2007

Abstract

Deterministic inventory theory provides streamlined optimization models that attempt to capture tradeoffs in managing the flow of goods through a supply chain. We will consider two well-studied deterministic inventory models, called the *one-warehouse multi-retailer problem* (OWMR) and its special case the *joint replenishment problem* (JRP), and give approximation algorithms with worst-case performance guarantees. That is, for each instance of the problem, our algorithm produces a solution with cost that is guaranteed to be at most 1.8 times the optimal cost; this is called a 1.8-approximation algorithm. Our results are based on an LP-rounding approach; we provide the first constant approximation algorithm for the OWMR problem and improve the previous results for the JRP problem.

*retsef@mit.edu. Sloan School of Management, MIT, Cambridge, MA, 02139. This research was partially conducted while the author was a Phd student in the ORIE department at Cornell University. Research supported partially by a grant from Motorola and NSF grants CCR-9912422&CCR-0430682.

†robin@orie.cornell.edu. School of ORIE, Cornell University, Ithaca, NY 14853. Research supported partially by a grant from Motorola and NSF grants DMI-0075627&DMI-0500263, and the Querétaro Campus of the Instituto Tecnológico y de Estudios Superiores de Monterrey.

‡shmoys@cs.cornell.edu. School of ORIE and Dept. of Computer Science, Cornell University, Ithaca, NY 14853. Research supported partially by NSF grants CCR-9912422&CCR-0430682&DMI-0500263.

§sviri@us.ibm.com. IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598.

1 Introduction

Deterministic inventory theory provides streamlined optimization models that attempt to capture tradeoffs in managing the flow of goods through a supply chain. We will consider two well-studied inventory models, the *one-warehouse multi-retailer problem* (OWMR) and its special case the *joint replenishment problem* (JRP). Using LP-rounding techniques we provide the first constant approximation algorithm for the OWMR problem. That is, for each instance of the problem, our algorithm produces a solution with cost that is guaranteed to be at most C times the optimal cost, for some constant $C > 1$. The constant C is called the *worst-case guarantee* of the algorithm. Moreover, when specialized to the JRP model, our LP-rounding approach provides worst-case guarantees that improve on previous approximation algorithms for this problem [12, 13].

As the name suggests, in the OWMR model there is one warehouse that orders a particular commodity from a supplier, in order to serve demand at N distinct retailers. We consider a discrete finite planning horizon of T periods, and are given the demand $d_{it} \geq 0$ required for each retailer $i = 1, \dots, N$, in each time period $t = 1, \dots, T$. There are two types of costs incurred: *ordering costs* (to model that there are fixed costs incurred each time the warehouse replenishes its supply on hand from the supplier, as well as the analogous cost for each retailer to be stocked from the warehouse) and *holding costs* (to model the fact that maintaining inventory, at both the warehouse and the retail store, incurs a cost). The aim of the model is to provide an optimization framework to balance the fact that ordering too frequently is inefficient for ordering costs, whereas ordering too rarely incurs excessive holding costs.

The dynamics of the OWMR model are as follows. At the beginning of each period s , each retailer i can place an order for any number of units from the warehouse, to replenish its on-hand inventory. The order is assumed to arrive instantaneously (this is without loss of generality), and can be used to satisfy demand in period s , or in subsequent periods. Any such order placed by retailer i incurs a *fixed* ordering cost K^i , which is independent of the size of the order and of the time period in which the order is placed. However, all orders placed by the different retailers, in each period s , must be satisfied only from the on-hand inventory at the warehouse in that period. So in turn, at the beginning of each period r the warehouse can place an order for any number of units from the supplier. This order is again assumed to arrive instantaneously, and can be used to satisfy retailer orders in period r , or in subsequent periods. Any such order of the warehouse in period r incurs a *fixed* ordering cost K_r^0 , which also is independent of the size of the order and the combination of items being ordered. All demands must be satisfied on time, i.e., any unit that is used by retailer i to satisfy its demand in period t , d_{it} , must be ordered by the warehouse from the supplier in some period r , and then by retailer i from the warehouse in some period s , where $r \leq s \leq t$. (In the inventory literature, these assumptions are usually referred to as “neither back orders nor lost sales are allowed”.) The goal is to find a feasible ordering policy that satisfies all demands on time with minimum total ordering and holding costs. Throughout the paper, we will use $[r, s]$ ($r \leq s$) to denote a pair of warehouse and retailer orders in periods

r and s , respectively. We note that while the warehouse ordering cost K_r^0 is time-dependent, the retailer ordering cost K^i is stationary over time. It is easy to show that if we allow it to be time-dependent, then the OWMR problem becomes as hard as set-cover problem (see [4] for the details). Thus, it is not likely that there exists an approximation algorithm with a sub-logarithmic worst-case guarantee [7].

The standard models for holding cost make two natural linearity assumptions: (1) the cost is proportional to the number of units of the commodity held, and (2) there is cost associated with holding from period t to $t+1$, which is then additive over the period held. We use a more general holding cost structure, extending the model that has been introduced by Levi, Roundy and Shmoys for the JRP [12, 13]. While still maintaining (1), we generalize (2) in a way that preserves the most useful properties of an optimal solution (as well as of an optimal solution to a natural LP relaxation), but captures much more general phenomena, such as the notion of perishable goods (where the holding cost becomes infinite, when the good is held too long). Capturing the right generalization is subtle here, due to the nature of the interaction between the two levels, and this is outlined in Section 2; we introduce, in essence, a holding cost h_{rs}^{it} associated with ordering one unit of the demand at retailer i for period t according to the pair $[r, s]$ which is assumed to satisfy certain natural monotonicity properties.

The one-warehouse multi-retailer problem is a generalization of several classical inventory models, such as the *single-item lot-sizing problem* (in which there is, in effect, only one retailer and the warehouse holding and ordering costs are 0) and the *joint replenishment problem* (JRP) (where, in effect, the holding cost at the warehouse is enormous, and hence each unit of demand can be assumed to be satisfied by an order $[s, s]$). The general OWMR model has been studied extensively, and plays a fundamental role in broader planning issues, such as the management of supply chains.

Arkin, Joneja and Roundy [1] have shown that OWMR is NP-hard even for the special case of the JRP, where the warehouse serves only as a cross-docking point (i.e., no inventory is ever held at the warehouse). Federgruen and Tzur [6] have proposed an interesting heuristic based on dynamic programming. However, for the theoretical analysis of the worst-case performance of their algorithm, they have assumed that the cost parameters and the demands are bounded by uniform constants. Chan, Muriel, Shen, Shimchi-Levi and Teo [4] have considered a variant of OWMR, in which the ordering costs are piecewise-linear functions, and the holding cost is linear and additive. They considered the class of *zero-inventory ordering* (ZIO) policies, in which the warehouse and retailers order if and only if their current on hand inventory is 0. They established the effectiveness of these policies, showing that the cost of the optimal ZIO policy is at most $\frac{4}{3}$ times the cost of the optimal policy. In [4] and in a subsequent paper by Shen, Simchi-Levi and Teo [15], they have proposed an integer program to find the optimal ZIO policy, which is NP-hard. Next they have developed heuristics to round the optimal solution of the LP relaxation to get an approximation algorithm for finding the best ZIO policy. However, the performance guarantee of their algorithm is $O(\log(N + T))$. For the problem we consider in this paper, it is well known that ZIO policies are optimal.

Recently, Levi, Roundy and Shmoys [12, 13] have presented a general primal-dual algorithmic frame-

work that solves the single-location lot-sizing problem and provides a 2-approximation for the JRP and the *assembly problem*, which is yet another classical inventory model (see [12, 13]). It is an open question whether the primal-dual approach can be extended to work in the more general OWMR problem considered in this paper. The main barrier seems to be the more complex structure involved with holding inventory in two "levels" (the warehouse and retailers), which does not seem to preserve several properties that are essential for the analysis in [12, 13].

For the problem we consider in this paper, it is well known that ZIO policies are optimal [19]. We propose a natural integer program to find the optimal policy, which is different from the one proposed in [4] and [15]. We first solve the LP relaxation to optimality, and then introduce techniques to round this optimal solution to a feasible solution for the OWMR problem, which can be proven to be near-optimal. The rounding is done in two phases. In the first phase we determine the warehouse orders; based on that, we determine the retailer orders in the second phase, and this is done separately for each retailer. Our algorithms are based on new dependent randomized rounding techniques, that are similar in spirit to those used for the *metric facility location problem* [16], but are able to exploit the additional special structure of the inventory model. Specifically, we show that the solution produced by the randomized algorithms has expected cost that is guaranteed to be at most 1.8 times the cost of an optimal solution to the OWMR problem. We then show how to derandomize these algorithms and this yields a deterministic 1.8-approximation algorithm for the OWMR problem. When specialized to the JRP problem our LP is identical to the one used by Levi, Roundy and Shmoys [12, 13]. Thus, the LP-rounding approach can be applied to the JRP problem and improves on their primal-dual 2-approximation for the JRP problem. [12, 13]. The inventory models that are discussed in this paper are usually solved in practice via integer programming solution methods, such as branch and bound. We believe that our techniques can be naturally incorporated into these method to generate good feasible solutions and enhance the computational procedures.

One note on the relation between deterministic inventory models and the facility location problem is in order. If one thinks of orders as facilities and demands as customers, then deterministic inventory models can be viewed as special facility location problems. Nevertheless, the inventory models we consider are significantly different, since the holding cost structure, which plays the role of the assignment costs, is asymmetric and does not obey the triangle inequality. These are both essential assumptions in all of the existing approximation algorithms for the metric facility location problem. It is interesting that the additional structure of these inventory problems is sufficient to extend some of these techniques. (For a survey on the approximation techniques that were applied to the metric facility location problem, see [16].)

