



HAL
open science

Fast Approximation Algorithms for the One-Warehouse Multi-Retailer Problem Under General Cost Structures and Capacity Constraints

Jean-Philippe Lucien Gayon, Guillaume Massonnet, Christophe Rapine,
Gautier Stauffer

► **To cite this version:**

Jean-Philippe Lucien Gayon, Guillaume Massonnet, Christophe Rapine, Gautier Stauffer. Fast Approximation Algorithms for the One-Warehouse Multi-Retailer Problem Under General Cost Structures and Capacity Constraints. *Mathematics of Operations Research*, 2017, 42 (3), pp.854-875. 10.1287/moor.2016.0830 . hal-01978078

HAL Id: hal-01978078

<https://hal.science/hal-01978078v1>

Submitted on 28 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fast approximation algorithms for the One-Warehouse Multi-Retailer problem under general cost structures and capacity constraints

Jean-Philippe Gayon

Univ. Grenoble Alpes, CNRS, G-SCOP, 38 000 Grenoble, France, jean-philippe.gayon@grenoble-inp.fr

Guillaume Massonnet

Laboratoire LS2N, Ecole des Mines de Nantes, 4 rue Alfred Kastler 44300 Nantes, France
guillaume.massonnet@mines-nantes.fr

Christophe Rapine

Université de Lorraine, laboratoire LGIPM, ile du Saulcy, 57045 Metz Cedex 01, France, christophe.rapine@univ-lorraine.fr

Gautier Stauffer

Univ. Grenoble Alpes, CNRS, G-SCOP, 38 000 Grenoble, France, gautier.stauffer@grenoble-inp.fr

We consider a well-studied multi-echelon (deterministic) inventory control problem, known in the literature as the One-Warehouse Multi-Retailer (OWMR) problem. We propose a simple and fast 2-approximation algorithm for this NP-hard problem, by recombining the solutions of single-echelon relaxations at the warehouse and at the retailers. We then show that our approach remains valid under quite general assumptions on the cost structures and under capacity constraints at some retailers. In particular, we present the first approximation algorithms for the OWMR problem with non-linear holding costs, truckload discount on procurement costs or with capacity constraints at some retailers. In all cases, the procedure is purely combinatorial and can be implemented to run in low polynomial time.

Key words: approximation algorithms, combinatorial algorithm, multi-echelon inventory problem, lot-sizing

Area of review: Discrete optimization

MSC2000 subject classification: Primary: 90B05, secondary: 68W25,90C59

OR/MS subject classification: Primary: Multi-item/echelon/stage < Inventory/production; Secondary:

Lot-sizing < Production smoothing < Inventory/production

History: Received June 19, 2015; revised December 22, 2015, and September 23, 2016; Accepted September 18, 2016

1. Introduction We study a two-level distribution system in which a central warehouse supplies several retailers facing final customer demands. The problem is known in the inventory literature as the (deterministic) *One-Warehouse Multi-Retailer* (OWMR) problem. We assume that the demands are deterministic and known over a discrete and finite planning horizon, and we want to minimize the distribution costs to flow the products from an external supplier to the final customers through the network. Two types of costs are incurred while the goods move through the system. Namely each location incurs procurement costs when it orders, and holding costs when it stores physical units in its stock. The objective is to find a planning for the orders at each location such that all the demands are fulfilled on time while minimizing the sum of the procurement and holding costs. The *Joint Replenishment Problem* (JRP) is a special case of the OWMR problem where the warehouse operates as a cross-docking station (i.e. holds no inventory). This can be captured in the model via a prohibitive holding cost at the warehouse.

This paper aims at designing efficient approximation algorithms for these problems. Recall that a polynomial time algorithm for a minimization problem is said to be a ρ -approximation, or to have a performance guarantee of ρ , if for any instance it delivers a solution of cost at most ρ times the optimal cost.

Literature Review. The OWMR and JRP problems have been studied extensively in the literature for continuous-review models. Schwarz (1973) is among the first to study the continuous-time version of the OWMR problem under constant-rate demands. While the complexity status of the problem under this setting is still open, Roundy (1985) introduced in its seminal work a 98%-effective algorithm which was later revisited and extended by Muckstadt and Roundy (1993) to other multi-echelon problems. More recently, Stauffer (2012) presented a simple 1.27-approximation based on the recombination of the optimal single-echelon strategies, which is also one of the basic ingredients of the approach presented in this paper. In the special case of the JRP, Nonner and Sviridenko (2013) proposed an EPTAS in a stationary continuous-time setting. We refer the reader to Aksoy and Erenguc (1998) and Khouja and Goyal (2008) for a detailed survey of the JRP.

In this article we focus on the discrete-time version of the problem, i.e. we consider periodic-review policies. Both the JRP and the OWMR problem are known to be NP-hard in this setting, see Arkin et al. (1989) or Nonner and Souza (2009). While it is not known if the standard versions of those problems are APX-hard (see Levi et al. (2008b)), the JRP with deadlines (a variant of the JRP problem where no holding cost is paid but the demands have to be served within certain intervals) was proven to be APX-hard first by Nonner and Souza (2009) and later by Bienkowski et al. (2014) in special cases. Only few papers propose algorithms or heuristics to solve the OWMR problem in its discrete-time version. Federgruen and Tzur (1999) have applied a time-partitioning heuristic, but considered bounded demands and parameters for the analysis of their worst-case bound. Chan et al. (2000) and later Shen et al. (2009) studied the class of Zero-Inventory-Ordering (ZIO) policies for the OWMR problem, i.e. in which the locations place orders only when their current inventory level is zero. Levi et al. (2006) proposed a 2-approximation algorithm for the JRP. Bienkowski et al. (2014) improved this result to a performance guarantee of 1.791 for the JRP, while Levi et al. (2008b) improved to a 1.8-approximation algorithm for the OWMR problem. Their results hold for shelf age dependent holding costs that are more general than linear holding costs (see Federgruen and Wang (2013) for a discussion on shelf age and level dependent holding costs). Furthermore in the special case of JRP with deadlines, Nonner and Souza (2009) and Bienkowski et al. (2013) have improved this result to a $5/3$ (respectively 1.574) approximation using a similar LP-rounding technique.

Contributions of this work. In contrast with the approximation algorithms presented in Levi et al. (2008b) and Nonner and Souza (2009) which exploit a linear integer formulation of the problem via LP-rounding techniques, we develop a combinatorial algorithm, based on a simple decomposition of the problem into single-echelon problems. This algorithm can be implemented to run in low polynomial time and yields a performance guarantee of two. The complexity of our algorithm can even be made linear for linear holding costs. We also introduce a new lower bound for this problem, as a byproduct of our analysis. While we do not match the best approximation guarantee of Levi et al. (2008b), we believe that the simplicity of the method and its computational complexity make it both a valuable theoretical approach and a practical tool. Besides, we demonstrate that our method can be applied to a significantly broader class of problems than the ones considered in the literature. In particular, we present the first approximation algorithms for the OWMR and JRP problems with non linear holding costs, general procurement costs (including Full Truck Load and Less than Truck Load cost structures) and capacity constraints at some retailers.

The rest of the paper is organized as follows. In §2, we present the model and important notation used throughout the rest of the paper. In §3, we introduce a natural decomposition of the OWMR problem into single-echelon subproblems on a simple model with a basic cost structure. We then propose a new method, called *uncrossing algorithm*, that recombines the solutions to these subproblems into a feasible solution for the original problem and prove that the overcost incurred by this solution is bounded by a constant factor.

The remaining sections show how the method can easily be adapted to numerous cost structures that capture more intricate cases. §4 shows how to modify the decomposition in order to tackle more complex holding cost structures, namely level dependent and shelf age holding costs. Specifically, we develop a 2-approximation when the holding costs are nonlinear functions of the inventory on hand, or when they satisfy an extension of the so-called Monge property introduced in Levi et al. (2008b). We then show in §5 that the algorithm can be adapted to general procurement costs that include the well-known FTL (*Full Truck Load*) and LTL (*Less than Truck Load*) transportation cost structures, under mild assumptions. Note that the FTL procurement cost structure is sometimes referred to as ordering with *soft capacities*. Finally we consider in §6 the case of a hard capacity constraint on the ordering of some specific retailers, defined as the W -retailers in the following.

Remark : For a quick overview of the core techniques and ideas presented in this paper, the reader can focus on sections §2 and §3. Note also that an early version of the results presented in sections §3 and §4 was published in the proceedings of SODA (Stauffer et al. 2011).

2. Model and notation We consider a two-level distribution network composed of one central warehouse (indexed by 0) and N retailers (indexed from 1 to N) facing customer demands. The planning horizon is finite, discretized into T time periods. Goods enter the system from an outside supplier of infinite capacity that replenishes the inventory of the warehouse, which in turn supplies the retailers. Assuming deterministic lead-time, we consider without loss of generality that the orders are delivered instantaneously from one location to another. Each retailer $i = 1, \dots, N$ faces in period $t = 1, \dots, T$ a deterministic demand d_t^i that must be satisfied on time, i.e. neither backlogging nor lost-sales are allowed.

We describe below the different costs incurred while products move across the network. The precise definitions and assumptions relative to the different cost structures will be given in the corresponding sections of the paper.

Holding costs. Each location can store goods in its inventory to serve demands – or, in the case of the warehouse, retailers orders – in future periods. Many models in the literature consider that a per-unit holding cost h^i is incurred to keep one unit of inventory in location i from the end of a period to the beginning of the next one. In what follows, this simple cost structure is referred to as the *linear* holding costs. In particular, we shall use this setting to explain the principles of our algorithms in §3. We also study two more general holding cost structures: Inventory *level dependent* holding costs in §4.1 and *shelf age* dependent holding costs in §4.2. In the level dependent case, the holding cost incurred in location i from period t to $t + 1$ is a nondecreasing function $h_t^i(\cdot)$ of its inventory level x . Linear holding costs are clearly a special case of level dependent holding costs for which $h_t^i(x) = h^i \cdot x$. In the shelf age case, the holding cost depends for each unit on the number of periods it has been stored in each location. Specifically, h_{rs}^{it} represents the total holding cost incurred to serve a unit of demand d_t^i by ordering it in period r at the warehouse and in period s at retailer i . Such a cost structure has been studied by Levi et al. (2008b). Linear holding costs are again a special case of shelf age dependent holding costs, where $h_{rs}^{it} = h^0(s - r) + h^i(t - s)$.

In the remainder of the paper, we assume that we can partition the set of retailers into two subsets I_J (J -retailers) and I_W (W -retailers) based on their holding cost. Specifically, it is cheaper to hold inventory at a J -retailer rather than at the warehouse, while it is the opposite for W -retailers (we will give a precise definition for each holding cost structure in the relevant sections). This assumption is classical in the literature: In particular in the case of linear holding costs we have $h^i < h^0$ if $i \in I_J$ and $h^i \geq h^0$ if $i \in I_W$. It is well-known that there exists an optimal solution that does not hold any stock at the warehouse for the J -retailers for the linear and shelf age holding cost structure. This dominance also holds under the level dependent holding costs we consider in §4.1, see Appendix A for a formal proof of this result.

Procurement costs. We denote by $p_t^i(q)$ the procurement cost incurred by location i to place an order of size q in period t . Many models in the literature focus on the special case of fixed ordering costs, in which $p_t^i(q) = K_t^i$ for all q , i.e. the procurement cost does not depend on the size q of the order. We shall use this simpler setting to present our algorithm in §3. In §5, we consider a more general ordering cost structure in which procurement costs $p_t^i(\cdot)$ model the well-known *Full Truck Load* (FTL) and *Less than Truck Load* (LTL) settings.

Note that Chan et al. (2000) have shown that if the procurement costs at the retailers vary over time, the OWMR problem is as hard to approximate as the set cover problem, even in the simple case of fixed ordering costs. Thus it is unlikely that there exists an approximation algorithm with constant guarantee unless $\mathcal{P} = \mathcal{NP}$, see Feige (1998). According to this result, we shall assume in this paper that the procurement costs at each retailer $i > 0$ are stationary, i.e. $p_t^i(\cdot) = p^i(\cdot)$ for all periods $t = 1, \dots, T$. In contrast, the procurement costs at the warehouse may be time-varying in our models.

Policies. We call indifferently *solution* or *policy* π a planning for the orders at each location. Any unit of product to satisfy a demand at retailer i in period t must be ordered in a period s at retailer i and in a period r at the warehouse such that $r \leq s \leq t$. In the remainder of this paper, we denote $[r, s]$ such a pair of orders, where the demand served with this pair will be clear from the context. Notice that if $d_t^i > 1$, the demand d_t^i can be served by different pair of orders. We say that a pair of orders $[r, s]$ is *valid* for a unit of demand d_t^i if and only if $r \leq s \leq t$. A policy is *feasible* if each unit of demand d_t^i is ordered via valid pairs of orders, for all $i = 1, \dots, N$ and $t = 1, \dots, T$.

We denote by $\mathcal{C}(\pi)$ the total cost incurred by a policy π over the planning horizon. This cost can be split into two parts: The total procurement cost $\mathcal{K}(\pi)$ and the total holding cost $\mathcal{H}(\pi)$. Thus we have $\mathcal{C}(\pi) = \mathcal{K}(\pi) + \mathcal{H}(\pi)$. The objective is to find a feasible policy minimizing the sum of the procurement costs and holding costs.

A policy is *First Come First Served* (FCFS) if units of product are consumed in the same order they are supplied. It is easy to check that for the cost structures considered in this paper there exists an optimal policy that is FCFS, therefore in what follows we restrict ourselves (w.l.o.g.) to such policies. For FCFS policies, we can also represent without ambiguity a policy π by a $(N + 1)$ -uplet $\pi = (\pi_0, \pi_1, \dots, \pi_N)$ where each π_i specifies the orders (time and quantity) for location i .

A policy is *Zero Inventory Ordering* (ZIO) if each location places orders only when its current inventory level is zero. Under a ZIO policy, a demand d_{it} is served by a single order. Under a non ZIO policy, a demand d_{it} might be served by several orders. The dominance of ZIO policies is well established when procurement costs do not depend on the size of the order, even under the various holding cost structures we consider. However, ZIO policies are not dominant when considering FTL and LTL procurement costs (see §5) or when the orders are capacitated (see §6).

