University of Salford MANCHESTER

Salford Business School

global thinking, sustainable practice, professional success

Working Paper: 342/11

SPARE PARTS MANAGEMENT: A DISTRIBUTION-BASED APPROACH

Lengu, D.	University of Salford, UK
Syntetos, A.A.	University of Salford, UK
Babai, M.Z.	BEM Bordeaux Management School, France

Salford Business School Working Paper (WP) Series

ISSN 1751-2700

Spare parts management: a distribution-based approach

Purpose

The management of spare part inventories is an issue of strategic concern for most industrial firms. Spare parts are known to have intermittent demand patterns and such patterns cause considerable problems with regards to forecasting and stock control due to their compound nature that renders the normality assumption invalid. A number of authors have suggested that compound distributions could be used to model intermittent demand patterns. There is however a lack of theoretical analysis and little relevant empirical evidence in support of these distributions. In this work, we are concerned with the value and the challenges involved in the use of compound distributions in an inventory management context.

Design/Methodology/Approach

We conduct a detailed empirical investigation on the goodness of fit of various compound distributions and we develop a distribution-based demand classification framework the validity of which is also assessed in empirical terms. The datasets used for the purposes of our research consist collectively of the individual demand histories of over 30,000 Stock Keeping Units (SKUs) from five different industrial sectors.

Findings

Our empirical investigation provides evidence in support of certain demand distributions in a real world context. A great number of empirical insights and implications are also offered serving a two-fold contribution: i) opening up the agenda for further research in this area; ii) assisting inventory and operations managers to make more informed decisions.

Research and Practical Implications

The work described in this paper should facilitate the task of selecting demand distributions in a real world spare parts inventory context. A conceptually appealing demand classification scheme is proposed and a linkage is also provided between the scheme and the potential qualitative attributes of SKUs falling with each category. We are also building on previous work to propose a comprehensive list of criteria to be used when selecting demand distributions. Finally, an extensive discussion is provided on parameter estimation related difficulties in this area. As such, our work should enable considerable further theoretical developments in the area of spare parts management and should also successfully inform relevant real world practices.

Originality/Value

This is an original large scale empirical study in a very important yet under-researched area in the Operations and Production Management literature. Technical developments are conceptually linked to qualitative considerations and the paper is intended to inform both the quantitative and the qualitative Operations Management academics and practitioners. To that end, all technicalities are kept away from the main body of the paper and they are separately presented in two Appendices to enable an uninterrupted coverage of the work by the non-quantitative audience of the Journal.

Classification: Research Paper

Keywords: Inventory Forecasting; Demand Distributions; Spare Parts Management

1. Introduction

Many industries rely on the effective management of spare parts, including aerospace and defence, transportation, telecommunications and information technology, utilities and durable goods suppliers. Spare parts are held by firms for internal use in the maintenance of tools and equipment. They are also held by suppliers at the retail or wholesale supply chain level for sale to customers. The costs associated with the inventory management of spare parts can be substantial. According to US Bancorp spare parts relate to a $700 billion annual expenditure that constitutes about 8 percent of the U.S. gross domestic product (Jasper, 2006). Given the very high level of inventory investments, it is clear that there is significant opportunity for cost-savings through better management.

The demand of spare part items is typically intermittent with demand orders arriving sporadically; the demand can also be highly variable as well as intermittent, in which case it is referred to as lumpy (Boylan and Syntetos, 2008). Kalchschmidt *et al.* (2006) have also defined lumpy demand as:

- variable, and therefore demand characterized by relevant fluctuations;
- sporadic, because the demand series are characterized by many periods of very low or no demand; and
- ‘nervous’, reflecting the low auto-correlation of the demand.

The area of inventory management has received a lot of attention in the Operations Management (OM) literature. Conventional inventory control approaches rely on a number of assumptions that are usually valid when demand is fast-moving. Demand over lead time is assumed to be normally distributed and standard forecasting methods are used to estimate the parameters of the normal distribution (see, for example, Adenso-Díaz, 1996). However, it has long been shown that such an assumption is invalid in a spare parts context where demand is

usually intermittent (Mak and Hung, 1993; Botter and Fortuin, 2000). Moreover, the intermittent nature of the demand makes it very difficult to forecast future requirements with much accuracy (Fortuin and Martin, 1999). This problem is exacerbated when the replenishment lead times are long. Blumenfeld *et al.* (1999) have demonstrated, amongst others, that the longer the lead times are, the higher the levels of inventory required in order to accommodate the demand uncertainty. Forecasting is an integral part of inventory management systems (Fildes and Beard, 1992). However, the challenges in forecasting intermittent demand have implications beyond inventory control; demand forecasts are also used in product development, production and supply chain planning.

Another important issue involved in inventory management is the categorisation of inventory items for the purpose of facilitating forecasting and stock control. When there is a large number of Stock Keeping Units (SKUs) it is not practical to evaluate them on an individual basis. In such cases, the SKUs will typically have to be categorised in order to facilitate decision-making and allow managers to focus their attention on the most important SKUs (however this is judged) (Teunter *et al.*, 2010). There have been a number of studies in the area of demand classification for inventory items with intermittent demand, including Partovi and Burton (1993), Botter and Fortuin (2000) and Syntetos *et al.* (2009). Recently, Syntetos *et al.* (2005) have proposed a classification scheme that categorises SKUs in a manner that facilitates demand forecasting.

The main objective of this study is to advance the current state of knowledge in inventory management of spare parts by bringing together the issues of distributional assumptions, forecasting and SKU classification. These issues will be linked together by using compound distributions to model demand during lead time. A number of authors (including Friend, 1960 and Kemp, 1967) have suggested that compound distributions (compound Poisson

distributions in particular) may provide a good fit for the demand distributions of such SKUs. Compound distributions are appealing because their underlying structure is similar to the demand-generating process associated with intermittent demand.

A top down approach will be used in order to identify compound distributions that may accommodate the distributional properties observed among SKUs with intermittent demand. Firstly, we will consider the shapes that frequency distributions of order sizes will usually take in an intermittent demand context. We will then propose a number of probability distributions that could be used to model such order sizes. Finally, we will introduce the assumption that demand orders arrive according to a Poisson process and, by bringing together the proposed order size distributions and the Poisson arrival process, we will obtain compound distributions that may be used to model intermittent demand. As part of this process, we also develop a demand classification framework. The categorisation in this framework will be motivated by a conceptual understanding of the distributional properties of the order sizes rather than a theoretically consistent match of every possible SKU in a particular category. This approach is different from the bottom up approaches that have previously been introduced in the area of intermittent demand management (for example by Syntetos *et al*., 2011). In the latter approaches, goodness of fit tests are first carried out for individual SKUs and the results of these tests are used towards the development of a possible classification scheme.

Our study also makes a number of further important contributions in the area including: (i) empirical analysis in order to assess whether compound distributions provide a good fit for spare part SKUs; (ii) highlighting a number of challenges related to parameter estimation and goodness-of-fit testing in the area of intermittent demand management; (iii) the development

of criteria that should be used when selecting distributions for modelling demand; (iv) deriving insights for practitioners and setting an agenda for further research..

The remainder of this paper is structured as follows. In the next section, we shall provide a brief overview of the literature on inventory management related issues for SKUs with intermittent demand. A detailed description of the demand classification scheme proposed by Syntetos *et al.* (2005) will also be given in this section. Compound distributions that may model the distributional properties associated with intermittent demand are considered in section 3. In that section we will also bring together the proposed compound distributions and the demand classification scheme proposed by Syntetos *et al.* (2005). In doing so, we will propose a new demand classification framework that may facilitate the process of the joint selection of the most appropriate compound Poisson distributions and forecasting methods. The empirical goodness of fit of these distributions is then assessed in section 4 on an extensive dataset of spare parts. We will also compare the relative levels of fit achieved by the compound Poisson distributions in the different categories of the proposed framework; this exercise allows us to assess the empirical validity of the proposed framework with respect to the selection of demand distributions. The practical and theoretical implications of our study are discussed in section 5. Finally, in section 6, we will provide the conclusions of this study and also identify a number of areas of future research.

2. Research background

There are a number of authors that have proposed strategies for managing SKUs with intermittent demand. Most of these strategies require, at some stage, that the SKUs are categorised. Fortuin and Martin (1999) proposed a framework based on which SKUs were categorised according to the price, the consumption (or usage) rate and the response times for repair (as agreed upon in the service contracts). The demand classification schemes proposed

by Williams (1984) and Eaves (2002) categorised SKUs according to the rate of demand arrival, the demand size variability and the lead time variability. Williams (op. cit.) also proposed a set of distributions that could be used to model the demand of the SKUs falling in the different categories. The ABC classification scheme has also been used by some authors (including Partovi and Burton, 1993, and Syntetos *et al.*, 2009) in the context of intermittent demand. In all these studies, the boundaries between the categories were decided based on the characteristics of the particular sample used in each study. It is therefore questionable whether any of these schemes would be effective when applied to data sets other that those considered during their development.

The purpose of a demand classification scheme is to assist inventory managers to identify the forecasting methods and inventory control policies that are more appropriate for the various SKUs included in a stock base. As Syntetos *et al.* (2005) have demonstrated, substantial improvements in forecasting accuracy may be obtained by first categorising SKUs with similar demand patterns and then identifying the best performing forecasting techniques in each category. Johnston and Boylan (1996) have also argued that, given that one of the purposes of such schemes is to facilitate forecasting, the logical approach would be first to compare alternative forecasting methods for the purpose of establishing regions of superior performance. The results of this comparison should then be used to inform the demand categorisation scheme. Johnston and Boylan (op. cit.) examined the conditions under which Single Exponential Smoothing is more accurate than Croston's method. According to Croston's method, the forecast is built from constituent elements of demand, namely, the demand size when demand occurs and the inter-demand interval (denoted by Z and p, respectively in this paper). Johnston and Boylan concluded, on the basis of simulation under a wide range of conditions, that Croston's method is more accurate (in terms of simulated

Mean Square Error, MSE) when the average inter-demand interval ($\overline{p}$) exceeds 1.25 review periods. The key contribution of this study, as far as classification is concerned, was the identification of the average inter-demand interval as a demand classification parameter.

Syntetos *et al.* (2005) built on the work of Johnston and Boylan (1996) by carrying out a formal mathematical analysis of the theoretical MSEs of three different methods: Single Exponential Smoothing, Croston method and the Syntetos-Boylan Approximation (SBA, Syntetos and Boylan, 2005) which is a bias-adjusted modification of Croston's method. Syntetos *et al.* (2005) identified the average inter-demand interval ($\overline{p}$) and the squared coefficient of variation of the demand sizes ($CV^2(Z)$) as two of the parameters that determined the relative performance of the three methods. Based on comparisons of the theoretical MSEs, they recommended that: (a) Croston's method should be used when $\overline{p} < 1.32$ and $CV^2(Z) < 0.49$, (b) the SBA should be used for all other combinations. This study, however, was based on the assumption that demand arrives according to a Bernoulli process and therefore the suggested cut-off boundaries between the methods may not necessarily apply under other demand arrival processes. The scheme (illustrated in Figure 1 below) provides operationalised definitions of demand patterns; SKUs are defined as Smooth, Erratic, Intermittent or Lumpy based upon the observed values of $\overline{p}$ and $CV^2(Z)$. For the remainder of the paper, we shall refer to this scheme as the *SBC scheme* (after Syntetos, Boylan and Croston).

