



HAL
open science

On the demand distributions of spare parts

Aris A Syntetos, Mohamed Zied Babai, Nezih Altay

► **To cite this version:**

Aris A Syntetos, Mohamed Zied Babai, Nezih Altay. On the demand distributions of spare parts. International Journal of Production Research, 2011, pp.1. 10.1080/00207543.2011.562561 . hal-00717973

HAL Id: hal-00717973

<https://hal.science/hal-00717973>

Submitted on 15 Jul 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



On the demand distributions of spare parts

Journal:	<i>International Journal of Production Research</i>
Manuscript ID:	TPRS-2010-IJPR-0522.R2
Manuscript Type:	Original Manuscript
Date Submitted by the Author:	06-Jan-2011
Complete List of Authors:	Syntetos, Aris; University of Salford, Salford Business School Babai, Mohamed; BEM Bordeaux Management School Altay, Nezh; DePaul University, College of Commerce
Keywords:	INVENTORY MANAGEMENT, FORECASTING
Keywords (user):	Demand distributions, Spare parts management

SCHOLARONE™
Manuscripts

On the demand distributions of spare parts*

A.A. Syntetos¹, M.Z. Babai^{2,3} and N. Altay⁴

¹ University of Salford, UK; ² BEM-Bordeaux Management School, France;

³ King Saud University, SA; ⁴ DePaul University, USA

Abstract

Spare parts have become ubiquitous in modern societies and managing their requirements is an important and challenging task with tremendous cost implications for the organisations that are holding relevant inventories. Demand for spare parts arises whenever a component fails or requires replacement and as such the relevant patterns are different from those associated with 'typical' stock keeping units. Such demand patterns are most often intermittent in nature, meaning that demand arrives infrequently and is interspersed by time periods with no demand at all. A number of distributions have been discussed in the literature for representing these patterns but empirical evidence is lacking. In this paper, we address the issue of demand distributional assumptions for spare parts management, conducting a detailed empirical investigation on the goodness-of-fit of various distributions and their stock control implications in terms of inventories held and service levels achieved. This is an important contribution from a methodological perspective, since the validity of demand distributional assumptions (i.e. their goodness-of-fit) is distinguished from their utility (i.e. their real world implications). Three empirical datasets are used for the purposes of our research that collectively consist of the individual demand histories of approximately 13,000 SKUs from the military sector (UK & USA) and the Electronics Industry (Europe). Our investigation provides evidence in support of certain demand distributions in a real world context. The natural next steps of research are also discussed and that should facilitate further developments in this area from an academic perspective.

Keywords: Demand Distributions; Inventory; Spare Parts Management; Empirical Study

* Acknowledgements

The research described in this paper has been partly supported by the Royal Society, UK (2007/Round 1, Int. Incoming Short Visits - North America). In addition, we would like to acknowledge the financial support provided by DePaul University that allowed a research visit of Aris Syntetos to DePaul in 2010.

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 (typically mean and variance). 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 almost invariably the case for service/spare parts. Such products have become ubiquitous in modern societies and managing their requirements is an important and challenging task with tremendous cost implications for the organisations that are holding relevant inventories.

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).

As mentioned above, intermittent demand items are typically spare parts; those may be engineering spares (e.g. Mitchell, 1962; Hollier, 1980; Strijbosch *et al.*, 2000), or parts kept at the wholesaling/retailing level (e.g. Sani, 1995; Syntetos *et al.*, 2009c). Demand for such items arises whenever a component fails or requires replacement as opposed to being generated according to buying behaviours of end-consumers (and the way that demand moves upstream in a supply chain). However, intermittent demand may indeed

1
2
3 be characterizing any SKU within the range of products offered by all organisations at
4
5 any level of the supply chain (e.g. Croston, 1972; Willemain *et al.*, 1994). Such items are
6
7 typically the so-called 'slow-movers' and although they contribute little to the sales of an
8
9 organisation they may collectively account for up to 60% of the total stock value
10
11 (Johnston *et al.*, 2003). Consequently, small improvements regarding their management
12
13 may be translated into substantial organisational savings. Such items are particularly
14
15 prevalent in the aerospace, automotive and IT sectors and they are often the SKUs at
16
17 greatest risk of obsolescence. As it will be discussed later in this section, the empirical
18
19 data available for the purposes of our investigation relate to spare parts; although such
20
21 SKUs constitute the focus of our research, natural extensions with regards to our findings
22
23 can be made to all products characterised by intermittent demand structures (please refer
24
25 also to the last section of our paper).
26
27
28
29
30
31

32
33 Research in the area of forecasting and stock control for intermittent demand items has
34
35 developed rapidly in recent years (e.g. Syntetos and Boylan, 2005; Chatfield and Hayya,
36
37 2007; Gutierrez *et al.*, 2007; Quintana and Leung, 2007; Altay *et al.*, 2008) with new
38
39 results implemented into software products because of their practical importance. Key
40
41 issues remaining in this area relate to (i) the further development of robust operational
42
43 definitions of intermittent demand for forecasting and stock control purposes and (ii) a
44
45 better modelling of the underlying demand characteristics for the purpose of proposing
46
47 more powerful estimators useful in stock control. Both issues link directly to the
48
49 hypothesised distribution used for representing the relevant demand patterns. Surprisingly
50
51 though, not much has been contributed in this area in the academic literature and this
52
53 constitutes the main purpose of our work.
54
55
56
57
58
59
60

1.1 Objectives and structure of the paper

Intermittent demand patterns are built from constituent elements, namely a demand arrival process and a distribution of the demand sizes, when demand occurs. As such, the very nature of such patterns necessitates, conceptually at least, the employment of compound distributions, such as the negative binomial distribution (NBD). Although a number of distributions have been put forward in the literature (please refer also to the second section of this paper) there is still more empirical evidence needed on the goodness-of-fit of these distributions to real data. That is to say, there is a need for more insights into the empirical validity of various distributional assumptions. In addition, a distinction needs to be made between the goodness-of-fit on demand per period data and the validity of a distributional assumption for representing lead-time demand. The latter is what matters in a stock control context and this distinction has been overlooked in the academic literature. Finally, it is also important to distinguish between the empirical validity and utility of a demand distributional assumption, the latter being related to the stock control implications of using a particular distribution. 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. All the above discussed issues are addressed as part of our work.

In particular, the objective of this paper is three-fold: first, we conduct an empirical investigation that enables the analysis of the goodness-of-fit of various two (or single) parameter statistical distributions used in the literature in the context of intermittent demand; second, we extend the analysis discussed above to assess the linkage between the validity of various demand distributional assumptions and their utility (i.e. actual

1
2
3 implications for stock control in terms of inventory volumes and service levels achieved);
4
5 third, we discuss the implications of our work for Operational Research (OR) theory and
6
7 practice and we provide an agenda for further research in this area. We use three
8
9 empirical datasets for the purposes of our analysis that collectively constitute the
10
11 individual demand histories of approximately 13,000 SKUs. Two datasets come from the
12
13 military sector (Royal Air Force, RAF UK and US Defense Logistics Agency, DLA) and
14
15 one from the Electronics industry. In all cases the SKUs are spare/service parts.
16
17
18
19

20
21 At this point it is important to note that some non-parametric procedures have also been
22
23 suggested in the literature to forecast intermittent demand requirements (e.g. Willemain
24
25 *et al.*, 2004; Porras and Dekker, 2008). Such approaches typically rely upon
26
27 bootstrapping procedures that permit a re-construction of the empirical distribution of the
28
29 data, thus making distributional assumptions and estimation of the moments of such
30
31 distributions redundant. Although it has been claimed that such approaches have an
32
33 advantage over parametric methodologies, more empirical research is needed to evaluate
34
35 the conditions under which one approach outperforms the other. In this paper, we will be
36
37 focusing solely on parametric approaches to inventory management.
38
39
40
41
42
43

44
45 The remainder of this paper is organized as follows. In Section 2, an introduction to
46
47 forecasting and stock control issues in the context of intermittent demand is presented
48
49 along with a review of the demand distributions discussed in the literature and/or used by
50
51 practitioners. In Section 3, we present the datasets used for the purpose of this empirical
52
53 investigation followed by a discussion on the statistical goodness-of-fit tests that we have
54
55 considered for analysing the performance of various distributions. The empirical validity
56
57
58
59
60

1
2
3 and utility of distributional assumptions is then explored in Sections 4 and 5 for demand
4 per period data and lead-time demands respectively. Finally, the conclusions of our
5 research along with its implications and some natural extensions for further work in this
6 area are presented in Section 6.
7
8
9
10
11

