ARTICLE IN PRESS
JID: EOR [m5G;July 29, 2018;10:55]

European Journal of Operational Research 000 (2018) 1–9

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

An exact static solution approach for the service parts end-of-life inventory problem

J. B. G. Frenk[a,*], S. Javadi[a], M. Pourakbar[b], S. O. Sezer[a]

[a] Faculty of Engineering and Natural Sciences, Sabancı University, 34956 Tuzla, Istanbul, Turkey
[b] Rotterdam School of Management, Erasmus University Rotterdam, 3000 DR Rotterdam, The Netherlands

ARTICLE INFO

Article history:
Received 13 July 2017
Accepted 22 June 2018
Available online xxx

Keywords:
Inventory
Service parts
End-of-life inventory
Poisson processes

ABSTRACT

This paper studies the spare parts end-of-life inventory problem that happens after the discontinuation of part production. A final ordering quantity is set such that the service process is sustained until all service obligations expire. Also, the price erosion of substitutable or new generation products over time makes it economically justifiable to consider switching to an alternative service policy for repair such as swapping the old product with a new one. This requires the joint optimization of the final order quantity and the time to switch from repair to an alternative service policy. To the best of our knowledge, the problem has not been optimally solved yet either in its static or dynamic formulation. In the current paper, we solve its static version as a bi-level optimization problem. We investigate the convexity of the objective function and give a computationally efficient algorithm to find an exact optimal solution up to any given numerical error level $\epsilon > 0$. We illustrate our approach on some numerical examples and compare our results with earlier works on this problem.

1. Introduction

With an ever increasing growth in technology development, obsolescence has become a major problem for many industries including military, consumer electronics, aerospace, etc. As an example in the consumer electronics sector, reportedly, HP suffered from a huge obsolescence cost because of end-of-life write-offs due to the short life cycles of PCs (Callioni, de Montgros, Slagmulder, Van Wassenhove, & Wright, 2005). The short life cycle of parts or products and early discontinuation of production have made the service part management very complicated. This brings certain issues for the Original Equipment Manufacturers (OEM). The main challenge in the service parts final phase inventory management is that, due to the discontinuation of production, the acquisition of parts is no longer guaranteed. Various tactical decisions can be made on how to deal with this challenge such as substituting another part for the obsolete one, obtaining the discontinued part from an after market manufacturer, redesigning the product or purchasing a sufficient volume of the obsolete part for the remaining life time; this is called a last-time or end-of-life buy and is one of the most popular policies in practice. In such a policy, right before the discontinuation of part production, a final order is placed with the manufacturer. Then this stock of service parts is used to repair/replace the defective items and meet the service requirements of customers. The central trade-off in this policy is how to optimize the inventory policy of service parts so that the risk of not meeting the service obligations is minimized while the chance of facing obsolete items and consequently writing off inventories is also controlled and kept at low levels.

In the consumer electronics market, we know that there exists a remarkable price erosion over time that could have significant implications for service parts inventory planning in the final phase. If the OEM plans to repair/replace all the defective items through a final order policy, then she has to plan in advance and purchase sufficiently many spare parts. Moreover, she needs to hold all the items in stock until the end of the product life cycle and provide a service to the consumer when a demand arrives. However, this process could be very costly. An alternative policy for this, considering the price erosion of these products, could be to offer consumers a new product at a discounted price or a voucher to purchase the next generation product when these products become sufficiently cheap. In this new setting, one should plan carefully the number of spare parts to place as the final order and also the timing of switching to such a more cost-effective alternative policy.

In this paper, we analyze the problem of determining an end-of-life inventory level and a time to switch to the alternative policy in order to minimize the total expected costs incurred during the final phase. We search for an optimal pair within the class of static

* Corresponding author.
E-mail addresses: frenk@sabanciuniv.edu (J.B.G. Frenk), sonyaj@sabanciuniv.edu (S. Javadi), mpourakbar@rsm.nl (M. Pourakbar), sezer@sabanciuniv.edu (S.O. Sezer).

https://doi.org/10.1016/j.ejor.2018.06.041

Please cite this article as: J.B.G. Frenk et al., An exact static solution approach for the service parts end-of-life inventory problem, European Journal of Operational Research (2018), https://doi.org/10.1016/j.ejor.2018.06.041

policies in which the switching time is determined initially at the beginning of the final phase. In this static formulation, the problem is a bi-level cost minimization problem. For each fixed switching time, we show that the expected operational cost function is convex in the order quantity. Hence, for a fixed switching time, the best order quantity can simply be found by the first order conditions. In the next step, we search for the best switching time. However, this problem does not exhibit useful properties for the sufficiency of the first order conditions. Therefore, we give a Lipschitz optimization procedure in order to determine an ϵ-optimal switching time. We leave the details to Sections 4 and 5.

The rest of the paper is organized as follows. Section 2 provides a brief literature review on the end-of-life inventory policies. In Section 3, we describe the problem formulation and introduce the cost components. In Sections 4 and 5, we analyze the objective function and investigate the convexity structure of the model. Using this structure we give an easily implementable algorithm to solve the problem numerically. In Section 6, we summarize the main dynamic heuristic developed in the recent paper Pourakbar, Frenk, and Dekker (2012) for this problem. Next, in Section 7 we illustrate our exact approach through some numerical experiments and compare our results with those produced by two main policies: (i) the policy of no switching and (ii) the heuristic of Pourakbar et al. (2012) described in Section 6. Finally, in Section 8 we give the concluding remarks.

2. Literature review

The service parts life-cycle spans through three main phases, namely, initial, normal and final phase. These phases correspond to the evolution of demand for service parts from the time that the product is introduced to the market until the last service contract or warranty expires. The majority of service parts inventory management studies relate to the normal phase where parts are still in production and could be replenished. In general, the final phase is the longest period within the life cycle of a service part. For instance, in the electronics industry this phase may last from four up to thirty years while the production of electronic appliances is normally terminated after less than two years (Teunter & Klein Haneveld, 2002). Despite its significance, the spare parts inventory management in the final phase has not been dealt with extensively in the literature. In the next few paragraphs we review some of the relevant studies.

Fortuin (1980, 1981) study the problem from a service level perspective. The author assumes an exponentially decreasing demand pattern and uses a normal approximation to derive expressions for several service levels. Then, a number of curves are derived that can be used to find the optimal order quantity to achieve a given service level. Van Kooten and Tan (2009) also look at service level considerations of a repairable service part in the final phase and use a transient Markovian model to find the optimal last time buy quantity. Another category of relevant studies aims at developing models to forecast the demand for spare parts during this phase. These models are then used to calculate order quantities; see for example Moore (1971), Ritchie and Wilcox (1977) and Hong, Koo, Lee, and Ahn (2008). The most relevant category of studies to the current paper consists of those papers which aim at finding the optimal final order quantity minimizing the expected future costs. As an example in this stream, Teunter and Fortuin (1999) find a near optimal solution for the final order quantity and further extend their work by introducing a dispose-down-to level policy which allows unused parts to be removed from stock before the end of the horizon. Similar problems can be found in Teunter and Fortuin (1998), Teunter and Klein Haneveld (1998, 2002), among others. More recently, Behfard, van der Heijden, Al Hanbali, and Zijm (2015) study a problem where next to the final order policy parts can also be extracted from a repair process. They develop a heuristic method to find a near-optimal last time buy quantity in the presence of an imperfect repair option of the failed parts that can be returned from the field.

