university of groningen / faculty of economics and business

2018010-OPERA

Forecasting and Inventory Control with Compound Poisson Demand Using Periodic Demand Data

July 2018

Dennis Prak
Ruud H. Teunter
Mohamed Z. Babai
Aris A. Syntetos
John E. Boylan

SOM is the research institute of the Faculty of Economics & Business at the University of Groningen. SOM has six programmes:
- Economics, Econometrics and Finance
- Global Economics & Management
- Innovation & Organization
- Marketing
- Operations Management & Operations Research
- Organizational Behaviour

Research Institute SOM
Faculty of Economics & Business
University of Groningen

Visiting address:
Nettelbosje 2
9747 AE Groningen
The Netherlands

Postal address:
P.O. Box 800
9700 AV Groningen
The Netherlands

T +31 50 363 7068/3815

www.rug.nl/feb/research

Forecasting and Inventory Control with Compound Poisson Demand Using Periodic Data

Dennis Prak
University of Groningen, Faculty of Economics and Business, Department of Operations
d.r.j.prak@rug.nl

Ruud H. Teunter
University of Groningen, Faculty of Economics and Business, Department of Operations

Mohamed Z. Babai
Kedge Business School, France

Aris A. Syntetos
Cardiff Business School, Cardiff University, United Kingdom

John E. Boylan
Lancaster University, United Kingdom

Forecasting and Inventory Control with Compound Poisson Demand Using Periodic Demand Data

Dennis Prak*[1], Ruud H. Teunter[1], Mohamed Z. Babai[2], Aris A. Syntetos[3], and John E. Boylan[4]

[1]*Department of Operations, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands*
[2]*Kedge Business School, 680 Cours de la Libération, 33405 Talence Cedex, France*
[3]*Panalpina Centre for Manufacturing and Logistics Research, Cardiff Business School, Cardiff University, Cardiff CF10 3EU, United Kingdom*
[4]*Department of Management Science, Lancaster University, LA1 4YX, United Kingdom*

Abstract

Problem Definition:
Although compound Poisson demand is a popular choice in the inventory control literature and practice, there exists hardly any guidance on obtaining its parameters from real demand data. The forecasting literature focuses on predicting period demand, which does not yield consistent estimates for the parameters of a compound Poisson process. 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 achieved service level.

Academic / Practical Relevance:
The inaccuracies of current parameter estimators make inventory calculations flawed, leading to severe deviations from target service levels and unnecessarily high inventory costs. Current commercial software packages wrongly fit compound Poisson processes based on period demand estimators such as Croston's method.

Methodology:
First, we stipulate the consequences of mis-applying Croston-type methods for estimating compound Poisson demand parameters. Second, we discuss our proposed method. The arrival rate is estimated using the fraction of periods without demand, and this combined with the average of period demand yields an estimate of the mean demand size. Finally, we compare that method to the standard MM estimator and ML in terms of estimation bias and achieved service level.

Results:
Under compound Poisson demand, Croston-based order levels dramatically overshoot the target fill rate. The proposed method outperforms standard MM and performs similarly to ML.

Managerial Implications:
Period forecasting techniques are not suitable for estimating compound Poisson demand parameters, and doing so leads to excessive inventories and thus costs. The proposed method should be preferred over standard MM, which is currently the norm in the inventory control literature.

1 Introduction

One of the most general and widely studied classes of demand models is that of compound Poisson demand. 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 and shown to often provide the best fit to demand data (Syntetos et al. 2012, 2013, Lengu et al. 2014, Turrini and Meissner 2018). Many inventory control models using compound Poisson demand have been developed during the past

*Corresponding author, d.r.j.prak@rug.nl

60 years, starting with Galliher et al. (1959) who considered the special case of stuttering Poisson demand (with Poisson arrivals and geometrically distributed demand sizes), and followed by Feeney and Sherbrooke (1966) and Archibald and Silver (1978), who covered order-up-to policies and (s, S) policies, respectively, for general compound Poisson processes. Compound Poisson demand is a standard choice in textbooks (Silver et al. 2017, Zipkin 2000, Axsäter 2015) and in commercial software (e.g. Forecast Pro, Oracle, SAP APO, SAS, and Slimstock).

Inventory control models with compound Poisson demand typically start from the assumption that the arrival rate and demand size parameters are known, and from there derive the optimal replenishment policy. The problem of how to obtain these parameters from a data set of observed demands is typically not discussed or even mentioned. The forecasting literature remarkably also does not answer this question, as this strand of literature is centered around producing point forecasts of demand per period or during the lead time. This is in line with practice, where demand data is typically stored periodically, for instance per day, week or month, and forecasting software uses period demand as input. However, as we will show, such period forecasting procedures do not provide consistent estimates for the arrival rate and average demand size of individual customers, as they do not distinguish between one large or several smaller orders arriving in a period, but rather look at the resulting total demand. Using them as parameter estimates for customer demand therefore leads to flawed inventory calculations and thereby to deviations from the target service level, even if a long demand history is available. In fact, most of the literature calculates order levels based on (implicit) normality assumptions or other ad-hoc procedures, completely ignoring the actual distribution that was fitted to the demand data.

Croston (1972) proposed to estimate the arrival rate of customers and the average demand size separately, via two separate exponential smoothing equations. This simple method is very popular in the forecasting literature and in practice, praised for its robustness, and many later forecasting papers discuss approaches in the same spirit. All five above-mentioned software packages implement Croston's method to forecast demand and fit compound Poisson demand distributions. However, Croston's method was developed for forecasting period demand, and is therefore not suitable for estimating the parameters of a compound Poisson process, which requires the arrival rate and average size of individual demands. Even though corrections have been suggested, these deal with the bias in inversing the estimated number of periods between two periods with positive demands, rather than with the discrepancy between period demands and individual demands. The first contribution of this paper is a quantification of the asymptotic bias that results from mis-using Croston's period demand interval and size estimates as demand parameters at the individual customer level, and of the consequences on achieved service levels when these parameter estimates are used in inventory calculations.

The few inventory control studies and books that do discuss how compound Poisson demand parameters can be obtained from empirical data suggest standard method-of-moments (MM) estimators. Ward (1978) and Axsäter (2015) do so for the Poisson-geometric (stuttering Poisson) and Poisson-logarithmic (negative binomial) demand distribution, respectively. The class of MM estimators results from setting certain distributional entities - typically central moments of the (full, compound) demand distribution, such as the mean or variance - equal to their sample equivalents, and solving the resulting equations to find expressions for the unknown parameters. This by definition yields consistent estimators (meaning that if the sample of historical demands would be infinitely large, then the true parameter values would be attained), but does not guarantee good estimation accuracy in finite samples of historical demands. In fact, as our results will confirm, using a 'higher' moment of demand per period (for instance its variance) makes the proposed MM estimators sensitive to very small or large (outlier) demand sizes, leading to slow convergence and a substantial estimation bias even for reasonably large sample sizes.

Another approach discussed in the broader statistics literature is that of using maximum likelihood (ML) estimators, which are guaranteed to be consistent and to have the lowest asymptotic variance. However, ML estimators are not guaranteed to perform well in finite samples, and have the serious drawback that they can often only be derived by numerically maximizing the likelihood function. This numerical complexity and lack of a closed-form solution is a possible explanation of why such procedures are unpopular in practice and typically left out in the inventory control and forecasting literature. Furthermore, both (standard) MM and ML procedures lack practical intuition, such as a separation of estimating the arrival rate and the average demand size.

