2011

Distributional assumptions for parametric forecasting of intermittent demand

Aris A Syntetos, *University of Salford*

M. Z Babai

David Lengu, *University of Salford*

Nezhir Altay, *DePaul University*

Metadata of the chapter that will be visualized in OnlineFirst

Book Title	Service Parts Management	
Series Title		
Chapter Title	Distributional Assumptions for Parametric Forecasting of Intermittent Demand	
Copyright Year	2011	
Copyright HolderName	Springer-Verlag London Limited	
Corresponding Author	Family Name	Syntetos
	Particle	
	Given Name	Aris A.
	Suffix	
	Division	
	Organization	University of Salford
	Address	Salford, UK
	Email	A.Syntetos@salford.ac.uk
Author	Family Name	Babai
	Particle	
	Given Name	M. Zied
	Suffix	
	Division	
	Organization	BEM Bordeaux Management School
	Address	Bordeaux, France
	Email	Mohamed-zied.babai@bem.edu
Author	Family Name	Lengu
	Particle	
	Given Name	David
	Suffix	
	Division	
	Organization	University of Salford
	Address	Salford, UK
	Email	d.lengu@salford.ac.uk
Author	Family Name	Altay
	Particle	
	Given Name	Nezih
	Suffix	
	Division	
	Organization	DePaul University
	Address	Chicago, IL, USA
	Email	naltay@depaul.edu
Abstract	Query to the author.

Chapter 2

Distributional Assumptions for Parametric Forecasting of Intermittent Demand

Aris A. Syntetos, M. Zied Babai, David Lengu and Nezih Altay

2.1 Introduction

Parametric approaches to stock control rely upon a lead-time demand distributional assumption and the employment of an appropriate forecasting procedure for estimating the moments of such a distribution. For the case of fast demand items the Normality assumption is typically sufficient. However, Stock Keeping Units (SKUs) often exhibit intermittent or irregular demand patterns that may not be represented by the normal distribution. This is perhaps not true when lead times are very long, in which case the Normality assumption may be plausible due to the Central Limit Theorem. This issue is further discussed later in this chapter.

Intermittent demand appears at random, with some time periods having no demand at all. Moreover, demand, when it occurs, is not necessarily for a single unit or a constant demand size. In the academic literature, intermittent demand is often referred to as lumpy, sporadic or erratic demand. A conceptual framework that serves the purpose of distinguishing between such non-normal demand patterns has been discussed by Boylan et al. (2007). A demand classification framework has also been presented by Lengu and Syntetos (2009) and this is

A. A. Syntetos (✉) and D. Lengu
University of Salford, Salford, UK
e-mail: A.Syntetos@salford.ac.uk

D. Lengu
e-mail: d.lengu@salford.ac.uk

M. Z. Babai
BEM Bordeaux Management School, Bordeaux, France
e-mail: Mohamed-zied.babai@bem.edu

N. Altay
DePaul University, Chicago, IL, USA
e-mail: naltay@depaul.edu

further discussed in Sect. 5 of the chapter. Intermittent demand items may be engineering spares (e.g. Mitchell 1962; Hollier 1980; Strijbosch et al. 2000), spare parts kept at the wholesaling/retailing level (e.g. Sani 1995), or any SKU within the range of products offered by all organisations at any level of the supply chain (e.g. Croston 1972; Willemain et al. 1994). Such items may collectively account for up to 60% of the total stock value (Johnston et al. 2003) and are particularly prevalent in the aerospace, automotive and IT sectors. They are often the items at greatest risk of obsolescence.

Research in the area of forecasting and stock control for intermittent demand items has developed rapidly in recent years with new results implemented into software products because of their practical importance (Fildes et al. 2008). Key issues remaining in this area relate to (i) the further development of robust operational definitions of intermittent demand for forecasting and stock control purposes and (ii) a better modelling of the underlying demand characteristics for the purpose of proposing more powerful estimators useful in stock control. Both issues link directly to the hypothesised distribution used for representing the relevant demand patterns. Surprisingly though, not much has been contributed in this area in the academic literature.

Classification for forecasting and stock control entails decisions with respect to an appropriate estimation procedure, an appropriate stock control policy and an appropriate demand distributional assumption. The subtle linkages between operationalized SKU classification procedures and distributional assumptions have not been adequately explored. In addition, the compound nature of intermittent demand necessitates, conceptually at least, the employment of compound distributions, such as the negative binomial distribution (NBD). Although this area has attracted some academic attention (please refer also to the second section of this chapter) there is still more empirical evidence needed on the goodness-of-fit of these distributions to real data.

The objective of this work is three-fold: first, we conduct an empirical investigation that enables the analysis of the goodness-of-fit of various continuous and discrete, compound and non-compound, two-parameter statistical distributions used in the literature in the context of intermittent demand; second, we critically link the results to theoretical expectations and the issue of classification for forecasting and stock control; third, we provide an agenda for further research in this area. We use three empirical datasets for the purposes of our analysis that collectively constitute the individual demand histories of approximately 13,000 SKUs. Two datasets come from the military sector (Royal Air Force, RAF UK and US Defense Logistics Agency, DLA) and one from the Electronics industry. In all cases the SKUs are spare/service parts.

At this point it is important to note that some non-parametric procedures have also been suggested in the literature to forecast intermittent demand requirements (e.g. Willemain et al. 2004; Porras and Dekker 2008). Such approaches typically rely upon bootstrapping procedures that permit a re-construction of the empirical distribution of the data, thus making distributional assumptions redundant. Although it has been claimed that such approaches have an advantage over

Fig. 2.1 Focus of the research

parametric methodologies, more empirical research is needed to evaluate the conditions under which one approach outperforms the other. In this chapter, we will be focusing solely on parametric forecasting. The focus of our research is presented in Fig. 2.1.