The SBC scheme is associated with two features that distinguish it from the other schemes discussed above. Firstly, the SBC scheme takes demand forecasting into account, thereby bringing together the demand classification and demand forecasting aspects of inventory management. Second, the categories in the scheme were derived based on the theoretical

comparative performance of forecasting methods and, as a result, they are generally applicable. The categories and the cut-off points in the scheme remain the same regardless of the data set, albeit conditioned to the three forecasting methods considered in that study and the assumption of a Bernoulli demand generation process.

Figure 1. Syntetos *et al.* (2005) classification scheme (SBC)

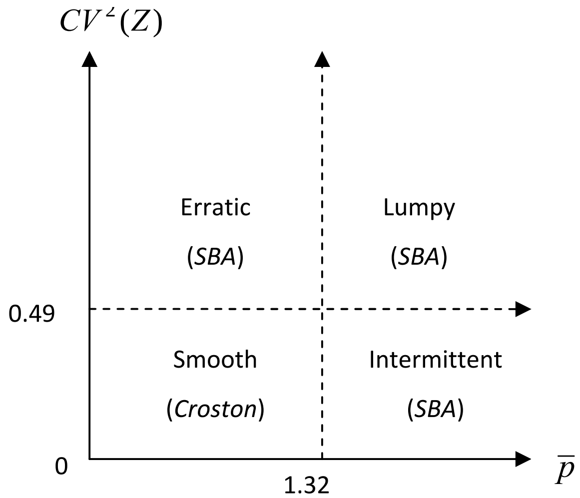

However, the SBC scheme does not offer any suggestions in relation to the probability distributions that could be used to model the SKUs falling in the different categories. In parametric approaches to stock control, a distribution for the demand during lead time will usually have to be specified in order to derive the stock control parameters. In this study, we will build on the SBC scheme by proposing compound distributions that could be used to model the demand of the SKUs falling in the different categories. The classification scheme developed in this study will therefore take into account the distributional patterns of the demand as well as the forecasting accuracy of various relevant estimators.

It is important to note that there are some non-parametric approaches that do not rely upon any distributional assumptions. Such distribution-free approaches include bootstrapping (e.g. Willemain *et al.*, 2004) and artificial neural networks (e.g. Gutierrez *et al.*, 2008). This study however will focus only on parametric approaches based on compound distributions.

3. Compound distributions

3.1 Compound Poisson distributions

In the context of intermittent demand, the demand arrival process can be reasonably modelled as a Bernoulli one if time is treated as a discrete variable. The Bernoulli process models whether an order arrives over any given unit time period or not. Demand orders arriving during each unit period are 'bucketed' and the aggregate demand over that period is known as the demand size. If demand arrives according to a Bernoulli process, then the inter-demand intervals will have a geometric distribution. Croston (1972) and Syntetos *et al.* (2005), among others, have modelled the demand occurrence process as a Bernoulli one.

If time is treated as a continuous variable, then demand arrival can be modelled as a Poisson process. The Poisson process models the arrival of individual demand orders; the orders are therefore not 'bucketed'. As a result, the Poisson process captures more information about the demand occurrence than the Bernoulli one. If demand arrives according to a Poisson process, then the intervals between transaction order arrivals will have an exponential distribution.

In this paper, we will assume that demand arrives according to a Poisson process; furthermore, we will assume that the order sizes (also known as 'transaction sizes') are distributed according to some arbitrary distribution. The distribution of demand during a fixed period of time will then have a compound Poisson distribution. Let us assume that demand has a compound Poisson distribution and let us denote sizes of the demand orders (i.e. transaction sizes) as X. Finally, let:

λ = the order arrival rate

$\mu = \mathrm{E}(X)$ = the mean of the transaction sizes

$\sigma^2 = \mathrm{Var}(X)$ = the variance of the transaction sizes

Y = the demand during a unit period of time.

Then the mean and variance of demand during a unit period of time are given respectively by (Satterthwaite, 1942):

$$E(Y) = \lambda\mu \tag{1}$$

$$Var(Y) = \lambda\left(\mu^2 + \sigma^2\right) \tag{2}$$

One of the appealing properties of compound Poisson distributions is that they are Lévy processes and, as such, they are infinitely divisible (Sato, 1999). Furthermore, a linear combination of a finite number of independent Lévy processes is again a Lévy process. The practical implication of this property is that, if the demand over a unit period of time (denoted as Y) is assumed to have a compound Poisson distribution, then the demand over a fixed period of length L (where L is a positive rational number) will also have a compound Poisson distribution.

A number of authors (e.g. Friend, 1960; Croston, 1972) have advocated the use of compound distributions to model intermittent demand patterns. The appeal of compound distributions stems from the fact that they can independently model the constituent elements of demand (transaction sizes and intervals between order arrivals). Feeney and Sherbrooke (1966) derived a simple analytic solution of the order-up-to level (under a base-stock policy) when demand follows a compound Poisson distribution. Ward (1978) proposed a regression model for calculating the reorder points of lumpy items. Watson (1987) examined the interactions between forecasting and inventory control in such a context. In the last two studies, demand was assumed to arrive according to a Poisson process and the transaction sizes were assumed

to have a geometric distribution. Other researchers that have used a compound Poisson distribution to model intermittent demand include: Adelson, 1966; Archibald and Silver, 1978; Naddor, 1978; Mitchell *et al.*, 1983; Matheus and Gelders, 2000; Özkaya *et al.*, 2006. In this study, we will make two main contributions to this body of work. Firstly, we will carry out empirical tests of goodness of fit for a number of compound Poisson distributions. Secondly, we will identify some of the challenges related to the use of compound Poisson distribution in parametric inventory management.

3.2 Transaction size distributions

In this section, we shall consider a number of distributions that could be used to model the transaction sizes. The term "transaction size" refers in this paper to the number of units in a distinct demand order. The term should not be confused with "demand size", which is the total numbers of units ordered during a given period of time. The distributions used to model transaction sizes should ideally provide good empirical fit but they should not be computationally demanding for use in practical settings. For SKUs with intermittent demand, the transaction size frequency distributions observed in practice are usually either monotonically decreasing or more centred (i.e. "mounded") but with a significant right skew. Boylan (1997) proposed three criteria for assessing the suitability of hypothesised demand distributions (regardless of the context of application): (a) A priori grounds for modelling demand, (b) The flexibility of the distribution to represent different types of demand, (c) Empirical evidence available in support of the distribution. The same criteria were adopted in this paper when selecting transaction sizes distributions. The three criteria are discussed in more detail below.

The first criterion (a priori grounds for modelling demand) relates to the intuitive appeal that a distribution may (or may not) have for representing the data under consideration. The

hypothesised distribution has to match the underlying structure of the transaction sizes, as understood by inventory managers. By their nature, transaction sizes are discrete and they have to be greater than zero. These properties would suggest that the hypothesised transaction size distributions should ideally be discrete distributions that are defined in the positive domain.

Flexibility (the second criterion) refers to robustness in terms of the ability of the distribution to cope with diverse transaction size profiles. For practical purposes, it would be more convenient to have a manageably small number of distributions that are collectively robust enough to cover a great majority of possible empirical scenarios. The third criterion requires that there should be corroborative empirical evidence, where possible, in support of the selected distributions. Unfortunately, there have only been a few empirical studies on the goodness of fit of distributions for intermittent demand items (specifically, Kwan, 1991; Eaves, 2002; Syntetos *et al.*, 2011). The findings of those studies will be used to inform our selection of the transaction size distributions.

In this paper, we add a fourth criterion – the selected distribution should have a probability distribution function that is computationally easy to work with in practice. The distribution functions' moments and the parameter estimates of the selected distributions should take functional forms that can be computed easily and quickly. The distributions should also have as few parameters as possible (ideally, one or two), otherwise it becomes harder for practitioners to get a good grasp of the relationship between the parameters and the probabilities or any statistics of interest.

The four transaction size distributions considered in this paper are the Geometric, Logarithmic series, Poisson and Pascal distributions. These distributions, for the most part,

satisfy the four criteria outlined above. All four distributions are discrete, and the Geometric and Logarithmic series distributions take only positive values. The distributions are also flexible in the sense that different levels of skewness may be obtained for all four of them by adjusting the parameters accordingly. As far as we are aware, as yet, there have been no goodness of fit studies carried out specifically for order sizes. However, the compound representations associated with the Geometric and Logarithmic series distributions have been found to provide good fit for demand during lead time (Syntetos et *al*., 2011). The probability functions of all four distributions can be easily computed in practice. All of the distributions have one parameter except the Pascal distribution which has two parameters.

As will be shown below, the four distributions are also associated with varied and distinct properties that naturally suggest a framework for classifying SKUs. Table 1 shows the properties of the distributions with respect to modality and variability (as measured by the squared coefficient of variation). The mode ($\tilde{m}(X)$) and the squared coefficient of variation ($CV^2(X)$) are arguably two statistics that may collectively best describe the shape of a transaction size distribution. The mode will help us determine whether the transaction size distribution is monotonically decreasing or more 'mounded'. The squared coefficient of variation will give us an idea about the relative spread of the distribution.

Table 1. Summary of the properties of the transaction size distributions

Transaction size distribution			**Compound Poisson distribution**
Name	$\tilde{m}(X)$	$CV^2(X)$	**Name**
Geometric	$\tilde{m}(X)=1$	$0 < CV^2(X) < 1$	Pólya-Aeppli (Stuttering Poisson)
Log series	$\tilde{m}(X)=1$	$0 < CV^2(X) < \infty$	Negative Binomial
Poisson	$\tilde{m}(X)\geq 1$	$0 < CV^2(X) \leq 1$	Neyman Type A
Pascal	$\tilde{m}(X)\geq 1$	$0 < CV^2(X) < \infty$	Poisson-Pascal

For each of the four distributions, the corresponding compound Poisson distribution is given in the final column of Table 1. Based on these differences, we developed a classification scheme (illustrated in Figure 2 below) according to which SKUs are categorised based on $\tilde{m}(X)$ and $CV^2(X)$. The distributions were assigned to the various categories sequentially. The definitions of the categories are given below and the characteristics of the SKUs that would fall in the different categories, as well as illustrative example for a trucking spare parts supplier, are given in Table 2:

Figure 2. Demand classification based on the properties of the transaction sizes

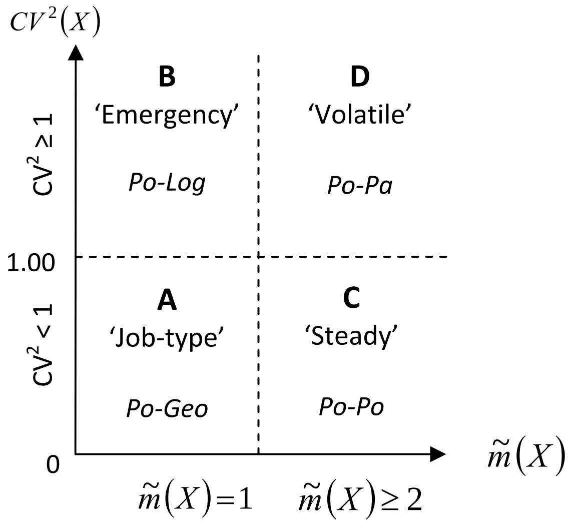

a) Category A ('job-type' items) – Contains all SKUs with $\tilde{m}(X)=1$ and $CV^2(X)<1$. The demand for each of these SKUs may be assumed to follow a Poisson-Geometric distribution (also known as the Pólya-Aeppli distribution). All four distributions could be used to model the transaction sizes in this category. However, the Geometric distribution was preferred to all the alternatives because it fully meets the criteria set out above. Unlike the Poisson and Pascal distributions which can take a value of zero, the Geometric distribution only takes positive values. Syntetos *et al.* (2011) also

found that the Poisson-Geometric distribution provided higher levels of frequency of fit than the Poisson-Logarithmic series distribution.