In the appendix, we also consider an important extension of the models above. In many real-life applications the ordering cost actually corresponds to *transportation cost*. Usually the transportation is based on trucks with a given capacity. We model this using *soft capacities*. Now we can order in batches each of capacity U , where for each batch we order (in a given period), we incur an *additional* fixed cost. We allow different batch capacities for the warehouse and the retailers, and then show how to extend the algo-

rithms developed for the OWMR problem to work in this more general model. In particular, we provide a 3.6-approximation algorithm for the OWMR and JRP problems with soft capacities, and a 2-approximation algorithm for the single-location lot-sizing (with time-dependent batch capacities). Here we are using ideas and techniques that were introduced by Jain and Vazirani in their seminal paper on the facility location problem [8]. Jin and Muriel also consider this model and propose several heuristics based on centralized and decentralized approaches [9].

As a by-product of our work, we prove lower bounds on the integrality gap of facility location inspired LP relaxations for several variants of the OWMR problem. Specifically, for the special case of the OWMR problem, the single-location lot-sizing problem, it can be shown that our rounding approach, when applied to the facility location-inspired linear program of this problem, yields an optimal solution. (As already mentioned, here we have only one retailer, and there is no warehouse. There is only one ordering cost K_s , in each period s , and holding costs as before.) This shows that the corresponding LP has an integer optima. Other proofs with very different styles were given by Krarup & Bilde [10], Barany, Van Roy, and Wolsey [2], Bertsimas, Teo and Vohra [3] and recently by Levi, Roundy and Shmoys [12]. Finally, for the single-location lot-sizing problem with soft capacities we show (in Appendix A) that the natural facility location inspired LP has an integrality gap of at most 2.

The rest of the paper is organized as follows. In Section 2, we discuss the holding cost structure that we use in this paper. In Section 3, we present an LP relaxation of the OWMR problem and discuss some of the properties of fractional optimal solutions of the LP. In Section 4, we describe our rounding algorithms and their worst-case analysis. The extension to soft capacities is considered in Appendix A. We then conclude with some open questions.

2 The holding cost structure

In most of the existing literature, the holding cost is modeled in the following way. For each period t , the warehouse and each retailer i have a per unit cost $h_t^i \geq 0$ ($i = 0, 1, \dots, N$) to hold one unit in inventory from period t to period $t + 1$. The holding cost incurred at the end of each period is a linear function of the on-hand inventory at the end of the period.

We model the holding cost in the following more general way. Consider a demand point (i, t) and a pair of potential orders $[r, s]$, where again r is the period in which the unit was ordered by the warehouse from the supplier, and s is the period in which it was ordered by retailer i from the warehouse ($r \leq s \leq t$). For each (i, t) and $[r, s]$, we let h_{rs}^{it} be the cost of holding one unit in the warehouse location over $[r, s]$, then sending it to retailer i (in period s), and holding it at the premises of retailer i over $[s, t)$. We assume that the holding cost parameters obey the following natural properties:

Property 1: *Non-negativity.* The parameters h_{rs}^{it} are assumed to be non-negative.

Property 2: *Monotonicity with respect to s .* Each retailer $i = 1, \dots, N$ has exactly one of the following

properties which applies to *all* demand points (i, t) (for $t = 1, \dots, T$). For each demand point (i, t) and warehouse order in period r ($r \leq t$), h_{rs}^{it} is *either* non-increasing in $s \in [r, t]$, *or* it is non-decreasing in $s \in [r, t]$. We partition the retailers into two sets accordingly. Let I_J be the set of retailers i such that h_{rs}^{it} is non-decreasing in s for each t and $r \leq t$ and call them *J-retailers*, and let I_W be the rest of the retailers, i.e., retailers i such that h_{rs}^{it} is non-increasing in s for each t and $r \leq t$, and call them *W-retailers*. In models with traditional holding cost structure, the set I_J corresponds to retailers for which it is cheaper to hold inventory at the retailer premises (i.e., $h_t^i \leq h_t^0$, for each t), and I_W corresponds to retailers for which it is cheaper to hold inventory at the warehouse (i.e., $h_t^i > h_t^0$, for each t). It is straightforward to see that in an optimal policy, the warehouse does not hold inventory of *J-retailers*. Instead, in each period in which the warehouse orders some amount of units for *J-retailers*, it is cheapest to distribute the complete amount immediately to these retailers. Thus, for these retailers it is sufficient to consider only pairs of orders $[s, s]$. Moreover, the joint replenishment problem is the special case where *all* of the retailers are *J-retailers*. We note that the partition of the retailers into these two types is a standard assumption in the literature. In particular, the OWMR problem is traditionally considered under the assumption that $h_t^i > h_t^0$ for each i and t .

Property 3: *Monotonicity with respect to r .* For each retailer i and some demand point (i, t) , fix the retailer order in some period s ($s \leq t$); we assume that h_{rs}^{it} is non-increasing in $r \in [1, s]$. Moreover, for each retailer $i \in I_J$ and a demand point (i, t) , we assume that, for each $r < r' \leq t$, the order $[r, r]$ is more expensive than the order $[r', r']$. This property captures the fact that in most of the common scenarios holding inventory for longer time is more expensive.

Property 4: *Monge Property.* For each demand point (i, t) with $i \in I_W$ and any four periods $r_2 < r_1 \leq s_2 < s_1 \leq t$, the inequality, $h_{r_2, s_1}^{it} + h_{r_1, s_2}^{it} \geq h_{r_2, s_2}^{it} + h_{r_1, s_1}^{it}$ is satisfied. This property implies that it is always cheapest for the warehouse to use a FIFO order in satisfying the orders of each retailer. As we shall show in Section 3, this property induces structural properties on the optimal solution of the corresponding LP relaxation.

One can easily verify that all of the above properties are satisfied under the traditional holding cost structure. Of course, the way we model the holding cost is much more general. In particular, it enables us to capture other very important phenomena, such as perishable commodities, where then the parameter h_{rs}^{it} can be equal to infinity. In addition, the fact that the holding costs are defined per demand point and not per period provides a transparent way to model situations in which serving different customers incurs different holding costs. We note that per unit ordering costs can be incorporated into the holding cost as long as we preserve the above mentioned properties. (We note that one can incorporate even demand-point-dependent ordering costs as long as the monotonicity properties above are preserved.)

3 A linear program

In this section, we will first present a natural formulation of the OWMR problem as an integer program. In the next section, we shall show how to round the optimal solution of the corresponding LP relaxation to a feasible solution for the OWMR problem, while increasing the cost by only a constant factor.

The formulation is based on the well-known fact that there exists an optimal solution to the OWMR problem in which each demand d_{it} is satisfied from a unique pair of orders $[r, s]$, where again $r \leq s \leq t$ (see [19] for details). By this we mean that the warehouse orders the entire demand d_{it} in some period $r \leq t$, and keeps it in inventory over the time interval $[r, s]$ ($r \leq s \leq t$). Then in period s , the entire demand d_{it} is ordered from the warehouse by retailer i and is kept in inventory (at the retailer's premises) until time t .

For each demand point (i, t) and a pair of orders $[r, s]$, such that $r \leq s \leq t$, we define $H_{rs}^{it} := h_{rs}^{it}d_{it}$ to be the total cost of providing the demand d_{it} from the pair of orders $[r, s]$. Let x_{rs}^{it} (for $r \leq s \leq t$) be a binary decision variable that is equal to 1 if demand point (i, t) (i.e., demand d_{it}) is satisfied by the pair of orders in periods r (warehouse order) and s (retailer i order). For each $i = 1, \dots, N$ and $s = 1, \dots, T$, let y_s^i be a binary decision variable that indicates whether retailer i placed an order in period s . Finally, let y_r^0 be a binary variable that indicates whether the warehouse placed an order in period r . This gives rise to the following LP relaxation:

$$\begin{aligned}
 \text{minimize} \quad & \sum_{r=1}^T y_r^0 K_r^0 + \sum_{i=1}^N \sum_{s=1}^T y_s^i K^i + \sum_{i=1}^N \sum_{t=1}^T \sum_{r,s:r \leq s \leq t} x_{rs}^{it} H_{rs}^{it} & (P) \\
 \text{subject to} \quad & \sum_{r,s:r \leq s \leq t} x_{rs}^{it} = 1, & i = 1, \dots, N, t = 1, \dots, T, d_{it} > 0, & (1) \\
 & \sum_{r:r \leq s} x_{rs}^{it} \leq y_s^i, & i = 1, \dots, N, t = 1, \dots, T, s = 1, \dots, t & (2) \\
 & \sum_{s:r \leq s \leq t} x_{rs}^{it} \leq y_r^0, & i = 1, \dots, N, t = 1, \dots, T, r = 1, \dots, t & (3) \\
 & x_{rs}^{it}, y_r^i \geq 0, & i = 0, \dots, N, s = 1, \dots, T, r = 1, \dots, s, & (4) \\
 & & t = s, \dots, T. &
 \end{aligned}$$

Constraint (1) ensures that each positive demand point (i, t) is fully satisfied no later than period t . Constraint (2) ensures that no demand d_{it} can be satisfied by a retailer order in period $s \leq t$ (and some warehouse order in period $r \leq s$), unless retailer i indeed placed an order in period s . Lastly, constraint (3) ensures that no demand point d_{it} can be satisfied by a warehouse order in period r (and some retailer order $r \leq s \leq t$), unless the warehouse placed an order in period r . It is straightforward to see that the corresponding integer program provides a correct formulation to the OWMR problem. Hence, the LP relaxation above provides a lower bound on the cost of any feasible solution to the OWMR problem. For the rest of this paper we let (\hat{x}, \hat{y}) and opt_{LP} be the optimal solution and the value of (P) , respectively.

We note again that for each retailer i in I_J , it suffices to consider only the variables x_{rs}^{it} with $r = s$ (the warehouse does not hold inventory of J -retailers).