The major operating characteristic of the inventory policies in these studies is that a final order is placed and all the future demand for spare parts are met either through the stock obtained by the final order or retrieved from the repair process. However, as argued earlier especially for consumer electronics, the price of the products may erode remarkably over time providing opportunities to offer alternative service policies for repair. Examples for such alternative policies are swapping the old defective product with a new one and offering customers vouchers or discounts on the purchase of a replacement product. This could possibly change the structure of the final phase inventory policy from setting the optimal final order quantity to jointly setting the ordering quantity and the time to switch to the alternative policy. A model in which these two decisions are considered simultaneously is first introduced by Pourakbar et al. (2012). The cited paper formulates the problem, highlights its importance, and gives practical dynamic heuristic methods for implementation. The choice of focusing on heuristic methods (rather than solving it) is understandable considering the difficulty of the problem; in its general dynamic formulation, the time to switch to the alternative policy is a stopping time of the point process of returned items and their conditions, and this makes the problem quite challenging. To the best of our knowledge, the optimal static and dynamic solutions of the problem have not been provided in the literature. Here, we give the static solution and this is the main contribution of the current paper. The dynamic solution, on the other hand, requires a separate and lengthy analysis, and is studied in a subsequent paper. Clearly, the static solution is much simpler to apply in practice. Also, it is appealing from a practical point of view as it allows a company to plan in advance for the alternative policy and for the scrapping of the leftover service parts inventories. Hence, it can easily be incorporated into the strategic plans of the company.

3. Description of the problem

In our optimization problem, the objective is to find an optimal (x, τ)-policy minimizing the expected total cost associated with the product until the end of its service life time. Here, x denotes the final order quantity and τ is the switching time from a *repair–replacement policy* to an eventually more cost-effective *alternative policy*. Both x and τ are static variables; their values are determined initially at time zero.

Following Pourakbar et al. (2012), we consider a model in which customers arrive according to a non-homogeneous Poisson process with a time-dependent intensity function λ. We let $T_n, n \in \mathbb{N}$ denote the arrival times of the customers and $R_n, n \in \mathbb{N}$ be a sequence of independent and identically distributed Bernoulli random variables indicating whether the arriving items are repairable. In particular,

$$R_n = \begin{cases} 1 & \text{if } n\text{th arriving item can be repaired} \\ 0 & \text{otherwise,} \end{cases}$$

and we have $\mathbb{P}(R_n = 1) = q$ and $\mathbb{P}(R_n = 0) = 1 - q$ for some $q \in (0, 1)$. Clearly, the counting process $N_0 = \{N_0(t) : t \geq 0\}$ of non repairable items is given by

$$N_0(t) = \sum_{n=1}^{\infty} (1 - R_n) 1_{\{T_n \leq t\}}, \quad t \geq 0,$$

and the counting process $N_1 = \{N_1(t) : t \geq 0\}$ of repairable items is represented by

$$N_1(t) = \sum_{n=1}^{\infty} R_n 1_{\{T_n \leq t\}}, \quad t \geq 0.$$

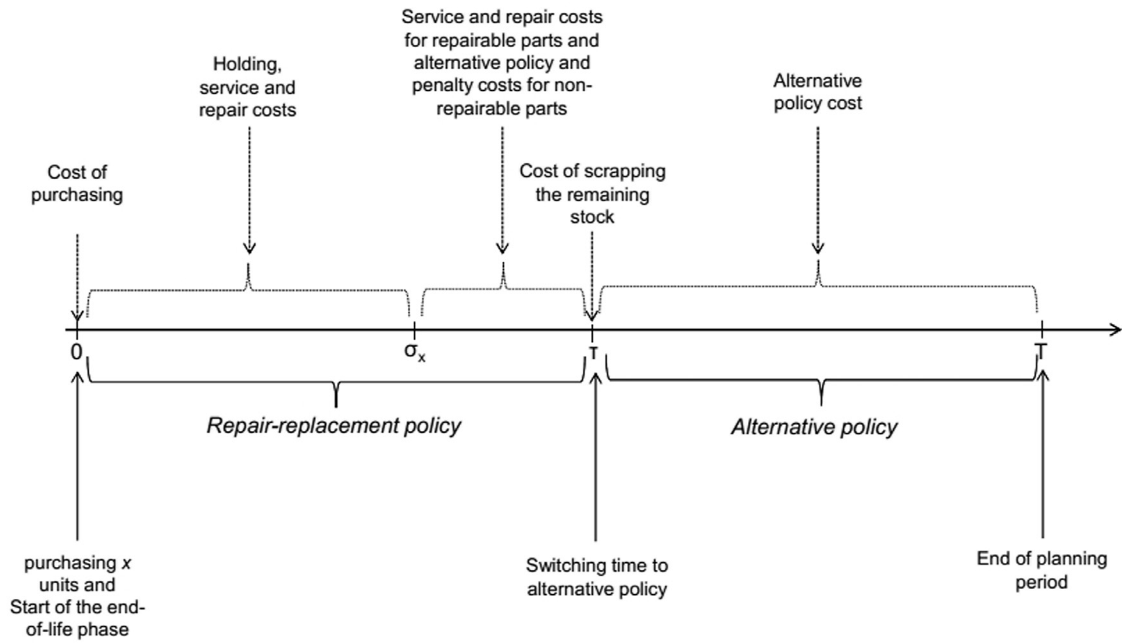

Fig. 1. Decisions and costs are illustrated over the time line [0, T] where σ_x denotes the inventory depletion time.

It follows from the independence of the random variables R_n, $n \in \mathbb{N}$ that the stochastic processes N_0 and N_1 are independent non-homogeneous Poisson processes with arrival rate intensities $\lambda_0 = (1-q)\lambda$ and $\lambda_1 = q\lambda$, respectively.

The cost of obtaining/producing the initial inventory is given by an increasing procurement cost function satisfying $c(0) = 0$ and $\lim_{x\to\infty} c(x) = \infty$. In this setting, we have an inventory holding cost which increases according to the cost rate $h > 0$ per item per unit of time. All costs are discounted back to time 0 with a discount rate $\delta > 0$.

Starting with x units, the OEM uses the repair–replacement policy until some time $\tau \le T$, where T denotes the time at which all the service obligations of the OEM expire. Under the repair–replacement policy, if an arriving item is repairable, it is repaired at some repair cost c_{re} plus some service cost c_{se}. Here, c_{se} captures the labor and other costs associated with serving a customer under the repair–replacement policy. If the item is non-repairable and the inventory level is non-zero, the item is replaced with a spare one from the inventory at the service cost c_{se} only. However, if a customer arrives with a non-repairable item and no spare part is available, then in addition to the cost of the alternative policy, which is given by a function c_a, the OEM incurs a higher service cost because of an unscheduled use of the alternative policy. This penalty cost is represented by a function p satisfying the natural assumption $p(u) > c_{se}$ for all $u \in [0, T]$.

At time τ, the OEM completely switches to the alternative policy and discards the existing inventory (if there is any) at a scrapping cost of $c_{scr} \ge 0$ per item. If a service request arrives at time $u \ge \tau$, then only the alternative policy is used at the cost $c_a(u)$.

Note that both p and c_a are non-increasing functions. That is, the alternative method as well as the penalty associated with an unscheduled use of it become cheaper over time because of the price erosion of the product. We take the natural assumption that the functions p and c_a are piecewise smooth; that is, they are differentiable on [0, T] except possibly at finitely many points, at which they can have (downward) discontinuities as well. For ease of notation, we denote the total cost of applying the alternative policy at time u before the switching time τ as $c_{ap}(u) := c_a(u) + p(u), u \ge 0$. Also, due to discarding the inventory at time τ at cost $c_{scr} \ge 0$ per unit and not keeping this inventory after time τ it is natural to assume that $h - \delta c_{csr} \ge 0$. Otherwise, one would keep an item in the inventory indefinitely rather than scrapping.

The timing of the decisions and the related costs are summarized graphically over a time line in Fig. 1.

