Journal of the Operational Research Society (2015) 66, 331–341

www.palgrave-journals.com/jors/

Inventory management of spare parts in an energy company

Mario Guajardo[1]*, Mikael Rönnqvist[1], Ann Mari Halvorsen[2] and Svein Inge Kallevik[2]

[1]NHH Norwegian School of Economics, Bergen, Norway; and [2]Statoil ASA, Sandsli, Norway

We address the problem of how to determine control parameters for the inventory of spare parts of an energy company. The prevailing policy is based on an (s, S) system subject to a fill rate constraint. The parameters are decided based mainly on the expert judgment of the planners at different plants. The company is pursuing to conform all planners to the same approach, and to be more cost efficient. Our work focuses on supporting these goals. We test seven demand models using real-world data for about 21 000 items. We find that significant differences in cost and service level may appear from using one or another model. We propose a decision rule to select an appropriate model. Our approach allows us to recommend control parameters for 97.9% of the items. We also explore the impact of pooling inventory for different demand sources and the inaccuracy arising from duplicate item codes.

Journal of the Operational Research Society (2015) **66**(2), 331–341. doi:10.1057/jors.2014.8

Published online 12 February 2014

Keywords: inventory; spare parts; demand model; fill rate; record inaccuracy; oil

1. Introduction

This research has been motivated in the context of Statoil ASA, one of the world's largest net sellers of crude oil and the second largest exporter of natural gas to Europe. This company holds inventory of about 200 000 spare part items in a number of locations spread in the Scandinavian region. Some of these parts are highly critical to assure safety and production at offshore platforms. At the same time, the storage of spare parts entails an important binding cost for the company. The prevailing inventory control policy is based on the min-max or (s, S) system. Despite several approaches have been proposed for this inventory policy early in the literature (eg, Veinott and Wagner, 1965; Ehrhardt, 1979; Federgruen and Zipkin, 1984), in practice the values of the control parameters are usually set in a rather arbitrary fashion (Silver *et al*, 1998, p 239). Recent work by Bacchetti and Saccani (2012) reveals the existence of a considerable gap between research and practice in spare parts management, pointing out that managers usually prefer to rely on their judgment or on simple models. The company in our case does not differ much from this, but it is in process of automating inventory management practices, in order to be more uniform across locations, to speed up the evaluations and to become more efficient. Our article describes the inventory management of spare parts at the company and focuses on how to select demand models and how to determine the inventory control parameters.

An overview of related literature on spare parts inventories can be found in Kennedy *et al* (2002). Applications in several contexts include air force organizations and commercial airlines (Eaves and Kingsman, 2004; Sherbrooke, 2004; Muckstadt, 2005), the service support of IBM in the US (Cohen *et al*, 1990), a chemical plant in Belgium (Vereecke and Verstraeten, 1994), a white goods manufacturer in Italy (Kalchschmidt *et al*, 2003) and a distributor of castors and wheels in Greece (Nenes *et al*, 2010). In the oil industry, we have found only one recent article, by Porras and Dekker (2008). Characterized by high service levels, customized equipment specifications with long lead times, and facility networks spread onshore and offshore, we believe the problem of inventory management of spare parts in this industry deserves attention from the research community and our article contributes by addressing a practical problem in this industry. Although a number of distributions have been discussed in the literature for modelling demand for spare parts, empirical evidence is lacking (Syntetos *et al*, 2012). We provide insights of how demand distributions behave in a real-world data set of about 21 000 items and propose a decision rule to select an appropriate distribution for each item. In order to carry out this selection, we test seven demand distributions using the chi-square goodness-of-fit test. We find that significant differences in cost and service level may appear from using one or another distribution. We also use the decision rule to analyse the impact of two main sources of savings: pooling and correcting inventory inaccuracy arising from having registered a same item with two different codes (a *duplicate*).

**Correspondence: Mario Guajardo, Department of Business and Management Science, NHH Norwegian School of Economics, Helleveien 30, Bergen N 5045, Norway.*

E-mail: mario.guajardo@nhh.no

The remainder of this article is organized as follows. In Section 2 we provide background of the company under study and its current practice on inventory of spare parts. In Section 3 we outline our problem settings. In Section 4 we present the different demand models used in this study. In Section 5 we propose the decision rule. In Section 6 we summarize our main numerical results. An evaluation of the impact of pooling and duplicates is presented in Section 7. Our concluding remarks are provided in Section 8.

2. Background

Statoil ASA is an international energy company headquartered in Norway, with presence in 42 countries. Its activities include extraction of petroleum, refinement, and production of gas and methanol. Several types of equipment are used in these activities. Inventory of spare parts is held to replace equipment, in order to assure proper operating conditions.

The company has operational responsibility for seven warehouses located along the Norwegian coast, which serve offshore installations. Within this structure of warehouses, there are 24 inventory plants (or *license inventories*). Statoil ASA is the majority owner of these, but there is also a set of other companies owning a share of them. Each offshore platform is assigned to one of the 24 plants. The inventory of spare parts for each platform is managed separately, even in the case where two different platforms are assigned to a same plant or to two different plants placed in the same warehouse. In the inventory system of the company, there are also other warehouses involved, serving on-shore installations.

The acquisition of spare parts is driven by consumption requirements triggered at the installations. The frequency of the requirement for a given spare part and the number needed are, in general, uncertain. For example, a spare part may not be required in 5 years but it may sometimes be needed a couple of times in few days.

The type of items can be anything from high value items required for production (eg, valves and compressors) to consumable products (smaller tools, such as gaskets, bolts and nuts). Some of these items are highly critical to assure safety and production. Some of them are also highly customized, with requirements specified on frame agreements with suppliers; thus, their construction could involve relatively long lead times.

All information of the spare parts is managed by the enterprise resource planning system SAP. In total, there are about 950 000 codes of items in the system, including warehouses serving onshore and offshore operations. Currently, the items in stock correspond to about 245 000 of those codes. A code in the system intends to serve as unique identifier of each type of items, but in practice a same type of item can actually appear in the database with more than one code. This may be, for example, because units of this type were obtained from different suppliers. A rough estimation indicates that between 10 and 30% of the item codes correspond to duplicates, and the number of distinct items actually in stock is around 200 000.

As a spare part is required at a platform and the corresponding inventory plant provides it, replenishment occurs based on an inventory system with either min or min and max allowable levels. Setting s equal to the min level minus 1 and S to the max level, this system can be described by a continuous review (s, S) policy, that is, every time the inventory drops to the reorder point s or lower, a new order is placed so as to reach the order-up-to level S. When only the min level is used, an order is placed every time there is a demand event in order to achieve a base-stock level, which is a particular case of the (s, S) known as $(S-1, S)$ policy.

