

A compound-Poisson Bayesian approach for spare parts inventory forecasting

Babai, M.Z.¹; Chen, H.²; Syntetos, A.A.³; Lengu, D.³

¹Kedge Business School (France), ²University of Portsmouth (UK), ³Cardiff University (UK)

Abstract

Spare parts are often associated with intermittent demand patterns that render their forecasting a challenging task. Forecasting of spare parts demand has been researched through both parametric and non-parametric approaches. However, little has been contributed in this area from a Bayesian perspective, and most of such research is built around the Poisson demand distributional assumption. However, the Poisson distribution is known to have certain limitations and, further, empirical evidence on the inventory performance of Bayesian methods is lacking. In this paper, we propose a new Bayesian method based on compound Poisson distributions. The proposed method is compared to the Poisson-based Bayesian method with a Gamma prior distribution as well as to a parametric frequentist method and to a non-parametric one. A numerical investigation (on 7,400 theoretically generated series) is complemented by an empirical assessment on demand data from about 3,000 stock keeping units in the automotive sector to analyse the performance of the four forecasting methods. We find that both Bayesian methods outperform the other methods with a higher inventory efficiency reported for the Poisson Bayesian method with a Gamma prior. This outperformance increases for higher demand variability. From a practical perspective, the outperformance of the proposed method is associated with some added complexity. We also find that the performance of the non-parametric method improves for longer lead-times and higher demand variability when compared to the parametric one.

Keywords: forecasting, inventory, intermittent demand, Bayesian method, empirical investigation.

1. Introduction

Demand forecasting and inventory control of spare parts, or any items associated with intermittent demand, are challenging tasks for inventory managers. This is mainly due to the compound nature of such demand patterns (demand occurrences are interspersed by intervals of no demand at all), which are often associated in practice with a lack of historical data to enable reliable estimates for the parameters of the assumed demand distribution (Syntetos et al., 2016; Babai et al., 2020; Ruiz et al., 2020). Different approaches have been proposed in the literature to address these issues, originating in the frequentist, Bayesian, and non-parametric domains. Parametric approaches, both frequentist and Bayesian, assume the lead-time demand distribution to be known; the former further assume that the parameters of that distribution are unknown (yet not subject to change over time) and need to be forecasted, and the latter regard the data as known, having been observed from the realised sample, and the parameters unknown being probabilistically described through the prior distribution. Under the non-parametric approach, no particular distribution is assumed for the lead-time demand (though the distribution is assumed to remain the same over time) and bootstrapping is the most commonly used method, which consists of building an empirical distribution based on sampling from past demand data. There is a large body of literature that deals with parametric and non-parametric forecasting approaches, extending as far back as the 1970s. The reader is referred to Syntetos et al. (2016) and Hasni et al. (2019a) for an overview on these approaches.

The focus of this paper is the Bayesian approach, which uses the Bayes Theorem to update a prior distribution (probabilities specified prior to data collection) into a posterior distribution (the probabilities following data analysis), by incorporating information (likelihoods) provided by the observed data. The Bayesian approach, with its own statistical theory, has been applied successfully in many areas and forecasting contexts. The ability to combine all information and

sources of uncertainty, and revise and update it as more data are acquired is particularly promising and appealing when (1) data are scarce and (2) there are considerable changes in the data.

In the spare parts forecasting context, the Bayesian approach is often used and justified by the fact that most of the parts do not exhibit a constant failure rate (Dekker et al., 2013; Boutselis and McNaught, 2019). In fact, in practice the usage context is unlikely to stay the same throughout the lifecycle of a system, which renders the failure rate of a part changing over time. Moreover, the failure rate is often unknown either due to the absence of operational data at the time of initial provisioning, or due to the lack of data when changes in environment occur. Under such changes and uncertain conditions, parametric frequentist and non-parametric forecasting approaches are not adequate due to the assumption of constant demand distribution parameters over time. It is also worth pointing out that the Bayesian approach provides practitioners with an opportunity to incorporate intuition and previous experience in a quantifiable form by selecting an appropriate choice of the prior distribution. Practitioners can select a prior with a value for the variance that reflects their perceived level of uncertainty about the distribution of demand. As demand evolves, the observed data are used to update the likelihood function and, as a result, the posterior distribution of demand.

There are a number of researchers that have used a Bayesian approach to forecast lead-time demand (e.g., Silver, 1965; Brown Jr and Rogers, 1973; De Wit, 1983; Azoury and Miller, 1984; Azoury, 1985; Karmarkar, 1994; Popovic, 1987; Dolgui and Pashkevich, 2008; Yelland, 2010). Within this research stream, Karmarkar (1994) developed a heuristic model that can be used to estimate the location of the percentiles directly in a Bayesian approach. The heuristic model thus tries to approximate upper-tail probabilities with a simple functional form while ignoring the general shape of the demand distribution. By estimating the location of the

percentiles directly, this approach seeks to avoid the pitfall of fitting a model using information about the centre of the distribution and then making inferences about the upper-tail probabilities. Karmarkar (1994) showed the better performance of the heuristic in achieving a target service level when compared to the case of Normal distribution with smoothed Mean Absolute Deviation (MAD) for estimating the demand (forecast error) variance.

Aronis et al. (2004) proposed a Bayesian model based on the work of Popovic (1987) in which demand has a Poisson distribution and the prior conjugate distribution is Gamma. This leads to a posterior distribution equivalent to the negative binomial distribution (NBD). Although a Poisson distribution, with only one parameter, is a limited representation of slow-moving demand, this is compensated by the choice of a Gamma prior with two parameters to incorporate more information, and the resulting posterior distribution of NBD is much better supported with empirical evidence (Syntetos and Boylan, 2006; Syntetos et al., 2013). From a theoretical perspective, the proposed Bayesian method has the advantage of offering a closed form expression of the lead-time demand distribution that can be used easily to calculate a base-stock policy's parameter. However, to the best of our knowledge the empirical performance of this Poisson-based Bayesian method has never been evaluated. Filling this gap constitutes one of the objectives of our paper. Dolgui and Pashkevich (2008) have proposed a Bayesian method using the Beta distribution as the prior and the Binomial distribution as the likelihood function, thus obtaining the beta-binomial distribution as the posterior distribution. However, this approach presents some challenges with respect to the estimation of the parameters of the posterior distribution (Dolgui and Pashkevich, 2008). In addition, there is no theoretical or empirical evidence in support of the Beta-Binomial distribution for intermittent demands (Eaves, 2002). This method will therefore not be considered further in this paper.

Hence, the motivation behind this research work is based on two arguments. First, the Bayesian approach is intuitively appealing when forecasting spare parts demand due to the appropriate assumption of demand distribution parameters changing over time, which reflects the changing failure rates. Second, there is strong empirical evidence in support of the compound Poisson demand assumption when modelling lead-time demand in a spare parts context (Syntetos et al., 2013; Lengu et al., 2014). It seems natural then to propose a Bayesian method based on such distribution. Although in this case there is no known conjugate prior that leads to a posterior distribution in a closed form, we use an approximate likelihood function to circumvent the issue. It enables us to calculate the order-up-to-level without much loss of information on the observed data and in a reasonable computational time. We will show that the proposed Bayesian method leads to a higher inventory efficiency than alternative parametric and non-parametric forecasting methods when a periodic order-up-to-level inventory control policy is considered under a cycle service level constraint. Hence, the contribution of this paper is three-fold:

1. We propose a new Bayesian method based on compound Poisson demand;
2. We evaluate the empirical performance of the Poisson-based Bayesian method presented in Aronis et al. (2004);
3. We compare the empirical performance of the two Bayesian methods to a parametric frequentist and a non-parametric one.