4. The objective function and its properties

The total cost of the OEM is the sum of its procurement cost and expected operation costs. The procurement cost is given by the function $c(x)$ as indicated in Section 3 above. The expected operation costs, on the other hand, consist of

(i) the expected inventory holding costs

$$h\mathbb{E}\left(\int_0^\tau e^{-\delta u}(x - N_0(u))^+ du\right) = h\int_0^\tau e^{-\delta u}\mathbb{E}[(x - N_0(u))^+]du$$

where the equality is due to Fubini's theorem;

(ii) the expected repair cost

$$\mathbb{E}\left(\int_0^\tau e^{-\delta u} c_{re}\, dN_1(u)\right) = c_{re}\int_0^\tau e^{-\delta u}\lambda_1(u)du$$

in which the equality follows after noting that the process $t \mapsto \int_{(0,t]} e^{-\delta u} dN_1(u) - \int_0^t e^{-\delta u}\lambda_1(u)du$ is a martingale (see Çınlar, 2011, Chapter 5.6 for Poisson processes and their martingales);

(iii) the expected service cost

$$\mathbb{E}\Big(\underbrace{\int_0^\tau e^{-\delta u} c_{se}\, dN_1(u)}_{\text{service cost for repairable items}} + \underbrace{\int_0^{\tau\wedge\sigma_x} e^{-\delta u} c_{se}\, dN_0(u)}_{\text{service cost for non-repairable items}}\Big)$$

$$= c_{se}\int_0^\tau e^{-\delta u}\lambda_1(u)du + c_{se}\mathbb{E}\left(\int_0^{\tau\wedge\sigma_x} e^{-\delta u}\lambda_0(u)du\right)$$

where $\sigma_x := \inf\{t > 0 : N_0(t) \ge x\}$ denotes the arrival time of the xth non-repairable item, and the first equality follows from the martingale property of $t \mapsto \int_{(0,t]} e^{-\delta u} dN_1(u) - \int_0^t e^{-\delta u}\lambda_1(u)du$ and Doob's stopping theorem applied to the martingale $t \mapsto \int_{(0,t]} e^{-\delta u} dN_0(u) - \int_0^t e^{-\delta u}\lambda_0(u)du$;

(iv) the expected total cost of using the alternative policy

$$\mathbb{E}\Big(\underbrace{\int_{\tau\wedge\sigma_x}^{\tau} e^{-\delta u}c_{ap}(u)dN_0(u)}_{\text{cost of the alternative policy before}\tau}\Big)+\mathbb{E}\Big(\underbrace{\int_{\tau}^{T} e^{-\delta u}c_a(u)dN(u)}_{\text{cost of the alternative policy after}\tau}\Big)$$

$$=\mathbb{E}\left(\int_{\tau\wedge\sigma_x}^{\tau} c_{ap}(u)e^{-\delta u}\lambda_0(u)du\right)+\int_{\tau}^{T} c_a(u)e^{-\delta u}\lambda(u)du$$

which is again due to the martingale property of the compensated Poisson integrals and Doob's stopping theorem;

(v) the expected scrapping cost

$$c_{scr}e^{-\delta\tau}\mathbb{E}((x-N_0(\tau))^+).$$

When we collect all the terms in (i)–(v) above, we obtain the total expected operation costs as

$$C(x,\tau)=\begin{cases} h\int_0^{\tau} e^{-\delta u}\mathbb{E}[(x-N_0(u))^+]du+c_{re}\int_0^{\tau} e^{-\delta u}\lambda_1(u)du\\ +c_{se}\int_0^{\tau} e^{-\delta u}\lambda_1(u)du+c_{se}\mathbb{E}\big(\int_0^{\tau\wedge\sigma_x} e^{-\delta u}\lambda_0(u)du\big)\\ +\mathbb{E}\big(\int_{\tau\wedge\sigma_x}^{\tau} c_{ap}(u)e^{-\delta u}\lambda_0(u)du\big)+\int_{\tau}^{T} c_a(u)e^{-\delta u}\lambda(u)du\\ +c_{scr}e^{-\delta\tau}\mathbb{E}((x-N_0(\tau))^+).\end{cases} \tag{1}$$

It follows from the chain rule for the process $t\mapsto e^{-\delta t}(x-N_0(t))$ that

$$\begin{aligned} e^{-\delta\tau}(x-N_0(\tau))^+ &= e^{-\delta(\tau\wedge\sigma_x)}(x-N_0(\tau\wedge\sigma_x))\\ &= x-\int_0^{\tau\wedge\sigma_x}\delta e^{-\delta u}(x-N_0(u))du-\int_0^{\tau\wedge\sigma_x} e^{-\delta u}dN_0(u)\\ &= x-\int_0^{\tau}\delta e^{-\delta u}(x-N_0(u))^+du-\int_0^{\tau\wedge\sigma_x} e^{-\delta u}dN_0(u), \end{aligned}\tag{2}$$

and by Fubini's theorem and Doob's stopping theorem we have

$$\begin{aligned} \mathbb{E}(e^{-\delta\tau}(x-N_0(\tau))^+) &= x-\mathbb{E}\left(\int_0^{\tau}\delta e^{-\delta u}(x-N_0(u))^+du\right)\\ &\quad -\mathbb{E}\left(\int_0^{\tau\wedge\sigma_x} e^{-\delta u}dN_0(u)\right)\\ &= x-\delta\int_0^{\tau} e^{-\delta u}\mathbb{E}(x-N_0(u))^+du\\ &\quad -\mathbb{E}\left(\int_0^{\tau\wedge\sigma_x} e^{-\delta u}\lambda_0(u)\,du\right). \end{aligned}\tag{3}$$

When we replace the last expectation in (1) with the expression in (3), we obtain after some simple re-arrangement of the terms the following alternative representation of the total expected operation costs

$$C(x,\tau)=\begin{cases} c_{scr}x+\mathbb{E}\big(\int_0^{\tau\wedge\sigma_x} e^{-\delta u}\lambda_0(u)[c_{se}-c_{scr}-c_{ap}(u)]du\big)\\ +\int_0^{\tau} e^{-\delta u}\lambda(u)[q(c_{se}+c_{re}-c_a(u))+(1-q)p(u)]du\\ +(h-\delta c_{scr})\int_0^{\tau} e^{-\delta u}\mathbb{E}((x-N_0(u))^+)du\\ +\int_0^{T} e^{-\delta u}c_a(u)\lambda(u)du.\end{cases}\tag{4}$$

Clearly by relation (4) we have

$$C(0,\tau)=\int_0^{\tau} e^{-\delta u}\lambda(u)[q(c_{se}+c_{re}-c_a(u))+(1-q)p(u)]du + \int_0^{T} e^{-\delta u}c_a(u)\lambda(u)du. \tag{5}$$

Since

$$\mathbb{E}\left(\int_0^{\tau\wedge\sigma_x} e^{-\delta u}\lambda_0(u)[c_{se}-c_{scr}-c_{ap}(u)]du\right) = \int_0^{\tau} e^{-\delta u}\lambda_0(u)[c_{se}-c_{scr}-c_{ap}(u)]\mathbb{P}(N_0(u)\le x-1)du,$$

we also deduce from relation (4) that the first order difference operator of the objective function with respect to x given by $\Delta_x C(x,\tau):=C(x+1,\tau)-C(x,\tau)$ equals

$$\Delta_x C(x,\tau)=\begin{cases} c_{scr}+\int_0^{\tau} e^{-\delta u}\lambda_0(u)[c_{se}-c_{scr}-c_{ap}(u)]\mathbb{P}(N_0(u)=x)du\\ +(h-\delta c_{scr})\int_0^{\tau} e^{-\delta u}\mathbb{P}(N_0(u)\le x)du,\end{cases}\tag{6}$$

for $x\in\mathbb{Z}_+$.