The second and main contribution of this paper is the presentation of an alternative MM estimator that retains Croston's core idea of separately estimating customer arrival rates and demand sizes, but (contrary to Croston's method that estimates period demand parameters) yields consistent estimators for compound Poisson demand parameters from period demand data. The proposed method bases the customer arrival rate on the fraction of periods without demand, and then calculates the expected demand size for which the expected period demand equals the observed average period demand. Contrary to standard MM estimators, no higher moments are involved, making this new method more robust to outliers and contributing to its higher estimation accuracy. Simultaneously, contrary to ML estimators, this new method does have 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 demonstrate for a wide range of intermittent demand scenarios that the proposed method has a better estimation accuracy than the standard MM estimator, and performs similarly to ML (and in some scenarios even slightly better). We study the service implications for the well-known order-up-to (base stock) inventory model when a fill rate service measure is used, and show that the proposed method converges equally fast to the target fill rate as ML, whereas standard MM converges very slowly and leads to order levels that are significantly too low for a wide range of scenarios and sample sizes. The idea of merging the logic and robustness of Croston-like estimators with the consistency of MM estimators relates naturally to demand forecasting and inventory control, but can yield interesting insights in the broader field of statistics.

The remainder of this paper is structured as follows. Section 2 discusses the relevant literature. Section 3 specifies the demand and inventory model and discusses the consequences of mis-using Croston's estimator for estimating compound Poisson demand parameters. Section 4 discusses the standard MM and ML benchmark parameter estimators from the literature and introduces the proposed estimator. Section 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 6 concludes.

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 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 estimation procedures, 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 is given by Axsäter et al. (1994) that replaces compound Poisson distributions by 'equivalent' Poisson distributions. Further extensions of inventory control models with compound Poisson demand are 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. Next to this, 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 2018). All conclude 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 Rao (2003), Axsäter (2003), Liu and Song (2012), Song and Zipkin (2013), and Shi et al. (2014).

Compound Poisson demand processes are obviously characterized by the demand rate (or degree of intermittency) 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 do so, twice applying exponential smoothing. Many others continued this research, as reviewed by Gardner (2006). Several papers show that Croston's method outperforms other simple methods such as single exponential smoothing (Ghobbar and Friend 2003, Eaves and Kingsman 2004, 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 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 and Boylan (2005) for Bernoulli period demand arrivals. For Poisson demand arrivals, Shale et al. (2006) derived an approximate mark-up term which was later proven incorrect and improved by Syntetos et al. (2015b). However, we stress that despite these corrections, the fundamental issue remains that Croston's method estimates period demand, and thus does not yield arrival rate and demand size estimates at the individual customer level. We stress that this discrepancy is not an error of Croston's method or 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 unaddressed. Most authors ignore 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) assumes an Erlang distribution of lead-time demand, while Schultz (1987) argues that traditional inventory control and forecasting methods are ill-suited for sporadic demands and uses a Croston-like forecasting procedure. However, he subsequently calculates the order level based on normality assumptions. Dunsmuir and Snyder (1989) use a more realistic demand distribution with a positive probability mass at zero to calculate the order level. Several more recent articles follow the logic of Schultz (1987) and forecast lead-time demand, estimate the forecast error, and subsequently implicitly assume normality to compute order levels. This includes empirical studies (Syntetos and Boylan 2006) and numerical studies where demands are explicitly generated from compound Poisson distributions (Syntetos et al. 2009, Teunter and Sani 2009). Altay et al. (2012) uses a power approximation, and Sani and Kingsman (1997) compare various approximate methods. Given the fact that demand data is almost always stored periodically, it is quite natural that so many authors directly estimate period or lead time demand rather than the underlying demand process. However, this has an important drawback. The (estimated) lead time demand distribution often does not allow exact inventory optimization, as for instance the so-called undershoot depends on individual demand sizes. Indeed, many of the reviewed optimization procedures for inventories under compound Poisson demand require estimates of the demand process parameters themselves.

Very few authors discuss how the unknown parameters of compound Poisson demand (or of demand processes in general) can be obtained from empirical demand data. Those who do, suggest the use of standard MM estimators. Ward (1978) gives 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) derives standard MM estimators for the Poisson-exponential distribution, albeit not for inventory control but for rainfall prediction. However, these estimators lack intuition and are not guaranteed to perform well in finite samples. 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 per period demand parameters. 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 furthermore study its inventory control performance. The demand model and inventory context are formally introduced in Section 3, where we also quantify the consequences of mis-using Croston's method as an estimator for compound Poisson demand parameters.

3 Demand Model, Inventory Model, and Croston's Method

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 Croston's period estimates as compound Poisson demand parameters at the customer level, quantify the resulting asymptotic bias of the parameter estimates and derive their effects on the achieved fill rate.

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 $\mu = 1/\theta$.

3.2 Inventory Model

We study a base stock inventory control model with a positive lead time of L periods, under a fill rate constraint. Base-stock inventory models are widely used in practice, as they are relatively easy to optimize (since they contain only one unknown parameter) and apply to any case where fixed ordering costs can be ignored. This holds for instance for most relatively expensive and slow moving service parts. The fill rate (the fraction of demand that can be serviced from on-hand inventory) 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

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

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.

3.3 Mis-application of Croston's Method

Croston's method estimates the time between two periods with positive demand and the size of a nonzero period demand both by exponential smoothing. Formally, given that we are currently at time t, observe a positive demand x_t, and have observed the previous positive demand q periods ago, Croston's method sets

$$\hat{\mu}_t = \alpha_1 x_t + (1 - \alpha_1)\hat{\mu}_{t-1},$$
$$\hat{p}_t = \alpha_2 q + (1 - \alpha_2)\hat{p}_{t-1},$$

where $0 < \alpha_1 < 1$ and $0 < \alpha_2 < 1$ are smoothing constants. If Croston's period estimates are wrongly taken as the estimates of the parameters of the demand distribution, then $\hat{\mu}_t$ is the estimate at time t for the mean μ of the compounding distribution, whereas $1/\hat{p}_t$ is the estimate for λ. As Syntetos et al. (2015b) show, using $1/\hat{p}_t$ as an estimate for λ leads to an inversion bias, as $\mathrm{E}[1/\hat{p}_t] \neq 1/\mathrm{E}[\hat{p}_t]$. Syntetos et al. (2015b) approximately correct for this bias using a second-order Taylor expansion, and derive the estimator $(1-\alpha_2/2)/\hat{p}_t$ 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 $\hat{\mu}_t$, 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,

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

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, $\hat{p}_t$ 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 $\mathrm{E}[\hat{p}_t] \to 1/(1-\exp(-\lambda))$ as $n \to \infty$, we can only approximate the limit of $\mathrm{E}[1/\hat{p}_t]$. We do so by a second-order Taylor expansion and find, using the mean and variance of the geometric distribution and the fact that $\hat{p}_t$ is an exponential smoothing estimator with asymptotic variance $\mathrm{Var}(Q)\alpha_2/(2-\alpha_2)$, that

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

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.

Table 1 quantifies the (approximate) asymptotic bias resulting from using either standard Croston's method, or the corrected version by Syntetos et al. (2015b) ('SBL'), 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.

Note from Table 1 that none of the results depend on μ, except for the achieved fill rates under the geometric compounding distribution. The reason is that for discrete 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. 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 average demand size is always overestimated. All asymptotic biases are increasing in λ and μ. 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 corrected 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 5.2. The achieved fill rates are increasing in λ, and for $\lambda = 1$, which is not an unrealistically high value, they are in the range of 98 to 99%, depending on the variant and compounding distribution. This implies that the inventory levels and associated costs are dramatically too high.

Table 1: *Asymptotic Biases and Achieved Fill Rates when Using Croston Type Methods for Compound Poisson Parameter Estimation*

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

The conclusion of this section is that mis-applying Croston's method to estimate compound Poisson demand parameters leads to severely overestimated average demand sizes, significantly overshot fill rates and correspondingly excessive inventory costs. For the special case of pure Poisson demand (where every customer has a fixed demand size of 1), Croston's period arrival rate and demand mean estimates are still inconsistent estimates for the arrival rate and mean demand size of individual customers. However, the estimate that Croston's method yields for the overall mean of period demand (including periods without demands), is $\hat{\mu}_t/\hat{p}_t$, which does converge to the true arrival rate of that Poisson distribution. Hence, in this special case there exists a way to convert Croston's period estimates into a consistent estimate of the demand parameter, whereas this is generally not possible.