3. A simple approximation algorithm This section introduces the main ideas of our technique on a simple model with stationary costs, i.e. each location i incurs a fixed ordering cost K^i and a per-unit holding cost h^i . Although simple, this cost structure renders the OWMR problem NP-hard, see Arkin et al. (1989). Recall that in this case, we have $h^i \geq h^0$ for all $i \in I_W$ and $h^i < h^0$ for all $i \in I_J$.

Our procedure works in two main phases, where the output of the first phase is used as an input for the second one. The first phase decomposes the OWMR problem into $N + 1$ single-echelon problems:

- Retailer i is considered as a single-echelon location facing its own demand d_t^i .
- The warehouse is regarded as a single-echelon, multi-item system facing for each period t a demand d_t^i for item i .

An optimal policy is computed for each single-echelon problem. In the second phase, we develop an *uncrossing algorithm* to recombine the $N + 1$ optimal single-echelon policies into a feasible solution for the OWMR problem. The rest of this section explains how the two phases can be adjusted to obtain a 2-approximation algorithm.

3.1. Decomposition of the problem We introduce two decompositions of the OWMR problem, which differ only by the way the holding costs are accounted in each problem.

A decomposition with full holding costs In this first decomposition, we consider the following single-echelon systems (S_i) :

(S_i) Each retailer i is considered as a single-echelon location facing demand d_t^i , with holding cost h^i and ordering cost K^i .

(S_0) The warehouse is regarded as a single-echelon, multi-item system facing for each period t a demand d_t^i for item i , with a fixed ordering cost K^0 . A different holding cost is incurred depending on which item (retailer) a unit serves: An item $i \in I_J$ incurs a holding cost h^i while an item $i \in I_W$ incurs a holding cost h^0 .

Observe that we consider real holding costs h^i at the retailers, and not echelon holding costs $(h^i - h^0)$ as it is classically done in the literature. The single-echelon problems (S_i) are commonly referred to as the *Uncapacitated Lot Sizing Problem*, which has been extensively studied in the literature since the seminal paper of Wagner and Whitin (1958), see for instance Brahimi et al. (2006) for a recent survey on single item lot-sizing problems or Pochet and Wolsey (2006) for polyhedral approaches. Problem (S_0) at the warehouse is multi-item while problems at retailers are single-item. There exists efficient algorithms to compute an optimal (ZIO) policy for these problems, see §3.4 for a detailed discussion.

Throughout the remainder of the paper, \mathcal{C}^* refers to the cost of an optimal policy for the OWMR problem and $\pi_0^*, \pi_1^*, \dots, \pi_N^*$ to the optimal policies for systems $(S_0), (S_1), \dots, (S_N)$, respectively. In addition, we denote $\mathcal{C}_i(\pi_i) = \mathcal{K}_i(\pi_i) + \mathcal{H}_i(\pi_i)$ the cost incurred by policy π_i in system (S_i) , with $\mathcal{K}_i(\pi_i)$ (resp. $\mathcal{H}_i(\pi_i)$) the total ordering cost (resp. holding cost) of policy π_i at location i .

The decomposition presented above induces two natural lower bounds for the problem. On the one hand, (S_0) is simply a OWMR problem with no fixed ordering cost for the retailers. One can note that the holding cost incurred to store a specific item i is equal to $\min\{h^0, h^i\}$ and therefore is lower than or equal to both its warehouse holding cost and its retailer holding cost. This ensures that (S_0) is a relaxation of the original problem. On the other hand, if we assume that there is no fixed ordering cost at the warehouse (which is yet another relaxation of our problem), then an optimal solution to this latter problem is simply the union of optimal solutions to the independent single-echelon systems (S_i) for $i = 1, \dots, N$. Both relaxations give a lower bound for the OWMR problem. We summarize this result in the following lemma.

LEMMA 1.

$$\mathcal{C}^* \geq \max \left\{ \mathcal{C}_0(\pi_0^*), \sum_{i=1}^N \mathcal{C}_i(\pi_i^*) \right\}$$

A decomposition with split holding costs Notice that in the natural decomposition above, the fixed ordering costs are naturally split between the different systems (S_i) , while the holding costs incurred by a particular unit in systems (S_0) and (S_i) may overlap. This explains the use of a maximum in the lower bound of Lemma 1. We shall aim at improving this result by considering a second decomposition, where holding costs are split between the two levels according to a *split parameter* $0 < \alpha < 1$. Specifically, we define the following single-echelon systems (\widehat{S}_i) :

(\widehat{S}_i) Each retailer i is considered as a single-echelon location facing demand d_t^i , with ordering cost K^i and holding cost $\widehat{h}^i = \alpha \cdot h^i$.

(\widehat{S}_0) The warehouse is regarded as a single-echelon, multi-item system facing for each period t a demand d_t^i for item i , with a fixed ordering cost K^0 . A different holding cost is incurred depending on which item (retailer) a unit serves: An item $i \in I_J$ incurs a holding cost $\hat{h}^0 = (1 - \alpha)h^i$ while an item $i \in I_W$ incurs a holding cost $\hat{h}^0 = (1 - \alpha)h^0$.

In the remainder of the paper, we denote by $\widehat{\mathcal{C}}_i(\pi_i)$ the cost incurred by policy π_i in system (\widehat{S}_i) and by $\widehat{\mathcal{H}}_i(\pi_i)$ the total holding cost in system (\widehat{S}_i) of a feasible policy π_i . Let also $\hat{\pi}_i^*$ be an optimal policy for (\widehat{S}_i). We have

$$\widehat{\mathcal{C}}_i(\pi_i) = \widehat{\mathcal{H}}_i(\pi_i) + \mathcal{K}_i(\pi_i).$$

The following lemma states that the sum of the optimal cost at the warehouse and at the retailers in subproblems (\widehat{S}_i) is a lower bound of any feasible policy for the OWMR problem:

LEMMA 2. *For any split parameter α such that $0 < \alpha < 1$, we have:*

$$\mathcal{C}^* \geq \widehat{\mathcal{C}}_0(\hat{\pi}_0^*) + \sum_{i=1}^N \widehat{\mathcal{C}}_i(\hat{\pi}_i^*).$$

PROOF. Consider an optimal policy $\pi^{\text{OPT}} = (\pi_0^{\text{OPT}}, \pi_1^{\text{OPT}}, \dots, \pi_N^{\text{OPT}})$ of cost \mathcal{C}^* for the OWMR instance. We assume w.l.o.g. that for all $i \in I_J$, every order of π_i^{OPT} is placed in an ordering period of π_0^{OPT} , hence no inventory is held at the warehouse for J -retailers.

Let x_{it}^{OPT} and $X_i^{\text{OPT}} = \sum_{t=1}^T x_{it}^{\text{OPT}}$ be the inventory level in period t and the cumulative inventory level over the entire planning horizon of π^{OPT} in location $i = 0, \dots, N$, respectively. Since policy π_i^{OPT} is feasible for system (\widehat{S}_i) for all i , we can evaluate the cost $\widehat{\mathcal{C}}_i(\pi_i^{\text{OPT}})$. One can use the inequality $h^0 \leq h^i \forall i \in I_W$ to bound the optimal cost as follows:

$$\begin{aligned} \mathcal{C}^* &= \sum_{i=0}^N \mathcal{K}_i(\pi_i^{\text{OPT}}) + h^0 X_0^{\text{OPT}} + \sum_{i=1}^N h^i X_i^{\text{OPT}} \\ &= \sum_{i=0}^N \mathcal{K}_i(\pi_i^{\text{OPT}}) + (1 - \alpha) \left[h^0 X_0^{\text{OPT}} + \sum_{i=1}^N h^i X_i^{\text{OPT}} \right] + \alpha \left[h^0 X_0^{\text{OPT}} + \sum_{i=1}^N h^i X_i^{\text{OPT}} \right] \\ &\geq \sum_{i=0}^N \mathcal{K}_i(\pi_i^{\text{OPT}}) + \left[(1 - \alpha)h^0 \left(X_0^{\text{OPT}} + \sum_{i \in I_W} X_i^{\text{OPT}} \right) + \sum_{i \in I_J} (1 - \alpha)h^i X_i^{\text{OPT}} \right] + \sum_{i=1}^N \alpha h^i X_i^{\text{OPT}} \\ &= \mathcal{K}_0(\pi_0^{\text{OPT}}) + \widehat{\mathcal{H}}_0(\pi_0^{\text{OPT}}) + \sum_{i=1}^N \left(\mathcal{K}_i(\pi_i^{\text{OPT}}) + \widehat{\mathcal{H}}_i(\pi_i^{\text{OPT}}) \right) \\ &= \widehat{\mathcal{C}}_0(\pi_0^{\text{OPT}}) + \sum_{i=1}^N \widehat{\mathcal{C}}_i(\pi_i^{\text{OPT}}). \end{aligned}$$

The proof then follows from the optimality of $\hat{\pi}_0^*, \hat{\pi}_1^*, \dots, \hat{\pi}_N^*$. □

3.2. The uncrossing algorithm From now on, we assume that we are given a feasible policy π_i to problem (S_i) (resp. (\widehat{S}_i)) for every $i = 0, \dots, N$. Since the single-echelon policies are computed independently, they define a policy $\pi = (\pi_0, \pi_1, \dots, \pi_N)$ that is not necessarily feasible for the OWMR problem. This issue arises if policy π_i orders a particular unit of demand d_t^i in period s while π_0 orders the same units in period $r > s$. We say that the pair of orders $[r, s]$ to serve a unit of demand d_t^i is *crossing* (or not feasible) if $r > s$. By extension, we say that a retailer order in period s is crossing if there exists some unit of demand d_t^i served by a pair $[r, s]$ such that $r > s$. In this section, we present the *uncrossing algorithm* that builds a feasible policy π^u for an OWMR

instance from single-echelon solutions to problems (S_i) (resp. (\widehat{S}_i)). We also show how to bound the cost of π^u .

Before giving the algorithm, we start by introducing some additional notation. We denote by $R = \{r_1, \dots, r_w\}$ the set of periods where the warehouse orders according to policy π_0 . For convenience we add to R two artificial periods $r_0 = 0$ and $r_{w+1} = T + 1$ corresponding to the beginning and the end of the planning horizon, with no ordering cost. For all $s \in \{1, \dots, T\}$, we define s^+ and s^- be the first period of R after s and the last period of R prior to s , respectively:

$$s^+ = \min \{r \in R : r \geq s\} \quad (1)$$

$$s^- = \max \{r \in R : r \leq s\} \quad (2)$$

When a pair of orders $[r, s]$ to serve a unit of demand d_t^i is crossing ($r > s$), the algorithm simply adds an additional order for retailer i at period s^+ to synchronize with the warehouse. This step suffices to obtain a feasible policy. In addition if i is a J -retailer and $s > r$, the algorithm shifts the retailer order from period s to period s^- . This transformation allows us to bound the holding costs in our analysis.

Uncrossing Algorithm

Input: A feasible policy π_i for each subproblem (S_i) (resp. (\widehat{S}_i)), $i = 0, \dots, N$

Output: A feasible policy π^u for the OWMR problem, defined as follows. Let r and s be the ordering periods to serve a unit of demand d_t^i in policy π_0 and π_i , respectively. Then the policy π^u serves this unit of demand with the pair of orders $[r^{it}, s^{it}]$, where:

$$[r^{it}, s^{it}] = \begin{cases} [s^+, s^+] & \text{if } r > s \\ [s^-, s^-] & \text{if } r \leq s \text{ and } i \in I_J \\ [r, s] & \text{if } r \leq s \text{ and } i \in I_W \end{cases} \quad (3)$$

The output policy π^u is feasible since $r^{it} \leq s^{it}$ for each unit of demand d_t^i . The algorithm works with ZIO or non ZIO policies. This will be useful in some of the extensions we consider for which the ZIO property is not dominant. When the input single echelon policies are ZIO, the uncrossing algorithm can be performed in $O(NT)$, as a demand d_t^i is served by a single pair of orders. The output policy is then also ZIO. Figure 1 illustrates the uncrossing algorithm on an example with two retailers by plotting the evolution of stocks over time.

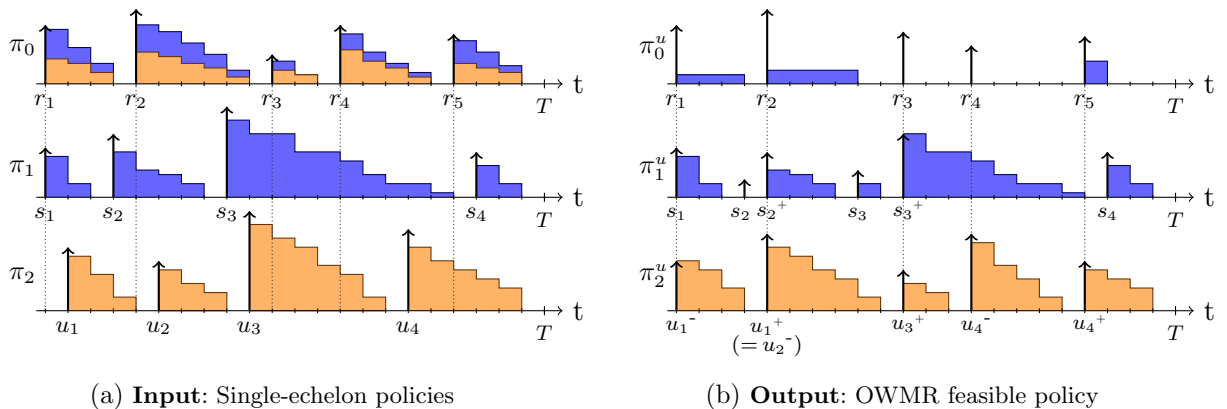


FIGURE 1. Uncrossing algorithm for $N = 2$, $I_W = \{1\}$ and $I_J = \{2\}$

By construction, each ordering period of π_i is replaced by at most two ordering periods at retailer i in π^u . More precisely, the set of ordering periods of retailer i in π^u is included in

$\{s^-, s^+ \mid s \text{ an ordering period of } \pi_i\}$ if i is a J -retailer, and in $\{s, s^+ \mid s \text{ an ordering period of } \pi_i\}$ if i is a W -retailer. Thus the number of ordering periods at each retailer is at most doubled by the uncrossing algorithm. Moreover policy π^u orders at the same periods as π_0 at the warehouse. Hence, together with the assumption of stationary ordering costs K_i , we can bound the procurement costs of policy π^u as follows:

$$\mathcal{K}(\pi^u) \leq \mathcal{K}_0(\pi_0) + 2 \sum_{i=1}^N \mathcal{K}_i(\pi_i). \quad (4)$$

We show now that the total holding cost incurred by policy π^u is at most the sum of the holding costs of the single-echelon policies in systems (S_i) . Consider a specific unit of demand d_i^i and let r and s be the ordering periods of this unit in the single-echelon policies π_0 and π_i , respectively. The holding cost incurred by this unit, in the single echelons systems, is equal to $h^i(t-s)$ in π_i and to $\min\{h^0, h^i\}(t-r)$ in π_0 . The holding cost h_{it}^u incurred by the same unit of demand in the final policy π^u is equal to:

$$h_{it}^u = h^0(s^{it} - r^{it}) + h^i(t - s^{it}) = \begin{cases} h^i(t - s^+) & \text{if } r > s \\ h^i(t - s^-) & \text{if } r \leq s \text{ and } i \in I_J \\ h^0(s - r) + h^i(t - s) & \text{if } r \leq s \text{ and } i \in I_W \end{cases} .$$

As $s^+ \geq s$ by definition and $s^- \geq r$ if $s \geq r$, it follows that $h_{it}^u \leq h^i(t-s) + \min\{h^0, h^i\}(t-r)$. Summing the holding costs over all retailers and periods gives:

$$\mathcal{H}(\pi^u) \leq \sum_{i=0}^N \mathcal{H}_i(\pi_i). \quad (5)$$

That is, the uncrossing algorithm can only decrease the total holding cost incurred by the policies. Using inequalities (4) and (5), we can now state the following lemma.