b) Category B ('emergency' items) – Contains all SKUs with $\tilde{m}(X)=1$ and $CV^2(X)\geq 1$. The demand for each of these SKUs may be assumed to follow a Poisson-Logarithmic series distribution (also known as the Negative Binomial distribution). While the Pascal distribution could also have been used to model the transaction sizes in this category, the Logarithmic series distribution was preferred because it only has one parameter (unlike the Pascal distribution which has two) and there is empirical evidence in its support (in Kwan, 1991; Syntetos *et al.*, 2011).

c) Category C ('steady' items) – Contains all SKUs with $\tilde{m}(X)>1$ and $CV^2(X)<1$. The demand for each of these SKUs may be assumed to follow a Poisson-Poisson distribution (also known as the Neyman type A distribution). The Pascal distribution could also be used to model the transaction sizes in this category but the Poisson distribution was preferred instead because it has only one parameter.

d) Category D ('volatile' items) – Contains the SKUs that have not been assigned to the three other categories. For each of the SKUs in this category, the demand will be assumed to follow a Poisson-Pascal distribution.

It should be noted that the Poisson-Pascal distribution cannot have, simultaneous, $\tilde{m}(X)>1$ and $CV^2(X)\geq 1$. Transaction size distributions are typically right-skewed and, as such, the mean is larger than the mode. $CV^2(X)$ is calculated by dividing the variance by the square of the mean. $CV^2(X)$ may therefore be raised by increasing the variance. However, if the mode is also increased at the same time, then the mean will increase as well thereby countering the rise in $CV^2(X)$. There are very few distributions that may have both $\tilde{m}(X)>1$ and

$CV^2(X) \geq 1$ (among them, the Beta-Binomial and the Beta-Negative Binomial discrete distributions as well as Lognormal, Inverse-Gaussian and the five-parameter Bi-Weibull continuous distributions). All these distributions have three or more parameters and the mathematical and computational challenges involved in using these distributions are quite considerable when compared to the four transaction size distributions proposed above. These challenges are bound to outweigh any improvements in modelling accuracy that may be obtained by using them in a real world context.

Table 2. Summary of the properties of the transaction size distributions

Category B (Emergency items) • Orders usually only for one unit but may occasionally be much greater than one. • Demand surges might arise because of emergencies or crises or as a result of extensive preventive maintenance operations carried out by large customers. • SKUs will typically have relatively long lives. • Examples: brake testing kits, fog lamps and lighting beacons following a snow storm.	**Category D (Volatile items)** • Orders will usually be for more than one item and they will vary considerably. • SKUs in this category include items that operate as a set but are sold individually. • These terms will usually have relatively low prices and short lives. Items are used in corrective maintenance (small orders) as well as preventive maintenance (typically, larger orders). As a result, the transaction size variability is high. • Examples: fasteners (bolts, screws, nuts, washers, clips, bearing).
Category A (Job-type items) • Orders usually only for one unit; rarely much greater than 1. • SKUs typically high-priced, highly-specialised items with relatively long lives. • Examples: cylinder heads, crankshafts, clutch kits, drive shafts.	**Category C (Steady items)** • Orders will usually be for more than one item. The transaction sizes do not vary much. • SKUs in this category include items that operate as a set but are sold individually. • Also likely to include standardised items that could be used across a wide range of machinery and makes. Markets for such items are likely to be homogeneous and broad (thus, the relatively low transaction size variability). • These terms will usually be moderately or highly priced with long lives. Items are used in corrective (but not preventive) maintenance (again, resulting in relatively low transaction size variability). • Examples: wiper blades, brake pads, pistons, piston rings and connecting rods for pistons.

The Pascal distribution is used to model the transaction sizes of the SKUs falling in category D because it could (at least in theory) be able to perform as well as any of the other three alternatives. The Pascal distribution provides a good approximation for each of the three other transaction size distributions. If the Pascal distribution is denoted by $Ne(r, p)$, where r is the number of successes and p is the probability of success, then the Geometric, Logarithmic series and Poisson distributions are all limiting forms of the Pascal distribution given the right choice of the parameter r (Katti and Gurland, 1961). However, unlike the other three distributions which have two parameters, the Pascal distribution has three parameters and it is comparatively more demanding in terms of computational effort. The Pascal distribution will therefore only be used when the three other distributions are not appropriate (i.e. in Category D).

3.3 A new integrated classification framework

The scheme in Figure 3 above cannot, as yet, be used in conjunction with the SBC scheme. In the scheme above, demand is assumed to arrive according to a continuous process (specifically, the Poisson process). The SBC scheme, in contrast, was developed based on the assumption that demand arrives according to a Bernoulli process which is a discrete one. The squared coefficients of variation in the two schemes are therefore not equivalent. An analogue of the SBC scheme may however be developed for the case where demand is assumed to arrive according to a Poisson process. When the cut-off boundaries in the SBC scheme ($\overline{p} = 1.32$ and $CV^2(Z) = 0.49$) are expressed in terms of the parameters λ and $CV^2(X)$, we have:

$$\lambda \cong 1.42 \tag{3}$$

$$CV^2(X) \cong \lambda\left[\frac{1}{1-e^{-\lambda}}(1.49)-1\right]-1 \tag{4}$$

The steps involved in deriving these cut-off boundaries are given in detail in *Appendix A*. The continuous analogue of the SBC scheme based on these new cut-off boundaries is illustrated in Figure 3 below.

Figure 3. Continuous analogue of the SBC scheme

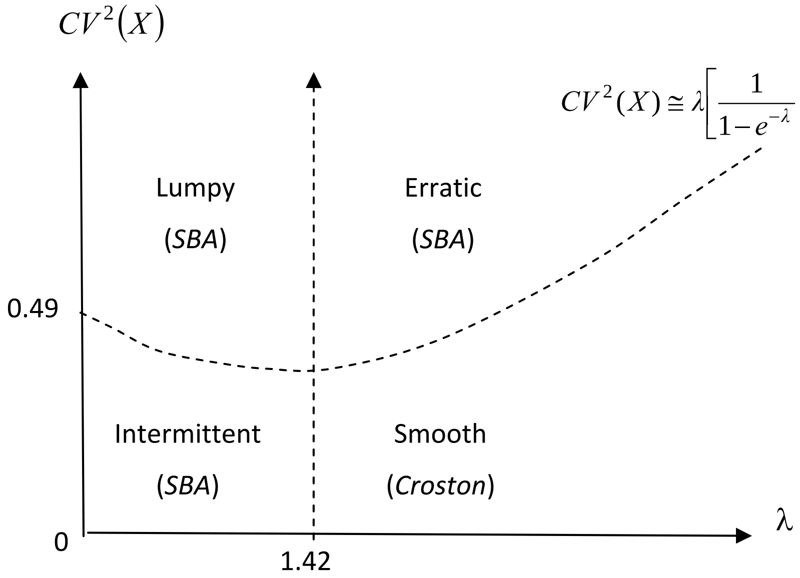

Figures 2 and 3 may be combined, as shown in Figure 4 below, in order to obtain a framework that takes into account both demand forecasting and distributional assumptions. The values of the mode, order arrival rate and the squared coefficient of variation of the transaction sizes may be easily calculated for each SKU in a given dataset. The SKUs may then be plotted against the three-dimensional framework.. The framework should identify the distribution that may be used to model the demand and the preferred estimator to forecast future requirements. For example for an SKU with $\tilde{m}(X)=1$, $\lambda=0.53$ and $CV^2(X)=1.37$, the demand will be assumed to have a negative binomial distribution and the SBA method will be used to forecast the relevant requirements. Thus, based on only three statistics, an inventory manager can identify a preferred distribution and the forecasting method needed to use for each SKU. This is a very important extension of previous work in this area both from an academic and practitioner perspective.

Figure 4. Demand classification framework

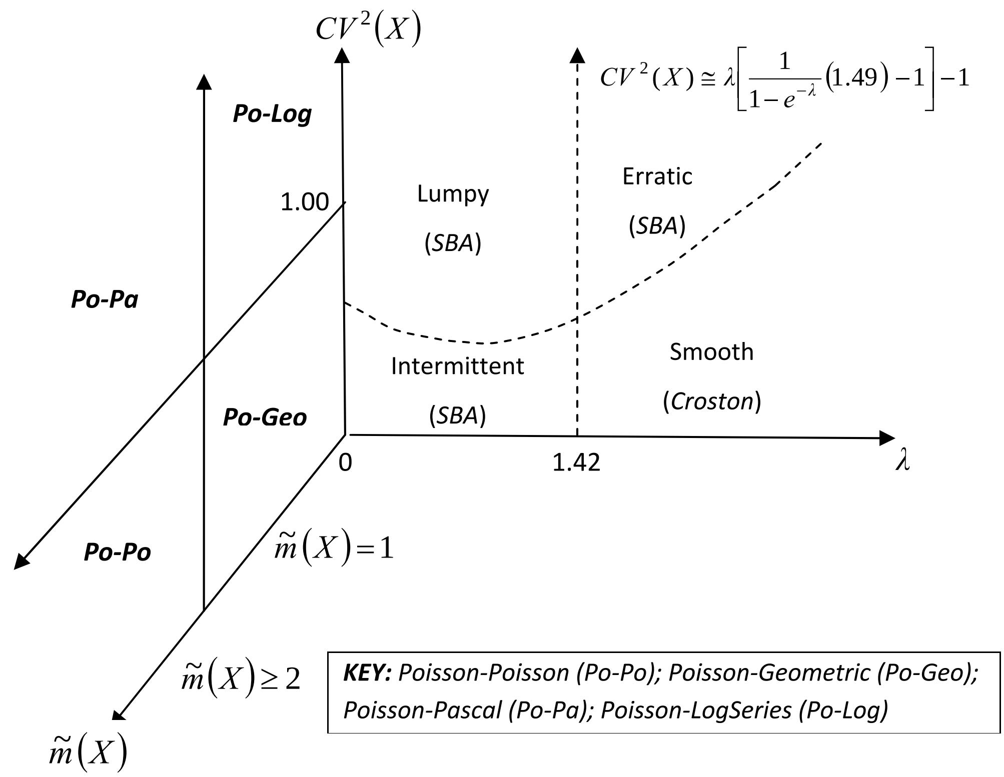

4. Empirical analysis

4.1 Empirical data sets

In this section, we shall carry out an empirical analysis to assess the validity of the theoretical propositions made in section 3. Goodness-of-fit tests will be carried out to assess whether the compound Poisson distributions proposed in section 3 provide a good fit for the demand distributions of empirical SKUs. We will also assess the effectiveness of the proposed demand classification scheme. The empirical data sets are made up of individual demand histories of nearly 30,000 spare part SKUs from five different industries. Table 3 below provides a summary description of the data sets.