The remainder of the paper is organised as follows. Section 2 presents the proposed Bayesian method as well as the three benchmark forecasting methods. In Section 3, we describe the numerical investigation conducted to compare the forecasting methods and present the numerical results. Section 4 is dedicated to an empirical investigation and the discussion of its findings. Conclusions and next steps of research are presented in Section 5.

2. Proposed Bayesian method and benchmarks

In this section, we first present the proposed compound Poisson Bayesian method, hereafter referred to as the CPB method, followed by the three considered benchmarks.

2.1 Proposed compound Poisson Bayesian (CPB) method

We assume that demand follows a compound Poisson distribution, which means that demand arrivals follow a Poisson process and the demand sizes are variable and characterised by a demand size distribution (a single parameter distribution is assumed for simplicity). This is in line with the compound nature of intermittent demand patterns we consider in this paper. We also assume that a set of observed demand data is available over n periods. The following notation is used throughout the paper to present:

λ : the Poisson distribution parameter (demand arrival rate)

θ : the demand size distribution parameter

$D = \{y_1, y_2, \dots, y_n\}$: the observed demand data at the end of period n

$f(\lambda)$: the prior distribution of the underlying parameter λ

$f(\theta)$: the prior distribution of the underlying parameter θ

$\mathcal{L}(\lambda, \theta | y_1, y_2, \dots, y_n)$: the likelihood function of the parameters λ and θ , given the observed data $\{y_1, y_2, \dots, y_n\}$

L : the replenishment lead-time

S_t : the order-up-to-level at period t .

If the two parameters (λ and θ) are assumed to be independent, then based on Bayes Theorem the posterior predictive distribution of demand y given the observed data $D = \{y_1, y_2, \dots, y_n\}$, denoted by $P(y|D)$, is given by:

$$P(y|D) = \int \int P(y|\lambda, \theta) \left(\frac{P(D|\lambda, \theta)f(\lambda, \theta)}{P(D)} \right) d\lambda d\theta$$

$$\begin{aligned}
&= \int \int P(y|\lambda, \theta) \left(\frac{P(D|\lambda, \theta)f(\lambda)f(\theta)}{\int \int P(D|\lambda, \theta)f(\lambda, \theta) d\lambda d\theta} \right) d\lambda d\theta \\
&\propto \int \int P(y|\lambda, \theta) P(D|\lambda, \theta) f(\lambda) f(\theta) d\lambda d\theta \\
&= \int \int P(y|\lambda, \theta) \mathcal{L}(\lambda, \theta|D) f(\lambda) f(\theta) d\lambda d\theta \\
&= \int \int P(y|\lambda, \theta) \mathcal{L}(\lambda, \theta|y_1, y_2, \dots, y_n) f(\lambda) f(\theta) d\lambda d\theta \tag{1}
\end{aligned}$$

Note that the missing proportionality constant $\int \int P(D|\lambda, \theta) f(\lambda, \theta) d\lambda d\theta$ can always be deduced from the fact that $P(y|D)$ is a probability density and it must therefore integrate to one.

The likelihood function $\mathcal{L}(\lambda, \theta|y_1, y_2, \dots, y_n) = \prod_{i=1}^n \mathcal{L}(\lambda, \theta|y_i)$ is a product of n compound Poisson distributions; the likelihood function is not a compound distribution (i.e. it is not conjugate) and it cannot be expressed in a simple form (for example, in terms of a sufficient statistic like the mean). As n increases, it becomes increasingly more difficult to derive this function. An alternative would be to use numerical integration methods based on analytic approximations or quadrature. However, the computational effort involved would be considerable. Therefore, in this paper we propose a different approach, which requires much less computational effort. This approach takes advantage of the fact that compound Poisson distributions are Levy processes and thus infinitely divisible (Sato, 1999). More specifically, let us suppose that $y_i \sim f(\lambda, \theta)$ where $f(\lambda, \theta)$ is the probability mass function of a compound Poisson distribution, λ is the arrival parameter and θ is the event size parameter. Furthermore, let n be a fixed rational number and let $T = \sum_{i=1}^n y_i$. Then $T \sim f(n\lambda, \theta)$.

The proof of this result for the individual compound Poisson distributions is summarised in *Appendix A*. More details of the proof can be found in Johnson et al. (2005).

Instead of taking the likelihood function $\mathcal{L}(\lambda, \theta | y_1, y_2, \dots, y_n)$ which is a function of all n observations $\{y_1, y_2, \dots, y_n\}$, we use the likelihood function $\mathcal{L}(\lambda, \theta | T)$ where $T = \sum_{i=1}^n y_i$. Our likelihood function will thus be a single compound Poisson distribution with parameters $n\lambda$ and θ . The likelihood function $\mathcal{L}(\lambda, \theta | T)$ only considers the total of the observations and, as such, it does not contain as much information as the function $\mathcal{L}(\lambda, \theta | y_1, y_2, \dots, y_n)$ which considers all the individual observations. In number theoretic terms, the different possible combinations of observations $\{y_1, y_2, \dots, y_n\}$ that sum up to T are simply partitions of the number.

It is important to point out that the use of the likelihood function $\mathcal{L}(\lambda, \theta | T)$ might involve some loss of information about the individual observations but the prior distributions should still allow us to get an accurate predictive distribution. Let us take, for example, the case where $T = \sum_{i=1}^{20} y_i = 100$. Considering just two extremes, the total T could have come from a single observation of 100 or from 20 observations of size 5. The likelihood function $\mathcal{L}(\lambda, \theta | T)$ places equal weight on these two outcomes. If the posterior predictive distribution is based only on this likelihood function, it would be fair to conclude that it would not be particularly informative. The posterior predictive distribution, however, will also incorporate our beliefs about the parameters through the prior distribution. Finally, it should be noted that in practice most companies store demand data periodically, which means that data are somehow aggregated over a certain time period, and it is not straightforward to estimate the parameters of a continuous compound Poisson process (Prak et al., 2018). Therefore, our idea of aggregating the demand data to calculate the likelihood function is in line with the research dealing with inventory models under the compound Poisson process.

For the purpose of the numerical and empirical investigation, the prior distribution used for the parameter λ is $f(\lambda) = e^{-\lambda}$, $0 < \lambda < \infty$. This distribution reflects the intermittent nature of demand arrivals and the domain of the distribution is the same as that of the parameter λ . The

prior distribution for the parameter θ is $f(\theta) = 1$. This is chosen for simplicity purposes in the numerical investigation.