LEMMA 3. *Given feasible single-echelon policies π_i to problems (S_i) , the uncrossing algorithm builds in time $O(NT)$ a feasible policy π^u for the OWMR problem with linear holding costs such that*

$$\mathcal{C}(\pi^u) \leq 2 \sum_{i=0}^N \mathcal{K}_i(\pi_i) + \sum_{i=0}^N \mathcal{H}_i(\pi_i) - \mathcal{K}_0(\pi_0).$$

3.3. Performance guarantees We start by considering the solution π^u obtained when the uncrossing algorithm is applied to the optimal single-echelon policies of subproblems (S_i) . In this case, Lemmas 1 and 3 imply that

$$\mathcal{C}(\pi^u) \leq \mathcal{C}_0(\pi_0^*) + 2 \sum_{i=1}^N \mathcal{C}_i(\pi_i^*) \leq 3 \cdot \max \left\{ \mathcal{C}_0(\pi_0^*), \sum_{i=1}^N \mathcal{C}_i(\pi_i^*) \right\} \leq 3 \cdot \mathcal{C}^*$$

and therefore the algorithm provides a 3-approximation. However, using the optimal single-echelon policies of subproblems (\hat{S}_i) as an input for the uncrossing algorithm leads to a better performance guarantee when the split parameter α is set to $1/2$. In fact, we prove that this procedure, referred to as the *Split & uncross algorithm* in the rest of the paper, is a 2-approximation for the OWMR with linear holding costs:

THEOREM 1 (Split & uncross algorithm). *For the OWMR problem with linear holding costs, the uncrossing algorithm has a performance guarantee of 2 when the input policies are the optimal single echelon policies of systems $(\hat{S}_0), (\hat{S}_1), \dots, (\hat{S}_N)$ and the split parameter is set to $\alpha = 0.5$.*

PROOF. Let $\hat{\pi}^u$ denote the output feasible policy of the uncrossing algorithm when the input policies are the optimal single echelon policies of systems $(\hat{S}_1), \dots, (\hat{S}_N)$. For all $i = 0, \dots, N$, the cost of policy $\hat{\pi}_i^*$ when applied to problem (S_i) is:

$$\mathcal{C}_i(\hat{\pi}_i^*) = \mathcal{H}_i(\hat{\pi}_i^*) + \mathcal{K}_i(\hat{\pi}_i^*) \leq \begin{cases} \frac{1}{\alpha} \hat{\mathcal{H}}_i(\hat{\pi}_i^*) + \mathcal{K}_i(\hat{\pi}_i^*) & \text{if } i \geq 1 \\ \frac{1}{1-\alpha} \hat{\mathcal{H}}_i(\hat{\pi}_i^*) + \mathcal{K}_i(\hat{\pi}_i^*) & \text{if } i = 0 \end{cases}.$$

Then we have from Lemma 3 (for all $0 < \alpha < 1$):

$$\begin{aligned} \mathcal{C}(\hat{\pi}^u) &\leq \mathcal{C}_0(\hat{\pi}_0^*) + \sum_{i=1}^N \mathcal{C}_i(\hat{\pi}_i^*) + \sum_{i=1}^N \mathcal{K}_i(\hat{\pi}_i^*) \\ &= \mathcal{H}_0(\hat{\pi}_0^*) + \mathcal{K}_0(\hat{\pi}_0^*) + \sum_{i=1}^N \mathcal{H}_i(\hat{\pi}_i^*) + 2 \sum_{i=1}^N \mathcal{K}_i(\hat{\pi}_i^*) \\ &\leq \frac{1}{1-\alpha} \hat{\mathcal{H}}_0(\hat{\pi}_0^*) + \mathcal{K}_0(\hat{\pi}_0^*) + \sum_{i=1}^N \frac{1}{\alpha} \hat{\mathcal{H}}_i(\hat{\pi}_i^*) + 2 \sum_{i=1}^N \mathcal{K}_i(\hat{\pi}_i^*) \\ &\leq \frac{1}{1-\alpha} \hat{\mathcal{C}}_0(\hat{\pi}_0^*) + \sum_{i=1}^N \max \left\{ 2, \frac{1}{\alpha} \right\} \hat{\mathcal{C}}_i(\hat{\pi}_i^*). \end{aligned}$$

Hence, from Lemma 2, the split & uncross algorithm has a performance guarantee of $\max \left\{ 2, \frac{1}{\alpha}, \frac{1}{1-\alpha} \right\}$. The minimum of this approximation ratio is 2 when $\alpha = 0.5$. \square

Notice that the ordering periods at the warehouse are not changed in $\hat{\pi}^u$. Thus cost $\mathcal{C}(\hat{\pi}^u)$ is in fact bounded by $2 \sum_{i=0}^N \mathcal{C}_i(\hat{\pi}_i^*) - \mathcal{K}_0(\hat{\pi}_0^*)$. It follows that we can get *a posteriori* a better guarantee than 2 for our algorithm: For instance if in the relaxation the total ordering costs at the warehouse represents $1/4$ of the total cost of $\hat{\pi}^*$, we get a 1.75-approximate solution.

3.4. Complexity analysis Besides its simplicity, the strength of this technique lies in its low computational complexity. Recall that the uncrossing algorithm can be performed in $O(NT)$ when the input single echelon policies are ZIO. It remains to determine the complexity of solving to optimality each single-echelon problem (S_i) (resp. (\hat{S}_i)) for $i = 0, \dots, N$. As the ZIO property is dominant, one can classically (see e.g. Zipkin (2000)) represent the single-echelon problem (S_i) (resp. (\hat{S}_i)) with a graph $G_i = (V_i, E_i)$ where the vertices are the different periods (i.e. $V_i = \{1, \dots, T+1\}$) and an edge $(s, t) \in E_i$ represents two consecutive ordering periods (or the last ordering period and $T+1$ if necessary), i.e. $E_i = \{(s, t) : 1 \leq s < t \leq T+1\}$. In addition, for $i \geq 1$ we add an artificial vertex labeled 0 to V_i and artificial edges $(0, t)$ to E_i for $t = 1, \dots, T+1$ to consider the possible zero demands at the beginning of the planning horizon. For the single-echelon problems corresponding to the retailers (i.e. $i \in \{1, \dots, N\}$), the length $l_{s,t}^i$ of each edge $(s, t) \in E_i$ is set to:

$$l_{s,t}^i = \begin{cases} K^i + \alpha \sum_{u=s}^{t-1} (u-s) h^i \cdot d_u^i & \text{if } s \geq 1 \\ 0 & \text{if } s = 0 \text{ and } \sum_{u=1}^{t-1} d_u^i = 0 \\ \infty & \text{otherwise} \end{cases}.$$

This length corresponds to the cost incurred for placing an order in period s and holding units to satisfy the demands of period s till period $t-1$. On the other hand, these lengths for the graph G_0 corresponding to the warehouse (multi-item uncapacitated lot-sizing problem) are set to:

$$l_{s,t}^0 = K^0 + (1-\alpha) \sum_{u=s}^{t-1} \left(\sum_{i \in I_W} (u-s) h^0 \cdot d_u^i + \sum_{i \in I_J} (u-s) h^i \cdot d_u^i \right).$$

Since h^i are nonnegative, these lengths are also nonnegative. One can then build a planning for the orders in the inventory problem (S_i) (resp. (\hat{S}_i)) from a path $0 = t_1^i < t_2^i < \dots < t_m^i = T + 1$ by ordering in every period corresponding to a vertex $1 \leq t_k^i \leq T$ in the path (see e.g. Zipkin (2000)). The optimal solution for the corresponding single-echelon problem is then simply the shortest path from 0 to $T + 1$ in G_i . Bellman (1958) algorithm solves such a problem in $O(T^2)$ time, leading to an overall complexity of $O(NT^2)$.

It is interesting to note that when the system has a linear holding cost structure, one can use advanced dynamic programming techniques to solve the single-echelon problems more efficiently. In the case of non-speculative motive, there even exist improved algorithms that find an optimal solution to the single-echelon problem in $O(T)$ time (See Federgruen and Tzur (1991), Wagelmans et al. (1992) and Aggarwal and Park (1993)). Using one of these techniques, it is possible to solve the $N + 1$ single-echelon problems with linear holding cost structures in time $O(NT)$ and the overall complexity of our algorithm then decreases to $O(NT)$.

To the best of our knowledge, the complexity of our algorithm outperforms any existing constant approximation algorithm for the OWMR problem.

3.5. A bad example In this section, we exhibit an instance of the OWMR problem for which the worst-case bound presented in Theorem 1 is tight. We consider a single warehouse that supplies two retailers over three periods, with a stationary linear holding cost structure. The parameters have the following values:

- $K^0 = 1, h^0 = 1$
- $K^1 = 0, h^1 = 0, d_1^1 = 1, d_2^1 = 0, d_3^1 = 0$
- $K^2 = z, h^2 = 2, d_1^2 = 0, d_2^2 = 1, d_3^2 = 1 + \varepsilon$, where $z \geq 2$ and $0 < \varepsilon < 1$

The optimal ZIO policies for the single-echelon systems with holding costs halved are unique and can be easily computed:

- The *warehouse* orders in periods 1 and 3.
- *Retailer 1* orders in period 1.
- *Retailer 2* orders in period 2.

Starting from these optimal single-echelon solutions, the uncrossing algorithm adds an order for retailer 2 in period 3 to obtain a solution for the original problem. Only one unit is held in the warehouse in period 1 (to serve d_2^2) and the final cost it incurs is:

$$\mathcal{C}(\pi^u) = 2K^0 + K^1 + 2K^2 + h^0 \times 1 = 2z + 3.$$

Another feasible policy π for this instance is such that the warehouse and retailer 1 order only in period 1, while retailer 2 orders only in period 2. The warehouse then holds demands d_2^2 and d_3^2 in period 1, before it supplies retailer 2 in period 2. Thus, it incurs a total cost of

$$\mathcal{C}(\pi) = K^0 + K^1 + K^2 + h^0 \times (2 + \varepsilon) + h^2 \times (1 + \varepsilon) = z + 5 + 3\varepsilon.$$

The worst case bound is then reached asymptotically as z goes to infinity:

$$\lim_{z \rightarrow \infty} \frac{\mathcal{C}(\pi^u)}{\mathcal{C}(\pi^*)} \geq \lim_{z \rightarrow \infty} \frac{\mathcal{C}(\pi^u)}{\mathcal{C}(\pi)} = \lim_{z \rightarrow \infty} \frac{2z + 3}{z + 5 + 3\varepsilon} = 2.$$

4. Extension to more general holding cost structures We now extend the result of the previous section to models with more general holding cost structures. Namely, we focus on two different classes of costs: The level dependent and the shelf-age holding cost structures. Basically, the former considers that the holding cost incurred in a period is an increasing function of the inventory on hand, while the latter assumes that this cost depends for each unit on the number of periods it is held in each location. Linear holding costs are a particular case of both cost structures.

In order to generalize Theorem 1, we shall show that the results of Lemmas 2 and 3 extend to these new models. That is, the sum of the optimal costs for the single-echelon subproblems is a lower bound for the original problem, on one hand, and the split & uncross algorithm at most doubles the total cost of these $N + 1$ policies, on the other hand.

4.1. Non-linear level dependent holding costs

Assumptions. We consider a level dependent holding cost structure: Recall that $h_t^i(x)$ represents the price to hold an amount x of products in stock at location i from period t to $t + 1$. Notice that the level dependent holding cost is memoryless, i.e. the price paid in period t only depends on the current stock level, no matter how long the products have been in stock. We show in what follows that it is possible to extend the main results of the previous section and obtain a 2-approximation for non-linear holding costs under quite weak assumptions, namely:

(LD1) Non-decreasing property. We assume that holding cost $h_t^i(\cdot)$ is non-decreasing with respect to the stock level for each location i and for each period t , that is $x \leq y \Rightarrow h_t^i(x) \leq h_t^i(y)$ for all $t = 1, \dots, T$. In addition, we assume that $h_t^i(0)$ is nonnegative.

(LD2) Sub-additivity property at the warehouse. We assume that holding cost $h_t^0(\cdot)$ at the warehouse is sub-additive with respect to the stock level in each period, that is $h_t^0(x + y) \leq h_t^0(x) + h_t^0(y)$ for all $t = 1, \dots, T$.

(LD3) Partition of the retailers. For each retailer i , either $h_t^i(x) \geq h_t^0(x)$ holds for any period t and any stock level x (i is a W -retailer), or $h_t^i(y + q) - h_t^i(y) \leq h_t^0(x + q) - h_t^0(x)$ holds for any period t and any quantities x, y and q (i is a J -retailer).