12 13 14 15 **2. Research background**

16
17 In this section, a brief review of the literature on issues related to parametric inventory
18 forecasting is presented. First we establish the need for estimating the mean and variance
19 of demand in an inventory control context, followed by a discussion of various
20 suggestions that have been made in the literature with regards to the hypothesised
21 distribution of demand patterns, focusing on the case of intermittence. The second sub-
22 section culminates with the selection of a number of distributions to be considered in the
23 empirical part of our work.
24
25
26
27
28
29
30
31
32

33 34 35 **2.1 Parametric forecasting**

36
37 Practical parametric approaches to inventory management rely upon estimates of some
38 essential demand distribution parameters. The decision parameters of the inventory
39 systems (such as the Re-Order Point or the Order-Up-To-Level) are then based on these
40 estimates. Different inventory systems require different variables to be forecasted. Some
41 of the most cited, for example (R, s, S) policies (Naddor, 1975; Ehrhardt and Mosier,
42 1984), require only estimates of the mean and variance of demand. (In such systems, the
43 inventory position is reviewed every R periods and if the stock level drops to the re-order
44 point s enough is ordered to bring the inventory position up to the order-up-to-level S .)
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 In other cases, and depending on the objectives or constraints imposed on the system,
4 such estimates are also necessary, although they do not constitute the 'key' quantities to
5 be determined. We may consider, for example, an (R, S) or an (s, Q) policy operating
6 under a fill-rate constraint – known as P_2 . (In the former case, the inventory position is
7 reviewed periodically, every R periods, and enough is ordered to bring it up to S . In the
8 latter case, there is a continuous review of the inventory position and as soon as that
9 drops to, or below, s an order is placed for a fixed quantity Q .) In those cases we wish to
10 ensure that $x\%$ of demand is satisfied directly off-the-shelf and estimates are required for
11 the probabilities of any demands exceeding S or s (for the (R, S) an (s, Q) policy
12 respectively). Such probabilities are typically estimated indirectly, based on the mean
13 demand and variance forecast in conjunction with a hypothesized demand distribution.
14 Similar comments apply when these systems operate under a different service driven
15 constraint: there is no more than $x\%$ chance of a stock-out during the replenishment cycle
16 (this service measure is known as P_1). Consequently, we need to estimate the $(100-x)th$
17 percentile of the demand distribution.
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39

40 In summary, parametric approaches to forecasting involve estimates of the mean and
41 variance of demand. In addition, a demand distribution needs also to be hypothesized, in
42 the majority of stock control applications, for the purpose of estimating the quantities of
43 interest. Issues related to the hypothesized demand distribution are addressed in the
44 following sub-section. For a detailed account on issues related to the estimation of mean
45 demand and its variance in an intermittent demand context, the interested readers are
46 referred to Syntetos and Boylan (2008) and Syntetos *et al.* (2009b).
47
48
49
50
51
52
53
54
55
56
57
58
59
60

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.

1
2
3 Regarding the compound Poisson distributions, the stuttering Poisson, which is a
4 combination of a Poisson distribution for demand occurrence and a geometric distribution
5 for demand size, has received the attention of many researchers (for example: Gallagher,
6 1969; Ward, 1978; Watson, 1987). Another possibility is the combination of a Poisson
7 distribution for demand occurrence and a normal distribution for demand sizes (Vereecke
8 and Verstraeten, 1994), although the latter assumption has little empirical support.
9 Particularly for lumpy demands, the demand size distribution is heavily skewed to the
10 right, rendering the normality assumption far from appropriate. Quenouille (1949)
11 showed that a Poisson-Logarithmic process yields a negative binomial distribution
12 (NBD). When event arrivals are assumed to be Poisson distributed and the order size is
13 not fixed but follows a logarithmic distribution, total demand is then negative binomially
14 distributed over time.
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32

33 Another possible distribution for representing demand is the gamma distribution. The
34 gamma distribution is the continuous analogue of the NBD and “*although not having a*
35 *priori support* [in terms of an explicit underlying mechanism such as that characterizing
36 compound distributions], *the gamma is related to a distribution which has its own*
37 *theoretical justification*” (Boylan, 1997: 168). The gamma covers a wide range of
38 distribution shapes, it is defined for non-negative values only and it is generally
39 mathematically tractable in its inventory control applications (Burgin and Wild, 1967;
40 Burgin, 1975; Johnston, 1980). Nevertheless if it is assumed that demand is discrete, then
41 the gamma can be only an approximation to the distribution of demand. At this point it is
42 important to note that the use of both NBD and gamma distributions requires estimation
43 of the mean and variance of demand only. In addition, there is empirical evidence in
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3 support of both distributions (especially the former) and therefore they are recommended
4
5 for practical applications.
6
7

8
9
10 If demand occurs as a Bernoulli process and orders follow the Logarithmic-Poisson
11
12 distribution (which is not the same as the Poisson-Logarithmic process that yields NBD
13
14 demand) then the resulting distribution of total demand per period is the log-zero-Poisson
15
16 (Kwan, 1991). The log-zero-Poisson is a three parameter distribution and requires a
17
18 rather complicated estimation method. Moreover, it was found by Kwan (1991) to be
19
20 empirically outperformed by the NBD. Hence, the log-zero Poisson cannot be
21
22 recommended for practical applications. One other compound binomial distribution
23
24 appeared in the literature is that involving normally distributed demand sizes (Croston,
25
26 1972; 1974). However, and as discussed above, a normality assumption is unrealistic and
27
28 therefore the distribution is not recommended for practical applications.
29
30
31
32
33
34

35 Despite the inappropriateness of the normal distribution for representing demand sizes it
36
37 may in fact constitute a reasonable assumption for lead time demand itself, when lead
38
39 times are long (see also Syntetos and Boylan, 2008). This is because long lead times
40
41 permit central limit theorem effects for the sum of demands over the corresponding
42
43 period, thus making the normality assumption more plausible. In addition, the assumption
44
45 of normality may prove to be good when the coefficient of variation (CV) of the
46
47 distribution of demand per period is small. Finally, algorithms based on normality are
48
49 simple to implement making the normal distribution a very commonly assumed one
50
51
52
53
54 among practitioners.
55
56
57
58
59
60

1
2
3 For very slow moving items, such as those commonly encountered in a military context
4
5 for example, the Poisson distribution is known to offer a very good fit and much of the
6
7 stock control theory in this area has been developed upon the explicit assumption that
8
9 demand per period is Poisson distributed (see, for example, Silver *et al.*, 1998). In this
10
11 case demand is assumed to arrive as a Poisson process couple with unit-sized
12
13 transactions. In an early work, Friend (1960) also discussed the use of a Poisson
14
15 distribution for demand occurrence, combined with demands of constant size. Vereecke
16
17 and Verstraeten (1994) presented an algorithm developed for the implementation of a
18
19 computerised stock control system for spare parts in a chemical plant. The demand was
20
21 assumed to occur as a Poisson process with a package of several pieces being requested
22
23 at each demand occurrence. The resulting distribution of demand per period was called a
24
25 'Package Poisson' distribution. The same distribution has appeared in the literature under
26
27 the name 'hypothetical SKU' (*h-SKU*) Poisson distribution (Williams, 1984), where
28
29 demand is treated as if it occurs as a multiple of some constant, or 'clumped Poisson'
30
31 distribution, for multiple item orders for the same SKU of a fixed 'clump size' (Ritchie
32
33 and Kingsman, 1985). The 'Package Poisson' distribution requires, as the Poisson
34
35 distribution itself, an estimate of the mean demand only.
36
37
38
39
40
41
42
43
44

45 The review of the literature presented above indicates that it is worthwhile assessing the
46
47 empirical validity and utility of the following demand distributional assumptions: i)
48
49 Poisson; ii) NBD; iii) Stuttering Poisson; iv) Gamma; and v) Normal. In the next sections
50
51 we conduct relevant tests and we comment on the plausibility of the relevant assumptions
52
53 for applications in an intermittent demand context.
54
55
56
57
58
59
60

3. Empirical investigation

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

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 1, followed by detailed descriptive statistics on the demand data series characteristics for each of the datasets presented in Tables 2-4. (At this point it should be noted that the time series considered have not been tested for stationarity.)