We now need to solve the optimization problem

$$\upsilon(P):=\inf_{x\in\mathbb{Z}_+,\tau\le T}\{c(x)+C(x,\tau)\}. \tag{P}$$

Since $c(0)=0$ it is obvious that $\upsilon(P)\le C(0,0)=\int_0^T e^{-\delta u}c_a(u)\lambda(u)\,du<+\infty$. Hence for

$$x_U:=\min\{x\in\mathbb{Z}_+ : c(x)>C(0,0)\}\ <+\infty$$

we obtain for every $x\ge x_U$ that

$$c(x)+C(x,\tau)\ge c(x)>C(0,0).$$

That is, we can restrict our search for an optimal x to $x=0,\dots,x_U-1$ and so

$$\upsilon(P)=\inf_{\tau\le T}\Phi(\tau) \tag{7}$$

with

$$\Phi(\tau):=\inf_{x\in\{0,\dots,x_U-1\}}\{c(x)+C(x,\tau)\} \tag{P_τ}$$

for every $\tau\in[0,T]$.

Below, we use a Lipschitz optimization procedure over the continuous variable τ, and thus we need to know how to identify an optimal order quantity for each given τ. Hence, it is important to determine for a given $\tau\le T$ under which conditions the function $x\mapsto C(x,\tau)$ is discrete convex. Clearly, $x\mapsto C(x,\tau)$ is discrete convex if and only if $x\mapsto\Delta_x C(x,\tau)$ is non-decreasing.

Before giving a simplified expression for $\Delta_x C(x,\tau)$ based on relation (4), we mention the following result for non-homogeneous Poisson processes.

Lemma 1. *Let N be a non-homogeneous Poisson process with arrival rate function β, and ψ be a piecewise differentiable function with finitely many points $\ell_1,\dots,\ell_m$ of lack-of-differentiability. Also let $\Delta\psi(\ell_i):=\psi(\ell_i)-\lim_{u\nearrow\ell_i}\psi(u)$ for $i\le m$. Then, for every $x\in\mathbb{Z}_+$ and $\tau\le T$ we have*

$$\begin{aligned} &\int_0^{\tau}\psi(u)\beta(u)\mathbb{P}(N(u)=x)du\\ &=\int_0^{\tau}\psi'(u)\mathbb{P}(N(u)\le x)du+\psi(0)-\psi(\tau)\,\mathbb{P}(N(\tau)\le x)\\ &\quad+\sum_{i\le m,\,\ell_i\le\tau}\Delta\psi(\ell_i)\,\mathbb{P}(N(\ell_i)\le x) \end{aligned}\tag{8}$$

where the integral $\int_0^{\tau}\psi'(u)\mathbb{P}(N(u)\le x)du$ denotes the sum

$$\sum_{i=0}^{m+1}\int_{\ell_i}^{\ell_{i+1}}\psi'(u)\mathbb{P}(N(u)\le x)1_{\{u\le\tau\}}du$$

with $\ell_0=0$ and $\ell_{m+1}=T$.

Proof. It is easy to see for a non-homogeneous Poisson process with an intensity function β that for every $k\in\mathbb{Z}_+$ the function $\varphi(u):=\mathbb{P}(N(u)\le k)=\sum_{j=0}^{k}e^{-\int_0^u\beta(s)ds}\frac{(\int_0^u\beta(s)ds)^j}{j!}$, for $u\ge0$, is differentiable and satisfies

$$\varphi'(u)=-\beta(u)\,\mathbb{P}(N(u)=k)$$

with the initial condition $\varphi(0)=\mathbb{P}(N(0)\le k)=1$. Applying now the chain rule gives

$$\psi(\tau)\varphi(\tau) - \psi(0) = \int_0^\tau \psi'(u)\varphi(u)\,du + \int_0^\tau \psi(u)\varphi'(u)\,du + \sum_{i\le m,\,\ell_i\le\tau} \Delta\psi(\ell_i)\varphi(\ell_i)$$
$$= \int_0^\tau \psi'(u)\varphi(u)\,du - \int_0^\tau \psi(u)\beta(u)\mathbb{P}(N(u)=k)\,du + \sum_{i\le m,\,\ell_i\le\tau} \Delta\psi(\ell_i)\,\mathbb{P}(N(\ell_i)\le k)$$

from which (8) follows after re-arranging the terms. □

Using Lemma 1 and relation (4) the next result follows easily. Below, the function c_{ap} will replace the function ψ in Lemma 1. Hence, for notational consistency, we let $\ell_1, \dots, \ell_m$ be the points of lack-of-differentiability of the function c_{ap}. Note that the difference $\Delta c_{ap}(\ell_i) := c_{ap}(\ell_i) - \lim_{u\nearrow\ell_i} c_{ap}(u)$ is strictly negative if there is a jump at the point ℓ_i.

Lemma 2. *For every $x\in\mathbb{Z}_+$ and $\tau\le T$ we have*

$$\Delta_x C(x,\tau) = \begin{cases} c_{se} - c_{ap}(0) + e^{-\delta\tau}(c_{ap}(\tau) + c_{scr} - c_{se})\mathbb{P}(N_0(\tau)\le x) \\ + \int_0^\tau e^{-\delta u}[h - c'_{ap}(u) + \delta(c_{ap}(u) - c_{se})]\mathbb{P}(N_0(u)\le x)du \\ - \sum_{i\le m,\,\ell_i\le\tau} e^{-\delta\ell_i}\Delta c_{ap}(\ell_i)\mathbb{P}(N_0(\ell_i)\le x). \end{cases} \quad (9)$$

Proof. Applying Lemma 1 with $\psi(u) = e^{-\delta u}[c_{se} - c_{scr} - c_{ap}(u)]$ gives

$$\begin{aligned} &\int_0^\tau e^{-\delta u}\lambda_0(u)[c_{se} - c_{scr} - c_{ap}(u)]\mathbb{P}(N_0(u)=x)du \\ &= -\int_0^\tau e^{-\delta u}c'_{ap}(u)\mathbb{P}(N_0(u)\le x)du \\ &\quad - \int_0^\tau \delta e^{-\delta u}[c_{se} - c_{scr} - c_{ap}(u)]\mathbb{P}(N_0(u)\le x)du \\ &\quad + [c_{se} - c_{scr} - c_{ap}(0)] \\ &\quad - e^{-\delta\tau}[c_{se} - c_{scr} - c_{ap}(\tau)]\mathbb{P}(N_0(\tau)\le x) \\ &\quad - \sum_{i\le m,\,\ell_i\le\tau} e^{-\delta\ell_i}\Delta c_{ap}(\ell_i)\mathbb{P}(N_0(\ell_i)\le x), \end{aligned} \quad (10)$$

and finally using (10) in (6) yields the desired result after straightforward simplifications. □

An immediate consequence of Lemma 2 is given by Lemma 3 below. The lemma establishes the convexity of the expected operation cost under the assumption that $c_{ap}(u) \ge c_{se} - c_{scr}$ for every $0\le u\le T$. Note that this condition is immediately satisfied under the natural sufficient condition $c_{scr}\ge 0$ (recall that $p(u) > c_{se}$ for all $u\le T$).