Based on these assumptions about the prior distributions, the posterior predictive distribution can be expressed as:

$$\begin{aligned} P(y|D) &\propto \int \int P(y|\lambda, \theta, D) f(\lambda, \theta|D) d\lambda d\theta = \int \int P(y|\lambda, \theta, D) e^{-\lambda} d\lambda d\theta \\ &\propto \int \int P(y|\lambda, \theta) P(D|\lambda, \theta) e^{-\lambda} d\lambda d\theta \end{aligned} \quad (2)$$

Without loss of generality, we assume that demand sizes follow a Geometric distribution, i.e. a Poisson-Geometric demand distribution is considered. The Geometric distribution is largely used in the inventory literature to model transaction sizes due to strong theoretical and empirical evidence in its support (Watson, 1987; Eaves, 2002; Teunter et al., 2010). Moreover, a strong goodness-of-fit of the Poisson-Geometric distribution (also known as the stuttering Poisson) is empirically demonstrated in the literature through investigations based on demand histories of more than 13,000 SKUs (Syntetos et al., 2013; Lengu et al., 2014). Hence, for the Poisson-Geometric demand distribution, we have:

$$\begin{aligned} P(y|D) &\propto \int \int \left[e^{-\lambda} (1 - \theta)^y \sum_{j=1}^y \binom{y-1}{j-1} \frac{[\lambda\theta/(1-\theta)]^j}{j!} \right] \times \\ &\quad \times \left[e^{-n\lambda} (1 - \theta)^T \sum_{j=1}^T \binom{T-1}{j-1} \frac{[n\lambda\theta/(1-\theta)]^j}{j!} \right] \times e^{-\lambda} d\lambda d\theta \end{aligned} \quad (3)$$

For the demand over lead-time period L (plus one review period, i.e. $L+1$ periods, to account for the periodic review in the inventory system we consider), the predictive distribution is given by:

$$\begin{aligned} P(y|D) &\propto \int \int \left[e^{-\lambda(L+1)} (1 - \theta)^y \sum_{j=1}^y \binom{y-1}{j-1} \frac{[\lambda(L+1)\theta/(1-\theta)]^j}{j!} \right] \times \\ &\quad \times \left[e^{-n\lambda} (1 - \theta)^T \sum_{j=1}^T \binom{T-1}{j-1} \frac{[n\lambda\theta/(1-\theta)]^j}{j!} \right] \times e^{-\lambda} d\lambda d\theta \end{aligned} \quad (4)$$

2.2 Benchmark methods

The inventory performance of the CPB method is compared to three different alternatives: (i) the Poisson Bayesian method put forward by Aronis et al. (2004), (ii) a parametric frequentist method, and (iii) a non-parametric forecasting method. These are presented below.

- **Poisson Bayesian method**

In the Aronis et al. (2004)'s Bayesian method, hereafter referred to as the PGB (Poisson Gamma Bayesian) method, the demand is Poisson distributed with parameter λ and the prior distribution of λ is Gamma with shape and scale parameters α and β respectively. The posterior predictive distribution of the demand per period when n demands y_i have been observed, is NBD with a shape parameter $\alpha + \sum_{i=1}^n y_i$ and scale parameter $\frac{\beta+n}{\beta+n+1}$ (Baker and Kharrat, 2018).

In order to evaluate the empirical performance of this method, we need to estimate the parameters α and β . To do so, we use the method of moments. Let $\hat{\alpha}$ and $\hat{\beta}$ be the estimates of α and β when n demands y_i are observed. Therefore, at each period the mean and variance of the demand, denoted by m_d and v_d , respectively, are given by

$$m_d = \frac{(\hat{\alpha} + \sum_{i=1}^n y_i) \left(1 - \frac{\hat{\beta} + n}{\hat{\beta} + n + 1}\right)}{\left(\frac{\hat{\beta} + n}{\hat{\beta} + n + 1}\right)} = \frac{\hat{\alpha} + \sum_{i=1}^n y_i}{\hat{\beta} + n} \quad (5)$$

and

$$v_d = \left(\hat{\alpha} + \sum_{i=1}^n y_i\right) \left(1 - \frac{\hat{\beta} + n}{\hat{\beta} + n + 1}\right) \left(\frac{\hat{\beta} + n}{\hat{\beta} + n + 1}\right)^2 = \frac{\hat{\alpha} + \sum_{i=1}^n y_i}{(\hat{\beta} + n)} + \frac{\hat{\alpha} + \sum_{i=1}^n y_i}{(\hat{\beta} + n)^2} \quad (6)$$

which leads to

$$\hat{\alpha} = \frac{m_d^2}{v_d - m_d} - \sum_{i=1}^n y_i = \frac{m_d^2}{v_d - m_d} - nm_d \quad (7)$$

$$\hat{\beta} = \frac{m_d^2}{v_d - m_d} - \sum_{i=1}^n y_i = \frac{m_d}{v_d - m_d} - n \quad (8)$$

Equations (7) and (8) are used to calculate the shape and scale parameters of the negative binomial distribution of the demand per period. In order to calculate the shape parameter of the negative binomial distribution of the demand over $(L+1)$ periods, the shape parameter should be multiplied by $(L+1)$. The scale parameter remains the same. Hence, under the PGB forecasting method, a closed form expression of the lead-time demand distribution can be used, which enables to easily calculate the inventory policy parameters. The detailed calculation of the inventory policy parameter under the PGB method will be provided further in Section 3.1.

- **Parametric and non-parametric forecasting methods**

The parametric frequentist forecasting method considered as benchmark in this paper is the Syntetos-Boylan Approximation, hereafter referred to as SBA (Syntetos and Boylan, 2005); it is the method with most empirical evidence in its support. It constitutes a bias-correction modification to Croston's method (Croston, 1972). The forecast of the demand at period t using SBA, denoted by F_t , is given by: $F_t = (1 - \alpha/2) (\hat{Z}_t / \hat{T}_t)$ where Z_t and T_t are the actual demand size and demand interval at period t respectively, \hat{Z}_t and \hat{T}_t are their respective estimates calculated using exponential smoothing and α a smoothing constant. Note that the update of the demand sizes and intervals take place only when demand occurs. The mean and variance of the lead-time demand are estimated by multiplying the forecast and the smoothed mean squared forecast error of the demand per period by $(L + 1)$.

The non-parametric method considered is the one proposed by Willemain, Smart and Schwarz (2004), hereafter referred to as WSS. WSS is a bootstrapping method that randomly samples, with replacement, $(L+1)$ demand values from historical information. A two-state (zero / non-zero) Markov process is used to model transition probabilities between the sampled $(L+1)$

demand values. The sampling procedure is replicated many times to build an empirical distribution of the lead-time demand.

The SBA and WSS methods have been shown in the literature to be (among) the best performing ones for intermittent demand patterns and their operation is described in more detail in Syntetos et al. (2015) and Hasni et al. (2019b). More details on the calculation of the inventory policy parameters under SBA and WSS and their implementation will be given in Section 3.1.

In the next two sections, we conduct both a numerical and an empirical investigation to analyse the inventory performance of the methods discussed thus far in the paper. The former allows for the consideration (by design) of a diverse and wide range of demand characteristics. The latter offers the ‘credibility’ associated with analysing empirical data.

