Practice-inspired contributions to inventory theory

Prak, Derk Rutger Jordi

Publication date:
2019

CHAPTER 4

Compound Poisson parameter estimation for inventory control

Abstract. *Most companies store demand data periodically (e.g. weekly) and make periodic demand forecasts, although inventories are controlled continuously. The same 'mismatch' can be observed in the literature; most forecasting methods are periodic, but inventory control theory often assumes the demand process to be continuous and given. Guidance on estimating the parameters of a continuous demand process from period demand data is lacking, in particular for the popular and well-studied compound Poisson class of demand. Commercial software packages wrongly fit compound Poisson processes based on period Size-Interval forecasting methods, leading to dramatic overshoots of the target service level and therefore too high inventory costs. Standard statistical methods have severe biases in finite samples (method-of-moments - MM) and/or are not available in closed form (maximum likelihood - ML). We propose an intuitive, consistent, closed-form MM alternative that dominates in terms of estimation accuracy and on-target inventory performance.*

This chapter is based on Prak et al. (2019): Prak, D.R.J., R.H. Teunter, M.Z. Babai, J.E. Boylan, A.A. Syntetos. 2019. Compound Poisson parameter estimation for inventory control - submitted for publication.

4.1 Introduction

Compound Poisson demand is one of the most general and widely studied classes of demand processes for inventory control. Using a Poisson process to model customer arrivals and some general compounding distribution to model their individual demands, it is flexible from a modeling perspective. Furthermore, it is shown to often provide the best fit to demand data, and many inventory control models using compound Poisson demand have been developed during the past 60 years, as we further discuss in the literature review of Section 4.2. Compound Poisson demand is a standard choice in textbooks (Zipkin, 2000; Silver et al., 2017; Axsäter, 2015) and in commercial software (e.g. Forecast Pro, Oracle, SAP APO, SAS, and Slimstock).

To apply compound Poisson demand processes in inventory control models, estimates of the (individual customer) arrival rate and demand size are required. In the inventory control literature, the problem of how to obtain these parameters from observed demands is typically not discussed and often not even mentioned. The forecasting literature remarkably also does not answer this question, as this is centered around producing forecasts of demand per period or during the lead time, rather than estimating the underlying demand process parameters. Depending on the specific organization of the demand forecasting (with data stored per transaction or per period) and inventory control (with continuous or periodic review) systems, four different scenarios may occur, which we summarize in Table 4.1. The left half of the table - where all individual transaction sizes and times are stored - allows for straightforward estimation of the unknown parameters, as can be found in standard statistics textbooks.

The right half of Table 4.1, to the contrary, is definitely not trivial to analyze, al-

Table 4.1: *Forecasting requirements for inventory control (compound Poisson demand)*

Forecasting (System) / Inventory (Control)	Per Transaction	Per Period
Continuous Review	Individual demand based on transaction data	**Individual demand based on periodic data**
Periodic Review	Period demand based on transaction data	Period demand based on periodic data

though highly relevant in practice. All companies that the authors of this paper have encountered store and forecast demand per period. This has several advantages. First, companies seek to reduce the amount of stored data (driven by e.g. cost considerations, privacy regulations, and preservation of tractability). Second, almost the entire forecasting literature is focused on forecasting period demand, and even temporal data aggregation to improve forecast accuracy is frequently suggested (see e.g. Nikolopoulos et al., 2011; Kourentzes et al., 2014). Third, most commercial forecasting software is based on period forecasting techniques that take period demand data as input to produce a forecast of period demand.

However, the question how to obtain compound Poisson parameter estimates from period demand data or forecasts is left almost completely un-addressed in the literature. In fact, a vast part of the forecasting literature even ignores that a demand process consists of individual customer arrivals and their demand sizes. As a result, order levels are calculated based on (implicit) normality assumptions or other ad-hoc procedures, completely ignoring the actual distribution that was fitted to the demand data.

In the bottom-right corner of Table 4.1, a specification of only the distribution of total demand per period may be enough to perform periodic review inventory control calculations. This depends on the exact assumptions underlying the model regarding e.g. the moments of demand arrival, incurring of inventory costs, and the lead time being a multiple of the review period length. We refer interested readers to e.g. Axsäter (2015) for a detailed discussion. We do not discuss it further here, as it is reasonably straightforward (under the right assumptions) to optimize periodic inventory decisions based on period demand forecasts.

Instead, the top-right corner of Table 4.1 is our particular area of focus, since it has hardly received any attention although it applies to many business settings. Contrarily to demand storage and forecasting, inventory control is often done on a (nearly) continuous basis. When a demand arrives that makes the inventory level drop to or below some reorder level, ERP or dedicated inventory control software systems will directly signal that a replenishment order is needed. Depending on the degree of automation, the order is either placed directly or after approval by a planner. Continuous review inventory control requires the modeling of individual customer arrivals and their demand sizes to calculate e.g. the reorder level undershoot arising from an individual transaction. Compound Poisson demand processes are preeminently suitable for this. However, a full specification of the demand process and all its parameters has to be obtained from the period demand observations, for which no guidance currently exists. This paper aims to fill this gap in the literature.

Separate forecasting of demand sizes and demand inter-arrival times may be achieved by so-called Size-Interval methods, of which the most well-known one was developed by Croston (1972). In Section 4.2 we will review the Size-Interval forecasting literature. Size-Interval methods are particularly popular for intermittent demand patterns where a significant number of periods show no demand, which is also the scenario in which compound Poisson distributions yield the largest advantages over other demand distributions. It is therefore not surprising that all five above-mentioned software packages implement Croston's method to forecast demand and fit compound Poisson demand processes. However, Size-Interval methods were developed for forecasting period demand, and are therefore not suitable for estimating the parameters of a compound Poisson process, which requires the arrival rate and average size of individual demands. By looking at period demands rather than individual demands, no distinction can be made between one large or several smaller orders arriving in a period. This leads to flawed inventory calculations and thereby to deviations from the target service level, even if a long demand history is available.

The first contribution of this paper is a quantification of the asymptotic bias that results from mis-using period demand interval and size estimates as demand process parameters, and of the consequences on achieved service levels. We discuss two variants of Croston's method, which are often referred to as benchmarks in the literature and most used in practice. We furthermore include another intuitive method that uses the average of all non-zero period demand sizes to estimate the mean demand size, and the average of all period inter-arrival times to estimate the arrival rate. We show that also this method does not yield consistent compound Poisson demand parameter estimates, although it does consistently estimate period demand.

The few inventory control studies and books that do discuss how compound Poisson demand parameters can be obtained from (periodic) empirical data suggest standard method-of-moments (MM) estimators, without motivation or comparison with other alternatives (see Section 4.2). The statistics literature furthermore discusses Maximum Likelihood (ML) methods. Both MM and ML are asymptotically unbiased, but no guidelines exist on if and when they show good performance in finite samples. Standard MM equates central moments of the compound distribution to their sample equivalents. This yields a closed-form solution, but the use of 'higher' moments (such as the variance) leads to slow convergence for irregular (e.g. intermittent) demand patterns. ML has asymptotically the lowest variance, but no closed-form solution. Its numerical complexity may explain why it is typically omitted in practice and applied research.

The second and main contribution of this paper is the presentation of an alternative MM estimator that retains the core idea of separately estimating customer arrival rates and demand sizes. It bases the customer arrival rate on the fraction of periods without demand, and then calculates the expected customer demand size for which the expected period demand equals the observed average period demand. As no 'higher' moments are involved, this method is particularly suitable for intermittent demand. Furthermore, this estimator has an intuitive, closed-form solution, making it easily applicable in practice and implementable in software.