The remainder of this chapter is organized as follows. In Sect. 2.2, a brief research background dealing with forecasting and stock control issues in the context of intermittent demand is presented along with a review on the demand distributions discussed in the literature and/or used by practitioners. In Sect. 2.3, we present the datasets used for the purpose of this empirical investigation, the statistical goodness-of-fit tests that we have performed and the empirical results. A critical discussion of the empirical findings follows in Sects. 2.4 and 2.5. Finally, the conclusions of our research along with some natural extensions for further work in this area are given in Sect. 2.6.

2.2 Research Background

In this section, a brief review of the literature on issues related to parametric forecasting of intermittent demand is presented. First we address the issue of estimating the mean and variance of intermittent demands, followed by a discussion of various suggestions that have been made in the literature with regards to the hypothesized distribution of such demands.

2.2.1 Parametric Forecasting

Practical parametric approaches to inventory management rely upon estimates of some essential demand distribution parameters. The decision parameters of the inventory systems (such as the re-order point or the order-up-to-level) are then based on these estimates.

Different inventory systems require different variables to be forecasted. Some of the most cited, for example (R, s, S) policies (Naddor 1975; Ehrhardt and Mosier 1984), require only estimates of the mean and variance of demand. (In such systems, the inventory position is reviewed every R periods and if the stock level drops to the re-order point s enough is ordered to bring the inventory position up to the order-up-to-level S .)

In other cases, and depending on the objectives or constraints imposed on the system, such estimates are also necessary, although they do not constitute the 'key' quantities to be determined. We may consider, for example, an (R, S) or an (s, Q) policy operating under a fill-rate constraint—known as P_2 . (In the former case, the inventory position is reviewed periodically, every R periods, and enough is ordered to bring it up to S . In the latter case, there is a continuous review of the inventory position and as soon as that drops to, or below, s an order is placed for a fixed quantity Q .) In those cases we wish to ensure that $x\%$ of demand is satisfied directly off-the-shelf and estimates are required for the probabilities of any demands exceeding S or s (for the (R, S) and (s, Q) policy, respectively). Such probabilities are typically estimated indirectly, based on the mean demand and variance forecast in conjunction with a hypothesized demand distribution. Nevertheless, and as discussed in the previous section, a reconstruction of the empirical distribution through a bootstrapping (non-parametric) procedure would render such forecasts redundant; this issue is further discussed in this Handbook in Chapter 6. Similar comments apply when these systems operate under a different service driven constraint: there is no more than $x\%$ chance of a stock-out during the replenishment cycle (this service measure is known as P_1). Consequently, we need to estimate the $(100 - x)$ th percentile of the demand distribution.

In summary, parametric approaches to forecasting involve estimates of the mean and variance of demand. In addition, a demand distribution needs also to be hypothesized, in the majority of stock control applications, for the purpose of estimating the quantities of interest. Issues related to the hypothesized demand distribution are addressed in the following sub-section. The estimation of the mean and variance of demand is addressed in Chapter 11 of this Handbook.

2.2.2 The Demand Distribution

Intermittent demand patterns are characterized by infrequent demands, often of variable size, occurring at irregular intervals. Consequently, it is preferable to model demand from constituent elements, i.e. the demand size and inter-demand

interval. Therefore, compound theoretical distributions (that explicitly take into account the size-interval combination) are typically used in such contexts of application. We first discuss some issues related to modelling demand arrivals and hence inter-demand intervals. We then extend our discussion to compound demand distributions.

If time is treated as a discrete (whole number) variable, demand may be generated based on a Bernoulli process, resulting in a geometric distribution of the inter-demand intervals. When time is treated as a continuous variable, the Poisson demand generation process results in negative exponentially distributed inter-arrival intervals.

There is sound theory in support of both geometric and exponential distribution for representing the time interval between successive demands. There is also empirical evidence in support of both distributions (e.g. Dunsmuir and Snyder 1989; Kwan 1991; Willemain et al. 1994; Janssen 1998; Eaves 2002). With Poisson arrivals of demands and an arbitrary distribution of demand sizes, the resulting distribution of total demand over a fixed lead time is compound Poisson. Inter-demand intervals following the geometric distribution in conjunction with an arbitrary distribution for the sizes, results in a compound binomial distribution.

Regarding the compound Poisson distributions, the stuttering Poisson, which is a combination of a Poisson distribution for demand occurrence and a geometric distribution for demand size, has received the attention of many researchers (for example: Gallagher 1969; Ward 1978; Watson 1987). Another possibility is the combination of a Poisson distribution for demand occurrence and a normal distribution for demand sizes (Vereecke and Verstraeten 1994), although the latter assumption has little empirical support. Particularly for lumpy demands, the demand size distribution is heavily skewed to the right, rendering the normality assumption far from appropriate. Quenouille (1949) showed that a Poisson-Logarithmic process yields a negative binomial distribution (NBD). When event arrivals are assumed to be Poisson distributed and the order size is not fixed but follows a logarithmic distribution, total demand is then negative binomially distributed over time.

Another possible distribution for representing demand is the gamma distribution. The gamma distribution is the continuous analogue of the NBD and “although not having a priori support [in terms of an explicit underlying mechanism such as that characterizing compound distributions], the gamma is related to a distribution which has its own theoretical justification” (Boylan 1997, p. 168). The gamma covers a wide range of distribution shapes, it is defined for non-negative values only and it is generally mathematically tractable in its inventory control applications (Burgin and Wild 1967; Burgin 1975; Johnston 1980). Nevertheless if it is assumed that demand is discrete, then the gamma can be only an approximation to the distribution of demand. At this point it is important to note that the use of both NBD and gamma distributions requires estimation of the mean and variance of demand only. In addition, there is empirical evidence in support of both distributions (especially the former) and therefore they are recommended for practical applications.

If demand occurs as a Bernoulli process and orders follow the Logarithmic-Poisson distribution (which is not the same as the Poisson-Logarithmic process that yields NBD demand) then the resulting distribution of total demand per period is the log-zero-Poisson (Kwan 1991). The log-zero-Poisson is a three parameter distribution and requires a rather complicated estimation method. Moreover, it was found by Kwan (1991) to be empirically outperformed by the NBD. Hence, the log-zero Poisson cannot be recommended for practical applications. One other compound binomial distribution appeared in the literature is that involving normally distributed demand sizes (Croston 1972, 1974). However, and as discussed above, a normality assumption is unrealistic and therefore the distribution is not recommended for practical applications.

