Estimating obsolescence risk from demand data - a case study

Willem van Jaarsveld* and Rommert Dekker

Econometric institute, Erasmus University Rotterdam,
P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

Report Econometric Institute EI 2010-03

January 18, 2010

Abstract

In this paper obsolescence of service parts is analyzed in a practical environment. Based on the analysis, we propose a method that can be used to estimate the risk of obsolescence of service parts. The method distinguishes groups of service parts. For these groups, the risk of obsolescence is estimated using the behavior of similar groups of service parts in the past. The method uses demand data as main information source, and can therefore be applied without the use of an expert's opinion. We will give numerical values for the risk of obsolescence obtained with the method, and the effects of these values on inventory control will be examined.

Keywords: Inventory; Spare parts; Obsolescence; Forecasting.

* Corresponding author: E-mail: vanjaarsveld@ese.eur.nl

1 Introduction

Giving good service is considered a requirement to remain competitive throughout industry. This requirement forces manufacturers to keep a stock of service parts, because this is often the only way in which defects of the product can be repaired fast. Furthermore, keeping stock is often needed to attain economies of scale in production or procurement. However, obsolescence of service parts is an important cost factor. Cattani & Souza (2003) report that scrapping of obsolete inventory can reduce profits by up to 1% each year.

In the literature, quite a few different approaches towards incorporating the possibility of obsolescence in inventory models are available. Brown, et al. (1964) have proposed two classes of discrete time models. The first class of models incorporates the risk of items becoming obsolete using a mortality distribution. In the second class of models Markov processes are used to model the risk of parts becoming obsolete. Moore (1971) develops a forecasting system to estimate the total requirement of consumable service parts. Furthermore, a dynamic programming inventory model is described to optimize the production runs. Ritchie & Wilcox (1977) develop a method to estimate the total requirement of service parts by using the sales data of the consumer products in which the service parts are used. Renewal theory is then used to develop an appropriate forecast for the relevant service parts. Song & Zipkin (1993) provide a continuous time framework for analysis of non-stationary demand processes. They remark that an important form of non-stationarity is the situation where demand can stop. Using the framework provided by Song & Zipkin (1993), Song & Zipkin (1996) investigate the effects of obsolescence on the inventory policy. They show that significant savings can be made by including the risk of obsolescence in the inventory decision. Cobbaert & van Oudheusden (1996) recognize the importance of stocks becoming obsolete in inventory control. They remark however that in practice, it is only possible to find a rough estimate for the probability that the part will become obsolete in the near future. This makes approaches that have a lot of parameters hard to implement. Therefore, they propose simple methods that only need a rough estimate for the risk that the part will become obsolete in the coming period. They argue that such an estimate can be given by an expert. Teunter & Fortuin (1999) consider the final order problem under the possibility of stock disposal. A dynamic programming formulation of the problem is derived in order to find the optimal policy. Hill, et al. (1999) consider an exponentially declining Poisson demand process. Dynamic programming is used to optimize the ordering process. Teunter & Klein Haneveld (2002) consider a model in which service parts can be obtained in two different ways. During a final production run, parts can be obtained at a low price. After this run the parts can only be obtained at an increased price. They find a series of order-up-to levels, which are decreasing in time, together with an optimal size for the initial order. Cattani & Souza (2003) study the effect of delaying the final order. They find that the manufacturer benefits from this delay, because it improves forecasts. On the other hand, the supplier will need an incentive to enact this delay, because an early final buy is beneficial for his turnover. Song & Lau (2004) construct an approximation for an EOQ model including obsolescence. The proposed solution relies on dynamic programming. Furthermore, their method requires sophisticated knowledge regarding the distribution of the time at which the part becomes obsolete. The problem of determining the final order quantity of repairable service parts is considered by van Kooten & Tan (Forthcoming). The parts cannot always be repaired, for they are sometimes condemned. The problem is modelled as a transient Markov chain. Also, an approximate model is presented that allows for more efficient calculations. Managerial insights are developed and a sensitivity analysis is performed.

This study focuses on obsolescence in service parts inventories. We concentrate on the main practical issue: quantification of the so-called risk of obsolescence. To our best knowledge, methods to estimate this risk are not available in the literature, as in the literature it is assumed that the parameters governing the evolution of the part towards obsolescence are known or can be estimated by an expert. The demand model on which we will concentrate is relatively simple; we use a so-called sudden death demand model with an exponentially distributed demand lifetime. For this model we describe a method that can be used to estimate the expectation of the demand lifetime using demand data.