4 Consistent Estimators

Having discussed the mis-application of Croston's method 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.

4.1 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_i 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.

4.2 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

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

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.

4.3 The Proposed Method

We propose an alternative MM estimator and choose its moments such that the intuition of Croston's method - explicitly separating the 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 to compare 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.

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 5.2 we discuss the implications for inventory control. Specifically, we compare the achieved fill rates in a base stock inventory model.

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 2018). 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 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 largely

Figure 1: *Averages of Demand Parameter Estimates*
Geometric Compounding Distribution, True Value $\lambda = 1/16$

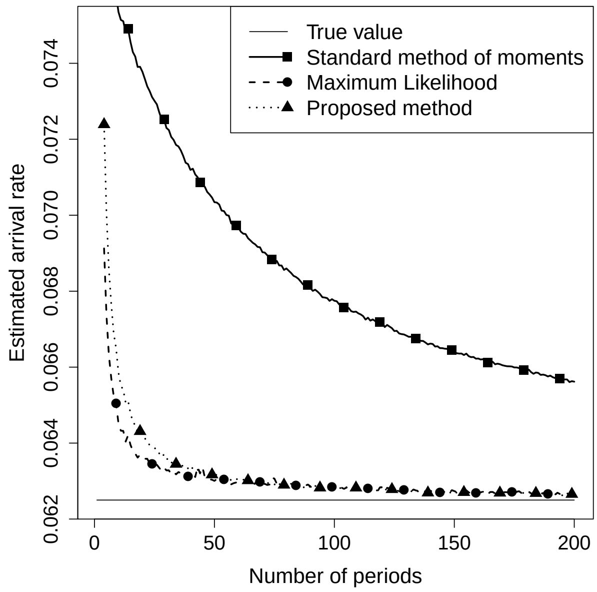

(a) *Estimated Arrival Rate for* $\mu = 2$

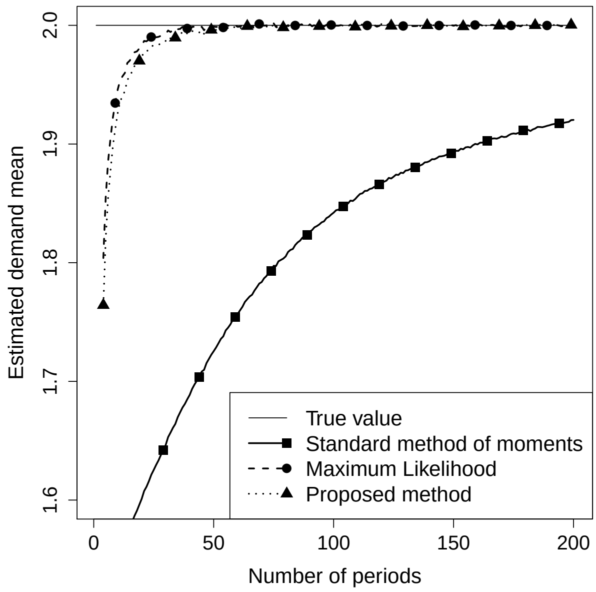

(b) *Estimated Demand Mean for* $\mu = 2$

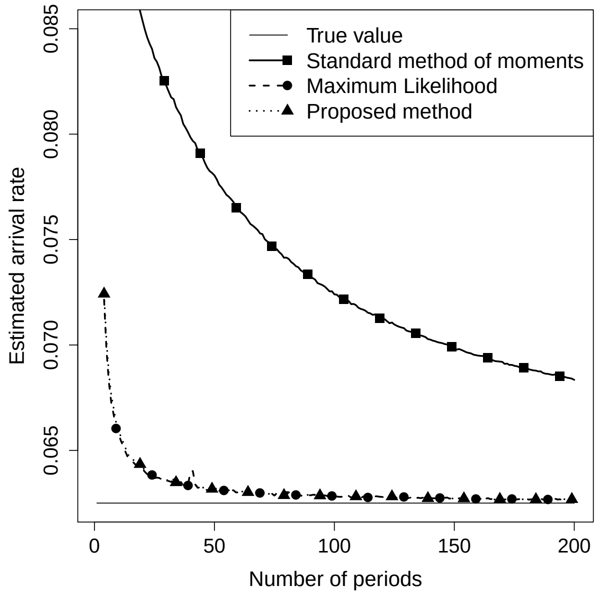

(c) *Estimated Arrival Rate for* $\mu = 5$

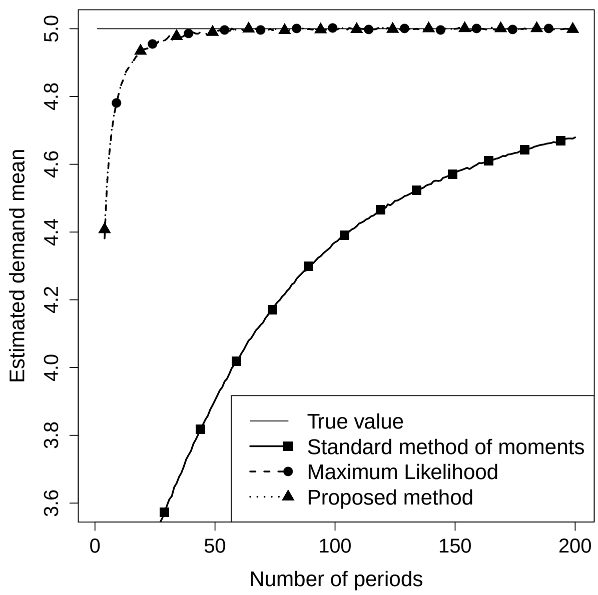

(d) *Estimated Demand Mean for* $\mu = 5$

Figure 2: *Averages of Demand Parameter Estimates,*
True Values $\lambda = 1/4$, $\mu = 5$

