

Tilburg University

Assessing the Effects of using Demand Parameters Estimates in Inventory Control

Janssen, E.; Strijbosch, L.W.G.; Brekelmans, R.C.M.

Publication date:
2006

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Janssen, E., Strijbosch, L. W. G., & Brekelmans, R. C. M. (2006). *Assessing the Effects of using Demand Parameters Estimates in Inventory Control*. (CentER Discussion Paper; Vol. 2006-90). Operations research.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Center



Discussion Paper

No. 2006–90

**ASSESSING THE EFFECTS OF USING DEMAND
PARAMETERS ESTIMATES IN INVENTORY CONTROL**

By Elleke Janssen, Leo Strijbosch, Ruud Brekelmans

September 2006

ISSN 0924-7815

Assessing the effects of using demand parameters estimates in inventory control

Elleke Janssen^{*†} Leo Strijbosch^{*} Ruud Brekelmans^{*}

Abstract

Inventory models need some specification of the distribution of demand in order to find the optimal order-up-to level or reorder point. This distribution is unknown in real life and there are several solutions to overcome this problem. One approach is to assume a distribution, estimate its parameters and replace the unknown demand parameters by these estimates in the theoretically correct model. Earlier research suggests that this approach will lead to underperformance, even if the true demand distribution is indeed the assumed one. This paper directs the cause of the underperformance and quantifies it in case of normally distributed demand. Furthermore the formulae for the order-up-to levels are corrected analytically where possible and otherwise by use of simulation and linear regression. Simulation shows that these corrections improve the attained performance.

Keywords: Unknown Demand Parameters, Inventory Control, Normal Distribution, Service Level Criterion

JEL-classification: C13, C53.

1 Introduction

Inventory control involves decisions on what to order, when, and in what quantity. The models dealing with these decisions need information about the distribution of demand during some period, e.g. the demand during lead time or during the review period. Bulinskaya (1990) discriminates between three situations: (a) the type of distribution is known, but its parameters are unspecified; (b) only several first moments of the demand distribution are known; and (c) there is no prior knowledge about the demand. The third situation is of course the most realistic and different approaches to deal with situation (c) have been proposed in literature. These approaches can be categorized into parametric and nonparametric methods. Examples of the first category are assuming a distribution and using Bayesian models; examples of the second category are using order statistics, the bootstrap procedure and kernel densities.

One of the most widespread approaches to deal with unknown demand is assuming a distribution, estimating its parameters and replacing the unknown parameters by its estimates in the theoretically correct formulae in which distribution and parameters are supposed to be known. Sani and Kingsman (1997), Artto and Pylkkänen (1999), Strijbosch et al. (2000) and Syntetos

^{*}Tilburg University, Department of Econometrics & Operations Research, Tilburg, The Netherlands.

[†]Corresponding author. Tel.: +31 13 466 3244; fax: +31 13 466 3280; e-mail: E.Janssen_1@uvt.nl.

and Boylan (2006) use this approach with different inventory models, while Kottas and Lau (1980) provide a short discussion on estimating the parameters needed for their model. Another parametric approach is the Bayesian approach; Azoury and Miller (1984), Azoury (1985) and Karmarkar (1994) are three examples of the Bayesian approach. Also Larson et al. (2001) use this approach, but they introduce a nonparametric form. Other nonparametric approaches involve order statistics, references include Lordahl and Bookbinder (1994) and Liyanage and Shanthikumar (2005), the bootstrap procedure, see e.g. Bookbinder and Lordahl (1989) and Fricker and Goodhart (2000), or using kernel densities, see Strijbosch and Heuts (1992).

Consider the approach of assuming a distribution and let the assumed distribution be the normal, since it is often used both in research and in practice, cf. Zeng and Hayya (1999). Furthermore, see Strijbosch and Moors (2006) for references to recent articles that involve the normal distribution. The assumption of normality is made because it yields tractable results and it seems to give quite good approximations when used on data with a low coefficient of variation (cf. Silver et al. (1998) and Zipkin (2000)). However, the normal distribution has two major disadvantages: it is symmetric and can take on negative values, while demand is nonnegative and often skewed to the right. This may not impose serious problems if the coefficient of variation is low (Zipkin, 2000) or if the demand during review/lead time consists of many individual and independent demands (Silver et al., 1998). However for high values of the coefficient of variation, these disadvantages get more important; hence Strijbosch and Moors (2006) suggested two simple modifications of the normal distribution. Tyworth and O'Neill (1997) and Lau and Lau (2003) investigate the (non)robustness of using the normal approximation with an (s, Q) inventory policy.

This paper will not deal with situation (c), but with the less realistic situation (a): the true distribution is known, but its parameters are unspecified, so that the effect of estimating them can be studied. Silver and Rahnama (1986, 1987) have investigated this effect in an (s, Q) inventory policy with a cost criterion. They constructed a function that determined the expected cost of estimating the demand distribution rather than knowing it and they concluded that this function is not symmetrical: underestimating causes larger costs than overestimating. In the second article they propose a method that deliberately biases the reorder point upwards. Strijbosch et al. (1997) and Strijbosch and Moors (2005) have investigated the same effect for an (R, S) inventory policy with a service level criterion. Both papers concluded that also in this case the order-up-to level needed to be biased upwards. Note that the (s, Q) and (R, S) policies are equivalent (see Silver et al., 1998), so the results of Silver and Rahnama (1986, 1987) apply for the (R, S) policy and the results of Strijbosch et al. (1997) and Strijbosch and Moors (2005) for the (s, Q) policy as well.

This paper will focus on two service level criteria within an (R, S) inventory policy with zero lead time, so every R time units the inventory is replenished up to the order-up-to level S and the order is delivered instantaneously. The i.i.d. demands during review periods have a normal distribution with mean μ and standard deviation σ , which leads to a coefficient of variation $\nu = \sigma/\mu$. As mentioned before the normal distribution could lead to negative demand and in our model this is interpreted as returned goods, so demand is actually net demand (demanded goods minus returned goods). In addition the goods returned by customers can be sent back to the supplier, thus the inventory level at the start of a review period will always equal S . Furthermore, demand during $t + 1$ consecutive review periods is assumed to be stationary, and the first t periodic review demands are used to estimate the mean and standard deviation of the demand in review period $t + 1$. We prefer using sample statistics instead of exponential

smoothing, since derivations are more tractable, while the conclusions will be similar.

The next section lists the notation used in this paper. Section 3 discusses the P_1 service criterion in short. An analytical correction of the order-up-to level is given for the case that only μ is unknown. Section 4 focusses on the P_2 service criterion. First two theoretical situations are considered for illustrative purposes: the cases that μ (μ and σ) are unknown, but σ and ν (ν) are known. The main Sections 4.3 and 4.4 treat the important case that these three parameters are all unknown. We show by simulation that just plugging in estimates leads to serious underperformance. Besides, we derive a correction function for the safety factor that nearly gives the desired fill rates. The last section concludes this research and provides directions for further research.