Despite the inappropriateness of the normal distribution for representing demand sizes it may in fact constitute a reasonable assumption for lead time demand itself, when lead times are long (see also Syntetos and Boylan 2008). This is because long lead times permit central limit theorem effects for the sum of demands over the corresponding period, thus making the normality assumption more plausible. In addition, the assumption of normality may also be likely to be good when the coefficient of variation (CV) of the distribution of demand per period is small. Finally, algorithms based on normality are simple to implement making the normal distribution a very commonly assumed one among practitioners.

For very slow moving items, such as those commonly encountered in a military context for example, the Poisson distribution is known to offer a very good fit and much of the stock control theory in this area has been developed upon the explicit assumption that demand per period is Poisson distributed (see, for example, Silver et al. 1998). In this case demand is assumed to arrive as a Poisson process couple with unit-sized transactions. In an early work, Friend (1960) also discussed the use of a Poisson distribution for demand occurrence, combined with demands of constant size. Vereecke and Verstraeten (1994) presented an algorithm developed for the implementation of a computerised stock control system for spare parts in a chemical plant. The demand was assumed to occur as a Poisson process with a package of several pieces being requested at each demand occurrence. The resulting distribution of demand per period was called a 'Package Poisson' distribution. The same distribution has appeared in the literature under the name 'hypothetical SKU' (h-SKU) Poisson distribution (Williams 1984), where demand is treated as if it occurs as a multiple of some constant, or 'clumped Poisson' distribution, for multiple item orders for the same SKU of a fixed 'clump size' (Ritchie and Kingsman 1985). The 'Package Poisson' distribution requires, as the Poisson distribution itself, an estimate of the mean demand only.

The short review of the literature presented above indicates that it is worthwhile testing the empirical goodness-of-fit of the following distributions: (i) Poisson; (ii) NBD; (iii) stuttering Poisson; (iv) Gamma; and (v) Normal. In the next section we conduct such tests and we comment on the plausibility of the relevant assumptions for applications in an intermittent demand context.

2.3 Empirical Investigation

In this section, we first describe the datasets used for the purposes of this empirical investigation, followed by a discussion of the statistical goodness-of-fit tests conducted and the empirical results.

2.3.1 Empirical Data

The empirical databases available for the purposes of our research come from the US Defense Logistics Agency (DLA), Royal Air Force (RAF) and Electronics Industry and they consist of the individual monthly demand histories of 4,588, 5,000 and 3,055 SKUs, respectively. Some information regarding these datasets is presented in Table 2.1, followed by detailed descriptive statistics on the demand data series characteristics for each of the datasets presented in Tables 2.2, 2.3, and 2.4. At this point it should be noted that the time series considered have not been tested for stationarity.

2.3.1.1 Statistical Goodness-of-Fit Tests

Two tests have been mainly used and discussed in the literature for checking statistically significant fit, namely: the Chi-Square test and the Kolmogorov–Smirnov (K–S) test (see, for example, Harnett and Soni 1991). These tests measure

Table 2.1 Empirical datasets

#	Country	Industry	No of SKUs	Time bucket	History length	Lead-time info	Cost info
1	USA	Military/ DLA	4,588	Month	60	No	No
2	UK	Military/ RAF	5,000	Month	84	Yes	Yes
3	Europe	IT	3,055 12,643	Month	48	Constant = 3	Yes

Table 2.2 Dataset #1—US Defense Logistics Agency

4,588 SKUs	Demand intervals		Demand sizes		Demand per period	
	Mean	SD	Mean	SD	Mean	SD
Min	1.000	0.000	1.000	0.000	0.083	0.279
25% _o	1.967	1.665	2.894	2.314	0.650	1.672
Median	3.278	3.236	5.375	5.142	1.750	3.749
75% _o	5.600	6.049	11.940	12.435	4.550	9.403
Max	14.750	24.597	1326.875	1472.749	783.917	1219.012

Table 2.3 Dataset #2—Royal Air Force

5,000 SKUs	Demand intervals		Demand sizes		Demand per period	
	Mean	SD	Mean	SD	Mean	SD
Min	3.824	0.000	1.000	0.000	0.036	0.187
25%	7.273	5.431	1.556	0.815	0.155	0.538
Median	9.000	6.930	3.833	3.062	0.369	1.452
75%	11.571	8.630	11.333	9.315	1.155	4.434
Max	24.000	16.460	668.000	874.420	65.083	275.706

Part of Dataset #2 has been used in the following study: Syntetos et al. (2009)

Table 2.4 Dataset #3—electronics

3,055 SKUs	Demand intervals		Demand sizes		Demand per period	
	Mean	SD	Mean	SD	Mean	SD
Min	1.000	0.000	1.000	0.000	0.042	0.245
25%	1.500	1.011	3.462	3.011	0.896	2.215
Median	2.556	2.285	5.900	6.220	2.104	4.501
75%	4.700	4.389	12.122	13.863	6.010	10.480
Max	24.000	32.527	5366.188	9149.349	5366.188	3858.409

Dataset #3 has been used in the following study: Babai et al. (2009)

the degree of fit between observed and expected frequencies. Problems often arise with the standard Chi-Square test through the requirement that data needs to be grouped together in categories to ensure that each category has an expected frequency of at least a minimum of a certain number of observations. Some modifications of this test have also been considered in the literature. A modified Chi-Square test has been developed for the purpose of testing the goodness-of-fit for intermittent demands (Eaves, 2002). This test differs in that boundaries are specified by forming a certain number of categories with similar expected frequencies throughout, rather than combining groups just at the margins. However, the implementation of this test requires the specification of the number of categories to be used. We encountered a difficulty in using the standard or modified Chi-Square test in our research, namely that of deciding how to specify the categories' intervals or the number of categories. On the other hand, the K-S test does not require grouping of the data in any way, so no information is lost; this eliminates the troublesome problem of categories' intervals specification.

