



HAL
open science

Approximation algorithms for deterministic continuous-review inventory lot-sizing problems with time-varying demand

Guillaume Massonnet, Jean-Philippe Gayon, Christophe Rapine

► **To cite this version:**

Guillaume Massonnet, Jean-Philippe Gayon, Christophe Rapine. Approximation algorithms for deterministic continuous-review inventory lot-sizing problems with time-varying demand. *European Journal of Operational Research*, 2014, 234 (3), pp.641-649. 10.1016/j.ejor.2013.09.037 . hal-00947092

HAL Id: hal-00947092

<https://hal.science/hal-00947092v1>

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.

Approximation algorithms for deterministic continuous-review inventory lot-sizing problems with time-varying demand

G. Massonnet^a, J.-P. Gayon^a, C. Rapine^b

^aGrenoble INP, Laboratoire G-SCOP, 46 avenue Félix Viallet, 38031 Grenoble Cedex, France.

^bUniversité de Lorraine, Laboratoire LGIPM, île du Saulcy, 57045 Metz Cedex 1, France.

Abstract

This work deals with the continuous time lot-sizing inventory problem when demand and costs are time-dependent. We adapt a cost balancing technique developed for the periodic-review version of our problem to the continuous-review framework. We prove that the solution obtained costs at most twice the cost of an optimal solution. We study the numerical complexity of the algorithm and generalize the policy to several important extensions while preserving its performance guarantee of two. Finally, we propose a modified version of our algorithm for the lot-sizing model with some restricted settings that improves the worst-case bound.

Keywords: Inventory theory, approximation algorithms, continuous-review policy, deterministic model

1. Introduction

The Economic Order Quantity (EOQ) problem deals with a single location that faces a demand of constant rate λ . In this model, costs are incurred when an order is placed as well as when units are physically held in the stock. The goal is to determine a continuous-review policy of minimal cost. More precisely, placing an order incurs a fixed order cost K and a linear order cost c while holding an inventory unit incurs a cost h per unit of time. This problem has been solved for a long time by Harris (1913) in the early twentieth century and popularized by Wilson (1934). Since then, many extensions and variations have been studied including production capacity, backorders, perishability, multi-echelon systems (see e.g. Zipkin (2000) for a state of the art).

One of the important limitations of the EOQ model is the assumption of time-independent parameters. In particular, it assumes a constant demand rate λ , which is not realistic in numerous practical situations. Many authors relax this assumption and consider a more general model with time-varying demand rate $\lambda(t)$. However they only solve the problem under very restrictive assumptions. Resh et al. (1976) consider a time proportional demand rate ($\lambda(t) = \alpha t$) with time-independent

cost parameters. Later, Donaldson (1977) proposes an optimal policy for a linear demand rate ($\lambda(t) = \alpha t + \beta$) and Barbosa & Friedman (1978) generalize to power-form demand rates ($\lambda(t) = \alpha t^\beta$, $\beta > -2$). This result is then extended by Henery (1979) to increasing log-concave demand patterns and Hariga (1994) who studies a more general model for any monotonic log-concave demand rate when shortages are allowed. Finally, Henery (1990) focuses on non monotonic demand patterns in the special case of cyclic demands.

These papers use a similar general approach that consists in finding the optimal policy for a fixed number of orders over the planning horizon, then to determine the optimal number of orders minimizing the total cost. This concept has been studied and generalized by Benkherouf & Gilding (2009), who show that the optimal cost over a finite planning horizon is a convex function of the number of orders. However, this method requires specific properties on the cost function and therefore the papers mentioned above restrict their attention to specific demand patterns and time-independent cost parameters. There is no general technique to determine an optimal continuous-review policy for time-varying parameters without the restrictive assumptions discussed above.

As the problem is difficult for more general demand patterns, the literature proposes several heuristics based on empirical methods such as the greedy (or myopic) approach. Silver (1979) introduces a heuristic where

Email addresses: guillaume.massonnet@g-scop.fr
(G. Massonnet), jean-philippe.gayon@g-scop.fr
(J.-P. Gayon), christophe.rapine@univ-lorraine.fr
(C. Rapine)

the length of a replenishment cycle is chosen such that the average cost is minimized (locally) on the replenishment cycle. This heuristic is inspired from the well known Silver & Meal (1973) heuristic designed for a periodic-review setting. Another one, widely used in practice, consists of averaging the demand rate and apply the EOQ formula with the average demand to compute the order times. Goyal & Giri (2003) consider an extension of Silver (1979) with backorders and time-varying demand rate $\lambda(t)$, production rate $\mu(t)$ and deterioration rate $\theta(t)$. They also use a greedy approach, optimizing the average cost on each cycle rather than globally. Many other heuristics and extensions have been considered in the literature. We refer the reader to Goyal & Giri (2001) and Bakker et al. (2012) for a review of deteriorating-inventory models with time-varying demand and to Teng et al. (2007) for some references on the Economic Production Quantity (EPQ) model with time-varying demand.

Another important heuristic is to discretize time and use dynamic programming techniques to solve the corresponding periodic-review problem. Indeed, there exists several polynomial time algorithms that are optimal for the periodic-review lot-sizing problem. In the discrete time version, the planning horizon is divided into n periods. In each period a demand occurs and an order can be placed to replenish the stock. Wagner & Whitin (1958) present an algorithm of complexity $O(n^2)$, while Aggarwal & Park (1993) propose a $O(n)$ algorithm to solve this problem using Monge arrays. However, it is very unlikely that the optimal solution for the discretized problem matches an optimal policy for the original continuous-review version.

More generally, the previous heuristics have an important drawback: They can perform arbitrarily bad (i.e. no guarantee of performance has been proven). Even in the periodic-review setting, Axsäter (1982) shows that the heuristic of Silver & Meal (1973) can perform arbitrarily bad. That is, on some instances, the cost of the heuristic is arbitrary large compared to the cost of the optimal policy. Similarly, the EOQ heuristic (that averages the demand and uses the EOQ formula) is also shown to perform arbitrarily bad (see Bitran et al. (1984)), even if the average demand is recalculated whenever an order is placed.

In this paper, we aim to derive tractable policies for the continuous lot-sizing problem with time-varying parameters that have provable performance guarantees. A policy is said to have a performance guarantee (or a worst-case guarantee) of α if its cost is at most α times the cost of an optimal policy. In the periodic-review setting with time-varying demand, Axsäter (1982) shows

that a policy that balances holding and set-up costs in each replenishment cycle has a performance guarantee of two. Bitran et al. (1984) extend the results of Axsäter to the case with time-varying holding costs. Recently, Van den Heuvel & Wagelmans (2010) have proven that in fact this worst-case guarantee cannot be improved for online heuristics, i.e. procedures that use a forward induction mechanism and cannot modify their past ordering periods.

In this paper, we apply the ideas of cost balancing introduced by Bitran et al. (1984) to the continuous-review problem and prove that the performance guarantee of two remains valid for this problem. We then show that the ideas developed are quite generic and can be extended to several important models from the inventory literature (production rate, deterioration rate, non-linear holding costs, models with shortages, time-varying order costs) while preserving the performance guarantee. Similarly to Bitran, our method uses a myopic mechanism to decide whether to place a new order or increase the quantity ordered in the previous one and hence is unlikely to achieve a worst-case guarantee lower than two. However, we introduce at the end of this paper an alternative procedure that enables us to improve the performance guarantee to $3/2$, with some additional assumptions.

Note that Hariga (1996) applies a similar cost balancing technique to the time-varying demand model with shortages and compares its performances to several other heuristics widely used in practice. Although his study presents numerical experiments and compare the performances of the different techniques, it focuses on time-independent cost parameters. Moreover, the theoretical performance and complexity are not discussed in his paper. Thus, to the best of our knowledge, this paper is the first to derive algorithms and prove their guarantee of performance for lot-sizing problems with time-varying parameters in a continuous-time setting.

The rest of the paper is organized as follows: in §2, we formally introduce the basic lot-sizing problem and useful notations to describe the solutions and prove their guarantee. §3 explains how to adapt the balancing policy of Bitran et al. (1984) for a continuous-time setting and prove that the resulting policy preserves the performance guarantee of 2. In §4, we introduce several extensions to the lot-sizing model and show how the balancing algorithm can be modified to solve them. Finally, we generalize the concept of cost balancing to present a new policy for the lot-sizing model and prove its worst-case guarantee is 1.5 in §5.

2. Assumptions and notations

We consider a single location that faces a continuous demand with rate $\lambda(t)$ at time t . We denote by $\Lambda(s, t) = \int_s^t \lambda(u)du$ the cumulated demand over $[s, t]$. Holding an inventory unit at time t incurs a cost $h(t)$ per unit of time. We assume that $\lambda(\cdot)$ and $h(\cdot)$ are piece-wise continuous functions. Placing an order at time t incurs a fixed order cost (or set-up cost) $K(t)$ and a linear order cost $c(t)$.

The inventory level at time 0, before placing the first order, is denoted by x_0 . For the majority of the models presented in this paper, it is dominant to place the first replenishment at time t_0 such that $x(t_0) = 0$ and $\lambda(t_0) > 0$: Hence we will assume w.l.o.g. that $x_0 = 0$ and $\lambda(0) > 0$, unless explicitly said otherwise. In what follows, a *policy* is defined as a set of rules to determine the ordering times and quantities for each instance of the problem. The objective is to find a policy minimizing costs over a finite horizon $[0, T]$ while satisfying all the demands: In particular, a policy that achieves the lowest possible cost for any instance of the problem is said to be optimal.