2 Notation

X_i	Demand during review period i
μ	Mean of demand during review period
σ	Standard deviation of demand during review period
ν	Coefficient of variation of demand during review period
t	Number of review periods with useful historical demand information
$\hat{\mu}_t$	Sample mean based on t observations
$\hat{\sigma}_t$	Sample standard deviation based on t observations
$\hat{\nu}_t$	Estimated coefficient of variation of demand during review
$\phi(\Phi)$	The pdf (cdf) of standard normal distribution
$G_u(x)$	Loss function of standard normal distribution
$(x)^+$	Short notation for $\max(x, 0)$
$S(\mu, \sigma, c)$	Order-up-to level as a function of mean μ , standard deviation σ and safety factor c , i.e., $S(\mu, \sigma, c) = \mu + \sigma c$
α	Desired P_1 -service level
c_α	Safety factor for P_1 -service level
β	Desired P_2 -service level
$c_\beta, c_\beta^T, \tilde{c}_\beta^T$	Safety factors for P_2 -service level
$\kappa(\nu, t, \beta)$	Correction needed to attain desired service level
$\hat{\kappa}(\nu, t, \beta)$	Estimate of correction needed to attain desired service level

3 P_1 service criterion

This section considers the P_1 service criterion. This criterion states that the fraction of periods in which inventory is sufficient to satisfy demand, is at least α . It is common to express the order-up-to level as a function of the mean μ , the standard deviation σ and a safety factor c_α .

Since demand is normally distributed, the order-up-to level should be as defined below (see e.g. Silver et al., 1998).

$$S(\mu, \sigma, c_\alpha) = \mu + c_\alpha \sigma$$

The safety factor c_α is $\Phi^{-1}(\alpha)$. S without arguments is used to denote the theoretically correct order-up-to level when all parameters are known, so $S = S(\mu, \sigma, c_\alpha)$ in case of a P_1 service criterion.

In practice the mean and variance are unknown, which means that S is unknown too. The common solution is to substitute the parameters μ and σ by its estimates. If we (unrealistically) assume that only μ is unknown and use the sample mean to estimate it, the resulting order-up-to level is $S(\hat{\mu}_t, \sigma, c_\alpha)$ with $\hat{\mu}_t = \frac{1}{t} \sum_{i=1}^t X_i$. This order-up-to level, although unbiased, will not meet the service requirements in the long run. Consider Figure 1. Since $S(\hat{\mu}_t, \sigma, c_\alpha)$ is normally




Figure 1: The PDF of demand during review X_i and order-up-to level $S(\hat{\mu}_t, \sigma, c_\alpha)$.

distributed, with mean S and variance σ^2/t , it is symmetric and a shift from S to the right is equally probable as a shift of the same size to the left. The shift to the right decreases the probability of having backorders with the darker area, while the shift of the same size to the left will increase the probability of having backorders with the lighter area. The surface of the lighter area is larger than the surface of the darker area and this implies that in the long run, the achieved service level will fall short of the desired one. This phenomenon is mathematically explained by considering the following.

$$P(X_{t+1} < S(\hat{\mu}_t, \sigma, c_\alpha)) = P\left(X_{t+1} - \frac{1}{t} \sum_{i=1}^t X_i < c_\alpha \sigma\right) = \Phi\left(\frac{c_\alpha}{\sqrt{1 + 1/t}}\right) < \Phi(c_\alpha) = \alpha$$

Now let Z be a general biased estimator of S and let $Z \sim N(S + b\sigma, (d^2 - 1)\sigma^2)$, so Z has a normal distribution. Note that $d \geq 1$. The probability of not having backorders in a review period is given below.

$$P(X_{t+1} < Z) = \Phi\left(\frac{c_\alpha + b}{d}\right)$$

In general this will not equal α , unless b and d are chosen according to the relation $b = (d - 1)c_\alpha$ (see also Strijbosch et al., 1997). This phenomenon is depicted in Figure 2. The gray surface




Figure 2: Attained P_1 service level ($\alpha = 0.95$) for estimator Z depending on b and d .

depicts the attained service level at values of $0 \leq b \leq 1.6$ and $1 \leq d \leq 1.5$. The fine grid depicts the desired service level. The attained service is below this level in some cases, while it is above in others. The attained service will only reach the required one on the solid line, for which obviously holds $b = (d - 1)c_\alpha$. Furthermore, if an unbiased estimator is used ($b = 0$, the dashed line), the desired service will only be reached if $d = 1$. That case corresponds to having an estimator with standard deviation equal to 0, which is not possible in practice. So, an unbiased normally distributed estimator will lead to underperformance.

Now consider the order-up-to level $S(\hat{\mu}_t, \sigma\tau, c_\alpha)$, where τ denotes $\sqrt{1 + 1/t}$. This order-up-to level is normally distributed with mean $\mu + c_\alpha\sigma\tau$ and variance σ^2/t , so it is a special case of Z with $d = \tau$ and $b = (d - 1)c_\alpha$. Hence the performance of this order-up-to level will be satisfactory. The sample mean $\hat{\mu}_t$ can be interpreted as a forecast of demand during the subsequent review period with forecast error variance $\sigma^2\tau^2$. So replacing the standard deviation of demand during review by the square root of the forecast error variance results in attaining the desired service. Note that σ is replaced by $\sigma\tau$, although this parameter is known. This may seem counterintuitive, but replacing σ only increases the expected value of the order-up-to level; it does not change its variance.

The next step is to assume that also σ is unknown, but since the P_1 criterion is of lesser interest compared to the P_2 criterion, only the latter will be considered in more depth.

4 P_2 service criterion

This section focusses on the P_2 service criterion, also known as the fill rate. This criterion states that at least a fraction β of total demand has to be satisfied immediately from stock. Using again that demand is normally distributed, the order-up-to level in this case should be (see e.g. Silver et al., 1998) as follows.

$$S(\mu, \sigma, c_\beta) = \mu + c_\beta\sigma \tag{1}$$

The safety factor c_β is given by $c_\beta = G_u^{-1}\left(\frac{1-\beta}{\nu}\right)$, where $G_u(\cdot)$ denotes the loss function of the standard normal distribution ($G_u(x) = E[(Y-x)^+]$ with Y a standard normal random variable), which is equal to $G_u(x) = \phi(x) - x\Phi(-x)$. As in the previous section, S without arguments again denotes the correct order-up-to level. Furthermore, the service level attained by using an order-up-to level Z is given below.

$$1 - \frac{E[(X_{t+1} - Z)^+]}{E[X_{t+1}]} = 1 - \frac{E[(X_{t+1} - Z)^+]}{\mu}$$

Note that if the order-up-to level Z is constructed using estimates instead of the true values of the mean, standard deviation and coefficient of variation, it is a random variable. When using a P_2 service criterion, one needs values of μ and σ , and in contrast to the P_1 criterion, also the coefficient of variation ν plays a role, as can be seen from the safety factor in (1). In practice these parameters have to be estimated. The next two sections illustrate the effect of estimating μ and σ in the case where the correct safety factor is used, i.e., when ν is known. Section 4.3 considers the more realistic case where also the coefficient of variation needs to be estimated, and hence the safety factor as well.