In an extensive numerical study, using both the geometric and exponential compounding distribution, we compare the proposed method to standard MM and Maximum Likelihood (ML). We demonstrate for a wide range of intermittent demand scenarios that the proposed method shows better estimation accuracy than the standard MM estimator, and performs similarly to ML (and in some scenarios even slightly better). We show that the fill rate achieved by the proposed method converges equally fast to the target as that achieved by ML, whereas standard MM leads to very slow convergence and order levels that are significantly too low for a wide range of scenarios and sample sizes.

The remainder of this paper is organized as follows. Section 4.2 discusses the relevant literature. Section 4.3 specifies the demand and inventory model and discusses the consequences of mis-using period Size-Interval methods for estimating compound Poisson demand parameters. Section 4.4 discusses the standard MM and ML benchmark parameter estimators from the literature and introduces the proposed estimator. Section 4.5 presents the numerical results of comparing these three methods, first in terms of parameter estimation accuracy, and then in terms of achieved fill rate in the base stock inventory model. Section 4.6 concludes.

4.2 Related literature

We first review the adoption and development of compound Poisson processes in the theoretical inventory control literature and their appraisal in empirical studies. Subsequently, we discuss the relevant demand forecasting literature, and explain why techniques presented there are not suitable for estimating the parameters of compound Poisson processes and therefore cannot serve as inputs for performing inventory control. We then discuss the few studies that have addressed suitable estimators, and provide background for the new procedure that we propose.

After Galliher et al. (1959) studied the special case of Poisson-geometric demand, Feeney and Sherbrooke (1966) (corrected by Chen et al., 2011) derived order-up-to policies under general compound Poisson demand. Archibald and Silver (1978) extended this work to general (s, S) policies. Order-up-to policies can be solved efficiently under compound Poisson demand, as shown by Sherbrooke (1968), Graves (1985), and Babai et al. (2011), who provided approximate procedures, and Axsäter (1990), who presented an exact method. Solving general batch ordering policies is more difficult, but an approximate method that replaces compound Poisson distributions by 'equivalent' Poisson distributions was given by Axsäter et al. (1994). Further extensions of inventory control models with compound Poisson demand were given by e.g. Forsberg (1995) and Axsäter (2000), who studied two-echelon inventory systems; Cheung (1996), who introduced stochastic lead times; and Bensoussan et al. (2005), who showed optimality of (s, S) policies when demand is a mixture of a compound Poisson process and either a diffusion process or constant demand. Also, there are studies that have compared the fit of several distributions to real demand data (Snyder et al., 2012; Syntetos et al., 2012, 2013; Lengu et al., 2014; Turrini and Meissner, 2019). All concluded that especially if products are not fast-moving, compound Poisson distributions provide a better fit for period demand than e.g. the normal or gamma distribution. Some of the many examples of research applying compound Poisson demand are Chaouch (2001), Axsäter (2003), Presman and Sethi (2006), Liu and Song (2012), Shi et al. (2014), Feng et al. (2015), and Ding et al. (2016).

Compound Poisson demand processes are obviously characterized by the demand rate and the demand size distribution. Correspondingly, in the forecasting literature a stream has developed on separately estimating these two components of demand, albeit for period demand rather than individual customer demands. Croston (1972) was the first to present such a Size-Interval method, twice applying exponential smoothing. Other authors continued this research, as reviewed by Gardner (2006). Several papers showed that Croston's method outperforms other simple methods such as single exponential smoothing (Ghobbar and Friend, 2003; Eaves and Kingsman, 2004; Queenan et al., 2007; Gutierrez et al., 2008), and is competitive with more complex non-parametric methods such as bootstrapping (Willemain et al., 2004; Teunter and Duncan, 2009; Syntetos et al., 2015a). Since Croston's method forecasts period demand from dividing the average non-zero demand size by the number of periods between positive demands, an inversion bias is present. This bias was first pointed out by Syntetos and Boylan (2001) and approximately corrected (using a Taylor series expansion) by Syntetos et al. (2005) for Bernoulli period demand arrivals. For Poisson period demand arrivals, Shale et al. (2006) derived an approximate mark-up term which was later corrected and improved by Syntetos et al. (2015b). However, we stress that despite these corrections, the fundamental issue remains that all Size-Interval forecasting methods estimate period demand, and thus do not yield arrival rate and demand size estimates at the individual customer level. We remark that this discrepancy is not an error of period forecasting techniques in general, but rather an indication that one should not derive demand distribution parameters from these at the individual customer level in order to make inventory calculations.

Interestingly, while inventory control theory for compound Poisson demand distributions has progressed during the past decades, the problem of estimating the demand process parameters remains almost un-addressed. Most empirical research ignores this problem completely, often using ad-hoc methods to calculate order levels based on the estimated mean and standard deviation of lead time demand. Watson (1987) assumed an Erlang distribution of lead time demand, while Schultz (1987) argued that traditional inventory control and forecasting methods are ill-suited for sporadic demands and used a Size-Interval forecasting procedure. However, he subsequently calculated the order level based on normality assumptions. Dunsmuir and

Snyder (1989) used a more realistic demand distribution with a positive probability mass at zero to calculate the order level. Several more recent articles followed the logic of Schultz (1987) and forecasted lead time demand, estimated the forecast error, and subsequently implicitly assumed normality to compute order levels (e.g. Chatfield et al., 2004; Syntetos et al., 2009a; Teunter and Sani, 2009; Wang et al., 2010; van Wingerden et al., 2014). Altay et al. (2012) used a power approximation, and Sani and Kingsman (1997) compared various approximate methods. Given the fact that demand data is almost always stored periodically, it is quite natural that so many authors directly estimated period or lead time demand rather than the underlying demand process. However, as explained in the previous section, in order to apply continuous review inventory control models, a complete specification of the demand process and its parameters is required instead.

Very few authors discussed how the unknown parameters of compound Poisson demand (or of demand processes in general) can be obtained from period demand data. Those who did, suggested the use of standard MM estimators. Ward (1978) derived these for Poisson-geometric demand, and Axsäter (2015) for Poisson-geometric and Poisson-logarithmic demand. For general compound Poisson demand such estimators are not discussed in the inventory control literature. Öztürk (1981) derived standard MM estimators for the Poisson-exponential distribution, albeit not for inventory control but for rainfall prediction.

We propose a different MM estimator for general compound Poisson demand that chooses its moments such that the fraction of periods without demand is used to estimate the customer arrival rate, and this estimate combined with the average demand in all periods is used to estimate the average demand size. The inspiration for this method comes from Anscombe (1950), Ehrenberg (1959), and Savani and Zhigljavsky (2006). Although, in line with most of the forecasting literature, they estimated period demand rather than the underlying compound demand process, they did exploit the intermittent nature of that process to do so. They considered period demand that has a negative binomial distribution, corresponding to a Poisson-logarithmic demand process, and showed that using the fraction of periods without demand leads to a simpler and more efficient procedure for estimating period demand. Our method is based on the same logic, but (a) we use it to estimate the parameters of the underlying compound Poisson demand process, (b) we suggest its use for any type of compound Poisson demand process, and (c) we do not only discuss its parameter estimation accuracy, but also study its inventory control performance. The demand model and inventory context are formally introduced in Section 4.3, where we also quantify the consequences of mis-using period Size-Interval techniques to estimate compound Poisson demand parameters.

4.3 Demand model, inventory model, and Size-Interval methods

In this section we specify the general compound Poisson demand model that we use throughout the paper, and the two particular compounding distributions on which we base our numerical study. Thereafter we discuss the base stock inventory model which we apply and the fill rate calculation. Finally, we discuss the consequences of mis-using period Size-Interval estimates as compound Poisson demand parameters at the customer level, quantify the resulting asymptotic bias of the parameter estimates and quantify their effects on the achieved fill rate.

4.3.1 Demand model

We consider a model where demand follows a compound Poisson distribution with some general compounding distribution. We observe n realizations of period demand from this compound Poisson distribution. That is, for $t = 1, 2, ..., n$, we observe realizations of the random variable $X_t = \sum_{i=1}^{K} D_i$, where K follows a Poisson distribution with mean λ and, independently, the D_i follow some distribution with mean μ. This set of demand realizations is used to estimate the parameters of the compound Poisson process.