When the demand rate is constant, the cost parameters are time-independent and the horizon is infinite, this problem reduces to the EOQ problem and can be solved analytically (see e.g. Zipkin (2000)). On the contrary, it is generally hard to solve the problem to optimality when parameters are time-dependent. The aim of this paper is to develop approximation algorithms for the time-varying version of the problem.

Without loss of generality, we assume that the procurement leadtime is null: All demands being deterministic, a deterministic positive leadtime simply shifts the decision earlier in time. We now introduce some useful notations and concepts for the rest of the paper. Let $x^P(u)$ be the inventory level at time u , under some policy P and let $C^P(s, t)$ be the cost incurred by a policy P over $(s, t]$. More precisely, $C^P(s, t)$ includes holding cost $\int_s^t h(u)x^P(u)du$ and set-up costs over $(s, t]$, excluding the set-up cost at time s , if any. Then, for any sequence $a_0 = 0 < a_1 < \dots < a_{n+1} = T$, the total cost C^P of policy P over the whole time horizon can be decomposed as

$$C^P = \kappa_0^P + C^P(0, T) = \kappa_0^P + \sum_{i=0}^n C^P(a_i, a_{i+1}) \quad (1)$$

, where κ_0^P is equal to $K(0)$ if P orders at time 0 and 0 otherwise. Note that since intervals $(a_i, a_{i+1}]$ partition the time horizon, equation (1) partitions the cost incurred by any policy according to the sequence $(a_i)_{i=0, \dots, n+1}$. In particular, the following proposition holds:

Proposition 1. *Let π be a feasible policy for the continuous-review problem and let $0 = a_0 < a_1 < \dots < a_n < a_{n+1} = T$ be a sequence of points in time. Then if $\kappa_0^\pi \leq \kappa_0^P$ and $C^\pi(a_i, a_{i+1}) \leq \alpha C^P(a_i, a_{i+1})$ for all $i = 0, \dots, n$ and all feasible policy P , π has a worst case guarantee of α .*

Proof. Summing all the partial costs as in equation (1) leads to $C^\pi \leq \alpha C^P$ for any feasible policy P . In particular, this inequality holds when P is an optimal policy and the proof follows. \square

Given a policy P , a (replenishment) cycle $(s, t]$ is defined as a time interval such that an order is placed at time s , an order is placed at time $t > s$ and no order is placed inbetween. In particular, if the sequence $(a_i)_{i=0, \dots, n+1}$ corresponds to the ordering times of policy P $P0 = s_0 < s_1 < \dots < s_n < s_{n+1} = T$ is the sequence of ordering times of policy P , one can decompose that given Finally, we say that a policy satisfies the *Zero Inventory Ordering* property, or is ZIO, if it orders only when its inventory level is zero. Note that under a ZIO policy P , the quantity ordered at the beginning of a cycle $(s, t]$ is exactly $\Lambda(s, t)$, the cumulated demand over the cycle. Moreover, the inventory level can be easily expressed in cycle $(s, t]$ as $x^P(u) = \Lambda(u, t)$. Hence the cumulated holding cost $H(s, t)$ of a ZIO policy over a cycle $(s, t]$ can be simply expressed as an integral of holding cost and demand functions.

$$H(s, t) = \int_s^t h(u)\Lambda(u, t)du$$

Note that since $h(\cdot)$ and $\Lambda(\cdot, \cdot)$ are both nonnegative piecewise continuous functions, $H(s, \cdot)$ is a nondecreasing continuous function for all s . We give below a summary of the notations used through the paper:

$\lambda(t)$	demand rate at time t
$\Lambda(s, t)$	cumulated demand over $[s, t]$
$\Lambda(t)$	cumulated demand from time 0 to t : $\Lambda[0, t]$
$h(t)$	per-unit holding cost at time t
$K(t)$	fixed order cost (or set-up cost) at time t
$c(t)$	linear order cost at time t
x_0	inventory level at time 0
T	time horizon
$x^P(u)$	inventory level at time u , under policy P
$C^P(s, t)$	cost of policy P over $(s, t]$ (time s excluded)
C^P	cost of policy P over $[0, T]$
$H^P(s, t)$	holding cost of policy P over cycle $(s, t]$
$H(s, t)$	holding cost of a ZIO policy over cycle $(s, t]$

3. An approximation algorithm

In this section, we present the central idea of our algorithm. Roughly, the main concept is to balance in each replenishment cycle the holding costs with the fixed order cost. For ease of understanding, we first make similar assumptions as Bitran et al. (1984) who have investigated a periodic-review version of this problem. More precisely, we assume that the fixed order cost and the linear order cost are time-independent, that is $K(t) = K$ and $c(t) = c$ for all $t \in [0, T]$. These assumptions will be relaxed in §4. Notice that without loss of generality, we can assume $c = 0$ as any policy incurs a linear order cost of exactly $c(\Lambda(0, T) - x_0)$ over the planning horizon to satisfy all the demands. Throughout the remainder of this section, we restrict our attention to policies that respect the Zero Inventory Ordering (ZIO) property, as it is dominant for the model we are interested in.

3.1. The balancing policy

The concept of cost balancing is well known in the inventory literature and can even be found in the classical EOQ problem, where all parameters are time-independent. Indeed, it is well known that for this special case the optimal policy balances exactly the holding costs with the order costs. More precisely, the optimal average cost C^* over an infinite horizon is

$$C^* = \underbrace{\sqrt{\lambda Kh/2}}_{\text{Holding costs}} + \underbrace{\sqrt{\lambda Kh/2}}_{\text{Fixed order costs}}$$

In the time-varying setting, we introduce the *balancing policy*, denoted BL in what follows, which balances in each cycle the two parts of the cost discussed above. The BL policy is inspired by the periodic-review policy proposed by Axsäter (1982) and Bitran et al. (1984). In fact, it is a ZIO policy whose order times are determined by a forward induction as follows: A first order is placed at the last moment when the inventory is nonnegative. If the BL policy places an order at time s , the next order time is then defined as the last time t such that

$$H(s, t) \leq K \quad (2)$$

Recall that $H(s, t)$ represents the holding cost incurred by a ZIO policy over cycle $(s, t]$. As $H(s, \cdot)$ is a nondecreasing continuous function and $H(s, s) = 0$, inequality (2) is tight for at least one point in $[s, T]$ if and only if $H(s, T) \geq K$ (see Figure 1). When the latter condition is not satisfied then s is the last order time of the BL policy and the quantity ordered at s is $\Lambda(s, T)$, such that all the remaining demand is ordered at s . More precisely, we define the BL policy as follows:

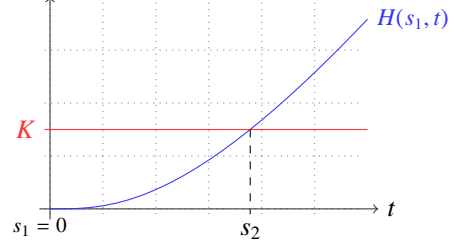


Figure 1: Computing the next ordering time s_2 when the current order is placed at time $s_1 = 0$ ($K = 1.5$, $h(t) = \frac{\ln(t+1)}{t+1}$, $\lambda(t) = \frac{2\ln(t+1)}{(t+1)}$).

Definition 1. *The BL policy is the ZIO replenishment policy whose order times $s_1 < s_2 < \dots < s_n$ are given by Algorithm 1.*

Algorithm 1 BL policy

```

set  $s_1 = 0$ 
set  $n \leftarrow 1$ 
while  $H(s_n, T) \geq K$  do
     $s_{n+1} \leftarrow \max \{t \leq T : H(s_n, t) = K\}$ 
     $n \leftarrow n + 1$ 
end while
return  $(s_1, \dots, s_n)$ 

```

Since the BL policy is ZIO, it is completely defined by the vector of order times (s_1, \dots, s_n) . By convention we set $s_0 = 0$ and $s_{n+1} = T$. Then for $i \geq 1$, the quantity ordered by the BL policy at time s_i is exactly the cumulated demand over $[s_i, s_{i+1}]$, i.e. $\Lambda(s_i, s_{i+1})$.

3.2. Performance guarantee of the BL policy

The BL policy is an interesting inventory policy both for its simplicity and for the quality of the solution obtained. In particular, we show in this section that the cost of the BL policy is at most twice the cost of an optimal policy. We first present a lower bound on the cost incurred by any policy between two instants $s < t$:

Lemma 1. *For any instants $s < t$ and any feasible policy P , we have*

$$C^P(s, t) \geq \min\{H(s, t), K\}$$

Proof. If policy P places an order on $(s, t]$, clearly $C^P(s, t) \geq K$. Otherwise, the inventory level $x^P(s)$ of policy P at instant s is at least $\Lambda(s, t)$ to prevent stockout on $[s, t]$. Since no order is placed on $(s, t]$, we have that for all $u \in (s, t]$, $x^P(u) \geq \Lambda(u, t)$ for $u \in (s, t]$ and thus $H^P(s, t) = \int_s^t h(u)x^P(u)du \geq \int_s^t h(u)\Lambda(u, t)du = H(s, t)$. The results follows. \square

We now use this result to state and prove the following theorem on the performance guarantee of the BL policy:

Theorem 1. *For piece-wise continuous functions $\lambda(\cdot)$ and $h(\cdot)$ and constant functions $K(\cdot)$ and $c(\cdot)$, the BL policy has a worst-case guarantee of two. That is, the cost incurred by the BL policy is at most twice the cost of an optimal policy.*

Proof. Let P be a feasible policy for the continuous-review lot-sizing problem considered and let s_1, \dots, s_n be the sequence of order times of the BL policy. We prove that on any time interval $(s_i, s_{i+1}]$, the cost incurred by the BL policy is at most twice the cost incurred by P .