4.1 Only μ unknown

This section assumes that only μ is unknown and thus that σ and ν are known. This is a purely theoretical assumption, since if σ and ν are known, μ can easily be determined. However, this case is interesting since even now the commonly used order-up-to level will not guarantee the desired service level. If the sample mean is used to estimate μ , the order-up-to level will be $S(\hat{\mu}_t, \sigma, c_\beta)$. Note that $S(\hat{\mu}_t, \sigma, c_\beta)$ is normally distributed with mean S and variance σ^2/t . Using $S(\hat{\mu}_t, \sigma, c_\beta)$ will not result in attaining the desired service level β in the long run, which is shown below¹. The Greek letter τ still denotes $\sqrt{1 + 1/t}$.

$$1 - \frac{E[(X_{t+1} - S(\hat{\mu}_t, \sigma, c_\beta))^+]}{\mu} = 1 - \frac{\sigma\tau G_u(c_\beta/\tau)}{\mu} < 1 - \nu\tau G_u(c_\beta) = 1 - \nu\tau \frac{1-\beta}{\nu} < \beta$$

Again, even in this simple case, using the most obvious estimator for S , however unbiased, leads to a lower service level than β . In case of the P_1 criterion, the problem of underperformance was solved by replacing the standard deviation σ by the square root of the forecast error variance, $\sigma\tau$. So it is obvious to apply the same adjustment in this case. Note however that σ has to be replaced twice: once explicitly and once implicitly, as the safety factor c_β depends on σ via ν . Now let us denote the new safety factor by c_β^τ .

$$c_\beta^\tau = G_u^{-1}\left(\frac{1-\beta}{\nu\tau}\right)$$

Obviously $c_\beta^\tau > c_\beta$. Upwards biasing the safety factor or the standard deviation has been mentioned in literature; see Section 1. However, to our knowledge, applying the factor τ both to the explicit and the implicit standard deviation has never been mentioned. Now consider the attained service level when using the order-up-to level $S(\hat{\mu}_t, \sigma\tau, c_\beta^\tau)$.

$$1 - \frac{E[(X_{t+1} - S(\hat{\mu}_t, \sigma\tau, c_\beta^\tau))^+]}{\mu} = 1 - \nu\tau G_u(c_\beta^\tau) = 1 - \nu\tau \frac{1-\beta}{\nu\tau} = \beta$$

¹Note that the derivative of $G_u(x)$ is $-\Phi(-x) < 0$, so $G_u(x)$ is a strictly decreasing function.

So using $S(\hat{\mu}_t, \sigma\tau, c_\beta^\tau)$ will indeed result in attaining the desired service level in the long run.

4.2 Both μ and σ unknown

Now assume that only ν is known, so σ has to be estimated using the sample standard deviation $\hat{\sigma}_t$. One could also use the known ν and the estimate of μ to estimate the value of σ , but since the next step is to assume that ν is also unknown, we do not use that approach. The sample standard deviation is determined by $\hat{\sigma}_t = \sqrt{\frac{1}{t-1} \sum_{i=1}^t (X_i - \hat{\mu}_t)^2}$. Substituting this estimate yields the order-up-to level $S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)$. Consider the expected value of this order-up-to level and note that $\hat{\sigma}_t^2$ is a unbiased estimator of σ^2 , so $E[\hat{\sigma}_t^2] = \sigma^2$. However from Jensen's Inequality it follows that $E[\hat{\sigma}_t] < \sigma$. This implies the following.

$$E[S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)] = E[\hat{\mu}_t + \hat{\sigma}_t c_\beta^\tau] = E[\hat{\mu}_t] + c_\beta^\tau E[\hat{\sigma}_t] < \mu + \sigma c_\beta^\tau = E[S(\hat{\mu}_t, \sigma\tau, c_\beta^\tau)]$$

Thus the expected value of $S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)$ is lower than the expected value of the order-up-to level in the case where σ is known. Furthermore, σ being unknown will result in extra variability, so the performance of $S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)$ will probably be lower than desired. Since we want to quantify the underperformance simulation runs have been performed for different values of t , ν and β . There are $n = 1,000,000$ samples of $t + 1$ normally distributed observations with mean $1/\nu$ and standard deviation 1 randomly generated. The mean μ and standard deviation σ need not to be varied, since the performance does not depend on these parameters separately (see appendix). The order-up-to levels are determined using $S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)$ and subsequently the attained service, denoted by $\hat{\beta}$ (see appendix for definition), is estimated. The estimates of the attained service are shown in Table 1. The left part of Table 1 displays the attained service

		$\hat{\beta}$			$\frac{(1 - \hat{\beta}) - (1 - \beta)}{1 - \beta} \cdot 100\%$		
		$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$	$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$
$\beta = 0.90$	$t = 2$	0.9007	0.8640	0.8252	-0.72	35.99	74.76
	$t = 6$	0.9008	0.8934	0.8849	-0.80	6.57	15.06
	$t = 10$	0.9006	0.8961	0.8919	-0.57	3.91	8.11
	$t = 15$	0.9002	0.8974	0.8945	-0.21	2.61	5.52
$\beta = 0.95$	$t = 2$	0.9386	0.8999	0.8625	22.80	100.15	174.97
	$t = 6$	0.9480	0.9397	0.9322	3.91	20.57	35.62
	$t = 10$	0.9489	0.9444	0.9398	2.25	11.13	20.36
	$t = 15$	0.9494	0.9462	0.9437	1.23	7.56	12.52
$\beta = 0.99$	$t = 2$	0.9679	0.9349	0.9045	221.12	551.08	854.99
	$t = 6$	0.9855	0.9799	0.9753	44.64	100.52	146.52
	$t = 10$	0.9876	0.9846	0.9824	24.06	54.07	75.64
	$t = 15$	0.9884	0.9867	0.9855	15.98	32.73	45.36

Table 1: Simulated performance of $S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)$ for different t , ν and β .

level in simulation. So if one considers the case that $\beta = 0.95$, $t = 2$ and $\nu = 0.2$, one can see

that the achieved service is 0.9386. The right part of the table gives the percentage deviation of the backordered demand, so in the example mentioned above backordering is 1.14% higher than the desired 5%, which is a deviation of almost 23%.

As one can clearly see from Table 1 the underperformance is worse if ν becomes larger. This is exactly what we expected to happen, since if ν is larger (*ceteris paribus*), the variability of the order-up-to level will be larger and hence, the expected amount of backorders will be larger and thus the achieved service will be less. The same line of reasoning applies to t being smaller. The only exception is $\nu = 0.2$ and $\beta = 0.90$, where the desired service is slightly exceeded and it does not really matter what value t has. From the right part of the table it is clear that if β is larger the percentage deviation gets larger as well. So, the underperformance is relatively larger when β is larger.