Lemma 3.1 *There exists an optimal solution to the OWMR problem where each J -retailer order is placed in a period in which there is also a warehouse order.*

Proof : Assume that there exists an optimal solution to the OWMR problem, in which a retailer order of some J -retailer i is placed in some period s' , where there is no warehouse order placed in that period. Let r' be the latest warehouse order placed by time period s' . That is, there is no warehouse order placed in each of the periods $r \in (r', s']$. Since the holding costs are monotonic in r (Property 2 of the holding costs) and the solution is assumed to be optimal, we can assume, without loss of generality, that all the units ordered by retailer i in period s' were ordered by the warehouse in period r' . Since i is a J -retailer, it is clear that canceling the order in period s' and placing instead a retailer order at r' will not increase the retailer ordering cost and will decrease the overall holding costs incurred by retailer i (Property 3 of the holding costs). The lemma then follows. ■

Consequently, for each retailer $i \in I_J$, we can adapt accordingly the constraints (1), (2) and (3). In particular, for each $i \in I_J$ and each period s , the modified constraints (2) and (3) are $x_{ss}^{it} \leq y_s^i$ and $x_{rr}^{it} \leq y_r^0$, respectively. It is easy to see that, in an optimal solution, we must have $y_s^i \leq y_s^0$, for each period $s = 1, \dots, T$. Next we discuss several structural properties of the optimal solution (\hat{x}, \hat{y}) that will be used throughout the rest of this paper.

3.1 Structural Properties of the Optimal Solution of (P)

The Monge Property. Recall the Monge property of the holding cost, i.e., Property 4 of h in Section 2. We say that a feasible solution (x, y) to (P) satisfies the Monge property, if $x_{rs}^{it} > 0$ ($r \leq s \leq t$) implies that $x_{\tilde{r}, \tilde{s}}^{it} = 0$ for any $[\tilde{r}, \tilde{s}]$ such that $\tilde{r} < r$ and $\tilde{s} > s$. Without loss of generality, we assume that (\hat{x}, \hat{y}) (the optimal solution of (P)) satisfies the Monge property. We note that because of the Monge property on the holding cost, any feasible solution to (P) can be converted in polynomial time to one that satisfies the Monge property and has no greater cost.

The Greedy Usage Property. We claim that there exists an optimal solution (\hat{x}, \hat{y}) to (P) with the property that, for each demand point (i, t) and a retailer- i order in period $s \leq t$, we have $\sum_{r \in [1, s]} \hat{x}_{rs}^{it} = \hat{y}_s^i$, except for possibly the earliest retailer- i order that fractionally serves (i, t) in the solution (\hat{x}, \hat{y}) . (By the earliest retailer- i order that fractionally serves (i, t) we mean \bar{s} , such that $\sum_{r \in [1, \bar{s}]} \hat{x}_{r\bar{s}}^{it} > 0$ and $\sum_{r \in [1, s]} \hat{x}_{rs}^{it} = 0$, for each $s < \bar{s}$.) In particular, define an *open fractional order* of the warehouse or some retailer i to be a period s with $\hat{y}_s^0 > 0$ or $\hat{y}_s^i > 0$, respectively. Consider some positive demand point (i, t) and the sequence of open retailer- i fractional orders in (\hat{x}, \hat{y}) over the time interval $[1, t]$. Intuitively, the *greedy usage property* means that in the optimal solution (\hat{x}, \hat{y}) , demand point (i, t) fractionally uses the open retailer (fractional) orders in a *greedy* manner from latest to earliest. One of the implications of this property is that each open fractional retailer- i order $s' \in [1, t]$ (i.e., $\hat{y}_{s'}^i > 0$), such that $\sum_{s \in [s', t]} \hat{y}_s^i < 1$, is fully used by (i, t) . That is,

$\sum_{r \in [1, s']} \hat{x}_{rs'}^{it} = \hat{y}_{s'}^i$. (In other words, Constraint (2) is tight.)

The greedy usage property follows from the monotonicity properties of h , specifically from Properties 2 and 3. For any feasible solution for (P) that does not satisfy the greedy usage property, there exists another feasible solution that does satisfy the property and has an objective values that is not higher. In particular, assume that there exist two retailer orders $s' < s$, such that for some (i, t) we have $\sum_{r \in [1, s]} \hat{x}_{rs}^{it} < \hat{y}_s^i$ and $\sum_{r \in [1, s']} \hat{x}_{rs'}^{it} > 0$. Let $r' \leq s'$ be such that $\hat{x}_{r's'}^{it} > 0$ and $\epsilon = \min\{\hat{x}_{r's'}^{it}, \hat{y}_s^i - \sum_{r \in [1, s]} \hat{x}_{rs}^{it}\}$. If $i \in I_W$, it is straightforward to verify that by increasing $\hat{x}_{r's}^{it}$ by ϵ and decreasing $\hat{x}_{r',s'}^{it}$ by ϵ , we get a feasible solution with objective values that is not higher. (This follows from Property 2.) If $i \in I_J$, then $r' = s'$ and by Property 2 it follows that increasing \hat{x}_{ss}^{it} by ϵ and decreasing $\hat{x}_{s',s'}^{it}$ by ϵ result in a feasible solution with no greater cost. (This follows from Property 3.)

4 The Random Shift Algorithms

In this section, we will show how to round the optimal solution of (P), denoted again by (\hat{x}, \hat{y}) , to a feasible solution to the OWMR problem with cost at most 1.8 times the optimal cost. We shall first describe two different randomized rounding procedures that we call *random shift with retailer two-sided push* and *random shift with retailer one-sided push*. Our rounding procedures run in two phases. In the first phase, we determine the warehouse orders, using a simple mechanism that we call *random shift*. In the second phase, we use the output of the first rounding phase to determine the orders of each retailer. This phase is done separately for each retailer. We shall show that the expected cost of each one of the algorithms is guaranteed to be at most twice the cost of an optimal policy for the OWMR problem. In the worst-case analysis we shall bound each part of the cost, i.e., the warehouse ordering cost, the retailer ordering cost and the holding cost, using the respective part in the cost incurred by the optimal fractional solution (\hat{x}, \hat{y}) . Moreover, we show that the algorithm with the cheapest expected cost among the two is guaranteed have expected cost at most 1.8 times the optimal cost of the OWMR problem. Finally, we describe how to derandomize the algorithms and get a deterministic approximation algorithm with worst-case performance guarantee of 1.8. That is, for each instance of the problem, the algorithm produces a solution that is guaranteed to be at most 1.8 times the cost of an optimal policy.

4.1 The Random Shift Procedure

We first describe the *random shift* procedure that is used in the first phase of the algorithms in which we decide in what periods to place warehouse orders. This simple randomized procedure is based on the values $\hat{y}_1^0, \dots, \hat{y}_T^0$.

For the description of the random shift procedure, consider the interval $(0, \sum_{r=1}^T \hat{y}_r^0]$, which corresponds to the total weight of open fractional warehouse orders in the optimal fractional solution (\hat{x}, \hat{y}) . Each period $m = 1, \dots, T$ is then associated with the respective interval $\hat{Y}_m^0 = (\sum_{r=1}^{m-1} \hat{y}_r^0, \sum_{r=1}^m \hat{y}_r^0]$, which is of length

\hat{y}_m^0 . In particular, some periods can correspond to empty intervals of length 0 (if $\hat{y}_m^0 = 0$). The input for this procedure is a *step parameter* $c \in (0, 1]$. Given c , choose a *shift parameter* α_0 uniformly at random from $(0, c]$. Let W be the smallest integer multiple of c that is greater than $\sum_{r=1}^T \hat{y}_r^0$. Specifically, W is the upper ceiling of the total accumulated weight of fractional warehouse orders in the optimal LP solution (\hat{x}, \hat{y}) scaled by $\frac{1}{c}$; that is, $W = \lceil \frac{1}{c} \sum_{r=1}^T \hat{y}_r^0 \rceil$. Note that the interval $(0, \sum_{r=1}^T \hat{y}_r^0]$ is contained in the interval $[0, cW]$. Within the interval $[0, cW]$ focus on the sequence of points $0, c, \dots, c(W-1)$. The shift parameter α_0 induces a sequence of what we call *warehouse shift points*. Specifically, the set of warehouse shift points is defined as $\{\alpha_0 + cw : w = 0, \dots, W-1\}$. This set is constructed through a *shift* of random length α_0 to the right of the points $0, c, \dots, c(W-1)$. Thus, there are W shift points that are all located within the interval $[0, cW]$. Observe that the sequence of warehouse shift points is a-priori random and is realized with the shift parameter α_0 (see Figure 4.1).




Figure 4.1: Each interval $(\]$ corresponds to some period r of length of \hat{y}_r^0 . The warehouse shift points (black bullets) are generated by shifting the points $0, c, 2c, \dots, (W-1)c$ to right by α_0 . A warehouse order is placed in period r , if there is at least one shift point within its corresponding interval (e.g., periods 1 and 3 in the picture).

The warehouse shift points determine the periods in which warehouse orders are placed. For each period $m = 1, \dots, T$, we place a warehouse order in that period if there is at least one shift point within the interval \hat{Y}_m^0 that is associated with m . That is, we place a warehouse order in period m , if for some integer $0 \leq w \leq W-1$ there exists a warehouse shift point $\alpha_0 + cw$ that falls within the interval \hat{Y}_m^0 .

Next we bound the expected warehouse ordering cost incurred by the random shift procedure.

Lemma 4.1 Consider the random shift procedure described above with input length parameter $c \in (0, 1]$.

Then, the total expected warehouse ordering cost of the random shift procedure, denoted by \mathcal{K}_0 is at most $\frac{1}{c}$ times the total warehouse ordering costs in the optimal LP solution. That is, $\mathcal{K}_0 \leq \frac{1}{c} \sum_{r=1}^T \hat{y}_r^0 K_r^0$.