It is straightforward to see that the linear holding costs satisfy properties (LD1), (LD2) and (LD3). In fact as long as shortages and backlogging are not allowed (i.e. inventory levels are constrained to be non-negative), the non-decreasing property appears to be fulfilled by any reasonable holding cost function. Sub-additivity is a common property in economical function that captures economies of scale. Note that the sum of two sub-additive functions is a sub-additive function, as well as the product by a positive constant. Functions like \sqrt{x} or $\lceil x/B \rceil$ are examples of sub-additive functions. The latter function may model the practical case where racks of capacity B are used to store the items and a cost is paid for each additional rack needed. Note also that a concave function that is non-negative at zero is sub-additive. The reverse is not true, $\lceil x/B \rceil$ is sub-additive but not concave.

Assumptions (LD1) and (LD2) on holding costs appear to be sufficiently weak to include a large variety of situations often met in practice. However, the more stringent condition (LD3) on the marginal holding costs has to be satisfied in order to partition the set of retailers into two subsets I_W and I_J . As for linear holding costs, property (LD3) states that a retailer i is in I_W if, for all periods, it is more expensive to hold inventory at the retailer than at the warehouse, i.e. $h_t^i(x) \geq h_t^0(x)$ for any period t and any stock level x . On the other hand, we would like to assess that if i is a J -retailer, there exists an optimal policy that does not hold any stock at the warehouse to supply its orders, i.e. in which every ordering period of retailer i is synchronized with a warehouse ordering period. Notice that this property may be imposed by practical considerations that force the warehouse to behave as a cross-dock facility for a specific subset of retailers. More generally, property (LD3) states that a retailer i belongs to I_J if its marginal holding cost is always less than or equal to the marginal holding cost at the warehouse. This condition on the marginal holding costs for a J -retailer ensures precisely that it is dominant to synchronize its orders with the warehouse (see Appendix A for the proof). For example, consider the case where the holding cost of the warehouse is of the form $\alpha \lceil x/B \rceil + \beta x$, which models a unitary holding cost plus a rack overcost as already discussed. A retailer i belongs to I_W if its holding cost is greater than $h^0(x)$ and to I_J if its marginal holding cost is at most β . For instance if retailer i has a linear holding cost γx , it belongs to set I_W if $\gamma \geq \alpha + \beta$ and to set I_J if $\gamma \leq \beta$.

Decomposition into subproblems (splitting phase). Similarly to the linear holding costs case, we decompose the OWMR problem into $N + 1$ independent single-echelon subsystems.

(\widehat{S}_i) Retailer i is considered as a single-echelon location facing demand d_t^i , with fixed ordering cost K^i and holding cost $\widehat{h}_t^i(\cdot) \equiv \frac{1}{2}h_t^i(\cdot)$ for all $t = 1, \dots, T$.

(\widehat{S}_0) The warehouse is regarded as a single-echelon, multi-item system with a fixed ordering cost K_r^0 in period r . Each retailer $i \in I_J$ plays the role of an item i , while item 0 represents the aggregated demand of all the retailers in I_W . We charge the units x on hand in period t at cost $\widehat{h}_t^0(\cdot) \equiv \frac{1}{2}h_t^0(\cdot)$ for item 0, and at cost $\widehat{h}_t^i(\cdot)$ for each item $i \in I_J$.

In a similar way as in §3.3 we solve independently to optimality the N single-echelon problems at the retailers and the multi-item problem at the warehouse to obtain a vector $\widehat{\pi}^* = (\widehat{\pi}_0^*, \dots, \widehat{\pi}_N^*)$ of optimal ZIO policies. We then rebuild a feasible policy π^u for the original problem using the uncrossing algorithm.

Analysis. We now show that the performance guarantee of the Split and Uncross algorithm remains the same under non-linear holding costs. The analysis follows the two steps defined by Lemmas 3 and 2. We prove that these results remain valid under the more general case of level dependent holding cost by combining similar arguments with properties (LD1), (LD2) and (LD3).

LEMMA 4. *Given independent single-echelon feasible ZIO policies π_i to problems (\widehat{S}_i), the uncrossing algorithm builds a feasible policy π^u to the OWMR problem with level dependent holding costs such that:*

$$\mathcal{C}(\pi^u) \leq 2 \left(\widehat{\mathcal{C}}_0(\pi_0) + \sum_{i=1}^N \widehat{\mathcal{C}}_i(\pi_i) \right).$$

A formal proof of this result is given in Appendix B for the slightly more general case where policies π_i at the retailers are PCO, see Definition 2 in §5, but we give the reader some insights here. First, it is easy to see that when recomposing independent uncrossing policies into a feasible policy π^u for OWMR, the stock level at the warehouse and at the W -retailers, in policy π^u , is less than or equal to the one in the independent policies in each period. Due to the non-decreasing property (LD1), the resulting policy pays less to hold inventory. The same holds for J -retailers except for periods where the algorithm shifts orders to synchronize with the warehouse. However we can show that the holding cost paid in π_0 can be used to pay for the extra holding costs incurred in π_i^u . Then, transforming policies π_i at the retailers into uncrossing policies at most doubles the total ordering cost, while changing from holding costs $\widehat{h}(\cdot)$ to $h(\cdot)$ simply doubles the total holding cost paid by the policy. This leads to the inequality expressed in Lemma 4.

The next step to complete our analysis is to extend the lower bound of Lemma 2 to non-linear level dependent holding costs. Sub-additivity of holding costs at the warehouse is required to prove the following result:

LEMMA 5. *Let \mathcal{C}^* be the cost of an optimal policy for the OWMR problem with level dependent holding cost structure and let $\widehat{\pi}_i^*$ be an optimal policy for single-echelon problem (\widehat{S}_i), for $i = 0, \dots, N$. If the holding costs satisfies properties (LD1), (LD2) and (LD3) then*

$$\mathcal{C}^* \geq \widehat{\mathcal{C}}_0(\widehat{\pi}_0^*) + \sum_{i=1}^N \widehat{\mathcal{C}}_i(\widehat{\pi}_i^*).$$

The proof uses quite similar arguments as the ones of Lemma 2. We combine these ideas with properties (LD1), (LD2) and (LD3) to obtain the lower bound. See Appendix C for a detailed proof. The following theorem is a corollary of Lemmas 4 and 5 and extends Theorem 1 to non-linear holding cost functions:

THEOREM 2. *The uncrossing algorithm applied to optimal single-echelon policies for systems (\widehat{S}_i) is a 2-approximation for the OWMR problem assuming non-decreasing level dependent holding costs at each location and sub-additive holding costs at the warehouse, see properties (LD1), (LD2) and (LD3). Its time complexity is $O(NT^3)$, provided that any cost $h_t^i(x)$ can be evaluated in constant time.*

PROOF. The time complexity corresponds to the resolution to optimality of the different (single-echelon) uncapacitated lot sizing problems. With general holding costs, the computation of the value of the arcs as described in §3.4 requires $O(T^3)$ operations for a retailer problem (\widehat{S}_i) and $O(NT^3)$ operations for the warehouse problem (\widehat{S}_0). The other steps are linear in the size of the OWMR problem, i.e. in $O(NT)$. \square

To the best of our knowledge this is the first approximation algorithm for the OWMR problem for this class of holding costs.

4.2. Shelf age dependent holding costs

The metric holding cost structure. Although level dependent holding cost structure is a classical assumption in the inventory literature, it remains restrictive and fails to modelize many practical situations. In Levi et al. (2008b), the authors introduce a shelf age dependent holding cost structure that captures additional phenomena such as perishable goods. In this case the cost of holding one unit of product depends on how long this specific unit has been physically held in the stock. Specifically, the cost incurred by a unit ordered with the pair $[r, s]$ to serve demand d_t^i incurs a total holding cost of h_{rs}^{it} . This cost includes in particular the cost of holding this unit from period r to s at the warehouse and from period s to t at retailer i . Notice that although we refer to them as holding costs, these parameters can also include per-unit ordering costs at the warehouse and at the retailer. The authors assume that these holding cost parameters obey the following so-called Monge properties (see Levi et al. (2008b) for a detailed discussion of these properties):

(SA1) *Non-negativity.* Parameters h_{rs}^{it} are nonnegative.

(SA2) *Monotonicity with respect to s .* Each retailer i is in exactly one of the two following situations: Either h_{rs}^{it} is non-increasing in $s \in [r, t]$ for each demand point (i, t) and warehouse order r , or h_{rs}^{it} is non-decreasing in $s \in [r, t]$ for each demand point (i, t) and warehouse order r . This property defines a partition of the set of retailers into two subsets: I_W and I_J , respectively. A retailer in I_W will be called for short a W -retailer, similarly a J -retailer if it belongs to I_J .

(SA3) *Monotonicity with respect to r .* For each demand point (i, t) and retailer order in period s , h_{rs}^{it} is non-increasing in $r \in [1, s]$. Moreover for each retailer i in I_J and demand point (i, t) , for each $r' < r \leq t$, the inequality $h_{r'r'}^{it} \geq h_{rr}^{it}$ holds.

(SA4) *Monge property.* For each demand point (i, t) with $i \in I_W$, and $r_2 < r_1 \leq s_2 < s_1 \leq t$, the inequality $h_{r_2, s_1}^{it} + h_{r_1, s_2}^{it} \geq h_{r_2, s_2}^{it} + h_{r_1, s_1}^{it}$ holds.

These properties are satisfied by linear holding costs. We call this holding cost structure *Monge holding costs*. Observe that as in the linear holding cost setting, I_W corresponds to the set of retailers for which it is cheaper to hold inventory at the warehouse while I_J refers to retailers for which it is cheaper to hold inventory at the retailer. In this paper we generalize the Monge holding costs by relaxing property (SA4) to property (SA4'):

(SA4') *Triangle inequality.* For each demand point (i, t) with $i \in I_W$, and $r \leq s \leq t$, the inequality $h_{rt}^{it} + h_{ss}^{it} \geq h_{rs}^{it}$ holds.

In what follows, we denote (SA) the set of assumptions (SA1), (SA2), (SA3) and (SA4'). Holding costs that satisfy the (SA) will be referred to as *metric holding costs* as property (SA4') resembles a triangle inequality. Indeed this inequality implies that the cost of holding a unit at the warehouse from period r to t plus the cost of holding this unit at the retailer from period s to t encapsulates the cost of keeping the unit at the warehouse from period r to s and then at the retailer from period s to t . It is clearly true for linear holding costs. The metric holding costs also ensure that each demand d_t^i can be served from a unique pair of orders $[r, s]$ in an optimal solution. Since demands are nonnegative by definition, this means one can include d_t^i into a parameter $H_{rs}^{it} = h_{rs}^{it} \cdot d_t^i$ such that if h_{rs}^{it} satisfies properties (SA1), (SA2), (SA3) and (SA4'), so does H_{rs}^{it} . Hence we shall assume w.l.o.g. that $d_t^i = 1$ for all i, t such that $d_t^i \geq 0$ in the remainder of this section.

Besides being more intuitive, one can see that the metric holding costs generalize the Monge holding costs, in the sense that Property (SA4) implies (SA4'): The triangle inequality is obtained by instantiating the Monge property with $r_2 = r$, $r_1 = s_2 = s$ and $s_1 = t$. In practice, they also address additional situations compared to their Monge counterpart. For example, any decomposable holding cost function of the form $H_{rs}^{it} := f_0(s - r) + f_i(s, t)$, where f_0 and f_i are non-negative functions, satisfy the metric holding cost assumptions when f_0 is non-decreasing. On the other hand, f_0 needs to be convex for the Monge condition to hold. This excludes cases where the items need a periodic maintenance operation that induces a fixed charge. For instance one can think of a plant distributor who hires a gardenener to water their goods every $\delta > 1$ periods: If this function takes the form $f_0(x) = \gamma \cdot \lceil \frac{x}{\delta} \rceil$, it is increasing and concave and hence is metric, but not Monge. Note that Levi et al. (2008b) explicitly use the Monge property (SA4) in their analysis to prove the better performance guarantee of their algorithm.

Decomposition into subproblems (splitting phase). We again introduce for all i, t and all pairs of orders $[r, s]$ an artificial holding cost parameter $\hat{h}_{rs}^{it} = \frac{1}{2}h_{rs}^{it}$ that we use to define the subproblems of the decomposition in the case of (shelf age) metric holding costs:

(\hat{S}_i) Retailer i is considered as a single-echelon location facing demand d_t^i , with ordering cost K^i and holding cost \hat{h}_{ss}^{it} for all $t = 1, \dots, T$ and ordering period $s \leq t$.

(\hat{S}_0) The warehouse is regarded as a single-echelon, multi-item system facing for each period t a demand d_t^i for item i , with a fixed ordering cost K_r^0 in period r . A different holding cost is incurred depending on which item (retailer) the units are intended to serve: if the units are ordered in period r , the corresponding holding cost is then \hat{h}_{rt}^{it} for $i \in I_W$ and \hat{h}_{rr}^{it} for $i \in I_J$.

Here again, the algorithm starts by solving independently to optimality the N single-echelon problems at the retailers and the multi-item problem at the warehouse to obtain a vector $\hat{\pi}^* = (\hat{\pi}_0^*, \dots, \hat{\pi}_N^*)$ of optimal ZIO policies and then uncrosses these solutions to build a feasible policy π^u for the OWMR problem with metric holding costs.

Analysis. The next lemma extends the results of Lemma 3 to an OWMR problem with metric holding cost structure. The proof is detailed in Appendix D.

LEMMA 6. *Given single-echelon feasible policies π_i to problems (\hat{S}_i), the uncrossing algorithm builds in time $O(NT)$ a feasible and uncrossing policy π^u for the OWMR problem with metric holding costs such that*

$$\mathcal{C}(\pi^u) \leq \mathcal{C}_0(\pi_0) + \sum_{i=1}^N \mathcal{C}_i(\pi_i) + \sum_{i=1}^N \mathcal{K}_i(\pi_i).$$

In order to obtain a performance guarantee of two, it remains to show that the sum of the costs of the optimal single echelon solutions to subproblems (\hat{S}_0), (\hat{S}_1), \dots , (\hat{S}_N) is a lower bound on the optimal cost.