Since it is difficult to consider this case analytically and it is not of practical interest (as ν is assumed to be known), we continue with the most practical case, in which all parameters are unknown.

4.3 All parameters — μ , σ and ν — unknown

In this section the demand is normally distributed with unknown mean and standard deviation, and also the coefficient of variation is unknown. Estimates are used for all of these unknowns, namely $\hat{\mu}_t$, $\hat{\sigma}_t$ and $\hat{\nu}_t = \hat{\sigma}_t / \hat{\mu}_t$. Since the safety factor c_β^τ depends on ν , this factor also has to be estimated; \hat{c}_β^τ denotes this estimate and it is defined as follows.

$$\hat{c}_\beta^\tau = \begin{cases} G_u^{-1} \left(\frac{1-\beta}{\hat{\nu}_t \tau} \right) & \text{if } \hat{\nu}_t > 0 \\ -\frac{1}{\hat{\nu}_t \tau} & \text{otherwise} \end{cases} \quad (2)$$

Note that ν is simply replaced by $\hat{\nu}_t$ if it is possible to do so. Since $\hat{\mu}_t$ may be negative (the demand values are generated using a normal distribution), $\hat{\nu}_t$ can be negative as well. In that case the function $G_u^{-1}(\cdot)$ has no outcome, as its domain is strictly positive. If $\hat{\mu}_t$ is negative, it means that the demand in the next period is forecasted to be negative. So inventory is not needed in that case and hence \hat{c}_β^τ is chosen in such a way that the resulting order-up-to level equals zero.

For this case the order-up-to level is even more complicated than in the previous section, so again it is not possible to get analytical results. Therefore simulation is applied; first to estimate the attained service level when using order-up-to level $S(\hat{\mu}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$ and second to find a correction to that order-up-to level that would assure that the desired service is reached more closely. The P_2 service using $S(\hat{\mu}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$ is estimated with help of $n = 1,000,000$ simulation runs for each combination of t , β and ν . Note that again the attained service only depends on ν and not on μ and σ separately; see the appendix for further details. This simulation is performed in the same way as described in the previous section; the only difference is that \hat{c}_β^τ is used instead of c_β^τ . The results, shown in Table 2, are based on the same samples as used in the previous section. In most cases the performance is indeed worse, as expected. In a few cases (t and ν both high and β not), the performance of the $S(\hat{\mu}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$ is slightly better than the performance of $S(\hat{\mu}_t, \hat{\sigma}_t \tau, c_\beta^\tau)$. Furthermore the same overall results as in Table 1 appear: the underperformance increases when t decreases and when ν increases. The underperformance also relatively increases when β increases.

		$\hat{\beta}$			$\frac{(1 - \hat{\beta}) - (1 - \beta)}{1 - \beta} \cdot 100\%$		
		$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$	$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$
$\beta = 0.90$	$t = 2$	0.8909	0.8520	0.8005	9.12	47.96	99.47
	$t = 6$	0.8993	0.8924	0.8873	0.74	7.59	12.67
	$t = 10$	0.8998	0.8958	0.8942	0.15	4.19	5.78
	$t = 15$	0.8998	0.8974	0.8961	0.24	2.64	3.90
$\beta = 0.95$	$t = 2$	0.9289	0.8885	0.8345	42.29	122.97	230.97
	$t = 6$	0.9456	0.9373	0.9317	8.72	25.39	36.60
	$t = 10$	0.9476	0.9432	0.9401	4.87	13.62	19.80
	$t = 15$	0.9486	0.9455	0.9440	2.88	9.06	11.96
$\beta = 0.99$	$t = 2$	0.9620	0.9264	0.8733	280.09	635.52	1166.97
	$t = 6$	0.9837	0.9779	0.9734	63.16	120.78	166.33
	$t = 10$	0.9865	0.9835	0.9818	34.66	64.96	82.14
	$t = 15$	0.9877	0.9860	0.9851	22.96	39.80	49.36

Table 2: Simulated performance of $S(\hat{\mu}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$ for different t , ν and β .

4.4 Correction of the order-up-to level

Now the underperformance is quantified, we also want to find a correction to $S(\hat{\mu}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$, such that the desired service is reached or at least approached more closely. This correction will depend on ν , t and β and is denoted by $\kappa(\nu, t, \beta)$. Such a correction function can be useful in practice, since it provides inventory managers with a simple tool to improve their easy-to-understand order-up-to levels. After considering several options to correct the order-up-to level, a dimensionless correction factor is determined in order to provide an upwards bias to $S(\hat{\mu}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$. The corrected order-up-to level is $S(\hat{\mu}_t + \kappa(\nu, t, \beta)\hat{\sigma}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$. Note that this will not be applicable in practice, since ν is unknown. However substituting ν by $\hat{\nu}_t$ again and using this correction will improve the attained service.

Simulation is used to estimate the correction needed for various values of ν , t and β . The two simulations performed earlier only used a limited number of values for the three parameters, but more values are needed to be able to estimate the correction using ν , t and β . So a new simulation is performed in which $n = 100,000$ samples of $t + 1$ observations are generated for each combination of t , ν and β . For each simulation run j ($j = 1, \dots, n$) the order-up-to level, denoted by S_j , is calculated using the first t observations and the $(t + 1)$ -th observation x_j is used to quantify the backorders that occurred. The correction needed can be determined using either the true value σ or the sample standard deviations $\hat{\sigma}_{tj}$ ($j = 1, \dots, n$) found in the simulation. On the one hand using the true value might be better, since it is the *true* value. On the other hand, a correction is needed since an estimate is used instead of the true value. So why not use the estimated value of the standard deviation? It is difficult to decide a priori which would result in a better correction formula, so both corrections, denoted by k_1 and k_2

respectively, are determined, by solving the equations below.

$$\sum_{j=1}^n (x_j - (S_j + k_1 \sigma))^+ = (1 - \beta) \sum_{j=1}^n x_j$$

$$\sum_{j=1}^n (x_j - (S_j + k_2 \hat{\sigma}_{tj}))^+ = (1 - \beta) \sum_{j=1}^n x_j$$

The values for k_1 and k_2 for all combinations of ν , t and β can be found by solving this equation numerically and the corrections are denoted by $k_i(\nu, t, \beta)$ ($i = 1, 2$). The corrections for different values of ν , t and β are depicted in Figure 3. For the first three graphs in Figures