The remainder of this paper is organized as follows. In Section 2 we will make qualitative observations of the obsolescence problem at the company. This discussion will serve as the primary motivation for our method. In Section 3 we give further motivation for the method by analyzing demand data of service parts. In Section 4 we describe the method, and give ideas on how it was implemented. In Section 5 we will draw conclusions, and give suggestions for future research.

2 Empirical analysis of obsolescence

Obsolescence of service parts can occur resulting from different causes. Depending on the particular cause of obsolescence, different modelling methodologies are more or less appropriate. The method we develop is appropriate for products with a long life cycle, and less appropriate for service parts used for consumer products with a short life cycle.

Examples of long life cycle products include baggage handling systems, aircrafts and machines for nanotechnology markets. In this section we will argue that for long life cycle products, obsolescence also occurs while the product is still being used.

During the final phase, the usage of the product gradually declines over a period of multiple decades. When the usage of the product has completely disappeared, service parts used in the product will be obsolete. For products with a *short* life cycle, this is the most important type of obsolescence. By making an accurate forecast of the installed base changes, the stock of service parts can be driven down before obsolescence occurs.

According to employees and managers at the company at which the study was performed, obsolescence of spare parts may also occur while the product is still in use. They experience that parts that are moving one year can be obsolete the next year, while the usage of the original product has not decreased significantly. In Section 3, we will show that the usage data of spare parts confirms that obsolescence can occur while the product is still in use. In consultation with the employees and management of the company the following three reasons for this to happen were identified:

- Changing operating conditions of original product.
- Changing maintenance policy on original product.
- Use of alternative parts.

The precise prediction of obsolescence that results from changing operating conditions is hard. It is often clear that a change of operation will change the consumption rates of service parts, but it is unclear in what way these consumption rates will change. For instance, if an aircraft changes operation to moving cargo rather than passengers, the consumption pattern will definitely change.

Getting a precise estimate for this change on the level of individual parts is often impossible or it requires a prohibitive amount of effort and know-how. Note that consumption rates may increase as well as decrease as a result of changing operating conditions.

Consumption rates of service parts are also affected by the maintenance policy. The maintenance policy may change if the downtime cost of the original product changes. When the product is young, downtime costs are often very high and a lot of effort is put into making the product as reliable as possible. Downtime costs tend to decrease as the product ages. As a result, comparably less effort is put into preventive maintenance. This affects the consumption rates. Another reason for the maintenance policy to change is that the mechanics performing the maintenance learn more about maintaining the product as they gain experience. This will change the manner in which they perform maintenance which, in turn, affects the service parts being used.

Alternative service parts can be an important reason for obsolescence. The willingness of customers to use alternative parts may also change with the aging of the product. Because using alternative service parts is often initiated by the customer or by third parties during the life cycle of the product, it is very hard for original equipment manufacturer to predict the effect of alternative parts on demand rates.

Summarizing, there are a number of reasons for spare parts to become obsolescent while the product is still in use. It is hard to precisely predict in advance the moment at which obsolescence occurs. Finally, it seems reasonable to assume that the reasons for obsolescence identified above also play a role for other long life cycle products.

3 Analysis of service part demand data

To get closer to the problem of obsolescence in practice, we will make a quantitative analysis of demand data of spare parts at the company. For the analysis we use a large data set consisting of the demand for all service parts used in a single type of product manufactured by an OEM. The product's original sell price was about $10 million. Each time a service part is needed, the date at which the part was needed, together with the part number and the quantity was registered. To get rid of interchangeability issues the data was preprocessed, treating different service parts that are interchangeable as a single service part. Also, parts which became obsolete for reasons that were known before the moment obsolescence occurred were filtered out.

The analysis is based on the discussion of obsolescence in practice in the previous section. We discussed that obsolescence of spare parts can also occur while the product is still being used. In the following, we will give additional motivation for the analysis.

In basic inventory control, the demand parameters of a demand model are determined based on historical usage. Based on the demand model and the parameters, the reorder points are determined. When the forecast is right, this leads to appropriate reorder points. In practical environments, demand fluctuations can cause the forecast to be wrong. In that case, the reorder points will not be the most suitable ones for the demand. To be able to adapt to demand changes, in most practical environments the forecasts are made repeatedly, and the most recent forecast is used to determine the reorder points. When demand changes, this procedure allows the stocks to be adapted appropriately. If demand increases, there may be some temporary availability problems. If demand decreases, the stocks are temporarily too high. However, if demand drops dead before we are able to adapt to the new situation, the stocks cannot be run down because