Since P is feasible, we clearly have $\kappa^B = \kappa^P$. On the last interval $(s_n, T]$, the cost incurred by BL is $H(s_n, T)$, which is lower than K by construction. According to Lemma 1, we have

$$\begin{aligned} C^P(s_n, T) &\geq \min\{H(s_n, T), K\} \\ &= H(s_n, T) \\ &= C^{BL}(s_n, T) \end{aligned}$$

Now consider an interval $(s_i, s_{i+1}]$, with $1 \leq i \leq n-1$. By construction, the BL policy incurs a holding cost $H(s_i, s_{i+1}) = K$ and thus $C^{BL}(s_i, s_{i+1}) = H(s_i, s_{i+1}) + K = 2K$. Due to Lemma 1, any feasible policy pays at least $\min\{H(s_i, s_{i+1}), K\} = K$ on $(s_i, s_{i+1}]$. In particular, we have that $C^P(s_i, s_{i+1}) \geq K$ and therefore

$$C^{BL}(s_i, s_{i+1}) \leq 2C^P(s_i, s_{i+1})$$

for all $i = 0, \dots, n-1$. The theorem then follows from Proposition 1. \square

Remark 1. *The BL policy is optimal in the case of the EOQ problem.*

Assume that $\lambda(t) = \lambda$ and $h(t) = h$ for all t . In this case, we have $\Lambda(s, t) = (t-s)\lambda$ and $H(s, t) = \lambda h(t-s)^2/2$. Solving Equation (2) yields

$$t - s = \sqrt{\frac{2K}{\lambda h}}$$

The time between orders is precisely the optimal order interval for the EOQ model (see e.g. Zipkin (2000)), which proves the optimality of the BL policy.

3.3. Numerical issues and complexity

In this section, we discuss how the order times can be computed by the balancing algorithm in practice. We

first turn our attention to the issue of computing the next order time t given the current order time s . From the definition of $H(s, \cdot)$, we have that $H(s, s) = 0$ and $H(s, \cdot)$ is a nondecreasing continuous function on $[s, T]$. Hence there exists at least one solution $t \in [s, T]$ to the equation $H(s, t) = K$ (unless $H(s, T) < K$). However, in most cases there is no analytical expression of t . One can then use classical root-finding algorithms, such as the bisection method, to compute an approximate value of t . The next ordering moment t' found by such techniques is such that $t' \in [t - \delta, t + \delta]$ for a given precision $\delta > 0$ and the final complexity depends on δ . For instance when using the bisection method, each computation of t' requires $O(\log(T/\delta))$ evaluations of function $H(s, \cdot)$.

For some classes of functions $h(\cdot)$ and $\lambda(\cdot)$, like polynomial, exponential, sinusoidal, the integral $H(s, t)$ can be expressed analytically. Otherwise, one can use classical numerical methods to evaluate this integral with some precision error. In what follows, we assume that $H(s, t)$ can be computed (exactly) in time $O(1)$ and we only focus on the complexity of determining the order times s_1, \dots, s_n of the BL policy. Notice that the number of order times of the BL policy, as for an optimal policy, may be arbitrary large (for instance considering an exponential holding cost $h(t) = e^t$). Thus the time complexity can not be bounded in general with respect to the instance size. For this reason we consider in what follows a natural restriction of the problem.

Assume that the demand function and the holding cost function are both upper bounded on $[0, T]$. That is, there exists $\hat{\lambda}, \hat{h}$ such that for all $t \in [0, T]$, $\lambda(t) \leq \hat{\lambda}$ and $h(t) \leq \hat{h}$. Then the following theorem holds:

Theorem 2. *For any $\varepsilon > 0$, the BL policy can be implemented to provide a solution with a performance guarantee of $(2 + \varepsilon)$ using $O(\Phi \log \Phi)$ evaluations of $H(\cdot, \cdot)$,*

where $\Phi = (1 + \frac{1}{\varepsilon}) \frac{\hat{\lambda} \hat{h}}{K} T^2$.

Proof. See Appendix Appendix A. \square

Remark 2. *In practice, $H(s, t)$ is often too complex to be computed exactly in constant time. However, if one uses a numerical method to compute $H(s, t)$ with maximum error $\gamma \geq 0$ in time $O(f(\gamma, T))$, this error can be somehow included in the general error ε of BL. Specifically, the precision for $H(\cdot, \cdot)$ must satisfy $\gamma < \varepsilon K$ and the total complexity is $O(\Phi' \log \Phi' f(\gamma, T))$, where*

$$\Phi' = \frac{1 + \varepsilon}{\varepsilon K - \gamma} \hat{\lambda} \hat{h} T^2.$$

3.4. Bad example

We should be able to present a quick bad example, where demand rate is null everywhere except on regular small intervals, and its limit when the length of the intervals tends to zero reach the worst-case guarantee. I found a simple bad example, just need to write it down clearly.

4. Extensions

The previous section introduces a cost balancing technique for a basic continuous lot-sizing model, but the underlying idea appears to be rather generic. In this section, we apply this concept to several important extensions of the lot-sizing problem. Although in most cases Algorithm 1 cannot be used directly, we adapt the policy to the specific settings of each extension and prove that the performance guarantee of two found in §3 remains valid for these more general models.

4.1. Nonlinear holding costs

The analysis used in §3 remains valid for nonlinear holding costs. In this model, the holding cost incurred depends on both the time moment and the current inventory on hand. Specifically, we denote $h(x, t)$ the holding cost value at t for holding x units in stock. The only (natural) assumption we make is that $h(x, t)$ is nondecreasing in x .

Essentially, one only has to check that Lemma 1 remains valid in this context to ensure the approximation result holds. As before, let $H(s, t)$ be the holding cost incurred on a cycle $(s, t]$ by a ZIO policy. We have $H(s, t) \equiv \int_s^t h(\Lambda(u, t), u) du$. Now consider a feasible policy P that does not order in the time interval $(s, t]$. The holding cost $H^P(s, t)$ incurred by P on $(s, t]$ satisfies:

$$\begin{aligned} H^P(s, t) &\equiv \int_s^t h(x^P(u), u) du \\ &= \int_s^t h(x^P(s) - \Lambda(s, u), u) du \\ &\geq \int_s^t h(\Lambda(s, t) - \Lambda(s, u), u) du \\ &= \int_s^t h(\Lambda(u, t), u) du = H(s, t) \end{aligned}$$

where the second equality is due to the fact that P does not order in (s, t) and subsequent inequalities come from the nondecreasing property of $h(\cdot, u)$. Therefore Lemma 1 is still valid.

The policy we propose in this case follows exactly Algorithm 1. This policy satisfies the ZIO property, thus the same arguments as the ones used in Theorem 1 hold and the performance guarantee of 2 remains valid in the general case of nonlinear holding costs.

4.2. Perishable products

Classical inventory models generally assume that one can store units indefinitely to meet future demands. However, many products do not satisfy this assumption in practice, e.g. because they deteriorate or become obsolete (see Goyal & Giri (2001) for a recent review on these models). In this section, we consider a model in which the inventory on hand at time t deteriorates at rate $\theta(t)$, incurring a per unit deterioration cost $a(t)$. In a cycle $(s, s']$, the inventory level $x(t)$ satisfies the following first-order differential equation:

$$\frac{dx(t)}{dt} + \theta(t)x(t) = -\lambda(t), \quad x(s') = 0$$

which can be solved easily, see e.g. Goyal & Giri (2003). Note that even though we only present our results for the simple inventory model derived from §3, this equation can also be solved with production rate $\mu(t)$ and backorders (see Goyal & Giri (2003)).

The sum of holding and deterioration costs over a time interval $[s, s']$ for a policy P can be expressed as:

$$H^P(s, s') = \int_s^{s'} [h(t) + a(t)\theta(t)]x^P(t)dt = \int_s^{s'} \tilde{h}(t)x^P(t)dt$$

where $\tilde{h}(t) = h(t) + a(t)\theta(t)$. Using this modified holding cost parameter, we can apply Algorithm 1 as in §3. The arguments used to prove the performance guarantee of BL remain valid in this more general context and thus the resulting algorithm is a 2-approximation for the model with perishable products.

4.3. Finite production rate

The model studied in §3 assumes that replenishments are instantaneous. We now relax this assumption and focus on a model where units are produced according to a time-dependent, piecewise continuous production rate $\mu(\cdot)$. In addition, we denote $\mu(s, s') = \int_s^{s'} \mu(z)dz$ the cumulative production capacity on $[s, s']$. If a production starts at time s for q units, all units are available on hand at time $s + t$ such that $\mu(s, t) = q$. Each production start incurs a time-independent fixed cost K (the nonspeculative linear production cost is omitted w.l.o.g. as before). The objective is again to minimize total costs over $[0, T]$ without stockout.

Although classical production models assume that the production rate $\mu(t)$ is larger than the demand rate $\lambda(t)$, for all t , we relax this constraint and consider a more general situation where the only assumption is that the production capacity is sufficient to satisfy the entire demand without stockout, that is $\mu(0, t) \geq \Lambda(0, t)$ for all t .