INSERT TABLE 1 ABOUT HERE

INSERT TABLE 2 ABOUT HERE

INSERT TABLE 3 ABOUT HERE

INSERT TABLE 4 ABOUT HERE

Table 2-4 indicate that the three datasets considered for the purpose of our empirical investigation reflect a wide range of possible underlying demand patterns. Datasets #1 and #3 are associated with a somewhat low degree of intermittence; the average demand intervals have a median value of approximately 3 and 2.5 respectively. However, there is

1
2
3 a relatively high degree of erraticness; demand sizes are highly variable with a median
4 standard deviation of approximately 5 (6) but a maximum one equal to 1472 (9149) for
5
6 dataset #1 (#3). On the other hand, dataset #2 consists of average demand intervals that
7
8 are very large since their median value is almost 9 months but the demand sizes are
9
10 associated with a comparatively lower degree of erraticness.
11
12
13
14
15

16 **3.2 Statistical goodness-of-fit tests**

17
18 Two tests have been mainly used and discussed in the literature for checking statistically
19 significant fit, namely: the Chi-Square test and the Kolmogorov-Smirnov (K-S) test (see,
20
21 for example, Harnett and Soni, 1991). These tests measure the degree of fit between
22
23 observed and expected frequencies. Problems often arise with the standard Chi-Square
24
25 test through the requirement that data needs to be grouped together in categories to ensure
26
27 that each category has an expected frequency of at least a minimum of a certain number
28
29 of observations. Some modifications of this test have also been considered in the
30
31 literature. A modified Chi-Square test has been developed for the purpose of testing the
32
33 goodness-of-fit for intermittent demands (Eaves, 2002). This test differs in that
34
35 boundaries are specified by forming a certain number of categories with similar expected
36
37 frequencies throughout, rather than combining groups just at the margins. However, the
38
39 implementation of this test requires the specification of the number of categories to be
40
41 used. We encountered a difficulty in using the standard or modified Chi-Square test in
42
43 our research, namely that of deciding how to specify the categories' intervals or the
44
45 number of categories. On the other hand, the K-S test does not require grouping of the
46
47 data in any way, so no information is lost; this eliminates the troublesome problem of
48
49 categories' intervals specification.
50
51
52
53
54
55
56
57
58
59
60

1
2
3 In an inventory context one could argue that measures based on the entire distribution can
4 be misleading (Boylan and Syntetos, 2006). A good overall goodness-of-fit statistic may
5 relate to the chances of low demand values, which can mask poor forecasts of the chances
6 of high-demand values. However, for inventory calculations, attention should be
7 restricted to the upper end of the distribution (say the 90th or 95th percentiles). The
8 development of modified goodness-of-fit tests for application in inventory control, and
9 even more specifically in an intermittent demand context, is a very important area but not
10 one considered as part of this research. (We return to this issue in the last section of the
11 paper where the natural next steps of research are discussed.) Consequently, we have
12 selected the K-S test for the purpose of assessing goodness-of-fit.
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

28 At this point, it should be noted that the K-S test assumes that the data is continuous and
29 the standard critical values are exact only if this assumption holds. Several researchers
30 (e.g. Noether, 1963, 1967; Walsh, 1963; Slakter, 1965) have found that the standard K-S
31 test is conservative when applied to data that is discrete. The standard exact critical
32 values provided for the continuous data are larger than the true exact critical values for
33 discrete data. Consequently, the test is less powerful if the data is discrete as in the case
34 of this research; it could result in accepting the null hypothesis at a given significance
35 level while the correct decision would have been not to accept the null hypothesis.
36
37
38
39
40
41
42
43
44
45
46
47
48

49 As discussed in the previous section, we are considering five distributions the fit of which
50 is tested on the demand data related to 12,643 SKUs. First, the distribution of the demand
51 per period has been considered, followed by tests on the lead-time demand. This
52
53
54
55
56
57
58
59
60

1
2
3 distinction is a very important one since the latter distribution may in fact be considerably
4
5 different than the former and yet is the one required for stock control purposes.
6
7

8
9
10 Critical values have been computed based on K-S statistical tables for 1% and 5%
11
12 significance levels. We consider that:

- 13
- 14 • There is a 'Strong Fit' if the P -value is less than the critical value for 5%;
- 15
- 16 • There is 'Good Fit' if the P -value is less than the critical value for 1% but larger
17 than the one for 5%;
- 18
- 19 • There is 'No Fit' if the P -value is larger than the critical value for 1%.
- 20
- 21
- 22
- 23
- 24
- 25
- 26

27 **4. Empirical results: demand per period**

28
29 In Table 5 we present the percentage of SKUs that satisfy the various degrees of
30
31 goodness-of-fit taken into account in our research, for each of the datasets and statistical
32
33 distributions considered.
34

35
36
37
38 As shown in Table 5, the discrete distributions, i.e. Poisson, NBD and Stuttering Poisson
39
40 provide, overall, a better fit than the continuous ones, i.e. Normal and Gamma. More
41
42 precisely, and with regards to 'Strong Fit', the Stuttering Poisson distribution performs
43
44 best in all three datasets considered in our research. This is followed by the NBD and
45
46 then by the Poisson distribution. On the other hand, the Normal distribution is judged to
47
48 be far from appropriate for intermittent demand items; as will be discussed in the next
49
50 section, this is partly, at least, due to the experimental structure employed for this part of
51
52 our analysis that relied upon the distribution of demand per time period rather than the
53
54 distribution of the lead time demand.
55
56
57
58
59
60

INSERT TABLE 5 ABOUT HERE

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

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

4.1 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). They compared, on a wide range of theoretically generated demand patterns, the Mean Squared Error performance of Croston's method that has been specifically designed for intermittent demand estimation, and Single Exponential Smoothing (SES) that is typically used for fast-moving items. They found that the former method performs better than the latter on demand patterns associated with an average inter-demand interval (p) greater than 1.25 forecast revision periods. Their contribution lies on the identification of the average inter-demand interval as a demand classification parameter to determine intermittence 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). In their work, they compared (assuming fixed lead-times) the theoretical MSE performance of three

1
2
3 estimators: Croston, SES and SBA (Syntetos-Boylan Approximation, Syntetos and
4
5
6 Boylan, 2005) and they identified regions where one method outperforms all others.
7
8
9

10 Both schemes serve the purpose of identifying the best performing estimation procedure
11
12 in each category; however, inventory control issues and demand distributional
13
14 assumptions were not addressed. Boylan *et al.* (2007) assessed the stock control
15
16 implications of the work conducted by Syntetos *et al.* (2005) by means of
17
18 experimentation on an inventory system developed by a UK-based software
19
20 manufacturer. The researchers demonstrated, empirically, the insensitivity of the p cut-
21
22 off value, for demand classification purposes, in the approximate range 1.18 – 1.86 periods.
23
24
25 In this section, we attempt to explore the potential linkages between demand
26
27 distributional assumptions and the classification scheme developed by Syntetos *et al.*
28
29 (2005). By doing so, we aim at providing some further insights into the potential
30
31 usefulness of such distributions in an intermittent demand context.
32
33
34
35
36
37

38 In Figure 1 we present for dataset #1 and each of the distributions considered, the SKUs
39
40 associated with a ‘Strong Fit’ as a function of the inter-demand intervals (p , expressed as
41
42 a number of zero demand periods: average interval between successive demand
43
44 occurrences) and the squared demand coefficient of variation (CV^2 , dimensionless
45
46 measure) of the demand sizes. (The results for the other two datasets lead to similar
47
48 insights and they are separately presented in *Appendix A* at the end of the paper.)
49
50
51
52
53

54 -----
55 **INSERT FIGURE 1 ABOUT HERE**
56 -----
57
58
59
60

1
2
3 As shown in Figure 1 (and in *Appendix A*) and theoretically expected, both the Stuttering
4 Poisson and the Negative Binomial distribution perform comparatively better for all the
5 datasets considered. This is true both for the SKUs with high inter-demand intervals (e.g.
6 SKUs with p being up to 12 in dataset #1 or SKUs with a p value up to 24 in datasets #2
7 and #3) and low demand intervals (e.g. SKUs with p values starting from 1 in datasets #1
8 and #3). Moreover, it should be noted that there is a strong fit of NBD and Stuttering
9 Poisson to all the SKUS that are also associated with a strong fit of the Poisson
10 distribution, which is expected since both distributions under concern are compound
11 Poisson ones. The SKUs where there is commonly a strong fit of those three distributions
12 are the ones characterized by relatively low CV^2 values. The Normal distribution seems to
13 have the narrowest CV^2 range among all tested distributions, showing a good fit roughly
14 between zero and one. This implies that the Normal distribution is not a good fit for
15 erratic data. Furthermore, the Normal distribution performs well for the SKUs with
16 relatively low inter-demand intervals (e.g. SKUs with p values close to 1 in datasets #1
17 and #3 and $p = 3.82$ in the dataset #2). However, there are also a few SKUs with high
18 inter-demand intervals (p going up to 12 in dataset #1, 24 in dataset #2 and 15 in dataset
19 #3) for which the Normal distribution provides a strong fit. Those latter SKUs have a
20 minimum CV^2 (i.e. $CV^2 = 0$) which can be explained by the fact that their demand is very
21 low (in most of the cases, the demand is equal to one) and can fit to the Normal
22 distribution with low mean (i.e. equivalently high values of p) and variance. Finally, in
23 addition to the SKUs where there is a fit to the Normal distribution (those with low
24 values of p), the Gamma distribution provides also a strong fit to the SKUs with very
25 high values of p (i.e. SKUs with an inter-demand interval going up to 12 periods in
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

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.