We focus in our numerical studies on the Poisson-geometric and Poisson-exponential distributions, thereby discussing one of the most frequently used discrete compound Poisson distributions in the inventory control literature and its continuous counterpart. In the former case, the individual demand sizes D_i have the (discrete) probability mass function

$$f_G(x) = (1 - \beta)\beta^{x-1},$$

for $x = 1, 2, ...$ and $0 < \beta < 1$, which has mean $\mu = 1/(1-\beta)$. In the latter case, the D_i have the probability density function

$$f_E(x) = \theta \exp(-\theta x),$$

for $x \geq 0$ and $\theta > 0$, with mean $\mu = 1/\theta$.

4.3.2 Inventory model

We study a continuous review base stock inventory control model with a positive fixed lead time of L periods and full backordering, under a fill rate constraint. Base stock inventory models are relatively easy to optimize and widely used in practice, in particular for spare parts management. We use a fill rate in our model (the fraction of demand that can be serviced from on-hand inventory), as this is one of the most used and realistic service measures from a customer perspective (Thomas, 2005).

We follow the computation of the fill rate for compound Poisson demand models by Axsäter (2015), and generalize it to handle also continuous compounding distributions. Specifically, denoting an individual demand by D and the inventory level by IL, for a given order level S, the fill rate is given by

$$FR(S) = \frac{\mathrm{E}_{IL}\left[\mathrm{E}_D\left[\min(\max(IL, 0), D)\right]\right]}{\mu},$$

which represents the expected fraction of a single order that can be fulfilled from on-hand stock at an arbitrary point in time. The distribution of D is the selected compounding distribution with mean μ, whereas the distribution of IL follows from the observation that the inventory level at time $t+L$ is equal to S minus the lead time demand. Since demand per period is compound Poisson with rate λ and demand size mean μ, the lead time demand follows the same distribution, but with rate λL. Denoting lead time demand by D_L, we find the relationship $\mathrm{P}(IL \leq S-x) = \mathrm{P}(D_L \geq x)$, which completes the fill rate calculation.

4.3.3 Mis-application of period Size-Interval methods

Period Size-Interval methods are designed to yield an estimate of the time between two periods with positive demand and the size of a non-zero period demand. We demonstrate here the consequences of mis-applying these as compound Poisson parameter estimators. We first discuss Croston's method and a 'bias-corrected' version of this method. Thereafter we discuss an intuitive method that uses the average of all non-zero period demands and the average of all times between periods with positive demand, which we refer to as 'Unweighted Averaging.' All discussed methods only update after a period with positive demand.

Croston's method estimates the time between two periods with positive demand and the size of a non-zero period demand both by exponential smoothing. Formally, given that we have observed n periods, the $n-$th period has a positive demand x_n, and we have observed the previous positive demand q_n periods ago, Croston's method sets

$$\hat{\mu}_n = \alpha_1 x_n + (1-\alpha_1)\hat{\mu}_{n-1},$$
$$\hat{p}_n = \alpha_2 q_n + (1-\alpha_2)\hat{p}_{n-1},$$

where $0 < \alpha_1 < 1$ and $0 < \alpha_2 < 1$ are smoothing constants. If $x_n = 0$, then $\hat{\mu}_n = \hat{\mu}_{n-1}$ and $\hat{p}_n = \hat{p}_{n-1}$.

If Croston's period estimates are wrongly taken as the estimates of the parameters of the demand distribution, then $\hat{\mu}_n$ is the estimate after n period observations for the mean μ of the compounding distribution, whereas $1/\hat{p}_n$ is the estimate for λ. As Syntetos et al. (2015b) show, using $1/\hat{p}_n$ as an estimate for λ leads to an inversion bias, as $\mathrm{E}[1/\hat{p}_n] \neq 1/\mathrm{E}[\hat{p}_n]$. This inversion bias does not diminish as the sample size goes to infinity, as Croston's method uses exponential smoothing and thus does not use all data points. Syntetos et al. (2015b) approximately correct for this bias using a second-order Taylor expansion, and derive the estimator $(1-\alpha_2/2)/\hat{p}_n$ for λ. The estimator for μ does not suffer from the inversion bias and remains unaltered.

It is straightforward to show that, with or without the inversion bias correction, Croston's method yields inconsistent parameter estimates when applied to period demand. Consider $\hat{\mu}_n$, which is unaffected by the inversion bias. Instead of estimating the size of an individual demand, it estimates the size of a positive period demand. Therefore,

$$\lim_{n\to\infty} \mathrm{E}[\hat{\mu}_n] = \mathrm{E}(X_n|X_n > 0) = \frac{\mu \sum_{k=1}^{\infty} k \frac{\exp(-\lambda)\lambda^k}{k!}}{1-\exp(-\lambda)} = \frac{\mu\lambda}{1-\exp(-\lambda)},$$

which is strictly larger than μ for all $\lambda > 0$ and $\mu > 0$. This limit approaches μ as $\lambda \to 0$, and $\mu\lambda$ as $\lambda \to \infty$. Observe that the estimator's limit (and thus its bias) is proportional to μ.

The asymptotic bias of the arrival rate is more cumbersome to quantify. Instead of estimating the time between two customer arrivals, $\hat{p}_n$ estimates the number of periods between two periods with positive demand. This number of periods (denoted by Q) follows a geometric distribution with success probability $1-\exp(-\lambda)$, mean $1/(1-\exp(-\lambda))$, and variance $\exp(-\lambda)/(1-\exp(-\lambda))^2$. Although $\mathrm{E}[\hat{p}_n] \to 1/(1-\exp(-\lambda))$ as $n \to \infty$, we can only approximate the limit of $\mathrm{E}[1/\hat{p}_n]$. We do so by a second-order Taylor expansion and find, using the mean and variance of the geometric distribution and the fact that $\hat{p}_n$ is an exponential smoothing estimator with asymptotic variance Var(Q)$\alpha_2/(2-\alpha_2)$, that

$$\lim_{n\to\infty} \mathrm{E}[1/\hat{p}_n] \approx \frac{1}{\mathrm{E}[Q]} + \frac{\alpha_2}{2-\alpha_2}\frac{\mathrm{Var}(Q)}{\mathrm{E}[Q]^3} = \left(1 + \frac{\alpha_2}{2-\alpha_2}\exp(-\lambda)\right)(1-\exp(-\lambda)).$$

As $\alpha_2 \to 0$, meaning that the weight given to old inter-arrival intervals is highest, this limit approaches $1-\exp(-\lambda)$, which is the result without inversion bias. In that scenario, λ is strictly underestimated for all $\lambda > 0$. The approximate limit of the estimator converges to 0 as $\lambda \to 0$, and converges to 1 as $\lambda \to \infty$. As α_2 increases, there are cases with low arrival rates where λ is actually overestimated, but as $\lambda \to 0$, the approximate limit approaches 0, and as $\lambda \to \infty$, it approaches 1. Obviously, the bias correction proposed by Syntetos et al. (2015b) scales the estimated arrival rate, and thus its expectation, by $(1-\alpha_2/2)$. As $\alpha_2 \to 0$, this effect diminishes. In conclusion, Croston's method generally underestimates the arrival rate, but this is somewhat counteracted by its inversion bias, leading to an overestimation especially for low values of λ and/or high values of α_2. However, (approximately) correcting for that inversion bias intensifies the underestimation.