3. Numerical investigation

3.1 Data and experimental settings

The theoretical generated dataset used for the purpose of the numerical investigation consists of 7,400 demand series where each demand series is 100 periods long. The demand series are randomly generated from different compound Poisson distributions: $Poisson(\lambda) - Geometric(\theta)$, $Poisson(\lambda) - Logarithmic\ Series(\varphi)$ and $Poisson(\lambda) - Poisson(\rho)$. These distributions are chosen in conjunction with a well-informed selection of parameters to generate demand data with a reasonably wide range of demand arrival rates and transaction size modality and variability. In order to ensure that demand arrival is intermittent, we have considered the values of the parameter λ in the range 0.05-1.95, step 0.1. The transaction size variability ranges from 0.05 to 4.95 and the transaction size modality from 1 to 25. A sample of five demand series is generated for each combination of the parameter values to reduce sampling error.

To evaluate the performance of the four forecasting methods, as commonly performed in the literature, we consider their stock control implications, which are reflected by the stock on hand and backorders and their achieved service level (Teunter and Duncan, 2008; Teunter et al., 2011; Khan et al., 2019; Klibi et al., 2018; Turrini and Meissner, 2019). We do so by considering an order-up-to-level (T, S) inventory control policy, where the optimal order-up-to-level is calculated to satisfy a target cycle service level (*CSL*). The *CSL* is the fraction of replenishment cycles in which all of the demand can be met from stock (Silver et al., 2017). Other service measures, such as the fill rate or the ready rate, are not considered because the bootstrapping WSS approach does not allow direct calculation of such measures. Under this policy, every T periods the inventory position is reviewed and an order is triggered if it is found to be below the order-up-to-level S (to raise it up to S). The order arrives after a lead-time L and any demand that is not satisfied from stock on-hand is backordered. Three lead-time values are considered in the numerical investigation $L = 1, 3$ and 5 . Three target *CSLs* are used, namely: $CSL = 85\%, 90\%$ and 95% .

We (reasonably) assume that the inventory review is made every period (i.e. $T = 1$). At each period, the sequence of events is as follows: demand occurs (assumed at the end of the period), net inventory levels are determined, a new order is placed and an order (placed L periods ago) is received.

Under the SBA forecasting method, the first 25 observations are used to initialise the forecasting method (using the mean demand over the within sample of 25 observations) and the next 25 observations are used to optimise the smoothing constants (to obtain the minimal mean square error). An out-of-sample with the last 50 observations is used to report performance. In order to calculate the optimal order-up-to-level, we assume that the lead-time demand follows a Negative Binomial distribution (NBD). Hence, at each period t the order-up-to-level S_t is

calculated as $S_t = \Phi_{SBA,L+1,t}^{-1}(CSL)$ where $\Phi_{SBA,L+1,t}^{-1}(\cdot)$ is the inverse of the cumulative NBD function of the demand over $L+1$ periods. The mean and variance of the NBD are calculated as $\mu_{SBA,L+1,t} = (L+1) * F_t$ and $\sigma_{SBA,L+1,t}^2 = (L+1) * MSE_t$ respectively, where F_t and MSE_t are the forecast and the smoothed mean squared forecast error per period calculated using SBA. Under the WSS method, a sampling with 1,000 replications is used to generate the empirical distribution as in Syntetos et al. (2015) and Hasni et al. (2019b). In both cases, a within sample of 50 periods of each demand series is considered to initialise the forecasting and inventory parameters. The determination of the order-up-to-level at any period t in the out-of-sample when WSS is considered is schematically represented in Figure 1. Since the empirical distribution is discrete and reconstructed upon observed lead-time values only, some percentiles are not readily available. In such cases, linear interpolation is undertaken in order to ‘estimate’ the values of interest.




Figure 1. Determination of S_t for the lead-time (+1) demand distribution.

Under the PGB approach, at each period t in the out-of-sample, if n demand observations y_i ($i = 1, 2, \dots, n$) are available with mean $m_{d,t}$ and variance $v_{d,t}$, we first calculate $\hat{\alpha}_t = \frac{m_{d,t}^2}{v_{d,t} - m_{d,t}}$ and $\hat{\beta}_t = \frac{m_{d,t}}{v_{d,t} - m_{d,t}} - n$. Hence, the order-up-to-level S_t is calculated as $S_t = \Phi_{L+1,t}^{-1}(CSL)$ where $\Phi_{L+1,t}(\cdot)$ is the cumulative NBD function (of the lead-time demand)

that has a shape parameter $r_t = (L + 1)(\hat{\alpha}_t + \sum_{i=1}^n y_i)$ and a scale parameter $p_t = \frac{\hat{\beta}_t + n}{\hat{\beta}_t + n + 1}$.

Note that the NBD function has equivalently a mean and variance given by $\mu_{L+1,t} = \frac{p_t r_t}{1-p_t}$ and

$\sigma_{L+1,t}^2 = \frac{p_t r_t}{(1-p_t)^2}$ respectively. Also note that under both the PGB and SBA approaches, when

NBD is used, the variance should be higher than the mean, therefore if the data give a variance

lower than the mean, we assume $\sigma_{L+1,t}^2 = 1.05 * \mu_{L+1,t}$ and $\sigma_{SBA,L+1,t}^2 = 1.05 * \mu_{SBA,L+1,t}$.

Under the CPB approach, if n demand observations y_i are available (at any period t) with a sum

$T = \sum_{i=1}^n y_i$, then the optimal order-up-to-level S_t is given by: $\sum_{y=0}^{S_t-1} P(y | D) \leq CSL \leq \sum_{y=0}^{S_t} P(y | D)$,

where $P(y|D)$ is calculated as described in Section 2.1 taking into account the normalizing constant.

At the end of each period t , the net stock i_t is calculated using: $i_t = i_{t-1} + Q_{t-L-1} - D_t$ where

Q_{t-L-1} is the order made at the end of period $t-L-1$ (received at the end of period $t-1$) and D_t is

the demand occurring at period t . The stock on hand $i_t^+ = \max(i_t, 0)$, the backorders $i_t^- =$

$\max(-i_t, 0)$ are calculated and the inventory position is updated using: $I_t = i_t + \sum_{i=1}^L Q_{t-i}$.

Then the order is calculated as: $Q_t = \max(S_t - I_t, 0)$.

If the within sample is composed of N_1 periods and the out-of-sample of N_2 periods, then for

each series the average stock on hand, \overline{SOH} , and average backorders, \overline{B} , are calculated using

(9). These averages per series are then calculated across all series for each forecasting method.

$$\overline{SOH} = \frac{1}{N_2} \sum_{t=N_1+1}^{t=N_1+N_2} i_t^+ \quad \text{and} \quad \overline{B} = \frac{1}{N_2} \sum_{t=N_1+1}^{t=N_1+N_2} i_t^- \quad (9)$$

At each period t in the out-of-sample (i.e. $t = 1..N_2$), we calculate CSL_t using (10) to indicate

that there is no backorder at period t .

$$CSL_t = \begin{cases} 1 & \text{if } i_t^- = 0 \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

Then, the average achieved *CSL* for each series is calculated using (11). An average is then calculated across all series for each forecasting method.

$$CSL = \frac{1}{N_2} \sum_{t=N_1+1}^{t=N_1+N_2} CSL_t \quad (11)$$

It should be noted that, when calculating the order-up-to-level S_t , the numerical evaluation of the predictive distribution in the CPB method is more computationally demanding than that of SBA. In fact, the former based on (4) requires a double numerical integration and more summations than the latter that is simply based on a simple numerical integration of the probability distribution function of NBD. These additional integrations and summations imply a higher complexity for the CPB method; it may be harder to be understood and implemented by practitioners than SBA.