5. Empirical results: lead-time demand

In this section, we are concerned with the goodness-of-fit of the various distributions considered in our work to lead time demand data. The empirical utility of the relevant demand distributional assumptions is also assessed by means of experimentation with a periodic order-up-to level system. At this point it is important to note that no lead time information has been made available to us for dataset #1 and consequently this has not been considered for this part of our research. All SKUs included in dataset #3 are associated with a fixed lead time of three periods (months).

INSERT TABLE 6 ABOUT HERE

The goodness-of-fit results are presented in Table 6. Although the separation among the tested distributions is not as clear as in the case of demand per period, the results indicate that NBD performs, overall, best with about 62% (weighted average) of 8,055 SKUs showing a strong fit. Stuttering Poisson comes in a close second. Despite the general expectation that the Normal distribution would be a good representative of lead-time demand due to the effect of Central Limit Theorem, our results do not demonstrate

1
2
3 dominance by the Normal distribution. On the contrary, the Normal distribution seems to
4 have a split between strong fit and no fit; only about 47% of the 8,055 SKUs showing a
5 strong fit for Normal, and about 43% showing no fit. Although the Gamma distribution
6 still compares unfavourably to the compound distributions, it does perform better than the
7 demand per period case. On the contrary, the Poisson distribution is associated with an
8 inferior performance when compared to the demand per period case; this was
9 theoretically expected since lead-times are overall rather lengthy resulting in high lead-
10 time demand cumulative sizes which cannot be represented by the Poisson distribution.
11 In the next sub-section, we design a simulation study to test the utility of these results.
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26

27 **5.1 Linking the goodness-of-fit to stock control performance**

28 In this sub-section we are concerned with the empirical utility of the demand
29 distributional assumptions discussed thus far in our paper. (The Poisson distribution will
30 not be considered in this part of the analysis as preliminary results indicated a very poor
31 performance as compared to that of all other distributions.) The implications of using
32 these distributions are evaluated through an order-up-to-level (R, S) policy that operates
33 under a Cycle Service Level (CSL) objective (proportion of demand satisfied directly
34 from stock). Mean demand is estimated by using the SBA estimator and the variance of
35 demand is estimated through the smoothed MSE procedure. By keeping all aspects of the
36 stock control system fixed and varying only the hypothesized demand distribution, the
37 effects of using such distributions may be evaluated and contrasted to each other.
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54

55 Three target CSL values have been considered: 90%, 95% and 99%. Performance is
56 evaluated by means of the average holding volumes and the average backlog volumes
57
58
59
60

(per time period for each SKU) as well as the achieved CSLs. In Tables 7 and 8, we show, for datasets #2 and #3 respectively, for every demand distribution and for every target CSL, the relevant quantities, averaged across all SKUs.

INSERT TABLE 7 ABOUT HERE

INSERT TABLE 8 ABOUT HERE

Overall, the results in Tables 7 and 8 indicate that the Normal distribution is the worst performing one for higher target CSLs (i.e. CSL = 95% and 99%). However, for a target CSL equal to 90% the Normal compares favourably to the other distributions. The Normal distribution may not approximate well the tail of the empirical distribution of data, but may perform very satisfactorily otherwise. For high target CSLs, the NBD provides the best results, followed by the Gamma distribution and Stuttering Poisson.

The results presented in the above tables are somewhat difficult to interpret in terms of deriving tangible conclusions on the comparative performance of the various distributions. Higher holding volumes are associated with higher achieved CSLs, making it difficult to appreciate the comparative merits of using one distribution instead of another. To that end, we have also considered the presentation of our results through the use of efficiency curves. In Figures 2 and 3 we plot (for datasets #2 and #3 respectively) the holding volumes required to support a certain achieved CSL for each of the three target values and each of the distributions considered in our investigation. In these figure, for a fixed holding volume, a curve that is further to the *x-axis* shows that the CSL is higher and indicates that the corresponding demand distribution is more efficient.

INSERT FIGURE 2 ABOUT HERE

INSERT FIGURE 3 ABOUT HERE

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

The efficiency curves show that the NBD and Gamma are the more efficient distributions, followed by the Stuttering Poisson. In fact, for a fixed holding volume, the former two distributions provide the highest CSL. Although the Stuttering Poisson provides a good overall fit (as demonstrated in the beginning of this section) the stock control performance of this distribution is somewhat inferior to that associated with Gamma (that performed less satisfactorily in terms of its goodness-of-fit results). This reinforces potentially earlier comments that we have made with respect to the importance of fitting the tail of the distribution as opposed to the entire distribution of lead time demand. In addition, the results demonstrate that the Normal distribution is the least efficient one for targets above 90% and this is in accordance with theoretical expectations.

In summary, our results for the lead-time demand distribution tests are overall different from the findings emerged from the demand-per-period analysis. The Poisson distribution was found to perform considerably better in the latter case whereas the opposite is true for the Gamma distribution. The Normal distribution proved to be a weak candidate for the demand per period distribution. It was found to perform better for lead-time demand (probably due to Central Limit Theorem effects), although it may still be judged as far from appropriate for implementation in a spare parts context. The results have also indicated that the goodness-of-fit performance on lead time demand may not translate to similar insights with regards to the stock control implications of using the various

1
2
3 distributions. Most noticeably, the Gamma distribution was found to compare favourably
4
5 to the Stuttering Poisson with regards to stock control whereas the goodness-of-fit tests
6
7 point to the opposite direction.
8
9

10 11 12 **6. Conclusions, implications and further research** 13

14
15 Practical parametric approaches to inventory management rely upon an explicit demand
16
17 distributional assumption. Although the Normal distribution is typically adequate for
18
19 'fast' demand items this is not true for spare parts that are typically characterised by
20
21 intermittent demand structures. Although a number of distributions have been discussed
22
23 in the academic literature for representing such patterns, evidence on the empirical
24
25 validity of the relevant assumptions (i.e. empirical goodness-of-fit) is lacking. Moreover,
26
27 the linkage between the validity and the utility of these distributional assumptions, the
28
29 latter referring to the actual stock control implications of using a particular distribution,
30
31 has not been explored in the literature. An exception to that is perhaps a PhD thesis
32
33 conducted at Lancaster University in the early 1990's (Kwan, 1991). Unfortunately, no
34
35 parts of that thesis have ever been published though in the academic literature. Finally an
36
37 important distinction to be made is that between the performance of a hypothesised
38
39 distribution when fitting demand per period data and that when applied to lead-time
40
41 demands, the latter being what matters in a real stock control system.
42
43
44
45
46
47
48

49
50 In this paper a number of distributions are first selected as potential candidates for
51
52 representing intermittent demands, on the basis of: i) theoretical arguments, ii) intuitive
53
54 appeal; iii) empirical support. These distributions are: the Poisson, Negative Binomial
55
56 Distribution (NBD), Stuttering Poisson, Normal and Gamma. The empirical goodness-of-
57
58
59
60

1
2
3 fit of these distributions is then assessed for both demand per period data and lead-time
4 demands by employing the Kolmogorov-Smirnov (K-S) test. Conclusions are being
5 drawn on the empirical validity of the various distributional assumptions followed by an
6 assessment of their stock control implications (i.e. their empirical utility), under a
7 periodic order-up-to-level system. The empirical databases available for the purposes of
8 our investigation come from the US DLA, RAF and Electronics Industry and they consist
9 of the individual monthly demand histories of 4,588, 5,000 and 3,055 SKUs respectively.
10
11
12
13
14
15
16
17
18
19