In an inventory context one could argue that measures based on the entire distribution can be misleading (Boylan and Syntetos 2006). A good overall goodness-of-fit statistic may relate to the chances of low demand values, which can mask poor forecasts of the chances of high-demand values. However, for inventory calculations, attention should be restricted to the upper end of the distribution (say the 90th or 95th percentiles). The development of modified goodness-of-fit tests for application in inventory control, and even more specifically in an intermittent demand context, is a very important area but not one considered as

part of this research. Consequently, we have selected the K–S test for the purpose of assessing goodness-of-fit.

The K–S test assumes that the data is continuous and the standard critical values are exact only if this assumption holds. Several researchers (e.g. Noether 1963, 1967; Walsh 1963; Slakter 1965) have found that the standard K–S test is conservative when applied to data that is discrete. The standard exact critical values provided for the continuous data are larger than the true exact critical values for discrete data. Consequently, the test is less powerful if the data is discrete as in the case of this research; it could result in accepting the null hypothesis at a given significance level while the correct decision would have been to reject the null hypothesis. Conover (1972) proposed a method for determining the exact critical levels for discrete data.

As discussed in the previous section, we are considering five distributions the fit of which is tested on the demand data related to 12,643 SKUs. 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 associated with the dataset 1. (Although this may be very restrictive regarding the performance of the normal distribution, this would still be expected to perform well on the time series that are associated with a small coefficient of variation of demand per period.)

Critical values have been computed based on K–S statistical tables for 1 and 5% significance levels. We consider that:

- There is a ‘Strong Fit’ if the P -value is less than both critical values;
- There is ‘Good Fit’ if the P -value is less than the critical value for 1% but larger than the one for 5%;
- There is ‘No Fit’ if the P -value is larger than both critical values.

2.3.1.2 Empirical Results

In Table 2.5 we present the percentage of SKUs that satisfy the various degrees of goodness-of-fit taken into account in our research, for each of the datasets and statistical distributions considered.

As shown in Table 2.5, the discrete distributions, i.e. Poisson, NBD and stuttering Poisson provide, overall, a better fit than the continuous ones, i.e. Normal and Gamma. More precisely, and with regards to ‘Strong Fit’, the stuttering Poisson distribution performs best in all three datasets considered in our research. This is followed by the NBD and then by the Poisson distribution. On the other hand, the normal distribution is judged to be far from appropriate for intermittent demand items; this is partly due to the experimental structure employed for the purposes of our investigation that relied upon the distribution of demand per time period rather than the distribution of the lead time demand.

Contrary to our expectations, the gamma distribution has also been found to perform poorly. This may be explained in terms of the inconsistency between the distribution under concern, which is continuous in nature, and the discreteness of the (demand) data employed in our goodness-of-fit tests. We return to this issue in

Table 2.5 Goodness-of-fit results

Dataset #	No of SKUs	Distribution	Percentage of SKUs (%)		
			Strong fit	Good fit	No fit
1	4,588	Poisson	39.45	5.51	55.04
		NBD	71.19	3.86	24.95
		Stuttering Poisson	84.18	3.64	12.18
		Normal	11.84	14.25	73.91
		Gamma	13.84	3.88	82.28
2	5,000	Poisson	59.84	2.94	37.22
		NBD	82.48	2.7	14.82
		Stuttering Poisson	98.64	0.48	0.88
		Normal	12.2	18.12	69.68
		Gamma	19.2	12.32	68.48
3	3,055	Poisson	32.64	7.4	59.96
		NBD	73.94	5.31	20.75
		Stuttering Poisson	79.05	4.49	16.46
		Normal	9.92	14.34	75.74
		Gamma	11.69	3.83	84.48

the last section of the chapter where the next steps of our research are discussed in detail.

2.4 Linking the Goodness-of-Fit to Demand Characteristics

Johnston and Boylan (1996) offered for the first time an operationalised definition of intermittent demand for forecasting purposes (demand patterns associated with an average inter-demand interval (p) greater than 1.25 forecast revision periods). The contribution of their work lies on the identification of the average inter-demand interval as a demand classification parameter rather than the specification of an exact cut-off value. Syntetos et al. (2005) took this work forward by developing a demand classification scheme that it relies upon both p and the squared coefficient of variation of demand sizes (CV^2), i.e. the contribution of their work lies in the identification of an additional categorisation parameter for demand forecasting purposes. Nevertheless, inventory control issues and demand distributional assumptions were not addressed. Boylan et al. (2007) assessed the stock control implications of the work discussed above by means of experimentation on an inventory system developed by a UK-based software manufacturer. The researchers demonstrated, empirically, the insensitivity of the p cut-off value, for demand classification purposes, in the approximate range 1.18–1.86 periods.

In this section, we attempt to explore the potential linkages between demand distributional assumptions and the classification scheme developed by Syntetos et al. (2005). In the following figures we present for dataset #1 and each of the distributions considered, the SKUs associated with a ‘Strong Fit’ as a function of

the inter-demand intervals (p) and the squared demand coefficient of variation (CV^2). The relevant results for the other two datasets are presented in the Appendix.

As shown in the figures presented below/in the Appendix and theoretically expected, both the stuttering Poisson and the Negative Binomial distribution perform comparatively better for all the datasets considered. This is true both for the SKUs with high inter-demand intervals (e.g. SKUs with p being up to 14.75 in dataset #1 or SKUs with a p value up to 24 in datasets #2 and #3) and low demand intervals (e.g. SKUs with p values starting from 1 in datasets #1 and #3). Moreover, it should be noted that there is a strong fit of NBD and stuttering Poisson to all the SKUs that are also associated with a strong fit of the Poisson distribution, which is expected since both distributions under concern are compound Poisson ones. The SKUs where there is commonly a strong fit of those three distributions are the ones characterized by relatively low CV^2 values (Figs. 2.2, 2.3, and 2.4).