3.2 Numerical results

For each forecasting method, target *CSL* and lead-time, we report in Table 1 the numerical results of the inventory performance. Table 1 shows the average stock on hand, the average backorders and the achieved *CSL*.

	Target <i>CSL</i>	<i>L</i> = 1			<i>L</i> = 3			<i>L</i> = 5		
		Stock on hand	Backorders	<i>CSL</i> (%)	Stock on hand	Backorders	<i>CSL</i> (%)	Stock on hand	Backorders	<i>CSL</i> (%)
WSS	85%	32.86	3.76	86.75	54.80	3.54	86.74	63.52	4.1	84.28
	90%	46.99	2.74	90.28	72.62	2.78	89.49	85.12	3.3	87.11
	95%	76.89	1.61	93.85	102.90	1.98	92.19	112.85	2.58	89.72
SBA	85%	33.84	3.73	83.92	54.30	3.54	82.47	69.57	3.57	80.74
	90%	47.09	2.70	88.94	69.67	2.69	87.08	85.75	2.79	85.32
	95%	72.00	1.62	93.92	96.66	1.73	92.08	167.6	1.84	91.12
CPB	85%	29.49	4.13	84.34	50.35	3.30	86.23	68.83	2.86	87.24
	90%	36.12	3.42	88.18	63.23	2.57	90.25	87.69	2.16	91.05
	95%	48.29	2.56	92.49	86.87	1.77	94.16	122.17	1.42	94.78
PGB	85%	32.00	3.65	87.65	52.14	3.16	88.84	68.63	2.91	89.33
	90%	44.45	2.61	91.52	68.49	2.26	92.18	87.68	2.08	92.53
	95%	68.21	1.49	95.53	98.00	1.31	95.76	121.18	1.21	95.84

Table 1. Inventory performance of the four forecasting methods

The results in Table 1 show that for $L = 1$ and 3, the CPB method results in the lowest stock on hand. For $L = 5$, the stock on hand of the CPB method relatively increases and WSS leads to the lowest stock on hand. Obviously, the increase in the stock on hand leads to a decrease of the backorders. By looking at the achieved *CSLs*, the results show that the PGB method is the only one that achieves all the Target *CSLs* for all lead-times and our proposed Bayesian method achieves the Target $CSL = 85\%$ and $CSL = 90\%$ for $L = 3$ and 5. Furthermore, it is important to mention that the target $CSL = 95\%$ cannot be achieved by the WSS regardless of the lead-time value. When the lead-time decreases, *CSL* achievement becomes easier, especially for $L = 1$. This is expected; it is more likely for the WSS method to achieve target *CSLs* for lower lead-times (Hasni et al., 2019b). With regard to the parametric frequentist method, the results show that SBA leads to an under-achievement of all targeted levels. The under-achievement is higher for high lead-times. This is also expected as SBA is associated with a (small) negative bias that leads to lower achieved *CSLs* (Syntetos et al., 2006).

Since the results show that an increase of the stock on hand is accompanied with a decrease of the backorders and an increase in the achieved *CSL*, the outperformance of a particular method cannot be concluded. In order to deduce a more conclusive comparison of the performance of the four forecasting methods, it is necessary to perform an additional analysis based on their inventory efficiency. To do so, in what follows the relative performance of the forecasting methods is summarised in terms of inventory efficiency curves, as in Teunter et al. (2010) and Hasni et al. (2019c). Figures 2-4 show efficiency curves comparing average stock on hand and average backorders. This means that a forecasting method is more efficient if for a certain stock on hand, it leads to lower backorders. Efficiency curves considering the average stock on hand and the achieved *CSL* will be further presented. Figures 2-4 are for the lead-time $L = 1, 3$ and 5,

respectively. In each figure, the closer the curve is to the x -axis, the more efficient the forecasting method is.




Figure 2. Stock on hand vs. backorders for $L = 1$




Figure 3. Stock on hand vs. backorders for $L = 3$




Figure 4. Stock on hand vs. backorders for $L=5$

The results show that, regardless of the lead-time value, the PGB method leads to the highest efficiency since it leads to the lowest backorders for a fixed stock on hand. This method is followed by our proposed CPB method and SBA. Note that the slight superiority of the PGB method compared to CPB is explained by the fact that the former uses a prior distribution (i.e. Gamma) with two parameters (leading to NBD), which offers a higher flexibility than the CPB method that is based on a single parameter prior distribution (i.e. Exponential). Note also that NBD, which is equivalent to a Poisson-Logarithmic distribution, has a strong empirical goodness-of-fit to lumpy demand patterns, which explains its high inventory performance (Syntetos et al. 2013). Moreover, these results show that the approximation made for the likelihood function in the CPB method implies some loss of performance, which limits the gain obtained from using the compound Poisson distribution when modelling the demand. The WSS bootstrapping method leads to the lowest performance. The results also show that the relative performance of our proposed Bayesian method increases with the lead-time. It is worth pointing out that when $L = 5$, the performance of our proposed Bayesian method becomes similar to that

of the PGB method for $CSL = 85\%$ and 90% . This performance improvement with the increase of the lead-time is attributed to the decrease of the variability of the lead-time demand, which renders the single parameter prior distribution sufficient to estimate the Poisson parameter in CPB. Note that the efficiency curves of our proposed CPB method, SBA and WSS almost overlap for the lead-time $L = 1$, which means that performance is very similar. However, for $L = 3$ and 5 , the outperformance of the two Bayesian methods becomes more obvious.

The relative performance of the forecasting methods can also be summarised in terms of efficiency curves between the average stock on hand and the achieved CSL . This means that a forecasting method is more efficient if for a certain stock on hand, the method leads to a higher achieved service level. The results for lead times $L=1, 3, 5$ are presented in Figures 5, 6 and 7, respectively. With this set of efficiency curves, the higher the curve, the more efficient a forecasting method is. This is because a higher curve correspond to a higher achieved CSL for a given stock on hand.




Figure 5. Stock on hand vs. achieved CSL for $L = 1$




Figure 6. Stock on hand vs. achieved CSL for $L = 3$




Figure 7. Stock on hand vs. achieved CSL for $L = 5$

The results in Figures 5-7 show both Bayesian methods to lead to the highest efficiency of stock on hand versus achieved CSL , with a slight advantage to the PGB method for higher lead-times. The WSS bootstrapping method comes third and the SBA fourth. The performance ranking of SBA and WSS is the reverse of what we observed earlier when we examined the efficiency curves for the stock on hand versus backorders. In that case, the SBA was second and the WSS

bootstrapping method last. A close look at the detailed numerical results over the series reveals that the superior performance of WSS in this case is explained by the fact that, although it leads to the highest backorders, the backorders occur less frequently than when SBA is used. Note that the higher achieved *CSLs* obtained by WSS compared to SBA confirms what is shown in Hasni et al. (2019c) when a highly variable demand is considered. The results also show that for lower lead-times (i.e. $L = 1$ or 3) and higher target *CSLs* (i.e. $CSL = 95\%$), SBA leads to higher efficiency than the WSS bootstrapping method.