As a consequence, a policy P for the problem is feasible if and only if it can satisfy the demands $\Lambda(s, t)$ on any time interval $[s, t]$, either by producing during this interval or by using its stock on hand at instant s . Observe that any policy cannot produce more than $\mu(s, t)$ units during the interval. Thus a sufficient condition for P to avoid stockouts is that for any time $t \geq s$, $x^P(s) \geq \Lambda(s, t) - \mu(s, t)$.

Property 1. Any feasible policy P satisfies

$$\forall s \in [0, T], x^P(s) \geq x_{min}(s) \quad (3)$$

with $x_{min}(s) = \max_{t \geq s} \{\Lambda(s, t) - \mu(s, t)\}$

When $t = s$, this condition simply implies that $x^P(s) \geq 0$. According to this new definition, we modify the initial condition and assume w.l.o.g. in this model that $x_0 = x_{min}(0)$ and $\lambda(0) > 0$.

Note that in the special case where $\mu(t) \geq \lambda(t)$ for all t , we have $x_{min}(s) = 0$ for all s and the ZIO property remains dominant. However in this more general production model, it is not necessarily the case. Consequently we focus on policies that order at time s only if their inventory level is equal to $x_{min}(s)$. In the remainder of this section, we call such policies MIO (Minimum Inventory Ordering) policies. In addition to this main property, the final stock of a MIO policy must be null ($x(T) = 0$). MIO policies are clearly dominant for our problem since for any non MIO policy, one can easily build a MIO policy with lower cost by delaying orders. Figure 2 illustrates a MIO policy.

For a MIO policy P , one can easily compute the stock evolution during a replenishment cycle $(s, t]$. First, note that by definition we have $x^P(s) = x_{min}(s)$ and $x^P(t) = x_{min}(t)$. Due to the continuous production process, cycle $(s, t]$ consists now of an active period $(s, u]$ (when production is on-going) and an idle period $(u, t]$ (when production is over). The idle period starts at time u , where u satisfies the following equation:

$$x_{min}(t) = x_{min}(s) + \mu(s, u) - \Lambda(s, t) \quad (4)$$

The value of the inventory level over a replenishment cycle $(s, t]$ is then easy to compute:

$$x^P(z) = \begin{cases} x_{min}(s) + \mu(s, z) - \Lambda(s, z) & \text{if } z \leq u \\ x_{min}(t) + \Lambda(z, t) & \text{otherwise} \end{cases}$$

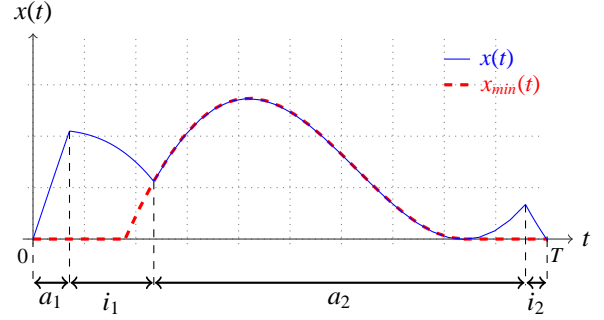


Figure 2: An example of a MIO policy for $T = 10$, $\lambda(t) = 2(1 - \cos(\frac{t}{2}))$, $\mu(t) = 2.5$. Time intervals a_1 and a_2 correspond to production (active) periods while i_1 and i_2 are idle periods.

Using equation (4) we can unify the expression of $x^P(z)$ as follows. Consider an instant $z \leq u$. We have $x^P(z) = [x_{min}(t) - \mu(s, u) + \Lambda(s, t)] + \mu(s, z) - \Lambda(s, z)$ and therefore $x^P(z) = x_{min}(t) - \mu(z, u) + \Lambda(z, t)$. Assuming $\mu(b, a) = 0$ for $a < b$, we get:

$$x^P(z) = x_{min}(t) - \mu(z, u) + \Lambda(z, t) \quad \forall z \in [s, t] \quad (5)$$

Let $H(s, t)$ be the holding cost incurred by a MIO policy over a replenishment cycle $(s, t]$. Since for any feasible policy P the stock level at any instant z is at least $x_{min}(z)$, the holding cost can now be split into two parts, $H(s, t) = H_1(s, t) + H_{min}(s, t)$, where:

$$H_1(s, t) \equiv \int_s^t h(z)(x(z) - x_{min}(z))dz \geq 0$$

is the *policy dependent part* of $H(s, t)$ and

$$H_{min}(s, t) \equiv \int_s^t h(z)x_{min}(z)du \geq 0$$

is the *independent part* of $H(s, t)$. Notice that due to Equation 3, $H_{min}(s, t)$ is incurred by any feasible policy for the problem we consider.

We now present an extension of the BL policy to this production model, called PB policy in what follows. The PB policy balances in each cycle the fixed order cost K with the policy dependent part of the holding cost. Formally, the algorithm for the production model works as follows:

Algorithm 2 Balancing policy for production models (PB)

```

set  $s_1 \leftarrow 0$ 
set  $n \leftarrow 1$ 
while  $H_1(s_n, T) \geq K$  do
     $s_{n+1} \leftarrow \max \{t \leq T : H_1(s_n, t) = K\}$ 
     $n \leftarrow n + 1$ 
end while
return  $(s_1, \dots, s_n)$ 

```

For $i = 1, \dots, n$, the quantity ordered by policy PB at time s_i is then $q_i = \Lambda(s_i, s_{i+1}) + x_{\min}(s_{i+1}) - x_{\min}(s_i)$ and the following theorem holds:

Theorem 3. *The PB policy has a worst-case guarantee of 2 for the general production model.*

Proof. See Appendix Appendix B. \square

4.4. A model with backlogging

One major constraint in the previous models is that demand has to be satisfied immediately. However, there exists more flexible models in which the demand does not necessarily have to be fulfilled on time. We now focus on such a model and present an approximation algorithm based on the same balancing idea as the one discussed in §3.

In this section, we assume backlogging is allowed. The model is similar to the one presented in §2, except that demand is not necessarily satisfied immediately anymore. We still assume that the policy incurs a fixed order cost K for each order placed while holding inventory induces a per-unit holding cost of $h(t)$ at time t . On the other hand, unmet demand is *backlogged*: Customers are willing to wait until the stock on hand is sufficient to fulfill their requirement. Thus backlogged demand is served with a subsequent order placed later in time and incurs a per-unit backlogging cost of $b(t)$ at time t .

Without loss of generality, we assume the units are consumed on a first-ordered first consumed basis. Moreover, we modify the initial conditions and simply consider that $x_0 \leq 0$. The balancing policy for models with backlogging uses an idea rather similar to the one used in the basic model: It balances the costs incurred between two consecutive order times with the fixed order cost K . The main difference is that instead of using only the holding costs, both holding and backlogging costs are considered. However, the policy decides to serve or backlog demand in a way that minimizes the sum of the costs incurred over the cycle. Each cycle $(s, t]$ is consequently split into two parts: $(s, u]$, during

which inventory is held, and $(u, t]$, when backorders accumulate.

In the remainder of the section, the balancing policy for models with backlogging is denoted BB. For a given policy P , let $B^P(s, t)$ be the backlogging cost incurred by P over $(s, t]$. In addition, we denote $B(s, t)$ the backlogging cost incurred over $(s, t]$ by a policy that does not order in (s, t) and whose inventory at time s is 0. Formally, the algorithm works as follows:

Algorithm 3 Balancing policy for models with backlogging (BB)

```

set  $s_0 \leftarrow 0, s_1 \leftarrow \max \{t \leq T : B(0, t) = K\}$ 
set  $n \leftarrow 1$ 
while  $H(s_n, T) > K$  do
     $s_{n+1} \leftarrow \max \left\{ t \leq T : \min_{u \in [s_n, t]} \{H(s_n, u) + B(u, t) = K\} \right\}$ 
     $n \leftarrow n + 1$ 
end while
return  $(s_1, \dots, s_n)$ 

```

Given $s < t$, the algorithm computes the quantity that minimizes the holding and backlogging costs incurred over $(s, t]$. This quantity has to satisfy all the backorders at s (before the order occurs) plus the total demand during the first part of the cycle (when inventory is physically held). Since satisfying the backlogged units at s does not induce any cost, the algorithm aims to find $u^* \in (s, t)$ such that function $f_{s,t}(u) = \int_s^u h(z)\Lambda(z, u)dz + \int_u^t b(z)\Lambda(u, z)dz$ is minimized. The first derivative of $f_{s,t}(\cdot)$ is the following expression:

$$\begin{aligned} \frac{df_{s,t}(u)}{du} &= \int_s^u h(z)\lambda(u)dz - \int_u^t b(z)\lambda(u)dz \\ &= \left(\int_s^u h(z)dz - \int_u^t b(z)dz \right) \lambda(u) \end{aligned}$$

As $\lambda(\cdot)$ is a nonnegative function, it is easy to see that for $s < t$, there exists $u^* \in (s, t)$ such that:

$$\begin{cases} \frac{df_{s,t}(u)}{du} \leq 0 & \text{for all } u \in (s, u^*] \\ \frac{df_{s,t}(u)}{du} \geq 0 & \text{for all } u \in (u^*, t] \end{cases}$$

Therefore $f_{s,t}(\cdot)$ is unimodal and reaches its minimum at $u^* \in (s, t)$ such that:

$$\int_s^{u^*} h(z)dz = \int_{u^*}^t b(z)dz$$

This property of functions $f_{s,t}(\cdot)$ is illustrated on Figure 3.