20
21 When we consider demand per period data, the results indicate that both the NBD and
22 Stuttering Poisson provide the most frequent fit. Both these distributions are compound in
23 nature, meaning that they account explicitly for a demand arrival process (Poisson) and a
24 different distribution for the transaction sizes (Log series and Geometric for the NBD and
25 Stuttering Poisson respectively). Despite previous claims, the Gamma distribution does
26 not perform very well and the same is true for the Normal distribution. This may be
27 attributed to the continuous nature of these distributions (since their fit is tested on
28 discrete observations) but also to the fact that demand per unit time period is considered
29 as opposed to lead time demand. The Poisson distribution provides a 'reasonable' fit and
30 this is theoretically expected for slow moving items.
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46

47 We have also attempted to link the goodness-of-fit of the various distributions to two key
48 characteristics that have been shown to be collectively sufficient for defining intermittent
49 demands (Syntetos *et al.*, 2005). These characteristics are the inter-demand interval and
50 the squared coefficient of variation of the demand sizes. Goodness-of-fit performance
51
52
53
54
55
56
57
58
59
60

1
2
3 was then analysed across all series as a function of those characteristics and some
4
5 interesting insights were generated into the comparative performance of the distributions.
6
7

8
9
10 Recognizing that the lead-time demand distribution is more important from an inventory
11
12 control perspective, we also tested the goodness-of-fit of the above mentioned
13
14 distributions on lead-time demand. NBD and Stuttering Poisson were found to perform
15
16 best also in this case; however the remaining distributions were associated with
17
18 differences in performance. The Poisson distribution was found to perform considerably
19
20 worse in that case and this may be explained in terms of the relatively long lead times
21
22 associated with our empirical dataset. The Gamma distribution though was found to
23
24 perform better for lead-time demand; similarly, the Normal distribution proved also to
25
26 perform better (Central Limit Theorem effects), although it may still be judged as far
27
28 from appropriate for implementation in a spare parts context.
29
30
31
32

33
34
35 Subsequently, simulation experiments using already proven forecasting (Syntetos-Boylan
36
37 Approximation) and inventory control methods ((R, S) type periodic review policy) were
38
39 used to assess the stock control implications of using the various distributions. Although
40
41 NBD was indeed confirmed as the best performing distribution, the results have also
42
43 indicated that the goodness-of-fit performance on lead time demand may not translate to
44
45 similar insights with regards to the stock control implications of using a particular
46
47 distribution. Most noticeably, the Gamma distribution was found to compare favourably
48
49 to the Stuttering Poisson with regards to stock control whereas the goodness-of-fit tests
50
51 pointed to the opposite direction.
52
53
54
55
56
57
58
59
60

1
2
3 In summary, the results of our empirical investigation suggest that the NBD performs
4 best in an inventory context. This is followed by the Gamma and Stuttering Poisson
5 distribution, the former found to outperform the latter in terms of their stock control
6 implications. We expect further academic studies to take note of our findings and build
7 spare parts management research on these distributions. We also expect these findings to
8 be of interest to practitioners dealing with spare parts and being interested in improving
9 automated procedures towards the most effective management of such items. At this
10 point we should also note that although spare parts constituted the focus of our research,
11 natural extensions with regards to our findings can be made to all products characterised
12 by intermittent demand structures.
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

29 Our research has demonstrated that the choice of a demand distribution has considerable
30 effects on stock control performance in a spare parts context. In addition, the work
31 presented in this paper has revealed a number of interesting themes for further research.
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

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 (e.g. Gamma) 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 paper in larger demand datasets coupled with the assessment of the empirical utility of various distributional assumptions under other stock control formulations should help advance knowledge in this area. Also, it would be worthwhile exploring the sensitivity of our findings under the realistic scenario of variable lead times (see, for example, Babai *et al.*, 2009 – and for continuity purposes also Babai and Dallery, 2009).
- Empirical studies are needed to determine the comparative merits of non-parametric approaches (such as bootstrapping) in a spare parts context. Such approaches do not rely upon any underlying distributional assumption. They rather consider the empirical distribution function and emphasise the direct estimation of the relevant percentiles (say the 95th or the 99th percentile). Although this is a very appealing property, the current empirical knowledge base is not extended enough to suggest how such approaches compare to parametric inventory management in real terms.
- In addition, and given that the inverse of the distribution function is the quantile function, another relevant approach to the issues discussed in this paper relates to *quantile forecasting*. Directly forecasting quantiles avoids the need for assumptions regarding the shape and spread of the underlying distribution (Taylor, 2007) and should be given more attention in the inventory literature.

- Finally, assessing Value at Risk (VaR) in inventory management is arguably also an area that has not been given sufficient attention. VaR is increasingly being used in financial management as a natural measure of the risk taken with a given position. In the framework of inventory management it can work as well (Tapiero, 2005) and is particularly relevant since VaR problems involve only one tail of the distribution (of the financial returns). Extensions into the distribution of inventory costs are feasible and intuitively appealing (Luciano *et al.*, 2003) and further work into the issue of the efficient computation of VaR would constitute an important addition to the current state of knowledge in the area of inventory control.

References

- Altay N, Rudisill F, Litteral LA, (2008) Adapting Wright's modification of Holt's method to forecasting intermittent demand. *International Journal of Production Economics* 111: 389-408.
- Babai MZ, Dallery Y, (2009) Dynamic versus static control policies in single stage production-inventory systems. *International Journal of Production Research* 47: 415-433.
- Babai MZ, Syntetos AA, Dallery Y, Nikolopoulos K, (2009) Periodic inventory control systems with non-stationary demand and lead time uncertainty: analysis and empirical investigation. *International Journal of Production Research* 47: 2461-2483.
- Babai MZ, Syntetos AA, Teunter R, (2010) On the empirical performance of (T, s, S) heuristics. *European Journal of Operational Research* 202: 466-472.
- 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: International Journal of Applied Forecasting* 4: 39-42.
- Boylan JE, Syntetos AA, Karakostas GC, (2007) Classification for forecasting and stock Control: a case study. *Journal of the Operational Research Society* 59: 473-481.
- Burgin TA, (1975) The gamma distribution and inventory control. *Operational Research Quarterly* 26: 507-525.
- Burgin TA, Wild AR, (1967) Stock control experience and usable theory. *Operational Research Quarterly* 18: 35-52.

1
2
3 Chatfield DC, Hayya JC, (2007) All-zero forecasts for lumpy demand: a factorial study.
4 International Journal of Production Research 45: 935-950.
5

6 Croston JD, (1972) Forecasting and stock control for intermittent demands. Operational
7 Research Quarterly 23: 289-304.
8

9 Croston JD, (1974) Stock levels for slow-moving items. Operational Research Quarterly 25:
10 123-130.
11

12 Dunsmuir WTM, Snyder RD, (1989) Control of inventories with intermittent demand.
13 European Journal of Operational Research 40: 16-21.
14

15 Eaves A, (2002) The Forecasting for the ordering and stock holding of consumable spare
16 parts. Unpublished PhD thesis, Lancaster University, UK.
17

18 Ehrhardt R, Mosier C, (1984) A revision of the power approximation for computing (s, S)
19 inventory policies. Management Science 30: 618-622.
20

21 Friend JK, (1960) Stock control with random opportunities for replenishment. Operational
22 Research Quarterly 11: 130-136.
23

24 Gallagher DJ, (1969) Two periodic review inventory models with backorders and stuttering
25 Poisson demands. AIIE Transactions 1: 164-171.
26

27 Gutierrez RS, Solis AO, Mukhopadhyay S, (2008), Lumpy demand forecasting using neural networks.
28 International Journal of Production Economics 111: 409-420.
29

30 Harnett DL, Soni AK, (1991) Statistical methods for business and economics (4th edition).
31 Addison Wesley, New York.
32

33 Hollier, R.H. (1980) The distribution of spare parts. International Journal of Production
34 Research 18: 665-675.
35

36 Janssen FBSLP, (1998) Inventory management systems; control and information issues.
37 Published PhD thesis, Centre for Economic Research, Tilburg University, The Netherlands.
38

39 Johnston FR, (1980) An interactive stock control system with a strategic management role.
40 Journal of the Operational Research Society 31: 1069-1084.
41

42 Johnston FR, Boylan JE, (1996) Forecasting for items with intermittent demand. Journal of
43 the Operational Research Society 47: 113-121.
44

45 Johnston FR, Boylan JE, Shale EA, (2003) An examination of the size of orders from
46 customers, their characterization and the implications for inventory control of slow moving
47 items. Journal of the Operational Research Society 54: 833-837.
48

49 Kwan HW, (1991) On the demand distributions of slow moving items. Unpublished PhD
50 thesis, Lancaster University, UK.
51

52 Luciano E, Peccati L, Cifarelli DM, (2003) VaR as a risk measure for multiperiod static
53 inventory models. International Journal of Production Economics 81-82: 375-384.
54

55 Mitchell GH, (1962) Problems of controlling slow-moving engineering spares. Operational
56 Research Quarterly 13: 23-39.
57

58 Naddor E, (1975) Optimal and heuristic decisions in single and multi-item inventory systems.
59 Management Science 21: 1234-1249.
60

1
2
3 Noether GE, (1963) Note on the Kolmogorov statistic in the discrete case. *Metrika* 7: 115-
4 116.