4. Empirical investigation

The dataset used for the purpose of the empirical investigation relates to the demand of spare parts from the automotive industry. The dataset is composed of 2,971 SKUs with a demand history of 24 months. The descriptive statistics of the demand dataset are summarised in Table 1. Hence, we show the minimum, the first quartile, the median, the third quartile and the maximum value of demand intervals, demand sizes and the demand per period of the SKUs under concern.

2,971 SKUs	Demand Intervals		Demand sizes		Demand per period	
	Mean	St. dev.	Mean	St. dev.	Mean	St. dev.
<i>Min.</i>	1.04	0.21	1.00	0.32	0.54	0.54
<i>25%ile</i>	1.10	0.30	2.05	1.14	1.46	1.32
<i>Median</i>	1.26	0.52	2.89	1.75	2.33	1.91
<i>75%ile</i>	1.41	0.73	5.00	3.30	4.17	3.44
<i>Max.</i>	2.00	1.59	193.75	89.10	129.17	81.48

Table 2. Descriptive statistics of the empirical dataset

The descriptive statistics in Table 2 show that the demand dataset is composed of many SKUs with a low degree of intermittence since the average demand intervals can be almost equal to 1. However, many SKUs exhibit a higher degree of intermittence and/or some lumpiness, e.g. in some cases the average demand size can be as high as 193 units. Information about the

number of SKUs and statistics related to different demand patterns in the dataset will be further presented in Section 5.

The lead-time for all SKUs is less than a month, therefore we assume in the empirical investigation that the lead-time $L = 1$ month.

Like in the numerical investigation, we consider an order-up-to policy with the same target *CSLs*. The demand history, composed of 24 months, is split into two parts: the within-sample is composed of 13 months and the remaining part (i.e. 11 months) is used for the out-of-sample performance evaluation. The settings of the forecasting methods are the same as previously described.

For the four forecasting methods, the efficiency curves of the average stock on hand versus the average backorders are given in Figure 8 and the efficiency curves of the average stock on hand versus the achieved *CSL* are given in Figure 9. The detailed numerical results of the average stock on hand, the average backorders and the achieved *CSL* are shown in *Appendix B*.




Figure 8. Empirical results: stock on hand vs. achieved CSL




Figure 9. Empirical results: stock on hand vs. backorders

The results in Figures 8-9 show that the two Bayesian methods lead to the highest efficiency since for a fixed stock on hand they lead to the lowest backorders and the highest achieved *CSLs*. The CPB method results in the highest achieved *CSL* whereas the PGB method is associated with slightly lower backorders. This means that although PGB leads to lower backorders, the backorders themselves occur more frequently than with CPB. The results also show that SBA is more efficient than WSS and it leads, for low target *CSLs*, to the same achieved *CSLs* as the PGB method. Note that the performance of WSS considerably decreases, especially compared to SBA, when the empirical data is used as opposed to theoretical data. This is expected since in the former case the lead-times are small (being equal to 1) and the demand variability is relatively low, which are two reasons contributing to the performance deterioration of bootstrapping, see, e.g., Syntetos et al. (2015). The overachievement of the target *CSL* by SBA compared to WSS when a low demand variability is considered confirms the findings of Hasni et al. (2019c).

Further, the empirical results show that the CPB method leads to an over-achievement of all target *CSLs*. However, when the PGB, SBA and WSS methods are used, they enable the

achievement of the targets $CSL = 85\%$ and $CSL = 90\%$, but not $CSL = 95\%$. It is also worth pointing out that due to the lower variability of the demand in the empirical dataset compared to the theoretical dataset, the empirical performance of CPB improves compared to PGB. This can be explained by the fact that the single parameter prior distribution used in the CPB (i.e., the exponential) becomes sufficient to estimate the Poisson demand process parameter. This is different to the case of the theoretical dataset where the demand variability is higher, which explains the relative higher performance of PGB due to the two-parameter distribution used for the Poisson demand process parameter. Finally, it should be noted that the empirical results reveal that SBA and CPB lead to almost the same achieved $CSLs$, which can be 2% higher than that of WSS. This over achievement also comes with lower backorders.

In order to better analyse the comparative performance of the forecasting methods with respect to the different SKUs' demand patterns, we have split the SKUs in the dataset into four categories according to the categorisation scheme proposed by Syntetos *et al.* (2005). We use two categorization criteria: the demand interval p and the squared coefficient of variation of demand sizes CV^2 . The first category includes the SKUs with a smooth demand, with the cutoff values $p \leq 1.32$ and $CV^2 \leq 0.49$. The second category includes the SKUs with an intermittent demand, with the cutoff values $p > 1.32$ and $CV^2 \leq 0.49$. The third category includes the SKUs with an erratic demand, with the cutoff values $p \leq 1.32$ and $CV^2 > 0.49$. The fourth category includes the SKUs with a lumpy demand, with the cutoff values $p > 1.32$ and $CV^2 > 0.49$. The number of SKUs in each category and its percentage in the dataset are shown in Figure 10.

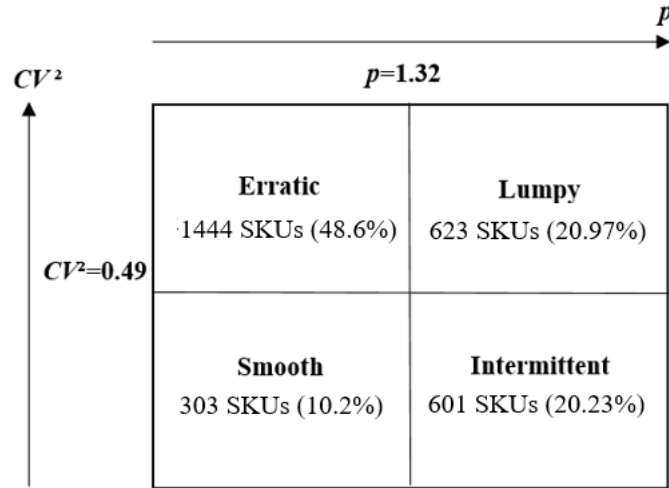


Figure 10. Number of SKUs per category and their percentages

Table 3 and Table 4 show the detailed inventory performance results of the smooth and lumpy SKUs respectively. The detailed results of the two other categories are reported in *Appendix C*. The stock on hand vs. *CSL* efficiency curves of the four forecasting methods in the four categories are presented in *Appendix D*.