While Croston's method and its variants give greater weight to the most recent demand sizes and intervals, this is not a requirement for Size-Interval methods. An alternative approach is to use equal weights, but retain the separate estimation of mean demand sizes and intervals. The Unweighted Averaging method does this by taking as the mean demand size the average of all positive period demand sizes, and as arrival rate the inverse of the average number of periods between two periods with positive demands. Given that we have observed n periods, of which n_p had a positive demand, the $n-$th period has a positive demand x_n, and we have observed the previous positive demand q_n periods ago, Unweighted Averaging sets

$$\hat{\mu}_n = \frac{1}{n_p} x_n + \frac{n_p - 1}{n_p} \hat{\mu}_{n-1},$$
$$\hat{p}_n = \frac{1}{n_p} q_n + \frac{n_p - 1}{n_p} \hat{p}_{n-1}.$$

It is straightforward to observe that, just like with both Croston variants,

$$\lim_{n\to\infty} \mathrm{E}[\hat{\mu}_n] = \mathrm{E}(X_n | X_n > 0) = \frac{\mu\lambda}{1 - \exp(-\lambda)}.$$

Croston's inversion bias associated with the arrival rate stems from the observation that for a random variable X, $\mathrm{E}[1/X] \neq 1/\mathrm{E}[X]$. The Unweighted Averaging method suffers from the same bias in finite samples. However, unlike Croston's method, it is asymptotically unbiased. Unweighted Averaging namely takes averages over all positive demands and inter-arrival times in the data set rather than applying diminishing weights via exponential smoothing. In this case,

$$\lim_{n\to\infty} \mathrm{E}[1/\hat{p}_n] = 1/\mathrm{E}[\hat{p}_n].$$

Therefore as

$$\mathrm{E}[\hat{p}_n] = \frac{1}{1 - \exp(-\lambda)},$$

it follows that

$$\lim_{n\to\infty} \mathrm{E}[\hat{\lambda}_n] = 1 - \exp(-\lambda).$$

The Unweighted Averaging estimator for λ converges to 0 as $\lambda \to 0$, and to 1 as $\lambda \to \infty$. This clearly shows the flaw in using period demand estimators as individual demand parameters. Even if the arrival rate is so high that at least 1 demand (and typically many more) occurs every period, then the estimated arrival rate can still not exceed 1, as no distinction can be made between 1 or several demands in a period. This leads to an underestimated arrival rate. At the same time, all individual demands in a period are wrongly added and considered as a single customer demand, leading to an overestimated mean demand size. In conclusion, none of the discussed estimators of the period demand size and inter-arrival times leads to consistent estimators for the parameters μ and λ of a compound Poisson demand process.

Table 4.2 quantifies the (approximate) asymptotic bias resulting from using standard Croston's method ('CROST'), the corrected version by Syntetos et al. (2015b) ('SBL'), or the Unweighted Averaging method ('UA'), for several parameter choices. Furthermore, we report the fill rate that would be achieved for Poisson-exponential and Poisson-geometric demand, under the true demand parameters, when the base stock order level is set to the minimum order level that achieves the target fill rate of 95% under the asymptotic values of the estimated parameters. We use a lead time of 2 periods. These settings are arbitrary but other settings lead to the same insights.

Note from Table 4.2 that none of the results depend on μ, except for the achieved fill rates under the geometric compounding distribution. The reason is that for dis-

Table 4.2: *Asymptotic biases and achieved fill rates when using period Size-Interval methods for compound Poisson parameter estimation (target fill rate 95%, lead time $L = 2$)*

True parameters			Asymptotic bias				Achieved fill rates (compounding distr.)					
							(Exponential)			(Geometric)		
λ	μ	α_2	λ CROST	λ SBL	λ UA	μ	Standard	SBL	UA	Standard	SBL	UA
1/16	2	0.1	+1.7%	-3.6%	-3.1%	+3.2%	95.5%	95.4%	95.4%	97.2%	97.2%	97.2%
1/16	2	0.5	+27.3%	-4.5%	-3.1%	+3.2%	95.8%	95.4%	95.4%	97.2%	97.2%	97.2%
1/16	5	0.1	+1.7%	-3.6%	-3.1%	+3.2%	95.5%	95.4%	95.4%	95.7%	95.7%	95.7%
1/16	5	0.5	+27.3%	-4.5%	-3.1%	+3.2%	95.8%	95.4%	95.4%	96.5%	95.7%	95.7%
1/4	2	0.1	-7.9%	-12.5%	-11.5%	+13.0%	96.5%	96.3%	96.3%	97.0%	97.0%	97.0%
1/4	2	0.5	+11.4%	-16.4%	-11.5%	+13.0%	97.1%	96.1%	96.3%	98.2%	97.0%	97.0%
1/4	5	0.1	-7.9%	-12.5%	-11.5%	+13.0%	96.5%	96.3%	96.3%	97.0%	96.4%	97.0%
1/4	5	0.5	+11.4%	-16.4%	-11.5%	+13.0%	97.1%	96.1%	96.3%	97.5%	96.4%	97.0%
1	2	0.1	-35.6%	-38.8%	-36.8%	+58.2%	98.6%	98.4%	98.5%	99.3%	99.0%	99.3%
1	2	0.5	-29.0%	-46.8%	-36.8%	+58.2%	98.9%	97.8%	98.5%	99.3%	98.5%	99.3%
1	5	0.1	-35.6%	-38.8%	-36.8%	+58.2%	98.6%	98.4%	98.5%	98.8%	98.6%	98.8%
1	5	0.5	-29.0%	-46.8%	-36.8%	+58.2%	98.9%	97.8%	98.5%	99.1%	98.2%	98.8%

crete demand distributions, fill rates cannot be achieved exactly, and the jumps in the achieved fill rate depend on all demand parameters. Other than that, μ is only a scaling parameter of the asymptotic biases and the corresponding fill rates. Obviously, α_2 only affects the results of both Croston variants. Another observation is that for low arrival rates and/or high smoothing parameter values, Croston's inversion bias results in an overall overestimation of λ. After the correction by Syntetos et al. (2015b), for all parameter settings an asymptotic underestimation is achieved. The Unweighted Averaging method gives the expected result of uniformly underestimating λ. Its performance falls in-between those of the original Croston's method and the corrected variant. The average demand size μ is always overestimated for all methods. All (absolute) percentage asymptotic biases are increasing in λ. Whereas for higher values of α_2 negative biases for λ are reduced and positive biases are enlarged, the correction has such a counteracting effect that the biases of the SBL arrival rate estimates are also increasing in α_2.

Interestingly, all fill rates are asymptotically too high, irrespective of whether λ is over- or underestimated. This is because the fill rate calculation depends more heavily on μ, as we will further discuss in Section 4.5.2. The achieved fill rates are increasing in λ, and for $\lambda = 1$, which is not an unrealistically high value, they are over 97.8% for all compounding distributions. This implies that the inventory levels and associated costs are dramatically too high. The conclusion of this section is that mis-applying period Size-Interval methods to estimate compound Poisson demand parameters leads to severely overestimated average demand sizes, significantly overshot fill rates and correspondingly excessive inventory costs.

4.4 Consistent estimators

Having discussed the mis-application of period forecasting methods for obtaining compound Poisson demand parameters, we move to consistent estimation methods. We revisit the compound Poisson demand model introduced in the previous section and discuss the standard MM estimator (from the inventory control literature), the ML estimator, and our proposed estimator.

4.4.1 The standard method-of-moments estimator

Standard MM estimators equate the mean and variance of the (compound Poisson) demand distribution to their corresponding sample equivalents, and solve the resulting equations for the unknown parameters. For the geometric compounding distribution this is described by Ward (1978) and Axsäter (2015), and for the exponential compounding distribution by Öztürk (1981). Denote the sample mean of the observations x_t by $\bar{x}$, and their sample variance by s^2. Then for the geometric compounding distribution we find

$$\hat{\lambda} = \frac{2\bar{x}^2}{\bar{x} + s^2} \text{ and } \hat{\mu} = \frac{\bar{x} + s^2}{2\bar{x}},$$

and for the exponential compounding distribution we find

$$\hat{\lambda} = \frac{2\bar{x}^2}{s^2} \text{ and } \hat{\mu} = \frac{s^2}{2\bar{x}}.$$

These estimators are by definition consistent, but may still have severe finite-sample biases and variances. The use of the sample variance makes these estimators sensitive to outliers and leads to slow convergence, as our numerical results will confirm.