The inventory control parameters s and S are determined separately for each platform. When a spare part is required at a given platform, it is provided from the corresponding plant. However, when the part is required in such number that it is not available from the stock assigned to this platform, the requirement can be fulfilled from stock assigned to another platform. This type of fulfilment occurs based on an informal agreement between the planners at the license inventories, but this is not considered when they decide the control parameters. If no plant has stock, the number of the spare part needed is ordered directly from the supplier (even when the parts are taken from stock assigned to another platform, an order is placed with the supplier so as to raise the corresponding inventory position to the order-up-to level).

In practice, the expert judgment of the inventory planners currently play a key role in the decision making rather than that of structured quantitative approach. Although the experience is essential, two members of the staff may provide different inventory control parameters for the same item. Furthermore, each inventory planner decides the control parameters separately; thus, the company lacks a more integrated approach that considers the possibilities of pooling the inventory for different platforms.

3. Problem settings

The primary problem is to find inventory control parameters of a continuous review (s, S) system, considering different demand models. For a given item stored for a given platform, the problem setting is characterized as follows:

- The goal is minimizing expected carrying cost plus ordering cost, subject to a service-level constraint.
- The service-level constraint we utilize is a lower bound β such that

$$1-\frac{\textit{average shortage per replenishment cycle}}{\textit{average demand per replenishment cycle}} \geqslant \beta. \quad (1)$$

The left-hand side above is the expression traditionally utilized to compute the *fill rate*, a performance measure defined as the fraction of demand that is fulfilled directly from the stock on-hand. As pointed out by Guijarro *et al* (2012), while this expression is an approximation to the fill rate, it is the most common method to compute it. Based on simulation, their numerical experiments reveal that this traditional approximation underestimates the simulated fill rate and, therefore, in our purpose to set inventory control parameters it would lead to a more conservative approach. It has been agreed in our project that this conservative approach is acceptable. The lower bound β represents the target fill rate, which depends on the item and the context in which it is used. The target values have been predefined by the company, based on the criticality of the items. Each criticality index is associated with a target value. For items appearing with different criticality indexes (because of the different contexts in which they can be used), the company prefers to consider the highest of the associated service levels.

- The demands for different spare parts are independent random variables.
- Unmet demand is backordered.
- The lead time is fixed. The procurement of spare parts is obtained from external suppliers, with whom the delivery times are agreed beforehand by frame agreements. The company also has knowledge on the setup time of the equipment. These times are reasonably fulfilled, thus supporting the fixed lead time assumption.
- Replenishment is carried out only with new parts acquired from the supplier; thus, in this paper, we do not study the possibility of repairing.

Note that the motivation to use the service-level constraint instead of a penalization cost on backorders comes from our discussion with managers and planners of Statoil ASA. Minimizing the average carrying and ordering costs, subject to such service-level constraint is a popular strategy in practice, since the shortage costs are usually difficult to assign. On the other hand, the service-level criterion is generally easier to state and interpret by practitioners. A number of references point out this fact, such as Cohen *et al* (1988, 1989), Bashyam and Fu (1998), and Chen and Krass (2001).

4. Demand modelling and fill rates

In the literature, there has been large interest on demand modelling and forecasting motivated for the intermittent and slow-moving consumption patterns observed in spare parts inventory (eg, Croston, 1972; Willemain *et al*, 2004; Syntetos and Boylan, 2006). Considering one or another demand behaviour affects the tractability of the fill rate. In order to deal with this, we distinguish between unit-size demand items (only demand events for one unit of the same item at a time can occur), lot-size demand items (events for either one or more units of the same item can occur) and binary demand items (each time a demand event occurs the amount demanded is the same). For unit-size demand, we test Poisson, negative binomial, gamma and gamma with probability mass at zero distributions; for lot-size demand, we test normal and gamma distributions; and for binary demand we test a package Poisson distribution. We choose these traditional distributions that have proven to be effective in case studies (Dunsmuir and Snyder, 1989; Kalchschmidt *et al*, 2003; Boylan *et al*, 2008; Porras and Dekker, 2008; Nenes *et al*, 2010). They are also easy to communicate to practitioners and to implement in software packages of common use.

Let A be the fixed cost per order, μ the average demand per unit of time, σ the standard deviation of demand, v the unit cost, r the carrying charge and β the target fill rate.

For a given item at a given plant following a given demand distribution, we determine s and S levels in two main steps. First, we set the difference $S-s$ as the economic order quantity (EOQ) $Q=\sqrt{(2A\mu)/(vr)}$, rounded to the closest positive integer. Fixing the order quantity, either by this or another formula, is commonly accepted as a reasonable approximation to keep the computational tractability of the problem (eg, Tijms and Groenevelt, 1984; Silver *et al*, 1998, p 331).

Second, given that value of $S-s$, we try to obtain the minimum reorder point s such that the fill rate β is satisfied. Then, we get $S=s+Q$. At the second step we evaluate the fill rate consecutively, starting from $s=0$. Let us call $\tilde{\beta}_s$ the approximated fill rate achieved for a given reorder point s. If the service level β is not satisfied, that is, if $\tilde{\beta}_s<\beta$, we try $s=s+1$ and so on, until we find the lowest value of s such that $\tilde{\beta}_s\geqslant\beta$.

Choosing s and S through this sequence of two steps does not necessarily drive to optimal values (in terms of expected costs). For achieving optimality, the problem should run in both variables simultaneously, but such an exact approach would highly complicate the problem. In fact, despite several decades of inventory research on (s, S) systems, finding optimal control parameters s and S subject to the fill rate service-level constraint remains as an open problem in the general case; most of the applications utilize approximation techniques specialized in particular settings (eg, Cohen *et al*, 1988; Schneider and Ringuest, 1990; Bashyam and Fu, 1998). The sequential computation of s and $Q=S-s$ as we use is a common approach in inventory control (eg, Vereecke and Verstraeten, 1994; Silver *et al*, 1998; Porras and Dekker, 2008).

In what follows, we will briefly overview how to calculate the fill rates for the seven demand models used in this study.

4.1. Unit-size demand

In the unit-size demand case it is relatively easy to compute $\tilde{\beta}_s$. Note that the denominator equals the order quantity $Q=S-s$ and the numerator depends on the lead time demand distribution.