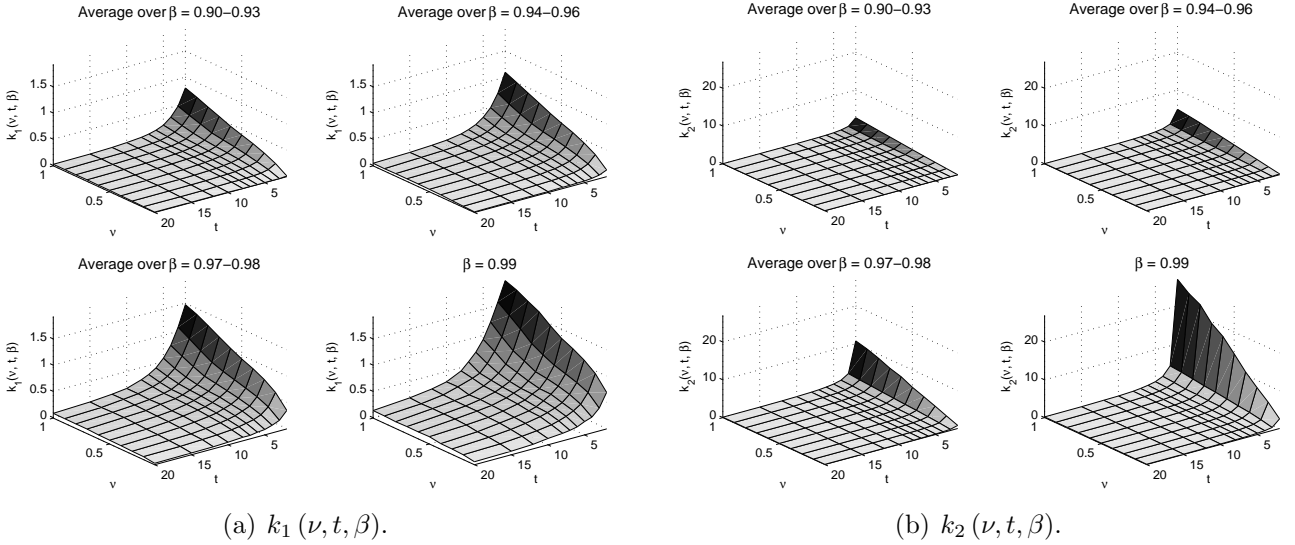


Figure 3: Graphs of corrections needed.


3(a) and 3(b) the corrections needed for different values of β are averaged, since they do not differ much. One sees that if ν becomes larger, the correction needs to be larger. The same is true if t decreases and if β increases. These conclusions correspond to the results shown in Tables 1 and 2.

The simulation resulted in some values $k_1(\nu, t, \beta)$ and $k_2(\nu, t, \beta)$, but a function that is able to estimate these values is more practical. Therefore linear regression is applied, following the method described by Strijbosch and Moors (1999). The estimation process is split in three steps:

1. Fix two of the three parameters and use linear regression to estimate the correction needed depending on the third parameter (denoted by q_1);
2. Fix one of the two parameters fixed in Step 1 and use the other (q_2) to estimate the coefficients found in Step 1;
3. Use the parameter fixed in Step 2 (q_3) to estimate the coefficients found in Step 2.

All six orderings in regressing the parameters are examined, both for $k_1(\nu, t, \beta)$ and $k_2(\nu, t, \beta)$. The three steps are discussed in more detail for $k_1(\nu, t, \beta)$ and with $q_1 = t$, $q_2 = \nu$ and $q_3 = \beta$.

Step 1 For every combination of ν and β the correction needed to attain the desired service is estimated. Figure 4(a) depicts $k_1(\nu, t, \beta)$ for three combinations of ν and β .



(a) $k_1(\nu, t, \beta)$ for some values of ν and β .

(b) $SSE(r)$ depending on the exponent r .

Figure 4: Graphical illustration of the process of choosing the correct value for r .

The graphs suggest a relation of the following form.

$$k_1(\nu, t, \beta) = \gamma_0(\nu, \beta, r) + \gamma_1(\nu, \beta, r)t^r + \varepsilon$$

The value of r depends on the shape of $k_1(\nu, t, \beta)$. The coefficients γ_0 and γ_1 are estimated using linear regression for values of r running from -20 up to 20 (in steps of 0.01) and this results in the regression equation $\hat{k}(\nu, t, \beta, r) = \hat{\gamma}_0(\nu, \beta, r) + \hat{\gamma}_1(\nu, \beta, r)t^r$. The exponent r was assigned the value that minimized $SSE(r)$, which is defined as below.

$$SSE(r) = \sum_{t \in \mathcal{T}} \sum_{\nu \in \mathcal{V}} \sum_{\beta \in \mathcal{B}} \left(k_1(\nu, t, \beta) - \hat{k}(\nu, t, \beta, r) \right)^2 \quad (3)$$

In this definition \mathcal{T} , \mathcal{V} and \mathcal{B} denote the sets of values used for t , ν and β respectively, which are given below.

$$\mathcal{T} = \{2, 3, 4, 5, 6, 7, 8, 10, 12, 15, 20\}$$

$$\mathcal{V} = \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$$

$$\mathcal{B} = \{0.90, 0.91, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, 0.99\}$$

$SSE(r)$ attains its minimum at $r = -1.16$ as Figure 4(b) shows. So this value is used to estimate the coefficients γ_0 and γ_1 for every combination of ν and β .

Step 2 The next step is to estimate the value of $\hat{\gamma}_0$ and $\hat{\gamma}_1$ for every value of β depending on ν . Using the same approach as described in Step 1 the following

relations are considered.

$$\begin{aligned}\hat{\gamma}_0(\nu, \beta, r) &= \eta_{00}(\beta, r_0) + \eta_{01}(\beta, r_0)\nu^{r_0} + \varepsilon \\ \hat{\gamma}_1(\nu, \beta, r) &= \eta_{10}(\beta, r_1) + \eta_{11}(\beta, r_1)\nu^{r_1} + \varepsilon\end{aligned}$$

For both $\hat{\gamma}_0$ and $\hat{\gamma}_1$ the best fitting exponents r_0 and r_1 are determined in the same way as described in Step 1. Using these exponents linear regression determines estimates for η_{00} , η_{01} , η_{10} and η_{11} , which are denoted by $\hat{\eta}_{00}$, $\hat{\eta}_{01}$, $\hat{\eta}_{10}$ and $\hat{\eta}_{11}$ and depend on the value of β and r_0 or r_1 .

Step 3 The final step is to estimate the value of $\hat{\eta}_{00}$, $\hat{\eta}_{01}$, $\hat{\eta}_{10}$ and $\hat{\eta}_{11}$ depending on β , thus the relations below appear.

$$\begin{aligned}\hat{\eta}_{00}(\beta, r_0) &= \xi_{000}(r_{00}) + \xi_{001}(r_{00})(1 - \beta)^{r_{00}} + \varepsilon \\ \hat{\eta}_{01}(\beta, r_0) &= \xi_{010}(r_{01}) + \xi_{011}(r_{01})(1 - \beta)^{r_{01}} + \varepsilon \\ \hat{\eta}_{10}(\beta, r_1) &= \xi_{100}(r_{10}) + \xi_{101}(r_{10})(1 - \beta)^{r_{10}} + \varepsilon \\ \hat{\eta}_{11}(\beta, r_1) &= \xi_{110}(r_{11}) + \xi_{111}(r_{11})(1 - \beta)^{r_{11}} + \varepsilon\end{aligned}$$

The independent variable is $1 - \beta$ instead of β , since it is used in the determination of the safety factor \hat{c}_β^τ , see equation (2). Furthermore, this results in better estimates. The values of the exponents are again found by minimizing the total sum of squares for errors. These exponents are used to estimate the coefficients in the four relations above.