5
6 Noether GE, (1967) *Elements of nonparametric statistics*. John Wiley & Sons, Inc., New
7 York.

8
9 Porras EM, Dekker R, (2008) An inventory control system for spare parts at a refinery: an
10 empirical comparison of different reorder point methods. *European Journal of Operational*
11 *Research* 184: 101-132.

12
13 Quenouille MH, (1949) A relation between the logarithmic, Poisson and negative binomial
14 series. *Biometrics* 5: 162-164.

15
16 Quintana R, Leung MT, (2007) Adaptive exponential smoothing versus conventional
17 approaches for lumpy demand forecasting: case of production planning for a manufacturing
18 line. *International Journal of Production Research* 45: 4937-4957.

19
20 Ritchie E, Kingsman BG, (1985) Setting stock levels for wholesaling: Performance measures
21 and conflict of objectives between supplier and stockist. *European Journal of Operational*
22 *Research* 20: 17-24.

23
24 Sani B, (1995) *Periodic inventory control systems and demand forecasting methods for low*
25 *demand items*. Unpublished PhD thesis, Lancaster University, UK.

26
27 Silver EA, Pyke DF, Peterson R, (1998) *Inventory management and production planning and*
28 *scheduling* (3rd edition). John Wiley & Sons, Inc., New York.

29
30 Slakter MJ, (1965) A comparison of the Pearson chi-square and Kolmogorov goodness-of-fit
31 tests with respect to validity. *Journal of the American Statistical Association* 60: 854-858.

32
33 Strijbosch LWG, Heuts RMJ, van der Schoot EHM, (2000) A combined forecast-inventory
34 control procedure for spare parts. *Journal of the Operational Research Society* 51: 1184-1192.

35
36 Syntetos AA, Babai MZ, Dallery Y, Teunter R, (2009a) Periodic control of intermittent
37 demand items: theory and empirical analysis. *Journal of the Operational Research Society* 60:
38 611-618.

39
40 Syntetos AA, Boylan JE, (2005) The accuracy of intermittent demand estimates. *International*
41 *Journal of Forecasting* 21: 303-314.

42
43 Syntetos AA, Boylan JE, (2008) Smoothing and adjustments of demand forecasts for inventory
44 control. *IMA Journal of Management Mathematics* 19: 175-192.

45
46 Syntetos AA, Boylan JE, Croston JD, (2005) On the categorisation of demand patterns.
47 *Journal of the Operational Research Society* 56: 495-503.

48
49 Syntetos AA, Boylan JE, Disney SM (2009b) Forecasting for inventory planning: a 50-year
50 review. *Journal of the Operational Research Society* 60: S149-S160.

51
52 Syntetos AA, Keyes M, Babai MZ, (2009c) Demand categorisation in a European spare parts
53 logistics network. *International Journal of Operations and Production Management* 29: 292-
54 316.

55
56 Tapiero C, (2005) Value at risk and inventory control. *European Journal of Operational Research*
57 *163: 769-775*.

58
59 Taylor JW, (2007) Forecasting daily supermarket sales using exponentially weighted quantile
60 regression. *European Journal of Operational Research* 178: 154-167.

1
2
3 Teunter R, Syntetos AA, Babai MZ, (2010) Determining order-up-to levels under periodic review
4 for compound binomial (intermittent) demand. *European Journal of Operational Research* 203:
5 619-624.
6

7 Vereecke A, Verstraeten P, (1994) An inventory management model for an inventory
8 consisting of lumpy items, slow movers and fast movers. *International Journal of Production*
9 *Economics* 35: 379-389.
10

11 Walsh JE, (1963) Bounded probability properties of Kolmogorov-Smirnov and similar
12 statistics for discrete data. *Annals of the Institute of Statistical Mathematics* 15: 153-158.
13

14 Ward JB, (1978) Determining re-order points when demand is lumpy. *Management Science*
15 24: 623-632.
16

17 Watson RB, (1987) The effects of demand-forecast fluctuations on customer service and
18 inventory cost when demand is lumpy. *Journal of the Operational Research Society* 38:75-82.
19

20 Willemain TR, Smart CN, Shockor JH, DeSautels PA, (1994) Forecasting intermittent
21 demand in manufacturing: a comparative evaluation of Croston's method. *International*
22 *Journal of Forecasting* 10: 529-538.
23

24 Willemain TR, Smart CN, Schwarz HF, (2004) A new approach to forecasting intermittent
25 demand for service parts inventories. *International Journal of Forecasting* 20: 375-387.
26

27 Williams TM, (1984) Stock control with sporadic and slow-moving demand. *Journal of the*
28 *Operational Research Society* 35, 939-948.
29

30 **Appendix A. Demand per period goodness-of-fit results (datasets #2, #3)**

31

32
33 -----
34 **INSERT FIGURE A1 ABOUT HERE**
35 -----
36

37
38 -----
39 **INSERT FIGURE A2 ABOUT HERE**
40 -----
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

List of tables

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

Table 1. Empirical datasets

4,588 SKUs	Demand Intervals		Demand Sizes		Demand per period	
	Mean	St. Deviation	Mean	St. Deviation	Mean	St. Deviation
Min	1.000	0.000	1.000	0.000	0.083	0.279
25%ile	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%ile	5.600	6.049	11.940	12.435	4.550	9.403
Max	12.000	24.597	1326.875	1472.749	783.917	1219.012

Table 2. Dataset #1 – US Defense Logistics Agency (DLA)

5,000 SKUs	Demand Intervals		Demand Sizes		Demand per period	
	Mean	St. Deviation	Mean	St. Deviation	Mean	St. Deviation
Min	3.824	0.000	1.000	0.000	0.036	0.187
25%ile	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%ile	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

Dataset #2 has been used in the studies conducted by Syntetos *et al.* (2009a) and Teunter *et al.* (2010)

Table 3. Dataset #2 – Royal Air Force (RAF)

3,055 SKUs	Demand Intervals		Demand Sizes		Demand per period	
	Mean	St. Deviation	Mean	St. Deviation	Mean	St. Deviation
Min	1.000	0.000	1.000	0.000	0.042	0.245
25%ile	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%ile	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 study conducted by Babai *et al.* (2010).

Table 4. Dataset #3 – Electronics

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

Table 5. Goodness-of-fit results: demand per period

Dataset #	No of SKUs	Distribution	Percentage of SKUs (%)		
			Strong Fit	Good Fit	No Fit
2	5,000	Poisson	49.08	4.88	46.04
		NBD	68.36	5.00	26.64
		Stuttering Poisson	68.10	6.24	25.66
		Normal	53.58	9.24	37.18
		Gamma	57.42	9.98	32.60
3	3,055	Poisson	23.24	6.91	69.85
		NBD	51.29	6.91	41.80
		Stuttering Poisson	50.25	6.38	43.37
		Normal	35.16	10.70	54.11
		Gamma	34.70	12.93	52.34

Table 6. Goodness-of-fit results: lead-time demand

Distribution		<i>Holding (volumes)</i>	<i>Backlog (volumes)</i>	<i>CSL (%)</i>
Normal	90%	30.277	0.383	85.9
	95%	36.298	0.344	87.4
	99%	47.798	0.293	89.7
Gamma	90%	23.208	0.461	84.1
	95%	33.644	0.365	87.4
	99%	72.270	0.250	91.1
Stuttering Poisson	90%	23.675	0.460	84.0
	95%	34.224	0.364	87.2
	99%	74.251	0.249	91.0
NBD	90%	23.102	0.464	84.7
	95%	33.582	0.365	88.8
	99%	72.208	0.248	92.9

Table 7. Stock control performance – Dataset #2

Distribution		<i>Holding (volumes)</i>	<i>Backlog (volumes)</i>	<i>CSL (%)</i>
Normal	90%	76.065	3.829	89.3
	95%	92.700	3.191	91.1
	99%	125.302	2.386	93.5
Gamma	90%	64.595	3.873	88.4
	95%	87.881	3.001	91.5
	99%	161.189	2.017	95.2
Stuttering Poisson	90%	66.253	3.872	88.4
	95%	89.094	2.997	91.4
	99%	161.253	2.012	95.1
NBD	90%	64.241	3.891	88.1
	95%	87.079	3.017	91.5
	99%	159.268	2.042	95.2