4.1.1. Poisson distribution. Let us call X the demand per unit of time and assume it follows a Poisson distribution of parameter μ, that is, $P(X=k)=(\mu^k e^{-\mu})/k!$ for all non-negative integer values of k. Then,

$$\tilde{\beta}_s = 1 - \frac{\sum_{k=s+1}^{\infty}(k-s)P(X_L=k)}{Q}, \tag{2}$$

where L is the lead time, and X_L and $\mu_L=\mu L$ are the demand during the lead time and its mean, respectively. Hence, $P(X_L=k)=((\mu L)^k e^{-\mu L})/k!$. In the particular case where $Q=1$, instead of approximating the fill rate by Equation (2), we use the well-known formula of the exact fill rate for $(S-1, S)$ systems with Poisson demand (eg, Muckstadt, 2005, p 52) as follows:

$$\tilde{\beta}_s = \sum_{k=0}^{S-1} P(X_L=k). \tag{3}$$

4.1.2. Negative binomial distribution. The negative binomial distribution is defined (see, eg, Graves, 1985) as $P(X=k)=\binom{r+k-1}{k}p^r(1-p)^k$ for all non-negative integer values of k, where r and p are parameters $(0<p<1)$. Empirical and theoretical support for this distribution in the context of spare parts has been provided in several articles, such as Syntetos and Boylan (2006), Boylan *et al* (2008) and Syntetos *et al* (2012). When lead time demand follows a negative binomial distribution, the parameters satisfy $p=\mu/\sigma^2$ and $r=L\mu^2/(\sigma^2-\mu)$, while for the fill rates we can use expressions (2)–(3) in analogous way as we did for the Poisson distribution.

4.1.3. Gamma distribution. The density function of a gamma random variable with scale parameter $\alpha>0$ and shape parameter $k>0$ is defined as $g(x)=\alpha^k x^{k-1}e^{-\alpha x}/\Gamma(k)$, $0\leqslant x<\infty$, where $\Gamma(k)=\int_0^\infty v^{k-1}e^{-v}\mathrm{d}v$. Its use on inventory control dates back to early years (Burgin, 1975; Snyder, 1984). If the lead time demand follows the gamma distribution, the approximated fill rate can be expressed as

$$\begin{aligned}\tilde{\beta}_s &= 1 - \frac{\int_s^\infty (x-s)g(x)\mathrm{d}x}{Q} \\ &= 1 - \frac{\frac{k}{\alpha}\left[1-G_{1,k+1}(\alpha s)\right]-s\left[1-G_{\alpha,k}(s)\right]}{Q},\end{aligned} \tag{4}$$

where $G_{\alpha,k}(x)=\int_0^x g(u)\mathrm{d}u$, $x>0$. The parameters of the gamma distribution in the lead time are estimated as $\alpha=\mu/\sigma^2$ and $k=L\mu^2/\sigma^2$.

4.1.4. Gamma distribution with mass probability at zero. As suggested by Dunsmuir and Snyder (1989), the approximated fill rate in this case can be expressed as

$$\tilde{\beta}_s = 1 - \frac{\left[\int_s^\infty (x-s)g_+(x)\mathrm{d}x\right]p}{Q}, \tag{5}$$

where p is a weight representing the probability of having positive demand in a unit of time, estimated from the data as the quotient between the number of months where demand was positive and the total number of months; and g_+ is the density function of a gamma random variable, for which we estimate scale parameter $\alpha_+=\mu_+/\sigma_+^2$ and shape parameter $k=L\mu_+^2/\sigma_+^2$, considering only the data for months with positive demand to calculate the mean μ_+ and the variance σ_+^2. Then, the term in brackets at the numerator of Equation (5) can be computed in the same way as we did in Equation (4).

4.2. Lot-size demand

As far as we know, in the general case of lot-size demand there are no exact methods to compute the left-hand side of constraint (1). In contrast to the unit-size case, the undershoot distribution makes the lot-size case more complex. We use the following approximation for the fill rate, taken from Tijms and Groenevelt (1984):

$$\tilde{\beta}_s \approx 1 - \frac{M(s)}{2\mu\left(S-s+\frac{\sigma^2+\mu^2}{2\mu}\right)}, \tag{6}$$

where the function $M(s)$ depends on the demand distribution. They derive this approximation by expressing the undershoot as the excess variable of a renewal process (Ross, 1996, p 120). The approximation is made under the assumption that the difference $S-s$ is sufficiently large compared with the average demand per unit of time (say, $S-s\geqslant 1.5\mu$). The experiments by Moors and Strijbosch (2002) confirm that this approximation behaves satisfactorily if the condition $S-s\geqslant 1.5\mu$ holds.

4.2.1. Normal distribution. For this distribution, Tijms and Groenevelt (1984) show that $M(s)=\sigma_\eta^2 J((s-\mu_\eta)/\sigma_\eta)-\sigma_\xi^2 J((s-\mu_\xi)/\sigma_\xi)$, where μ_η and σ_η are the mean and standard deviation of demand in the lead time plus a unit of time; μ_ξ and σ_ξ are the mean and standard deviation of demand in the lead time; and $J(x)=\int_x^\infty (u-x)^2\varphi(u)\mathrm{d}u=(1+x^2)[1-\Phi(x)]-x\varphi(x)$ is an expression in terms of the standard normal density $\varphi(x)=\frac{e^{-x^2/2}}{\sqrt{2\pi}}$ and its cumulative function $\Phi(x)=\int_{-\infty}^x \varphi(u)\mathrm{d}u$.

4.2.2. Gamma distribution. For this distribution, Tijms and Groenevelt (1984) show that

$$\begin{aligned}M(s) &= \sigma_\eta^2(a_\eta+1)\left[1-F_{a_\eta+2}(b_\eta s)\right]-2su_\eta\left[1-F_{a_\eta+1}(b_\eta s)\right] \\ &\quad + s^2\left[1-F_{a_\eta}(b_\eta s)\right]-\sigma_\xi^2(a_\xi+1)\left[1-F_{a_\xi+2}(b_\xi s)\right] \\ &\quad + 2s\mu_\xi\left[1-F_{a_\xi+1}(b_\xi s)\right]-s^2\left[1-F_{a_\xi}(b_\xi s)\right],\end{aligned}$$

where $a_\eta=\mu_\eta^2/\sigma_\eta^2$ and $b_\eta=\mu_\eta/\sigma_\eta^2$ (analogously for a_ξ and b_ξ) and $F_k(x)$ is the distribution function of a gamma with scale parameter 1 and shape parameter $k=\mu^2/\sigma^2$, that is, $F_k(x)=[1/\Gamma(k)]\int_0^x e^{-u}u^{k-1}\mathrm{d}u$, $x>0$.