LEMMA 7. *Let \mathcal{C}^* be the cost of an optimal policy for the OWMR problem with metric holding cost structure, and $\hat{\pi}_i^*$ be an optimal policy for the single-echelon problem (\hat{S}_i), $i = 0, \dots, N$. We have:*

$$\mathcal{C}^* \geq \hat{\mathcal{C}}_0(\hat{\pi}_0^*) + \sum_{i=1}^N \hat{\mathcal{C}}_i(\hat{\pi}_i^*).$$

One can prove this result by simply breaking apart an optimal policy to the original problem into feasible policies for the problems (\hat{S}_i), in a similar fashion as the proof of Lemma 5. A formal proof is given in Appendix E.

Since Phase 2 of the procedure applies exactly the uncrossing algorithm introduced in the previous section, one can use similar arguments as the one introduced in §3, combined with Lemmas 6 and 7, to establish the following theorem:

THEOREM 3. *The uncrossing algorithm based on optimal policies to the decomposition (\widehat{S}_i) has a performance guarantee of 2 and a time complexity of $O(NT^2)$ for the one-warehouse multi-retailer problem with metric holding costs.*

PROOF. One can represent the single-echelon subproblems with a graph as in the linear case. The length of each edge $l_{s,t}^i$ in the graph G_i representing retailer i is defined as

$$l_{s,t}^i = \begin{cases} K^i + \sum_{u=s}^{t-1} H_{ss}^{iu} & \text{if } s \geq 1 \\ 0 & \text{if } s = 0 \text{ and } \sum_{u=1}^{t-1} d_u^i = 0 \\ \infty & \text{otherwise} \end{cases} .$$

Similarly, the value of each edge in the graph G_0 is defined as $l_{s,t}^0 = K^0 + \sum_{u=s}^{t-1} \left(\sum_{i \in I_W} H_{su}^{iu} + \sum_{i \in I_J} H_{ss}^{iu} \right)$. Hence these problems can be solved in $O(T^2)$ time using Bellman algorithm, leading to an overall complexity of $O(NT^2)$. \square

5. Extension to more general procurement cost structures In the previous sections we have assumed a fixed cost per order at each location. We now aim to generalize our approach to more general procurement costs $p_t^i(x)$ charged at location i to order an amount x of units in period t . In the literature, two procurement cost structures have been extensively studied, due to their practical relevance: The FTL (Full Truck Load) and the LTL (Less than Truck Load) procurement costs (Li et al. 2004). Indeed, supplies are often delivered by batch, corresponding to truck capacity, and it is common to pay a fixed cost per truck plus a variable cost according to the actual load of the truck. In the inventory terminology, such costs are also referred to as multiple setup costs, stepwise costs, truckload discount or soft capacities. We define these costs as follows:

FTL. The cost to order $x > 0$ units in period t at location i is $p_t^i(x) = K_t^i + \lceil x/B^i \rceil k^i$. That is, a fixed cost k^i is charged for each batch of size B^i used to supply the order, in addition to a fixed ordering cost K_t^i .

LTL. The cost to order $x > 0$ units in period t at location i is $p_t^i(x) = K_t^i + \lfloor x/B^i \rfloor k^i + f^i(x - B^i \lfloor x/B^i \rfloor)$, where f is a non-decreasing function such that $f^i(0) = 0$ and $f^i(x) \leq k^i$ for $x \leq B^i$. That is, in addition to a fixed ordering cost K_t^i , a fixed cost k^i is charged for each full batch of size B^i used to supply the order and the (last) partially loaded batch is charged according to f . We call f the *LTL freight cost function*.

Notice that we restrict ourself to the case of a stationary batch size B^i and a stationary fixed cost per batch k^i at each location. In the remainder of this section we assume also that the fixed ordering cost at the retailers is stationary, that is $K_t^i = K^i$ for all $i = 1, \dots, N$. Observe that a fixed ordering cost is a special case of FTL/LTL procurement costs with fixed cost per batch equal to zero. In the following, we will consider 2-linearly sandwich procurement costs. We define a λ -linearly sandwiched function as follows.

DEFINITION 1. Given a positive λ , functions $p_t(\cdot), t = 1, \dots, T$ are said to be λ -linearly sandwiched if there exists some constants $(A_t)_{t=1,T}$ and a constant b such that for every period t we have:

$$A_t + bx \leq p_t(x) \leq \lambda(A_t + bx) \quad \forall x > 0 \quad (6)$$

The following proposition provides conditions for FTL/LTL procurement costs to be 2-linearly sandwiched. The proof is given in Appendix F.

LEMMA 8.

- Any FTL procurement costs $p_t^i(\cdot)$ are 2-linearly sandwiched.
- LTL procurement costs $p_t^i(\cdot)$ are 2-linearly sandwiched if $k^i \leq K_t^i$ for all periods t .

Assumptions on the procurement costs. Although we are mainly interested by FTL/LTL cost structures, our results hold for the following more general procurement cost structure:

- (P1) *Non-decreasing.* For each location i , $p_t^i(\cdot)$ is non decreasing.
- (P2) *Stationarity at retailers.* For each retailer i , $p_t^i(\cdot) = p^i(\cdot)$.
- (P3) *Sub-additivity for retailers.* For each retailer i , $p^i(\cdot)$ is sub-additive.
- (P4) *2-linearly sandwiched at the warehouse.* $p_t^0(\cdot)$ is 2-linearly sandwiched.

This set of assumptions will be referred to as (P). The FTL/LTL procurement costs clearly satisfy (P1) (P2) and (P3), assuming that $K_t^i = K^i$ for each retailer i . Our motivation for requiring (P4), i.e. that the procurement costs at the warehouse to be 2-linearly sandwiched, is algorithmic. Indeed, to the best of our knowledge, no polynomial time algorithm has been proposed for the multi-item uncapacitated lot sizing problem with multiple setup, while an efficient solution to this latter problem is central for our decomposition approach. Using this assumption, we are able to derive easily a 2-approximation for this latter problem, by relaxing the problem at the warehouse, which is sufficient to build a 2-approximation for the OWMR problem, as we detail in the next section.

We assume that the holding costs are either shelf age dependent at all locations or level dependent at all locations. For level (resp. shelf age) dependent holding costs, we assume that the assumptions (LD) in §4.1 (resp. (SA) in §4.2) are satisfied.

Decomposition into subproblems (Splitting phase). We assume from now that the procurement costs satisfy assumptions (P). In particular the procurement cost $p_t^0(\cdot)$ at the warehouse is 2-linearly sandwiched between some affine procurement costs $(A_t + bx)$ and $2(A_t + bx)$.

We consider the following decomposition into subproblems:

(\widehat{S}_i) Retailer i is considered as a single-echelon location facing demand d_t^i , with procurement cost $p^i(\cdot)$. The holding costs are split as in §4.

(\widehat{S}_0) The warehouse is regarded as a single-echelon, multi-item system with a fixed ordering cost A_r in period r . The holding costs are split as in §4.

Observe that (\widehat{S}_0) is identical to the single-echelon problems we have considered for the warehouse in §4 and thus an optimal policy can be found efficiently. Another consequence is that ZIO policies remain dominant for the problem (\widehat{S}_0) .

This is not necessarily true for a problem (\widehat{S}_i) associated with a retailer, since an optimal policy may order a fraction of a subsequent demand to complete a batch. However we can show that the so called *Positive Consumption Ordering* (PCO) policies are dominant. Remind that we assume without loss of generality that demands are satisfied in a FCFS manner in each policy. Hence for a given period t we can define without ambiguity the *consumption in t from (an order in period) $s \leq t$* as the number of units ordered in s that are used to serve demands till period t included.

DEFINITION 2. We say that a policy is a *Positive Consumption Ordering* (PCO) policy if for each order s , the consumption in period s from order s is strictly positive.

That is, a PCO policy orders in a period s only if at least a fraction of the quantity ordered is used to serve the demand of period s . Clearly ZIO policies are a special case of PCO policies that require in addition that all units serving a specific demand are ordered in the same period. Also note that due to the FCFS discipline, the units on hand at the beginning of a period t have all been ordered in the previous ordering period. This simple observation on PCO policies is in fact the main property used in our analysis. It is easy to see that PCO policies are dominant for each system (\widehat{S}_i) : For any $i > 0$, there exists an optimal policy that orders in a period s only if the consumption of the order in period s is not null. Otherwise the order in period s can simply be postponed, maintaining a feasible policy without increasing the total cost (the procurement costs are stationary and the holding costs can only decrease when an order is delayed).

Performance guarantee The split & uncross algorithm can be applied directly to the optimal PCO policies $\hat{\pi}_i^*$ for subproblems (\hat{S}_i) . As in the previous sections, it builds a feasible uncrossing policy π_i^u for each retailer i . The analysis of the previous sections asserts that the holding cost incurred by policy π^u is at most twice the total holding costs of the policies $\hat{\pi}_i^*$ for the problems (\hat{S}_i) . This holds with level dependent or shelf age dependent holding costs (see §4.1 and §4.2).

It remains to show that the procurement cost $\mathcal{K}(\pi^u)$ incurred by policy π^u is still at most twice the sum of the procurement costs incurred by the policies $\hat{\pi}_i^*$. We first prove a technical lemma which states that in the case of linearly sandwiched procurement costs, the simple fact that a single-echelon policy χ' orders in the same periods as an other one χ is sufficient to bound its procurement costs, regardless of the quantities ordered by both policies.

LEMMA 9 (Sandwiched procurement costs). *Consider a single-echelon system with λ -linearly sandwiched procurement costs $p_t(\cdot)$. Let χ and χ' be two feasible policies that order respectively quantities z_t and z'_t in period t , such that $z_t = 0 \Rightarrow z'_t = 0$. That is, policy χ' can order in period t only if policy χ orders. Then the overall procurement cost of χ' is at most λ times the overall procurement cost of χ :*

$$\sum_{t=1}^T p_t(z'_t) \leq \lambda \sum_{t=1}^T p_t(z_t)$$

PROOF. The proof is immediate from the definition of linearly sandwiched procurement costs. Let Z (resp. Z') be the set of periods when policy χ (resp. χ') orders. By definition we have $\sum_{t=1}^T p_t(z_t) \geq \sum_{t \in Z} A_t + b \sum_{t=1}^T z_t$ and $\sum_{t=1}^T p_t(z'_t) \leq \lambda (\sum_{t \in Z'} A_t + b \sum_{t=1}^T z'_t)$. Since policy χ' can order only if policy χ orders, we have $Z' \subseteq Z$. In addition any feasible policy must order exactly the total demands on the horizon, that is $\sum_{t=1}^T z_t = \sum_{t=1}^T z'_t$. The result follows. \square

This allows us to bound the total procurement cost of policy π^u at each location as follows.

LEMMA 10. $\mathcal{K}_i(\pi^u) \leq 2\mathcal{K}_i(\hat{\pi}_i^*)$ for all $i = 0, \dots, N$.

PROOF. For the warehouse, one can note that even though the uncrossing operation may change drastically the quantities ordered between $\hat{\pi}_0^*$ and π_0^u , both policies share (by construction) the same ordering periods. Therefore, the inequality is straightforward from Lemma 9 and property (P4) for $i = 0$.

Now let i be a W -retailer and consider a crossing ordering period s in policy $\hat{\pi}_i^*$. Let $q_s > 0$ and $q_{s+} \geq 0$ be the quantity ordered by $\hat{\pi}_i^*$ in period s and s^+ , respectively. Note that $q_{s+} > 0$ if and only if s^+ is an ordering period of $\hat{\pi}_i^*$. Let $x \geq 0$ be the stock level at the beginning of period s^+ , before receiving q_{s+} . As $\hat{\pi}_i^*$ is PCO (and FCFS) and s is a crossing order, $\hat{\pi}_i^*$ does not place any order between s and s^+ and thus $x \leq q_s$.

In the uncrossed policy π^u , only $q_s - x$ units are ordered in period s and $q_{s+} + x$ units are ordered in period s^+ . Since the procurement cost $p^i(\cdot)$ is sub-additive and non-decreasing, we have

$$p^i(q_s - x) + p^i(q_{s+} + x) \leq p^i(q_s) + p^i(q_{s+}) + p^i(x) \leq 2p^i(q_s) + p^i(q_{s+}).$$

If i is a J -retailer, the uncrossing algorithm can be seen in two steps. First, uncross the crossing orders as for W -retailer by adding an order at s^+ . This step at most doubles the procurement cost (same arguments as for W -retailer). Second, synchronize orders that are not crossing with s^- . This second step can only decrease the ordering cost at retailer i as the procurement costs are sub-additive and stationary. \square

Lemma 10 implies that $\mathcal{K}(\pi^u) \leq 2 \sum_{i=0}^N \mathcal{K}_i(\hat{\pi}_i^*)$. We can state the following theorem:

THEOREM 4. *The split and uncross algorithm with procurement costs that satisfy assumptions (P) has a performance guarantee of two for the one-warehouse multi-retailer problem, either with level dependent holding costs (that satisfy assumptions (LD)) or with shelf age dependent holding costs (that satisfy assumptions (SA)).*

Complexity analysis. The complexity of the split and uncross algorithm depends on the complexity of the uncrossing algorithm and on the complexity to solve the single echelon problems to optimality. The uncrossing algorithm is unchanged. When the input policies are PCO and FCFS, it can be implemented in $O(NT)$, as a demand d_t^i is served by at most two pairs of orders. The warehouse single-echelon problem (\widehat{S}_0) is similar to the previous sections and can be solved in $O(NT)$ for linear holding cost, $O(NT^2)$ for metric holding cost and $O(NT^3)$ for level dependent holding cost.

At retailer i , the complexity to solve the single echelon problem (\widehat{S}_i) depends on the procurement cost function under consideration. Li et al. (2004) propose a $O(T^3 \log T)$ algorithm for the single-echelon single-item problem with nondecreasing concave holding cost and non-decreasing LTL freight cost function f . The complexity can be reduced to $O(T^3)$ for linear holding costs using Monge arrays.

6. Capacity constraint at the W -retailers In this section, we show how to adapt our technique to an OWMR problem in which the W -retailers are subject to capacity constraints. We prove that a slightly modified uncrossing algorithm leads to similar results as the ones obtained in §3 (uncapacitated, linear holding costs and fixed ordering costs). The case of FTL procurement costs and of more general holding cost structures is then briefly discussed at the end of the section.