303 SKUs	Target CSL	Stock On Hand	Backorders	CSL
WSS	85%	3.41	0.09	95.38%
	90%	4.11	0.07	96.67%
	95%	5.22	0.05	97.39%
SBA	85%	2.56	0.14	90.88%
	90%	3.14	0.08	94.57%
	95%	4.12	0.03	97.81%
CPB	85%	3.92	0.04	97.45%
	90%	5.33	0.00	99.67%
	95%	8.13	0.00	99.97%
PGB	85%	2.55	0.15	90.64%
	90%	3.07	0.08	94.45%
	95%	3.95	0.03	97.75%

Table 3. Empirical results: detailed inventory performance of smooth SKUs

623 SKUs	Target CSL	Stock On Hand	Backorders	CSL
WSS	85%	8.06	0.65	86.76%
	90%	10.35	0.48	90.77%
	95%	13.34	0.32	94.21%
SBA	85%	6.15	0.87	85.23%
	90%	8.05	0.68	88.60%
	95%	11.70	0.46	92.24%
CPB	85%	7.03	0.79	87.90%
	90%	9.24	0.58	91.87%
	95%	13.66	0.34	95.80%
PGB	85%	6.82	0.80	86.28%
	90%	8.69	0.62	89.41%
	95%	11.94	0.39	93.16%

Table 4. Empirical results: detailed inventory performance of lumpy SKUs

The results in Table 3 show that when the intermittence is very low and the demand is smooth, the four forecasting methods lead to an overachievement of the target *CSL*. Figure D1 shows that WSS leads though to the lowest efficiency in terms of the achieved *CSL* for a fixed stock on hand. The results in Table 4 show that for SKUs with a lumpy demand, WSS leads to higher achieved *CSLs* compared to SBA and PGB and a higher efficiency for high target *CSLs*, which is expected since WSS is known to be associated with good inventory performance for lumpy demands (Hasni et al., 2019b). However, as shown in Figure D3, WSS is associated with the lowest inventory efficiency. For the intermittent demand category, the four efficiency curves (as shown in Figure D2) almost overlap, which means that there is no clear outperformance of a particular forecasting method. Note that the CPB is associated with the highest achieved *CSLs* for all categories as well as the highest inventory efficiency.

5. Conclusion

Several Bayesian approaches have been developed in the inventory forecasting literature to deal with spare parts demand or intermittent demand items in general. However, such research has been mainly based on the Poisson demand assumption and is associated with very little

empirical validation (Syntetos et al., 2013; Lengu et al., 2014). We have attempted in this paper to overcome this limitation by developing a Bayesian-based method under the assumption that demand follows a compound Poisson distribution.

We have conducted a numerical investigation with a theoretically generated dataset composed of 7,400 demand series as well as an empirical experiment based on 2,971 spare parts from the automotive industry. Four forecasting methods have been included in the investigations, namely: Syntetos and Boylan Approximation (SBA), the WSS bootstrapping method, the Poisson-based Bayesian method and our proposed compound-Poisson based Bayesian method. Inventory efficiency curves show the two Bayesian methods leading, overall, to a close performance whilst they outperform SBA and WSS. The numerical results also show that the relative performance of our proposed method increases with the lead-time. The Poisson Bayesian method has led in some cases to superior performance compared to our proposed Bayesian method. This is explained by the fact that the approximation used for the likelihood function in the CPB method implies some loss of performance that is higher than the gain obtained from using the compound Poisson distribution to model the demand rather than the Poisson distribution. The superiority is also due to the fact that PGB uses a prior distribution (i.e. Gamma) with two parameters, which offers a higher flexibility than our proposed CPB method that is based on a prior distribution (i.e. Exponential) with one parameter. Hence, an interesting avenue for further research would be to test the performance of our proposed method with a prior distribution with two parameters such as Gamma.

We have also performed an empirical comparative study of the forecasting methods by considering different demand categories in the dataset. The empirical results show that our proposed CPB method is associated with the highest achieved *CSLs* for all demand categories as well as the highest inventory efficiency. We have shown that for smooth demand the four

forecasting methods lead to an overachievement of the target *CSL* with WSS having the lowest stock on hand vs. *CSL* efficiency. However, for SKUs with a lumpy demand, the performance of WSS increases, especially for higher target *CSL*s. For the intermittent demand category, the four forecasting methods lead to a similar inventory efficiency.

To conclude this research, it should be noted that our proposed CPB method is associated with a higher computational time compared to the PGB one, especially for some SKUs with high demand sizes. An interesting idea to extend this research would be to develop a numerical method to calculate the optimal inventory levels in a more computationally affordable way. It is also worth pointing out that despite the empirical and theoretical higher performance of the proposed Bayesian method, it is computationally more demanding and it is certainly more complicated and difficult to be understood by practitioners than parametric frequentist methods, particularly the SBA one. Hence, our recommendation to managers is to analyse if this outperformance is worth the considerable added complexity when selecting a forecasting method.

References

- Aronis K.P., Magou, L., Dekker, R., Tagaras, G. (2004). Inventory control of spare parts using a Bayesian approach: A case study, *European Journal of Operational Research*, 154, 730–739.
- Azoury, K.S. and Miller, B.L. (1984). A comparison of the optimal ordering levels of Bayesian and non-Bayesian inventory models, *Management Science*, 30 (8), 993-1003.
- Azoury, K.S. (1985). Bayes solution to dynamic inventory models under unknown demand distribution, *Management Science*, 31 (9), 1150-1160.
- Babai, M.Z., Tsadiras, A., Papadopoulos, C. (2020). On the empirical performance of some new neural network methods for forecasting intermittent demand, *IMA Journal of Management Mathematics*, doi:10.1093/imaman/dpaa003.
- Baker, R., Kharrat, T. (2018). Event count distributions from renewal processes: fast computation of probabilities, *IMA Journal of Management Mathematics*, 29 (4), 415–433.
- Boutselis, P. and McNaught, K. (2019). Using Bayesian Networks to forecast spares demand from equipment failures in a changing service logistics context, *International Journal of Production Economics*, 209, 325-333.

- Brown Jr, G.F. and Rogers, W.F. (1973). A Bayesian approach to demand estimation and inventory provisioning, *Naval Research Logistics Quarterly*, 20 (4), 607-624.
- Croston, J.D., (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly*, 23 (3), 289-304.
- De Wit, J.R. (1983). Inventory problems with slow moving items: a Bayesian approach, *The Statistician*, 32 (1/2), 201-206.
- Dekker, R., Pinç, C., Zuidwijk, R., Jalil, M.N. (2013). On the use of installed base information for spare parts logistics: a review of ideas and industry practice, *International Journal of Production Economics*, 143 (2), 536–545.
- Dolgui, A. and Pashkevich, M. (2008). On the performance of binomial and beta-binomial models of demand forecasting for multiple slow-moving inventory items, *Computers & Operations Research*, 35 (3), 893-905.
- Eaves, A.H.C. (2002). Forecasting for the ordering and stock-holding of consumable spare parts. Unpublished PhD thesis, Lancaster University, Lancaster, UK.
- Hasni, M., Aguir, M.S., Babai, M.Z., Jemai, Z. (2019a). Spare Parts Demands Forecasting: A Review on Bootstrapping Methods, *International Journal of Production Research*, 57 (15), 4791-4804.
- Hasni, M., Babai, M.Z., Aguir, M.S., Jemai, Z. (2019b). An investigation on bootstrapping forecasting methods for intermittent demands, *International Journal of Production Economics*, 209, 20-29.
- Hasni, M., Aguir, M.S., Babai, M.Z., Jemai, Z. (2019c). On the performance of adjusted bootstrapping methods for intermittent demand forecasting, *International Journal of Production Economics*, 216, 145-153.
- Johnson, N.L., Kemp, A.W., and Kotz, S. (2005), *Univariate discrete distributions*, Wiley-Interscience, Hoboken, New Jersey.
- Karmarkar, U.S. (1994). A robust forecasting technique for inventory and leadtime management, *Journal of Operations management*, 12 (1), 45-54.
- Khan, M.H., Ahmed, S., Hussain, D. (2019). Analysis of Bullwhip effect: A Behavioral Approach. *Supply Chain Forum: An International Journal*, 20(4), 310-331.
- Klibi, W., Rice, J.B., Urciuoli, L. (2018). Special dossier: quantifying supply chain resilience, *Supply Chain Forum: An International Journal*, 19(4), 253-254.
- Lengu, D., Syntetos, A.A., Babai, M.Z. (2014). Spare parts management: linking distributional assumptions to demand classification, *European Journal of Operational Research*, 235 (3), 624-635.
- Popovic, J.B. (1987). Decision making on stock levels in cases of uncertain demand rate, *European Journal of Operational Research*, 32 (2), 276-290.
- Prak, D., Teunter, R. H., Babai, M. Z., Syntetos, A. A., and Boylan, J. E. (2018). Forecasting and Inventory Control with Compound Poisson Demand Using Periodic Demand Data. (SOM Research Reports; Vol. 2018, No. 010). Groningen: University of Groningen, SOM research school.