Furthermore, the normal distribution performs well for the SKUs with relatively low inter-demand intervals (e.g. SKUs with p values close to 1 in datasets #1 and #3 and $p = 3.82$ in the dataset #2). However, there are also a few SKUs with high inter-demand intervals (p going up to 12 in dataset #1, 24 in dataset #2 and 15 in dataset #3) for which the normal distribution provides a strong fit. Those latter SKUs have a minimum CV^2 (i.e. $CV^2 = 0$) which can be explained by the fact that their demand is very low (in most of the cases, the demand is equal to zero and one) and can fit to the normal distribution with low mean (i.e. equivalently high values of p) and variance. As shown in Figs. 2.5 and 2.6, in addition to the SKUs where there is a fit to the normal distribution (those with low values of p), the

Fig. 2.2 Dataset #1—goodness-of-fit results for the Poisson distribution

Fig. 2.3 Dataset #1—goodness-of-fit results for the NBD

Fig. 2.4 Dataset #1—goodness-of-fit results for the stuttering Poisson

Fig. 2.5 Dataset #1—goodness-of-fit results for the normal distribution

Fig. 2.6 Dataset #1—goodness-of-fit results for the gamma distribution

gamma distribution provides also a strong fit to the SKUs with very high values of p (i.e. SKUs with an inter-demand interval going up to 15 periods in dataset #1 and 24 periods in datasets #2 and #3) and high CV^2 values (i.e. SKUs with CV^2 up to 6 in dataset #1, $CV^2 = 10$ in the dataset #2 and $CV^2 = 8$ in the dataset #3). This is also expected since the gamma distribution is known to be very flexible in terms of its mean and variance, so it can take high values for its p and CV^2 and can be reduced to the normal distribution for certain parameters of the mean and the variance.

Based on the goodness-of-fit results presented in this section, we have attempted to derive inductively an empirical rule that suggests which distribution should be used under particular values of the inter-demand interval and squared coefficient of variation of the demand sizes. That is to say, we have explored the possibility of extending the classification scheme discussed by Syntetos et al.

Fig. 2.7 Demand distributional assumptions: an inductive classification rule

(2005) to demand distributional assumptions. An inductive *Rule* has been developed (please refer to Fig. 2.7) based on the reported empirical performance of the distributions considered in our research in relation to specific values of p and CV^2 . The *Rule* suggests appropriate regions for the selection of these distributions, i.e. Normal is used for SKUs with ‘low’ p and CV^2 ; Poisson is used for SKUs with low CV^2 , Gamma is used for SKUs with ‘extreme’ values of p and CV^2 , NBD and stuttering Poisson (SP) is used for the other ranges. The stock control implications of using such a rule were evaluated through the use of the Syntetos–Boylan Approximation (Syntetos and Boylan, 2005) for forecasting purposes and the standard order-up-to-level stock control policy for a specified target cycle service level. Inventory volumes and achieved service levels were compared against those obtained from the same inventory management system that relies though upon a single demand distributional assumption, i.e. NBD. However, the results indicated no superior empirical performance of the ‘new’ approach. This may be explained in terms of the construction of the goodness-of-fit testing that considers the entire demand distribution whereas stock control performance is explicitly dependant upon the fit on the right-hand tail of the distributions. This is an important issue in Inventory Management and one that has not received adequate attention in the academic literature. We return to this issue in the last section of this chapter.

2.5 Theoretical Expectations

Lengu and Syntetos (2009) proposed a demand classification scheme based on the underlying demand characteristics of the SKUs (please refer to Fig. 2.8). SKUs are first categorised as *non-qualifying* if the variance of the demand per period is less than the mean or *qualifying* if the variance is at least equal to the mean. Compound Poisson distributions can be used to model the demand series of qualifying SKUs but they are regarded as not appropriate for modelling the demand of non-qualifying SKUs. Let us assume that demand is generated from a compound Poisson

Fig. 2.8 Compound Poisson modelling of intermittent series

model (i.e. demand ‘arrives’ according to a Poisson process and, when demand occurs, it follows a specified distribution). If we let μ denote the mean demand per unit time period and σ^2 denote the variance of demand per unit period of time, then

$$\mu = \lambda \mu_z \quad (1)$$

$$\sigma^2 = \lambda (\mu_z^2 + \sigma_z^2) \quad (2)$$

where λ is the rate of demand occurrence, and μ_z and σ_z^2 the mean and variance, respectively, of the transaction size when demand occurs. Note that

$$\frac{\sigma^2}{\mu} = \frac{\lambda (\mu_z^2 + \sigma_z^2)}{\lambda \mu_z} \geq 1 \quad (3)$$

since $\mu_z^2 \geq \mu_z$ (the transaction size is at least of 1 unit) and σ_z^2 is non-negative. The compound Poisson demand model is therefore not appropriate for SKUs associated with $\sigma^2/\mu < 1$ (non-qualifying). Note that the actual rate of demand occurrence λ does not affect the classification of SKUs as to whether they are qualifying or not.

2.5.1 Poisson-Geometric Compound Distribution (stuttering Poisson)

The Geometric distribution $\text{Ge}(\pi_G)$ is a discrete monotonically decreasing distribution with $0 \leq \text{CV}^2 \leq 1$ and mode $\tilde{m} = 1$. It can model transaction sizes that are

usually equal to one but can also take higher values. The Poisson-Geometric compound distribution also accommodates the case of clumped demand since the Poisson distribution is a special case of the Poisson-Geometric distribution. Specifically, if the parameter of the Geometric distribution $\text{Ge}(\pi_G)$ is 1, then the transaction size can only take one value (transaction size 1). With the transaction size being clumped, the demand model is now reduced to a standard Poisson distribution. In the empirical goodness-of-fit tests, the Poisson-Geometric distribution provided the most frequent fit of all the distributions considered (see Table 2.5).