Limitations on the size of orders is a common restriction met by companies in practice. Such situations may arise from physical constraints on the transportation mode (size of a carrier) or from the maximum number of units a location can handle when it receives an order. Florian and Klein (1971) were among the first to introduce such capacity constraints by considering a single-echelon lot-sizing problem in which the quantity ordered in period t cannot exceed a given capacity $C_t \geq 0$. In the constant capacity case where $C_t = C$ for each period t , they proposed a polynomial-time algorithm in $O(T^4)$ to compute an optimal solution to the problem when holding costs are concave functions. For linear holding costs, the complexity has been improved to $O(T^3)$ by van Hoesel and Wagelmans (1996) (for concave variable procurement costs) and to $O(T^2 \log T)$ by Van Vyve (2007) (for linear variable procurement costs). However, when the capacities are time-dependent, Florian et al. (1980) and Bitran and Yanasse (1982) have shown that the capacitated lot-sizing problem is \mathcal{NP} -hard, even with no holding cost. Very few papers have proposed exact or approximation algorithms for capacitated multi-echelon systems. van Hoesel et al. (2005) present a polynomial time algorithm for serial systems, where the capacity occurs only at the first level. In particular, its time complexity is in $O(T^5)$ in the case of 2 levels and linear holding costs. Recently, Levi et al. (2008a) propose an approximation algorithm for the capacitated multi-item lot-sizing problem. This is a special case of the OWMR problem where the fixed costs are zero at all retailers and the capacity constraints appear only at the warehouse.

Assumptions. The assumptions are the same as in §3. Holding costs are linear at each location and fixed ordering costs are stationary at the retailers. In addition we assume, contrary to the rest of the paper, that the fixed ordering costs at the warehouse are stationary, that is $K_t^0 = K^0$ for all periods. Moreover, in each period t , a W -retailer i cannot order a quantity greater than its capacity C^i . The warehouse orders and the J -retailers orders remains unbounded.

Smooth capacitated problem. Observe that neither ZIO nor PCO policies (see §5 for a definition) are dominant for a capacitated lot-sizing problem, since it may be necessary to order at full capacity during several periods to serve a large subsequent demand. However, we establish below that PCO policies remain dominant for the so-called *smooth* capacitated instances. We say that an instance is *smooth* if the demand at a period never exceeds the capacity. It is easy to see that any OWMR instance can be transformed (in linear time) into a smooth capacitated instance without modifying the set of feasible solutions as noticed by Bitran and Yanasse (1982): Consider a W -retailer i and a period t such that $d_t^i > C^i$. Any feasible policy must order at least $\omega = d_t^i - C^i$




FIGURE 2. Uncrossing of a problematic crossing order s at a W -retailer. Only orders are represented at the warehouse.

units of the demand in some previous periods and physically hold them its stock from period $t-1$ to t . Thus we can consider the equivalent instance with a demand $\bar{d}_t^i = C^i$ and $\bar{d}_{t-1}^i = d_{t-1}^i + \omega$. Notice that for any policy π , the costs incurred on these two instances differ by a positive constant $h^i\omega$ independent of π . Therefore any policy with a performance guarantee α for the modified instance also has a guarantee α for the initial instance. We have the following result (the proof is given in Appendix G):

LEMMA 11. *PCO policies are dominant for the smooth capacitated lot-sizing problem.*

Modified uncrossing algorithm. Similarly to the previous sections, we define for each W -retailer i the (smooth capacitated) subproblem (\hat{S}_i) as a single-echelon problem with capacity constraint C^i , a fixed ordering cost K^i and holding costs $\hat{h}^i = \frac{1}{2}h^i$. The definition of the subproblems for the J -retailers and the warehouse are identical to the decomposition used in §3. Since PCO policies are dominant, one may consider to apply directly the uncrossing algorithm. However, the resulting policy may be infeasible. To see why, consider a PCO policy $\hat{\pi}_i$ for the W -retailer i . Let $q_t \leq C^i$ be the quantity ordered in period t by policy $\hat{\pi}_i$. On one hand, the uncrossed policy built by the algorithm essentially orders the same quantities at the same periods, if they are not crossing, and thus satisfies the capacity constraint as $\hat{\pi}_i$ is feasible for (\hat{S}_i) . On the other hand, if s is a crossing ordering period in policy $\hat{\pi}_i$, the algorithm diminishes the quantity q_s ordered in period s by its stock level, say x , at the beginning of period $s+$, and orders these units in period $s+$. Since π_i is PCO, the quantity x cannot exceed $q_s - 1$ units, and thus is lower than capacity C^i . The problem arises if $s+$ is already an ordering period of policy $\hat{\pi}_i$. We say that such crossing order is *problematic*. In this situation the uncrossing policy π^u must order at period $s+$ the quantity x plus the quantity q_{s+} , which may be as large as $2C^i - 1$.

To overcome this difficulty, we adapt the uncrossing algorithm by adding some ordering periods also at the warehouse. This motivates the stationarity assumption made on K^0 . More precisely, if there exists a problematic crossing order s at a W -retailer, then we add the ordering period $r' = s+ - 1$ at the warehouse. Let R' be the resulting set of ordering periods at the warehouse. Clearly we have at most doubled the number of ordering periods, that is $|R'| \leq 2|R|$. Notice that a problematic crossing order s relatively to the set R of warehouse orders becomes a non problematic crossing order relatively to the set R' if $s < r'$. In this case, we uncross the ordering period s using the additional warehouse order r' , by ordering the number x of units in stock in policy $\hat{\pi}_i$ at the beginning of period r' , both at the warehouse and at the retailer i at period r' . If it happens that $s = r'$, we simply order the same quantities at the retailer i in periods s and $s+$ as in policy $\hat{\pi}_i$, since in this case s is not crossing with respect to the set R' . In both case the quantities ordered do not exceed the capacity of the retailer. Figure 2 illustrates the adaptation of the uncrossing algorithm on a problematic crossing order.

It is easy to check that the resulting policy is feasible, that is, satisfies the capacity constraints for all W -retailers. Observe that the PCO assumption is necessary: Basically it ensures that the

inventory level at the beginning of a period is always lower than the capacity constraint, since for each order (which is limited to C^i units), at least one unit is consumed in its period under a FCFS discipline. Using Lemma 11 and the analysis of previous sections, we obtain in corollary that the OWMR problem can still be approximated with a performance guarantee of 2 by the split & uncross algorithm. The time complexity is dominated by the computation of the optimal PCO policies for the capacitated lot-sizing problems at the W -retailers.

THEOREM 5. *The (modified) split and uncross algorithm is a 2-approximation for the OWMR problem with a constant capacity constraint at W -retailers, linear holding costs and stationary fixed ordering costs. Its time complexity is $O(NT^2 \log T)$.*

Though Theorem 5 extends in a simple way Theorem 1 using the properties of the PCO policies, the result does not hold if a capacity constraint appears at a J -retailer or at the warehouse. The reason is that the uncrossing algorithm synchronizes the orders of J -retailers with the warehouse, which may result into a large order quantity violating the capacity constraint at the retailer and/or the warehouse. Even if set I_J is empty, uncrossing an order at a W -retailer generates a new order at the next warehouse ordering period. This uncrossing order may violate the warehouse capacity, as many W -retailers may need to uncross with the same warehouse order. As Levi et al. (2008a) mention in their paper, other new ideas are certainly needed to be able to take into account fixed item ordering cost with a capacity at the warehouse.

Extension to FTL procurement costs. We turn our attention to the more general case of FTL procurement costs introduced in §5. We only need two ingredients to extend the uncrossing algorithm while preserving a performance guarantee of 2: The single echelon lot-sizing problems (S_i) at the W -retailers should be polynomially solvable, and the PCO policies should be dominant for the smooth instances to ensure the feasibility of the uncrossed policy. To ensure the latter point, we restrict ourself to the case where the capacity C^i is a multiple of the batch size B^i at each retailer. We have the following result (see Appendix G for the proof):

LEMMA 12. *PCO policies are dominant for smooth capacitated lot-sizing problem with FTL cost structure if the capacity is a multiple of the batch size.*

Though restrictive, we think that this condition is of practical relevance. Notice that if the capacity is not a multiple of the batch size, PCO policies are not dominant anymore: Consider the instance with $C = 3$, $B = 2$ and demands $d = (0, 3, 3)$ to satisfy. All the costs but the fixed cost per batch are null. We set $k = 1$. A PCO policy must orders 4 batches, while the optimal policy orders only 3 (full) batches. The single echelon lot-sizing problem with FTL cost structure and linear holding costs remains polynomially solvable in presence of a stationary capacity, see Akbalik and Rapine (2012). In the case where the capacity is a multiple of the batch size, the authors proposed a dynamic programming algorithm of time complexity $O(T^4)$. It results that the OWMR problem with FTL cost structure and capacitated W -retailers can be approximated with a guarantee of 2 by the uncrossing algorithm in time $O(NT^4)$ for linear holding costs at the W -retailers.

7. Perspectives and recent developments The split & uncross procedure presented in this paper has proven to be highly adaptable to various types of costs for the OWMR problem. In fact, it is fairly easy to see that most of the cost structures presented in this paper can be combined without deteriorating the worst-case bound. We also believe that the techniques introduced in this paper can be adapted to deal with even more complex inventory control settings. Indeed it was already proven recently that those ideas could be extended to provide efficient approximation algorithms in the context of the OWMR problem with backlogs and lost sales, see Gayon et al. (2016), and for general distribution problems over several layers, see Stauffer (2016). We are convinced that other applications will follow.

Appendix A: Proof of the dominance of synchronized policies for J -retailers, § 4.1

Consider a J -retailer i satisfying property (LD3) defined in §4.1: For any period t and any quantities x , y and q , the inequality $h_t^i(y+q) - h_t^i(y) \leq h_t^0(x+q) - h_t^0(x)$ holds. We prove that it is dominant for i to synchronize its orders with the warehouse. Assume that a policy π does not satisfy this property and let s be the first period when retailer i orders while the warehouse does not. Let r be the previous ordering period of the warehouse and q be the quantity ordered in period s by retailer i . We denote by x_t and y_t the stock on hand at the warehouse and at the retailer, respectively. We now consider the alternative policy π' that shifts the order of period s to period r , ordering the quantity q . The stock levels of policy π' on interval $[r, s)$ are $x'_t = x_t - q$ and $y'_t = y_t + q$. If we compute the holding costs paid by π and π' on $[r, s)$, we get $\sum_{t=r}^{s-1} (h_t^0(x_t) + h_t^i(y_t))$ and $\sum_{t=r}^{s-1} (h_t^0(x_t - q) + h_t^i(y_t + q))$, respectively. The inequality on marginal costs ensures that the cost incurred by π' is less than or equal to the cost of π . Before period r and after period s both policies are identical by construction. Using an interchange argument, we conclude that it is dominant to synchronize all the orders of retailer i with the warehouse.

Appendix B: Proof of Lemma 4, §4.1 We consider that we are given a single-echelon feasible policy π_i for each problem (\widehat{S}_i) . The proof of Lemma 4 for the uncrossing algorithm is similar to the one of Lemma 3. We only need to prove that (5) still holds with level dependent holding costs, that is the total holding cost incurred by π^u is at most the sum of the holding costs incurred by the policies π_i 's. We prove here a slightly more general result than Lemma 4: we require only policy π_0 to be ZIO; for the retailers we assume that policy π_i is PCO, see Definition 2 §5. In particular a ZIO policy is a PCO policy.

We first introduce some notation. For $i = 0, \dots, N$, we denote by x_{it}^u the inventory level at location i in the final policy π_i^u in period t . Similarly for $i = 1, \dots, N$ let x_{it} be the inventory level of policy π_i in period t . We use a different notation for policy π_0 at the warehouse, in order to distinguish between item 0 and items $i \in I_J$. Recall that item 0 in (\widehat{S}_0) stands for the aggregated demands of all the W -retailers. Specifically, for $i \in \{0\} \cup I_J$, we denote by y_{it} the inventory level of policy π_0 for item i in period t .

In the following we consider a specific unit of demand d_t^i and we let r and s be the ordering periods of this unit in the single-echelon policies π_0 and π_i , respectively. Recall that the ordering periods $[r^{it}, s^{it}]$ in π^u are defined by (3). Assume that this unit of demand is hold in stock at a location j at some period τ in policy π^u . We distinguish in the following between the case where i is a W -retailer and i is a J -retailer.