Table 8. Stock control performance – Dataset #3

List of figures

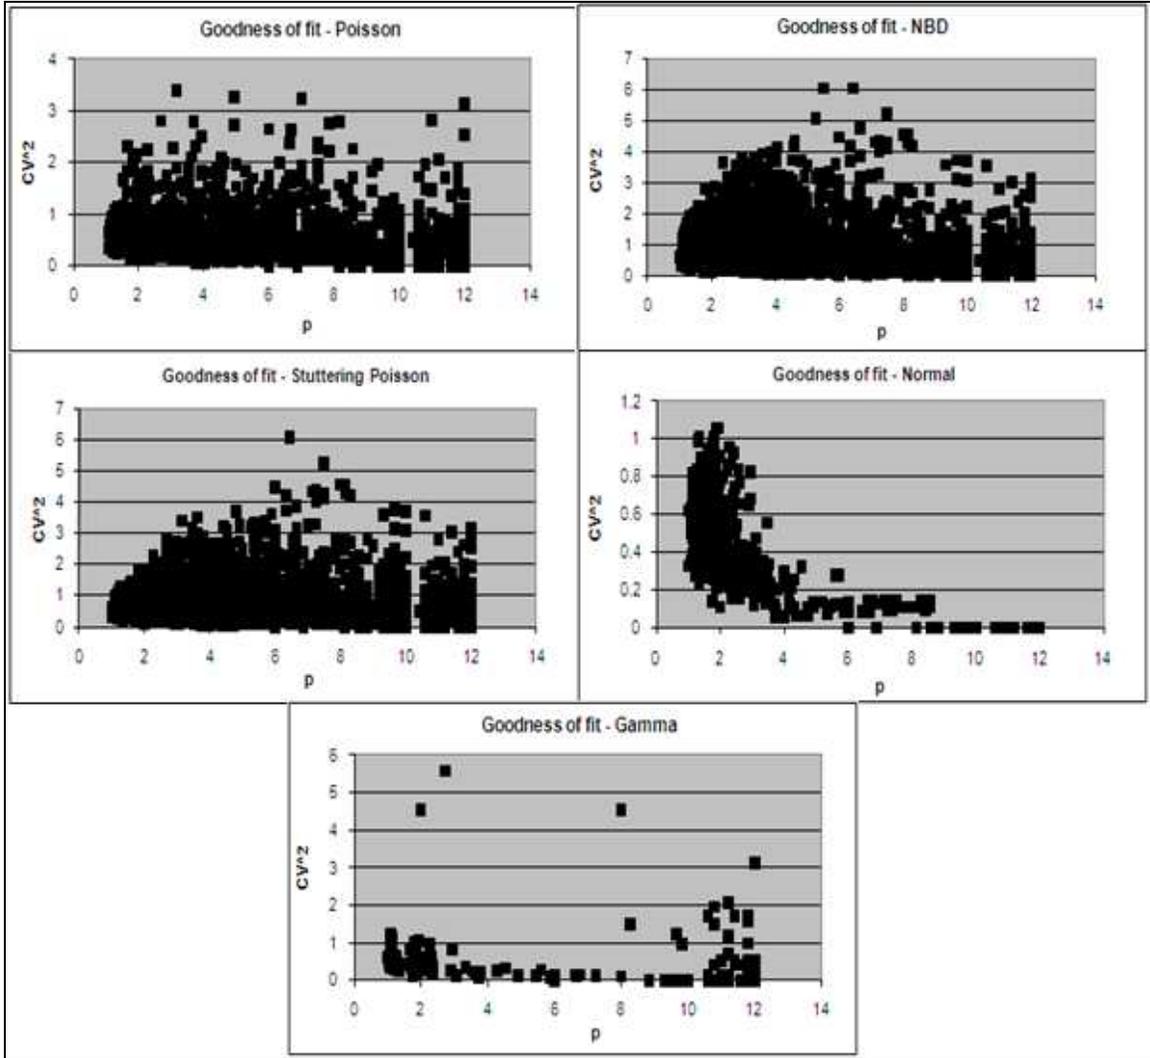


Figure 1. Dataset #1 - Goodness-of-fit results

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

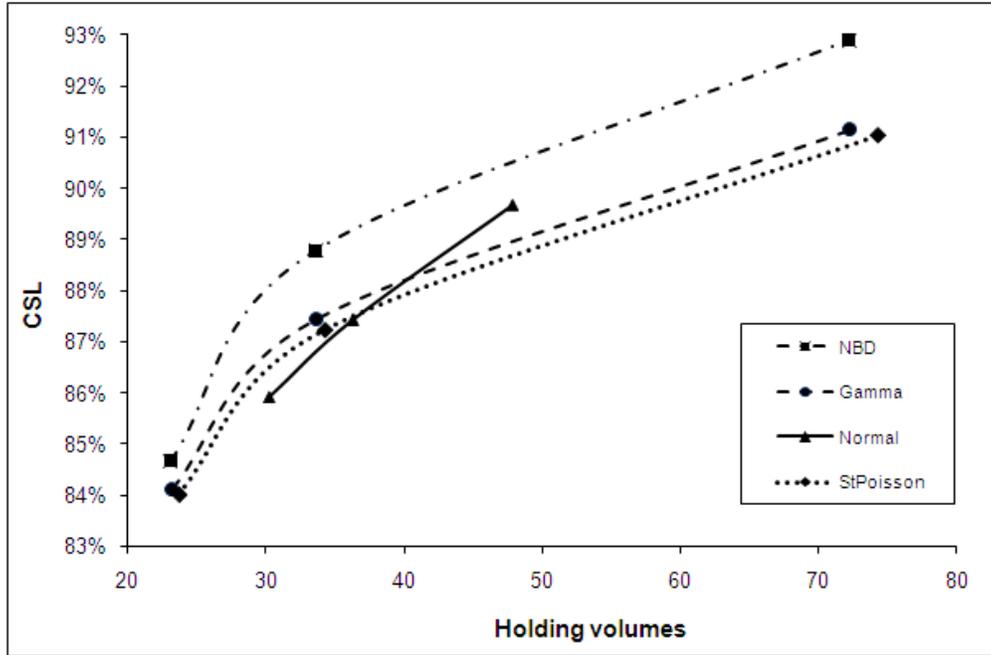


Figure 2. Efficiency curves: CSL against Holding Volume (Dataset #2)

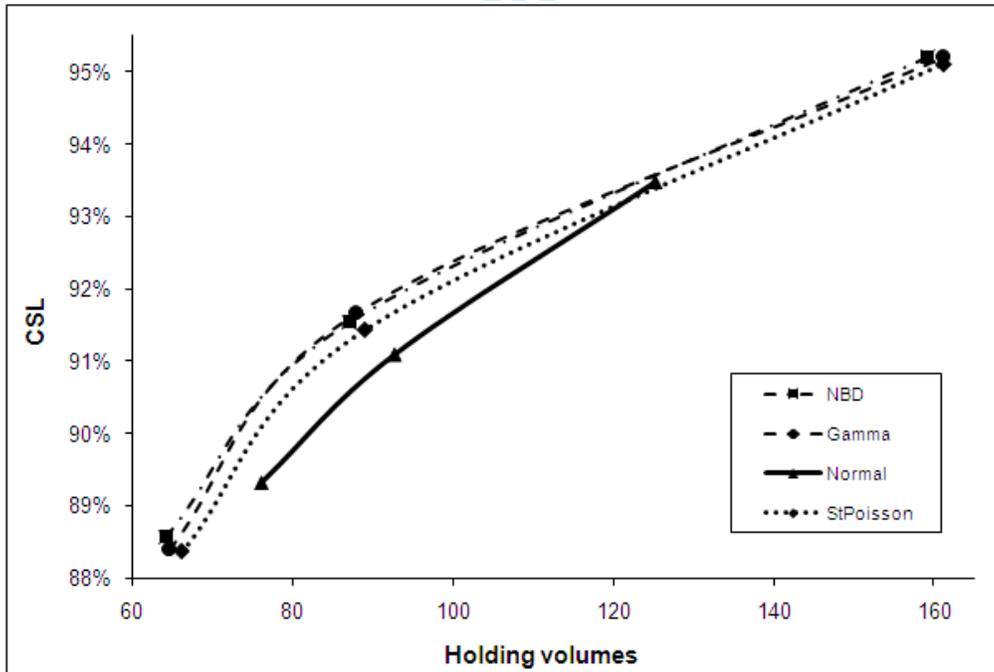


Figure 3. Efficiency curves: CSL against Holding Volume (Dataset #3)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