Detailed demand information, at transaction order level, was available for datasets 5 and 6. However, in the case of datasets 1, 2, 3 and 4, the demand data was recorded at fixed monthly intervals and information on individual demand orders was not available.

Table 3. Summary description of the empirical data sets

Dataset	Industry	No of SKUs	Time buckets	History length (months)
1	Defence (Air force)	5,000	Monthly	84
2	Defence (Navy)	4,588	Monthly	60
3	Automotive	3,000	Monthly	24
4	Electronics	3,055	Monthly	48
5	Commercial airlines	496	Transaction level	28
6	Domestic Appliances	14,881	Transaction level	60

Additional statistics on the characteristics of the demand series in each of the datasets are presented in Tables 4. Five (rather arbitrarily defined) categories are considered for the demand arrival rate (π) and the squared coefficient of variation of the demand sizes. The demand arrival rate π represents the number of time periods during which there was a demand expressed as a proportion of the total numbers of time periods in the demand history. The percentages indicated represent the proportions of SKUs (in the corresponding dataset) that fall within each category.

Most of the SKUs in datasets 1 (specifically, 99.4% of the SKUs) have very low demand rates and they constitute what is very often termed to as ‘slow intermittent’ SKUs. In contrast, the SKUs in dataset 3 are ‘faster’ intermittent. With respect to demand size variability, the majority of SKUs in data sets 1, 2, 3 and 4 have a $CV^2(Z)$ value equal to or less than 0.5. However, in datasets 5 and 6, most of SKUs have a $CV^2(Z)$ value greater than 1. These statistics show that there is considerable diversity among the empirical data sets with respect to demand intermittence and demand size variability. The compound Poisson distributions will therefore be tested against empirical SKUs with very diverse distributional profiles.

Table 4. Summary statistics of the empirical data sets

π	$0.0 < \pi \leq 0.2$	$0.2 < \pi \leq 0.4$	$0.4 < \pi \leq 0.6$	$0.6 < \pi \leq 0.8$	$0.8 < \pi \leq 1.0$
Dataset 1	*99.40%*	*0.60%*	-	-	-
Dataset 2	*29.38%*	*33.57%*	*20.64%*	*11.12%*	*5.30%*
Dataset 3	-	-	*10.40%*	*45.47%*	*44.13%*
Dataset 4	*23.47%*	*29.42%*	*17.70%*	*13.56%*	*15.86%*
Dataset 5	*7.86%*	*30.04%*	*29.23*	*17.34*	*15.52*
Dataset 6	*61.09%*	*16.16%*	*8.43%*	*5.54%*	*8.78%*
$CV^2(Z)$	**$0.0 < CV^2 \leq 0.5$**	**$0.5 < CV^2 \leq 1.0$**	**$1.0 < CV^2 \leq 5.0$**	**$5.0 < CV^2 \leq 10.0$**	**$10.0 < CV^2$**
Dataset 1	*52.96%*	*25.46%*	*20.82%*	*0.74%*	*0.02%*
Dataset 2	*54.10%*	*25.48%*	*19.83%*	*0.57%*	*0.02%*
Dataset 3	*75.90%*	*19.63%*	*4.40%*	*0.03%*	*0.03%*
Dataset 4	*56.27%*	*24.75%*	*18.55%*	*0.43%*	-
Dataset 5	*2.02%*	*11.29%*	*65.12%*	*17.14%*	*4.44%*
Dataset 6	*2.51%*	*3.87%*	*25.40%*	*20.96%*	*47.26%*

4.2 Goodness-of-fit tests in the context of intermittent demand

The goodness of fit test used in this paper is the Kolmogorov-Smirnov test (or K-S test, in short). Other goodness of fit tests were also considered. Pearson's χ^2 test is a well-known goodness of fit test that places observations in categories and compares the observed and expected frequencies in the each of the categories. This test is easy to use but is associated with some requirements/'rules' (given in Cochran, 1952; Birnbaum, 1962; Roscoe and Byars, 1971; Kendall *et al.*, 1987; Cramer, 1999) that specify the minimum and average expected frequencies for the categories. The data in our study mostly failed to meet these requirements. The intermittent nature of our data meant that the demand was zero in most periods and there were very few demand observations. In most of the cases, we could not create more than two viable categories and, as a result, we could not carry out a valid χ^2 test.

The Cramer von Mises and Anderson-Darling goodness of fit tests could potentially have also been used. However, whereas the K-S is distribution-free (i.e. the critical values are independent of the hypothesised distribution), the critical values of the Cramer von Mises and

Anderson-Darling tests will depend on the hypothesised distribution. As a result, different tables of the critical values must be calculated for each of the proposed distributions. The computational effort involved in deriving the critical values (by using, for example, Monte Carlo methods) would be very time-consuming.

As it was discussed in sub-section 3.2 there have been only a few empirical studies on goodness of fit in the area of intermittent demand management. As far as we are aware, the only studies in this area are by Kwan (1991), Eaves (2002) and Syntetos *et al.* (2011). The negative binomial distribution performed well in Kwan's study but the sample used in that study was rather small (only 86 SKUs). Eaves (op. cit.) carried out goodness of fit tests on a larger set of spare/service parts SKUs (6,795 units). However, the goodness-of-fit tests were carried out separately for the inter-demand intervals and demand sizes. It is important to note that one cannot assume that a compound distribution will provide high levels of fit simply because the constituent distributions provide high levels of frequency of fit for the corresponding transaction sizes and inter-demand intervals. The results presented by Eaves (2002) therefore do not contribute much to our study. In this paper, we will take a direct approach. Goodness-of-fit tests will be carried out in order to assess the compound Poisson distributions and not just their constituent parts (transaction sizes and inter-demand intervals).

Syntetos *et al.* (2011) have also assessed the goodness of fit of a number of distributions, including two of the distributions examined in this study (the Poisson-Geometric and Poisson-Logarithmic series distributions). The goodness-of-fit test used in that study was the K-S test and the empirical database was made up of three of the datasets (Datasets 1, 2 and 4) used in our study. They found that the Poisson-Geometric distribution outperformed the Poisson-Logarithmic series distributions. However, in the process of carrying out the tests, Syntetos *et al.* derived the number of categories based on the hypothesised distribution and

not the empirical one. As a result, the number of categories was too high and consequently the test was excessively 'liberal' (please refer also to the next sub-section). A communication is currently under preparation to appear in the *International Journal of Production Research* where the above study was published for the purpose of pointing out the limitation under concern

4.3 Goodness-of-fit: analysis and results

The goodness-of-fit test used in this study is the K-S test with the significance level set at 5%. The distribution of the demand per period has been considered rather than the distribution of the lead-time demand; this is due to the lack of information on the actual lead times for the SKUs in Datasets 2, 5 and 6. It could also be argued that this is the correct approach to take if our goal is to assess whether the hypothesised distribution accurately models demand. The lead-time demand data will depend not just on the demand process (which is what we are interested in) but also on a number of other factors related to the firm in question (such as the geographical positioning of the suppliers, the periodic or continuous nature of the replenishment policies being used etc). For example, two firms might face the same demand series but due to the above discussed reasons their lead-time data will not be the same. Demand-per-period data therefore provide a more accurate and objective representation of the demand process in the context of our study.

In the K-S tests, the empirical distribution function (EDF) for each SKU was taken as the cumulative frequency distribution of the demand for the SKU under concern and the fitted distribution was the cumulative distribution function (CDF) of the hypothesised compound Poisson distribution. The parameters of the hypothesised distribution were estimated from the observed demand data using the following two methods:

a) The method of moments, using the first two moments (or MM, in short);

b) The method of mean and zero frequency (or M&Z, in short). With this method, the estimates are derived by equating: (i) the sample mean and the population mean; (ii) the observed and expected probabilities of zero observations. This method has been used in a number of studies including Katti and Gurland (1962), Bowman and Shenton (1967) and Shenton and Bowman (1977).

In the case of the Poisson-Pascal distribution, an additional moment is required under either method in order to obtain the third parameter. The formulae for the parameter estimators under each of these methods are given in the *Appendix B*.

In this paper, a parameter estimator is referred to as *domain compliant* if the values of the estimator will always fall within the domain of the relevant parameter. Taking the well-known normal distribution $N\left(\mu,\sigma^2\right)$ with $-\infty < \mu < \infty$ and $0 \le \sigma^2$ as an example, the sample mean is a domain compliant estimator of the parameter μ since the average will always fall within the domain of the parameter. The sample variance is also a domain compliant estimator of the parameter σ^2. For the purposes of this study, the value of the parameter estimate has to fall within the domain of the relevant parameter otherwise it is not possible to obtain a valid fitted CDF.

In the case of the compound Poisson distributions, the parameter estimation methods given above are not necessarily associated with domain compliant estimators. The MM estimators will provide meaningful estimates so long as the empirical data satisfies the theoretical relationship between the relevant moments. For the four compound Poisson distributions discussed in this paper, the variance is always greater than equal to the mean (Johnson *et al.*, 2005). The MM estimates will therefore fall outside the domain if the sample variance

(denoted as s_y^2) is less than the sample mean (denoted as $\bar{y}$). This can be easily seen in *Appendix B*. Most of the MM estimators fall outside the relevant domain whenever $s_y^2 < \bar{y}$. The method of mean and zero frequency also fails in the few cases that there are no periods with zero demand (i.e. when the observed zero frequency, f_0, is equal to 0).

Parameter estimators that are not domain compliant present a practical challenge in goodness-of-fit tests; if any of the derived estimates falls outside the relevant domain, what conclusion do we draw with regards to the goodness-of-fit? For example, if the observed sample variance is less than the sample mean, then this might be a genuine reflection of the fact that the underlying demand distribution is under-dispersed. Alternatively, the underlying demand distribution might actually be over-dispersed, but the observed sample variance might be less than the sample mean simply as a result of sampling error. Without knowing the underlying distribution, it is not possible to know what the right conclusion should be.

Domain incompliance is a challenge not only in K-S goodness-of-fit tests, but also for every procedure that relies on parameter estimation, including parametric inventory management. Stock control parameters such as the reorder point and the order-up-to level are derived based on the distribution of demand during lead-time. If the parameter estimators that are used are domain incompliant, then the parameters estimates obtained will sometimes fall outside the relevant domain; such estimates are meaningless and they would not provide us with a valid distribution for the lead time demand.

The problem of domain incompliance is exacerbated in the case of intermittent demand. Intermittent demand is often characterised by only a small number of demand observations; as the number of demand observations decreases, the variance of the parameter estimates

(and thus, the probability that the estimate will fall outside the domain) will usually increase. It is worth pointing out that domain incompliance will usually not be a problem for fast-moving items; the demand for such items is typically assumed to be normally distributed and as discussed above the parameter estimators for this distribution are domain compliant.

In this study, we have made some restrictions in order to ensure that the parameter estimates do not fall outside the domain. In the case of MM estimates,

a) whenever the sample variance (s_y^2) is less than or equal to the sample mean ($\bar{y}$), the sample variance is increased and made equal to $1.05\bar{y}$. This is similar to the approach adopted by Kwan (2002).

b) Furthermore, in the case of the Poisson-Pascal distribution, the sample variance has to fall within the following interval:

$$\frac{\bar{y}+\sqrt{\bar{y}^2+8m_3\bar{y}}}{4} < \sigma_y^2 < \frac{-\bar{y}+\sqrt{\bar{y}^2+4\bar{y}(m_3-\bar{y})}}{2}$$

(where $\bar{y}$, s_y^2 and m_3 are the sample mean, sample variance and the sample central third moment respectively). The derivation of this restriction is given in *Appendix B*.