Proof : For each $w = 0, \dots, W - 1$ the interval $(cw, c(w + 1))$ generates at most one warehouse order. Moreover, in each interval $(cw, c(w + 1))$, there is exactly one warehouse shift point that is uniformly distributed over the interval. Thus, the expected cost of the warehouse order generated by the interval $(cw, c(w + 1))$ is at most $\frac{1}{c} \sum_{m=1}^T |\hat{Y}_m^0 \cap (cw, c(w + 1))| K_m^0$. It follows that the overall expected warehouse ordering cost is at most

$$\frac{1}{c} \sum_{w=0}^{W-1} \sum_{m=1}^T |\hat{Y}_m^0 \cap (cw, c(w + 1))| K_m^0 = \frac{1}{c} \sum_{m=1}^T K_m^0 \sum_{w=0}^{W-1} |\hat{Y}_m^0 \cap (cw, c(w + 1))| = \frac{1}{c} \sum_{m=1}^T \hat{y}_m^0 K_m^0,$$

where the last equality follows from the fact that each interval \hat{Y}_m^0 is partitioned by the intervals $(cw, c(w + 1))$. The proof of the lemma then follows. \blacksquare

Let $\mathcal{T}_W := \{r_1 < r_2 < \dots < r_M\}$ be the set of periods of the warehouse orders as determined in the first phase of the algorithm using the random shift procedure. Note that constraints (1) and (3) imply that if t is the earliest period with a positive demand point (i.e., the earliest demand point with $d_{it} > 0$), then $\sum_{r=1}^t \hat{y}_r^0 \geq 1$. Moreover, by the properties of the random shift procedure described above, it is straightforward to verify that there will be at least one warehouse order placed in the interval $[1, t]$, i.e., $r_1 \in [1, t]$. This implies that each positive demand point can be satisfied by at least one warehouse order that is placed earlier in time.

Once we decide upon the warehouse orders, then the OWMR problem decomposes into N single-location, single-item lot-sizing problems. These problems can be solved optimally using dynamic programming (see, for example, [18, 5]) to achieve the minimum overall retailer ordering cost and holding cost under the assumption that warehouse orders are placed at $r_1 < r_2 < \dots < r_M$. The collection of the solutions to these single-location problems provides a solution to the OWMR problem. However, as part of the worst-case analysis, we next describe two algorithms that use the random shift in the first phase, but determine the retailer orders in the second phase using randomized rounding procedures that are applied to each retailer i separately. We shall analyze the worst-case expected performance of these algorithms. The corresponding algorithms might not yield the optimal solution with respect to the warehouse orders placed in phase one. Nevertheless, we shall show that, regardless of the instance of the problem, the algorithm with the cheapest expected cost among the two is guaranteed to produce a solution with expected cost at most 1.8 times the cost of an optimal solution to the OWMR problem. In particular, this is true for the solutions obtained by dynamic programming.

4.2 The Random Shift Algorithm with Two-sided Retailer Push Algorithm

Throughout the rest of the paper, we shall refer to the random shift algorithm with two-sided retailer push as *Algorithm 1*. As we have already mentioned, Algorithm 1 has two phases. The first phase is the random

shift procedure described above with step parameter $c = 1$. Consider again $\mathcal{T}_W := \{r_1 < r_2 < \dots < r_M\}$, the set of warehouse orders placed in the first phase of the algorithm.

Next we consider each retailer i separately ($i = 1, \dots, N$), and determine its orders using what we call *two-sided push* procedure. First, the algorithm generates a sequence of (random) *retailer- i shift points* in a way similar to the way in which warehouse shift points are constructed. Let W_i be the upper ceiling of the accumulated weight of fractional retailer orders in the LP solution; that is, $W_i = \lceil \sum_{s=1}^T \hat{y}_s^i \rceil$. Similar to the random warehouse shift procedure above, choose a retailer shift parameter α_i uniformly at random from $(0, 1]$ and construct a sequence of W_i retailer- i shift points $\{\alpha_i + w : w = 0, \dots, W_i - 1\}$ (recall that $c = 1$). In contrast to the warehouse shift points, the retailer- i shift points are used to determine only *tentative retailer- i orders*. The reason is that placing retailer orders depends also on the output of the first phase, in which warehouse orders are determined. For each period $m = 1, \dots, T$, we say that there is a tentative retailer order placed in period m , if there is a retailer- i shift point within the interval $(\sum_{s=1}^{m-1} \hat{y}_s^i, \sum_{s=1}^m \hat{y}_s^i]$.

The tentative orders are used to determine the *permanent retailer orders*. The way in which this is done depends on whether retailer i is a J -retailer or a W -retailer. Suppose that there is a tentative retailer- i order placed in some period m , then one of the following two cases applies:

Case I: Retailer i is a J -retailer. Recall that, without loss of generality, for J -retailers we restrict attention only to policies in which retailer- i orders are placed only in periods with warehouse orders. That is, permanent retailer- i orders in the second phase must be placed in periods $s \in \mathcal{T}_W$, where again \mathcal{T}_W is the set of periods in which warehouse orders were placed in the first phase of the algorithm. Since we place retailer orders only in periods $s \in \mathcal{T}_W$, if $m \notin \mathcal{T}_W$ we wish to *push* this tentative retailer order to periods in which we have already placed warehouse orders in the first phase of the algorithm. In particular, for each tentative retailer order, we place up to two permanent retailer orders: one order is placed in the latest period with a warehouse order in \mathcal{T}_W prior to period m , if such an order exists (i.e., the tentative order is ‘pushed’ to be earlier in time); a second order is placed in the earliest period with a warehouse order in \mathcal{T}_W after period m (i.e., the tentative order is ‘pushed’ to be later in time), if such an order exists. In other words, we place permanent retailer- i orders in $\max\{r \in \mathcal{T}_W : r \leq m\}$ and $\min\{r \in \mathcal{T}_W : r \geq m\}$.

Case II: Retailer i is a W -retailer. In this case we can place a permanent retailer order in each period m for which there is a tentative order. However, we also place a *second* permanent retailer order in the earliest period in \mathcal{T}_W (strictly) after m , if such warehouse order exists. That is, we place one permanent order at m and possibly a second permanent order in $\min\{r \in \mathcal{T}_W : r > m\}$.

The reason that we push tentative retailer orders both earlier and later in time will be made clear in the following discussion. Intuitively, we place additional retailer orders to guarantee that the holding costs incurred by each demand point (i, t) are not too high compared to the holding costs this demand point incurs in the fractional optimal solution (\hat{x}, \hat{y}) . (This property of the algorithm will be used in the proof of Lemma 4.3 below.)

Let \mathcal{T}_i be the set of permanent retailer- i orders placed by Algorithm 1. As we have already observed, there is a warehouse order placed prior to the earliest period with a positive demand point. We claim that the sets \mathcal{T}_W and \mathcal{T}_i (for $i = 1, \dots, N$) induce a feasible solution to the OWMR problem. That is, for each demand point (i, t) , the solution produced by Algorithm 1 has at least one pair of warehouse-retailer orders $[r, s]$ that can serve (i, t) , i.e., $r \in \mathcal{T}_W$, $s \in \mathcal{T}_i$ and $r \leq s \leq t$. In particular, each demand point (i, t) , is satisfied by the cheapest pair of orders $[r, s]$, such that $r \in \mathcal{T}_W$ and $s \in \mathcal{T}_i$. The proof of this claim is discussed in Lemma 4.3 below, in which we show that not only such a pair of warehouse-retailer orders exists, but that the holding costs incurred by (i, t) under Algorithm 1 are bounded.

From Lemma 4.1 above it follows that the total expected warehouse ordering cost of Algorithm 1 is bounded by $\sum_{r=1}^T \hat{y}_r^0 K_r^0$. Next we bound the total expected retailer ordering cost, which is denoted by \mathcal{K}_I . The proof is identical to the proof of Lemma 4.1 above.

Lemma 4.2 *The total expected retailer ordering cost of Algorithm 1 is at most twice the total retailer ordering costs in (\hat{x}, \hat{y}) , the optimal solution of the LP. That is, $\mathcal{K}_I \leq 2 \sum_{i=1}^N \sum_{s=1}^T \hat{y}_s^i K^i$.*

Finally, we wish to bound the total expected holding costs incurred by Algorithm 1, which is denoted by \mathcal{H} . Each demand point (i, t) is considered separately (for $i = 1, \dots, N$ and $t = 1, \dots, T$), and its expected holding cost is bounded using the holding cost that this demand point incurs in the optimal LP solution (\hat{x}, \hat{y}) . In particular, focus on some demand point (i, t) , and let $\hat{H}^{it} = \hat{H}$ be the random holding cost that Algorithm 1 incurs in satisfying this demand point. (Since the following discussion is focused on a fixed demand point, we simplify the notation and omit the superscript it whenever possible.) We wish to bound $E[\hat{H}]$, the expectation of \hat{H} .

Service points. Consider demand point (i, t) , and let $\mathcal{S}^{it} = \mathcal{S}$ be set of all pairs of warehouse and retailer- i orders, which fractionally serve (i, t) in the optimal LP solution (\hat{x}, \hat{y}) . Specifically, let $\mathcal{S} = \{[r_m, s_m] : \hat{x}_{r_m, s_m}^{it} > 0\}$. Without loss of generality, assume that $\mathcal{S} = \{[r_m, s_m] : m = 1, \dots, L\}$, where $h_{r_1, s_1}^{it} \leq h_{r_2, s_2}^{it} \leq \dots \leq h_{r_L, s_L}^{it}$. That is, the order pairs $[r_1, s_1], \dots, [r_L, s_L]$ are sorted in an increasing order according to the per-unit holding costs that they incur. We call these L pairs of warehouse-retailer orders the *service points* of (i, t) . However, since the solution (\hat{x}, \hat{y}) is assumed to have the Monge Property, we conclude that $[r_m, s_m] \leq [r_{m'}, s_{m'}]$, i.e., $r_m \leq r_{m'}$ and $s_m \leq s_{m'}$, for each $1 \leq m' < m \leq L$. Moreover, if i is a J -retailer, we have $s_m = r_m$, for each $m = 1, \dots, L$. To simplify notation, for each $m = 1, \dots, L$, we use H_m to denote $H_{r_m, s_m}^{it} = h_{r_m, s_m}^{it} d_{it}$, assuming $H_1 \leq H_2 \leq \dots \leq H_L$. Thus, the holding cost incurred by (i, t) in the optimal LP solution (\hat{x}, \hat{y}) can be expressed as

$$\sum_{m=1}^L \hat{x}_{r_m, s_m}^{it} H_m.$$

Next we lower bound the probability $Pr(\hat{H} \leq H_m)$ that holding cost incurred by (i, t) in the solution produced by Algorithm 1 is at most H_m , for each $m = 1, \dots, L$. This will then be used to bound the overall

expected holding costs incurred by demand point (i, t) .