Using the steps discussed above to estimate $k_1(\nu, t, \beta)$ results in a R^2 (determination coefficient) of 0.9897, so the fit is good. The fit of all possible orderings of adding the parameters and for both $k_1(\nu, t, \beta)$ and $k_2(\nu, t, \beta)$ are denoted in Table 3. So using $k_2(\nu, t, \beta)$ instead of

q_1	q_2	q_3	R^2 of $k_1(\nu, t, \beta)$	R^2 of $k_2(\nu, t, \beta)$
t	ν	β	0.9897	0.9960
t	β	ν	0.9894	0.9962
ν	t	β	0.9828	0.9963
ν	β	t	0.9872	0.9889
β	t	ν	0.9876	0.9958
β	ν	t	0.9840	0.9900

Table 3: R^2 found for different orderings of adding parameters.

$k_1(\nu, t, \beta)$ results in slightly better values of R^2 , although the determination coefficients for both corrections are very good. Next a simulation is performed to determine which provides the best performance. Again the samples of size $n = 1,000,000$ are used, now to determine the performance of $S(\hat{\mu}_t + \hat{\kappa}_1(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$ and $S(\hat{\mu}_t + \hat{\kappa}_2(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$. The function $\hat{\kappa}_1(\hat{\nu}_t, t, \beta)$ is based on $k_1(\nu, t, \beta)$ with ordering $q_1 = t$, $q_2 = \nu$ and $q_3 = \beta$, while $\hat{\kappa}_2(\hat{\nu}_t, t, \beta)$ is based on $k_2(\nu, t, \beta)$ with ordering $q_1 = \nu$, $q_2 = t$ and $q_3 = \beta$. Also the performance of these order-up-to levels is independent of σ and μ ; see appendix. The results are shown Tables

		$\hat{\beta}$			$\frac{(1 - \hat{\beta}) - (1 - \beta)}{1 - \beta} \cdot 100\%$		
		$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$	$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$
$\beta = 0.90$	$t = 2$	0.8945	0.8730	0.8449	5.47	27.04	55.12
	$t = 6$	0.8998	0.8998	0.9005	0.17	0.23	-0.48
	$t = 10$	0.8995	0.8984	0.8985	0.49	1.63	1.53
	$t = 15$	0.8990	0.8973	0.8953	1.04	2.69	4.67
$\beta = 0.95$	$t = 2$	0.9387	0.9143	0.8882	22.65	71.36	123.59
	$t = 6$	0.9497	0.9479	0.9480	0.59	4.29	4.10
	$t = 10$	0.9497	0.9488	0.9485	0.69	2.47	3.08
	$t = 15$	0.9496	0.9483	0.9479	0.81	3.35	4.14
$\beta = 0.99$	$t = 2$	0.9732	0.9525	0.9332	168.11	375.34	568.07
	$t = 6$	0.9883	0.9865	0.9858	17.40	35.19	41.89
	$t = 10$	0.9894	0.9886	0.9890	6.25	13.75	9.73
	$t = 15$	0.9897	0.9895	0.9900	3.32	5.24	0.48

Table 4: Simulated performance of $S(\hat{\mu}_t + \hat{\kappa}_1(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^T)$ for different t , ν and β .

		$\hat{\beta}$			$\frac{(1 - \hat{\beta}) - (1 - \beta)}{1 - \beta} \cdot 100\%$		
		$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$	$\nu = 0.2$	$\nu = 0.5$	$\nu = 0.8$
$\beta = 0.90$	$t = 2$	0.8990	0.9034	0.9007	0.99	-3.43	-0.69
	$t = 6$	0.8986	0.8968	0.8972	1.45	3.24	2.78
	$t = 10$	0.8986	0.8990	0.9025	1.44	1.03	-2.50
	$t = 15$	0.8984	0.9005	0.9045	1.63	-0.47	-4.48
$\beta = 0.95$	$t = 2$	0.9508	0.9515	0.9506	-1.63	-3.06	-1.29
	$t = 6$	0.9491	0.9479	0.9491	1.86	4.30	1.82
	$t = 10$	0.9506	0.9527	0.9557	-1.13	-5.34	-11.50
	$t = 15$	0.9516	0.9551	0.9597	-3.15	-10.16	-19.40
$\beta = 0.99$	$t = 2$	0.9901	0.9899	0.9898	-1.01	0.67	1.66
	$t = 6$	0.9890	0.9880	0.9879	9.91	19.83	21.36
	$t = 10$	0.9910	0.9908	0.9915	-9.88	-8.45	-15.06
	$t = 15$	0.9921	0.9928	0.9937	-20.52	-27.55	-36.99

Table 5: Simulated performance of $S(\hat{\mu}_t + \hat{\kappa}_2(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^T)$ for different t , ν and β .

4 and 5. When comparing Table 2 on the one hand to Tables 4 and 5 on the other it is clear that the achieved service level is closer to the desired one when using the corrections in most cases. Comparing Tables 4 and 5 shows that in most cases the order-up-to level based on $k_2(\nu, t, \beta)$

performs better, although the differences are not that obvious if t is large. However if t is small, $S(\hat{\mu}_t + \hat{\kappa}_1(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$ performs far worse compared to $S(\hat{\mu}_t + \hat{\kappa}_2(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$. The first does not reach the desired service if t is small, while it does approximately if t is large. The performance of $S(\hat{\mu}_t + \hat{\kappa}_2(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$ hardly depends on the values of t , ν and β , so $\hat{\kappa}_2(\hat{\nu}_t, t, \beta)$ seems to be the best correction. The most extreme deviations from the desired service using this correction are -0.0032 (displayed in italic in Table 5) and +0.0097 (displayed in bold), so the attained service is close to the desired one. As an extra check, the performance of using $\hat{\kappa}_2(\hat{\nu}_t, t, \beta)$ is simulated more elaborate and the results are similar. Finally, the formula for $\hat{\kappa}_2(\nu, t, \beta)$ is given below.

$$\begin{aligned} \hat{\kappa}_2(\nu, t, \beta) = & \\ & \left[(-0.0669 + 0.00305 \cdot (1 - \beta)^{-0.95}) + (-185.124 - 6.359 \cdot (1 - \beta)^{-1.00}) t^{-9.17} \right] \\ & + \left[(0.335 - 5.671 \cdot (1 - \beta)^{1.41}) + (-3.841 + 4.541 \cdot (1 - \beta)^{-1.03}) t^{-4.19} \right] \nu^{0.90} \end{aligned}$$

5 Conclusions and further research

Conclusions

This paper has investigated a common approach in inventory management of dealing with the unknown distribution of demand. This approach is to assume a distribution, estimate its parameters using historical demand information and replace the parameters in the theoretically correct inventory model by its estimates. Assuming a distribution provides tractable results, but it can be too rigid to represent the real demand. However, whether or not a specific distribution should be assumed is outside the scope of this paper.