4.4.2 The maximum likelihood estimator

The ML estimator for the two demand parameters can be obtained by maximizing the log-likelihood function with respect to these parameters. For both demand distributions that we study, the ML estimator does not exist in closed form, and has to be found by a numerical search procedure. Given the sample $x_1, ..., x_n$ of period demand observations, the log-likelihood is defined as

$$\mathcal{L}(\lambda, \mu | x_1, ..., x_n) = \sum_{t=1}^{n} \log(f(x_t | \lambda, \mu)),$$

where $f(\cdot|\lambda, \mu)$ denotes the probability mass function (for the geometric compounding distribution) or the probability density function (for the exponential compounding distribution) given λ and μ.

The probability mass function of the Poisson-geometric distribution is given by

(Balakrishnan et al., 2017)

$$f_{PG}(x) = \exp(-\lambda) \sum_{i=0}^{x} \frac{1}{i!} \binom{x-1}{i-1} (\lambda/\mu)^i \left(1 - \frac{1}{\mu}\right)^{x-i},$$

and the probability density function of the Poisson-exponential distribution is given by (Öztürk, 1981)

$$f_{PE}(x) = \exp(-\lambda - x/\mu) \sum_{i=0}^{\infty} \left(\frac{(\lambda/\mu)^i x^{i-1}}{i!(i-1)!} \right),$$

for $x > 0$, whereas this distribution has a positive probability mass of $\exp(-\lambda)$ at $x = 0$.

Since the probability mass function and probability density function already contain finite and infinite summations, and the log-likelihood (which is again a summation of these functions over the sample size) is to be optimized by a search procedure, it is evident that ML is the most cumbersome and computationally demanding. However, ML is guaranteed to provide a consistent estimator that asymptotically has the lowest variance, although it offers no guarantee for good performance in finite samples.

4.4.3 The proposed method

We propose an alternative MM estimator and choose its moments such that the intuition of the Size-Interval forecasting literature - explicitly separating estimation of the customer arrival rate and the average demand size - is employed. Denote by N_0 the (stochastic) number of periods in the sample of size n in which no demand occurred. The corresponding realization of that number of periods is denoted by n_0. The occurrence of no demand in a period is Bernoulli distributed with probability $\exp(-\lambda)$, and therefore the number of periods without demand out of n periods in total follows a binomial distribution with parameters n and $\exp(-\lambda)$. Therefore

$$\mathrm{E}[N_0] = n \exp(-\lambda).$$

We equate this sample moment to its expectation to find an estimator for the arrival rate as follows:

$$n\,\exp(-\hat{\lambda}) = n_0,$$

leading to

$$\hat{\lambda} = -\ln(n_0/n).$$

Using $\hat{\lambda}$ and the expected total period demand $\lambda\mu$, we solve for the mean demand size:

$$\hat{\lambda}\hat{\mu} = \bar{x},$$

leading to

$$\hat{\mu} = -\frac{\bar{x}}{\ln(n_0/n)}.$$

It is evident that to calculate these estimates, one only needs to obtain the fraction of periods without demand and the overall mean demand during all periods. Furthermore, whereas the standard MM estimator and the ML estimator have to be derived specifically for every compounding distribution, this new method is generally applicable to any compounding distribution with one parameter. If the compounding distribution contains several parameters, then one can simply add a moment equation for the variance of total demand per period, or if necessary for higher moments. Irrespective of the number of parameters to be estimated, the explicit use of the fraction of periods without demand to estimate the arrival rate always implies that the highest order of moments used by this method is one less than for the standard MM approach.

The simple structure of the arrival rate estimator allows the use of Taylor's theorem to approximate its finite-sample bias. Since N_0 is binomially distributed with parameters n and $\exp(-\lambda)$, we have that $\mathrm{E}[N_0/n] = \exp(-\lambda)$ and $\mathrm{Var}[N_0/n] =$

$\exp(-\lambda)(1-\exp(-\lambda))/n$. Using a second-order Taylor expansion, we find

$$\begin{aligned} \mathrm{E}\left[\hat{\lambda}\right] &\approx \lambda + \frac{\exp(\lambda)-1}{2n}, \\ \mathrm{Var}\left[\hat{\lambda}\right] &\approx \frac{\exp(\lambda)-1}{n}. \end{aligned}$$

It is evident that as $\lambda \to 0$, both the bias and the variance tend to 0, whereas as λ increases, both increase exponentially. This indicates that the best performance of this estimator is achieved for relatively low arrival rates, i.e. for intermittent demand patterns. When the sample size goes to infinity, the bias and variance go to 0, as expected. The correlations of the different sample moments such as the average, variance, and number of periods without demand make it impossible to derive simple, closed-form approximations in a similar fashion for the other estimators, but the numerical study in the next section will aid the comparison of the finite-sample performance of all estimators.

Two scenarios for the observed sample of historical demands deserve special attention. The first is the case where no demands have occurred at all, so that only periods without demand are observed. In practice one will not fit any demand distribution if this is the case, but in (our) numerical experiments with many repetitions it does occasionally occur. All methods are then undefined and to enable a fair comparison, we set in this case the estimated arrival rate to 0 for all methods and do not define an estimate for the demand size. The other extreme scenario is that where no periods have occurred without demand. Again, the practical relevance of this case is rather minor as a compound Poisson process is typically used to describe slow-moving demand, and so one would expect some degree of intermittency in the observed period demands. However, to cope with this unlikely event, we set the new method equal to the standard MM estimator if it does happen in our numerical investigation. The same approach can be used in software implementation: if all observations are non-zero, then use standard MM estimator; else, use new MM estimator

4.5 Numerical study: comparing consistent estimators

To analyze the performance of the different estimators that were discussed in the previous section, we perform a numerical study. For both the geometric and exponential compounding distribution, for various choices of the arrival rate λ and the average demand size μ, and for sample sizes $n = 4, 5, ..., 200$, each time we draw 1,000,000 times a period demand history from the corresponding compound Poisson distribution. So each draw consists of n period demands, each of which may be zero or positive. For each draw, we calculate the parameter estimates for λ and μ using (i) the standard MM estimator, (ii) the ML estimator, and (iii) the proposed alternative MM estimator. We present per parameter combination, per sample size, and per estimator, the average of all obtained estimates. In the next subsection we compare the estimation accuracy of the different estimators. In Section 4.5.2 we discuss the implications for inventory control. Specifically, we compare the achieved fill rates in a base stock inventory model.

4.5.1 Estimation accuracy

We present the averages of the demand parameter estimates for $\lambda = 1/16, 1/4$ and 1. For $\lambda = 1/16$ we initially discuss the scenarios $\mu = 2$ and $\mu = 5$, but thereafter we restrict attention to $\mu = 5$, as the results are very similar. The setting $\lambda = 1/16$ corresponds to a demand occurring on average once every 15 periods, whereas $\lambda = 1$ corresponds to approximately 1 out of 3 periods without demand. These values cover a wide range of intermittent demand patterns similar to those classified in the empirical literature (Syntetos et al., 2012, 2013; Lengu et al., 2014; Turrini and Meissner, 2019). As the results differ only marginally between the geometric and exponential compounding distribution, we focus here on the former, and select one scenario to compare the results with those under the latter.