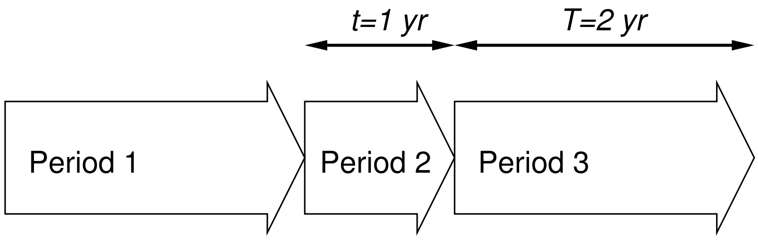

Figure 1: From the demand data, three time periods were selected as shown.

there is no demand to accomplish this. The stocks become obsolete. As discussed before, we will focus on this last issue.

We are thus mainly interested in how often demand drops dead before we are able to run the stocks down. To this end, we selected from the demand data three time periods as shown in Figure 1. Assume a forecast is made for the usage in period 2 at the beginning of that period. The demand data in period 1 is used to make this forecast, and based on the forecast, we determine a reorder point which is used in period 2. At the end of period 2, we might want to adapt the reorder point based on a new forecast. We examine if there is any demand in period 3 which would allow us to drive the stocks down in case we would want to decrease the reorder point. Clearly, the outcomes of this analysis depend on the time interval for which we do the analysis. The obsolescence problem might increase or decrease for this application as time passes. To assess this effect, we will do two analyses, which differ because we let the time periods start 2 years earlier in the second analysis.

The length of period 1 is chosen in such a way that the total usage of the underlying product is the same for period 1 and period 3, and it is only slightly shorter than the length of period 3 (2 years), because the usage of the underlying product is more or less constant. The total usage of service parts is also equal over period 1 and 3. The service parts were grouped based on the demand in period 1. For each of the groups it was assessed how many parts did not have any demand in period 3. It is important to note that the fact that a part does not have any demand in period 3 does not necessarily mean that the part is obsolescent. We will make a distinction between parts that are obsolescent, and parts that are not obsolescent but have zero demand in period 3 because of statistical variation.

To distinguish between zero demand caused by statistical variation and zero demand because of obsolescence, we will compare the fraction of parts in a group for which there is zero demand in period 3 with the fraction of parts in the group which should have zero demand if demand followed a (compound) Poisson process in period 3, with a rate following the forecast based on the number of orders in period 1. The compound Poisson process was chosen as a benchmark because its memoryless property is a reasonable property to ask for when demand originates from a number of independently maintained products. Moreover, it is the standard demand process used for modelling spare parts demand, and therefore it represents a suitable starting point for a more refined approach including obsolescence.

The results of the analysis are shown in Table 1. This table shows that most parts are slow moving, in the sense that they were used only a few times in period 1. This is typical for service

Number of orders in period 1	Number of parts	Fraction of parts with no demand in period 3	Expected fraction of parts with zero demand (Poisson)
1	5630	57.5% (55.4%)	$e^{-1} \approx 36.8\%$
2	2434	35.2% (34.4%)	$e^{-2} \approx 13.5\%$
3	1340	18.2% (20.5%)	$e^{-3} \approx 5.0\%$
4	809	13.1% (14.0%)	$e^{-4} \approx 1.8\%$
5	690	6.8% (7.9%)	$e^{-5} \approx 0.7\%$
6	482	5.0% (5.1%)	$e^{-6} \approx 0.2\%$
7	401	4.0% (3.3%)	$e^{-7} \approx 0.1\%$
8	292	1.4% (3.0%)	$e^{-8} \approx 0.0\%$
9	259	0.8% (1.0%)	$e^{-9} \approx 0.0\%$
≥ 10	1664	0.2% (0.2%)	$\leq e^{-10} \approx 0.0\%$

Table 1: Outcomes of the analysis of the demand for service parts. The numbers in parenthesis show the outcome when we let the analysis start two years earlier, shifting all time periods by two years.

parts in complex products. By looking at the third and fourth column of Table 1 we conclude that demand drops dead more often than would be expected based on the stationary (compound) Poisson process assumption. By comparing the numbers in the third column with the numbers in parenthesis in that column, we conclude that the obsolescence problem does not appear to be growing or shrinking significantly in two years.

The data confirms that obsolescence of service parts also occurs while the product is still in use. This results in high costs, because any money spent to obtain unused service parts is lost. These costs are costs *in addition to* costs of tied up capital and warehouse costs. An approach sometimes used in industry is therefore to add obsolescence costs to the holding costs of service parts. This implies an assumption of the same risk of obsolescence for all parts. Table 1 shows that this assumption is not very precise: there are large differences in the risk of service parts becoming obsolete even for parts used in the same product type. It seems that slow moving parts become obsolete more often than fast moving parts. In the remainder of this paper, we will devise a method to include this knowledge in an inventory model.