Lemma 3. *If $c_{ap}(u)\ge c_{se} - c_{scr}$ for every $u\le T$, then for every τ the function $x\mapsto C(x,\tau)$ is discrete convex and*

$$\lim_{x\uparrow\infty} \Delta_x C(x,\tau) = c_{scr}e^{-\delta\tau} + \delta^{-1}h(1 - e^{-\delta\tau}) > 0. \quad (11)$$

Proof. As $c_{ap}(\tau)\ge c_{se} - c_{scr}$ we obtain that the function

$$x\mapsto e^{-\delta\tau}(c_{ap}(\tau) + c_{scr} - c_{se})\mathbb{P}(N_0(\tau)\le x)$$

is non-decreasing in x. Moreover since (i) $h - \delta c_{scr}\ge 0$, (ii) c_{ap} non-increasing (recall that both c_a and p are non-increasing), and (iii) $c_{ap}(\tau)\ge c_{se} - c_{scr}$, we obtain for every $u\le\tau$

$$\begin{aligned} h - c'_{ap}(u) + \delta(c_{ap}(u) - c_{se}) &\ge h - c'_{ap}(u) + \delta(c_{ap}(\tau) - c_{se}) \\ &\ge (h - \delta c_{scr}) - c'_{ap}(u) \ge 0. \end{aligned}$$

This shows that the integral term in (9) is non-decreasing in x. Finally, because c_{ap} is non-increasing, the discrete sum in (9) is also non-decreasing in x. As a result, all the components are non-decreasing in x, and so it follows that the function $x\to\Delta_x C(x,\tau)$ is non-decreasing showing $x\mapsto C(x,\tau)$ is discrete convex. To establish the limit result, we observe that $\mathbb{P}(N_0(u)=x)$ converges to 0 as $x\to\infty$ pointwise for $u\le T$. This implies by the bounded convergence theorem that the first integral in (6) vanishes as x goes to ∞. For $u\le T$, we also obtain $\mathbb{P}(N_0(u)\le x)\nearrow 1$ as $x\to+\infty$, and again by monotone convergence theorem, the second integral in (6) converges to $\int_0^\tau e^{-\delta u}du = \frac{1}{\delta}(1 - e^{-\delta\tau})$. Then the limit result in (11) follows after a straightforward simplification. □

If the function c is also discrete convex then under the assumptions of Lemma 3 it follows that an optimal order quantity $x(\tau)$ of the optimization problem (P_τ) is given by

$$x(\tau) = \min\{x\in\mathbb{Z}_+ : c(x+1) - c(x) + \Delta_x C(x,\tau)\ge 0\}. \quad (12)$$

To compute $\Delta_x C(x,\tau)$ for every $x\in\mathbb{Z}_+$ we observe by relation (9) and introducing

$$\Lambda_0(s) := (1-q)\Lambda(s) = (1-q)\int_0^s \lambda(u)du, \quad s\ge 0,$$

that

$$\Delta_x C(0,\tau) = \begin{cases} c_{se} - c_{ap}(0) + e^{-\delta\tau}[c_{ap}(\tau) + c_{scr} - c_{se}]e^{-\Lambda_0(\tau)} \\ + \int_0^\tau e^{-\delta u}[h - c'_{ap}(u) + \delta(c_{ap}(u) - c_{se})]e^{-\Lambda_0(u)}du \\ - \sum_{i\le m,\,\ell_i\le\tau} e^{-\delta\ell_i}\Delta c_{ap}(\ell_i)e^{-\Lambda_0(\ell_i)} \end{cases} \quad (13)$$

and for $x\in\mathbb{Z}_+$

$$\Delta_x^{(2)}C(x,\tau) := \Delta_x C(x+1,\tau) - \Delta_x C(x,\tau)$$
$$= \begin{cases} e^{-\delta\tau}(c_{ap}(\tau) + c_{scr} - c_{se})e^{-\Lambda_0(\tau)}\frac{\Lambda_0(\tau)^{x+1}}{(x+1)!} \\ + \int_0^\tau e^{-\delta u}[h - c'_{ap}(u) + \delta(c_{ap}(u) - c_{se})] \\ \times e^{-\Lambda_0(u)}\frac{\Lambda_0(u)^{x+1}}{(x+1)!}du \\ - \sum_{i\le m,\,\ell_i\le\tau} e^{-\delta\ell_i}\Delta c_{ap}(\ell_i)e^{-\Lambda_0(\ell_i)}\frac{\Lambda_0(\ell_i)^{x+1}}{(x+1)!}. \end{cases} \quad (14)$$

Hence for a given value of τ, the optimal order quantity $x(\tau)$ and the solution of optimization problem (P_τ) can be computed, due to the convexity properties, by evaluating the increments in (13) and (14).

5. Optimization over τ-space

To apply now in an efficient way a Lipschitz optimization over the decision variable τ we construct a finite increasing sequence of points $0 = \tau_1 < \cdots < \tau_N = T$ such that

$$\inf_{\tau\in[\tau_n,\tau_{n+1}]} \Phi(\tau) \ge \Phi(\tau_n) - \epsilon$$

for every $n\le N$ for a given $\epsilon > 0$. Before constructing this sequence we first introduce for any function $f:[0,T]\to\mathbb{R}$ its so-called sup-norm given by

$$\| f \|_\infty := \sup_{0\le u\le T} | f(u) | .$$

Also introduce the functions $f_i:[0,T]\to\mathbb{R},\ i=1,2$ given by

$$\begin{aligned} f_1(u) &= e^{-\delta u}(c_{se} - c_{scr} - c_{ap}(u)), \\ f_2(u) &= e^{-\delta u}[q(c_{se} + c_{re} - c_a(u)) + (1-q)p(u)]. \end{aligned} \quad (15)$$

When $c_{ap}(u)\ge c_{se} - c_{scr}$ for every $u\le T$, the function f_1 is non-positive and non-decreasing.

Lemma 4. *If $c_{ap}(u) \geq c_{se} - c_{scr}$ for every $u \leq T$, then*

$$\Phi(\tau+s) - \Phi(\tau) \geq ((1-q)f_1(\tau) - \| f_2 \|_\infty)(\Lambda(\tau+s) - \Lambda(\tau)) \tag{16}$$

for any $\tau, s > 0$.

Proof. The order quantity $x(\tau+s)$ solving $\Phi(\tau+s)$ is also a feasible solution for $\Phi(\tau)$ (see (P_τ)). Hence, using the expected cost function $C(x, \tau)$ in (4), we have

$$\begin{aligned}\Phi(\tau+s) - \Phi(\tau) &\geq C(x(\tau+s), \tau+s) - C(x(\tau+s), \tau) \\ &= \begin{cases} \mathbb{E}\left(\int_{\tau\wedge\sigma_{x(\tau+s)}}^{(\tau+s)\wedge\sigma_{x(\tau+s)}} \lambda_0(u) f_1(u) du\right) \\ + \int_\tau^{\tau+s} \lambda(u) f_2(u) du \\ + (h - \delta c_{scr}) \int_\tau^{\tau+s} e^{-\delta u} \mathbb{E}(x(\tau+s) - N_0(u))^+ du \end{cases} \\ &\geq \mathbb{E}\left(\int_{\tau\wedge\sigma_{x(\tau+s)}}^{(\tau+s)\wedge\sigma_{x(\tau+s)}} \lambda_0(u) f_1(u) du\right) \\ &\quad + \int_\tau^{\tau+s} \lambda(u) f_2(u) du \end{aligned} \tag{17}$$

with $f_i, i = 1, 2,$ listed in (15). Since by our assumptions the function f_1 is non-positive and non-decreasing, we obtain

$$\begin{aligned}\mathbb{E}\left(\int_{\tau\wedge\sigma_{x(\tau+s)}}^{(\tau+s)\wedge\sigma_{x(\tau+s)}} \lambda_0(u) f_1(u) du\right) &\geq f_1(\tau)\,\mathbb{E}\left(\int_{\tau\wedge\sigma_{x(\tau+s)}}^{(\tau+s)\wedge\sigma_{x(\tau+s)}} \lambda_0(u) du\right) \\ &= f_1(\tau)\,\mathbb{E}(N_0((\tau+s)\wedge\sigma_{x(\tau+s)}) \\ &\quad - N_0(\tau\wedge\sigma_{x(\tau+s)})) \\ &\geq f_1(\tau)\,\mathbb{E}(N_0(\tau+s) - N_0(\tau)) \\ &= (1-q) f_1(\tau)[\Lambda(\tau+s) - \Lambda(\tau)). \end{aligned}$$

Similarly, we also have

$$\int_\tau^{\tau+s} \lambda(u) f_2(u) du \geq -\| f_2 \|_\infty (\Lambda(\tau+s) - \Lambda(\tau)).$$