Figure 4.1 shows the results for $\lambda = 1/16$. This scenario corresponds to a 94% probability of having no demand in a period. All methods yield consistent estimates for both parameters, but show large differences in convergence speeds. Especially the standard MM estimator converges slowly. For $\mu = 2$ and $n = 200$, its arrival rate estimate is still 5% too high, whereas its demand mean estimate is 4% too low. For $\mu = 5$ these relative errors are even larger (9% and 6%, respectively). The proposed

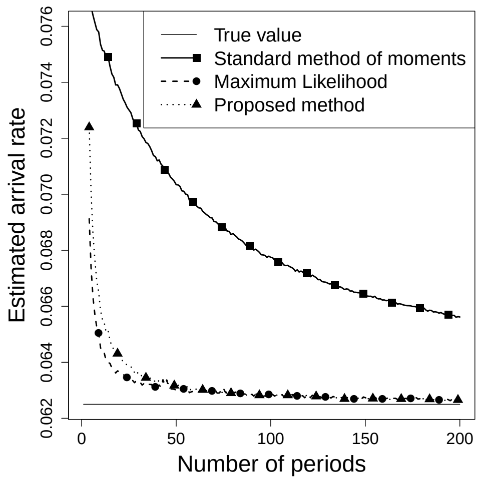

(a) *Estimated arrival rate for* $\mu = 2$

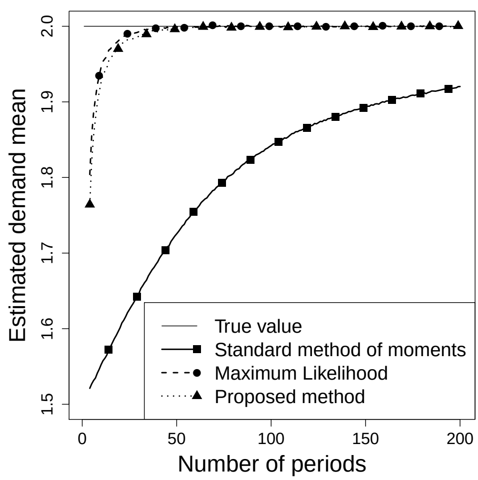

(b) *Estimated demand mean for* $\mu = 2$

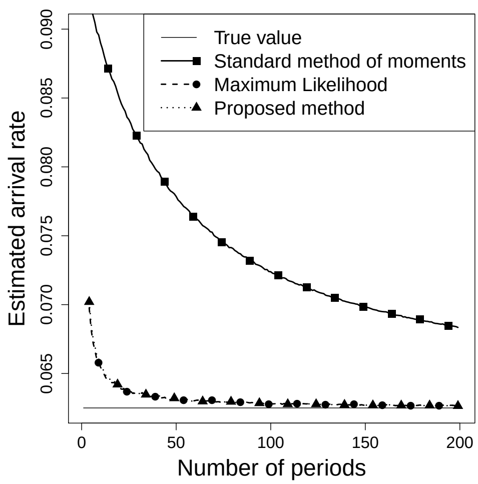

(c) *Estimated arrival rate for* $\mu = 5$

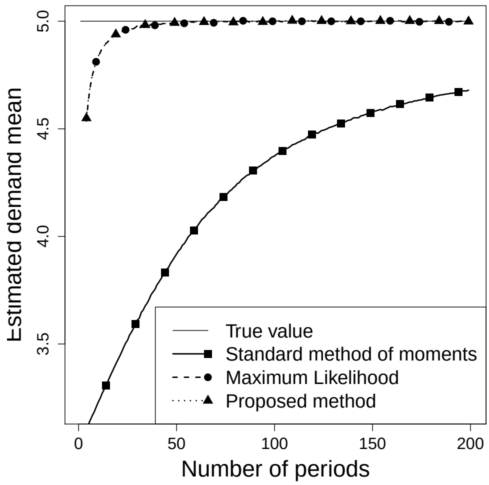

(d) *Estimated demand mean for* $\mu = 5$

Figure 4.1: *Averages of demand parameter estimates. Geometric compounding distribution, true value* $\lambda = 1/16$

method and the ML estimator perform almost identically, and are considerably more accurate than standard MM. For $\mu = 2$ and $n \geq 50$ the difference with the true parameters is below 1%, and also for smaller sample sizes these two methods differ only slightly. For $\mu = 5$ and $n = 25$, both are approximately 17 times more accurate than standard MM in estimating the arrival rate, and 30 times in estimating the demand mean. Interestingly, whereas for $\mu = 2$ the ML estimator performs slightly better than the proposed estimator, their performance is practically identical for $\mu = 5$. We therefore conclude that the proposed estimator is considerably more accurate than the standard MM estimator. The proposed method's performance is comparable to that of the ML estimator, which is considerably more complex to apply.

Next we discuss the results for $\lambda = 1/4$ and $\mu = 5$ (see Figure 4.2). This scenario corresponds to a 78% probability of having no demand in a period. We first discuss the results for the geometric compounding distribution, and then compare them with those for an exponential compounding distribution. We again observe that the standard MM estimator converges slowest, and still shows a significant deviation from the true value for $n = 200$ (3.3% for λ and 1.6% for μ). The proposed method and the ML estimator are again the most accurate and perform almost identically. The arrival rate estimates of ML and the proposed method converge more slowly than their demand mean estimates. It is worthwhile to notice that the estimation accuracy of the three estimators for μ increases more quickly in the true value of the arrival rate than those for λ, whereas the estimation accuracy for μ is already higher than that for λ. This will also affect the inventory control performance, as we will see in Section 4.5.2. For samples of at least 75 observations, they are both within 0.1% of the true demand mean in this scenario, whereas the arrival rate still shows a 0.3% error for $n = 200$. In this case the proposed estimator actually (slightly) outperforms ML. This is an interesting observation of the fact that whereas ML has asymptotically the lowest variance, no guarantees can be given about its performance in finite samples. The results for the exponential distribution are very similar, with the main difference being that the standard MM estimator performs worse in this case (a bias of 4.5% for λ and 2% for μ, for $n = 200$). In conclusion, whereas the standard MM estimator converges more quickly here than in the previous scenario, it is still outperformed considerably by the proposed estimator and ML.

In Figure 4.3, the arrival rate is increased to 1, corresponding to a probability