4 The method

In Section 2 we argued that obsolescence can occur while the product is still being used, because of sudden changes in operating conditions, maintenance policy, or the introduction of alternative parts. We will develop a demand model that includes this possibility of obsolescence. Moreover, we will develop a method to extract from Table 1 the parameters needed to apply this method.

Song & Zipkin (1993) provide a framework which incorporates the possibility of sudden changes in demand rate. They let the demand rate in a Poisson demand model depend on the state of an underlying Markov process, representing the 'state of the world'. In light of the discussion in Section 2, their model is attractive because state transitions in their model could correspond to changes in operating conditions, maintenance policy, or the introduction of alternative parts. The memoryless property in the model makes that these transitions cannot be predicted, but when they occur the model allows us to have them affect the demand intensity. The framework by Song & Zipkin leaves much freedom in determining the precise structure of the underlying Markov chain. The framework would allow for a part to visit multiple underlying 'states of the world', before ending up to become obsolescent.

For models in which multiple states have a positive demand rate, parameters can be estimated using techniques from hidden Markov theory (see e.g. Rabiner (1988)). While these techniques have proven powerful, a lot of data is needed for successful application and implementing the hidden Markov algorithms requires a lot of effort.

For the purpose of making the model aware of the risk of stocking slow moving parts, such a model is however not needed. Instead, we propose the use of a simpler, two-state model, with only one state in which the part is moving. In the other state, demand for the part has dropped dead. This last state is assumed to be an absorbing state.

The advantages of such a simple Markov model are threefold. Firstly, it is possible to estimate the parameters for this model in a far simpler manner, and less data suffices for estimating the parameters. Also, the method adds only a single parameter in comparison to more standard demand models, and this parameter has a simple, intuitive interpretation. Finally, when the parameters of the model are known, the optimization of this model is far simpler.

Using a simpler model also brings disadvantages. The two state model described above models decreases in demand, but no increases. If we use this model to accommodate for the increased probability of zero demand in Table 1, the model will predict an overall decrease in demand in period 3 with respect to period 1. In practice, the demand over all parts remains about equal in period 3 with respect to period 1. More states would allow us to model the possibility of significant demand increases, as well as allowing us to keep the total demand (in the model) over all parts equal while modelling the increased variability in the demand of individual parts.

In this contribution, we focus on the simpler model. We believe that while this model still can be improved, it represents a significant step towards application of models including obsolescence. Most importantly, the model improves on the standard Poisson model because of its awareness of the possibility of obsolescence.

The model

We assume that the demand rate depends on the state of a continuous time Markov process. We assume this Markov process has two states x_0 and x_1. In state x_1 the demand is healthy, in state x_0 demand has dropped dead. State x_0 is an absorbing state. Apart from the parameters that govern the demand in state x_1, the Markov process introduces one additional parameter that corresponds to the rate at which the system moves from the first to the second state. This parameter will be denoted by ψ. Furthermore, we assume that as long as the system is in state x_1, demand will follow a (compound) Poisson process with rate λ and compounding distribution D. We assume $P(D > 0) = 1$. The state of the Markov chain at time t will be denoted by $X(t) \in \{x_0, x_1\}$. The demand in the interval (t, t') will be denoted by $C(t, t')$.

Estimating the parameters

If we want to apply the model at the mentioned company we need an estimate for the parameters λ, D and ψ for each part. In this section, we develop such a method. To obtain the parameters

λ and D we use demand data from this forecasting period. λ can be estimated by using the number of orders in this period. D can be fitted by using the mean and the variance of the size of the orders in this period and by subsequently fitting on these values some distribution that is deemed appropriate. Now, the parameter ψ, that can be identified with the short term risk of obsolescence, remains to be determined.

To determine ψ we assume that the short term future behavior of parts with a certain number of orders in the forecasting period will be similar to the short term future behavior of parts with the same number of orders in a similar period in the past. The reason for this assumption is that we have observed that the numbers in the third column of Table 1 do not depend to a great degree on the point in time at which we let the first time interval start.

We will thus use the information we gathered for the different groups in Table 1 to estimate the obsolescence rate ψ. Table 1 however does not give an estimate for the parameter ψ, but an estimate for the probability of zero demand in a certain time interval, i.e. period 3 in Figure 1. To arrive at the parameter ψ we will calculate the probability of zero demand for this same period in our demand model, given the demand in period 1. For ease of notation we will fix the time origin at the end of period 1. We are then interested in the probability that there is no demand in the interval $(t, t+T)$.