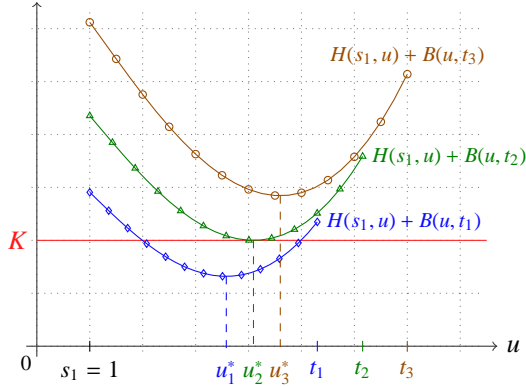


Figure 3: Computing u^* for different values of t ($t_1 = 5.3$, $t_2 = 6.156$, $t_3 = 7$) when the previous order is placed at time $s_1 = 1$ ($K(t) = 2$, $h(t) = 1$, $p(t) = 1.5$, $\lambda(t) = \frac{e^{\sqrt{t+1}}}{4(t+1)} \left(1 - \frac{1}{\sqrt{t+1}}\right)$). The next order chosen by the BB policy is $s_2 = t_2$ since the minimum cost incurred over $(s_1, t_2]$ is exactly equal to K . Furthermore the quantity ordered at s_1 increases the inventory level up to $\Lambda(s_1, u_2^*)$.

If $s < t$ are two consecutive order times of the BB policy, the quantity ordered at time s increases the inventory level up to $\Lambda(s, u^*)$. Note that consequently the remaining $\Lambda(u^*, t)$ units are backlogged and served at time t . Finally if the last order computed by the algorithm is $s_n < T$, the BB policy orders up to $\Lambda(s_n, T)$ at time s_n . Otherwise ($s_n = T$), it orders up to 0 at time T and satisfies all the demand backlogged in $(s_{n-1}, s_n]$. We now state and prove the following theorem on the performance guarantee of the balancing algorithm for models with backlogging:

Theorem 4. *The BB policy for continuous lot-sizing models with backlogging has a worst-case guarantee of two.*

Proof. See Appendix Appendix C. \square

In practice, one can use classical search algorithms such as the Fibonacci search technique (see Kiefer (1953)) to compute the minimum $u^* \in (s, t)$ in $O(\log(T/\delta))$, where δ is the chosen precision. Thus assuming that $\lambda(\cdot)$ (resp. $h(\cdot)$, $p(\cdot)$) is bounded by a constant $\hat{\lambda}$ (resp. \hat{h} , \hat{p}), one can choose Ω as in §3.3 and compute the order times in $O(\Omega(\log \Omega)^2)$ (using the precision $\delta = T/\Omega$) while bounding the error in the cost by a chosen gap ε .

4.5. Time-dependent order costs

In §3 we restricted our attention to time-independent fixed order costs. We now relax this assumption and allow for monotonic time-varying fixed order costs and

nonspeculative linear order costs. Specifically, we assume that function $K(t)$ is continuous in time and either nonincreasing or nondecreasing with t , while $c(\cdot)$ satisfies for all $ds > 0$:

$$c(s + ds) \leq c(s) + \int_s^{s+ds} h(u)du \quad (6)$$

In other words, given two potential order times $s_1 < s_2$ to serve a specific demand at $t \geq s_2$, it is always cheaper to order at s_2 , rather than order earlier at s_1 and hold the corresponding units for a longer period. One can notice that inequality (6) resemble the well-known assumption of nonspeculative motivations often encountered in discrete lot-sizing models.

Intuitively, the replenishment rule used by the policy consists in balancing in each cycle the holding costs and the linear order costs incurred with the minimum fixed order cost over this cycle. Due to our assumption, this minimum fixed order cost is realized at one of the extremities of the cycle. That is, if $K(\cdot)$ is nonincreasing (resp. nondecreasing) the holding cost $H(s, t) + c(s)\Lambda(s, t)$ incurred over a cycle $(s, t]$ is balanced with $K(s)$ (resp. $K(t)$). The procedure is then applied in a forward or in a backward manner depending on the variation of $K(\cdot)$. Notice that Bitran et al. (1984) also define a forward and a backward method for the discrete time problem when ordering costs are time-independent. In our case, these two algorithms are defined as follows:

Algorithm 4 Forward balancing policy for nonincreasing order costs

```

set  $s_1 \leftarrow 0$ 
set  $n \leftarrow 1$ 
while  $c(s_n)\Lambda(s_n, T) + H(s_n, T) \geq K(T)$  do
     $s_{n+1} \leftarrow \max \{t \leq T : c(s_n)\Lambda(s_n, t) + H(s_n, t) = K(t)\}$ 
     $n \leftarrow n + 1$ 
end while
return  $(s_1, \dots, s_n)$ 

```

Remark 3. *When $K(\cdot)$ is nondecreasing and $x_0 > 0$, placing an order before the initial stock is depleted may reduce the total cost incurred by the policy. Hence in this case, we relax the initial condition discussed in §2 and only assume $x_0 \geq 0$.*

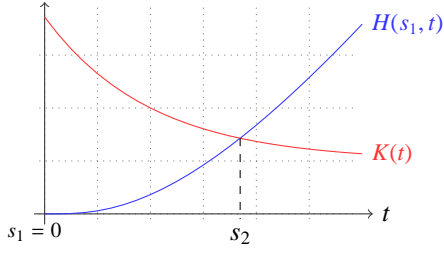


Figure 4: $K(t) = e^{1-\frac{t}{2}}$, $h(t) = \frac{\ln(t+1)}{t+1}$, $\lambda(t) = \frac{2\ln(t+1)}{(t+1)}$.

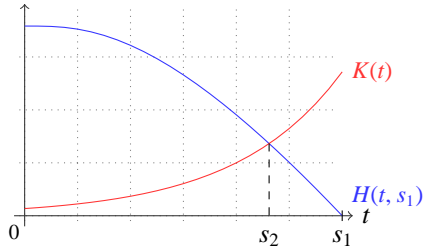


Figure 5: $K(t) = e^{1-\frac{6-t}{2}}$, $h(t) = \frac{\ln(t+1)}{t+1}$, $\lambda(t) = \frac{2\ln(t+1)}{(t+1)}$.

Algorithm 5 Backward balancing policy for nondecreasing order costs

```

set  $s_0 = T$ 
set  $t_0 \leftarrow \max \{t \leq T : \Lambda(0, t) \leq x_0\}$ 
set  $n \leftarrow 0$ 
while  $c(t_0)\Lambda(t_0, s_n) + H(t_0, s_n) \geq K(t_0)$  do
     $s_{n+1} \leftarrow \min \{t \geq t_0 : c(t)\Lambda(t, s_n) + H(t, s_n) = K(t)\}$ 
     $n \leftarrow n + 1$ 
end while
if  $s_n > t_0$  then
     $s_{n+1} \leftarrow \min \{0 \leq t \leq t_0 : c(t)\Lambda(t_0, s_n) + H(t, s_n) - H(t_0, t_0) \leq K(t)\}$ 
     $n \leftarrow n + 1$ 
end if
return  $(s_n, \dots, s_1)$ 

```

Figures 4 and 5 show how the forward and backward algorithms compute the next and the previous order, respectively.

Note that in the case of the backward algorithm, the order times are computed in a reverse order: $0 \leq s_n \leq \dots \leq s_1 \leq T$. Moreover, the ZIO rule is not necessarily respected for the first order. Let x_0 be the initial inventory on hand and let $t_0 = \max \{t \leq T : \Lambda(0, t) \leq x_0\}$: If $s_n < t_0$, the quantity ordered at time s_n is $q_n = \Lambda(t_0, s_{n-1})$. All the other orders are of size $q_i = \Lambda(s_i, s_{i-1})$.

Theorem 5. *For time-dependent order costs, the forward*

ward (resp. backward) balancing policy has a worst-case guarantee of two when $K(\cdot)$ is a continuous nonincreasing (resp. nondecreasing) function.

Proof. See Appendix Appendix D □

5. Improvement of the performance guarantee

In this section, we present a generalization of the balancing technique used in §3 and use it to improve the performance guarantee from 2 to 3/2. The idea remains to balance the costs incurred over a time interval $(s, s']$ with order costs but we now consider the possibility to place additional orders in the interval. We present this modified algorithm in the simple situation of Section 3, in which order costs are time-independent and holding costs are linear.

5.1. A generalized balancing policy

Consider a time interval $(s, t]$ and let k be a positive integer. We define $G_k(s, t)$ as the minimum cost over (s, t) incurred by a feasible policy ordering at most k times in (s, t) .