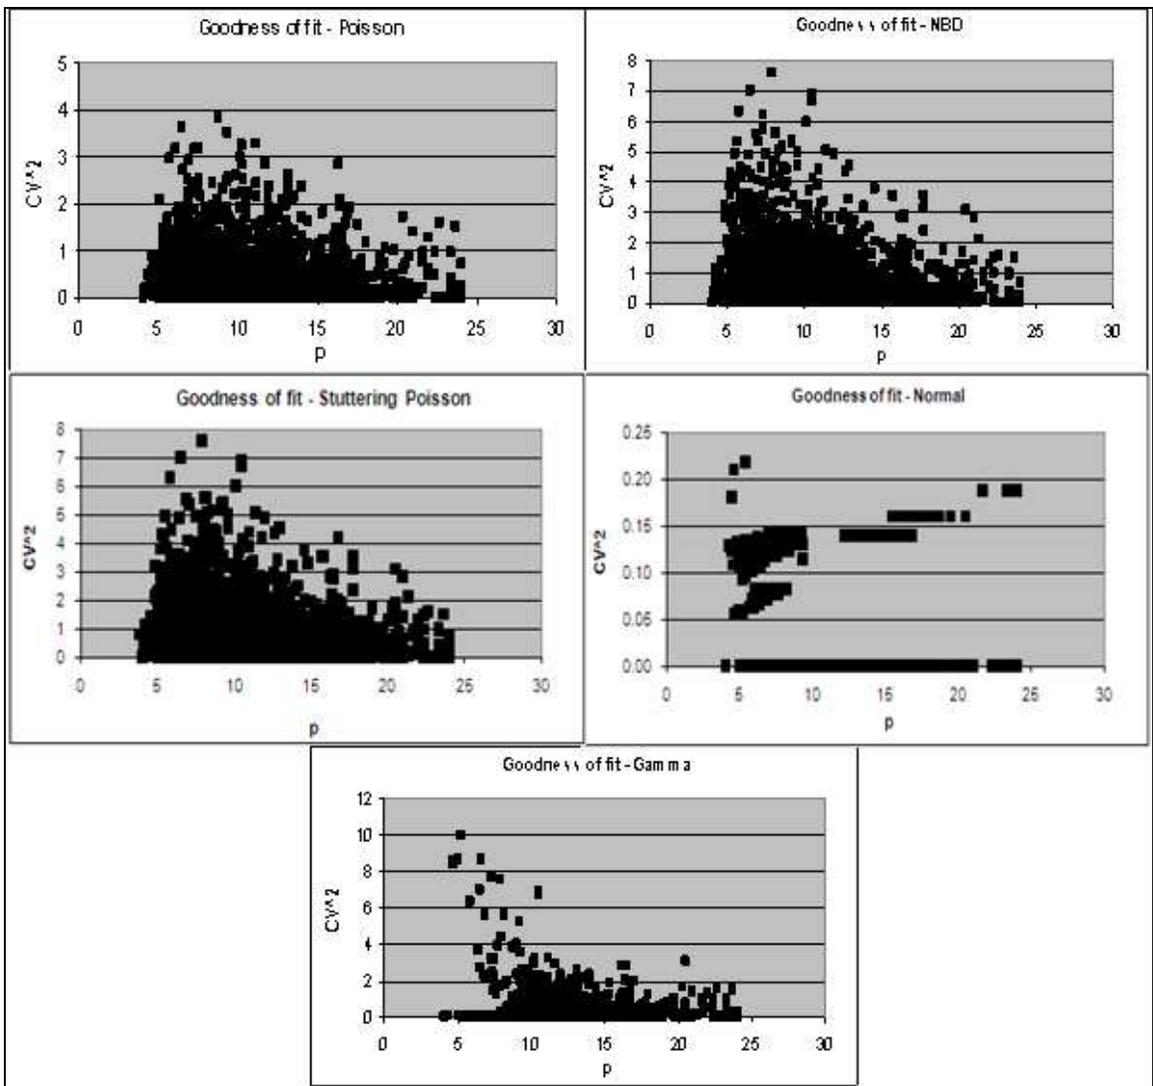


Figure A1. Dataset #2 - Goodness-of-fit results

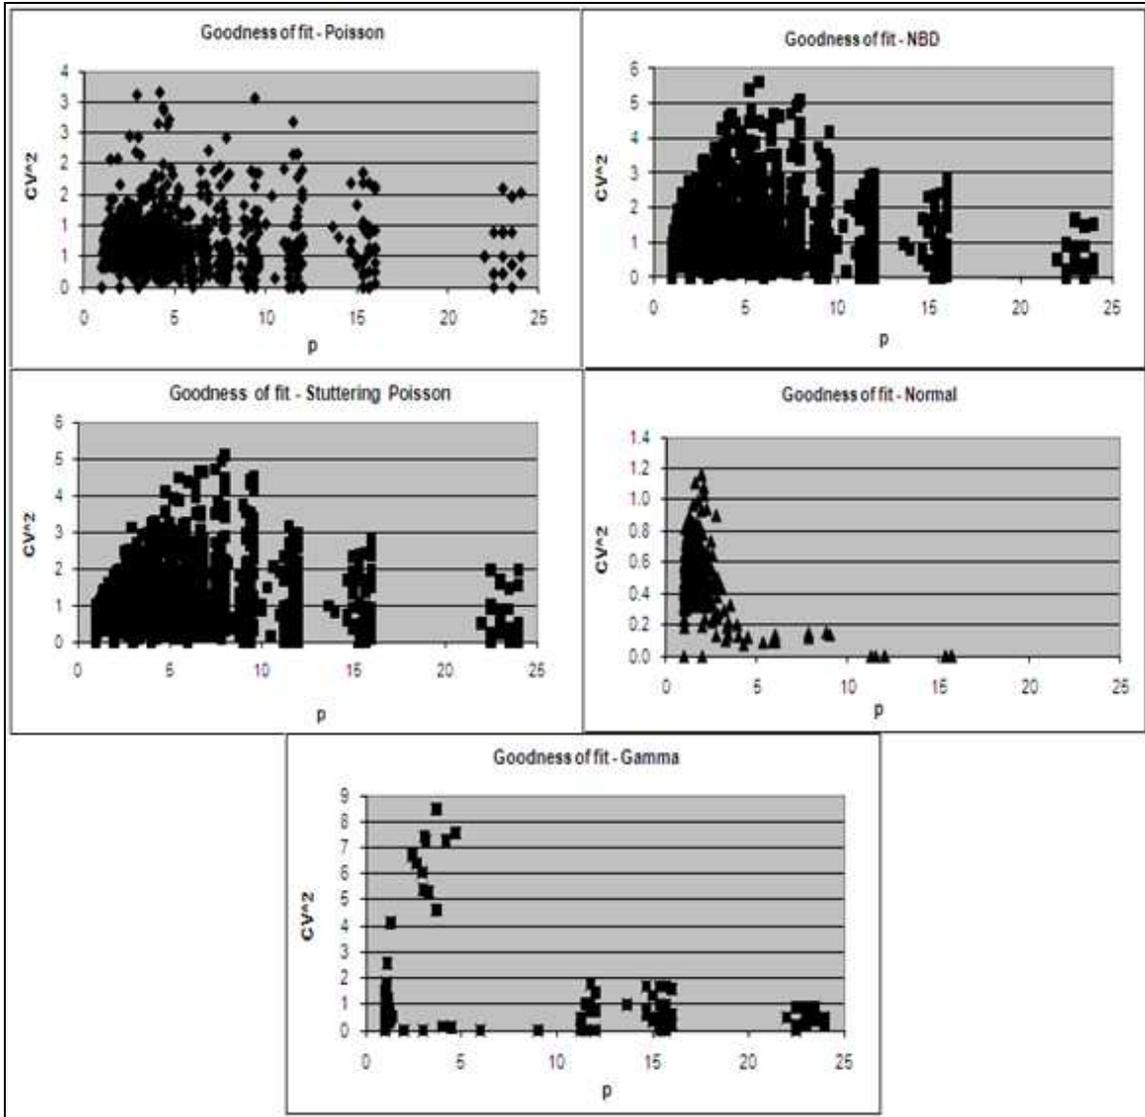


Figure A2. Dataset #3 - Goodness-of-fit results

Only