Lemma 4.3 For each $m = 1, \dots, L$, the probability that the holding cost incurred by (i, t) under Algorithm 1 is at most H_m is at least $(\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it})^2$; that is, $Pr(\hat{H} \leq H_m) \geq (\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it})^2$.

Proof : We have already mentioned that given the sets \mathcal{T}_W and \mathcal{T}_i , each demand point (i, t) is served from the cheapest possible pair of warehouse-retailer orders. Moreover, if there exist $r \in \mathcal{T}_W$ and $s \in \mathcal{T}_i$, such that $r \geq r_m$, $s \geq s_m$ and $r \leq s \leq t$, then the holding cost incurred by (i, t) is at most H_m . (We have already seen that for any two pairs of orders $[r, s]$ and $[r', s']$, such that $r \leq r'$ and $s \leq s'$, we have $h_{r's'}^{it} \leq h_{rs}^{it}$.)

Consider now the event, in which there is a warehouse shift point within the interval $[r_m, t]$ and a retailer- i shift point within $[s_m, t]$. This implies that there is a warehouse order within $[r_m, t]$ and a tentative retailer- i order within $[s_m, t]$. Since $s_m \geq r_m$ it follows that within the interval $[r_m, t]$ there is at least one warehouse order either earlier or later than the retailer- i tentative order within $[s_m, t]$. However, for each the retailer tentative order, Algorithm 1 aims to generate a permanent retailer order earlier *and* later in time. It follows that there must exist a pair of warehouse-retailer orders $[r, s]$, such that $r \in \mathcal{T}_W$, $s \in \mathcal{T}_i$, $r \geq r_m$, $s \geq s_m$ and $r \leq s \leq t$. It is now sufficient to lower bound the probability of the latter event.

By the construction of Algorithm 1, it follows that the probability to place a warehouse order within $[r_m, t]$ is equal $\min\{1, \sum_{r \in [r_m, t]} \hat{y}_r^0\}$ and that the probability to have a tentative retailer order within $[s_m, t]$ is equal $\min\{1, \sum_{s \in [s_m, t]} \hat{y}_s^i\}$. However, warehouse orders are determined, independently of the tentative retailer orders. This implies that

$$Pr(\hat{H} \leq H_m) \geq \min\{1, \sum_{r \in [r_m, t]} \hat{y}_r^0\} \cdot \min\{1, \sum_{s \in [s_m, t]} \hat{y}_s^i\}.$$

Finally, Constrains (2) and (3), respectively, imply that $\min\{1, \sum_{s \in [s_m, t]} \hat{y}_s^i\} \geq \sum_{u=1}^m \hat{x}_{r_u, s_u}^{it}$ and that $\min\{1, \sum_{r \in [r_m, t]} \hat{y}_r^0\} \geq \sum_{u=1}^m \hat{x}_{r_u, s_u}^{it}$. The proof of the lemma then follows. \blacksquare

Lemma 4.3 above implies that under Algorithm 1, demand point (i, t) is served by a pair of orders $[r', s']$, such that $r_L \leq r'$ and $s_L \leq s'$. (Observe that $(\sum_{u=1}^L \hat{x}_{r_u, s_u}^{it})^2 = 1$.) In particular, it implies that the set \mathcal{T}_W and \mathcal{T}_i indeed induce a feasible solution. Moreover, we can express $E[\hat{H}]$ as

$$\sum_{[r, s]: r_L \leq r, s_L \leq s} H_{rs}^{it} Pr(\hat{H} = H_{rs}^{it}), \quad (5)$$

where again $Pr(\hat{H} = H_{rs}^{it})$ denotes the corresponding probability that under Algorithm 1 demand point (i, t) is served by the pair of orders $[r, s]$.

Given (5) above, it is straightforward to derive an upper bound on the expected holding cost incurred by demand point (i, t) under Algorithm 1. Let $H_0 = 0$ and observe that

$$\begin{aligned}
E[\hat{H}] &= \sum_{[r,s]: r_L \leq r, s_L \leq s} H_{rs}^{it} Pr(\hat{H} = H_{rs}^{it}) \\
&\leq H_1 Pr(H_0 \leq \hat{H} \leq H_1) + \sum_{m=2}^L H_m Pr(H_{m-1} < \hat{H} \leq H_m) \\
&= H_1 Pr(\hat{H} \leq H_1) + \sum_{m=2}^L H_m [Pr(\hat{H} \leq H_m) - Pr(\hat{H} \leq H_{m-1})] \\
&= H_L + \sum_{m=1}^{L-1} Pr(\hat{H} \leq H_m) [H_m - H_{m+1}].
\end{aligned} \tag{6}$$

The inequality in (6) above follows from the fact that, for each $m = 1, \dots, L$, we weight the probability $Pr(H_{m-1} < \hat{H} \leq H_m)$ by H_m , which is the highest holding cost within this range. The first equality follows from the fact that $Pr(\hat{H} < 0) = 0$ and the identity $Pr(H_{m-1} < \hat{H} \leq H_m) = Pr(\hat{H} \leq H_m) - Pr(\hat{H} \leq H_{m-1})$. The last equality follows from Lemma 4.3, in which we show that $Pr(\hat{H} \leq H_L) = 1$. Moreover, observe that the term $\sum_{m=1}^{L-1} Pr(\hat{H} \leq H_m) [H_m - H_{m+1}]$ above is non-positive, since $H_m - H_{m+1} \leq 0$. This implies that if we consider (6) above, but, for each $m = 1, \dots, L-1$, replace $Pr(\hat{H} \leq H_m)$ with a lower bound on that probability, then the upper bound developed in (6) is still maintained.

In particular, Lemma 4.3 and (6) above imply that

$$\begin{aligned}
E[\hat{H}] &\leq H_1 (\hat{x}_{r_1, s_1}^{it})^2 + \sum_{m=2}^L H_m \left[\left(\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} \right)^2 - \left(\sum_{u=1}^{m-1} \hat{x}_{r_u, s_u}^{it} \right)^2 \right] \\
&= \sum_{m=1}^L H_m \left[\left(\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} \right)^2 - \left(\sum_{u=1}^{m-1} \hat{x}_{r_u, s_u}^{it} \right)^2 \right].
\end{aligned} \tag{7}$$

The inequality follows because, for each $m = 1, \dots, L-1$, we replace $Pr(\hat{H} \leq H_m)$ in (6) by the lower bound established in Lemma 4.3 above, and Constraint (1) implies that $(\sum_{u=1}^L \hat{x}_{r_u, s_u}^{it})^2 = 1$.

The holding cost function. To conclude the analysis, we next introduce the *holding cost function* $\bar{H}^{it}(\beta) = \bar{H}(\beta)$. This function is defined for each demand point (i, t) according to the optimal LP solution (\hat{x}, \hat{y}) . For a given value of $\beta \in (0, 1]$, let $m(\beta)$ be the unique index, such that $\beta \in (\sum_{u=1}^{m(\beta)-1} \hat{x}_{r_u, s_u}^{it}, \sum_{u=1}^{m(\beta)} \hat{x}_{r_u, s_u}^{it}]$. (Constraint (1) above implies that $\sum_{u=1}^L \hat{x}_{r_u, s_u}^{it} = 1$.) We call $m(\beta)$ the *the β -index* of (i, t) and $[r_{m(\beta)}, s_{m(\beta)}]$ the *β -point* of the service points $[r_1, s_1], \dots, [r_L, s_L]$ to reflect the fact that this is the first service point by which the accumulated β fraction of the demand (i, t) is satisfied in the optimal LP solution (\hat{x}, \hat{y}) . Then we define $\bar{H}(\beta) = H_{m(\beta)}$. The function $\bar{H}(\beta)$ is a step function with steps starting at the points $0, \hat{x}_{r_1, s_1}^{it}, \sum_{u=1}^2 \hat{x}_{r_u, s_u}^{it}, \dots, \sum_{u=1}^{L-1} \hat{x}_{r_u, s_u}^{it}$ and step heights H_1, \dots, H_L , respectively. Moreover, the integral of $\bar{H}(\beta)$ over $(0, 1]$ is equal to the holding costs incurred by (i, t) in the LP optimal solution (\hat{x}, \hat{y}) . That is,

$$\int_0^1 \bar{H}(\beta) d\beta = \sum_{u=1}^L \hat{x}_{r_u, s_u}^{it} H_u.$$

We note that Shmoys, Tardos and Aardal [17] have used a similar function to $\bar{H}(\beta)$ in their paper that provides the first constant approximation algorithm for the classical metric facility location problem. Next we shall describe another application of this function. In particular, we use the density function \bar{H} to bound the expected holding costs incurred by (i, t) under Algorithm 1.

Inequality (7) and the properties of the function $\bar{H}(\beta)$ imply

$$\begin{aligned} E[\hat{H}] &\leq \sum_{m=1}^L H_m \left[\left(\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} \right)^2 - \left(\sum_{u=1}^{m-1} \hat{x}_{r_u, s_u}^{it} \right)^2 \right] = \\ &2 \int_0^1 \beta \bar{H}(\beta) d\beta \leq 2 \int_0^1 \bar{H}(\beta) d\beta \leq 2 \sum_{u=1}^L \hat{x}_{r_u, s_u}^{it} H_u. \end{aligned} \quad (8)$$

The equality follows from the properties of $\bar{H}(\beta)$, being a step function. (In particular, on each of the intervals $(\sum_{u=1}^{m-1} \hat{x}_{r_u, s_u}^{it}, \sum_{u=1}^m \hat{x}_{r_u, s_u}^{it}]$ we take the integral $H_m \int_{\sum_{u=1}^{m-1} \hat{x}_{r_u, s_u}^{it}}^{\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it}} \beta$.) The second inequality follows from the fact that we integrate over $[0, 1]$. This implies the following lemma.