- Sato, K.I. (1999). Lévy processes and infinitely divisible distributions, Cambridge University Press, Cambridge.
- Silver, E.A. (1965). Bayesian determination of the reorder point of a slow moving item, *Operations Research*, 13 (6), 989-997.
- Silver, E., Pyke, D., Thomas, D.J. (2017). Inventory and Production Management in Supply Chains, 4th edition, CRC Press (Francis and Taylor Group), New York (USA).
- Syntetos, A.A., Babai, M.Z., Gardner, E.S. (2015). Forecasting intermittent inventory demands: simple parametric methods vs. bootstrapping, *Journal of Business Research*, 68 (8), 1746-1752.
- Syntetos, A.A., Boylan, J.E. (2005). The accuracy of intermittent demand estimates, *International Journal of Forecasting*, 21, 303 - 314.
- Syntetos, A.A., Boylan, J.E., Croston, J.D. (2005). On the categorization of demand patterns, *Journal of the Operational Research Society*, 56, 495–503.
- Syntetos, A.A., Boylan, J.E. (2006). On the stock-control performance of intermittent demand estimators. *International Journal of Production Economics*, 103(1), 36–47.
- Syntetos, A. A., Babai, Z., Boylan, J.E., Kolassa, S., Nikolopoulos, K. (2016). Supply Chain Forecasting: Theory, Practice, their Gap and the Future, *European Journal of Operational Research*, 252 (1), 1-26.
- Syntetos, A.A., Lengu, D., Babai, M.Z. (2013). A note on the demand distributions of spare parts, *International Journal of Production Research*, 51(21), 6356-6358.
- Teunter, R., Syntetos, A.A., Babai, M.Z. (2010). Determining order-up-to levels under periodic review for compound binomial (intermittent) demand, *European Journal of Operational Research*, 203, 619-624.
- Turrini, L., Meissner, J. (2019). Spare parts inventory management: New evidence from distribution fitting, *European Journal of Operational Research*, 273, 118–130.
- Watson, R.B. (1987). The effects of demand-forecast fluctuations on customer service and inventory cost when demand is lumpy. *Journal of the Operational Research Society*, 38, 75-82.
- Willemain, T.R., Smart, C.N., and Schwarz, H.F. (2004). A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting*, 20(3), 375-387.
- Yelland, P.M. (2010). Bayesian forecasting of parts demand, *International Journal of Forecasting*, 26 (2), 374-396.

Appendix A. Moment generating function of the sum of compound Poisson distributions

The probability generating function of a $Poisson(\lambda) - Geometric(\theta)$ distribution is given by:

$$P[Y = y] = e^{-\lambda}(1 - \theta)^y \sum_{j=1}^y \binom{y-1}{j-1} \frac{[\lambda\theta/(1-\theta)]^j}{j!}$$

and the moment generating function of the $Poisson(\lambda) - Geometric(\theta)$ distribution, denoted by $M_Y(t)$, is given by:

$$M_Y(t) = \exp \left[\frac{\lambda}{(1-\theta)} \left(\frac{\theta}{[1 - (1-\theta)e^t]} - 1 \right) \right]$$

Let y_i follows a $Poisson(\lambda) - Geometric(\theta)$ distribution and let $T = \sum_{i=1}^n y_i$. Then the moment generating function of T denoted by $M_S(t)$ is given by:

$$\begin{aligned} M_S(t) &= \prod_{i=1}^n M_Y(t) = \prod_{i=1}^n \exp \left[\frac{\lambda}{(1-\theta)} \left(\frac{\theta}{[1 - (1-\theta)e^t]} - 1 \right) \right] \\ &= \exp \left[\frac{n\lambda}{(1-\theta)} \left(\frac{\theta}{[1 - (1-\theta)e^t]} - 1 \right) \right] \end{aligned}$$

The last term is the moment generating function of the $Poisson(n\lambda) - Geometric(\theta)$ distribution.

Appendix B. Detailed empirical results of inventory performance (all SKUs)

2,971 SKUs	Target CSL	Stock on hand	Backorders	CSL (%)
WSS	85%	6.61	0.76	87.49
	90%	8.28	0.63	90.12
	95%	10.62	0.52	92.31
SBA	85%	5.78	0.69	86.55
	90%	7.30	0.52	90.26
	95%	10.04	0.35	93.95
CPB	85%	7.42	0.49	91.74
	90%	10.02	0.32	95.43
	95%	15.38	0.16	98.20
PGB	85%	6.13	0.62	87.34
	90%	7.60	0.46	90.91
	95%	10.10	0.28	94.57

Appendix C. Detailed empirical results of inventory performance of SKUs with intermittent and erratic demand patterns

601 SKUs	Target CSL	Stock On Hand	Backorders	CSL
WSS	85%	3.00	0.05	96.43%
	90%	3.68	0.03	98.18%
	95%	4.73	0.01	99.58%
SBA	85%	2.04	0.13	91.36%
	90%	2.58	0.08	94.58%
	95%	3.46	0.03	97.55%
CPB	85%	2.95	0.05	96.32%
	90%	4.06	0.02	98.79%
	95%	6.22	0.00	99.94%
PGB	85%	2.04	0.13	91.01%
	90%	2.50	0.09	94.07%
	95%	3.24	0.04	96.96%

Table C1. Detailed empirical results of inventory performance of SKUs with an intermittent demand

1,444 SKUs	Target CSL	Stock On Hand	Backorders	CSL
WSS	85%	8.16	1.24	82.42%
	90%	10.17	1.06	85.12%
	95%	13.04	0.91	87.40%
SBA	85%	7.85	0.96	84.34%
	90%	9.82	0.73	88.42%
	95%	13.31	0.50	92.48%
CPB	85%	10.18	0.65	90.29%
	90%	13.83	0.39	94.69%
	95%	21.47	0.19	98.15%
PGB	85%	8.29	0.85	85.58%
	90%	10.20	0.62	89.51%
	95%	13.46	0.39	93.52%

Table C2. Detailed empirical results of inventory performance of SKUs with an erratic demand

Appendix D. Efficiency curves of the forecasting methods for the four demand categories




Figure D1. Empirical results: Stock on hand vs. achieved CSL (Smooth demand)




Figure D2. Empirical results: Stock on hand vs. achieved CSL (Intermittent demand)




Figure D3. Empirical results: Stock on hand vs. achieved CSL (Lumpy demand)




Figure D4. Empirical results: Stock on hand vs. achieved CSL (Erratic demand)