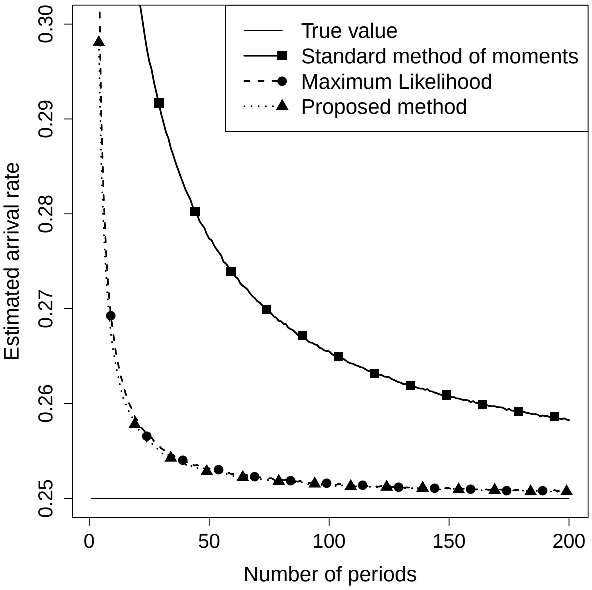

(a) *Geometric Compounding Distribution,*
Estimated Arrival Rate

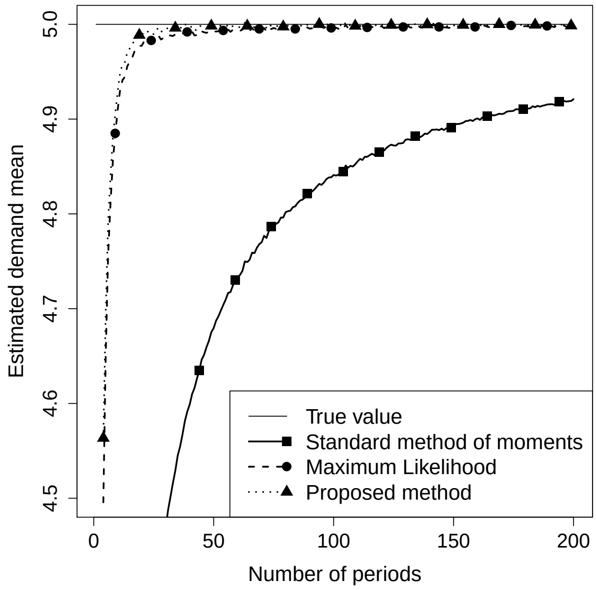

(b) *Geometric Compounding Distribution,*
Estimated Demand Mean

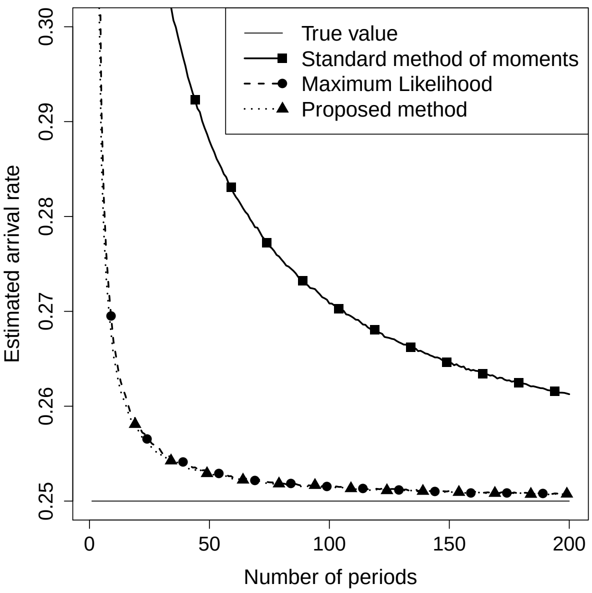

(c) *Exponential Compounding Distribution,*
Estimated Arrival Rate

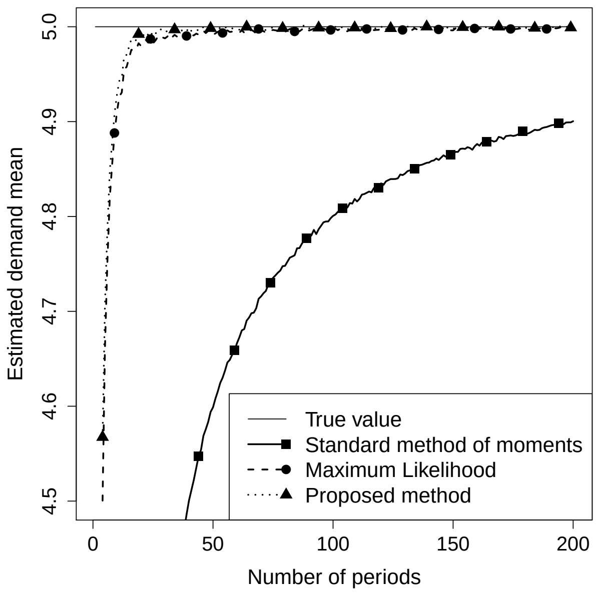

(d) *Exponential Compounding Distribution,*
Estimated Demand Mean

Figure 3: *Averages of Demand Parameter Estimates*
Geometric Compounding Distribution, True Values $\lambda = 1$, $\mu = 5$

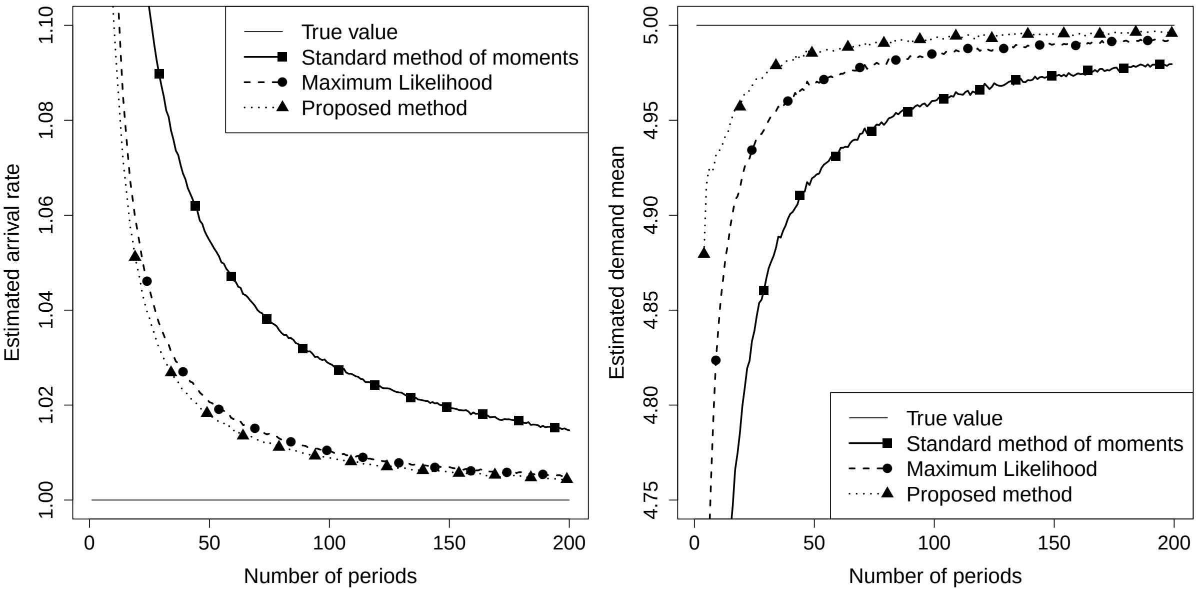

(a) *Estimated Arrival Rate* **(b)** *Estimated Demand Mean*

different 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 method and the ML estimator perform almost identically, and are considerably more accurate than ML. 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 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 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 3, the arrival rate is increased to 1, corresponding to a probability 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.

5.2 Inventory Control

For the base stock inventory model that we described in Section 3.2, we study the service effects when demand parameters are obtained according to the 3 methods discussed in Section 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 this 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 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 this as well and 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 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 also this method is more accurate now, it nevertheless converges more slowly than the proposed method and ML.

In Figure 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 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 val-

Figure 4: *Achieved Fill Rates for an Exponential Compounding Distribution with Target Fill Rate 95%*