Firstly, consider that i is a W -retailer. The definition of pair $[r^{it}, s^{it}]$ implies in particular that $s^{it} \geq s$. That is, each unit at the retailer is ordered later (or at the same period) in π^u than in π_i . As a consequence we have $x_{i\tau}^u \leq x_{i\tau}$ for all periods τ . Therefore from (LD1) we have for all $t = 1, \dots, T$

$$\mathcal{H}_i(\pi^u) \leq \mathcal{H}_i(\pi_i). \quad (7)$$

At the warehouse, notice that only units supplying W -retailers are kept in stock in policy π^u , since all the pairs of orders for demands of J -retailers are synchronized by construction. Consider a period τ such that the unit of demand d_i^i is kept at the warehouse in policy π^u . Notice that necessarily we have $r^{it} \leq \tau < s^{it}$. For this situation to happen for a W -retailer, we must have $[r^{it}, s^{it}] = [r, s]$, see (3). It results in particular that $r \leq \tau < t$, that is, this unit is also in stock at period τ in policy π_0 . We can conclude that $x_{0\tau}^u \leq y_{0\tau}$ for all periods τ , and thus

$$\mathcal{H}_0(\pi^u) \leq \sum_{\tau=1}^T h_0(y_{0\tau}). \quad (8)$$

Secondly, consider that i is a J -retailer. We prove that the holding cost incurred by π_i^u is lower than or equal to the the holding cost incurred by π_i plus the holding cost incurred by item i in policy π_0 . We count the holding cost of π_i^u for each ordering interval $[u, u^+]$ of the warehouse, from the beginning of period u till the beginning of period u^+ . Observe that these intervals form a partition of the time horizon, therefore we only need to prove that on each interval the holding costs paid by π^u is at most the holding costs paid by π_i or by π_0 for item i . Let u and u^+ be 2 consecutive ordering periods at the warehouse in policy π^u . Consider that the unit of demand (i, t) is in stock in π_i^u at a period τ , $u \leq \tau < u^+$. By construction, see Definition (3) of ordering pairs of π^u , we have $r^{it} = s^{it} \geq \min\{r, s\}$. Notice that s^{it} is synchronised with the ordering periods of the warehouse, thus for the unit to be in stock at period τ , we necessarily have $s^{it} \leq u$. Since (3) implies in particular that $s^{it} \geq s^-$, it involves that $s < u^+$. We now distinguish between 2 cases:

- **case 1:** Policy π_i does not order in periods $u + 1, \dots, u^+ - 1$. Hence either $s \leq u$ or $s \geq u^+$. Since we have noticed that necessarily $s < u^+$, we are in the situation where $s \leq u$. In other words, the unit is also in stock at period τ in policy π_i . This is true for any units and any instant τ , hence we can conclude that:

$$x_{i\tau}^u \leq x_{i\tau} \quad \forall \tau, \quad u \leq \tau < u^+$$

- **case 2:** There exists a period v , $u < v < u^+$ where policy π_i orders. Consider again one unit of demand (i, t) in stock in π^u at the retailer at a period τ , $u \leq \tau < u^+$. We first prove that necessarily $t < u^+$, that is, the unit serves a demand before the next ordering period of the warehouse. Assume for the sake of contradiction that $t \geq u^+$. On one hand, since policy π_0 is ZIO, we must have $r \geq u^+$. On the other hand, we have noticed that $s < u^+$. Thus we are in the situation where $s < r$, and

s^{it} is defined as period s^+ . Since $s^{it} = s^+ \leq u$, we can conclude that $s \leq u$. This contradicts the assumption that π_i is PCO, since all units ordered before period u should be consumed by the end of period $v < u^+ \leq t$. As a consequence, for the unit to be in stock at period τ in policy π_i^u , we must have $t < u^+$. Clearly, this unit should be ordered no later than period t in any feasible policy. It implies that in policy π_0 , this unit is ordered no later than period t , that is, we have $r \leq u$. In other words, the unit is also in stock at period τ in policy π_0 . This is true for any units and any instant τ , hence we can conclude that

$$x_{i\tau}^u \leq y_{i\tau} \quad \forall \tau, u \leq \tau < u^+$$

Thus, depending on the warehouse intervals $[u, u^+)$ considered, the stock level $x_{i\tau}^u$ is lower or equal either to the stock level in policy π_i or to the stock of item i in policy π_0 . Hence we obtain that:

$$\mathcal{H}_i(\pi^u) \leq \mathcal{H}_i(\pi_i) + \sum_{\tau=1}^T h_{\tau}^i(y_{i\tau}). \quad (9)$$

Summing inequalities (7), (8) and (9), we obtain that:

$$\begin{aligned} \mathcal{H}(\pi^u) &= \mathcal{H}_0(\pi^u) + \sum_{i \in I_W} \mathcal{H}_i(\pi^u) + \sum_{i \in I_J} \mathcal{H}_i(\pi^u) \\ &\leq \sum_{\tau=1}^T h_{\tau}^0(y_{0\tau}) + \sum_{i \in I_W} \mathcal{H}_i(\pi_i) + \sum_{i \in I_J} \left(\mathcal{H}_i(\pi_i) + \sum_{\tau=1}^T h_{\tau}^i(y_{i\tau}) \right) \\ &\leq \sum_{i=0}^N \mathcal{H}_i(\pi_i). \end{aligned}$$

Hence (5) still holds and Lemma 4 follows.

Appendix C: Proof of Lemma 5, §4.1 Consider an optimal policy $\pi^{\text{OPT}} = (\pi_0^{\text{OPT}}, \pi_1^{\text{OPT}}, \dots, \pi_N^{\text{OPT}})$ of cost \mathcal{C}^* for the OWMR instance. To prove the result, we exhibit feasible policies $\tilde{\pi}_i$ for each system (\hat{S}_i) such that $\hat{\mathcal{C}}_0(\tilde{\pi}_0) + \sum_{i=1}^N \hat{\mathcal{C}}_i(\tilde{\pi}_i)$ is less than or equal to \mathcal{C}^* . Lemma 5 then follows from the optimality of each policy $\hat{\pi}_i^*$ for system (\hat{S}_i) .

According to (LD3), one can choose policy π^{OPT} such that no inventory is held at the warehouse for J -retailers. We denote by x_{it}^{OPT} the inventory level of π^{OPT} in period t in location i and we let $x_t^e \equiv x_{0t}^{\text{OPT}} + \sum_{i \in I_W} x_{it}^{\text{OPT}}$ be the aggregated stock level of the warehouse and W -retailers altogether, which represents the echelon stock of the system when set I_J is empty. We now specify a feasible policy $\tilde{\pi}_i$ for each system (\hat{S}_i) :

- For each $i > 0$, policy $\tilde{\pi}_i$ is identical to policy π^{OPT} restricted to retailer i : It orders the same quantities in the same periods.

• For the multi-item problem (\widehat{S}_0) , policy $\tilde{\pi}_0$ orders when policy π^{OPT} places its warehouse orders. In addition, the quantity ordered in a period r corresponds to the cumulative demand of retailers in I_W until the next ordering period (excluded), plus all the quantities ordered in period r by retailers in I_J .

We first prove that the sum of the costs of policies $\tilde{\pi}_i$ applied to systems (\widehat{S}_i) is lower than the optimal cost \mathcal{C}^* for the original problem.

It is straightforward to see that the total ordering cost paid by policies $\tilde{\pi}_i$ is exactly the total ordering cost paid by policy π^{OPT} , since they order at the same periods. Therefore we focus on the holding costs incurred by policies $\tilde{\pi}_i$. We denote by \tilde{x}_{it} the inventory level at retailer i in period t for policy $\tilde{\pi}_i$, $i > 0$, and by \tilde{y}_{it} the inventory level of item i for policy $\tilde{\pi}_0$ at the warehouse. Recall that item 0 represents the aggregation of the W -retailers. For this product, we have for each period t :

$$\begin{aligned} \hat{h}_t^0(\tilde{y}_{0t}) &= \frac{1}{2}h_t^0(\tilde{y}_{0t}) \leq \frac{1}{2}h_t^0(x_t^e) \\ &\leq \frac{1}{2}h_t^0(x_{0t}^{\text{OPT}}) + \frac{1}{2} \sum_{i \in I_W} h_t^0(x_{it}^{\text{OPT}}) \\ &\leq \frac{1}{2}h_t^0(x_{0t}^{\text{OPT}}) + \frac{1}{2} \sum_{i \in I_W} h_t^i(x_{it}^{\text{OPT}}). \end{aligned}$$

The first inequality comes from (LD1), the non-decreasing property of $h_t^0(\cdot)$. Indeed in any feasible policy the stock level at the warehouse plus the stock level at a subset of retailers has to be sufficient to satisfy the demands of these retailers until the next ordering period of the warehouse (which is precisely \tilde{y}_{0t} if we focus on product 0), which implies that $\tilde{y}_{0t} \leq x_t^e$. The second inequality directly comes from (LD2), the sub-additivity of h_t^0 . The definition of set I_W implies the last inequality.

Now if we restrict our attention to items $i \in I_J$ of system (\widehat{S}_0) , we have by construction $\tilde{y}_{it} = x_{it}^{\text{OPT}}$ in each period. Hence the holding cost $\hat{h}_t^i(\tilde{y}_{it})$ paid for such a product $i \in I_J$ in period t is exactly $\frac{1}{2}h_t^i(x_{it}^{\text{OPT}})$. Adding the holding costs paid for all products altogether in system (\widehat{S}_0) , we obtain for each period t that:

$$\hat{h}_t^0(\tilde{y}_{0t}) + \sum_{i \in I_J} \hat{h}_t^i(\tilde{y}_{it}) \leq \frac{1}{2}h_t^0(x_{0t}^{\text{OPT}}) + \frac{1}{2} \sum_{i=1}^N h_t^i(x_{it}^{\text{OPT}}).$$

Finally the stock level in each system (\widehat{S}_i) , is identical to the stock level at retailer i in policy π^{OPT} . Hence we also have for each period t :

$$\hat{h}_t^i(\tilde{x}_{it}) = \frac{1}{2}h_t^i(x_{it}^{\text{OPT}}) \quad \forall i = 1, \dots, N.$$

As a consequence for each period t , the total holding cost paid by all the policies $\tilde{\pi}_i$ in (\widehat{S}_i) is at most $\frac{1}{2}h_t^0(x_{0t}^{\text{OPT}}) + \sum_{i=1}^N h_t^i(x_{it}^{\text{OPT}})$, which is a lower bound on the holding cost paid by policy π^{OPT} , since h_t^0 takes only positive values. Summing the holding costs paid by all the policies $\tilde{\pi}_i$, we obtain for each period t that:

$$\begin{aligned} \sum_{i=0}^N \widehat{\mathcal{H}}(\widehat{\pi}_i^*) &= \widehat{h}_t^0(\widehat{y}_{0t}) + \sum_{i \in I_J} \widehat{h}_t^i(\widehat{y}_{it}) + \sum_{i=1}^N \widehat{h}_t^i(\widehat{x}_{it}) \\ &\leq \frac{1}{2}h_t^0(x_{0t}^{\text{OPT}}) + \sum_{i=1}^N h_t^i(x_{it}^{\text{OPT}}) \\ &\leq \sum_{i=0}^N h_t^i(x_{it}^{\text{OPT}}) \\ &= \mathcal{H}(\pi^{\text{OPT}}). \end{aligned}$$

The lemma follows.

Appendix D: Proof of Lemma 6, §4.2 The proof is similar to the one of Lemma 3. We only need to prove that (5) still holds, that is, $\mathcal{H}(\pi^u) \leq \sum_{i=0}^N \mathcal{H}_i(\pi_i)$. Consider a specific unit of demand d_t^i and let r and s be the ordering periods of this unit in the single-echelon policies π_0 and π_i , respectively. The ordering periods in π_u are again defined by (3).

Case 1: i is a W -retailer. If $r > s$, the final holding cost incurred by policy π^u is $h_{s^+s^+}^{it}$ and we have from properties (SA1), (SA2) and (SA3) :

$$h_{s^+s^+}^{it} \leq h_{ss}^{it} \leq h_{rt}^{it} + h_{ss}^{it}. \quad (10)$$

If $r \leq s$, π^u incurs a holding cost of h_{rs}^{it} and property (SA4) ensures that

$$h_{rs}^{it} \leq h_{rt}^{it} + h_{ss}^{it}. \quad (11)$$

Case 2: i is a J -retailer. If $r > s$, π^u incurs again a holding cost $h_{s^+s^+}^{it}$. Properties (SA1) and (SA3) implies

$$h_{s^+s^+}^{it} \leq h_{ss}^{it} \leq h_{rr}^{it} + h_{ss}^{it}. \quad (12)$$

If $r \leq s$, the holding cost incurred by π^u is equal to $h_{s^-s^-}^{it}$. Using property (SA3), we obtain

$$h_{s^-s^-}^{it} \leq h_{rr}^{it} \leq h_{rr}^{it} + h_{ss}^{it}. \quad (13)$$

In all cases, the holding cost incurred in the final policy to serve demand d_t^i is lower than the sum of the holding costs incurred in the single-echelon policies. Hence relation (5) still holds.

Appendix E: Proof of Lemma 7, §4.2 Similarly to the proof of Lemma 2, consider an optimal policy $\pi^{\text{OPT}} = (\pi_0^{\text{OPT}}, \pi_1^{\text{OPT}}, \dots, \pi_N^{\text{OPT}})$ of cost \mathcal{C}^* for the OWMR instance, for which (w.l.o.g) no inventory is held at the warehouse for J -retailers. Let $\tilde{\pi}_i$ be the single-echelon policy for system (\widehat{S}_i) that orders the same quantities of the same items as π_i^{OPT} in the same periods. Clearly we have $\mathcal{K}_i(\tilde{\pi}_i) = \mathcal{K}_i(\pi_i^{\text{OPT}})$ for all $i = 0, \dots, N$.

Now, consider a specific demand d_t^i and assume π^{OPT} uses the pair of orders $[r, s]$ to serve this demand, incurring a holding cost H_{rs}^{it} . Then by construction $\tilde{\pi}_0$ and $\tilde{\pi}_i$ order d_t^i in periods r and s , respectively. If $i \in I_J$, we have $r = s$ and the holding costs incurred by $\tilde{\pi}_0$ and $\tilde{\pi}_i$ to serve this demand in problems (\widehat{S}_0) and (\widehat{S}_i) are both equal to $\frac{1}{2}H_{rr}^{it} = \frac{1}{2}H_{ss}^{it}$. If $i \in I_W$, then the holding costs incurred by $\tilde{\pi}_0$ and $\tilde{\pi}_i$ to serve this demand in problems (\widehat{S}_0) and (\widehat{S}_i) are equal to $\frac{1}{2}H_{rr}^{it}$ and $\frac{1}{2}H_{ss}^{it}$, respectively. Recall that from the monotonicity properties (SA2) and (SA3), we have that H_{rs}^{it} is greater than both of these quantities and thus $H_{rs}^{it} \geq \frac{1}{2}H_{rr}^{it} + \frac{1}{2}H_{ss}^{it}$ for all i, t . Summing over all demands, we obtain that $\mathcal{H}(\pi^{\text{OPT}}) \geq \widehat{\mathcal{H}}_0(\tilde{\pi}_0) + \sum_{i=1}^N \widehat{\mathcal{H}}_i(\tilde{\pi}_i)$. Therefore we have:

$$\begin{aligned} \mathcal{C}^* &= \mathcal{K}(\pi^{\text{OPT}}) + \mathcal{H}(\pi^{\text{OPT}}) \\ &\geq \sum_{i=0}^N \mathcal{K}_i(\tilde{\pi}_i) + \widehat{\mathcal{H}}_0(\tilde{\pi}_0) + \sum_{i=1}^N \widehat{\mathcal{H}}_i(\tilde{\pi}_i) \\ &= \widehat{\mathcal{C}}_0(\tilde{\pi}_0) + \sum_{i=1}^N \widehat{\mathcal{C}}_i(\tilde{\pi}_i) \end{aligned}$$

and the lower bound follows from the optimality of $\hat{\pi}_0^*, \hat{\pi}_1^*, \dots, \hat{\pi}_N^*$ for problems $(\widehat{S}_0), (\widehat{S}_1), \dots, (\widehat{S}_N)$.