2.5.2 Poisson-Logarithmic Series Compound Distribution (NBD)

The Logarithmic series distribution $\text{Log}(\pi_L)$ is a discrete monotonically decreasing distribution with an unbounded CV^2 and $\tilde{m} = 1$. Just like the Geometric distribution, the Logarithmic distribution can model transaction sizes that are constant or monotonically decreasing. However, unlike the Geometric distribution the parameter CV^2 does not have an upper bound. The Poisson-Logarithmic series compound distribution is therefore more flexible and can accommodate SKUs with exceptionally large transaction sizes. In the empirical goodness-of-fit tests, the Poisson-Logarithmic series distribution provided the second most frequent fit after the stuttering Poisson distribution.

The work discussed in this section has been developed under the assumption that demand arrivals follow a Poisson process. Similar results would be obtained if demand was assumed to occur according to a Bernoulli process since when the probability of more than one occurrence per period is negligible the Poisson and Bernoulli distributions are nearly identical. In such cases, the Poisson distribution, $P_0(\lambda)$, is approximately equal to the Bernoulli distribution with:

$$P(0) = \exp(-\lambda) \quad \text{and} \quad P(1) = 1 - \exp(-\lambda).$$

2.5.3 Non-Qualifying SKUs

While qualifying SKUs can be reasonably modelled using compound distributions, modelling non-qualifying SKUs is more challenging. Adan et al. (1995) proposed using a Binomial distribution-based model for what is termed as non-qualifying SKUs for the purposes of our research. Note that for the binomial distribution $\text{Bi}(n, p)$, $\sigma^2/\mu = npq/np = q < 1$; the binomial distribution can therefore accommodate non-qualifying SKUs. We are not aware of any empirical studies conducted to determine whether the model proposed by Adan et al. may provide adequate fit for real-life demand series. Moreover, it is not possible from that

model to distinguish between the demand occurrence process and the transaction size distribution. Such a model could however be useful for modeling slow-moving non-qualifying SKUs and we will consider it in the next steps of our research.

2.5.4 Other Considerations

The normal distribution and the gamma distribution seem to be the least promising of all the distributions considered in the empirical part of this chapter. For either distribution, the variance can be less than, equal to or larger than the mean. The two distributions can therefore be used to model both qualifying and non-qualifying SKUs. Furthermore, the normal distribution and the gamma distributions have been studied extensively and tables of the critical values for both distributions are widely available. However, in the empirical study, the two distributions provided the least frequent fit and there is no clear pattern associated with the SKUs for which the distributions provided a good fit. The normal distribution and the gamma distribution might be convenient to use but that should be contrasted to their rather poor empirical performance.

As we have mentioned in Sect. 2.2, that the K-S test assumes that the data is continuous and the test is less powerful if the data is discrete as in the case of this research. The standard exact critical values provided for the continuous data are larger than the true exact critical values for discrete data. Conover (1972) and Pettitt and Stephens (1977) proposed a method for determining the exact critical levels for the K-S test for discrete data. Choulakian et al. (1994) proposed a method of calculating the critical values of the Cramér-von Mises test and the Anderson-Darling test for discrete data. These tests have one significant drawback because of their sensitivity: their critical values are very much dependent upon the model being tested. Different tables of the critical values are therefore required for each demand model being tested. Steele and Chaselling (2006) have compared the power of these different goodness-of-fit tests for discrete data but their study was not extensive enough to indicate which test is the most powerful for our purposes.

2.6 Conclusions and Further Research

Parametric approaches to forecasting rely upon an explicit demand distributional assumption. Although the normal distribution is typically adequate for 'fast' demand items this is not true when demand is intermittent. Some research has been conducted with regards to the hypothesised distributions needed for representing such patterns and a number of distributions have been put forward as potential candidates on the basis of: (i) theoretical arguments, (ii) intuitive appeal; (iii) empirical support. A review of the literature though reveals that: (i) more empirical

studies are required in order to develop our understanding on the adequacy of these distributions under differing underlying intermittent demand structures; (ii) there is some scope for linking demand distributional assumptions to classification for forecasting and stock control purposes. Both these issues are explored as part of the research work presented in this chapter. The empirical databases available for the purposes of our investigation come from the US DLA, RAF and Electronics Industry and they consist of the individual monthly demand histories of 4,588, 5,000 and 3,055 SKUs, respectively.

The empirical goodness-of-fit of five distributions (of demand per period) has been assessed by means of employing the Kolmogorov–Smirnov (K–S) test. These distributions are: Poisson, Negative Binomial Distribution (NBD), stuttering Poisson, Normal and Gamma. The results indicate that both the NBD and stuttering Poisson provide the most frequent fit. Both these distributions are compound in nature, meaning that they account explicitly for a demand arrival process (Poisson) and a different distribution for the transaction sizes (Log series and Geometric for the NBD and stuttering Poisson, respectively). Despite previous claims, the gamma distribution does not perform very well and the same is true for the normal distribution. This may be attributed to the continuous nature of these distributions (since their fit is tested on discrete observations) but also to the fact that we model demand per unit time period as opposed to lead time demand. Upon reflection, this is viewed as a limitation of our work since lead time demand could have been considered for two of the three datasets available to us (in those cases the actual lead time was available). If that was the case, both the Normal and gamma distribution would be associated potentially with a better performance. The Poisson distribution provides a ‘reasonable’ fit and this is theoretically expected for slow moving items.

Some recent work on the issue of demand classification (Syntetos et al. 2005) has focused on both the demand arrival pattern and distribution of the demand sizes. In this chapter, we have attempted empirically to link the goodness-of-fit of the above discussed distributions to the classification scheme proposed by Syntetos et al. (2005). Although some of the results were matched indeed by relevant theoretical expectations this was not the case when the inventory implications of the proposed scheme were considered. Goodness-of-fit tests focus on the entire demand distribution whereas stock control performance is explicitly dependant upon the fit on the right-hand tail of a distribution. This is an important issue in Inventory Management and one that has not received adequate attention in the academic literature. The empirical results discussed above have also been contrasted to some theoretical expectations offered by a conceptual demand classification framework presented by Lengu and Syntetos (2009). The framework links demand classification to some underlying characteristics of intermittent demand patterns and although it seems capable of explaining a number of empirical results it may not be utilized in an operationalised fashion yet.