Lemma 4.4 *Let \mathcal{H} denote the total expected holding costs incurred by Algorithm 1. Then these costs are at most twice the total holding costs incurred in the optimal LP solution (\hat{x}, \hat{y}) . That is,*

$$\mathcal{H} \leq 2 \sum_{i=1}^N \sum_{t=1}^T \sum_{r, s: r \leq s \leq t} \hat{x}_{rs}^{it} H_{rs}^{it}.$$

Lemmas 4.1, 4.2 and 4.4 imply the following theorem.

Theorem 4.5 *The total expected cost $\mathcal{K}_0 + \mathcal{K}_I + \mathcal{H}$ incurred by Algorithm 1 is guaranteed to be at most twice the cost of an optimal policy for the OWMR problem. Thus, Algorithm 1 is a randomized 2-approximation for the OWMR problem and its special case the JRP problem.*

Proof : Let opt denote the value of an optimal policy for the OWMR problem and opt_{LP} be the optimal value of (P) . Lemmas 4.1, 4.2 and 4.4 imply that

$$\mathcal{K}_0 + \mathcal{K}_I + \mathcal{H} \leq \sum_{r=1}^T \hat{y}_r^0 K_r^0 + 2 \sum_{i=1}^N \sum_{s=1}^T \hat{y}_s^i K^i + 2 \sum_{i=1}^N \sum_{t=1}^T \sum_{r, s: r \leq s \leq t} \hat{x}_{rs}^{it} H_{rs}^{it} \leq 2opt_{LP} \leq 2opt.$$

The proof of the theorem then follows. ■

4.3 The Random Shift Algorithm with Retailer One-Sided Push

Next we describe the *random shift algorithm with retailer one-sided push* that we refer to as *Algorithm 2*. Note that the inequality in the proof of Theorem 4.5 above is not balanced, in that the warehouse ordering part in the LP solution is weighted by 1, while the retailer ordering cost and holding cost parts are weighted by 2. The idea underlying Algorithm 2 is to place warehouse orders more frequently. This incurs additional warehouse ordering costs, but decreases the retailer ordering costs and holding costs incurred. Thus, Algorithm 2 creates a different balance between the different parts of the total cost.

Like Algorithm 1, Algorithm 2 also runs in two phases. In the first phase we again determine the warehouse orders, now applying the random shift procedure described above but with a step parameter $c \in (0, 0.5]$. (We will later set c so as to appropriately balance between Algorithm 1 and Algorithm 2.) Let \mathcal{T}_W be again the set of periods of the warehouse orders placed by the algorithm. It is readily verified that Algorithm 2 places warehouse orders more frequently than Algorithm 1, which uses a step parameter equal to 1. Moreover, from Lemma 4.1 above, we conclude that the total expected warehouse ordering costs of Algorithm 2 is at most $\frac{1}{c}$ times the warehouse ordering costs in the LP solution. That is, $\mathcal{K}_0 \leq \frac{1}{c} \sum_{r=1}^T \hat{y}_r^0 K_r^0$. We have already mentioned that once the warehouse orders are determined one can solve for each of the retailers separately to minimize the resulting retailer ordering and holding costs. However, we next describe the second phase of the algorithm as part of the worst-case analysis.

In the second phase of the algorithm we determine the retailer orders, and this is again done separately for each retailer. First we generate retailer- i shift points and tentative retailer orders in a way similar to what was described above for Algorithm 1, but with a different step parameter that is coordinated with the step parameter used in the first phase of the algorithm. Specifically, set the retailer step parameter to be equal $1 - c$. Since $c \in (0, 0.5]$ the retailer step parameter $1 - c$ is greater than c . The permanent retailer orders are again placed according to whether retailer i is a J -retailer or a W -retailer. Suppose that there is a tentative retailer- i order placed in some period m . Then one of the following two cases applies:

1. *Retailer i is a J -retailer.* In this case we simply push the tentative order to the earliest period in \mathcal{T}_W later in time, if such an order exists. That is, we place the permanent retailer- i order in $\min\{r \in \mathcal{T}_W : r \geq m\}$.
2. *Retailer i is a W -retailer.* In this case, we simply place a permanent retailer order in period m .

Observe that in Algorithm 2, tentative retailer orders are ‘pushed’ only to be later in time.

The next lemma bounds the total expected retailer ordering costs incurred by Algorithm 2. The proof is similar to that of Lemma 4.1.

Lemma 4.6 *Let \mathcal{K}_I be the overall expected retailer ordering costs incurred by Algorithm 2. Then these costs are at most $\frac{1}{1-c}$ times the total retailer ordering costs incurred in the LP optimal solution (\hat{x}, \hat{y}) . That is, $\mathcal{K}_I \leq \frac{1}{1-c} \sum_{i=1}^N \sum_{s=1}^T \hat{y}_s^i K^i$*

For each $i = 1, \dots, N$, let \mathcal{T}_i be the set of periods of permanent retailer- i orders placed by the algorithm. Once again we claim that together with \mathcal{T}_W they induce a feasible solution to the OWMR, in which each demand point is served from the cheapest possible pair of warehouse-retailer orders. (Similar to the discussion of Algorithm 1, it is sufficient to show that, for each positive demand point (i, t) , there exists a pair of warehouse-retailer orders $[r, s]$ that can serve demand point (i, t) , such that $s \in \mathcal{T}_i$ and $r \in \mathcal{T}_W$.) In fact, in Lemma 4.7 we shall show that under Algorithm 2, each demand point (i, t) , is served from such a pair of warehouse-retailer orders.

Next we wish to bound the overall expected holding costs incurred by Algorithm 2. Similar to the analysis of Algorithm 1, we consider each demand point separately, and bound the expected holding costs it incurs under Algorithm 2 using the respective holding cost it incurs in (\hat{x}, \hat{y}) . The first step is to lower bound the probabilities $Pr(\hat{H} \leq H_m)$, for each $m = 1, \dots, L$, where \hat{H} is the holding cost demand point (i, t) incurs in the solution obtained by Algorithm 2, and $Pr(\cdot)$ is the corresponding probability induced by the Algorithm 2.

Lemma 4.7 For each $m = 1, \dots, L$,

$$Pr(\hat{H} \leq H_m) \geq \max\left\{0, \frac{\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} - c}{1 - c}\right\}.$$

Proof : Recall that $\sum_{u=1}^L \hat{x}_{r_u, s_u}^{it} = 1$. For each given value of $\beta \in (0, 1]$, let again $m(\beta)$ be the β -index of (i, t) , i.e., the unique index such that $\beta \in (\sum_{u=1}^{m(\beta)-1} \hat{x}_{r_u, s_u}^{it}, \sum_{u=1}^{m(\beta)} \hat{x}_{r_u, s_u}^{it}]$.

Note that, for each $m < m(c)$, the lower bound above trivially holds, since $\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} < c$ and the lower bound is equal to 0. Consider now some $m \geq m(c)$, such that $\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} > c$. In particular, we have already seen that $r_m \leq r_{m(c)}$ and $s_m \leq s_{m(c)}$.

Moreover, we have already seen that if (i, t) is served by a pair of warehouse-retailer orders $[r, s]$ such that $r \geq r_m$ and $s \geq s_m$, then the holding cost it incurs is at most H_m . Thus, we focus on this event and show that the probability it occurs is at least

$$\max\left\{0, \frac{\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} - c}{1 - c}\right\}.$$

First consider the case, in which i is a J -retailer. Focus on the event where there is a warehouse order within the interval $[r_{m(c)}, t]$ and a tentative retailer- i within the interval $[s_m, s_{m(c)}]$. Since $s_{m(c)} = r_{m(c)}$ and Algorithm 2 aims to shift tentative retailer orders later in time, it follows that indeed (i, t) will be served by a pair of orders $[r, s]$ such that $r \geq r_m$ and $s \geq s_m$. It is now sufficient to lower bound the probability of the latter event.

However, since $\sum_{r \in [r_{m(c)}, t]} \hat{y}_r^0 \geq c$, there is a warehouse order placed within the time interval $[r_{m(c)}, t]$ with probability 1. Moreover, we claim that the probability of having a tentative retailer- i order within the time interval $[s_m, s_{m(c)}]$ is at least $\frac{1}{1-c} (\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} - c)$. Since warehouse orders are determined independently of tentative retailer orders the proof of this case will follow. This can be seen as follows:

$$\begin{aligned} \sum_{u=m(c)}^m \hat{x}_{r_u, s_u}^{it} &= \sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} - \sum_{u=1}^{m(c)-1} \hat{x}_{r_u, s_u}^{it} \\ &\geq \sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} - c. \end{aligned}$$

However, Constant (2) implies that the cumulative weight of retailer- i fractional orders over the time interval $[s_m, s_{m(c)}]$, i.e., $\sum_{u=m(c)}^m \hat{y}_{s_u}^i$, is at least $\sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} - c$, from which the claim follows. (Recall that the step parameter in second phase of Algorithm 2 is $1 - c$.)

Next consider the case, in which i is a W -retailer, let $\mu = \sum_{u=1}^m \hat{x}_{r_u, s_u}^{it} - c > 0$. Focus on the event where there is a warehouse order within the interval $[r_m, r_{m(\mu)}]$ and a retailer- i order within $[s_{m(\mu)}, t]$. Since $r_{m(\mu)} \leq s_{m(\mu)}$, it follows that indeed (i, t) will be served by a pair of orders $[r, s]$ such that $r \geq r_m$ and $s \geq s_m$. Using similar arguments we conclude that there is a retailer order placed within the time interval $[s_{m(\mu)}, t]$ with probability at least $\mu/(1-c)$ and that there is a warehouse order placed within the time interval $[r_m, r_{m(\mu)}]$ with probability 1. Moreover, the above two events are again independent of each other.