Note that a policy that minimizes the cost $G_k(s, t)$ is ZIO. Thus for $k = 0$ we have $G_0(s, t) = H(s, t)$ and clearly, the BL policy incurs on each of its ordering interval $(s, s']$ a cost $G_0(s, s')$. We can generalize Lemma 1 to the following result:

Lemma 2. *Let $k \in \mathbb{N}$, $s < t$ and P a feasible policy. We have*

$$C^P(s, t) \geq \min\{G_k(s, t), (k+1)K\}$$

Proof. The proof is immediate due to the definition of G_k . If policy P places at most k orders inside the time interval, then $C^P(s, t) \geq G_k(s, t)$. Otherwise $C^P(s, t) \geq (k+1)K$. □

For a given positive integer k , we propose the following balancing policy, denoted BL_k in the remainder of this section. Its principle is to balance $G_k(s, t)$ with $(k+1)K$ in each of its replenishment cycle $(s, t]$. That is, given an instant s , the policy computes the largest instant $t \leq T$ such that $G_k(s, t) \leq (k+1)K$. Notice that in this case policy BL_k orders at instants s and t and possibly up to k times in the interval (s, t) . We call *main* order times the instants s and t , while the other order times in (s, t) are called *secondary* order times. Although main and secondary orders play the same role in the final solution, we differentiate these two types of orders to simplify the definition and the proof of performance of the algorithm. Similarly to BL, the first (main) order of BL_k is placed at the last instant s_1 such that the

initial inventory x_0 can satisfy the demands on $[0, s_1]$. The following result generalizes Theorem 1:

Theorem 6. *For $K(\cdot)$ and $c(\cdot)$ constant functions and $h(\cdot)$ and $\lambda(\cdot)$ a piece-wise continuous function, the BL_k policy has a worst-case guarantee of $\frac{k+2}{k+1}$.*

Proof. The proof is very similar to the one of Theorem 1 and only the main arguments are given here. Let P be a feasible policy and let s_1, \dots, s_n be the sequence of the main order times of the BL_k policy. First notice that any feasible policy P places an order at time 0 and therefore $\kappa_0^{\text{BL}k} = K = \kappa_0^{\text{P}}$. Now consider an interval $(s_i, s_{i+1}]$, with $0 \leq i \leq n-1$. By construction, the BL_k policy incurs a cost $C^{\text{BL}k}(s_i, s_{i+1}) \leq G_k(s_i, s_{i+1}) + K = (k+2)K$ in this time interval. Due to Lemma 2, policy P pays at least $\min\{G_k(s_i, s_{i+1}), (k+1)K\} = (k+1)K$. It results that $C^{\text{BL}k}(s_i, s_{i+1}) \leq \frac{k+2}{k+1}C^{\text{P}}(s_i, s_{i+1})$. Then consider the last interval. On $[s_n, T]$, we have $C^{\text{BL}k}(s_n, T) = G_k(s_n, T)$, which is lower than $C^{\text{P}}(s_n, T)$ due to Lemma 2. As a consequence the inequality $C^{\text{BL}k}(s_i, s_{i+1}) \leq \frac{k+2}{k+1}C^{\text{P}}(s_i, s_{i+1})$ holds for every interval $(s_i, s_{i+1}]$ and the proof follows. \square

5.2. Complexity analysis

Contrarily to the balancing policy presented in §3, the family of algorithms presented in the previous section does not fall within the class of myopic policies and enables us to design efficient procedures to solve the problem introduced in §2. In particular, the resulting performance guarantee can be as close as aimed to the cost of an optimal policy. However the problem of determining (analytically or numerically) the minimal cost $G_k(s, t)$ on an interval (s, t) is likely to require a large computing effort, especially for large values of k . In this section, we first discuss the computation of $G_k(s, t)$ for an arbitrary k . Then, for $k=1$, we show that the BL_1 policy can be computed in $\mathcal{O}(\Phi(\log \Phi)^2)$ with some restrictions on the demand and holding cost functions.

Arbitrary k

When k secondary orders are placed in interval (s, t) at times v_1, \dots, v_k , the inventory holding cost of a ZIO policy on interval (s, t) is

$$f(v_1, \dots, v_k) = \sum_{i=0}^k \int_{v_i}^{v_{i+1}} h(x)\Lambda(x, v_{i+1})dx$$

where $v_0 = s$ and $v_{k+1} = t$.

The problem is then to minimize f with the constraint that $v_0 = s \leq v_1 \leq \dots \leq v_{k+1} = T$. We have for

$$1 \leq i \leq k$$

$$\frac{\partial f}{\partial v_i}(v_1, \dots, v_k) = \lambda(v_i)h(v_{i-1}, v_i) - h(v_i)\Lambda(v_i, v_{i+1})$$

where $h(x, y) = \int_x^y h(t)dt$.

Hence the optimal secondary order points satisfy the following equality for $1 \leq i \leq k$:

$$\lambda(v_i)h(v_{i-1}, v_i) = h(v_i)\Lambda(v_i, v_{i+1}) \quad (7)$$

Note that this condition is necessary but not sufficient since there may exist several solutions to Equation (7). For time-independent holding cost $h(t) = h$, condition (7) reduces to a simpler condition that was established by Barbosa & Friedman (1978): $(v_i - v_{i-1})\lambda(v_i) = \Lambda(0, v_{i+1}) - \Lambda(0, v_i)$. However even in this restrictive setting, the problem of solving equation (7) is difficult, except for very specific demand functions (linear, power-form).

$k=1$

We now turn our attention to the case $k=1$, where the problem is to find a single optimal intermediary order. In this case, there is a single variable v_1 that we rename v for ease of understanding. If we assume that the demand function $\lambda(\cdot)$ and the holding cost function $h(\cdot)$ are once derivative, we can compute the second derivative of $f(\cdot)$:

$$\frac{d^2 f}{dv^2}(v) = \lambda'(v)h(s, v) + 2\lambda(v)h(v) - h'(v)\Lambda(v, t)$$

As a consequence, $f(\cdot)$ is convex if and only if the following condition is satisfied:

$$\lambda'(v)h(s, v) + 2\lambda(v)h(v) - h'(v)\Lambda(v, t) \geq 0 \quad (8)$$

Unfortunately, there exists numerous demand functions that do not satisfy inequality (8), even when the holding cost function is constant (see for example Figure 6). However, we now exhibit two cases satisfying inequality (8). Both cases relax some assumptions made in the existing literature.

Case 1: $h(\cdot)$ and $\lambda(\cdot)$ once derivative, $h(\cdot)$ non-increasing and $\lambda(\cdot)$ non-decreasing.

In this case $-h'(\cdot)$ and $\lambda'(\cdot)$ are positive and the inequality (8) is trivially satisfied. These assumptions relax the framework used in Resh et al. (1976), Donaldson (1977) and Henery (1979) to any increasing demand function.

Case 2: $h(\cdot)$ once derivative, non-increasing and $\lambda(t) = \alpha t^\beta$, with $\alpha \geq 0$ and $\beta \geq -2$.

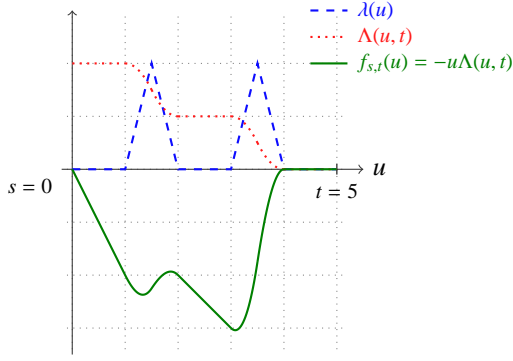


Figure 6: A bad example for function $f_{s,t}(\cdot)$ with triangular demand, $s = 0$, $t = 5$ and $h = 1$.

In this case $-h'(\cdot)$ is nonnegative and $h(s, v) \geq (v - s)h(v)$. The second derivative of $f(\cdot)$ then becomes

$$\begin{aligned} \frac{d^2 f(v)}{dv^2} &= \alpha\beta v^{\beta-1}h(s, v) + 2\alpha v^\beta h(v) - h'(v)\Lambda(v, t) \\ &\geq \alpha h(v) (\beta(v - s)v^{\beta-1} + 2v^\beta) \\ &\geq \alpha h(v) ((\beta + 2)v^\beta - \beta s v^{\beta-1}) \end{aligned}$$

Hence for $\beta \geq 0$ we have

$$\frac{d^2 f(v)}{dv^2} \geq 2\alpha h(v)v^\beta$$

and for $-2 \leq \beta < 0$, we have

$$\frac{d^2 f(v)}{dv^2} \geq \alpha h(v)(\beta + 2)v^\beta$$

In both cases, $\frac{d^2 f(v)}{dv^2} \geq 0$ and $f(\cdot)$ is convex.

Note that this setting extends the problem studied in Barbosa & Friedman (1978) to a time-varying holding cost function.

As a consequence, BL_1 provides solutions that are both computationally tractable and provably close to the optimal, while relaxing assumptions made in the existing literature. The following theorem summarizes the performance guarantee and the complexity of the technique in practice.

Theorem 7. Assume that $\lambda(\cdot)$ and $h(\cdot)$ are bounded, once derivative and satisfy inequality (8) for all $s < v < t$. Then for any $\varepsilon > 0$, the BL_1 policy can be implemented to provide a solution with a performance guarantee of $(3/2 + \varepsilon)$ using $O(\Phi(\log \Phi)^2)$, where Φ is defined as in 3.3.

Proof. The proof is similar to the one of Theorem 2, except that in this case for every t candidate for the next order time, the algorithm has to compute the optimal intermediary order v . For $s < t$, one can use the convexity of $f(\cdot)$ and apply the bisection method to compute the optimal secondary order v^* in $O(T/\delta)$ for a chosen precision δ . Hence for an arbitrary $\varepsilon > 0$, one can define Φ as in §3.3 and choose a precision $\delta = T/\Phi$. The final complexity of BL_1 for a maximum error of ε is then $O(\Phi(\log \Phi)^2)$. \square

6. Conclusion

The EOQ formula for time-independent lot-size models has been extensively studied and is now considered as a classical result in the literature. On the contrary, the optimal policy for its time-varying extension is still an open question and the existing research is limited to restricted cases where demand has a specific pattern (i.e. linear, power form, monotonic log-concave, etc.).