As for the M&Z estimates,

a) the observed zero frequency, f_0, was bound within the range $1 \le f_0 \le N-1$, where N is the length of the demand series. Note that $f_0 = N$ represents the trivial case of a demand series that does not have any periods with positive demand; $f_0 = 0$ leads to computation problems (specifically, taking logarithms of zero).

b) Furthermore, in the case of the Poisson-Pascal distribution, whenever the estimated parameter p is less than 0, we have assumed that $s_y^2 = 1.05 \times -\frac{\bar{y}^2}{ln\left(f_0/N\right)}$. The derivation of this restriction is also given in *Appendix B*.

The additional restriction arises in the case of the Poisson-Pascal distribution (for both estimators) necessarily as a consequence of the fact that the distribution has one more parameter. With these restrictions in place, the parameter estimates will always fall within the relevant domain and the probability distributions obtained from these estimates will be valid. The fitted CDF obtained under these restrictions is simply our best effort on obtaining a valid hypothesised distribution that provides a close fit to the empirical data. The CDF fitted in this manner will still fail to provide significant fit if there is little agreement between the empirical data and the hypothesised compound Poisson distribution.

The empirical data used in our study is highly varied as indicated in Table 3 and it is not clear which of the two parameter estimation methods would perform best for our data. Goodness-of-fit tests have therefore been carried out using both methods. The results presented in Table 5 indicate, per dataset and under each parameter estimation method, the percentage of SKUs for which a distribution was found to provide a significant fit. The levels of frequency of fit achieved by the four compound Poisson distributions are quite high (for all four distributions and under each parameter estimation method, the level of frequency of fit was at least 70%). The proposed compound Poisson distributions therefore provided significant fit for most of the empirical demand data used in our study. With respect to the parameter estimation methods, a comparison of the levels of frequency of fit achieved reveals that, overall, there is little to choose between MM and M&Z estimators. MM estimators performed better than M&Z estimators in some cases but they performed worse in others. However, in the case of

the Poisson-Pascal distribution, the MM estimators consistently attained higher levels of frequency fit when compared to the M&Z estimators; the differences in the levels of frequency of fit achieved varied from as little as 0.54% (Dataset 1) to as much as 16.33% (Dataset 5). The choice of parameter estimation method may therefore make a difference.

Table 5. Goodness-of-fit results for the compound Poisson distributions - % fit

	Po-Geo		**Po-Log**		**Po-Po**		**Po-Pa**	
	MM	M&Z	MM	M&Z	MM	M&Z	MM	M&Z
Dataset 1	*99.48%*	*99.70%*	*94.80%*	*99.52%*	*98.34%*	*99.86%*	*99.28%*	*98.74%*
Dataset 2	*92.00%*	*95.44%*	*89.12%*	*94.12%*	*82.41%*	*88.03%*	*90.78%*	*79.60%*
Dataset 3	*98.20%*	*98.17%*	*97.80%*	*95.40%*	*95.37%*	*92.07%*	*97.50%*	*81.23%*
Dataset 4	*89.33%*	*91.52%*	*92.80%*	*91.26%*	*76.53%*	*85.04%*	*87.53%*	*72.70%*
Dataset 5	*84.07%*	*86.69%*	*82.66%*	*81.05%*	*72.18%*	*77.42%*	*85.08%*	*68.75%*
Dataset 6	*93.85%*	*94.47%*	*92.88%*	*93.77%*	*89.12%*	*91.02%*	*93.37%*	*87.11%*

MM – Method of moments

M&Z – Method of mean and zero frequency

4.4 Validity of the proposed demand classification framework

In this subsection, we will assess the empirical validity of the proposed demand classification framework. In order to carry out this assessment, demand data that includes information on the individual transaction orders is required. However, demand data at transaction order level was available only for datasets 5 and 6. Consequently, the empirical validity of the framework will be assessed using the nearly 15,000 SKUs coming from these two datasets.

Our assessment of the empirical validity will be restricted to the distributional of the framework. The SKUs in our empirical data set will be categorised using the proposed framework and the levels of fit achieved by the different distributions in each category will be compared. The empirical validity of the framework with respect to forecasting accuracy is not assessed in this paper. In addition to the original study by Syntetos *et al.* (2005), another recent investigation (Heinecke *et al.*, 2011) shows the good empirical performance of the

SBC scheme. The cut-off points in the proposed framework were derived directly from the SBC scheme under the assumption that demand arrival follows a Poisson process rather than a Bernoulli one (which was the case in the SBC scheme). One would expect that there would be little difference in terms of performance in the two cases. The accuracy of the cut-off boundaries in the framework is also to be tested as part of our future research through empirical and simulation-based analyses.

The SKUs in datasets 5 and 6 were first categorised according to the mode and variability of the transaction sizes as proposed by the relevant framework. Goodness of fit tests were then carried out for each of the four compound Poisson distributions and the K-S statistics were calculated accordingly. Finally, sign tests were carried out, for each pair of distributions and in each category, to test the hypothesis that there is no difference in the goodness of fit achieved by the two distributions. For each pair of distributions, the difference between the K-S statistics achieved by the two distributions was calculated and the proportion of negative differences (denoted as $\hat{d}$) were calculated. If the null hypothesis that there is no difference is correct, then we would expect that $\hat{d}$ would not be statistically different from 0.50.

The results of the sign tests are given in Table 6. The difference between the K-S statistics achieved by each pair of distributions was calculated as follows:

K-S statistic achieved by the distribution given in the corresponding row minus (-) the K-S statistic achieved by the distribution given in the corresponding column.

N represents the number of SKUs that fell into the given category and s is the sample standard deviation of $\hat{d}$. Confidence intervals were constructed for $\hat{d}$ under the assumption that $\hat{d}$ is normally distributed. Such an assumption is justified given the high number of SKUs in each category (Berry, 1941; Esseen, 1956). The pair-wise comparisons between the

distributions are performed simultaneously; a multiple-comparison correction is therefore required in order ensure that the overall confidence level (in this case, 95%) is maintained. If a multiple-comparison correction was not used, then the Type I error (i.e. the probability of incorrectly rejecting the null hypothesis) could be significantly higher than 0.05 (Williams and Abdi, 2010). The multiple-comparison correction used in this study was the Bonferroni correction (Benjamini and Hochberg, 1995).

To allow for easier interpretation, we have expressed the results in this paper in terms of the variable $\hat{\delta} = \hat{d}$ - 0.50 . Under this translation, the confidence intervals under the null hypothesis are centred at the value $\hat{\delta} = 0$ which is more convenient than the original value of $\hat{d} = 0.50$. The confidence intervals in table 6 can be interpreted as follows:

a) If the lower limit of the confidence interval is greater than 0, then we conclude that the distribution given in the row label outperformed the distribution given in the column label in the given category.
b) Alternatively, if the upper limit of the confidence interval is less than 0, we conclude that the distribution given in the column label outperformed the distribution given in the row label in the given category.
c) A conclusion that there is no difference between the two distributions is obtained if 0 falls within the confidence interval.

Table 7 presents the conclusions that are drawn from the observed confidence intervals. The results that are relevant when assessing the validity of the proposed demand classification scheme are highlighted in bold. The results presented in the two tables largely agree with the suggestions in the proposed framework. The Poisson-Logarithmic series and the Poisson-Poisson distributions outperformed all the alternatives in categories B and C respectively. The Poisson-Geometric distribution also performed at least as good as each of the other three

distributions in category A. However, contrary to what the proposed framework suggests, the Poisson-Pascal distribution was not the best performing one in category D.

Table 6. Comparison of goodness-of-fit in the four categories (confidence intervals)

Category B (N = 2,834; *s* = 0.01018)

	Po-Log	Po-Po	Po-Pa
PoGeo	**-0.12, -0.07**	0.41, 0.46	-0.05, -0.01
PoLog		**0.29, 0.34**	**0.19, 0.23**
PoPo			-0.50, -0.45

Category D (N = 856; *s* = 0.01612)

	Po-Log	Po-Po	Po-Po
PoGeo	0.06, 0.11	0.44, 0.48	**0.04, 0.08**
PoLog		0.19, 0.24	**0.01, 0.05**
PoPo			**-0.51, -0.47**

$CV^2(X) = 1$

Category A (N = 6,552; *s* = 0.00692)

	Po-Log	Po-Po	Po-Pa
PoGeo	**0.13, 0.17**	**-0.01, 0.03**	**-0.04, 0.01**
PoLog		-0.08, -0.04	-0.09, -0.04
PoPo			-0.05, -0.01

Category C (N = 5,128; *s* = 0.00607)

	Po-Log	Po-Po	Po-Pa
PoGeo	0.46, 0.51	**-0.34, -0.29**	-0.41, -0.36
PoLog		**-0.47, -0.42**	-0.50, -0.45
PoPo			**0.11, 0.15**

$\tilde{m}(X) = 1$ $\tilde{m}(X) \geq 2$

Table 7. Comparison of goodness-of-fit in the four categories (conclusions)

Category B

PoGeo<PoLog	PoGeo>PoPo	PoGeo<PoPa
	PoLog>PoPo	**PoLog>PoPa**
		PoPo<PoPa

Category D

PoGeo>PoLog	PoGeo>PoPo	**PoGeo>PoPa**
	PoLog>PoPo	**PoLog>PoPa**
		PoPo<PoPa

$CV^2(X) = 1$

Category A

PoGeo>PoLog	**No difference**	**No difference**
	PoLog<PoPo	PoLog<PoPa
		PoPo<PoPa

Category C

PoGeo>PoLog	**PoGeo<PoPo**	PoGeo<PoPa
	PoLog<PoPo	PoLog<PoPa
		PoPo>PoPa

$\tilde{m}(X) = 1$ $\tilde{m}(X) \geq 2$

The poor performance of the Poisson-Pascal distribution could possibly be best explained by the greater number of restrictions placed on the parameters estimates. In theory, the Poisson-Pascal distribution should provide greater flexibility when compared to the other three

distributions. However, this would only be the case in practice if the parameter estimators of the distribution were domain compliant. When parameter estimators are not domain compliant, restrictions will have to be placed in order to ensure that the parameter estimates fall within the domain. These restrictions affect not only the domain incompliant estimator in question but also all estimators that are correlated with the restricted estimator. The restrictions on the parameter estimates effectively reduce the ability of the hypothesised distribution to give a good fit. Instead of picking from the set of hypothesised distributions with every possible parameter estimate combination, we are only selecting from the set of hypothesised distributions with parameter estimates combinations that satisfy the restrictions.

Despite the poor relative performance of the Poisson-Pascal distribution in category D, there are a number of points worth considering before we can draw any final conclusions on the validity of the framework. First, even though the Poisson-Pascal distribution was not the best performing distribution in category D, the distribution achieved a level of frequency of fit of 73.3% which, based on the results given in table 5, is not particularly low. Secondly, there were considerably fewer SKUs (only 856) in category D when compared to the other three categories. Category D however covers a much larger space than the other three categories. Category A is bounded in terms of both $\tilde{m}(X)$ and $CV^2(X)$ and categories B and C are bounded in terms of one of the two variables. Category D is unbounded in both $\tilde{m}(X)$ and $CV^2(X)$ and, given the larger space that category D covers, it would have been helpful if there were more SKUs in that category.