The work presented in this chapter has revealed a number of interesting themes for further research. Distributional assumptions play a critical role in practical inventory management applications and further work on the following

issues should prove to be valuable both from a theoretical and practitioner perspective:

- The development of modified goodness-of-fit tests for application in inventory control, and even more specifically in an intermittent demand context, is a very important area. In particular, putting more emphasis on the right-hand tail of the distribution seems appropriate for stock control applications.
- Quantifying the effect that the inconsistency between the discrete nature of demand data and the continuous nature of certain distributions may have on goodness-of-fit statistics constitutes an interesting research question.
- The inconsistency between the discrete nature of demand observations and the implicit assumption of continuous data employed by various goodness-of-fit tests should be further explored.
- Replication of the analysis conducted in this chapter in larger demand datasets coupled with the assessment of the goodness-of-fit of various distributions to the lead time demand as opposed to demand per period should help advance knowledge in this area.

Acknowledgements The research described in this chapter has been partly supported by the Engineering and Physical Sciences Research Council (EPSRC, UK) grants no. EP/D062942/1 and EP/G006075/1. More information on the former project may be obtained at <http://www.mams.salford.ac.uk/CORAS/Projects/Forecasting/>. In addition, we acknowledge the financial support received from the Royal Society, UK: 2007/Round 1 Inter. Incoming Short Visits—North America.

2.7 Appendix

2.7.1 Goodness-of-Fit Results

Figures A1, A2, A3, A4, A5, A6, A7, A8, A9, and A10

Fig. A1 Dataset #2—goodness-of-fit results for Poisson distribution

2 Distributional Assumptions for Parametric Forecasting of Intermittent Demand

19

Fig. A2 Dataset #2—goodness-of-fit results for the NBD

Fig. A3 Dataset #2—goodness-of-fit results for the stuttering Poisson

Fig. A4 Dataset #2—goodness-of-fit results for the normal distribution

Fig. A5 Dataset #2—goodness-of-fit results for gamma distribution

Fig. A6 Dataset #3—goodness-of-fit results for the Poisson distribution

Fig. A7 Dataset #3—goodness-of-fit results for the NBD

Fig. A8 Dataset #3—goodness-of-fit results for the stuttering Poisson

Fig. A9 Dataset #3—goodness-of-fit results for the normal distribution

Fig. A10 Dataset #3—
goodness-of-fit results for
gamma distribution

References

- Adan I, van Eenige M, Resing J (1995) Fitting discrete distributions on the first two moments. *Probab Eng Inf Sci* 9:623–632
- Babai MZ, Syntetos AA, Teunter R (2009) On the empirical performance of (T, s, S) heuristics. *Eur J Oper Res* (in press)
- Boylan JE (1997) The centralisation of inventory and the modelling of demand. Unpublished PhD thesis, University of Warwick, UK
- Boylan JE, Syntetos AA (2006) Accuracy and accuracy-implication metrics for intermittent demand. *FORESIGHT: Int J Appl Forecast* 4:39–42
- Boylan JE, Syntetos AA, Karakostas GC (2007) Classification for forecasting and stock control: a case study. *J Oper Res Soc* 59:473–481
- Burgin TA (1975) The gamma distribution and inventory control. *Oper Res Q* 26:507–525
- Burgin TA, Wild AR (1967) Stock control experience and usable theory. *Oper Res Q* 18:35–52
- Choulakian V, Lockhart RA, Stephens MA (1994) Cramér-von Mises statistics for discrete distributions. *Can J Stat* 22:125–137
- Conover WJ (1972) A Kolmogorov goodness-of-fit test for discontinuous distributions. *J Am Stat Assoc* 67:591–596
- Croston JD (1972) Forecasting and stock control for intermittent demands. *Oper Res Q* 23:289–304
- Croston JD (1974) Stock levels for slow-moving items. *Oper Res Q* 25:123–130
- Dunsmuir WTM, Snyder RD (1989) Control of inventories with intermittent demand. *Eur J Oper Res* 40:16–21
- Eaves A (2002) The forecasting for the ordering and stock holding of consumable spare parts. Unpublished PhD thesis, Lancaster University, UK
- Ehrhardt R, Mosier C (1984) A revision of the power approximation for computing (s, S) inventory policies. *Manag Sci* 30:618–622
- Fildes R, Nikolopoulos K, Crone S, Syntetos AA (2008) Forecasting and operational research: a review. *J Oper Res Soc* 59:1150–1172
- Friend JK (1960) Stock control with random opportunities for replenishment. *Oper Res Q* 11:130–136
- Gallagher DJ (1969) Two periodic review inventory models with backorders and stuttering Poisson demands. *AIIE Trans* 1:164–171
- Harnett DL, Soni AK (1991) *Statistical methods for business and economics*, 4th edn. Addison Wesley, New York
- Hollier RH (1980) The distribution of spare parts. *Int J Prod Res* 18:665–675
- Janssen FBSLP (1998) Inventory management systems; control and information issues. Published PhD thesis, Centre for Economic Research, Tilburg University, The Netherlands
- Johnston FR (1980) An interactive stock control system with a strategic management role. *J Oper Res Soc* 31:1069–1084
- Johnston FR, Boylan JE, Shale EA (2003) An examination of the size of orders from customers, their characterization and the implications for inventory control of slow moving items. *J Oper Res Soc* 54:833–837