Applying these inequalities in (17) we obtain the desired result. □

For $0 \leq s \leq \tau \leq T$, introduce the function

$$I(\tau, s) := \big((1-q)f_1(s) - \| f_2 \|_\infty\big)(\Lambda(\tau) - \Lambda(s)).$$

Clearly for $\tau = s$, we have $I(s, s) = 0$, and this function is non-increasing and non-positive for $\tau > s$. We now construct for a given $\epsilon > 0$ a finite sequence $(\tau_n)_{n=1}^N \subseteq [0, T]$ as follows. We let $\tau_1 = 0$, and define for $n \in \mathbb{N}$ iteratively

$$\tau_{n+1} = \begin{cases} \min\{\tau_n \leq \tau \leq T : I(\tau, \tau_n) \geq -\varepsilon\} & \text{if the set is nonempty} \\ T & \text{if the set is empty.} \end{cases} \tag{18}$$

In this procedure, N is the first index for which we have $\tau_N = T$. For each $n \leq N$, we solve the optimization problem (P_{τ_n}) and then we select τ_n with the smallest $\Phi(\tau_n)$ value. That is, we solve the problem

$$\upsilon_\epsilon(P) := \min_{n \leq N} \Phi(\tau_n). \tag{19}$$

By Lemma 4, the selected switching time τ_n and the order quantity $x(\tau_n)$ form an ϵ-optimal solution for our end-of-life inventory problem.

We list these steps as an algorithm below.

Algorithm to compute the optimal solution and the optimal objective value

1. Construct a finite increasing sequence of points $0 \leq \tau_n \leq T$ using (18).
2. For $n \leq N$,
   (a) compute $C(0, \tau_n)$ and $\Delta_x C(0, \tau_n)$ (see relations (5) and (13)) and check whether the first order condition in relation (12) is satisfied. If so, for given τ_n, set $x(\tau_n) = 0$. Otherwise go to step (b).
   (b) For $x = 1$ until the first order condition in relation (12) is satisfied, compute the second difference $\Delta_x^{(2)} C(x-1, \tau_n)$ (see relation (14)) and set

   $$C(x, \tau_n) = C(x-1, \tau_n) + \Delta_x C(x-1, \tau_n)$$

   with

   $$\Delta_x C(x, \tau_n) = \Delta_x C(x-1, \tau_n) + \Delta_x^{(2)} C(x-1, \tau_n).$$

   Set $x(\tau_n) = x$.
3. Search for τ_n with the smallest $\Phi(\tau_n)$ value. For that value of n, set $(\tau_n, x(\tau_n))$ as the ϵ-optimal policy.

6. Dynamic programming approach presented in Pourakbar et al. (2012)

In Pourakbar et al. (2012) a backward discrete-time dynamic program is developed to find a static order quantity and a dynamic switching rule. First, the time line is divided into equally spaced points $\{0, 1, 2, \ldots, T-1, T\}$. Then, starting from time T and moving backward, a DP equation is solved at every time point to find whether it is optimal to retain the repair–replacement policy (and postpone the decision of switching to an alternative policy) until the next period or to start serving customers using the alternative policy at the beginning of that period.

When we have x many units in the inventory at time $k \leq T$, the incremental cost of switching to the alternative policy (discounted to time k) is given by

$$A(x, k) := c_{scr} x + \int_k^T e^{-\delta(u-k)} \lambda(u) c_a(u) du.$$

The operating cost of using the repair–replacement policy for one more period is

$$B(x, k) = \begin{cases} h \int_k^{k+1} e^{-\delta(u-k)} \mathbb{E}[(x - (N_0(u) - N_0(k)))^+] du \\ + \int_k^{k+1} e^{-\delta(u-k)} \lambda(u) q (c_{se} + c_{re}) du \\ + \int_k^{k+1} e^{-\delta(u-k)} \lambda_0(u) c_{ap}(u) du \\ + \mathbb{E}\left(\int_k^{(k+1)\wedge\hat\sigma_x} e^{-\delta(u-k)} \lambda_0(u)(c_{se} - c_{ap}(u)) du\right) \end{cases}$$

where $\hat\sigma_x$ denotes the arrival of the xth non-repairable item after time k.

Then, if we let $V(x, k)$ be the value function of this DP program at time $k \leq T$, we have the recursive relation

$$\begin{aligned} V(x, k) = \min \Big\{ & A(x, k), \ B(x, k) \\ & + \mathbb{E} V\big((x - (N_0(k+1) - N_0(k)))^+, k+1\big) \Big\}, \end{aligned} \tag{20}$$

for $k \leq T - 1$, and we have the boundary condition $V(T, x) = c_{scr} x$ at $k = T$. At $k = 0$, we solve the problem

$$\min_x \ c(x) + V(x, 0)$$

to find the initial order quantity. In the implementation, we start with this order quantity and switch to the alternative policy at the first time k when $V(X_k, k) = A(X_k, k)$ with X_k denoting the inventory level at that time.

Note that the DP approach of Pourakbar et al. (2012) actually solves the optimal stopping problem

$$\inf_{x \in \mathbb{Z}_+, \, \tau \in \mathbb{F}_1} c(x) + C(x, \tau) \tag{21}$$

where $\mathbb{F}_1$ is the set of all stopping times (of the observation filtration) taking values in the set $\{0, 1, \ldots, T\}$. Hence, it can be regarded as a relaxation of the original dynamic problem in which the stopping time can take any values on the set [0, T].

The heuristic dynamic approach of Pourakbar et al. (2012) summarized above has a number of drawbacks compared to our static solution. First, solving the recursive equation in (20) is computationally quite expensive, whereas our static approach simply relies on the straightforward evaluation of the difference operators in (13) and (14). Secondly, as any other dynamic policy, the switching time is now a random variable; it depends on the realization of the inventory process over time and so it is not known at time zero. This may create problems when the OEM needs to make strategic plans in advance (for example, at time zero). Thirdly, unlike its continuous nature in the original formulation of the problem, in the DP heuristic of Pourakbar et al. (2012) the switching moment is discrete and can take the values $0, 1, 2, \ldots, T$ only.

Despite its drawbacks, this heuristic (as a relaxation) has the major advantage that it approximates the value function of the true optimal stopping problem. Hence, although it seems not possible to make a generalization, in many cases we expect it to yield lower expected cost values than our static solution. We verify this numerically in the next section where we report the cost values for various problem instances. Yet, the relative improvement of this dynamic heuristic method relative to the static solution is within 1% to 2% in general (with an overall average of 1.35%).

7. Numerical experiments

In this numerical section, following Pourakbar et al. (2012), we consider a setting in which the cost functions are as follows: $c(x) = c_p x$, $c_a(u) = c_a(0)e^{-\gamma u}$, for some $c_a(0) > c_{se}$, and $\gamma > 0$. We also have $p(u) = p$, for some $p > c_{se}$ and we take $c_{scr} \geq 0$. For this particular choice of functions it is clear that the function c_{ap} is decreasing and $c_{ap}(u) \geq c_{se} - c_{scr}$, and hence by Lemma 3 the function $x \mapsto c(x) + C(x, \tau)$ is discrete convex for every $\tau > 0$. Also for these particular chosen functions, the functions f_1 and f_2 in (15) become

$$f_1(u) = e^{-\delta u}[c_{se} - c_{scr} - p - c_a e^{-\gamma u}],$$
$$f_2(u) = e^{-\delta u}[q(c_{se} + c_{re}) + (1 - q)p - q c_a e^{-\gamma u}].$$

Introducing for every $\rho \geq 0$ the function $J_\rho : [0, T] \to \mathbb{R}_+$ given by

$$J_\rho(\tau) = \mathbb{E}\left(\int_0^\tau e^{-\rho u} dN(u)\right) = \int_0^\tau e^{-\rho u}\lambda(u)du,$$

it follows by relation (5) that for the above choice of functions

$$C(0, \tau) = [q(c_{se} + c_{re}) + (1 - q)p]J_\delta(\tau) - q c_a J_{\delta+\gamma}(\tau) + c_a J_{\delta+\gamma}(T). \tag{22}$$