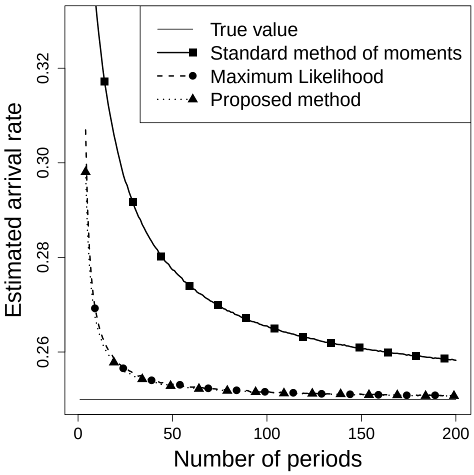

(a) *Geometric distribution, estimated arrival rate*

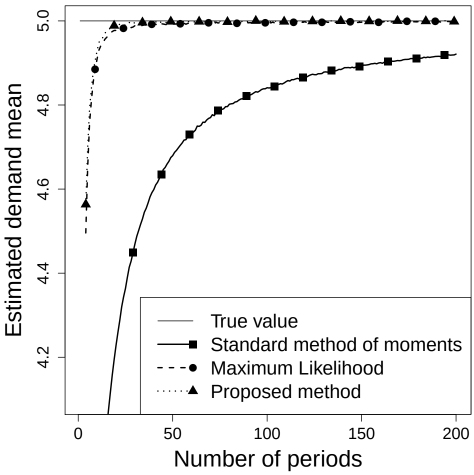

(b) *Geometric distribution, estimated demand mean*

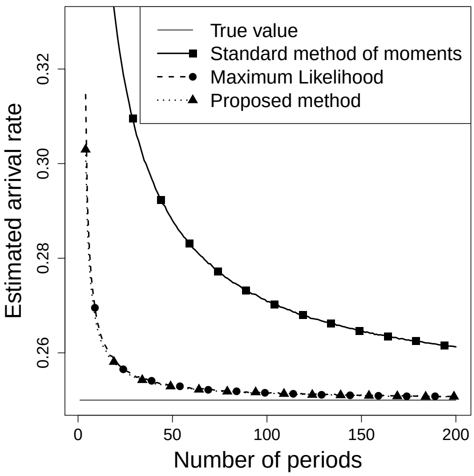

(c) *Exponential distribution, estimated arrival rate*

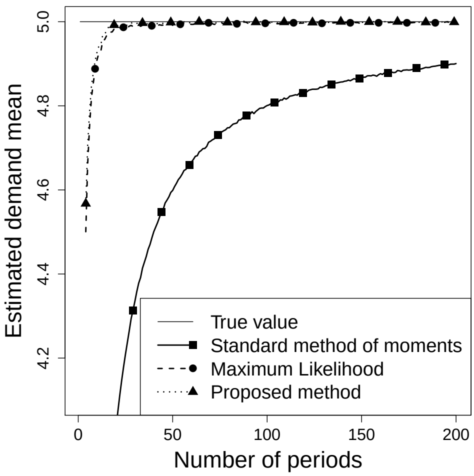

(d) *Exponential distribution, estimated demand mean*

Figure 4.2: *Averages of demand parameter estimates. True values* $\lambda = 1/4$ *and* $\mu = 5$

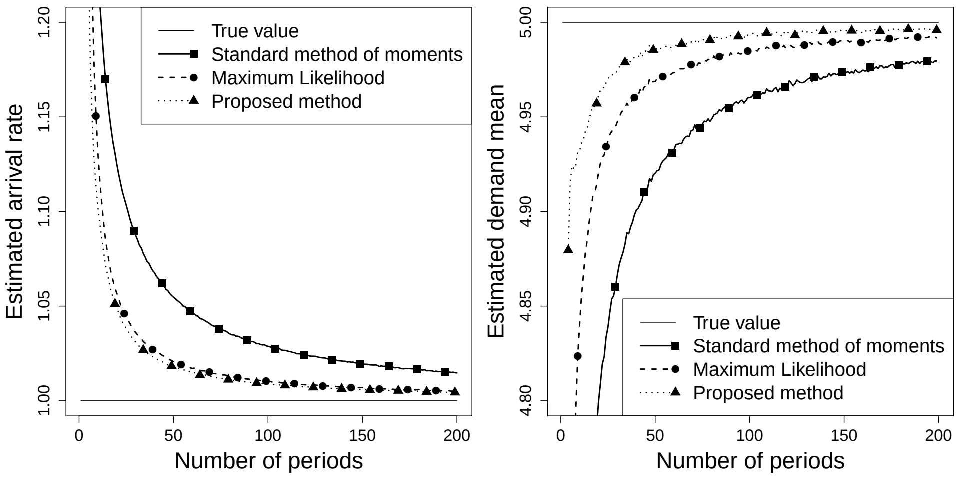

(a) *Estimated arrival rate* **(b)** *Estimated demand mean*

Figure 4.3: *Averages of demand parameter estimates. Geometric compounding distribution, true values $\lambda = 1$ and $\mu = 5$*

of 37% that a period has zero demand. The standard MM estimator still performs worst (with an error of 1.5% in λ and 0.4% in μ for $n = 200$), but the gap with the proposed estimator and ML is smaller than for lower arrival rates. The proposed method again outperforms ML, now by a larger margin. Also in this scenario it is confirmed that when the arrival rate increases, the accuracy of the estimated demand mean increases more quickly than that of the estimated arrival rate.

We conclude that, especially for relatively low values of λ, the standard MM estimator performs very poorly, even for large sample sizes. The new estimator shows fast convergence in all these scenarios, either comparable to or slightly better than the ML estimator. We remark that if λ is larger than approximately 2 (corresponding to a 14% chance of observing a period without demand), then the standard MM estimator has a slight performance advantage over both ML and the proposed estimator. This can be explained by the fact that in such cases the specific moment that the proposed method utilizes for estimating λ (the fraction of periods without demand) becomes very small. Arguably, such scenarios are less relevant in practice, as compound Poisson distributions are most popular for intermittent demand patterns.

4.5.2 Inventory control

For the base stock inventory model that we described in Section 4.3.2, we study the service effects when demand parameters are obtained according to the 3 methods discussed in Section 4.4. We use a target fill rate of 95%. Based on the average estimates for the various methods, for a certain scenario and sample size, we search for the smallest order level S that achieves at least the target fill rate. We report the fill rates that are achieved by these order levels under the true demand parameters. Please note that the achieved fill rate increases monotonically with the order level and thus with the inventory holding costs. Contrary to the previous subsection, we limit our discussion here to the exponential compounding distribution to simplify interpretation of the results, as the target fill rate can be achieved exactly for a continuous demand distribution but not for a discrete one.

Figure 4.4a shows the results for $\lambda = 1/16$, $\mu = 5$, and $L = 2$. Consistently with the results of the previous subsection, the MM and ML estimators all lead to fill rates that converge to the target. The proposed estimator and the ML estimator perform almost identically (in accordance with the results of the previous subsection) and converge much more quickly than the standard MM estimator. The order level (and achieved fill rate) that is set in a given scenario by any of the three methods is driven by the (slight) overestimation of λ (leading to overshooting the order level and fill rate) and the underestimation of μ (leading to an undershoot). In this scenario the underestimation of μ dominates for all three methods, leading to fill rates that converge to their targets from below. Although the overestimation of λ is more severe, this parameter only affects the inventory level distribution (via the lead time demand), whereas μ affects also the distribution of a single customer's demand size, which are both needed for the fill rate computation. Whereas the proposed method and ML are already at the 95% fill rate target for a sample of 50 observations, standard MM only attains 93.8% even for $n = 200$.

An interesting phenomenon occurs when $\lambda = 1/4$ (see Figure 4.4b). In this case, the overestimation of λ and underestimation of μ by the proposed method and ML almost cancel each other out at the order level calculation, leading to a realized fill rate very close to its target over the entire range of sample sizes. Standard MM still leads to fill rates that are approaching their optima from below, and although this