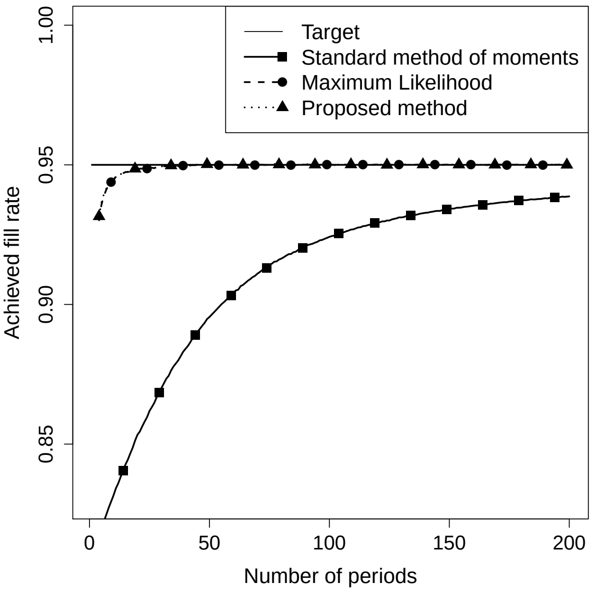

(a) *True Values* $\lambda = 1/16$, $\mu = 5$, *and* $L = 2$

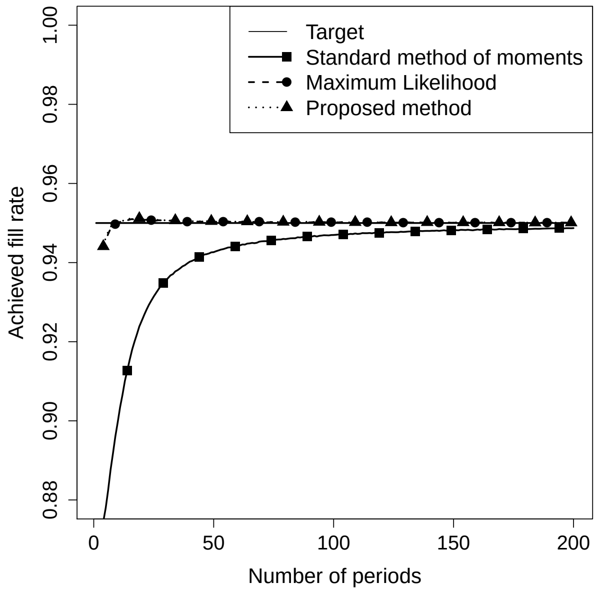

(b) *True Values* $\lambda = 1/4$, $\mu = 5$, *and* $L = 2$

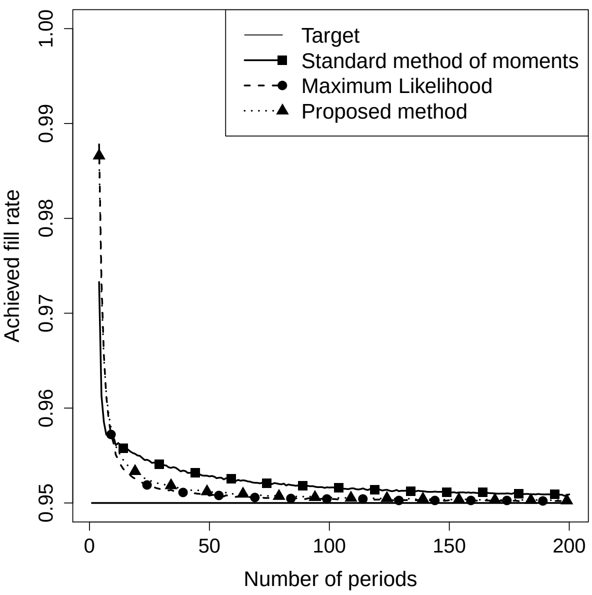

(c) *True Values* $\lambda = 1$, $\mu = 5$, *and* $L = 2$

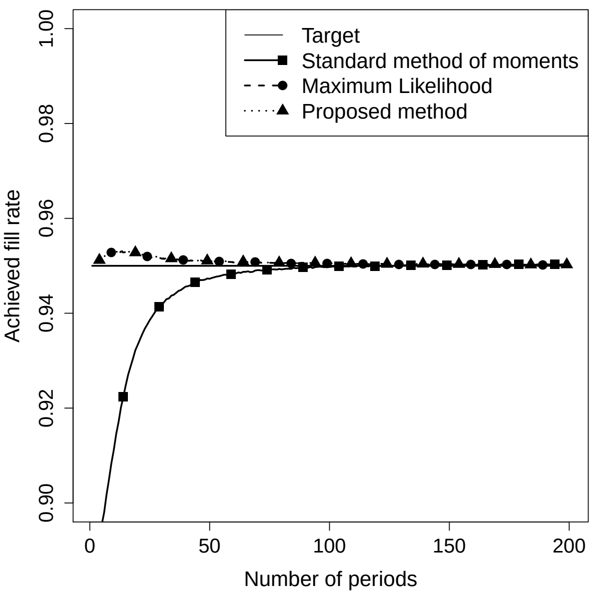

(d) *True Values* $\lambda = 1/4$, $\mu = 5$, *and* $L = 8$

ues 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.

6 Conclusions

Inventory models using a compound Poisson demand distribution cannot be applied until (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 methods based on Croston (1972) do. Mis-using Croston's period estimates 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.

The inventory control literature provides, apart from sparse mentions of standard MM estimators, hardly any guidance on obtaining demand parameters from periodic demand data, although companies typically store historical demand observations per period (e.g. day, week, or month), 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 distribution. 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 mean of period demand is used to obtain the mean of an individual demand size. This reflects the intuition of Croston's method in the sense that arrival rate and demand mean are estimated separately.

Our results show that, for the distributions examined, and for intermittent demand, the proposed estimator outperforms 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. 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 high accuracy of the proposed method also implies that the performance loss due to storing demands periodically rather than individually is minimal if this estimator is used.

The major managerial insights of this research are twofold. Firstly, despite current practice, standard forecasting techniques for period demand should not be used to obtain compound Poisson parameter estimates for performing inventory control. This leads to drastically excessive inventory levels and costs, and overshooting of target fill rates. Secondly, order levels (and thereby average inventories) can be reduced by using the proposed estimator. A decision rule results that is easily implementable in software packages: one should select the proposed estimator if the observed demand pattern is intermittent, and standard MM otherwise.

This research could be expanded by studying our proposed method for different compounding distributions, also with more than one parameter. 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 in different period lengths, leading to demand patterns with higher and lower degrees of intermittency. 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.

References

Altay, N., L. A. Litteral, F. Rudisill. 2012. Effects of correlation on intermittent demand forecasting and stock control. *International Journal of Production Economics* **135**(1) 275–283.

Anscombe, F. J. 1950. Sampling theory of the negative binomial and logarithmic series distributions. *Biometrika* **37**(3/4) 358–382.

Archibald, B. C., E. A. Silver. 1978. (s, S) policies under continuous review and discrete compound Poisson demand. *Management Science* **24**(9) 899–909.

Axsäter, S. 1990. Simple solution procedures for a class of two-echelon inventory problems. *Operations Research* **38**(1) 64–69.

Axsäter, S. 2000. Exact analysis of continuous review (R, Q) policies in two-echelon inventory systems with compound poisson demand. *Operations Research* **48**(5) 686–696.

Axsäter, S. 2003. A new decision rule for lateral transshipments in inventory systems. *Management Science* **49**(9) 1168–1179.

Axsäter, S. 2015. *Inventory control*. 3rd ed. Springer International Publishing, Basel.

Axsäter, S., R. Forsberg, W. F. Zhang. 1994. Approximating general multi-echelon inventory systems by Poisson models. *International Journal of Production Economics* **35**(1) 201–206.

Babai, M. Z., Z. Jemai, Y. Dallery. 2011. Analysis of order-up-to-level inventory systems with compound Poisson demand. *European Journal of Operational Research* **210**(3) 552–558.

Balakrishnan, N., Y. YE, R.A. AL-Jarallah. 2017. Likelihood inference for type I bivariate Pólya-Aeppli distribution. *Communications in Statistics* **46**(8) 6436–6453.

Bensoussan, A., R. H. Liu, S. P. Sethi. 2005. Optimality of an (s, S) policy with compound Poisson and diffusion demands: A quasi-variational inequalities approach. *SIAM Journal on Control and Optimization* **44**(5) 1650–1676.

Chen, J., P. L. Jackson, J. A. Muckstadt. 2011. Exact analysis of a lost sales model under stuttering Poisson demand. *Operations Research* **59**(1) 249–253.

Cheung, K. L. 1996. On the $(S-1, S)$ inventory model under compound Poisson demands and i.i.d. unit resupply times. *Naval Research Logistics (NRL)* **43**(4).

Croston, J. D. 1972. Forecasting and stock control for intermittent demands. *Operational Research Quarterly* **23**(3) 289–303.