In this paper, we introduce a cost balancing technique and apply it to continuous-review inventory models when demands and cost parameters are time-varying. We prove that the resulting algorithm has a worst-case guarantee of two for the lot-size problem with time-varying parameters. In addition, we show that the underlying idea is rather generic and adapt it to many important extensions, such as perishable products, production systems or models with backlog (among others). For specific demand patterns, we also improve upon the performance guarantee of two and propose a 1.5 approximation algorithm for increasing or power form demand. Finally, we show that the complexity of our algorithm depends mainly on the chosen precision for the approximation ratio. The resulting running time makes it an interesting practical tool for decision makers.

One natural direction for further research is to experiment the cost balancing technique on numerical examples and compare its performances with existing optimal policies or heuristics. Furthermore, we believe that the concept is very generic and can also be applied to multi-echelon systems under mild restrictions. In particular, a first step in this direction would be to study how the algorithm can be extended to series systems.

Appendix A. Proof of Theorem 2

Assume the current order is placed at time $s \geq 0$ and that $H(s, T) \geq K$. The BL algorithm defines the next order time as $t = \max\{u \leq T : H(s, u) \leq K\}$. Numerically, using the bisection method with a precision δ , the

next order time t' is found in $[t - \delta, t + \delta]$. We choose the precision δ of the bisection method to be:

$$\delta = \frac{K}{\lambda \hat{h} T} \frac{\varepsilon}{1 + \varepsilon} = \frac{T}{\Phi}$$

In addition we require that two consecutive order times are separated at least by δ , that is $t' \geq s + \delta$. Therefore the BL policy orders at most T/δ times on the planning horizon. If the computational time to evaluate $H(s, t)$ is $O(1)$ for any $s \leq t$, the final complexity of the BL policy is then $O((T/\delta) \log(T/\delta)) = O(\Phi \log \Phi)$.

It remains to prove that applying the bisection method with the precision δ leads to the expected approximation factor of $2 + \varepsilon$. Recall that the next order time t' found by the bisection method belongs to $[t - \delta, t + \delta]$, even if we impose $t' \geq s + \delta$. For conciseness we denote by $\Delta H(s, r)$ the difference $H(s, r + \delta) - H(s, r)$ for any instant $r \geq s$. According to Lemma 1 and using the fact that $H(s, t) = K$, we can bound the cost ratio between BL and any feasible policy P on (s, t') as follows:

Case 1. $t' \in [t - \delta, t]$:

$$\begin{aligned} \frac{C^{\text{BL}}(s, t')}{C^{\text{P}}(s, t')} &\leq \frac{H(s, t') + K}{H(s, t')} \\ &\leq \frac{H(s, t - \delta) + K}{H(s, t - \delta)} \\ &= \frac{2H(s, t) - (H(s, t) - H(s, t - \delta))}{H(s, t) - (H(s, t) - H(s, t - \delta))} \\ &= \frac{2H(s, t) - \Delta H(s, t - \delta)}{H(s, t) - \Delta H(s, t - \delta)} \\ &= 2 + \frac{\Delta H(s, t - \delta)}{K - \Delta H(s, t - \delta)} \end{aligned}$$

Case 2. $t' \in (t, t + \delta]$:

$$\begin{aligned} \frac{C^{\text{BL}}(s, t')}{C^{\text{P}}(s, t')} &\leq \frac{H(s, t + \delta) + K}{K} \\ &= \frac{2K + (H(s, t + \delta) - H(s, t))}{K} \\ &= 2 + \frac{\Delta H(s, t)}{K} \end{aligned}$$

We next bound the holding cost $\Delta H(s, r)$ for any instants $s \leq r \leq T - \delta$:

$$\begin{aligned} H(s, r + \delta) - H(s, r) &= \int_s^{r+\delta} h(u) (\Lambda(r + \delta) - \Lambda(u)) du \\ &\quad - \int_s^r h(u) (\Lambda(r) - \Lambda(u)) du \\ &= \int_s^r h(u) (\Lambda(r + \delta) - \Lambda(r)) du \\ &\quad + \int_r^{r+\delta} h(u) (\Lambda(r + \delta) - \Lambda(u)) du \\ &\leq (\Lambda(r + \delta) - \Lambda(r)) \int_s^{r+\delta} h(u) du \\ &\leq (\Lambda(r + \delta) - \Lambda(r)) \int_0^T h(u) du \\ &\leq \lambda \delta \hat{h} T \end{aligned}$$

Let $\alpha = \frac{\lambda \hat{h} T}{K} \delta = \frac{\varepsilon}{1 + \varepsilon}$. The previous inequality shows that $\Delta H(s, r) \leq \alpha K$ for any instants s, r with $s \leq r \leq T - \delta$. Using this upper bound, we obtain in the two previous cases:

Case 1. Since $x \mapsto x/(K - x)$ is clearly a nondecreasing function of x on $[0, K)$ we have:

$$\begin{aligned} \frac{C^{\text{BL}}(s, t')}{C^{\text{P}}(s, t')} &\leq 2 + \frac{\alpha K}{K - \alpha K} = 2 + \frac{\alpha}{1 - \alpha} \\ &= 2 + \varepsilon \end{aligned}$$

Case 2. We have directly, using that $x/(1 + x) \leq x$ for any $x > 0$:

$$\begin{aligned} \frac{C^{\text{BL}}(s, t')}{C^{\text{P}}(s, t')} &\leq 2 + \frac{\alpha K}{K} = 2 + \alpha \\ &\leq 2 + \varepsilon \end{aligned}$$

Therefore in both cases the performance ratio is bounded by $2 + \varepsilon$. In particular, the previous inequalities are true for P an optimal policy and the result follows.

Appendix B. Proof of Theorem 3

Consider a feasible policy P for the production problem and two consecutive order time instants s and t of PB. Similarly to Lemma 1, we establish that the cost $C^{\text{P}}(s, t)$ incurred by P on time interval $(s, t]$ is at least $K + H_{\min}(s, t)$.

As already noticed, the holding cost of any feasible policy on $(s, T]$ is at least $H_{\min}(s, t)$. Thus the inequality clearly holds if P orders (i.e. starts a production) in the

time interval $(s, t]$. Conversely assume that P does not order on $(s, t]$. Let $u \in (s, t]$ be the end of the active period of PB in the time interval. We establish that the stock level in policy P is greater or equal to the stock level in policy PB at any point in time of $(s, t]$. That is, the balancing policy PB carries the minimum possible inventory for a policy that does not order inside the time interval $(s, t]$. First note that for any $z \in [u, t]$, we have $x^P(z) \geq x^{PB}(z)$ as P satisfies demand in $[u, t]$ without stockouts. In addition, since $x^P(s) \geq x_{min}(s) = x^{PB}(s)$, $x^P(u) \geq x^{PB}(u)$ and PB continuously produces in $[s, u]$, we also have $x^P(z) \geq x^{PB}(z)$ for all $z \in [s, u]$. As a result, the holding cost incurred by policy P on $(s, T]$ is at least $H(s, t) = K + H_{min}(s, t)$.

Thus for any policy P we have $C^P(s, t) \geq K + H_{min}(s, t)$ and a proof similar to the one of Theorem 1 leads to the same approximation guarantee of two.

Appendix C. Proof of Theorem 4

Consider a feasible policy P for the problem with shortages and let $s_1 < \dots < s_n$ be the sequence of orders found by the BB policy. We prove that the total cost C^{BB} incurred by BB on the time horizon is at most twice C^P , the total cost incurred by policy P . We use again the notation $C^P(s, t)$ to denote the total cost incurred by P over $(s, t]$, including backlogging costs.

We start by noticing that the BB policy places its first order at time s_1 , such that $B(0, s_1) = K > 0$: Hence $s_1 > 0$ and $\kappa_0^{BB} = 0$. For the last cycle, we distinguish between two cases, depending on whether $s_n < T$ or $s_n = T$. If $s_n < T$, the BB policy incurs a total cost of $H(s_n, T) \leq K$ on $(s_n, T]$. As policy P is feasible, either it orders on $(s_n, T]$ or it does not and then its inventory level at s_n is at least $\Lambda(s_n, T)$. Therefore it incurs on $(s_n, T]$ a cost of at least $C^P(s_n, T) \geq H(s_n, T) = C^{BB}(s_n, T)$. If $s_n = T$, the BB policy incurs a cost

$$C^{BB}(s_{n-1}, s_n) = K + \min_{u \in [s_{n-1}, s_n]} \{H(s_{n-1}, u) + B(u, T)\} \leq 2K$$

Note that since BB orders at T , the algorithm ensures that $H(s_{n-1}, T) > K$ by construction. Therefore, if P orders in $(s_{n-1}, T]$ it incurs an order cost of K , while if it does not it incurs a holding cost greater than K . Consequently we have $C^{BB}(s_{n-1}, T) \leq 2C^P(s_{n-1}, T)$.

Now let s_i and s_{i+1} be two consecutive order times of the BB policy, with $s_{i+1} < T$. The total cost incurred by

BB over $(s_i, s_{i+1}]$ is

$$C^{BB}(s_i, s_{i+1}) = \min_{u \in [s_i, s_{i+1}]} \{H(s_i, u) + B(u, s_{i+1})\} + K = 2K$$

On the other hand, either P orders in $(s_i, s_{i+1}]$ and incurs a cost of at least K or it does not and then there exists $u^P \in [s_i, t]$ such that $x^P(z) \geq 0$ for all $z \in [s_i, u^P]$ and $x^P(z) \leq 0$ for all $z \in (u^P, s_{i+1}]$. Hence P incurs a holding and backlogging cost of at least $\min_{u \in [s_i, s_{i+1}]} \{H(s_i, u) + B(u, s_{i+1})\} = K$. Therefore we have $C^{BB}(s_i, s_{i+1}) = 2K \leq 2C^P(s_i, s_{i+1})$.