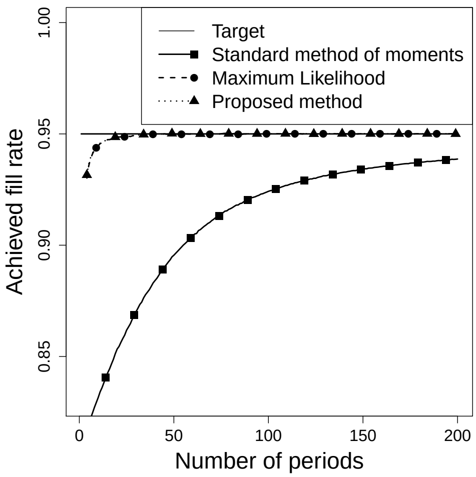

(a) *True values* $\lambda = 1/16$, $\mu = 5$, *and* $L = 2$

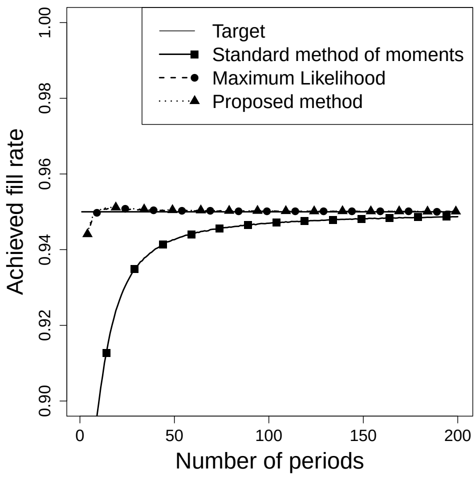

(b) *True values* $\lambda = 1/4$, $\mu = 5$, *and* $L = 2$

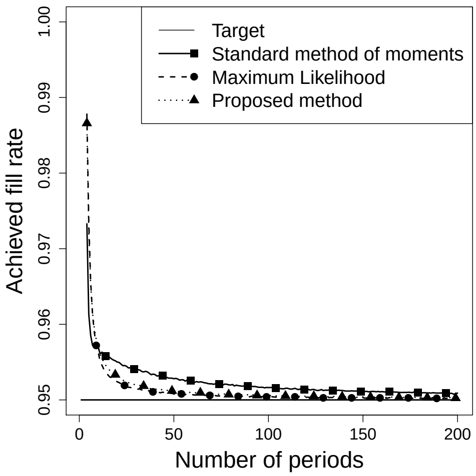

(c) *True values* $\lambda = 1$, $\mu = 5$, *and* $L = 2$

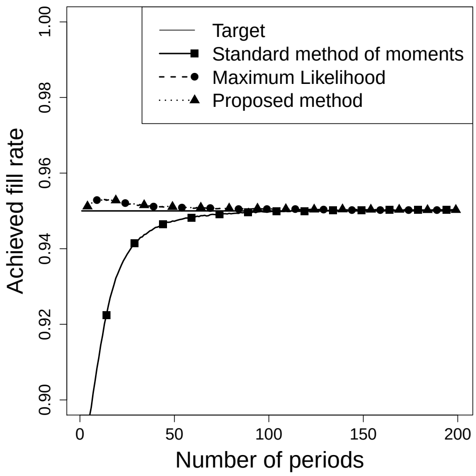

(d) *True values* $\lambda = 1/4$, $\mu = 5$, *and* $L = 8$

Figure 4.4: *Achieved fill rates for an exponential compounding distribution with target fill rate 95%*

method is also more accurate now, it nevertheless converges more slowly than the proposed method and ML.

In Figure 4.4c the arrival rate is increased to 1. All methods now lead to fill rate overshoots rather than undershoots, indicating that the overestimation of λ dominates the fill rate effects. Standard MM still performs worst and ML performs best, but only marginally better than the proposed method. This is an interesting result, as in terms of estimation accuracy ML was actually outperformed by the proposed method. The explanation is also here that the estimation errors in λ and μ counteract each other in the order level calculation. The differences between the three approaches are diminishing.

In the last scenario that we discuss, λ is reset to 1/4, but the lead time is quadrupled to 8 periods. Figure 4.4d shows the results. Comparing this scenario with that where $\lambda = 1/4$ and $L = 2$, we observe that standard MM actually performs better under a longer lead time. This can be explained by the fact that for this parameter setting the underestimation of μ dominates the overestimation of λ, implying that if by extending the lead time more weight is given to the estimation error in λ, then the overall negative effect on the fill rate is counteracted. For ML and the proposed method this is not the case, as for the majority of the studied sample sizes the achieved fill rate was already (slightly) above the target fill rate.

Summarizing the results of all scenarios, we find that the standard MM estimators as applied in the inventory control literature show poor estimation accuracy and slow convergence, especially for low values of the arrival rate, i.e. for highly intermittent demand patterns with relatively many periods without demand. ML and the proposed method yield more accurate estimates, and therefore lead to fill rates that are closer to their targets. Whereas they are similar in convergence speed, the proposed estimator has the advantage that it is available in a simple, closed form. The effect of a method's estimation inaccuracy on the achieved fill rate depends mainly on the proportions of its estimation error in the arrival rate and the mean demand size. For low arrival rates, the proposed method, ML, and standard MM all converge to the optimal fill rate from below, whereas for higher arrival rates this pattern shifts to a convergence from above. In conclusion, as long as the demand pattern shows at least some intermittency, the proposed estimator outperforms standard MM, which is the norm in the inventory control literature. We end this section by remarking that although some particular results (such as the way in which calculations depend on μ and λ) are specific to the fill rate service measure, experiments with different service measures have shown a similar connection between estimation accuracy and achieved service.

4.6 Conclusion

Inventory models using a compound Poisson demand process cannot be applied unless estimates of the demand parameters have been obtained. The focus of the demand forecasting literature is not on obtaining such estimates or on performing inventory control, but on forecasting period or lead time demand. Period demand forecasts cannot be directly transformed into compound Poisson parameter estimates, even if the forecasting procedure explicitly yields separate estimates for the time between periods with positive demands and the average (period) demand size, as Size-Interval methods do. Mis-using these to obtain demand parameters, as several leading commercial software packages do, leads - also asymptotically - to severely biased estimates, overshot fill rates, and therefore excessive inventories and associated costs. It is important to observe that this flaw is not specific to Size-Interval methods, but occurs if any period forecasting procedure is used to estimate individual customer demand parameters. That is, the same inconsistency will occur for example if instead of the demand interval, the probability of a positive period demand is estimated.

The inventory control literature provides, apart from sparse mentions of standard MM estimators, hardly any guidance on obtaining demand parameters from period demand data, although companies typically store historical demand observations per period and software packages also use period demand as input for their forecasts. We have addressed this lack of connection between the forecasting and inventory control literature, and proposed an alternative estimator that takes a sample of period demands as input and uses them to estimate the parameters of the compound Poisson process. The estimator uses the method-of-moments principle, but with a non-standard choice of moments that reflects the intermittent nature of the demand. The fraction of periods without demand is exploited to calculate the customer arrival rate, and this estimate together with the overall average of period demand is used to obtain the mean of an individual demand size. This reflects the intuition of Size-Interval methods in the sense that arrival rate and demand mean are estimated separately.

Our results show that the proposed estimator outperforms the standard MM estimator in terms of convergence speed, and its deviation from the true value of both parameters is smaller over the entire range of sample sizes considered. Furthermore, the proposed estimator performs very similarly to the ML estimator, which does not exist in closed form and has to be found by a numerical search. In some scenarios the proposed estimator even slightly outperforms ML. Using the standard MM estimator can more than double the fraction of demand that cannot be serviced. The results furthermore show that for the proposed method, the estimation errors of the arrival rate and mean demand size are for low arrival rates in such proportions that the target fill rate is approached from below, whereas for medium and high arrival rates it is approached from above. For an arrival rate of 1/4 (corresponding to on average 1 positive demand every 5 periods), both errors almost cancel each other out, leading to achieving very close to the target fill rate for all sample sizes.

A theoretically appealing alternative solution to the studied parameter estimation problem would be to estimate the parameters directly from transactional data (individual demand sizes and times between demands), leading also to unbiased parameter estimates. However, in practice companies typically do not store data at the individual transaction level, because of cost considerations, tractability, and because current demand forecasting and inventory control software packages work with period demands instead. Furthermore, in the demand forecasting literature, temporal aggregation of demand observations is actually recommended. That this is undesirable from an inventory control perspective is not touched upon in that stream of literature. The presented approach can directly be implemented into current software packages, and minimizes the efficiency loss due to periodic data storage compared to the mentioned alternative estimators. It is nevertheless worthwhile to perform an evaluation in which the efficiency of demand parameter estimation and its effect on customer service based on transactional data is compared with the current practice of estimating based on period demand data.

This research could be expanded by studying our proposed method for different compounding distributions, also with multiple parameters. Our estimator will use higher sample moments in such cases, but the highest moment used by our method will always be one less than that of the standard MM estimator. Important empirical research avenues are to compare the new estimator and other methods on real life data sets, and to study the effect of bucketing the demand observations into different period lengths, leading to demand patterns with higher and lower degrees of intermittency, validating the above-described relationship between the arrival rate and achieved fill rate. Finally, an overall connection is lacking in the broader sense between the forecasting literature on the one hand, and the inventory control literature on the other hand. For example, the quality of a certain forecasting method is typically assessed via loss measures such as the Mean Square Error or Mean Absolute Deviation, whereas from an inventory control perspective the resulting cost or service level is important. A study comparing different forecasting methods from this perspective may lead to insights that are different from those currently known in the forecasting literature.