Dunsmuir, W. T. M., R. N. Snyder. 1989. Control of inventories with intermittent demand. *European Journal of Operational Research* **40**(1) 16–21.

Eaves, A. H. C., B. G. Kingsman. 2004. Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society* **55**(4) 431–437.

Ehrenberg, A. S. C. 1959. The pattern of consumer purchases. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **8**(1) 26–41.

Feeney, G. J., C. C. Sherbrooke. 1966. The $(S-1,\ S)$ inventory policy under compound Poisson demand. *Management Science* **12**(5) 391–411.

Forsberg, R. 1995. Optimization of order-up-to-S policies for two-level inventory systems with compound Poisson demand. *European Journal of Operational Research* **81**(1) 143–153.

Galliher, H. P., P. M. Morse, M. Simond. 1959. Dynamics of two classes of continuous-review inventory systems. *Operations Research* **7**(3) 362–384.

Gardner, E. S. 2006. Exponential smoothing: The state of the art - part II. *International Journal of Forecasting* **22**(4) 637–666.

Ghobbar, A. A., C. H. Friend. 2003. Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model. *Computers & Operations Research* **30**(14) 2097–2114.

Graves, S. C. 1985. A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Science* **31**(10) 1247–1256.

Gutierrez, R.l S., A. O. Solis, S. Mukhopadhyay. 2008. Lumpy demand forecasting using neural networks. *International Journal of Production Economics* **111**(2) 409–420.

Lengu, D., A. A. Syntetos, M. Z. Babai. 2014. Spare parts management: Linking distributional assumptions to demand classification. *European Journal of Operational Research* **235**(3) 624–635.

Liu, F., J. Song. 2012. Good and bad news about the (S, T) policy. *Manufacturing & Service Operations Management* **14**(1) 42–49.

Öztürk, A. 1981. On the study of a probability distribution for precipitation totals. *Journal of Applied Meteorology* **20**(12) 1499–1505.

Rao, U. S. 2003. Properties of the periodic review (R, T) inventory control policy for stationary, stochastic demand. *Manufacturing & Service Operations Management* **5**(1) 37–53.

Sani, B., B. G. Kingsman. 1997. Selecting the best periodic inventory control and demand forecasting methods for low demand items. *Journal of the Operational Research Society* **48**(7) 700–713.

Savani, V., A. A. Zhigljavsky. 2006. Efficient estimation of parameters of the negative binomial distribution. *Communications in Statistics - Theory and Methods* **35**(5) 767–783.

Schultz, C. R. 1987. Forecasting and inventory control for sporadic demand under periodic review. *Journal of the Operational Research Society* **38**(5) 453–458.

Shale, E. A., J. E. Boylan, F. R. Johnston. 2006. Forecasting for intermittent demand: the estimation of an unbiased average. *Journal of the Operational Research Society* **57**(5) 588–592.

Sherbrooke, C. C. 1968. Metric: A multi-echelon technique for recoverable item control. *Operations Research* **16**(1) 122–141.

Shi, J., M. N. Katehakis, B. Melamed, Y. Xia. 2014. Production-inventory systems with lost sales and compound poisson demands. *Operations Research* **62**(5) 1048–1063.

Silver, E. A., D. F. Pyke, D. J. Thomas. 2017. *Inventory and production management in supply chains*. 4th ed. CRC Press, Boca Raton.

Snyder, R. D., J. K. Ord, A. Beaumont. 2012. Forecasting the intermittent demand for slow-moving inventories: A modelling approach. *International Journal of Forecasting* **28**(2) 485–496.

Song, J., P. Zipkin. 2013. Supply streams. *Manufacturing & Service Operations Management* **15**(3) 444–457.

Syntetos, A. A., M. Z. Babai, N. Altay. 2012. On the demand distributions of spare parts. *International Journal of Production Research* **50**(8) 2101–2117.

Syntetos, A. A., M. Z. Babai, Y. Dallery, R. H. Teunter. 2009. Periodic control of intermittent demand items: theory and empirical analysis. *Journal of the Operational Research Society* **60**(5) 611–618.

Syntetos, A. A., M. Z. Babai, E. S. Gardner. 2015a. Forecasting intermittent inventory demands: simple parametric methods vs. bootstrapping. *Journal of Business Research* **68**(8) 1746–1752.

Syntetos, A. A., M. Z. Babai, S. Luo. 2015b. Forecasting of compound Erlang demand. *Journal of the Operational Research Society* **66**(12) 2061–2074.

Syntetos, A. A., J. E. Boylan. 2001. On the bias of intermittent demand estimates. *International Journal of Production Economics* **71**(1) 457–466.

Syntetos, A. A., J. E. Boylan. 2005. The accuracy of intermittent demand estimates. *International Journal of Forecasting* **21**(2) 303–314.

Syntetos, A. A., J. E. Boylan. 2006. On the stock control performance of intermittent demand estimators. *International Journal of Production Economics* **103**(1) 36–47.

Syntetos, A. A., D. Lengu, M. Z. Babai. 2013. A note on the demand distributions of spare parts. *International Journal of Production Research* **51**(21) 6356–6358.

Teunter, R. H., L. Duncan. 2009. Forecasting intermittent demand: a comparative study. *Journal of the Operational Research Society* **60**(3) 321–329.

Teunter, R. H., B. Sani. 2009. Calculating order-up-to levels for products with intermittent demand. *International Journal of Production Economics* **118**(1) 82–86.

Thomas, D. J. 2005. Measuring item fill-rate performance in a finite horizon. *Manufacturing & Service Operations Management* **7**(1) 74–80.

Turrini, L., J. Meissner. 2018. Spare parts inventory management: new evidence from distribution fitting. *European Journal of Operational Research* doi:https://doi.org/10.1016/j.ejor.2017.09.039. In press.

Ward, J. B. 1978. Determining reorder points when demand is lumpy. *Management Science* **24**(6) 623–632.

Watson, R. B. 1987. The effects of demand-forecast fluctuations on customer service and inventory cost when demand is lumpy. *Journal of the Operational Research Society* **38**(1) 75–82.

Willemain, T. R., C. N. Smart, H. F. Schwarz. 2004. A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting* **20**(3) 375–387.

Zipkin, P. H. 2000. *Foundations of inventory management*. 1st ed. McGraw-Hill, New York.

List of research reports

13001-EEF: Kuper, G.H. and M. Mulder, Cross-border infrastructure constraints, regulatory measures and economic integration of the Dutch – German gas market

13002-EEF: Klein Goldewijk, G.M. and J.P.A.M. Jacobs, The relation between stature and long bone length in the Roman Empire

13003-EEF: Mulder, M. and L. Schoonbeek, Decomposing changes in competition in the Dutch electricity market through the Residual Supply Index

13004-EEF: Kuper, G.H. and M. Mulder, Cross-border constraints, institutional changes and integration of the Dutch – German gas market

13005-EEF: Wiese, R., Do political or economic factors drive healthcare financing privatisations? Empirical evidence from OECD countries

13006-EEF: Elhorst, J.P., P. Heijnen, A. Samarina and J.P.A.M. Jacobs, State transfers at different moments in time: A spatial probit approach

13007-EEF: Mierau, J.O., The activity and lethality of militant groups: Ideology, capacity, and environment

13008-EEF: Dijkstra, P.T., M.A. Haan and M. Mulder, The effect of industry structure and yardstick design on strategic behavior with yardstick competition: an experimental study

13009-GEM: Hoorn, A.A.J. van, Values of financial services professionals and the global financial crisis as a crisis of ethics

13010-EEF: Boonman, T.M., Sovereign defaults, business cycles and economic growth in Latin America, 1870-2012

13011-EEF: He, X., J.P.A.M Jacobs, G.H. Kuper and J.E. Ligthart, On the impact of the global financial crisis on the euro area

13012-GEM: Hoorn, A.A.J. van, Generational shifts in managerial values and the coming of a global business culture