We assume that the part is still moving at the end of period 1. This assumption serves as an approximation, because if the orders in period 1 are early in this period, it is possible that the Markov chain has already moved to the state indicating obsolescence. However, taking this into account will mean that the probability of zero demand in period 3 will depend on the moment that the last demand for the part was incurred. This means that the probability will differ for different parts in the same group, which is something that would greatly complicate the estimation of ψ later on.

Based on this assumption, the probability can be calculated to be:

$$\begin{aligned} P(C(t,t+T)=0|X(0)=x_1) &= 1-P(C(t,t+T)>0|X(0)=x_1), \\ &= 1-P(C(t,t+T)>0;X(t)=x_1|X(0)=x_1), \\ &= 1-P(C(t,t+T)>0|X(t)=x_1)P(X(t)=x_1|X(0)=x_1), \\ &= 1-P(C(0,T)>0|X(0)=x_1)e^{-\psi t}. \end{aligned} \tag{1}$$

The first equality follows from the assumption that demand is non-negative. The second equality is obtained by conditioning on $X(t)$ and noting that $P(C(t,t+T)>0;X(t)=x_0)=0$, and the third equality follows from the Markov property. In the last equality we use again the Markov property.

We now need an expression for the term $P(C(0,T)>0|X(0)=x_1)$. We can obtain such an expression by conditioning on the type of the first event after 0. This can either be a transition of the Markov chain to state x_0 (with probability $\psi/(\lambda+\psi)$), or a demand for the service part (with probability $\lambda/(\lambda+\psi)$). In the former case we know that $C(0,T)=0$, in the latter case we have $C(0,T)>0$ if the event occurs before T (we need $P(D>0)=1$). Based on this argument, we have:

$$P(C(0,T)>0|X(0)=x_1)=\frac{\lambda}{\lambda+\psi}\left(1-e^{-(\psi+\lambda)T}\right).$$

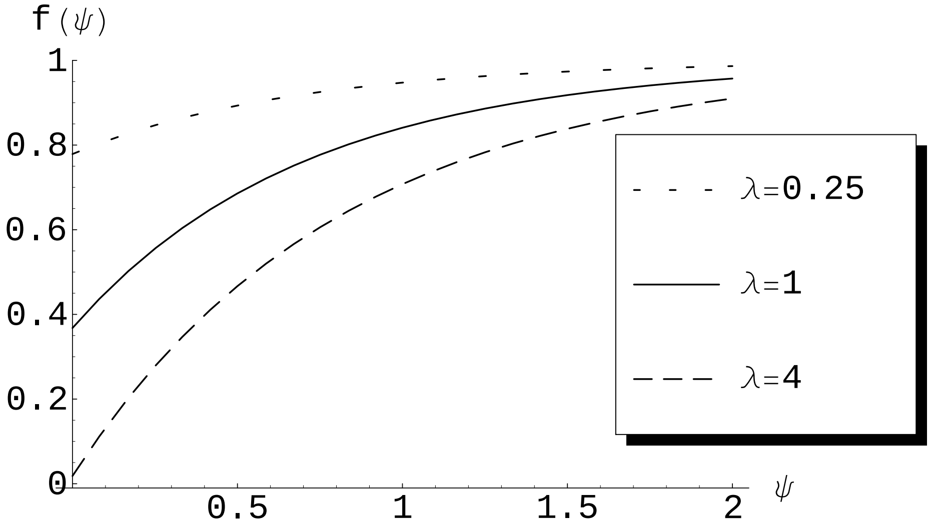

Figure 2: The function $f_{t,T,\lambda}(\psi)$, for different values of λ, as given by (2). We fix $T = t = 1$.

Using this expression in (1) we have:

$$P(C(t, t+T) = 0 | X(0) = x_1) = 1 - \frac{\lambda}{\lambda + \psi}\left(1 - e^{-(\psi+\lambda)T}\right) e^{-\psi t}.$$

We now consider the following set of functions, indexed by $t, T, \lambda \in (0, \infty)$:

$$f_{t,T,\lambda} : [0, \infty) \to \mathbb{R} : \psi \mapsto 1 - \frac{\lambda}{\lambda + \psi}\left(1 - e^{-(\psi+\lambda)T}\right) e^{-\psi t}. \tag{2}$$

Some of these functions are plotted in Figure 2. We want to use the inverse of these functions in conjunction with the information in Table 1 to get an estimate for ψ. From this figure, it seems clear that the functions have a uniquely defined inverse. To prove this, we need the following lemma:

Lemma 1. *For every $t, T, \lambda \in (0, \infty)$, the following holds.*

(i) *The function $f_{t,T,\lambda}$ is continuous on its domain $[0, \infty)$ and differentiable on $(0, \infty)$.*

(ii) *$f_{t,T,\lambda}(0) = e^{-\lambda T}$, and $\lim_{\psi \to \infty} f_{t,T,\lambda}(\psi) = 1$.*

(iii) *For any $\psi, \psi' \in [0, \infty)$ with $\psi < \psi'$, we have $f_{t,T,\lambda}(\psi) < f_{t,T,\lambda}(\psi')$.*

Proof of Lemma 1. Continuity and differentiability immediately follow from the fact that the function is composed of functions that are continuous and differentiable. For (ii), the value of f at $\psi = 0$ can be easily checked. The limit value can be obtained by checking the limit values of the individual terms. We will prove (iii) by showing that

$$\frac{\lambda}{\lambda + \psi}\left(1 - e^{-(\psi+\lambda)T}\right) \tag{3}$$

is strictly decreasing. This can be checked by checking the derivative:

$$\begin{aligned}
&\frac{\partial}{\partial\psi}\left(\frac{\lambda}{\lambda+\psi}\left(1-e^{-(\lambda+\psi)T}\right)\right)\\
&\quad=\frac{\lambda T}{\lambda+\psi}e^{-(\lambda+\psi)T}-\frac{\lambda}{(\lambda+\psi)^2}\left(1-e^{-(\lambda+\psi)T}\right)\\
&\quad=\frac{\lambda}{(\lambda+\psi)^2}e^{-(\lambda+\psi)T}\left(T(\lambda+\psi)-e^{(\lambda+\psi)T}+1\right)\\
&\quad=\frac{\lambda}{(\lambda+\psi)^2}e^{-(\lambda+\psi)T}\left(\sum_{i=2}^{\infty}\frac{-\left(T(\lambda+\psi)\right)^i}{i!}\right).
\end{aligned}$$

The derivative is negative because it is the product of a strictly positive function and a (converging) sum of strictly negative terms. (3) is thus decreasing because it is continuous and it has a negative derivative. □

Based on Lemma 1, we have the following theorem.

Theorem 1. *The function $f_{t,T,\lambda}$ has a unique inverse*

$$f^{-1}_{t,T,\lambda} : [e^{-\lambda T}, 1) \to [0, \infty) : p \mapsto f^{-1}_{t,T,\lambda}(p). \tag{4}$$

In particular, $f_{t,T,\lambda} \circ f^{-1}_{t,T,\lambda}$ is the identity function on $[e^{-\lambda T}, 1)$.

Proof of Theorem 1. Existence follows from (i) and (ii) of Lemma 1 and the intermediate value theorem. Uniqueness follows from (iii) of Lemma 1. □

We were not able to find a closed form formula for the inverse function given by (4), but the function $f_{t,T,\lambda}(\psi)$ as well as its derivative can be evaluated for every $\psi \in (0, \infty)$. We were thus able to numerically evaluate the function given by (4) using the Newton-Raphson method (Press, et al. 2007, Section 9.4).

Now we turn back to the problem of determining the parameter ψ for the different parts. For the groups given in Table 1 we have an estimate for the probability of zero demand in period 3: the observed fraction with zero demand in each group. Depending on the size of the groups and possible dependence issues this estimate is more or less accurate.

Because most groups are quite large, we will use the fraction of parts with zero demand as an estimate for the probability of zero demand, and we use as an estimate for ψ the unique value that gives exactly this estimated probability of zero demand, defined by the function given in (4). An estimate for the value of ψ obtained in this way is given in in Table 2.

We are thus able to assign a value of ψ to the different groups, based on the number of parts that dropped dead within these groups. When forecasting to determine a stock policy, we will assign these values to the parts based on the number of orders that these parts have in the forecasting period.