- Kwan HW (1991) On the demand distributions of slow moving items. Unpublished PhD thesis, Lancaster University, UK
- Lengu D, Syntetos AA (2009) Intermittent demand: classification and distributional assumptions. Working Paper (WP) 333/09, Management and Management Sciences Research Institute (MaMS RI), University of Salford, UK
- Mitchell GH (1962) Problems of controlling slow-moving engineering spares. *Oper Res Q* 13:23–39
- Naddor E (1975) Optimal and heuristic decisions in single and multi-item inventory systems. *Manag Sci* 21:1234–1249
- Noether GE (1963) Note on the Kolmogorov statistic in the discrete case. *Metrika* 7:115–116
- Noether GE (1967) Elements of nonparametric statistics. Wiley, New York
- Pettitt AN, Stephens MA (1977) The Kolmogorov–Smirnov goodness-of-fit statistic with discrete and grouped data. *Technometrics* 19:205–210
- Porras EM, Dekker R (2008) An inventory control system for spare parts at a refinery: an empirical comparison of different reorder point methods. *Eur J Oper Res* 184:101–132
- Quenouille MH (1949) A relation between the logarithmic, Poisson and negative binomial series. *Biometrics* 5:162–164
- Ritchie E, Kingsman BG (1985) Setting stock levels for wholesaling: performance measures and conflict of objectives between supplier and stockist. *Eur J Oper Res* 20:17–24
- Sani B (1995) Periodic inventory control systems and demand forecasting methods for low demand items. Unpublished PhD thesis, Lancaster University, UK
- Silver EA, Pyke DF, Peterson R (1998) Inventory management and production planning and scheduling, 3rd edn. Wiley, New York
- Slakter MJ (1965) A comparison of the Pearson chi-square and Kolmogorov goodness-of-fit tests with respect to validity. *J Am Stat Assoc* 60:854–858
- Stephens MA (1974) EDF statistics for goodness of fit and some comparisons. *J Am Stat Assoc* 69:730–737
- Stephens MA (1976) Asymptotic results for goodness-of-fit statistics with unknown parameters. *Ann Stat* 4:357–369
- Stephens MA (1977) Goodness of fit for the extreme value distribution. *Biometrika* 64:583–588
- Strijbosch LWG, Heuts RMJ, van der Schoot EHM (2000) A combined forecast-inventory control procedure for spare parts. *J Oper Res Soc* 51:1184–1192
- Syntetos AA, Boylan JE (2005) The accuracy of intermittent demand estimates. *Int J Forecast* 21:303–314
- Syntetos AA, Boylan JE (2008) Smoothing and adjustments of demand forecasts for inventory control. *IMA J Manag Math* 19:175–192
- Syntetos AA, Boylan JE, Croston JD (2005) On the categorisation of demand patterns. *J Oper Res Soc* 56:495–503
- Syntetos AA, Babai MZ, Dallery Y, Teunter R (2009) Periodic control of intermittent demand items: theory and empirical analysis. *J Oper Res Soc* 60:611–618
- Vereecke A, Verstraeten P (1994) An inventory management model for an inventory consisting of lumpy items, slow movers and fast movers. *Int J Prod Econ* 35:379–389
- Walsh JE (1963) Bounded probability properties of Kolmogorov–Smirnov and similar statistics for discrete data. *Ann Inst Stat Math* 15:153–158
- Ward JB (1978) Determining re-order points when demand is lumpy. *Manag Sci* 24:623–632
- Watson RB (1987) The effects of demand-forecast fluctuations on customer service and inventory cost when demand is lumpy. *J Oper Res Soc* 38:75–82
- Willemain TR, Smart CN, Shocker JH, DeSautels PA (1994) Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. *Int J Forecast* 10:529–538
- Willemain TR, Smart CN, Schwarz HF (2004) A new approach to forecasting intermittent demand for service parts inventories. *Int J Forecast* 20:375–387
- Williams TM (1984) Stock control with sporadic and slow-moving demand. *J Oper Res Soc* 35:939–948

Book: 182963_1_En
Chapter: 2

Author Query Form

Please ensure you fill out your response to the queries raised below and return this form along with your corrections

Dear Author

During the process of typesetting your chapter, the following queries have arisen. Please check your typeset proof carefully against the queries listed below and mark the necessary changes either directly on the proof/online grid or in the ‘Author’s response’ area provided below

Query	Details required	Author’s response
1.	Please provide abstract for online publication only, not in PDF.	
2.	Please check and confirm the authors and their respective affiliations are correctly identified and amend if necessary.	

MARKED PROOF

Please correct and return this set

Please use the proof correction marks shown below for all alterations and corrections. If you wish to return your proof by fax you should ensure that all amendments are written clearly in dark ink and are made well within the page margins.

<i>Instruction to printer</i>	<i>Textual mark</i>	<i>Marginal mark</i>
Leave unchanged	... under matter to remain	Ⓟ
Insert in text the matter indicated in the margin	⧵	New matter followed by ⧵ or ⧵ [Ⓢ]
Delete	/ through single character, rule or underline or ⎓ through all characters to be deleted	⧻ or ⧻ [Ⓢ]
Substitute character or substitute part of one or more word(s)	/ through letter or ⎓ through characters	new character / or new characters /
Change to italics	— under matter to be changed	↵
Change to capitals	≡ under matter to be changed	≡
Change to small capitals	≡ under matter to be changed	≡
Change to bold type	~ under matter to be changed	~
Change to bold italic	≈ under matter to be changed	≈
Change to lower case	Encircle matter to be changed	≡
Change italic to upright type	(As above)	⧻
Change bold to non-bold type	(As above)	⧻
Insert 'superior' character	/ through character or ⧵ where required	Y or Y under character e.g. Y or Y
Insert 'inferior' character	(As above)	⧵ over character e.g. ⧵
Insert full stop	(As above)	⊙
Insert comma	(As above)	,
Insert single quotation marks	(As above)	Y or Y and/or Y or Y
Insert double quotation marks	(As above)	Y or Y and/or Y or Y
Insert hyphen	(As above)	⎓
Start new paragraph	⌞	⌞
No new paragraph	⌞	⌞
Transpose	⌞	⌞
Close up	linking ○ characters	○
Insert or substitute space between characters or words	/ through character or ⧵ where required	Y
Reduce space between characters or words		↑