4.3. Binary (or clumped) demand

Some items in our database correspond to items that have historically presented binary or *clumped* demand, defined by Boylan *et al* (2008) as intermittent items for which demand, when it occurs, is constant. In order to deal with the clumped items, we use a package Poisson distribution (Vereecke and Verstraeten, 1994; Nenes *et al*, 2010), in which the demand events per unit of time are modelled by a Poisson distribution and the size of each demand event corresponds to a fixed value. Let us call $\overline{\mu}$ this fixed demand value when a demand event occurs and define the modified order quantity $\overline{Q} = \overline{\mu}\lceil \frac{Q}{\overline{\mu}} \rceil$. This modified order quantity $\overline{Q}$ is the actual order quantity under demand events of constant size $\overline{\mu}$. Let us call h the undershoot, which is computed as $h = \overline{Q} - Q$. Also, let us define a modified reorder point $\overline{s} = max(0, s-h)$, which is the actual stock on-hand when a replenishment order is placed.

Let us assume the lead time L is a positive integer; when the actual lead time is fractional, we round it up to the closest greater integer.

Let us assume there are k periods with demand during the lead time; thus $k\overline{\mu}$ units are demanded. There is shortage if $k\overline{\mu} \geqslant \overline{s}+1$. Let us define $\overline{k} = \lceil \frac{\overline{s}+1}{\overline{\mu}} \rceil$. If $\overline{k} \leqslant L$, when there are k demand events such that $\overline{k} \leqslant k \leqslant L$, the shortage is equal to $k\overline{\mu} - \overline{s}$. Therefore, the expected shortage during the lead time is $\sum_{k=\overline{k}}^{L} (k\overline{\mu} - \overline{s}) P(X_L = k)$, where X_L is the number of demand events during the lead time, modelled by a Poisson distribution. Assuming that the number of demand events per unit of time follows a Poisson distribution with parameter p, we get

$$\tilde{\beta}_s = 1 - \frac{\sum_{k=\overline{k}}^{L} (k\overline{\mu} - \overline{s}) \frac{(pL)^k e^{-pL}}{k!}}{\overline{Q}}. \tag{7}$$

Table 1 Illustrative data

Item	μ	σ	μ_+	σ_+	$N_{D>0}$	$N_{D>1}$	L	β
M1	0.16	0.48	1.22	0.63	9	1	0.33	0.95
M2	0.03	0.17	1.00	0.00	2	0	0.50	0.95
M3	0.04	0.27	1.50	0.50	2	1	0.33	0.95
M4	0.03	0.17	1.00	0.00	2	0	10.20	0.97
M5	0.15	0.55	1.43	1.05	7	1	0.17	0.95
M6	0.28	0.73	1.73	0.86	11	6	6.47	0.97
M7	0.04	0.21	1.00	0.00	3	0	6.67	0.97
M8	0.04	0.27	1.50	0.50	2	1	1.17	0.95
M9	1.73	7.57	29.00	13.00	4	4	0.47	0.95

5. Selecting a demand model

Seven alternatives to model demand can provide seven highly different outcomes. Table 1 presents data on nine items over a period of 67 months. $N_{D>0}$ denotes the number of months with demand greater than zero and $N_{D>1}$ the number of months with demand greater than one. The resulting values of s and S in each case are reported in Table 2.

From this example, we emphasize that important differences can arise for some items depending on the demand model used. For instance, the Poisson model leads to a reorder point $s = 1$ for item M9, which is the one with highest volatility, while the other models lead to much higher reorder points, in a range from 8 to 65. Among the two lot-size demand models, the normal distribution leads to lower reorder points than the gamma distribution for all items, which is expected since the normal allows negative values. In other cases, different demand models lead to equal or similar results. For instance, the four unit-size models produce the same results for items M1, M4 and M5, while the Poisson, unit-size gamma and package Poisson distributions produce the same results for items M2 and M4.

A natural question is then how to select only one demand model for each item. The suggestions in the literature are rather fuzzy. Silver *et al* (1998) refer to a number of distributions and as a rule of thumb they recommend: the Poisson distribution for slow-moving items; the normal distribution for items such

Table 2 Results of the illustrative data example

Item	*Unit-size*								*Lot-size*				*Binary*	
	Poisson		*NBinomial*		*Gamma*		*Gamma$_0$*		*Normal*		*Gamma*		*Package*	
	s	*S*	*s*	*S*	*s*	*S*	*s*	*S*	*s*	*S*	*s*	*S*	*s*	*S*
M1	0	4	0	4	0	4	0	4	1	5	2	6	—	—
M2	0	1	—	—	0	1	—	—	1	2	2	3	0	1
M3	0	1	0	1	0	1	0	1	1	2	3	4	—	—
M4	2	3	—	—	2	3	—	—	2	3	3	4	2	3
M5	0	1	0	1	0	1	0	1	1	2	4	5	—	—
M6	5	6	5	6	8	9	14	15	6	7	9	10	—	—
M7	1	2	1	2	2	3	—	—	2	3	3	4	2	3
M8	0	1	0	1	0	1	1	2	1	2	4	5	—	—
M9	1	9	5	13	10	18	8	16	14	22	65	73	—	—

'—' indicates that the distribution could not be applied because the data is not consistent with the definition of such distribution.

that the coefficient of variation σ/μ is less than 0.5; if this coefficient is greater than 0.5 it might be desirable to use another distribution such as the gamma. In a recent case study, Nenes *et al* (2010) suggest the use of the chi-square goodness-of-fit test to ensure that the Poisson or gamma with mass at zero distributions give acceptable representations of the demand behaviour for some items. Given the large amount of items in our case and the variety of demand behaviours that they present, we proceed through some classification criteria before selecting a distribution. First, we check whether the EOQ is either equal to or greater than one. Second, we classify the item according to its historical demand behaviour. If there was at most one period per year with demand greater than one (ie, if $N_{D>1} \leqslant 1$), then we classify the demand as unit-size behaviour; if there were two or more periods per year with demand greater than one (ie, if $N_{D>1} \geqslant 2$), then we classify the demand as lot-size behaviour; if there was no variation in the size of demand events ($\sigma_+ = 0$), then we classify the demand as clumped behaviour. Third, we consider the ratio $r = |\sigma^2 - \mu|/\mu$, which gives a measure of the relative difference of the variance σ^2 with respect to the mean μ. Fourth, we consider the empirical support for each distribution, by performing a chi-square goodness-of-fit test over the historical data. Then, we propose the following decision rule.