13013-EEF: Samarina, A. and J.E. Sturm, Factors leading to inflation targeting – The impact of adoption

13014-EEF: Allers, M.A. and E. Merkus, Soft budget constraint but no moral hazard? The Dutch local government bailout puzzle

13015-GEM: Hoorn, A.A.J. van, Trust and management: Explaining cross-national differences in work autonomy

13016-EEF: Boonman, T.M., J.P.A.M. Jacobs and G.H. Kuper, Sovereign debt crises in Latin America: A market pressure approach

13017-GEM: Oosterhaven, J., M.C. Bouwmeester and M. Nozaki, The impact of production and infrastructure shocks: A non-linear input-output programming approach, tested on an hypothetical economy

13018-EEF: Cavapozzi, D., W. Han and R. Miniaci, Alternative weighting structures for multidimensional poverty assessment

14001-OPERA: Germs, R. and N.D. van Foreest, Optimal control of production-inventory systems with constant and compound poisson demand

14002-EEF: Bao, T. and J. Duffy, Adaptive vs. eductive learning: Theory and evidence

14003-OPERA: Syntetos, A.A. and R.H. Teunter, On the calculation of safety stocks

14004-EEF: Bouwmeester, M.C., J. Oosterhaven and J.M. Rueda-Cantuche, Measuring the EU value added embodied in EU foreign exports by consolidating 27 national supply and use tables for 2000-2007

14005-OPERA: Prak, D.R.J., R.H. Teunter and J. Riezebos, Periodic review and continuous ordering

14006-EEF: Reijnders, L.S.M., The college gender gap reversal: Insights from a life-cycle perspective

14007-EEF: Reijnders, L.S.M., Child care subsidies with endogenous education and fertility

14008-EEF: Otter, P.W., J.P.A.M. Jacobs and A.H.J. den Reijer, A criterion for the number of factors in a data-rich environment

14009-EEF: Mierau, J.O. and E. Suari Andreu, Fiscal rules and government size in the European Union

14010-EEF: Dijkstra, P.T., M.A. Haan and M. Mulder, Industry structure and collusion with uniform yardstick competition: theory and experiments

14011-EEF: Huizingh, E. and M. Mulder, Effectiveness of regulatory interventions on firm behavior: a randomized field experiment with e-commerce firms

14012-GEM: Bressand, A., Proving the old spell wrong: New African hydrocarbon producers and the 'resource curse'

14013-EEF: Dijkstra P.T., Price leadership and unequal market sharing: Collusion in experimental markets

14014-EEF: Angelini, V., M. Bertoni, and L. Corazzini, Unpacking the determinants of life satisfaction: A survey experiment

14015-EEF: Heijdra, B.J., J.O. Mierau, and T. Trimborn, Stimulating annuity markets

14016-GEM: Bezemer, D., M. Grydaki, and L. Zhang, Is financial development bad for growth?

14017-EEF: De Cao, E. and C. Lutz, Sensitive survey questions: measuring attitudes regarding female circumcision through a list experiment

14018-EEF: De Cao, E., The height production function from birth to maturity

14019-EEF: Allers, M.A. and J.B. Geertsema, The effects of local government amalgamation on public spending and service levels. Evidence from 15 years of municipal boundary reform

14020-EEF: Kuper, G.H. and J.H. Veurink, Central bank independence and political pressure in the Greenspan era

14021-GEM: Samarina, A. and D. Bezemer, Do Capital Flows Change Domestic Credit Allocation?

14022-EEF: Soetevent, A.R. and L. Zhou, Loss Modification Incentives for Insurers Under ExpectedUtility and Loss Aversion

14023-EEF: Allers, M.A. and W. Vermeulen, Fiscal Equalization, Capitalization and the Flypaper Effect.

14024-GEM: Hoorn, A.A.J. van, Trust, Workplace Organization, and Comparative Economic Development.

14025-GEM: Bezemer, D., and L. Zhang, From Boom to Bust in de Credit Cycle: The Role of Mortgage Credit.

14026-GEM: Zhang, L., and D. Bezemer, How the Credit Cycle Affects Growth: The Role of Bank Balance Sheets.

14027-EEF: Bružikas, T., and A.R. Soetevent, Detailed Data and Changes in Market Structure: The Move to Unmanned Gasoline Service Stations.

14028-EEF: Bouwmeester, M.C., and B. Scholtens, Cross-border Spillovers from European Gas Infrastructure Investments.

14029-EEF: Lestano, and G.H. Kuper, Correlation Dynamics in East Asian Financial Markets.

14030-GEM: Bezemer, D.J., and M. Grydaki, Nonfinancial Sectors Debt and the U.S. Great Moderation.

14031-EEF: Hermes, N., and R. Lensink, Financial Liberalization and Capital Flight: Evidence from the African Continent.

14032-OPERA: Blok, C. de, A. Seepma, I. Roukema, D.P. van Donk, B. Keulen, and R. Otte, Digitalisering in Strafrechtketens: Ervaringen in Denemarken, Engeland, Oostenrijk en Estland vanuit een Supply Chain Perspectief.

14033-OPERA: Olde Keizer, M.C.A., and R.H. Teunter, Opportunistic condition-based maintenance and aperiodic inspections for a two-unit series system.

14034-EEF: Kuper, G.H., G. Sierksma, and F.C.R. Spieksma, Using Tennis Rankings to Predict Performance in Upcoming Tournaments

15001-EEF: Bao, T., X. Tian, X. Yu, Dictator Game with Indivisibility of Money

15002-GEM: Chen, Q., E. Dietzenbacher, and B. Los, The Effects of Ageing and Urbanization on China's Future Population and Labor Force

15003-EEF: Allers, M., B. van Ommeren, and B. Geertsema, Does intermunicipal cooperation create inefficiency? A comparison of interest rates paid by intermunicipal organizations, amalgamated municipalities and not recently amalgamated municipalities

15004-EEF: Dijkstra, P.T., M.A. Haan, and M. Mulder, Design of Yardstick Competition and Consumer Prices: Experimental Evidence

15005-EEF: Dijkstra, P.T., Price Leadership and Unequal Market Sharing: Collusion in Experimental Markets

15006-EEF: Anufriev, M., T. Bao, A. Sutin, and J. Tuinstra, Fee Structure, Return Chasing and Mutual Fund Choice: An Experiment

15007-EEF: Lamers, M., Depositor Discipline and Bank Failures in Local Markets During the Financial Crisis

15008-EEF: Oosterhaven, J., On de Doubtful Usability of the Inoperability IO Model

15009-GEM: Zhang, L. and D. Bezemer, A Global House of Debt Effect? Mortgages and Post-Crisis Recessions in Fifty Economies

15010-I&O: Hooghiemstra, R., N. Hermes, L. Oxelheim, and T. Randøy, The Impact of Board Internationalization on Earnings Management

15011-EEF: Haan, M.A., and W.H. Siekman, Winning Back the Unfaithful while Exploiting the Loyal: Retention Offers and Heterogeneous Switching Costs

15012-EEF: Haan, M.A., J.L. Moraga-González, and V. Petrikaite, Price and Match-Value Advertising with Directed Consumer Search

15013-EEF: Wiese, R., and S. Eriksen, Do Healthcare Financing Privatisations Curb Total Healthcare Expenditures? Evidence from OECD Countries

15014-EEF: Siekman, W.H., Directed Consumer Search

15015-GEM: Hoorn, A.A.J. van, Organizational Culture in the Financial Sector: Evidence from a Cross-Industry Analysis of Employee Personal Values and Career Success

15016-EEF: Te Bao, and C. Hommes, When Speculators Meet Constructors: Positive and Negative Feedback in Experimental Housing Markets

15017-EEF: Te Bao, and Xiaohua Yu, Memory and Discounting: Theory and Evidence

15018-EEF: Suari-Andreu, E., The Effect of House Price Changes on Household Saving Behaviour: A Theoretical and Empirical Study of the Dutch Case