5. Implications for the Operations Management theory and practice

As others (e.g. Fortuin and Martin, 1999; Botter and Fortuin, 2000; Syntetos *et al.*, 2009) have already pointed out, the management of spare parts and other inventory items with intermittent demand is a difficult task. A number of authors have argued that compound

distributions could be used to model such intermittent demand patterns. However, there have been very few empirical studies in this area. The main contribution of this paper relates to a detailed empirical investigation on the viability of using compound Poisson distributions to model intermittent demand. Goodness of fit tests were carried out for various compound Poisson distributions and the challenges involved in using such distributions were explored.

Compound Poisson processes have a structure that is similar to the demand-generating process associated with intermittent demand – events (in this case demand orders) arrive sporadically and the size of the events is variable. The likeness between compound Poisson processes and the demand arrival processes typically observed among spare parts will have an intuitive appeal to inventory managers. Goodness of fit tests were carried out in this study for four different compound Poisson distributions: i) Poisson-Pascal; ii) Poisson-Poisson; iii) Poisson-Log Series and iv) Poisson-Geometric. The empirical demand data used in these tests was extensive and consisted of the demand histories of more than 30,000 spare parts SKUs. All four distributions were found to provide high levels of frequency of fit. In general, distributions that provide high levels of fit provide high levels of stock control performance. These results suggest that managers should consider using compound Poisson distributions to model demand of spare parts and other items with similar intermittent demand patterns.

Compound Poisson distributions also model the transaction sizes independently of the order arrival process. Orders are assumed to arrive according to a Poisson process but different distributions could be used to model the transaction order sizes. Different compound Poisson distributions could therefore be used to model SKUs with differing transaction size profiles. In the area of inventory management, there is wide agreement that effective classification can lead to substantial improvements in performance. In this paper, we proposed a framework that assigns different compound Poisson distributions to SKUs with differing transaction size

properties. The framework classifies SKUs based on the modality and variability of the observed transactions sizes and it can greatly facilitate the process of selecting distributional models for items with intermittent demand. The framework has been assessed for its empirical validity in terms of the goodness of fit. The results suggest that the framework is most effective in assigning the best-fitting distribution to SKUs falling in the different categories. The scheme does also offer a qualitative characterisation of the SKUs that may fall within each of the categories. The linkage between the technical attributes of the classification scheme and the qualitative attributes of the SKUs assigned to the various categories should be of great value to practitioners operating in this area. However, it is true to say that further tests are required in order to assess the effectiveness of the scheme in terms of its stock control performance.

One of issues considered in this study is the need for a hierarchical list of criteria that should be used when selecting distributions for modelling demand. The most important criterion is that the hypothesised distribution has to match the underlying structure of demand as understood by the inventory managers. But based on the challenges encountered in this study, it seems that the next most important criterion should be the mathematical tractability of the distribution. If the distribution is to be useful in practical settings, then it needs to have a probability function that is easy to compute using readily available software packages such as Excel. In the context of intermittent demand, distributions with large number of parameters should be avoided as much as possible. For a given demand pattern, as the number of parameters increases, the degrees of freedom (the number of independent observations in a sample that are available to estimate parameters) falls. The accuracy of the parameter estimates will therefore deteriorate as the number of parameters increases. This is particularly a problem in the case of intermittent demand. In general, the accuracy of the parameter estimates will improve as the samples becomes more diverse. When demand is intermittent,

there is little diversity in the observations (most of the observations are zeros). Finally, mathematical tractability in terms of the domain compliance of the parameter estimators is also an issue worth considering. Some distributions might seem appealing in theory but, if they have domain incompliant estimators, they might not perform as well (for example, the Poisson-Pascal distribution in this study).

After mathematical tractability, the next most important criterion is corroborative empirical evidence. However, relevant empirical evidence might be hard to come by and occasionally the findings in different studies might contradict one another. The final criterion should be the flexibility of the distribution. While flexibility might be desirable, this is an issue that can be easily resolved by simply increasing the number of distributions in order to ensure that there is a distribution to accommodate each of the possible demand profiles. While this might seem inconvenient, the challenges encountered in this study suggest that it might be worthwhile to sacrifice flexibility for mathematical tractability.

6. Conclusions and further work

Demand classification is an important operational issue in the management of spare part inventory items. Demand classification facilitates decision-making with respect to forecasting and stock control and enables managers to focus their attention on the SKUs considered most important. In this paper, we carried out goodness-of-fit tests to assess whether compound Poisson distributions provide a good fit to SKUs with intermittent demand. An empirical data set of nearly 30,000 spare part SKUs from five different industries was used in these tests. The compound Poisson distributions were found to provide good fit for most of the SKUs in the empirical data set. We have also proposed a demand classification framework that categorises SKUs based on the order arrival rate and the mode and variability of the observed transaction sizes. The framework facilitates the process of selecting forecasting methods and

distributional models for items with intermittent demand. The framework was also tested for its empirical validity but only with respect to the distributional properties of demand. The results suggest that the scheme is very effective in the sense that the proposed compound distributions provide high levels of frequency of fit for SKUs that fall within the assigned categories. A linkage was also provided between the scheme and the potential qualitative attributes of SKUs falling within each category and we have also built on previous work to propose a comprehensive list of criteria to be used when selecting demand distributions. Finally, an extensive discussion has been provided on parameter estimation related difficulties in this area. As such, we feel that our work should enable further theoretical developments in the area of spare parts management and should successfully inform relevant real world practices.

In the next steps of our research we plan to assess the empirical validity of the classification framework in terms of its implications for forecast accuracy. The framework will also need to be assessed for effectiveness in terms of stock control performance. It is the latter issue to receive priority since empirical studies have already been performed in terms of forecasting, albeit with an emphasis on a Bernoulli rather than a Poisson demand arrival process. Further work and empirical studies on the performance of non-parametric approaches (like Bootstrapping for example) and the way such approaches compare to the more 'traditional' distribution-based inventory control considered in this paper should also contribute significantly towards extending the current state of knowledge in this area.

References

Adelson, R. M. (1966), "Compound poisson distributions," *Operational Research,* Vol. 17 No. 1, pp. 73-75.

Adenso-Díaz, B. (1996), "How many units will be short when stockout occurs?," *International Journal of Operations & Production Management,* Vol. 16 No. 4, pp. 112-118.

Archibald, B. C. and Silver, E. A. (1978), "(s, S) policies under continuous review and discrete compound poisson demand," *Management Science,* Vol. 24 No. 9, pp. 899-909.

Benjamini, Y. and Hochberg, Y. (1995), "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal Statistical Society. Series B (Methodological),* Vol. 57 No. 1, pp. 289-300.

Berry, A.C. (1941), "The Accuracy of the Gaussian Approximation to the Sum of Independent Variates," *Transactions of the American Mathematical Society* , Vol. 49 No. 1, pp. 122–136.

Birnbaum, Z. W. (1962), *Introduction to probability and mathematical statistics,* 1st ed., Harper, New York, NY.

Blumenfeld, D. E., Daganzo, C. F., Frick, M. C., and Gonsalvez, D. J. A. (1999), "Impact of manufacturing response time on retailer inventory," *International Journal of Operations & Production Management,* Vol. 19 No. 8, pp. 797-811.

Botter, R. and Fortuin, L. (2000), "Stocking strategy for service parts–a case study," *International Journal of Operations & Production Management,* Vol. 20 No. 6, pp. 656-674.

Bowman, K. O. and Shenton, L. (1967), *Remarks on estimation problems for the parameters of the Neyman Type A distribution*, Union Carbide Corporation, Report ORNL-4102, Oak Ridge, TN.

Boylan, J. E. (1997), "*The centralisation of inventory and the modelling of demand",* unpublished Ph.D. Thesis, University of Warwick, Coventry.

Boylan, J. E. and Syntetos, A. A. (2008) "*Forecasting for inventory management of service parts*", in K. A. H. Kobbacy and D. N. P. Murthy (Ed.), *Complex System Maintenance Handbook*, Springer Verlag, NY, pp. 479-506.

Cochran, W. G. (1952), "The Chi-Squared test of goodness of fit," *The Annals of Mathematical Statistics,* Vol. 23 No. 3, pp. 315-345.

Cramer, H. (1999), *Mathematical methods of statistics,* 19th ed., Princeton University Press, Princeton, NJ.

Croston, J. D. (1972), "Forecasting and stock control for intermittent demands," *Operational Research,* Vol. 23 No. 3, pp. 289-303.

Eaves, A. H. C. (2002), "*Forecasting for the ordering and stock-holding of consumable spare parts",* PhD thesis, Lancaster University, Lancaster.

Esseen, C. (1956), "A moment inequality with an application to the central limit theorem," *Skandina isk Aktuarietidskr*, Vol 39, pp. 160–170.

Feeney, G. J. and Sherbrooke, C. C. (1966), "The (S-1, S) inventory policy under compound poisson demand," *Management Science,* Vol. 12 No. 5, pp. 391-411.

Fildes, R. and Beard, C. (1992), "Forecasting systems for production and inventory control," *International Journal of Operations & Production Management,* Vol. 12 No. 5, pp. 4-27.

Fortuin, L. and Martin, H. (1999), "Control of service parts," *International Journal of Operations & Production Management,* Vol. 19 No. 9, pp. 950-971.

Friend, J. K. (1960), "Stock control with random opportunities for replenishment," *OR,* Vol. 11 No. 3, pp. 130-136.

Gutierrez, R. S., Solis, A. O., and Mukhopadhyay, S. (2008), "Lumpy demand forecasting using neural networks," *International Journal of Production Economics,* Vol. 111 No. 2, pp. 409-420.

Jasper, J. B. (2006), *Quick Response Solutions, FedEx Critical Inventory Logistics Revitalized*, FedEx, Memphis, TN.

Johnson, N. L., Kemp, A. W., and Kotz, S. (2005), *Univariate discrete distributions,* 3rd ed., Wiley-Interscience, Hoboken, NJ.

Johnston, F. R. and Boylan, J. E. (1996), "Forecasting for items with intermittent demand," *Journal of the Operational Research Society,* Vol. 47 No. 1, pp. 113-121.

Kalchschmidt, M., Verganti, R., and Zotteri, G. (2006), "Forecasting demand from heterogeneous customers," *International Journal of Operations & Production Management,* Vol. 26 No. 6, pp. 619-638.

Katti, S. and Gurland, J. (1961), "The Poisson Pascal distribution," *Biometrics,* Vol. 17 No. 4, pp. 527-538.

Katti, S. and Gurland, J. (1962), "Efficiency of certain methods of estimation for the negative binomial and the Neyman type A distributions," *Biometrika,* Vol. 49 No. 1/2, pp. 215-226.

Kemp, C. D. (1967), "'Stuttering - Poisson' distributions," *Journal of the Statistical and Social Inquiry Society of Ireland,* Vol. 21 No. 5, pp. 151-157.

Kendall, M. G., Stuart, A., Ord, J. K., and O'Hagan, A. (1987), *Kendall's advanced theory of statistics,* vol. 1, Oxford University Press, New York, NY.