The proof follows from Proposition 1.

Appendix D. Proof of Theorem 5

In this proof, we modify the notation introduced in §2 as follows: For $s < t$ and P a policy for the problem, let $C^P(s, t)$ be the sum of the fixed order cost incurred by P in $(s, t]$ plus the linear order cost and holding cost incurred by the units used to serve demands in $(s, t]$. For instance, assume that P orders at time $v \in (s, t]$ to satisfy demands in $[u, t]$ and at time $r < s$ to serve demands in (s, v) . Then we have:

$$C^P(s, t) = c(r)\Lambda(s, v) + c(v)\Lambda(v, t) + \int_r^s h(x)\Lambda(s, v)dx + \int_s^v h(x)\Lambda(x, v)dx + \int_v^t h(x)\Lambda(x, t)dx$$

Note that this is an extension of the original definition of $C^P(\cdot, \cdot)$ and thus all the previous proofs remain valid with this new definition.

We start by proving the result for the forward algorithm. Let P be a feasible policy for the problem with nonincreasing order costs. We again have $\kappa_0^{BL} = K(0) = \kappa_0^P$ for any feasible policy P . Now let $i < n$ and focus on the ordering cycle $(s_i, s_{i+1}]$. We bound the cost incurred by policy P over $(s_i, s_{i+1}]$ as follows:

Case 1. P places an order at time $u \in (s_i, s_{i+1}]$: It incurs an order cost $K(u) \geq K(s_{i+1})$.

Case 2. P does not order in $(s_i, s_{i+1}]$: It incurs a holding cost of at least $H(s_i, s_{i+1})$. Moreover, P orders the units used to serve the demands in $(s_i, s_{i+1}]$ at a previous point in time $s' < s_i$, incurring an additional linear order cost and holding cost of $(c(s') + \int_{s'}^s h(u)du) \Lambda(s_i, s_{i+1})$. According to equation 6, we thus have:

$$\begin{aligned}
C^P(s_i, s_{i+1}) &= \left(c(s') + \int_{s'}^s h(u) du \right) \Lambda(s_i, s_{i+1}) \\
&\quad + H(s_i, s_{i+1}) \\
&\geq c(s_i) \Lambda(s_i, s_{i+1}) + H(s_i, s_{i+1}) \\
&\geq K(s_{i+1})
\end{aligned}$$

Thus for any $i < n$ and ordering cycle $(s_i, s_{i+1}]$, we have:

$$C^{\text{BL}}(s_i, s_{i+1}) = 2K(s_{i+1}) \leq 2C^P(s_i, s_{i+1})$$

Finally, one can use similar arguments to prove that we have

$$C^{\text{BL}}(s_n, T) = c(s_n) \Lambda(s_n, T) + H(s_n, T) \leq C^P$$

and the cost incurred by the balancing policy over the entire planning horizon can be bounded as follows:

$$\begin{aligned}
C^{\text{BL}}(0, T) &= C^{\text{BL}}(0, t_0) + \sum_{i=1}^{n-1} C^{\text{BL}}(s_i, s_{i+1}) + C^{\text{BL}}(s_n, T) \\
&= C^{\text{BL}}(0, t_0) + \sum_{i=1}^{n-1} 2K(s_{i+1}) + C^{\text{BL}}(s_n, T) \\
&\leq C^P(0, t_0) + \sum_{i=1}^{n-1} 2C^P(s_i, s_{i+1}) + C^P(s_n, T) \\
&\leq 2C^P(0, T)
\end{aligned}$$

For the backward algorithm, the cost accounting for a cycle $(s, t]$ of the balancing policy is slightly modified since we now account for the order cost in period s and the holding cost over $[s, t)$. As a consequence, a cycle is denoted $[s, t)$ in the remaining of the proof.

Let P be a feasible policy for the problem with nondecreasing order costs. First, consider the case where the first order is place earlier than t_0 : $s_n < t_0$. By definition we have:

$$\begin{aligned}
C^{\text{BL}}(t_0, s_{n-1}) &= K(s_n) + c(s_n) \Lambda(t_0, s_{n-1}) \\
&\quad + H(s_n, s_{n-1}) - H(s_n, t_0) \\
&\leq 2K(s_n)
\end{aligned}$$

On the other hand, either P orders at some time instant $u \in [s_n, s_{n-1})$ and incurs a fixed order cost $K(u) \geq K(s_n)$, or it does not and the units used to serve demands in $(t_0, s_{n-1}]$ are ordered in some previous period $s' < s_n$.

Therefore we have:

$$\begin{aligned}
C^P(t_0, s_{n-1}) &= \left(c(s') + \int_{s'}^{s_n} h(u) du \right) \Lambda(t_0, s_{n-1}) \\
&\quad + H(s_n, s_{n-1}) - H(s_n, t_0) \\
&\leq c(s_n) \Lambda(t_0, s_{n-1}) \\
&\quad + H(s_n, s_{n-1}) - H(s_n, t_0) \\
&\leq K(s_n)
\end{aligned}$$

For $i \leq n$, we bound the cost incurred by policy P over the cycle $[s_i, s_{i-1})$ using similar arguments as the forward case: Either P places an order in $[s_i, s_{i-1})$ and incurs an order cost of at least $K(s_i)$ or it does not and thus incurs a linear order cost and holding cost of at least $K(s_i)$ to serve demands in $[s_i, s_{i-1})$. Therefore we have:

$$\begin{aligned}
C^{\text{BL}}(0, T) &= C^{\text{BL}}(0, s_n) + \sum_{i=1}^n C^{\text{BL}}(s_i, s_{i-1}) \\
&\leq C^P(0, s_n) + \sum_{i=1}^n 2K(s_i) \\
&\leq C^P(0, s_n) + 2 \sum_{i=1}^n C^P(s_i, s_{i-1}) \\
&\leq 2C^P(0, T)
\end{aligned}$$

and the proof follows.

References

- Aggarwal, A., & Park, J. (1993). Improved algorithms for economic Lot-Size problems. *Operations Research*, 41, 549–571.
- Axsäter, S. (1982). Worst case performance for lot sizing heuristics. *European Journal of Operational Research*, 9, 339–343.
- Bakker, M., Riezebos, J., & Teunter, R. H. (2012). Review of inventory systems with deterioration since 2001. *European Journal of Operational Research*, 221, 275–284.
- Barbosa, L., & Friedman, M. (1978). Deterministic inventory lot size models - a general root law. *Management Science*, 24, 819–826.
- Benkherouf, L., & Gilding, B. H. (2009). On a class of optimization problems for finite time horizon inventory models. *SIAM J. Control Optim.*, 48, 993–1030.
- Bitran, G., Magnanti, T., & Yanasse, H. (1984). Approximation methods for the uncapacitated dynamic lot size problem. *Management Science*, 30, 1121–1140.
- Donaldson, W. (1977). Inventory replenishment policy for a linear trend in demand – an analytical solution. *Operational Research Quarterly*, 28, 663–670.
- Goyal, S., & Giri, B. (2001). Recent trends in modeling of deteriorating inventory. *European Journal of Operational Research*, 134, 1–16.
- Goyal, S., & Giri, B. (2003). The production-inventory problem of a product with time varying demand, production and deterioration rates. *European Journal of Operational Research*, 147, 549–557.

- Hariga, M. A. (1994). The lot-sizing problem with continuous time-varying demand and shortages. *Journal of the Operational Research Society*, 45, 827–837.
- Hariga, M. A. (1996). Lot-sizing heuristics for continuous time-varying demand and shortages. *Computers & Ops Res.*, 23, 1211–1217.
- Harris, F. (1913). How many parts to make at once. *Factory, the Magazine of management*, 10, 135–136.
- Henery, R. (1979). Inventory replenishment policy for increasing demand. *Journal of the Operational Research Society*, 30, 611–617.
- Henery, R. (1990). Inventory replenishment policy for cyclic demand. *Journal of the Operational Research Society*, 41, 639–643.
- Van den Heuvel, W., & Wagelmans, A. (2010). Worst-case analysis for a general class of online lot-sizing heuristics. *Operations Research*, 58, 59–67.
- Kiefer, J. (1953). Sequential minimax search for a maximum. *Proc. Amer. Math. Soc.*, 4, 502 – 506.
- Resh, M., Friedman, M., & Barbosa, L. (1976). On a General Solution of the Deterministic Lot Size Problem with Time-Proportional Demand. *Operations research*, 24, 718–725.
- Silver, E. (1979). A simple inventory replenishment rule for a linear trend in demand. *Journal of the Operational Research Society*, 30, 71–75.
- Silver, E., & Meal, H. (1973). A heuristic for selecting lot size quantities for the case of a deterministic time-varying demand rate and discrete opportunities for replenishment. *Production and Inventory Management*, 2nd Quarter, 64–74.
- Teng, J., Ouyang, L., & Chen, L. (2007). A comparison between two pricing and lot-sizing models with partial backlogging and deteriorated items. *International Journal of Production Economics*, 105, 190 – 203.
- Wagner, H., & Whitin, T. (1958). Dynamic version of the economic lot size model. *Management Science*, 5, 89–96.
- Wilson, R. (1934). A scientific routine for stock control. *Harvard business review*, 13, 116–128.
- Zipkin, P. (2000). *Foundations of Inventory Management*. New York: McGraw-Hill.