Also by relation (13) it follows that

$$\Delta_x C(0, \tau) = \begin{cases} c_{se} - c_a - p + e^{-\delta\tau}(p + c_{scr} - c_{se} + c_a e^{-\gamma\tau})e^{-\Lambda_0(\tau)} \\ +[h + \delta(p - c_{se})]\int_0^\tau e^{-\delta u}e^{-\Lambda_0(u)}du \\ +c_a(\gamma + \delta)\int_0^\tau e^{-(\delta+\gamma)u}e^{-\Lambda_0(u)}du, \end{cases} \tag{23}$$

and for every $x \in \mathbb{N}$,

$$\Delta_x^{(2)} C(x - 1, \tau) = \begin{cases} e^{-\delta\tau}(p + c_{scr} - c_{se} + c_a e^{-\gamma\tau})e^{-\Lambda_0(\tau)}\frac{(\Lambda_0(\tau))^x}{x!} \\ +[h + \delta(p - c_{se})]\int_0^\tau e^{-\delta u}e^{-\Lambda_0(u)}\frac{\Lambda_0(u)^x}{x!}du \\ +c_a(\gamma + \delta)\int_0^\tau e^{-(\delta+\gamma)u}e^{-\Lambda_0(u)}\frac{\Lambda_0(u)^x}{x!}du. \end{cases} \tag{24}$$

We can then implement our algorithm in Section 5 with these explicit computations.

In our numerical experiments, we consider a parameter setting comparable to that in Pourakbar et al. (2012). We take the timeline as 66 time units (months), and in our base case scenario, we set the parameters as $c_p = 225$, $c_{scr} = 35$, $c_{se} = 20$, $c_{re} = 30$, $h = 2.25$, $q = 0.5$, $c_a(0) = 645$, $\gamma = 0.03$, $p = 280$, $\delta = 0.0035$. For the arrival intensity, we consider a piecewise constant function which takes values λ, $\beta\lambda$, $\beta^2\lambda$ over the intervals $[0, \frac{T}{3}]$, $[\frac{T}{3}, \frac{2T}{3}]$, $[\frac{2T}{3}, T]$ respectively. In the base case, we have $\beta = 0.5$ and we set the value of λ so that the total expected demand over [0, T] is $10T$ (that is, 10 defective items are returned per unit time).

To see the sensitivity of our results, we vary the values of the cost parameters by several folds and the results are summarized in Table 1. In the table, the row with $\gamma = 0.03$ corresponds to the base case.

In Table 1, along with the results of our static policy and the dynamic approach of Pourakbar et al. (2012), we also report the order quantity and the expected cost associated with the policy in which we do not switch to the alternative policy at all. This problem and its variations are studied in Teunter and Fortuin (1998, 1999), Teunter and Klein Haneveld (1998, 2002), and it fits well to situations where such an alternative policy does not exists or is expensive to employ before the end of horizon T. We still report these values to emphasize the importance of the option to switch to the alternative policy. In the table, P1 corresponds to this policy. P2 is our static policy, and P3 is the dynamic policy of Pourakbar et al. (2012). In each row, we also report the percentage improvements in the expected costs. More precisely, Pi vs Pj reports

$$\frac{\text{Exp. cost of Pi} - \text{Exp. cost of Pj}}{\text{Exp. cost of Pi}}\,\%$$

We observe that, in all cases, P1 performs the worst as expected. On average, P2 offers a reduction of 9.81% and P3 gives a reduction of 11.01% relative to P1. Compared to P2, P3 performs slightly better (average of all cases is 1.35%). As noted in the previous section, P3 is a heuristic for the optimal dynamic policy. Hence, it is expected to perform better than the static policy. As a dynamic policy, it has the ability to adjust to declining inventory levels and prices over time dynamically. However, this relative improvement is not very high and is within 2% in most cases.

When we compare the results for various γ values, we get intuitive observations. For example, in all policies, the costs are monotone in γ as expected. Recall that when γ is very low, the alternative policy remains still expensive, and the policy is not an economically viable option. Indeed, we see that when $\gamma = 0.005$ and 0.01, it is never optimal to employ the alternative policy and all three policies give the same result. On the other hand, when γ increases, the alternative policy becomes more attractive. In the table, as γ increases, the reduction in the costs by P2 and P3 (relative to P1) becomes more apparent, and P2 and P3 start with significantly lower number of items initially as they can later take advantage of the declining prices. Similarly, the improvement of P3 relative to P2 is also increasing as γ increases. As γ increases, the switching decision in the static policy is taken sooner as expected.

As the penalty term p increases, all policies start with higher number of items and the costs are increasing. Also, the relative improvement of P2 and P3 compared to P1 increases as both policies can switch to the alternative policy, and this allows them to avoid incurring penalty costs. The relative improvement of P3 compared to P2 also increases since P3 can adjust the timing of the switching decision dynamically based on the available inventory. When p is small, this has a lower effect as expected. When we increase p, the static policy naturally switches to the alternative policy sooner.

As the holding cost h increases, all policies start with less items and the costs are increasing. The percentage improvement in costs in P2 and P3 compared to P1 is also monotone in h. This is also intuitive since in P1 we have to keep the inventory until the end

Table 1
Sensitivity results for various problem parameters.

	P1		P2			P3				
	x	Exp. cost	x	τ	Exp. cost	x	Exp. cost	**P1 vs P2** (%)	**P1 vs P3** (%)	**P2 vs P3** (%)
$\boldsymbol{\gamma}$										
0.005	325	116432.2	325	66	116432.2	325	116432.2	0.00	0.00	0.00
0.01	321	115309	321	66	115309	321	115309	0.00	0.00	0.00
0.03	296	111213.2	252	39.07	104538.5	255	102711.3	6.00	7.64	1.75
0.05	261	107772.3	198	25.67	87138.36	201	85473.53	19.15	20.69	1.91
0.1	206	102723	113	13.6	58155.69	115	57020	43.39	44.49	1.95
$\boldsymbol{p}$										
140	248	104741.7	244	42.24	102254.1	247	102178.6	2.38	2.45	0.07
200	272	108257.4	249	40.13	103487.3	255	102577	4.41	5.25	0.88
350	307	112830.8	254	38.61	105213.7	256	102797.8	6.75	8.89	2.30
450	316	114410	256	38.21	105960.8	256	102904.8	7.39	10.06	2.88
$\boldsymbol{h}$										
0.5625	313	101045.6	276	43.96	97428.36	276	95456.03	3.58	5.53	2.02
1.125	308	104590.2	267	42.11	99964.15	269	98021.21	4.42	6.28	1.94
4.5	265	122317.4	231	34.85	112305	235	110749.4	8.19	9.46	1.39
9	220	137791.1	205	29.77	124617.2	209	123609.4	9.56	10.29	0.81
$\boldsymbol{c_a(0)}$										
161.25	226	105914.8	0	0	55787.13	0	55787.13	47.33	47.33	0.00
322.5	266	108699.8	164	19.8	87036.23	169	85491.19	19.93	21.35	1.78
1290	313	113579.1	305	60.13	113183.1	306	111825.9	0.35	1.54	1.20
2580	324	116027	324	66	116027	324	116027	0.00	0.00	0.00
$\boldsymbol{c_{scr}}$										
-30	297	111202.9	253	38.87	104371.5	255	102703	6.14	7.64	1.60
17.5	296	111210.5	252	38.94	104496.7	255	102709.4	6.04	7.64	1.71
70	296	111218.6	252	39.27	104617.6	255	102714.4	5.94	7.65	1.82
140	296	111229.4	251	39.34	104757.7	255	102718.1	5.82	7.65	1.95
$\boldsymbol{q}$										
0	598	191442.2	390	23.69	156704	395	153212	18.15	19.97	2.23
0.2	477	159344.6	344	28.84	137820.9	344	134764.8	13.51	15.43	2.22
0.4	356	127254.8	286	35.11	116434.1	288	114126.7	8.50	10.32	1.98
0.6	236	95174.62	214	43.89	91618.49	217	90409.7	3.74	5.01	1.32
0.8	116	63102.07	114	59.6	62875.66	115	62710.52	0.36	0.62	0.26
1	0	30441.36	0	0	30441.36	0	30441.36	0.00	0.00	0.00
$\boldsymbol{\beta}$										
0.125	318	108671.9	314	41.98	108182.9	312	106834.8	0.45	1.69	1.25
0.25	310	109276	294	40.19	107357	295	105511	1.76	3.45	1.72
1	258	114737.8	184	38.54	97831.64	188	96281.93	14.73	16.09	1.58
1.5	228	116809	140	38.48	92417.23	143	91077.65	20.88	22.03	1.45
2	207	118029.1	110	38.28	88471.14	113	87283.99	25.04	26.05	1.34