1. If $EOQ > 1$, then:
   - (1.1) If the demand has historically behaved as clumped, choose the package Poisson model if it was not rejected in the test.
   - (1.2) If the demand has historically behaved as unit-size, then: (1.2.1) choose the Poisson model if $r \leqslant 0.1$ and if it was not rejected in the test and obtained higher p-value than the other three unit-size distributions; otherwise (1.2.2) choose the negative binomial distribution if it was not rejected in the test and obtained higher p-value than the other three unit-size distributions; otherwise (1.2.3) choose from the gamma distribution and the gamma with mass at zero the distribution such that it was not rejected in the test and it obtained the highest p-value among them.
   - (1.3) If the demand has historically behaved as lot-size and $S - s \geqslant 1.5\mu$, then choose one of the two lot-size demand models such that it was not rejected in the test and it obtained the highest p-value among them.
   - (1.4) If neither of the previous steps leads to the selection of a demand model, then leave the item for further managerial revision.
2. If $EOQ = 1$, then:
   - (2.1) Choose between the Poisson and negative binomial distributions the one that was not rejected in the test and achieved the highest p-value between the two of them; otherwise:
   - (2.2) if $r \leqslant 10$, then follow the rules 1.2.3, 1.3 and 1.4;
   - (2.3) if $r > 10$, then leave the item for further managerial revision.

Step 1 addresses the items with $EOQ > 1$. Steps 1.1, 1.2 and 1.3 discern on the demand behaviour of these items, according to unit-size, lot-size and clumped definitions. Steps 1.2.1, 1.2.2 and 1.2.3 discern on the fitness of the four unit-size distribution to the data. In the chi-square goodness-of-fit test, the null hypothesis is that the data follows a given distribution. Since it is possible that this hypothesis is not rejected for more than one distribution, we use the p-value as a measure of how much agreement there is between the data and the null hypothesis (see, eg, Wonnacott and Wonnacott, 1985, p 263). Step 1.3 incorporates these fitness criteria for items with lot-size behaviour. The condition $S - s \geqslant 1.5\mu$ is incorporated for ensuring the approximated fill rate for lot-size demand (expression (6)) is well behaved. Step 1.4 is stated for those items which did not fulfil any of the previous conditions.

Step 2 addresses the items with $EOQ = 1$. These usually correspond to the most important items, in terms of cost and criticality, and present unit-size behaviour. Therefore, we attempt to treat them with more accurate calculation. This is pursued in step 2.1, where only the discrete distributions and the exact fill rate expression (3) are allowed. In step 2.2 we allow these items to be modelled by continuous distributions and their corresponding fill rate approximations, as far as $r \leqslant 10$. This inequality is incorporated for avoiding these continuous distributions being selected for items with extremely high variance in comparison with the mean. In these cases, which rarely occurred, we rather leave the item for managerial revision as stated in step 2.3.

6. Numerical results

We present a summary of the results obtained, using a database with historical demand on 21 448 items over a period of 78 months. Our computational implementation is coded in Matlab R2011b on an Intel Core2 Duo 2.27GHz processor with 8GB of RAM, supported by a spreadsheet in Excel. The implementation runs reasonably fast, providing results in few minutes for the whole set of 21 448 items. The results of applying the chi-square goodness-of-fit test are summarized in Table 3. A first observation is that the normal distribution is rejected for the majority of the items, except for a 0.4% of them. There is previous literature pointing out the limitations of the normal distribution to model demand of spare parts

Table 3 Outcome of the chi-square goodness-of-fit test over a set of 21 448 items

Distribution	*Non-rejected*	*Percentage*
Poisson	16 059	74.9
Negative binomial	16 070	74.9
Gamma	20 669	96.4
Gamma mass zero	3806	17.7
Normal	82	0.4
Package Poisson	2686	12.5

(eg, Dunsmuir and Snyder, 1989; Vereecke and Verstraeten, 1994; Dolgui and Pashkevich, 2008; Syntetos *et al*, 2009) and our work with real-world data provides additional evidence of that. In contrast to the normal, the Poisson distribution, the negative binomial distribution and the gamma distribution are not rejected for the majority of the items, accounting for 74.9, 74.9 and 96.4%, respectively. For only 0.1%, equivalent to 24 items, all the distributions were rejected.

We then apply our selection criteria for demand models. Table 4 presents a summary of the frequency in which each distribution is selected. The negative binomial distribution is the one selected for most of the items, with a frequency of 41.5%. It is followed by the Poisson and gamma distributions, selected with a frequency of 20.7 and 18.2%, respectively. We select the lot-size demand behaviour with gamma distribution for 11.9% of the items and the package Poisson for 5.5%. A significantly lower number of items are selected for normal distributions, equivalent to only 0.2%, while the gamma with mass at zero is selected for none of the items. Note that we have been able to recommend control parameters for 21 001 items, while for only 447 items, which corresponds to 2.1%, our approach does not result in any recommendation and rather leaves the items for further managerial evaluation. Most of these 447 items presented a lot-size demand behaviour and both the normal and gamma distributions were rejected by the test or, if at least one of these two distributions was not rejected, the coefficient $(S-s)/\mu$ was less than 1.5, thus not sufficiently large to apply the approximated fill rate of the lot-size demand case.

Table 4 Number of items for which each demand model is selected

Distribution	*Recommendations*	*Percentage*
Poisson	4429	20.7
Negative binomial	8906	41.5
Gamma	3899	18.2
Gamma mass zero	0	0
Gamma lot-size	2541	11.9
Normal	43	0.2
Package Poisson	1183	5.5
Total	21 001	97.9

After the demand models have been selected for each item, we determine the inventory control parameters. We compare the performance of our recommendations in terms of costs and service level with another four alternatives, as shown in Table 5. Alternative A1 corresponds to setting the inventory control parameters assuming for all items that the demand follows a Poisson distribution, while alternatives A2 and A3 assume for all the items gamma and normal distributions, respectively. Alternative A4 corresponds to a mixed selection strategy, where slow-moving items are modelled by a Poisson distribution, while fast-moving items with *low* coefficient of variation (namely, $\sigma/\mu \leqslant 0.5$) are modelled by a normal distribution and fast-moving items with *high* coefficient of variation ($\sigma/\mu > 0.5$) are modelled by a gamma distribution. We define the slow-moving items as those having an average demand per year lower than three units, while fast-moving items as those having an average demand per year equal or greater than three units. Based on their criticality, the items are divided in four main categories that we will call C1, C2, C3 and C4, where a higher index category means higher criticality.