We have assumed that the demand during review truly is normally distributed. In steps the information about the mean and variance is reduced. First, only the mean is assumed to be unknown and using the just mentioned common approach results in the order-up-to levels $S(\hat{\mu}_t, \sigma, c_\alpha)$ for a P_1 criterion and $S(\hat{\mu}_t, \sigma, c_\beta)$ for a P_2 criterion. Both order-up-to levels will not ensure that the desired service is reached. This can be resolved by replacing the standard deviation σ with the square root of the forecast error variance $\sigma\tau$, where $\tau = \sqrt{1 + 1/t}$.

Second, also the standard deviation becomes unknown, although the coefficient of variation is assumed to be known. This case is less tractable and, therefore, only the more interesting P_2 criterion is considered for this case. The order-up-to level using the correction factor τ becomes $S(\hat{\mu}_t, \hat{\sigma}_t \tau, c_\beta^\tau)$, for which we have shown that the expected value will be too low. Furthermore, simulation has shown that indeed the desired service is not reached when using this order-up-to level.

Finally, also the coefficient of variation ν is unknown and in that case only simulation is used to find that the performance is again worse compared to the case when ν is known. We have developed a correction to the order-up-to level with help of simulation and linear regression. Using this correction, being a function of ν , β and t , ensures that the desired service is reached closely in general: the largest deviations are 0.0032 below the desired service and 0.0097 above.

It can be concluded that simply replacing the parameters in a theoretical correct inventory control model by its estimates results in underperformance, even if the true distribution belongs to the assumed distribution family. With a simple correction the achieved service can be improved and using the correction function $\hat{\kappa}_2(\nu, t, \beta)$ results in closely reaching the desired

service.

Further research

This paper focussed on normally distributed demand within an (R, S) inventory policy with zero lead time. Similar results could also be derived for other inventory control policies, e.g. an (R, s, S) or (R, s, Q) policy, or for other demand distribution, e.g. a gamma distribution. This is one direction for further research. Another direction lies in the assumption of stationary demand. This paper assumed that demand during $t+1$ consecutive review periods is stationary, but in most real life situations the demand pattern changes over time. We could investigate how well our method, which assumes stationary demand, works in such a situation. In this paper only the last t observations are used to estimate the demand parameters, thus non-stationarity is taken into account to some extent. This forecasting method is known as the moving average. Another forecasting method, exponential smoothing, is used often in real life, as it reacts better to changes in the demand pattern compared to the moving average. Using another forecasting method provides a third direction for further research.

A Independence of achieved performance of μ and σ

This appendix shows that the achieved performance depends on the quotient of σ and μ (and thus on the coefficient of variation ν), but that it is independent of μ and σ separately. Three order-up-to levels, $S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)$, $S(\hat{\mu}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$ and $S(\hat{\mu}_t + \hat{\kappa}_i(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t\tau, \hat{c}_\beta^\tau)$, are considered. The first is discussed in Section 4.2, the second in Section 4.3 and the third in Section 4.4.

A.1 $S(\hat{\mu}_t, \hat{\sigma}_t\tau, c_\beta^\tau)$

The service achieved in simulation is denoted by $\hat{\beta}(\mathbf{S}_n)$ where \mathbf{S}_n is the vector of the n order-up-to levels determined for the n samples. It is defined as given below.

$$\hat{\beta}(\mathbf{S}_n) = 1 - \frac{\sum_{j=1}^n (x_j - S_j)^+}{\sum_{j=1}^n x_j}$$

In this definition x_j denotes the observation that is used to check the order-up-to level obtained in the j -th simulation run (n runs in total). $S_j = S(\hat{\mu}_{tj}, \hat{\sigma}_{tj}\tau, c_\beta^\tau)$ is the order-up-to level determined in the j -th simulation, where $\hat{\mu}_{tj}$ and $\hat{\sigma}_{tj}$ are defined as follows.

$$\hat{\mu}_{tj} = \frac{1}{t} \sum_{i=1}^t x_{ij} \quad \text{and} \quad \hat{\sigma}_{tj} = \sqrt{\frac{1}{t-1} \sum_{i=1}^t (x_{ij} - \hat{\mu}_{tj})^2}$$

In the above x_{1j}, \dots, x_{tj} are the t demand observations that are used in the j -th simulation run to determine the order-up-to level.

Since $x_j \sim N(\mu, \sigma^2)$, $x_j^* = \frac{x_j - \mu}{\sigma} \sim N(0, 1)$ and hence x_j^* is indeed independent of μ and σ . Note that the same holds for $x_{ij}^* = (x_{ij} - \mu)/\sigma$. Also $S_j^* = (S_j - \mu)/\sigma$ is independent of μ and σ , as is shown below.

$$S_j^* = \frac{\frac{1}{t} \sum_{i=1}^t x_{ij} + \hat{\sigma}_{tj} c_\beta^\tau \tau - \mu}{\sigma} = \frac{1}{t} \sum_{i=1}^t \left(\frac{x_{ij} - \mu}{\sigma} \right) + \frac{\hat{\sigma}_{tj}}{\sigma} c_\beta^\tau \tau = \frac{1}{t} \sum_{i=1}^t x_{ij}^* + \hat{\sigma}_{tj}^* c_\beta^\tau \tau$$

In the above $\hat{\sigma}_{tj}^* = \hat{\sigma}_{tj}/\sigma$. It suffices to show that $\hat{\sigma}_{tj}^*$ is independent of μ and σ .

$$\begin{aligned}\hat{\sigma}_{tj}^* &= \frac{\hat{\sigma}_{tj}}{\sigma} = \frac{\sqrt{\frac{1}{t-1} \sum_{i=1}^t (x_{ij} - \frac{1}{t} \sum_{i=1}^t x_{ij})^2}}{\sigma} = \sqrt{\frac{1}{t-1} \sum_{i=1}^t \left(\frac{x_{ij} - \mu}{\sigma} - \frac{1}{t} \sum_{i=1}^t \frac{x_{ij} - \mu}{\sigma} \right)^2} \\ &= \sqrt{\frac{1}{t-1} \sum_{i=1}^t \left(x_{ij}^* - \frac{1}{t} \sum_{i=1}^t x_{ij}^* \right)^2}\end{aligned}$$

Note that $x_j = x_j^* \sigma + \mu$ and $S_j = S_j^* \sigma + \mu$. If these are substituted in the definition for the simulated performance the following result appears.

$$\begin{aligned}\hat{\beta}(\mathbf{S}_n) &= 1 - \frac{\frac{1}{n} \sum_{j=1}^n (x_j^* \sigma + \mu - (S_j^* \sigma + \mu))^+}{\frac{1}{n} \sum_{j=1}^n x_j^* \sigma + \mu} = 1 - \frac{\frac{1}{n} \sum_{j=1}^n \sigma (x_j^* - S_j^*)^+}{\frac{1}{n} \sum_{j=1}^n x_j^* \sigma + \mu} \\ &= 1 - \frac{\frac{1}{n} \sum_{j=1}^n (x_j^* - S_j^*)^+}{\frac{1}{n} \sum_{j=1}^n x_j^* + \nu^{-1}}\end{aligned}$$

Thus the performance only consists of terms independent of μ and σ . It does, however, clearly depend on ν : directly through ν^{-1} in the denominator and indirectly through c_β^τ in S_j^* .