Kwan, H. W. (1991), "*On the demand distributions of slow moving items",* unpublished Ph.D. thesis, Lancaster University, Lancaster.

Mak, K. and Hung, C. (1993), "A Simple Model for Computing (s, S) Inventory Policies when Demand is Lumpy," *International Journal of Operations & Production Management,* Vol. 6 No. 2, pp. 62-68.

Matheus, P. and Gelders, L. (2000), "The (R, Q) inventory policy subject to a compound Poisson demand pattern," *International Journal of Production Economics,* Vol. 68 No. 3, pp. 307-317.

Mitchell, C. R., Rappold, R. A., and Faulkner, W. B. (1983), "An analysis of Air Force EOQ data with an application to reorder point calculation," *Management Science,* Vol. 29 No. 4, pp. 440-446.

Naddor, E. (1978), "Sensitivity to distributions in inventory systems," *Management Science,* Vol. 24 No. 16, pp. 1769-1772.

Özkaya, B. Y., Gürler, Ü., and Berk, E. (2006), "The stochastic joint replenishment problem: A new policy, analysis, and insights," *Naval Research Logistics (NRL),* Vol. 53 No. 6, pp. 525-546.

Partovi, F. Y. and Burton, J. (1993), "Using the analytic hierarchy process for ABC analysis," *International Journal of Operations & Production Management,* Vol. 13 No. 9, pp. 29-44.

Roscoe, J. T. and Byars, J. A. (1971), "An investigation of the restraints with respect to sample size commonly imposed on the use of the chi-square statistic," *Journal of the American Statistical Association,* Vol. 66 No. 336, pp. 755-759.

Sato, K. I. (1999), *Lévy processes and infinitely divisible distributions,* Cambridge University Press, Cambridge.

Satterthwaite, F. (1942), "Generalized Poisson distribution," *The Annals of Mathematical Statistics,* No. pp. 410-417.

Shenton, L. R. and Bowman, K. O. (1977), *Maximum likelihood estimation in small samples,* C. Grifiin & Co., London.

Syntetos, A. A., Babai, M. Z., and Altay, N. (2011), "On the demand distribution of spare parts," advance online publication, 10.1080/00207543.2011.562561.

Syntetos, A. A. and Boylan, J. E. (2005), "The accuracy of intermittent demand estimates," *International Journal of Forecasting,* Vol. 21 No. 2, pp. 303-314.

Syntetos, A. A., Boylan, J. E., and Croston, J. D. (2005), "On the categorization of demand patterns," *Journal of the Operational Research Society,* Vol. 56 No. 5, pp. 495-503.

Syntetos, A. A., Keyes, M., and Babai, M. (2009), "Demand categorisation in a European spare parts logistics network," *International Journal of Operations & Production Management,* Vol. 29 No. 3, pp. 292-316.

Teunter, R. H., Babai, M. Z., and Syntetos, A. A. (2010), "ABC classification: service levels and inventory costs," *Production and Operations Management,* Vol. 19 No. 3, pp. 343-352.

U.S. General Accounting Office (G.A.O.) (1997), *Defense inventory: Spare and repair parts inventory costs can be reduced*, (*Letter report, 17 January, GAO/NSIAD-97–47*), G.A.O.,Washington, D.C.

Ward, J. B. (1978), "Determining reorder points when demand is lumpy," *Management Science,* Vol. 24 No. 6, pp. 623-632.

Watson, R. B. (1987), "The effects of demand-forecast fluctuations on customer service and inventory cost when demand is lumpy," *Journal of the Operational Research Society,* Vol. 38 No. 1, pp. 75-82.

Willemain, T. R., Smart, C. N., and Schwarz, H. F. (2004), "A new approach to forecasting intermittent demand for service parts inventories," *International Journal of Forecasting,* Vol. 20 No. 3, pp. 375-387.

Williams, T. M. (1984a), "Stock control with sporadic and slow-moving demand," *Journal of the Operational Research Society,* Vol. 35 No. 10, pp. 939-948.

Williams, L. J. and Abdi, H. (2010) "*Post-hoc comparisons*", in N. J. Salkind*, et al.* (Ed.), *Encyclopedia of Research Design*, Sage Publications, Inc, Thousand Oaks, CA, pp. 1060 - 1067.

Appendix A: Classification cut-off values under a Poisson arrival process

If the demand data is available at transaction order level, the demand can be modelled using a continuous-time (e.g. Poisson) or discrete-time (e.g. Bernoulli) process. Furthermore, continuous-time demand processes can be easily converted into discrete-time processes – time can be divided into fixed discrete intervals (i.e. unit periods such as weeks or months) and the demand orders arriving in each such period can be 'bucketed' (aggregated). However, when the orders are bucketed, information on the sizes of the individual orders as well as their arrival times is lost. Once the orders have been bucketed, it is no longer possible to obtain statistics such as the order arrival rate and the squared coefficient of variation of the transaction orders. It is however still possible to obtain the averages of such statistics over a single period.

Syntetos *et al.* (2005) identified cut-off boundaries that delineated regions over which the Syntetos-Boylan Approximation (SBA) method outperforms Croston's forecasting method. When Syntetos *et al.* developed these cut-off boundaries, they modelled demand as a discrete (Bernoulli) process. The challenge in this paper is to convert the demand arrival and variability rates from that study to the equivalent average arrival and variability rates under a continuous arrival process. In this paper, we have used the following notation for the transaction order sizes and the demand over a single period:

X : individual transaction sizes under a Poisson arrival process;

Y : the demand over a single period when demand arrival is modelled as a Poisson process;

λ: the transaction order arrival rate when demand arrival is modelled as a Poisson process;

U: the demand over a single period under a Bernoulli demand arrival process; note that, in this case, we consider all periods (that is, periods with no demand as well as periods with positive demands);

Z: the demand over a single period under a Bernoulli demand arrival process, but in this case considering only those periods during which demand occurs;

π: the demand arrival rate when demand arrival is modelled as a Bernoulli process.

A.1 Demand arrival rates

Under the Bernoulli demand arrival process, the probability that demand occurs during a single period is equal to π (i.e. $\mathrm{P}(U > 0) = \pi$). Alternatively, if demand is assumed to follow a compound Poisson distribution, then the probability that demand occurs during a single period is given by $\mathrm{P}(Y > 0) = 1 - e^{-\lambda}$:

– $\mathrm{P}(Y > 0) = 1 - e^{-\lambda}$, if the domain of the transaction size distribution does not include zero (for example, in the case of the Logarithmic series distribution or the geometric distribution of the form used in this paper).

– $\mathrm{P}(Y > 0) = 1 - \mathrm{C}e^{-\lambda} \cong 1 - e^{-\lambda}$, if the transaction size distribution has a domain that includes zero (for example, in the case of the Poisson or Pascal series distribution).

In the latter case, C is slightly greater than 1 and the approximation is close for the two distributions in question. In particular:

– In the case of the Poisson-Poisson distribution, $\mathrm{C} = e^{\lambda Exp\,(-\rho)}$ where λ is the demand arrival rate and ρ is the parameter of the Poisson-distributed transaction sizes. For example, for transaction sizes with a mode equal to 2 (i.e. $\rho = 2$), $\mathrm{C} = 1.07$ for a

moderately intermittent demand series with $\lambda = 0.5$. C quickly tends to 1 as λ decreases and ρ as increases.

- For the Poisson-Pascal distribution, $\mathrm{C} = e^{\lambda p^r}$ where r and p are the parameters of the Pascal-distributed transaction sizes. The mode of the Pascal distribution is equal to $m = (r-1)(1-p)/p$. Thus if we express C in terms of r and m, we get $\mathrm{C} = Exp\left(\lambda\left(\frac{r-1}{m+r-1}\right)^r\right)$. The Pascal distribution takes positive integral values. For transaction sizes with a mode equal to 2 (i.e. $m = 2$), the term $\left(\frac{r-1}{m+r-1}\right)^r$ is monotonically increasing with an asymptotic limit equal to $Exp(-m)$. Thus, for a moderately intermittent demand series with $\lambda = 0.5$, $\mathrm{C} \leq 1.07$. The upper bound C will quickly tend to 1 as λ decreases and m as increases.

The approximation $\mathrm{P}(Y > 0) \cong 1 - e^{-\lambda}$ is therefore bound to be sufficiently accurate. In any case, it is important to keep in mind that the boundary is only a guide for practical purposes. Even if the true distribution of demand was known, there is always a probability that the SKU will at some point cross the boundary momentarily purely as a result of sampling error. The SKUs might therefore be assigned to different categories at various points in time.

The variables Y and U both represent the demand over a single period, the only difference being that Y represents the demand in a continuous arrival model whereas U represents the demand in a discrete arrival model. Given that $\mathrm{P}(U > 0) \equiv \mathrm{P}(Y > 0)$, it follows that:

$$\pi \approx 1 - e^{-\lambda} \Rightarrow \lambda \approx -ln(1-\pi) \tag{A.1}$$

The cut-off boundary $\overline{p} = 1.32$ in the SBC scheme was derived under the assumption that inter-demand intervals are geometrically distributed. If the average of the geometrically-distributed inter-demand intervals is $\overline{p} = 1.32$, there the demand arrival rate under the associated Bernoulli arrival process is given by $\pi = 1/\overline{p} = 0.76$ Based on equation A.1 above, this demand arrival rate translates to a cut-off value of $\overline{\lambda} = 1.42$ in the case of the Poison arrival process.

A.2 Demand size and transaction size variability

As it has been previously pointed out in sub-section A.1 above, demand over a single review period is the same whether the demand is modelled as a continuous process or as a discrete process. It thus follows that $CV^2(Y) \equiv CV^2(U)$. However, in both the SBC scheme and the framework proposed in this paper, the squared coefficients of variation are calculated by considering only those instances during which demand occurs. The squared coefficient of variation used in the proposed framework ($CV^2(X)$) relates only to the transaction sizes. In the SBC scheme, the squared coefficient of variation ($CV^2(Z)$) is calculated using only the demand sizes from the review periods with positive demand.

As Figure A.1 below shows, the two coefficients ($CV^2(X)$ and $CV^2(Z)$) can be linked by if they are expressed in terms of the squared coefficient of variation of demand over a single review period (i.e. $CV^2(U)$ and $CV^2(Y)$). In the rest of this subsection, we will make a number of arguments (as indicated by the arrows) that link $CV^2(X)$ and $CV^2(Z)$.