In order to quantify the impact on inventory costs, we use some standard definitions from the inventory literature (see, eg, Silver *et al*, 1998). We calculate the safety stock κ as the reorder point s minus the expected demand during the lead time. We also calculate the expected number of orders per year N_A as $\mu/[S-s+E(U)]$, where $E(U)$ is the expected undershoot which can be approximated by $(\sigma^2+\mu^2)/2\mu$, and for the average inventory on-hand $\bar{I}$, we use the rather simple approximation $\kappa+(S-s)/2$ (see, eg, Strijbosch *et al*, 2000).

Then, we calculate the total expected cost as $C_T = C_r + C_A$, where C_r is the expected carrying cost and C_A is the expected ordering cost, calculated respectively as $C_r = \bar{I}vr$ and $C_A = AN_A$. Owing to confidentiality reasons, we have normalized the total cost results by setting at 100 the cost of our recommendation and expressing the cost achieved by the other alternatives in relative terms to ours. The service levels, on its hand, are calculated assuming that the actual demand follows the distribution as selected by our rule. Similarly, the number of violations by each alternative corresponds to the number of items such that the service-level constraint is violated if the inventory control parameters would be set by such alternative and the actual demand follows the distribution as selected by our rule. The outcomes reveal that assuming a Poisson

Table 5 Performance of our recommendation and other four alternatives in terms of costs and four service-level categories

	Total cost	*Average service level*				*Nr. of violations*			
		C1	*C2*	*C3*	*C4*	*C1*	*C2*	*C3*	*C4*
Recommendation	100	96.1	97.1	97.8	98.5	0	0	0	0
A1: All Poisson	93.2	93.8	95.6	96.4	96.5	1443	114	247	895
A2: All gamma	122.8	94.1	96.5	97.2	97.6	1414	96	205	643
A3: All normal	108.5	96.5	98.3	98.7	98.7	1276	95	207	717
A4: Mixed	118.5	94.1	96.5	97.2	97.6	1415	96	205	676

distribution for all the items implies 6.8% lower costs, but at the risk of 2699 items with unsatisfied service levels. In particular, for the items in category C4 (those with highest criticality indexes) violations to the service-level bounds are incurred for 895 items, achieving an average of 96.5% of service level, while our recommendation achieves a 98.5% for this category of items. On the other hand, assuming that all items follow gamma or normal distributions lead to total cost 22.8 and 8.5% higher than our recommendation. The amount of items whose service-level bound is violated by using alternative A2 corresponds to 2358, of which 643 belong to category C4, while for alternative A3 the number of violations is 2295, of which 717 belong to category C4. The mixed alternative A4 drives to 2392 items with violations and 18.5% higher total cost. Our recommendation also outperforms alternative A4 in terms of average service level, achieving greater percentages in all categories; in particular, for category C4 our recommendation achieves 98.5% of average service level, while alternative A4 achieves 97.6%.

These comparisons highlight the relevance of setting inventory control parameters by using one or another demand model. Especially when a large number of items are involved, some of them very expensive and critical to assure safety and production, small percentage of savings or gain in service levels can become significant.

7. Duplicates and pooling

Determining the control parameters by our quantitative approach instead of the expert judgment does not only allow us to evaluate the trade-off between service levels and inventory costs, but also to analyse the impact of certain projects as we will study in this section.

7.1. Duplicates

A duplicate means that the same item has been registered in the database with different codes. In the event of there being few items to manage, it would be relatively easy to identify which items correspond to duplicates and the inventory staff could quickly realize any discrepancies between the information in the system and the actual stock. However, when the number of items is in the order of a hundred thousand and several people access the system in different locations, the duplicates could create a more difficult puzzle to deal with. Evidence from the retail sector (DeHoratius *et al*, 2008) has pointed out that inventory record inaccuracy is a significant problem in practice, and the elimination of inventory inaccuracy can reduce supply chain costs as well as the out-of-stock level (Fleisch and Tellkamp, 2005). In these references, the inaccuracies refer to discrepancy between the recorded inventory quantity and the actual inventory quantity physically available on the shelf. In our case, the discrepancies arise from the duplicated codes, which have been estimated to be roughly between 10 and 30% of the item codes in the system. This may lead to inventory levels higher than desirable when, for example, one item with a given code in the system reaches its reorder point, triggering a replenishment order while the same item with different code is still available. On the other hand, when a stockout of an item occurs, the realized service level could have been better if the item was actually available but with a different code and it was not realized by the staff.

We implement a procedure in order to analyse the cost impact of the spare parts duplicated in the system. Assessing this impact is relevant in practice, since the company is pursuing to implement a system to identify and fix the duplicates. Let us assume there is 10% of duplicates. Then, our procedure is based on aggregating one out of each 10 items in the system. First, we sort the database by prices, from the most expensive item to the cheapest one. Then, we merge the 10th item with the 9th item, the 19th with the 20th and so on, as if they would correspond to the same item. The merging of two items consists of aggregating their historical demand (obtaining the resulting average, standard deviation, etc) and calculating for the merged item a new lead time, unit cost and order cost as the average of the corresponding values from the two separate items. As for the criticality index, if the two items appeared with different values, we use the highest one. We implemented this procedure for 1639 items stored for one of the main platforms. In total, 318 of these 1639 items were merged into 159 items. Then, we use our approach to compute the inventory control parameters and compare the results obtained with the case where the items were not merged. We found that in the merged case the safety stock decreased by 3% and the average number of orders decreased by 15% (while the inventory on-hand increased 7%, but mainly for low value items). These figures translate into savings of 10% of the expected cost per year.

We performed a similar analysis varying the percentage of merged items from 0 to 30%, increasing it by 5% in each run. This analysis is motivated by the lack of knowledge of exactly how many items are duplicated in the system. If we assume that a certain percentage of items are duplicated and that they are corrected by merging them, we can compute the savings compared with the total expected cost when all the original items in the database are considered per separate. As shown in Figure 1, the merger of up to 30% of the items led to savings of about 14% in total expected cost.

If there are between 10 and 30% of duplicates in the system, our computations indicate on average there would be 12% of savings if they would be corrected.