There is, however, one issue with determining a value for ψ for different groups. This is that the function $f^{-1}_{t,T,\lambda}(p)$ is only defined for $p \in [e^{-\lambda T}, 1)$. If $p < e^{-\lambda T}$, we can thus not readily give an estimate for ψ. Note that this only happens if less parts have zero demand in period 3 than would be expected based on the Poisson assumption. This would indicate that, at least in our

Number of orders in period 1	Number of parts	Fraction of parts with no demand in period 3	$f^{-1}_{t,T,\lambda}(p)$ (see (4))
1	5630	57.5%	$f^{-1}_{1,2,0.5}(0.554) \approx 0.22/\text{yr}$
2	2434	35.2%	$f^{-1}_{1,2,1.0}(0.316) \approx 0.17/\text{yr}$
3	1340	18.2%	$f^{-1}_{1,2,1.5}(0.196) \approx 0.10/\text{yr}$
4	809	13.1%	$f^{-1}_{1,2,2.0}(0.097) \approx 0.08/\text{yr}$
5	690	6.8%	$f^{-1}_{1,2,2.5}(0.090) \approx 0.05/\text{yr}$
6	482	5.0%	$f^{-1}_{1,2,3.0}(0.057) \approx 0.04/\text{yr}$
7	401	4.0%	$f^{-1}_{1,2,3.5}(0.025) \approx 0.03/\text{yr}$
8	292	1.4%	$f^{-1}_{1,2,4.0}(0.009) \approx 0.01/\text{yr}$
9	259	0.8%	$f^{-1}_{1,2,4.5}(0.016) \approx 0.01/\text{yr}$
≥ 10	1664	0.2%	*

Table 2: Estimates of the obsolescence risk ψ by the model, using the data for the groups presented in Table 1.

framework, obsolescence is not a problem. In that case we set $\psi = 0$, and the model reduces to the compound Poisson model.

Another issue is that we cannot calculate a value for ψ for the group consisting of parts with 10 or more orders, because there is no single value for λ available. We could solve this problem by making different groups for 10, 11, ... orders but the groups will become quite small. This means that statistical deviation will become more and more important, and results obtained with the method will have less and less value. However, for groups with a large number of orders obsolescence is not a big issue, as only 0.2% of the parts did not have any demand in the second period. We therefore set $\psi = 0$ for parts with 10 or more orders.

Implementation

The precise manner in which the method was used is not the focus of the paper. We will however give a short overview of the implementation at the company, because it offers some insights in the value of the method.

As discussed in this paper, we obtained estimates for λ, D and ψ for each part. Also, shortage costs, ordering costs, holding costs, and obsolescence costs were defined in consultation with the management. The former three costs are relatively standard, the latter costs are the costs of parts becoming obsolete. At the particular company these costs consisted of the costs to obtain the service parts and the costs of scrapping them. Subsequently, recommendations for the reorder point and the order-up-to point were given by minimizing the total expected costs. This resulted in reorder points and order up to points that could be imported in the ERP system. These recommendations were followed most of the time by the inventory controllers.

In comparison to the approach where the obsolescence costs are spread evenly over all parts by including them as a constant factor in the holding cost, the model has a clear advantage. This advantage lies in the fact that the model knows that stocking slow moving parts is more costly than stocking faster moving parts, because of the higher risk of obsolescence. In comparison to the simpler approach, the model will thus stock more faster moving parts, and less slower moving parts. This improvement was also recognized by the inventory controllers. Simulation results

		Cost estimates of naive model			Cost estimates of sophisticated model			
Part	Stock	Holding	Back order	Total	Holding	Obsol.	Back order	Total
S	2	1104	7566	8669	883	1700	7566	10149
S	3	2023	1704	3727	1619	2550	1704	**5872**
S	4	3004	317	**3322**	2403	3400	317	6121
S	5	4001	50	4051	3201	4250	50	7501
F	12	5049	3610	8659	4040	0	3610	7649
F	13	6022	1639	**7661**	4818	0	1639	6457
F	14	7010	704	7713	5608	0	704	**6311**
F	15	8004	286	8290	6403	0	286	6689

Table 3: The costs calculated for part S and part F using both the naive and the sophisticated model. The minimal costs for both parts according to both models are indicated in bold.

using real demand data on which we will not report in detail also indicate that including the risk on obsolescence improves the recommendation.

Illustration of the advantage of the method

In order to illustrate the manner in which the knowledge of obsolescence risk can improve the recommendations given by the model, we will give some results on two hypothetical, but realistic parts S and F. Both parts have a price of 4000, and the costs for a part becoming obsolete are 5000, equal to the price of the part plus some cost for scrapping it. Both parts have a leadtime of 1 year. Both parts are demanded only in quantity 1, so we assume Poisson demand for both. We assume full back-ordering, and a back order cost of 365×200 per part per year. For simplicity, we assume a base stock policy. Now, part S is a slow mover, as it has had 2 orders in the last two years, while part F has had 14 orders in the last two years.

We proceed to find cost estimates for different reorder points according to two different models. The *naive* model is a model in which holding costs of 25% are taken into account for both parts, in which 5% obsolescence cost is naively included. This gives an annual holding cost of $25\% \times 4000 = 1000$ for both parts. We assume no holding costs are only paid for parts in the pipeline.