Figure A1. Relationship between the squared coefficients of variation in the SBC scheme and the proposed framework

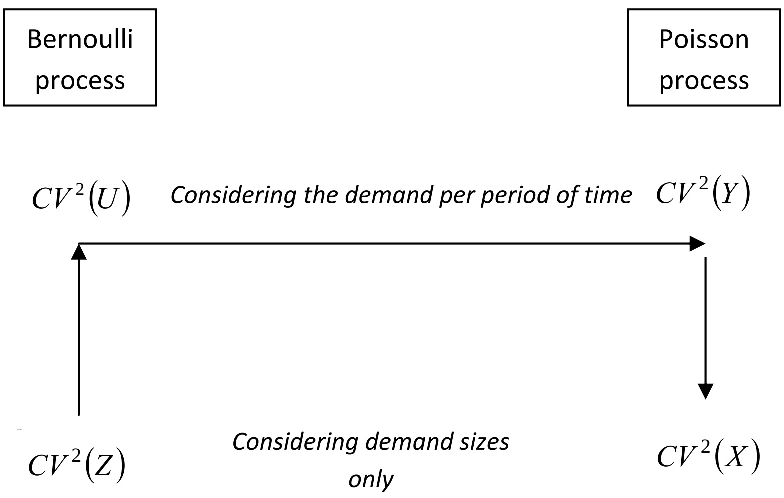

The squared coefficient of variation in the SBC scheme is calculated by considering only the demand sizes in the periods during which demand occurs. If we denote the demand sizes in such periods by Z, then the mean, variance and squared coefficient of variation of Z are given by

$$E(Z) = \sum_{i=1} i \times \mathrm{P}(Z = i) \qquad E(Z^2) = \sum_{i=1} i^2 \mathrm{P}(Z = i)$$

$$Var(Z) = E(Z^2) - [E(Z)]^2$$

$$CV^2(Z) = \frac{Var(Z)}{[E(Z)]^2} = \frac{E(Z^2) - [E(Z)]^2}{[E(Z)]^2} = \frac{E(Z^2)}{[E(Z)]^2} - 1 \tag{A.2}$$

Notice that the mean and variance of the demand sizes have been obtained by taking expectations only for positive demand sizes (i.e. $i = 1,2,3,\ldots$).

Let us now consider the demand sizes in all periods, including the periods when demand does not occur. Let us also denote the demand sizes in this case by U and the demand arrival rate by π. We then have

$$E(U)=\sum_{i=0} i\times \mathrm{P}(U=i)=0\times \mathrm{P}(U=0)+\sum_{i=1} i\times \mathrm{P}(U=i)$$
$$=0\times(1-\pi)+\pi\sum_{i=1} i\times \mathrm{P}(Z=i)=\pi E(Z)$$

$$E(U^2)=\sum_{i=0} i^2\times \mathrm{P}(U=i)=0^2\times \mathrm{P}(U=0)+\sum_{i=1} i^2\times \mathrm{P}(U=i)$$
$$=0^2\times(1-\pi)+\pi\sum_{i=1} i^2\times \mathrm{P}(Z=i)=\pi E(Z^2)$$

$$Var(U)=E(U^2)-[E(U)]^2=\pi E(Z^2)-[\pi E(Z)]^2$$

$$CV^2(U)=\frac{Var(U)}{[E(U)]^2}=\frac{E(U^2)-[E(U)]^2}{[E(U)]^2}=\frac{\pi E(Z^2)-[\pi E(Z)]^2}{[\pi E(Z)]^2}$$
$$=\frac{E(Z^2)}{\pi[E(Z)]^2}-1=\frac{1}{\pi}\left(\frac{E(Z^2)}{\pi[E(Z)]^2}-1\right)+\frac{1}{\pi}-1$$
$$=\frac{1}{\pi}\left(CV^2(Z)\right)+\frac{1}{\pi}-1=\frac{1}{\pi}\left(CV^2(Z)+1\right)-1 \qquad \text{(A.3)}$$

As it has been noted at the beginning of this subsection, $CV^2(Y)\equiv CV^2(U)$. Finally, according to equations (1) and (2) in section 3,

$$CV^2(Y)=\frac{Var(Y)}{[E(Y)]^2}=\frac{\lambda(\mu^2+\sigma^2)}{[\lambda\mu]^2}=\frac{(\mu^2+\sigma^2)}{\lambda\mu^2}=\frac{1}{\lambda}\left(1+CV^2(X)\right) \qquad \text{(A.4)}$$

Given that $CV^2(Y)\equiv CV^2(U)$, we obtain the following result by equating the terms on the left-hand sides of equations (A.3) and (A.4),

$$CV^2(Y)\equiv CV^2(Z)$$
$$\frac{1}{\lambda}\left(1+CV^2(X)\right)\equiv\frac{1}{\pi}\left(CV^2(Z)+1\right)-1$$
$$\frac{1}{\lambda}\left(1+CV^2(X)\right)\cong\frac{1}{1-e^{-\lambda}}\left(CV^2(Z)+1\right)-1$$
$$1+CV^2(X)\cong\lambda\left[\frac{1}{1-e^{-\lambda}}\left(CV^2(Z)+1\right)-1\right]$$
$$CV^2(X)\cong\lambda\left[\frac{1}{1-e^{-\lambda}}\left(CV^2(Z)+1\right)-1\right]-1$$

After substituting in the cut-off value $CV^2(Z)=0.49$ from the SBC scheme, we have

$$CV^2(X)\cong\lambda\left[\frac{1}{1-e^{-\lambda}}(1.49)-1\right]-1$$

Appendix B: Parameter estimation

Let $\bar{y}$ be the sample mean of demand, s_y^2 be the sample variance of demand and f_0 be the observed number of zero values. Then the moment estimators of the parameters of the compound Poisson distributions are given by:

Poisson-Geometric distribution

Method of first two moments:

$$\hat{\lambda} = \frac{2\bar{y}^2}{s_y^2 + \bar{y}}, \qquad \hat{\theta} = \frac{2\bar{y}}{s_y^2 + \bar{y}} = 1 - \frac{s_y^2 - \bar{y}}{s_y^2 + \bar{y}} \tag{B.1}$$

Method of mean and zero frequency:

$$\lambda^* = -ln\left(\frac{f_0}{N}\right), \qquad \theta^* = \frac{\lambda^*}{\bar{y}} \tag{B.2}$$

where λ and θ are the parameters of the Poisson and Geometric distributions respectively.

Poisson-Logarithmic series distribution

Method of first two moments:

$$\hat{r} = \frac{\bar{y}^2}{s_y^2 - \bar{y}}, \qquad \hat{p} = \frac{\bar{y}}{s_y^2} \tag{B.3}$$

Method of mean and zero frequency:

$$r^* = \frac{ln\left(\frac{f_0}{N}\right)}{ln\left(p^*\right)}, \qquad \bar{y} = \frac{ln\left(\frac{f_0}{N}\right)}{ln\left(p^*\right)}\left(\frac{1}{p^*} - 1\right) \tag{B.4}$$

when the Poisson-Logarithmic series distribution is expressed as a Negative Binomial distribution $Ne\,(r, p)$. The parameters of the negative binomial distribution can be expressed in terms of the parameters of the corresponding Poisson and logarithmic series distributions as

$$r = \frac{-\lambda}{ln(1-\varphi)}, \qquad p = (1-\varphi) \tag{B.5}$$

where λ and φ are the parameters of the Poisson and Logarithmic series distributions respectively.

<u>Poisson-Poisson distribution</u>

Method of first two moments:

$$\hat{\lambda} = \frac{\bar{y}}{\hat{\rho}}, \qquad \hat{\rho} = \frac{s_y^2 - \bar{y}}{\bar{y}} \tag{B.6}$$

Method of mean and zero frequency:

$$\frac{\bar{y}}{ln\left(\frac{f_0}{N}\right)} = \frac{\rho^*}{e^{-\rho^*} - 1}, \qquad \bar{y} = \lambda^* \rho^* \tag{B.7}$$

where λ is the parameter of the Poisson arrival process and ρ is the parameter of the Poisson-distributed transaction sizes.

<u>Poisson-Pascal distribution</u>

Per Katti and Gurland (1962),

Method of first two moments:

$$\hat{r} = \frac{\left(m_3 - 3s_y^2 + 2\bar{y}\right)\bar{y}}{\left(m_3 - 3s_y^2 + 2\bar{y}\right)\bar{y} - \left(s_y^2 - \bar{y}\right)^2} - 2, \quad \hat{p} = \frac{\bar{y}(\hat{r}+1)}{\left(s_y^2 - \bar{y}\right) + \bar{y}(\hat{r}+1)}, \quad \hat{\lambda} = \frac{\bar{y}}{\hat{r}\hat{p}} \tag{B.8}$$

where m_3 is the sample central third moment.

Thus,

$$\hat{r} > 0 \quad \Leftrightarrow \quad \left(m_3 - 3s_y^2 + 2\bar{y}\right)\bar{y} > \left(s_y^2 - \bar{y}\right)^2 > \tfrac{1}{2}\left(m_3 - 3s_y^2 + 2\bar{y}\right)\bar{y}$$

Separating the inequality,

$$\left(m_3 - 3s_y^2 + 2\bar{y}\right)\bar{y} > \left(s_y^2 - \bar{y}\right)^2 \quad \Leftrightarrow \quad \frac{-\bar{y} - \sqrt{\bar{y}^2 + 4\bar{y}(m_3 - \bar{y})}}{2} < s_y^2 < \frac{-\bar{y} + \sqrt{\bar{y}^2 + 4\bar{y}(m_3 - \bar{y})}}{2}$$

$$\left(s_y^2 - \bar{y}\right)^2 > \tfrac{1}{2}\left(m_3 - 3s_y^2 + 2\bar{y}\right)\bar{y} \quad \Leftrightarrow \quad s_y^2 < \frac{\bar{y} - \sqrt{\bar{y}^2 + 8m_3\bar{y}}}{4} \quad \text{or} \quad s_y^2 > \frac{\bar{y} + \sqrt{\bar{y}^2 + 8m_3\bar{y}}}{4}$$

and as such:

$$\frac{\bar{y} + \sqrt{\bar{y}^2 + 8m_3\bar{y}}}{4} < \sigma_y^2 < \frac{-\bar{y} + \sqrt{\bar{y}^2 + 4\bar{y}(m_3 - \bar{y})}}{2}$$

Method of mean and zero frequency:

$$\begin{aligned} &\frac{1-p^*}{p^*} ln\left(1 + \left(\frac{s_y^2 - \bar{y}}{\bar{y}} - \frac{1-p^*}{p^*}\right)\frac{ln\left(\frac{f_0}{N}\right)}{\bar{y}}\right) + \left(\frac{s_y^2 - \bar{y}}{\bar{y}} - \frac{1-p^*}{p^*}\right) ln\left(\frac{1}{p^*}\right) = 0, \\ &\quad r^* = \frac{\left(s_y^2 - \bar{y}\right)p^*}{\bar{y}\left(1 - p^*\right)} - 1, \quad \lambda^* = \frac{\bar{y}}{r^* p^*} \end{aligned} \tag{B.9}$$

The solution for p^* exists as long as

$$1 + \left(\frac{s_y^2 - \bar{y}}{\bar{y}} - \frac{1-p^*}{p^*}\right)\frac{ln\left(\frac{f_0}{N}\right)}{\bar{y}} > 0 \quad \Leftrightarrow \quad p^* < \frac{\bar{y} ln\left(\frac{f_0}{N}\right)}{\bar{y}^2 + s_y^2 \, ln\left(\frac{f_0}{N}\right)}$$

Thus,

$$0 < p^* < 1 \quad \text{if} \quad -\frac{\bar{y}^2}{ln\left(\frac{f_0}{N}\right)} < s_y^2$$

Finally, it is important to point out that for each set of parameter estimators given above, the parameter estimators in the set are correlated. For example, in the case of the moment of mean and zero frequency estimators given for the Poisson-Geometric distribution, θ^* is a function of λ^*. In other cases, such as the method of first two moments given for the Poisson-Geometric distribution, the two estimators ($\hat{\lambda}$ and $\hat{\theta}$) are both a function of the same statistics ($\bar{y}$ and s_y^2). Thus, restrictions placed on any one estimator will necessarily affect all the other estimators in the set.