15019-EEF: Bijlsma, M., J. Boone, and G. Zwart, Community Rating in Health Insurance: Trade-off between Coverage and Selection

15020-EEF: Mulder, M., and B. Scholtens, A Plant-level Analysis of the Spill-over Effects of the German *Energiewende*

15021-GEM: Samarina, A., L. Zhang, and D. Bezemer, Mortgages and Credit Cycle Divergence in Eurozone Economies

16001-GEM: Hoorn, A. van, How Are Migrant Employees Manages? An Integrated Analysis

16002-EEF: Soetevent, A.R., Te Bao, A.L. Schippers, A Commercial Gift for Charity

16003-GEM: Bouwmeerster, M.C., and J. Oosterhaven, Economic Impacts of Natural Gas Flow Disruptions

16004-MARK: Holtrop, N., J.E. Wieringa, M.J. Gijsenberg, and P. Stern, Competitive Reactions to Personal Selling: The Difference between Strategic and Tactical Actions

16005-EEF: Plantinga, A. and B. Scholtens, The Financial Impact of Divestment from Fossil Fuels

16006-GEM: Hoorn, A. van, Trust and Signals in Workplace Organization: Evidence from Job Autonomy Differentials between Immigrant Groups

16007-EEF: Willems, B. and G. Zwart, Regulatory Holidays and Optimal Network Expansion

16008-GEF: Hoorn, A. van, Reliability and Validity of the Happiness Approach to Measuring Preferences

16009-EEF: Hinloopen, J., and A.R. Soetevent, (Non-)Insurance Markets, Loss Size Manipulation and Competition: Experimental Evidence

16010-EEF: Bekker, P.A., A Generalized Dynamic Arbitrage Free Yield Model

16011-EEF: Mierau, J.A., and M. Mink, A Descriptive Model of Banking and Aggregate Demand

16012-EEF: Mulder, M. and B. Willems, Competition in Retail Electricity Markets: An Assessment of Ten Year Dutch Experience

16013-GEM: Rozite, K., D.J. Bezemer, and J.P.A.M. Jacobs, Towards a Financial Cycle for the US, 1873-2014

16014-EEF: Neuteleers, S., M. Mulder, and F. Hindriks, Assessing Fairness of Dynamic Grid Tariffs

16015-EEF: Soetevent, A.R., and T. Bružikas, Risk and Loss Aversion, Price Uncertainty and the Implications for Consumer Search

16016-HRM&OB: Meer, P.H. van der, and R. Wielers, Happiness, Unemployment and Self-esteem

16017-EEF: Mulder, M., and M. Pangan, Influence of Environmental Policy and Market Forces on Coal-fired Power Plants: Evidence on the Dutch Market over 2006-2014

16018-EEF: Zeng,Y., and M. Mulder, Exploring Interaction Effects of Climate Policies: A Model Analysis of the Power Market

16019-EEF: Ma, Yiqun, Demand Response Potential of Electricity End-users Facing Real Time Pricing

16020-GEM: Bezemer, D., and A. Samarina, Debt Shift, Financial Development and Income Inequality in Europe

16021-EEF: Elkhuizen, L, N. Hermes, and J. Jacobs, Financial Development, Financial Liberalization and Social Capital

16022-GEM: Gerritse, M., Does Trade Cause Institutional Change? Evidence from Countries South of the Suez Canal

16023-EEF: Rook, M., and M. Mulder, Implicit Premiums in Renewable-Energy Support Schemes

17001-EEF: Trinks, A., B. Scholtens, M. Mulder, and L. Dam, Divesting Fossil Fuels: The Implications for Investment Portfolios

17002-EEF: Angelini, V., and J.O. Mierau, Late-life Health Effects of Teenage Motherhood

17003-EEF: Jong-A-Pin, R., M. Laméris, and H. Garretsen, Political Preferences of (Un)happy Voters: Evidence Based on New Ideological Measures

17004-EEF: Jiang, X., N. Hermes, and A. Meesters, Financial Liberalization, the Institutional Environment and Bank Efficiency

17005-EEF: Kwaak, C. van der, Financial Fragility and Unconventional Central Bank Lending Operations

17006-EEF: Postelnicu, L. and N. Hermes, The Economic Value of Social Capital

17007-EEF: Ommeren, B.J.F. van, M.A. Allers, and M.H. Vellekoop, Choosing the Optimal Moment to Arrange a Loan

17008-EEF: Bekker, P.A., and K.E. Bouwman, A Unified Approach to Dynamic Mean-Variance Analysis in Discrete and Continuous Time

17009-EEF: Bekker, P.A., Interpretable Parsimonious Arbitrage-free Modeling of the Yield Curve

17010-GEM: Schasfoort, J., A. Godin, D. Bezemer, A. Caiani, and S. Kinsella, Monetary Policy Transmission in a Macroeconomic Agent-Based Model

17011-I&O: Bogt, H. ter, Accountability, Transparency and Control of Outsourced Public Sector Activities

17012-GEM: Bezemer, D., A. Samarina, and L. Zhang, The Shift in Bank Credit Allocation: New Data and New Findings

17013-EEF: Boer, W.I.J. de, R.H. Koning, and J.O. Mierau, Ex-ante and Ex-post Willingness-to-pay for Hosting a Major Cycling Event

17014-OPERA: Laan, N. van der, W. Romeijnders, and M.H. van der Vlerk, Higher-order Total Variation Bounds for Expectations of Periodic Functions and Simple Integer Recourse Approximations

17015-GEM: Oosterhaven, J., Key Sector Analysis: A Note on the Other Side of the Coin

17016-EEF: Romensen, G.J., A.R. Soetevent: Tailored Feedback and Worker Green Behavior: Field Evidence from Bus Drivers

17017-EEF: Trinks, A., G. Ibikunle, M. Mulder, and B. Scholtens, Greenhouse Gas Emissions Intensity and the Cost of Capital

17018-GEM: Qian, X. and A. Steiner, The Reinforcement Effect of International Reserves for Financial Stability

17019-GEM/EEF: Klasing, M.J. and P. Milionis, The International Epidemiological Transition and the Education Gender Gap

2018001-EEF: Keller, J.T., G.H. Kuper, and M. Mulder, Mergers of Gas Markets Areas and Competition amongst Transmission System Operators: Evidence on Booking Behaviour in the German Markets

2018002-EEF: Soetevent, A.R. and S. Adikyan, The Impact of Short-Term Goals on Long-Term Objectives: Evidence from Running Data

2018003-MARK: Gijsenberg, M.J. and P.C. Verhoef, Moving Forward: The Role of Marketing in Fostering Public Transport Usage

2018004-MARK: Gijsenberg, M.J. and V.R. Nijs, Advertising Timing: In-Phase or Out-of-Phase with Competitors?

2018005-EEF: Hulshof, D., C. Jepma, and M. Mulder, Performance of Markets for European Renewable Energy Certificates

2018006-EEF: Fosgaard, T.R., and A.R. Soetevent, Promises Undone: How Committed Pledges Impact Donations to Charity

2018007-EEF: Durán, N. and J.P. Elhorst, A Spatio-temporal-similarity and Common Factor Approach of Individual Housing Prices: The Impact of Many Small Earthquakes in the North of Netherlands

2018008-EEF: Hermes, N., and M. Hudon, Determinants of the Performance of Microfinance Institutions: A Systematic Review

2018009-EEF: Katz, M., and C. van der Kwaak, The Macroeconomic Effectiveness of Bank Bail-ins

2018010-OPERA: Prak, D., R.H. Teunter, M.Z. Babai, A.A. Syntetos, and J.E. Boylan, Forecasting and Inventory Control with Compound Poisson Demand Using Periodic Demand Data www.rug.nl/feb