■

Lemma 4.7 above implies that (i, t) is served by a pair of warehouse-retailer orders $[r, s]$, such that $r \geq r_L$ and $s \geq s_L$, with probability 1. This implies that feasibility of the solution induced by the sets \mathcal{T}_W and \mathcal{T}_i . In particular, Inequality (6) above is valid. (Observe that Inequality (6) does not depend on Algorithm 1, but only on the fact that with probability 1 (i, t) is served by a pair of warehouse-retailer orders $[r, s]$, such that $r \geq r_L$ and $s \geq s_L$.) Similar to the analysis of algorithm 1, the upper bound obtained by Inequality (6) is maintained if, for each $m = 1, \dots, L-1$, one replaces $Pr(\hat{H} \leq H_m)$ by a respective lower bound.

Using the lower bounds obtained in Lemma 4.7, we get that

$$\begin{aligned} E[\hat{H}] &\leq \frac{1}{1-c} \sum_{u=m(c)}^L H_u \hat{x}_{r_u, s_u}^{it} \leq \frac{1}{1-c} \sum_{u=m(c)+1}^L \hat{x}_{r_u, s_u}^{it} H_u + H_{m(c)} \frac{\sum_{m=1}^{m(c)} \hat{x}_{r_m, s_m}^{it} - c}{1-c} \\ &\leq \frac{1}{1-c} \sum_{u=1}^L \hat{x}_{r_u, s_u}^{it} H_u. \end{aligned}$$

We have obtained the following lemma.

Lemma 4.8 *Let \mathcal{H} denote the overall expected holding cost incurred by Algorithm 2 with a step parameter $c \in (0, 0.5]$. Then \mathcal{H} is at most $\frac{1}{1-c}$ times the holding costs incurred by the optimal LP solution (\hat{x}, \hat{y}) . That is, $\mathcal{H} \leq \frac{1}{1-c} \sum_{i=1}^N \sum_{t=1}^T \sum_{r, s: r \leq s \leq t} \hat{x}_{rs}^{it} H_{rs}^{it}$.*

Lemmas 4.6 and 4.8 imply the following theorem.

Theorem 4.9 *The overall expected costs $\mathcal{K}_0 + \mathcal{K}_I + \mathcal{H}$ incurred by Algorithm 2 with a step parameter $c \in (0, 0.5]$ is at most*

$$\frac{1}{c} \sum_{r=1}^T \hat{y}_r^0 K_r^0 + \frac{1}{1-c} \sum_{i=1}^N \sum_{s=1}^T \hat{y}_s^i K^i + \frac{1}{1-c} \sum_{i=1}^N \sum_{t=1}^T \sum_{r, s: r \leq s \leq t} \hat{x}_{rs}^{it} H_{rs}^{it}.$$

It is readily verified that for $c = 0.5$ Algorithm 2 is a randomized 2-approximation for the OWMR problem.

4.4 Combining Algorithms 1 and 2

Next we use Algorithm 1 and Algorithm 2 together. Specifically, we shall show that taking the algorithm with the minimum expected cost among algorithms 1 and 2 yields an improved expected worst-case guar-

antee of 1.8. We shall achieve this by choosing the step parameter of Algorithm 2 to be $c = 1/3$. Using the fact that $\min\{a, b\} \leq \lambda a + (1 - \lambda)b$, for each $0 \leq \lambda \leq 1$, we apply Theorems 4.5 and 4.9 (with $c = 1/3$) and take $\lambda = 3/5$ to conclude that the solution with the smaller expected cost has expected value at most

$$1.8 \left(\sum_{r=1}^T \hat{y}_r^0 K_r^0 + \sum_{i=1}^N \sum_{s=1}^T \hat{y}_s^i K^i + \sum_{i=1}^N \sum_{t=1}^T \sum_{r,s:r \leq s \leq t} \hat{x}_{rs}^{it} H_{rs}^{it} \right) = 1.8opt_{LP} \leq 1.8opt.$$

Theorem 4.10 *There exists a randomized 1.8-approximation algorithm for the OWMR problem and its special case the JRP problem.*

Finally, we describe how to derandomize the algorithms and get deterministic approximation algorithms with the same guarantee. We have already mentioned that once the warehouse orders are determined, the problem decomposes to N single-retailer subproblems that can be solved to optimality via dynamic programming. Thus, the expected cost of the solutions obtained by dynamic programming is at most the expected cost of algorithms 1 and 2. It is now enough to show how to derandomize the first phase of the algorithms. However, it is readily verified that in the random shift procedure described above, there is only a polynomial number of values of α_0 that yield distinct sets of warehouse orders. Specifically, there are $O(\frac{T}{c})$ such points, where c is again the step parameter being chosen. (Observe that we can restrict attention only to sequences of shift points in which one of them is at the right edge of an interval \hat{Y}_m^0 or $(cw, c(w+1)]$. Thus, the number of different sequences of shift points to be considered is bounded by $\lceil T/C \rceil$.) It is readily verified that these values can be easily enumerated. (For the analysis we have considered the values $c = 1$ and $c = \frac{1}{3}$.) Specifically, the cost of the solution obtained by taking the best (cheapest) choice of warehouse orders is at most the expected cost over all choices.

Theorem 4.11 *There exists a deterministic 1.8-approximation algorithm for the OWMR problem and its special case the JRP problem.*

5 Conclusions

In this paper, we have demonstrated how strong LP relaxations can be used to construct provably near-optimal solutions for a class of classical deterministic inventory models. We have focused on the classical model, known as the one-warehouse multi-retailer problem, and designed the first constant factor approximation algorithms for several variants of this model. This is yet another example of the potential of LP-based approximation methods as a tool to "attack" inventory problems. We find that this direction is important, both theoretically and practically. From the practical aspect, we point out that in many real-life cases, these inventory problems are solved as integer programs. This emphasizes the importance of techniques that enable us to efficiently round fractional solutions to good feasible integer solutions, as a tool for enhancing the computational procedures for solving these IP's.

We believe that it would be interesting to test the typical quality of the solutions that our algorithms generate on different inputs and compare them to other known heuristics.

A very interesting theoretical open question is related to the approximability of the OWMR problem. The problem is proven to be NP-hard, since the special case of the JRP is NP-hard [1]. However, we know of no approximability hardness result and one can not even exclude the existence of a polynomial-time approximation scheme (i.e., one might be able to design a ρ -approximation algorithm for any $\rho > 1$). In addition, the analysis presented in this paper is not tight, that is, we do not have bad examples on which the performance of the proposed algorithms matches the upper bound of 1.8.

Finally, it seems that LP-based approximation techniques can be used to provide high quality solutions to a class of classical inventory models. It is most interesting to see whether these techniques can be applied to more complicated inventory models.

Acknowledgments We thank the Associate Editor and the two anonymous referees for their constructive comments, which significantly improved the exposition of the paper.

Preliminary presentations of part of the results in this paper were given in [11].

A Appendix - Transportation Costs and Batch Capacities

In this appendix, we consider the OWMR problem, but with batch capacities and a correspondingly more complex cost structure. Instead of ordering costs we consider transportation costs. Usually, transportation is based on trucks with given capacities. We model this in the following way. For each warehouse-retailer i segment we consider trucks/batches each with capacity U^i and cost K^i ($i = 1, \dots, N$). In addition, in the supplier-warehouse segment, we consider trucks/batches each with capacity U^0 and cost K^0 . Each order consists now of several *complete* batches, say w ($w \in \mathbb{Z}_+$), that can provide a capacity of wU^i units and incurs a cost of wK^i ($i = 0, \dots, N$). The batch-based ordering structure is sometimes called ordering with *soft capacities*.

We first modify the linear program (P) presented in Section 3 to capture the new model. For each period $r = 1, \dots, T$ we add the constraint $\sum_{i=1}^N \sum_{t \geq r} \sum_{s \in [r,t]} x_{rs}^{it} d_{it} \leq y_r^0 U^0$. For each $i = 1, \dots, N$ and $s = 1, \dots, T$ we add the constraint $\sum_{t \geq s} \sum_{r \leq s} x_{rs}^{it} d_{it} \leq y_s^i U^i$. Observe that the variables y_s^i and y_r^0 can be larger than 1. Consider now the corresponding dual program:

$$\text{maximize} \quad \sum_{i=1}^N \sum_{t=1}^T b_t^i \quad (\text{D1})$$

$$\text{subject to} \quad b_t^i \leq H_{rs}^{it} + l_{st}^i + z_{rt}^i + \delta_{is} d_{it} + \theta_r d_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T, \quad (9)$$

$$s = 1, \dots, t, \quad r = 1, \dots, s.$$

$$\sum_{t=s}^T l_{st}^i + \delta_{is} U^i \leq K^i, \quad i = 1, \dots, N, \quad s = 1, \dots, T, \quad (10)$$

$$\sum_{i=1}^N \sum_{t=r}^T z_{rt}^i + \theta_r U^0 \leq K^0, \quad r = 1, \dots, T, \quad (11)$$

$$l_{st}^i, z_{rt}^i, \delta_{is}, \theta_r \geq 0, \quad i = 1, \dots, N, \quad t = 1, \dots, T, \quad (12)$$

$$s = 1, \dots, t, \quad r = 1, \dots, t.$$

Note that weak duality implies that, each feasible solution $(b, l, z, \delta, \theta)$ for the above dual program (D1) provides a lower bound on the optimal cost of the primal LP; thus, provides a lower bound on the optimal cost of the OWMR problem with batch capacities. Suppose now that we set the value of each dual variable δ_{is} to be equal to $\frac{K^i}{2U^i}$, and of each dual variable θ_r to be equal to $\frac{K^0}{2U^0}$, and consider the induced modified LP. It is straightforward to verify that the modified LP is the dual program of the following primal LP (recall that $H_{rs}^{it} = h_{rs}^{it} d_{it}$):