Other than the relative order in the sorted-by-price list, one could add more conditions when merging items, such as their similarity in other attributes. We have performed additional runs, merging two consecutive items in the sorted-by-price list only if their unit prices differ by no more than 0.1% and their lead times differ by no more than 0.1%. Under these criteria, 386 items (equivalent to 23%) are merged into 193 items, resulting in savings of 8.2% in comparison with the total expected costs when they are not merged. If instead we would

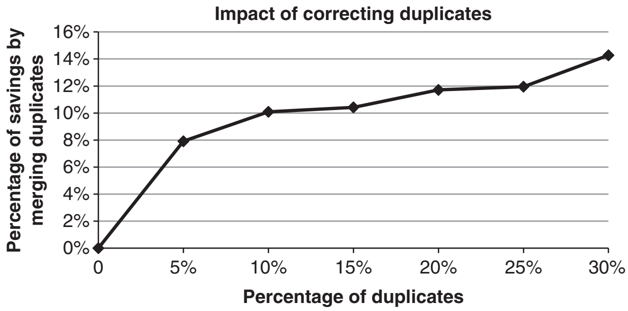

Figure 1 Percentage of savings in total expected cost as a function of the percentage of duplicated items, if the duplicates are corrected.

consider the cost of the case with merged items as basis, the impact of duplicates corresponds to an increase of 9.0% in total expected costs.

7.2. Pooling

Pooling is a traditional mechanism of collaboration among inventory managers, although not necessarily taken into account when deciding the inventory control policies. Kukreja and Schmidt (2005) and Kranenburg and van Houtum (2009) report problems where different warehouses do not consider collaboration when deciding their policies but in practice they supply each other, which is the same practice at the company in our case. Our situation is similar to the one of Kukreja and Schmidt (2005), namely a single-echelon problem with a multi-location, continuous-review inventory system in which complete pooling is permitted and each location utilizes an (s, S) policy. However, in our case, some of the different inventory locations can actually be placed in the same warehouse, but just keeping stock for different platforms as we described in the introduction. Particularly these separate stocks at a same warehouse present an important opportunity for pooling. In order to analyse the potential savings of pooling, we carry out computations for 10 platforms, which are served by six inventory plants located in the same warehouse. We focus on 952 items that presented consumption in at least two of these 10 platforms.

We first determine inventory control parameters s and S for each platform separately. Then, we also solve the case where the warehouse would be managed centrally as one inventory pool serving the demand from the 10 platforms, instead of separating the stocks for different platforms. The data we use for the pool case is the aggregated demand data from the 10 platforms. For a given item, we use the same service-level bound in the separated planning and the integrated planning. If the item presents different service-level bounds for different platforms, in the pool planning we round-up the levels, that is, we use the greatest of these service-level bounds for this item.

From the results so obtained, we summarize important measures in Table 6. The pooling approach results in 9% less

Table 6 Safety stock (κ), expected number of orders (N_A) and average inventory on-hand ($\bar{I}$) under separated planning and pool planning

	κ	N_A	$\bar{I}$
Separated	11 009	1424	23 869
Pool	9979	1047	17 872
Difference	−9%	−26%	−25%

safety stock, 26% fewer expected number of orders and 25% less average inventory on-hand. These reductions mean total annual savings of about 21% when comparing the pool to the separated planning case. Since we have rounded-up the service-level bounds and have limited the analysis to only 952 items in only one of the seven warehouses, the potential savings by pooling for the whole problem appear to be promising.

8. Concluding remarks

We have studied a problem of spare parts inventory in Statoil ASA, a large energy company. Our attention focused on how to set control parameters for the (s, S) continuous review system, subject to the fill rate service-level constraint. Selecting an appropriate demand model plays a relevant role when solving this problem. We have studied the fitness of seven demand models to the data of about 21 000 items. Based on this fitness, we have proposed a decision rule to select a demand model for each item. The demand models selected have proved to be useful for modelling a vast amount of items, while at the same time keeping the computations reasonably simple and quick. Our approach allowed us to recommend control parameters for 97.9% of all tested items.

Although our article has been inspired in the case of Statoil ASA, our contribution connects with the more general research agenda on inventory, which we would like to summarize in four main points. First, we have used seven demand models and defined rules to select one of them for each item based on the support provided from the data, rather than setting a demand model beforehand. In line with Syntetos *et al* (2012), we believe that the development of goodness-of-fit tests for application in inventory control of intermittent demand items is an interesting future research issue. They use the Kolmogorov-Smirnov test, which is not appropriate when the parameters are estimated from the sample (as it is in our case) and it only applies to continuous distributions. The chi-square goodness-of-fit test, on its hand, may fail on discerning whether a distribution is rejectable or not when the demand values have presented small frequencies in the history. Developing a specialized test for spare parts could help to bridge the gap between practice and literature.

Second, we have explored the benefits of pooling in a problem with a single-echelon, multi-location, continuous-review inventory system in which complete pooling is permitted and each location utilizes an (s, S) policy. As noted by Kukreja and Schmidt (2005), this problem has not received special attention from the literature, despite its complexity and relevance in practice. A related problem is how the costs of the pool should be shared. Recent literature has introduced the use of game theory to address this issue (Wong *et al*, 2007; Karsten *et al*, 2012), opening a promising research line within the literature on spare parts inventories.

Third, we have introduced the correction of duplicate codes as another issue in large inventory data systems. To our knowledge, this has not been addressed in previous literature and we believe it is worthy of in-depth exploration in the context of spare parts, as has been done with related issues of record inaccuracy on retail inventories (Fleisch and Tellkamp, 2005; DeHoratius *et al*, 2008).

Fourth, our work has highlighted the importance of the inventory management of spare parts in the oil and gas industry, which so far has not been a main protagonist in related literature. Characterized by high service levels, geographical networks with interaction between onshore and offshore facilities and items with customized specifications and long lead times, we believe there are important chances of improving the inventory practices in the industry, with opportunities for achieving great impact on savings while at the same time maintaining or increasing service levels.

Acknowledgements—We are grateful to discussants and audiences at ISIR Summer School Istanbul 2011, INFORMS Annual Meeting Charlotte 2011, EurOMA Cambridge 2011, Statoil Kompetansedager Bergen 2011, NHH Seminar Geilo 2012, EURO Vilnius 2012 and ILS Quebec 2012 for their valuable comments. We thank the anonymous reviewers, whose comments improved the final version of the article. We also thank Vivienne Knowles and Vegard Engeset for their help in bringing this article to fruition.