of the time horizon (or the first time it is depleted prior to T). With higher values of h, under P2 we switch to the alternative policy earlier. Higher values of h also forces the dynamic policy to take the (dynamic) switching decision earlier. Hence as h increases, both P2 and P3 switches sooner and their relative difference becomes smaller (reduces from 2.02% to 0.81%).

As $c_a(0)$ increases, the alternative policy becomes clearly less attractive. Indeed when $c_a(0)$ is very high (see the case with $c_a(0) = 2580$), under all policies we never switch to the alternative policy. On the other hand, when $c_a(0)$ is very low (see the case with $c_a(0) = 161.25$), under P2 and P3 we do not use any inventory and immediately switch to the alternative policy. Needless to say such a setting is not realistic, but we include it here to validate this intuitive result numerically. Again, as expected, when $c_a(0)$ increases, P2 yields higher values for the policy switching time.

When we compare the results for different scrapping cost values, we also have natural observations regarding to the monotonicity of order quantities and costs. However, we observe that the results are not very sensitive to the cost of scrapping. This is intuitive considering that all policies decide on the initial order quantity by taking the probabilistic structure of the demand into account (so that not many items are scrapped at the end).

The results are quite sensitive, on the other hand, to value of the repair probability q. As q decreases, we start with more items in the initial inventory and costs are increasing. Repairing an item is a relatively cheap way of serving a customer. Therefore as q increases, we observe that the static policy switches to the alternative policy later. In the extreme case when $q = 1$, it is never optimal to place an item in the inventory and switch to the alternative policy. In this case, the expected total cost can be explicitly computed as $\int_0^{66} (c_{se} + c_{re}) e^{-\delta u} \lambda(u) du$. In our numerical setting, this number is indeed the reported value 30441.36 (after rounding up). These figures highlight the importance of producing reliable products in managing the end-of-life inventory. Clearly, producing highly reliable products also has its own costs (which occurs during the initial phase of the product life cycle). Hence one should find a balance between these costs, and this should be studied in a separate model as our focus here is on the management of the end-of-life inventory.

Finally, recall that β is the change factor in the demand rate from one interval to the next (while keeping the total expected demand fixed). In reality, we observe fewer and less frequent arrivals over time as noted in Pourakbar et al. (2012). Hence, β values less than 1 are more realistic. In our experiments, we also consider values bigger than 1 for completeness. For low values of β, most of the demand arrives in the beginning of the planning horizon (when the price of the alternative policy is still high) and that is why we place more items in the initial inventory. In this case, all policies perform similarly and the expected costs are comparable. The relative improvement of P2 and P3 compared to P1 becomes significant as β gets higher, which is the case when more demand arrives later. For P1, the expected cost is higher when β increases; when the demand arrives later P1 is forced to keep the inventory and the total cost increases. For P2 and P3, we observe the

Please cite this article as: J.B.G. Frenk et al., An exact static solution approach for the service parts end-of-life inventory problem, European Journal of Operational Research (2018), https://doi.org/10.1016/j.ejor.2018.06.041

opposite behavior. When the demand arrives later, as these policies can take advantage of the declining prices of the alternative policies, P2 and P3 can serve more customers with this cheaper solution. Hence the total costs are decreasing. The policy switching time of P2 is not very sensitive to β. To derive an early switching time, the prices (in the alternative policy) should be sufficiently low. That is why the sensitivity of the policy change time is high with respect to the parameter γ but low with respect to the parameter β.

8. Conclusion

This paper addresses the end-of-life inventory problem of service parts with an alternative service policy in a static setting where the optimal ordering quantity and the policy switching time are both decided at the beginning of the planning horizon. Exploiting the convexity properties of the cost function with respect to the order quantity x, we develop a simple and computationally efficient algorithm to find the optimal solution. We compare the results with a dynamic heuristic in which the policy switching time takes values on the set $\{0, 1, \ldots, T\}$, and the switching decision is made dynamically at one of those time points depending on the remaining inventory. In our numerical experiments, the results indicate that the dynamic policy brings little cost advantage over the static policy. The static policy is clearly much simpler to apply in practice. It allows for the advance planning of the alternative policy and disposal of leftover service parts inventory. Hence, the policy can be embedded much easily into the strategic plans of the OEM. These actions (i.e., switching to a new policy and scrapping the inventory) are more difficult to implement in the dynamic approach since the decisions are taken instantaneously.

References

Behfard, S., van der Heijden, M. C., Al Hanbali, A., & Zijm, W. H. M. (2015). Last time buy and repair decisions for spare parts. *European Journal of Operational Research, 244*(2), 498–510.

Callioni, G., de Montgros, X., Slagmulder, R., Van Wassenhove, L. N., & Wright, L. (2005). Inventory-driven costs. *Harvard Business Review, 83*(3), 135–141.

Çınlar, E. (2011). *Probability and stochastics*. New York: Springer Verlag.

Fortuin, L. (1980). The all-time requirement of spare parts for service after sales – Theoretical analysis and practical results. *International Journal of Operations & Production Management, 1*(1), 59–70.

Fortuin, L. (1981). Reduction of the all-time requirement for spare parts. *International Journal of Operations & Production Management, 2*(1), 29–37.

Hong, J. S., Koo, H.-Y., Lee, C.-S., & Ahn, J. (2008). Forecasting service parts demand for a discontinued product. *IIE Transactions, 40*(7), 640–649.

Moore, J. J. R. (1971). Forecasting and scheduling for past-model replacement parts. *Management Science, 18*(4-part-i), B-200.

Pourakbar, M., Frenk, J. B. G., & Dekker, R. (2012). End-of-life inventory decisions for consumer electronics service parts. *Production and Operations Management, 21*(5), 889–906.

Ritchie, E., & Wilcox, P. (1977). Renewal theory forecasting for stock control. *European Journal of Operational Research, 1*(2), 90–93.

Teunter, R. H., & Fortuin, L. (1998). End-of-life service: A case study. *European Journal of Operational Research, 107*(1), 19–34.

Teunter, R. H., & Fortuin, L. (1999). End-of-life service. *International Journal of Production Economics, 59*, 487–497.

Teunter, R. H., & Klein Haneveld, W. K. (1998). The final order problem. *European Journal of Operational Research, 107*, 35–44.

Teunter, R. H., & Klein Haneveld, W. K. (2002). Inventory control of service parts in the final phase. *European Journal of Operational Research, 137*(3), 497–511.

Van Kooten, J. P. J., & Tan, T. (2009). The final order problem for repairable spare parts under condemnation. *Journal of the Operational Research Society, 60*(10), 1449–1461.

Please cite this article as: J.B.G. Frenk et al., An exact static solution approach for the service parts end-of-life inventory problem, European Journal of Operational Research (2018), https://doi.org/10.1016/j.ejor.2018.06.041