$$\text{minimize} \quad \sum_{i=0}^N \sum_{s=1}^T \frac{1}{2} Y_s^i K^i + \sum_{i=1}^N \sum_{t=1}^T \sum_{r,s:r \leq s \leq t} X_{rs}^{it} d_{it} (h_{rs}^{it} + \frac{1}{2} (\frac{K^i}{U^i} + \frac{K^0}{U^0})) \quad (\text{P2})$$

$$\text{subject to} \quad \sum_{r,s:r \leq s \leq t} X_{rs}^{it} = 1, \quad i = 1, \dots, N, t = 1, \dots, T, d_{it} > 0, \quad (13)$$

$$\sum_{r:r \leq s} X_{rs}^{it} \leq Y_s^i, \quad i = 1, \dots, N, t = 1, \dots, T, s = 1, \dots, t \quad (14)$$

$$\sum_{s:r \leq s \leq t} X_{rs}^{it} \leq Y_r^0, \quad i = 1, \dots, N, t = 1, \dots, T, r = 1, \dots, t \quad (15)$$

$$X_{rs}^{it}, Y_r^i \geq 0, \quad i = 0, \dots, N, s = 1, \dots, T, \quad (16)$$

$$r = 1, \dots, s, t = s, \dots, T.$$

However, (P2) is an LP relaxation of an uncapacitated OWMR problem, where the ordering and the holding cost parameters are modified accordingly. Thus, the modified dual program has an optimal solution which we denote by $(\hat{b}, \hat{l}, \hat{z})$. In particular, by strong duality $\sum_{i=1}^N \sum_{t=1}^T \hat{b}_t^i$ is equal to the optimal value of (P2) that we denote by opt_{LP2} . It is also clear that the modified holding cost parameters $h_{rs}^{it} + \frac{1}{2} (\frac{K^i}{U^i} + \frac{K^0}{U^0})$ still obey all the assumptions discussed in Section 2. Hence, we can use the algorithms described in Section 4 to find an integer solution to this uncapacitated OWMR problem, denoted by (\bar{X}, \bar{Y}) , with the following property:

$$\sum_{i=0}^N \sum_{s=1}^T \frac{1}{2} \bar{Y}_s^i K^i + \sum_{i=1}^N \sum_{t=1}^T \sum_{r,s:r \leq s \leq t} \bar{X}_{rs}^{it} d_{it} (h_{rs}^{it} + \frac{1}{2} (\frac{K^i}{U^i} + \frac{K^0}{U^0})) \leq 1.8 opt_{LP2} =$$

$$1.8 \sum_{i=1}^N \sum_{t=1}^T \hat{b}_t^i$$

Next we define a feasible solution to the original OWMR problem with batch capacities. For each (i, t) and $r \leq s$, we set $\bar{x}_{rs}^{it} = \bar{X}_{rs}^{it}$. For each $r = 1, \dots, T$, we set $\bar{y}_r^0 = \lceil \frac{1}{U^0} (\sum_{i=1}^N \sum_{t \geq r} \sum_{s \in [r,t]} \bar{X}_{rs}^{it}) \rceil$. For each $i = 1, \dots, N$ and $s = 1, \dots, T$, we set $\bar{y}_s^i = \lceil \frac{1}{U^i} (\sum_{t \geq s} \sum_{r \leq s} \bar{X}_{rs}^{it}) \rceil$. It follows that $\bar{y}_r^0 \leq \bar{Y}_r^0 + \frac{1}{U^0} (\sum_{i=1}^N \sum_{t \geq r} \sum_{s \in [r,t]} \bar{X}_{rs}^{it})$ (for each $r = 1, \dots, T$), and $\bar{y}_s^i \leq \bar{Y}_s^i + \frac{1}{U^i} (\sum_{t \geq s} \sum_{r \leq s} \bar{X}_{rs}^{it})$ (for each $i = 1, \dots, N$ and $s = 1, \dots, T$). This implies that:

$$\sum_{i=0}^N \sum_{s=1}^T \bar{y}_s^i K^i + \sum_{i=1}^N \sum_{t=1}^T \sum_{r,s:r \leq s \leq t} \bar{x}_{rs}^{it} H_{rs}^{it} \leq$$

$$2 \left(\sum_{i=0}^N \sum_{s=1}^T \frac{1}{2} \bar{Y}_s^i K^i + \sum_{i=1}^N \sum_{t=1}^T \sum_{r,s:r \leq s \leq t} \bar{X}_{rs}^{it} d_{it} (h_{rs}^{it} + \frac{1}{2} (\frac{K^i}{U^i} + \frac{K^0}{U^0})) \right) \leq$$

$$3.6 \sum_{i=1}^N \sum_{t=1}^T \hat{b}_t^i$$

Finally, we claim that $\sum_{i=1}^N \sum_{t=1}^T \hat{b}_t^i$ provides a lower bound on the optimal cost of the original OWMR problem with batch capacities. It is sufficient to show that (D1) has a feasible solution with objective value

$\sum_{i=1}^T \sum_{t=1}^T \hat{b}_t^i$. However, by setting $\hat{\delta}_{is} = \frac{K^i}{2U^i}$ and $\hat{\theta}_r = \frac{K^0}{2U^0}$, the solution $(\hat{b}, \hat{l}, \hat{z})$ is mapped to a feasible solution $(\hat{b}, \hat{l}, \hat{z}, \hat{\delta}, \hat{\theta})$ to (D1) with the same objective value. We now conclude the following theorem:

Theorem A.1 *The algorithm provides a 3.6-approximation algorithm for the OWMR and the JRP problems with transportation costs and batches.*

Finally, we observe that by using dynamic programming the single-location lot-sizing problem can be solved optimally for any cost parameters h_{st} . In turn, this yields a 2-approximation algorithm for the single-location lot-sizing problem with batches. Here we can allow a time-dependent ordering cost K_s and batch size U_s ($s = 1, \dots, T$). As a by product we also prove a lower bound of 2 on the integrality gap of the corresponding natural LP-relaxations. For the case where $U_s = U$ (uniform batch size), Pochet and Wolsey [14] have shown that the problem can be solved optimally using dynamic programming (applied directly to the problem). Moreover, they have observed a linear program for this problem with integrality property (i.e., an LP that describes the convex hull of the integer solutions). This LP has an exponential number of constraints, but these constraints are separable.

Theorem A.2 *For the single-location lot-sizing problem with time-dependent batch size there exists a 2-approximation algorithm.*

Theorem A.3 *The facility location inspired LP for the single-location lot-sizing problem with time-dependent batch size has an integrality gap of at most 2.*

References

- [1] E. Arkin, D. Joneja, and R. Roundy. Computational complexity of uncapacitated multi-echelon production planning problems. *Operations Research Letters*, 8:61–66, 1989.
- [2] I. Bárány, T. J. Van Roy, and L. A. Wolsey. Uncapacitated lot-sizing: the convex hull of solutions. *Mathematical Programming Study*, 22:32–43, 1984.
- [3] D. Bertsimas, C. Teo, and R. Vohra. On dependent randomized rounding algorithms. *Operations Research Letters*, 25:105–114, 1999.
- [4] A. Chan, A. Muriel, Z.-J. Shen, D. Simchi-Levi, and C.-P. Teo. Effectiveness of zero inventory ordering policies for an one-warehouse multi-retailer problem with piecewise linear cost structures. *Management Science*, 48:1446–1460, 2000.
- [5] A. Federgruen and M. Tzur. A simple forward algorithm to solve general dynamic lot-sizing models with n periods in $O(n \log n)$ or $O(n)$ time. *Management Science*, 37:909–925, 1991.

- [6] A. Federgrun and M. Tzur. Time-partitioning heuristics: Application to one warehouse, multi-item, multi-retailer lot-sizing problems. *Naval Research Logistics*, 46:463–486, 1999.
- [7] U. Feige. A threshold of $\ln(n)$ for approximating set-cover. *Journal of ACM*, 45:634–652, 1998.
- [8] K. Jain and V. V. Vazirani. Approximation algorithms for metric facility location and k -median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM* 48, pages 274–296, 2001.
- [9] Y. Jin and A. Muriel. Single-warehouse multi-retailer inventory systems with full truckload shipments. Working Paper, 2005.
- [10] J. Krarup and O. Bilde. Plant location, set covering and economic lot sizing: an $O(mn)$ algorithm for structural problems. In *Numerische Methoden Bei Optimierungsaufgaben*, volume 3, pages 155–180, 1977.
- [11] R. Levi, R. Roundy, and D. Shmoys. A constant approximation algorithm for the one-warehouse multi-retailer problem. Technical Report TR1408, ORIE Department, Cornell University, 2004. Submitted, extended abstract will appear in SODA 2005.
- [12] R. Levi, R. O. Roundy, and D. B. Shmoys. Primal-dual algorithms for deterministic inventory problems. Technical Report TR1042, ORIE Department, Cornell University, 2004. To appear in Math of OR.
- [13] R. Levi, R. O. Roundy, and D. B. Shmoys. Primal-dual algorithms for deterministic inventory problems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 353–362, 2004.
- [14] Y. Pochet and L. A. Wolsey. Lot-sizing with constant batches: *Formulation* and valid inequalities. *Mathematics of Operations Research*, 18:767–785, 1993.
- [15] Z. J. Shen, D. Simchi-Levi, and C. P. Teo. Approximation algorithms for the single-warehouse multi-retailer problem with piecewise linear cost structures. url: citeseer.nj.nec.com/439759.html.
- [16] D. B. Shmoys. The design and analysis of approximation algorithms: facility location as a case study. In S. Hosten, J. Lee, and R. Thomas, editors, *Trends in Optimization. AMS Proceedings of Symposia in Applied Mathematics*, 2004.
- [17] D. B. Shmoys, E. Tardos, , and K. I. Aardal. Approximation algorithms for facility location problems. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 265–274, 1997.
- [18] H. M. Wagner and T. M. Whitin. Dynamic version of the economic lot sizing model. *Management Science*, 5:89–96, 1958.

[19] P. H. Zipkin. *Foundations of inventory management*. The McGraw-Hill Companies, Inc, 2000.