References

Bacchetti A and Saccani N (2012). Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega* **40**(6): 722–737.

Bashyam S and Fu MC (1998). Optimization of (s, S) inventory systems with random lead times and a service level constraint. *Management Science* **44**(12): S243–S256.

Boylan JE, Syntetos AA and Karakostas GC (2008). Classification for forecasting and stock control: A case study. *Journal of the Operational Research Society* **59**(4): 473–481.

Burgin TA (1975). The gamma distribution and inventory control. *Operational Research Quarterly* **26**(3): 507–525.

Chen FY and Krass D (2001). Inventory models with minimal service level constraints. *European Journal of Operational Research* **134**(1): 120–140.

Cohen MA, Kleindorfer PR and Lee HL (1988). Service constrained (s, S) inventory systems with priority demand classes and lost sales. *Management Science* **34**(4): 482–499.

Cohen MA, Kleindorfer PR and Lee HL (1989). Near-optimal service constrained stocking policies for spare parts. *Operations Research* **37**(1): 104–117.

Cohen M, Kamesam PV, Kleindorfer P and Lee HL (1990). Optimizer: IBM's multi-echelon inventory system for managing service logistics. *Interfaces* **20**(1): 65–82.

Croston JD (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly* **23**(3): 289–303.

DeHoratius N, Mersereau AJ and Schrage L (2008). Retail inventory management when records are inaccurate. *Manufacturing & Service Operations Management* **10**(2): 257–277.

Dolgui A and Pashkevich M (2008). On the performance of binomial and beta-binomial models of demand forecasting for multiple slow-moving inventory items. *Computers & Operations Research* **35**(3): 893–205.

Dunsmuir WTM and Snyder RD (1989). Control of inventories with intermittent demand. *European Journal of Operational Research* **40**(1): 16–21.

Eaves AHC and Kingsman BG (2004). Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society* **55**(4): 431–437.

Ehrhardt R (1979). The power approximation for computing (s, S) inventory policies. *Management Science* **25**(8): 777–786.

Federgruen A and Zipkin P (1984). An efficient algorithm for computing optimal (s, S) policies. *Operations Research* **32**(6): 1268–1285.

Fleisch E and Tellkamp C (2005). Inventory inaccuracy and supply chain performance: A simulation study of a retail supply chain. *International Journal of Production Economics* **95**(3): 373–385.

Graves SC (1985). A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Science* **31**(10): 1247–1256.

Guijarro E, Cardós M and Babiloni E (2012). On the exact calculation of the fill rate in a periodic review inventory policy under discrete demand patterns. *European Journal of Operational Research* **218**(2): 442–447.

Kalchschmidt M, Zotteri G and Verganti R (2003). Inventory management in a multi-echelon spare parts supply chain. *International Journal of Production Economics* **81–82**: 397–413, http://www.sciencedirect.com/science/journal/09255273/81-82.

Karsten F, Slikker M and Houtum GV (2012). Inventory pooling games for expensive, low-demand spare parts. *Naval Research Logistics* **59**(5): 311–324.

Kennedy WJ, Patterson JW and Fredendall LD (2002). An overview of recent literature on spare parts inventories. *International Journal of Production Economics* **76**(2): 201–215.

Kranenburg AA and van Houtum GJ (2009). A new partial pooling structure for spare parts networks. *European Journal of Operational Research* **199**(3): 908–921.

Kukreja A and Schmidt CP (2005). A model for lumpy demand parts in a multi-location inventory system with transshipments. *Computers & Operations Research* **32**(8): 2059–2075.

Moors JJA and Strijbosch LWG (2002). Exact fill rates for (R, s, S) inventory control with gamma distributed demand. *Journal of the Operational Research Society* **53**(11): 1268–1274.

Muckstadt JA (2005). *Analysis and Algorithms for Service Parts Supply Chain*. Springer Series in Operations Research and Financial Engineering. Springer: New York.

Nenes G, Panagiotidou S and Tagaras G (2010). Inventory management of multiple items with irregular demand: A case study. *European Journal of Operational Research* **205**(2): 313–324.

Porras E and Dekker R (2008). An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods. *European Journal of Operational Research* **184**(1): 101–132.

Ross SM (1996). *Stochastic Processes*. John Wiley & Sons: New York.

Schneider H and Ringuest JL (1990). Power approximation for computing (s, S) policies using service level. *Management Science* **36**(7): 822–834.

Sherbrooke CC (2004). *Optimal Inventory Modeling of Systems Multi-Echelon Techniques*. Kluwer Academic Publishers: Boston, MA.

Silver EA, Pyke DF and Peterson R (1998). *Inventory Management and Production Planning and Scheduling*. John Wiley & Sons: New York.

Snyder RD (1984). Inventory control with the gamma probability distribution. *European Journal of Operational Research* **17**(3): 373–381.

Strijbosch LWG, Heuts RMJ and van der Schoot EHM (2000). A combined forecast-inventory control procedure for spare parts. *Journal of the Operational Research Society* **51**(10): 1184–1192.

Syntetos A and Boylan J (2006). On the stock control performance of intermittent demand estimators. *International Journal of Production Economics* **29**(103): 36–47.

Syntetos A, Keyes M and Babai M (2009). Demand categorisation in a European spare parts logistics network. *International Journal of Operations & Production Management* **29**(3): 292–316.

Syntetos AA, Babai MZ and Altay N (2012). On the demand distributions of spare parts. *International Journal of Production Research* **50**(8): 2101–2117.

Tijms HC and Groenevelt H (1984). Simple approximations for the reorder point in periodic and continuous review (s, S) inventory systems with service level constraints. *European Journal of Operational Research* **17**(2): 175–190.

Veinott AF and Wagner HM (1965). Computing optimal (s, S) inventory policies. *Management Science* **11**(5): 525–552.

Vereecke A and Verstraeten P (1994). An inventory management model for an inventory consisting of lumpy items, slow movers and fast movers. *International Journal of Production Economics* **35**(1–3): 379–389.

Willemain TR, Smart CN and Schwarz HF (2004). A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting* **20**(3): 375–387.

Wong H, Oudheusden DV and Cattrysse D (2007). Cost allocation in spare parts inventory pooling. *Transportation Research Part E* **43**(4): 370–386.

Wonnacott RJ and Wonnacott TH (1985). *Introductory Statistics*. John Wiley & Sons: New York.

Received 11 February 2013;
accepted 13 January 2014 after one revision