In the *sophisticated* model we take into account the obsolescence cost in a sophisticated manner. Using Table 2, we obtain the obsolescence rate $\psi_S = 0.17$ for part S, and the obsolescence rate $\psi_F = 0$ for part F. Because we include the obsolescence costs in a more sophisticated way, we leave out the 5% obsolescence cost in the holding cost and work with a holding cost of 20%. This gives us a holding cost of $20\% \times 4000 = 800$ for both parts. Based on the obsolescence risk, we can calculate the expected lifetime of the part. At the end of this lifetime, the parts on stock or on order will have to be scrapped. By dividing the total scrapping costs over the expected number of years until obsolescence, annual obsolescence costs can be calculated. All other costs are also computed as the average annual costs until the moment of obsolescence, by using the steady state distribution of the inventory position and the properties of the Poisson process.

In Table 3, we show the costs estimates of stocking different quantities according to the two models. Both models will give a recommendation for the reorder point by minimizing their cost estimates.

The sophisticated model is aware of the high risk of stocking on slow moving parts. Therefore, it decides to stock only 3 on the slow moving part (S). The naive model will stock 4 for this part, ignoring the high risk of obsolescence. Based on the sophisticated model, we estimate that the additional costs for ignoring this risk are 249 on average annually, until the moment of obsolescence.

Something similar happens with the faster moving part (F). The sophisticated model knows there is no significant risk of this item becoming obsolescent, and therefore stocking on the part is relatively cheap. It will therefore stock 14 for this part. The naive model uses a higher obsolescence cost for this part, not knowing that this part will probably never become obsolescent. Therefore, it stocks to conservatively, which will cost an additional 146 annually based on the estimate by the sophisticated model.

5 Conclusions and extensions

We have presented a method that can be used to estimate the risk of obsolescence using demand data. The method is based upon observations in the demand data of service parts that are used in products with a long life cycle. In principle, the method can be applied by any company with sufficient data for a sufficient number of parts, and products with long life cycles. However, more research is needed to find out if other companies have similar demand patterns for service parts. In particular, it would be interesting to find out whether a similar analysis as the one used in Section 3 gives similar results at other companies, in the sense that the number of parts in each group that have zero demand in the second period exceed the number of parts that should have zero demand according to the Poisson model. The method was implemented at the company, and the resulting order suggestions were in general followed by the inventory controllers.

It would be interesting to extend the method to Markov models with more than 2 states, examples of which are considered by Song & Zipkin (1996). While this allows us to model demand increases as well as decreases, multiple states will greatly complicate the estimation of the model parameters from the demand data. The theory of hidden Markov models (see e.g. Rabiner (1988)) might prove useful in this respect. Another direction for future research is getting rid of the assumption that the part is still moving at the end of period 1, while still retaining a simple procedure to estimate ψ.

References

G. W. Brown, et al. (1964). 'Dynamic modelling of inventories subject to obsolescence'. *Management Science* **11**:51–63.

K. D. Cattani & G. C. Souza (2003). 'Good buy? Delaying end-of-life purchases'. *European Journal of Operations Research* **146**:216–228.

K. Cobbaert & D. van Oudheusden (1996). 'Inventory models for fast moving items subject to "sudden death" obsolescence'. *International Journal of Production Economics* **44**:239–248.

R. M. Hill, et al. (1999). 'Stock replenishment policies for a stochastic exponentially-declining demand process'. *European Journal of Operational Research* **116**:374–388.

J. R. Moore (1971). 'Forecasting and scheduling for past-model replacement parts'. *Management Science* **18**:B200–B213.

W. H. Press, et al. (2007). *Numerical recipes.* Cambridge University Press, 3 edn.

L. R. Rabiner (1988). 'A tutorial on hidden Markov models and selected applications in speech recognition'. *Proceedings of the IEEE* **77**:256–286.

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

J. Song & P. H. Zipkin (1993). 'Inventory control in a fluctuating demand environment'. *Operations Research* **41**:351–370.

J. Song & P. H. Zipkin (1996). 'Managing inventory with the prospect of obsolescence'. *Operations Research* **44**:215–222.

Y. Song & H. C. Lau (2004). 'A periodic-review inventory model with application to the continuous review obsolescence problem'. *European Journal of Operational Research* **159**:110–120.

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

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

J. P. van Kooten & T. Tan (Forthcoming). 'The final order problem for repairable spare parts under condemnation'. *Journal of the Operational Research Society* .