Appendix F: Proof of Lemma 8, §5 We prove here that any FTL procurement cost and any LTL procurement costs satisfying conditions of Lemma 8 is 2-linearly sandwiched. Recall that we restrict ourself to a stationary batch size B and a stationary fixed cost per batch k .

First consider a FTL procurement cost: $FLL_t(x) = K_t + \lceil x/B \rceil k$. We define $A_t = K_t + k/2$ and $b = \frac{k}{2B}$. On one hand, for any quantity $x > 0$, we clearly have $FLL_t(x) \geq K_t + \frac{k}{B}x$ and $FLL_t(x) \geq K_t + k$. It results that $FLL_t(x) \geq A_t + bx$. On the other hand $FLL_t(x) \leq K_t + (\frac{x}{B} + 1)k$ implies that $FLL_t(x) \leq 2A_t + 2bx$. Thus $FLL_t(x)$ is linearly sandwiched by $A_t + bx$.

Now consider a LTL procurement cost: $LTL_t(x) = K_t + \lceil x/B \rceil k + f(x - \lceil x/B \rceil)$, with $k \leq K_t$. Recall that we assume that freight cost function is non-decreasing on $[0, B]$ with $f(0) = 0$ and $f(B) = k$. We define two constants λ and μ as follows:

$$\begin{aligned} \lambda &= \inf \left\{ f(u) - \frac{k}{2B}u \mid u \in [0, B] \right\} \\ \mu &= \sup \left\{ f(u) - \frac{k}{B}u \mid u \in [0, B] \right\} \end{aligned}$$

Consider a positive quantity x . For conciseness we denote by $u(x)$ the quantity $x - \lfloor x/B \rfloor$ ($u(x) \in [0, B]$). We thus have $LTL_t(x) = K_t + \frac{k}{B}(x - u(x)) + f(u(x))$. On one hand, by definition of λ , we can write that $LTL_t(x) \geq K_t + \frac{k}{B}(x - u(x)) + \lambda + \frac{k}{2B}u(x)$. It follows that $LTL_t(x) \geq K_t + \lambda + \frac{k}{2B}x - \frac{k}{2B}u(x)$. Since $u(x) \leq x$, we have $LTL_t(x) \geq K_t + \lambda + \frac{k}{2B}x$. On the other hand, by definition of μ , we have $LTL_t(x) \leq K_t + \frac{k}{B}x + \mu$. We define $A_t = K_t + \lambda$ and $b = \frac{k}{2B}$. From what precedes, procurement costs $LTL_t(x)$ are 2-linearly sandwiched if the inequality $\mu \leq K + 2\lambda$ holds. Since we assume that $k \leq K_t$ for each period, it is sufficient to establish that $\mu - k \leq 2\lambda$.

To show that $\mu - k \leq 2\lambda$, let $\varepsilon > 0$ be a fixed value. By definition of μ , there necessary exists at least one value $y \in [0, B]$ such that $f(y) - \frac{k}{B}y \geq \mu - \varepsilon$. For any $u \leq y$, using the fact that $0 \leq f(u) \leq k$, we have:

$$f(u) - \frac{k}{2B}u \geq 0 - \frac{k}{2B}y \geq \frac{1}{2}(\mu - \varepsilon - f(y)) \geq \frac{1}{2}(\mu - \varepsilon - k).$$

For any $u \geq y$ we have, since f is non-decreasing:

$$f(u) - \frac{k}{2B}u \geq f(y) - \frac{k}{2B}u \geq (\mu - \varepsilon + \frac{k}{B}y) - \frac{k}{2B}u \geq \mu - \varepsilon - \frac{1}{2}k.$$

Observe that $\mu \geq 0$ since $f(0) = 0$. It results that for all $u \in [0, B]$ we have $f(u) - \frac{k}{2B}u \geq \frac{1}{2}(\mu - k) - \varepsilon$. By definition it implies that $2\lambda \geq \mu - k - 2\varepsilon$. Since the inequality holds for any $\varepsilon > 0$, we obtain that $2\lambda \geq \mu - k$, which concludes the proof.

Appendix G: Proof of Lemmas 11 and 12, §6 Consider an instance of a smooth capacitated lot-sizing problem, with planning horizon T and capacity C . We assume the problem has a feasible solution: That is, $\sum_{t=1}^s d_t \leq C \cdot s$ for all $s = 1, \dots, T$. We prove that PCO policies are dominant under a FTL cost structure if the capacity is a multiple of the batch size B . We denote by k the fixed cost of a batch. Notice that Lemma 11, without FTL procurement cost, is a particular case by taking $B = C$ and $k = 0$.

Consider a feasible policy π for the problem that does not satisfy the PCO property and let q_t be the quantity ordered by π in period t . We focus on u , the first ordering period of π that is not PCO. It involves that π orders a quantity $q_u > 0$ in period u and no unit of q_u is consumed at u . We can also assume that the total quantity ordered by π matches the total demand, therefore as u is not PCO we have:

$$\sum_{t=u}^T q_t \leq \sum_{t=u+1}^T d_t. \quad (14)$$

We use an interchange argument to transform π into a feasible policy π' of lower cost and which is PCO on periods $\{1, \dots, u\}$. By repeating the transformation on the planning horizon, we can then establish the dominance. The idea is to push forward the q_u units while respecting the capacity constraint. For each subsequent period, as many pushed units as possible are ordered in the period,

and the remaining units are pushed again to the next period till all the q_u units have been ordered. Specifically, let δ_t be the quantity from order u pushed in period t . We have:

$$\delta_t = \begin{cases} q_u & \text{if } t = u \\ (\delta_{t-1} + q_t - C)^+ & \text{otherwise} \end{cases}, \quad (15)$$

where x^+ is defined as $\max\{0, x\}$. Note that since π is feasible, $q_t \leq C$ for all $t = 1, \dots, T$ and therefore δ_t is nonincreasing from (15). It implies in particular that $\delta_t \leq q_u \leq C$.

We have to prove that (i) the resulting policy π' is feasible, and (ii) the cost incurred is at most the cost of π . Policy π' is clearly feasible, since the remaining capacity from period $u + 1$ till the end of the planning horizon is sufficient to satisfy all subsequent demands, that is $\sum_{t=u+1}^T C \geq \sum_{t=u+1}^T d_t$. This is a direct consequence of Inequality 14.

To compare the cost of π and π' , let v be the first period such that $\delta_v = 0$. Notice that policy π' orders at full capacity in periods $u + 1$ till $v - 1$. The holding costs incurred in π' are not greater than the holding costs incurred by π , since all the pushed units are ordered later in time. Thus, we only need to compare the number of ordering periods (where a fixed cost K is incurred) and the number of fractional batches over the time interval $[u, \dots, v]$. Clearly, policy π orders in all these periods, except possibly in period v , since $q_t = 0$ implies that $\delta_t = 0$. Since policy π' does not order in period u , its number of ordering periods is at most the ones of π . Now let us turn attention to the number of fractional batches. Since all the periods $u + 1$ till $v - 1$ order at full capacity, at most one fractional batch may appear in policy π' , possibly in period v . Thus we can restrict to the case where π orders only full batches over $[u, \dots, v]$. It implies that q_t is a multiple of B for all these periods. Using Equation 15, we obtain immediately that d_{v-1} is a multiple of v . As a consequence the quantity $q'_v = q_v + \delta_{v-1}$ ordered in π' at period v is also a multiple of B . We can conclude that the number of fractional batches used in policy π' is always lower or equal to the number of fractional batches used in policy π . The proof follows.

References

- Aggarwal A, Park JK (1993) Improved algorithm for economic lot-size problems. *Oper. Res.* 41(3):549–571.
- Akbalik A, Rapine C (2012) Polynomial time algorithms for the constant capacitated single-item lot sizing problem with stepwise production cost. *Oper. Res. Letters* 40(5):390–397.
- Aksoy Y, Erenguc SS (1998) Multi-item inventory models with co-ordinated replenishments: A survey. *International J. Oper. Production Management* 8:63–73.
- Arkin E, Joneja D, Roundy R (1989) Computational complexity of uncapacitated multi-echelon production planning problems. *Oper. Res. Letters* 8:61–66.
- Bellman R (1958) On a routing problem. *Quarterly of Applied Mathematics* 16:87–90.
- Bienkowski M, Byrka J, Chrobak M, Dobbs N, Nowicki T, Sviridenko M, Świrszcz G, Young NE (2013) Approximation algorithms for the joint replenishment problem with deadlines. *Automata, Languages, and Programming*, 135–147 (Springer).
- Bienkowski M, Byrka J, Chrobak M, Jež L, Nogneng D, Sgall J (2014) Better approximation bounds for the joint replenishment problem. *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms* 42–54.

- Bitran G, Yanasse H (1982) Computational complexity of the capacitated lot size problem. *Management Sci.* 28(10):1174–1185.
- Brahimi N, Dauzère-Peres S, Najid NM, Nordli A (2006) Single item lot sizing problems. *Eur. J. Oper. Res.* 168:1–16.
- Chan LMA, Muriel A, Shen ZJM, Simchi-Levi D, Teo CP (2000) Effective zero-inventory-ordering policies for the single-warehouse multiretailer problem with piecewise linear cost structures. *Management Sci.* 48(11):1446–1460.
- Federgruen A, Tzur M (1991) A simple forward algorithm to solve the general dynamic lot-sizing models with n periods in $o(n \log n)$ or $o(n)$ time. *Management Sci.* 37(8):909–925.
- Federgruen A, Tzur M (1999) Time-partitioning heuristics: Application to one-warehouse, multi-item, multi-retailer lot-sizing problems. *Naval Research Logistics* 46:463–486.
- Federgruen A, Wang M (2013) Inventory models with shelf age and delay dependent inventory costs. *MSOM Conference 2013, Fontainebleau, France*.
- Feige U (1998) A threshold of $\log(n)$ for approximating set-cover. *J. of the ACM* 45:634–652.
- Florian M, Klein M (1971) Deterministic production planning with concave costs and capacity constraints. *Management Sci.* 18:12–20.
- Florian M, Lenstra JK, Rinnooy Kan AHG (1980) Deterministic production planning: Algorithms and complexity. *Management Sci.* 26(7):669–679.
- Gayon JP, Massonnet G, Rapine C, Stauffer G (2016) Constant approximation algorithms for the one warehouse multiple retailers problem with backlog or lost-sales. *European Journal of Operational Research* 250(1):155 – 163, ISSN 0377-2217, URL <http://dx.doi.org/http://dx.doi.org/10.1016/j.ejor.2015.10.054>.
- Khouja M, Goyal S (2008) A review of the joint replenishment problem literature: 1989-2005. *Eur. J. Oper. Res.* 186:1–16.
- Levi R, Lodi A, Sviridenko M (2008a) Approximation algorithms for the capacitated multi-item lot-sizing problem via flow-cover inequalities. *Math. Oper. Res.* 33:461–474.
- Levi R, Roundy R, Shmoys D (2006) Primal-dual algorithms for deterministic inventory problems. *Math. Oper. Res.* 31(2):267–284.
- Levi R, Roundy R, Shmoys D, Sviridenko M (2008b) A constant approximation algorithm for the one-warehouse multiretailer problem. *Management Sci.* 54(8):763–776.
- Li CL, Hsu V, Xiao WQ (2004) Dynamic lot sizing with batch ordering and truckload discounts. *Operations Research* 52(4):639–654.
- Muckstadt J, Roundy R (1993) *Analysis of multistage production systems*, volume 4, Logistics of Production and Inventory of *Handbooks in OR & MS*, chapter 2, 59–131 (North Holland Amsterdam).
- Nonner T, Souza A (2009) A 5/3-approximation algorithm for joint replenishment with deadlines. *Proc. 3rd Int. Conf. Combinatorial Optimization and Applications, Huangshan, China*, 24–35.
- Nonner T, Sviridenko M (2013) An efficient polynomial-time approximation scheme for the joint replenishment problem. *Proc. 16th Conf. on Integer Programming and Combinatorial Optimization (IPCO'13), Valparaíso, Chile*, 314–323.
- Pochet Y, Wolsey L (2006) *Production planning by mixed integer programming* (Springer).
- Roundy R (1985) A 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management Sci.* 31:1416–1430.
- Schwarz LB (1973) A simple continuous review deterministic one-warehouse n-retailer inventory problem. *Management Sci.* 19:555–566.
- Shen ZJ, Shu J, Simchi-Levi D, Teo CP, Zhang J (2009) Approximation algorithms for general one-warehouse multi-retailer systems. *Naval Research Logistics* 56:642–658.
- Stauffer G (2012) Using the economical order quantity formula for inventory control in one-warehouse multi-retailer systems. *Naval Research Logistics* 59(3-4):285–297.

- Stauffer G (2016) Constant factor approximation algorithms for deterministic k -echelon inventory control in pure distribution networks. Technical Report GSCOP-STA-16-1, Grenoble Institute of Technology (Grenoble INP) - GSCOP, 46 av. Félix Viallet, 38000 Grenoble France.
- Stauffer G, Massonnet G, Rapine C, Gayon JP (2011) A simple and fast 2-approximation for the one-warehouse multi-retailers problem. *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 67–79.
- van Hoesel C, Wagelmans A (1996) An $\mathcal{O}(T^3)$ algorithm for the economic lot-sizing problem with constant capacities. *Management Sci.* 42(1):142–150.
- van Hoesel S, Romeijn HE, Morales DR, Wagelmans A (2005) Integrated lot-sizing in serial supply chains with production capacities. *Management Sci.* 51(11):1706–1719.
- Van Vyve M (2007) Algorithms for single-item lot-sizing problems with constant batch size. *Math. Oper. Res.* 32(3):594–613.
- Wagelmans A, Van Hoesel C, Kolen A (1992) Economic lot sizing: an $O(n \log n)$ algorithm that runs in linear time in the wagner-whitin case. *Oper. Res.* 40(S1):145–156.
- Wagner H, Whitin T (1958) Dynamic version of the economic lot size problem. *Management Sci.* 5(1):89–96.
- Zipkin P (2000) *Foundations of Inventory Control* (McGraw-Hill/Irwin).