A.2 $S(\hat{\mu}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$

The only thing that changes with respect to the previous section is that now ν has to be estimated (else \hat{c}_β^τ cannot be found). Using the same line of reasoning results in $\hat{\sigma}_t^* = \hat{\sigma}_t/\sigma$ and $\hat{\mu}_t^* = (\hat{\mu}_t - \mu)/\sigma$ being both independent of μ and σ . So $\hat{\nu}_t = \hat{\sigma}_t/\hat{\mu}_t$ is independent of μ and σ , as is shown below.

$$\hat{\nu}_t = \frac{\hat{\sigma}_t}{\hat{\mu}_t} = \frac{\hat{\sigma}_t^* \sigma}{\hat{\mu}_t^* \sigma + \mu} = \frac{\hat{\sigma}_t^*}{\hat{\mu}_t^* + \nu^{-1}}$$

As stated before \hat{c}_β^τ is defined as follows.

$$\hat{c}_\beta^\tau = \begin{cases} G_u^{-1} \left(\frac{1-\beta}{\hat{\nu}_t \tau} \right) & \text{if } \hat{\nu}_t > 0 \\ -\frac{1}{\hat{\nu}_t \tau} & \text{otherwise} \end{cases}$$

Thus both for $\hat{\nu}_t > 0$ and $\hat{\nu}_t \leq 0$, \hat{c}_β^τ purely depends on terms independent of μ and σ . Hence, \hat{c}_β^τ is independent of μ and σ and it is already shown that the other terms in $(S(\hat{\mu}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau) - \mu)/\sigma$ are. Hence the performance will again be independent of μ and σ .

A.3 $S(\hat{\mu}_t + \hat{\kappa}_i(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau)$

So again $(S(\hat{\mu}_t + \hat{\kappa}_i(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau) - \mu)/\sigma$ ($i = 1, 2$) should be independent of μ and σ . This order-up-to level can be rewritten as below.

$$S(\hat{\mu}_t + \hat{\kappa}_i(\hat{\nu}_t, t, \beta) \hat{\sigma}_t, \hat{c}_\beta^\tau) = S(\hat{\mu}_t, \hat{\sigma}_t \tau, \hat{c}_\beta^\tau) + \hat{\kappa}_i(\hat{\nu}_t, t, \beta) \hat{\sigma}_t$$

Hence it suffices to show that $\hat{\kappa}_i(\hat{\nu}_t, t, \beta) \hat{\sigma}_t/\sigma$ is independent of μ and σ . This is easily seen, if one realizes that $\hat{\nu}_t$ and $\hat{\sigma}_t/\sigma$ are independent of μ and σ . Thus the performance of this order-up-to level is independent of μ and σ .

References

- Artto, K. A. and Pykkänen, E. (1999), An effective procedure for the distribution of magazines, *International Transactions in Operational Research* (6), 289–310.
- Azoury, K. S. (1985), Bayes solution to dynamic inventory models under unknown demand distribution, *Management Science* **31**(9), 1150–1160.
- Azoury, K. S. and Miller, B. L. (1984), A comparison of the optimal ordering levels of bayesian and non-bayesian inventory models, *Management Science* **30**(8), 993–1003.
- Bookbinder, J. H. and Lordahl, A. E. (1989), Estimation of inventory reorder level using the bootstrap statistical procedure, *IIE Transactions* **21**, 302–312.
- Bulinskaya, E. V. (1990), Inventory control in case of unknown demand distribution, *Engineering Costs and Production Economics* **19**, 301–306.
- Fricker, R. D. and Goodhart, C. A. (2000), Applying a bootstrap approach for setting reorder points in military supply systems, *Naval Research Logistics* **47**, 459–478.
- Karmarkar, U. S. (1994), A robust forecasting technique for inventory and leadtime management, *Journal of Operations Management* **12**, 45–54.
- Kottas, J. F. and Lau, H.-S. (1980), The use of versatile distribution families in some stochastic inventory calculations, *The Journal of the Operational Research Society* **31**(5), 393–403.
- Larson, C. E., Olson, L. J. and Sharma, S. (2001), Optimal inventory policies when the demand distribution is not known, *Journal of Economic Theory* **101**, 281–300.
- Lau, H.-S. and Lau, A. H.-L. (2003), Nonrobustness of the normal approximation of lead-time demand in a (Q, R) system, *Naval Research Logistics* **50**, 149–166.
- Liyanage, L. H. and Shanthikumar, J. F. (2005), A practical inventory policy using operational statistics, *Operations Research Letters* **33**, 341–348.
- Lordahl, A. E. and Bookbinder, J. H. (1994), Order-statistic calculation, costs and service in an (s, Q) inventory system, *Naval Research Logistics* **41**, 81–97.
- Sani, B. and Kingsman, B. G. (1997), Selecting the best periodic inventory control and demand forecasting methods for low demand items, *Journal of the Operational Research Society* **48**, 700–713.
- Silver, E. A., Pyke, D. F. and Peterson, R. (1998), *Inventory Management and Production Planning and Scheduling*, third edn, John Wiley & Sons, New York.
- Silver, E. A. and Rahnema, M. R. (1986), The cost effects of statistical sampling in selecting the reorder point in a common inventory model, *Journal of the Operational Research Society* **37**(7), 705–713.
- Silver, E. A. and Rahnema, M. R. (1987), Biased selection of the inventory reorder point when demand parameters are statistically estimated, *Engineering Costs and Production Economics* **12**, 283–292.

- Strijbosch, L. W. G. and Heuts, R. M. J. (1992), Modelling (s, Q) inventory systems: Parametric versus non-parametric approximations for the lead time demand distribution, *European Journal of Operational Research* **63**, 86–101.
- Strijbosch, L. W. G., Heuts, R. M. J. and van der Schoot, E. H. M. (2000), A combined forecast–inventory control procedure for spare parts, *Journal of the Operational Research Society* **51**, 1184–1192.
- Strijbosch, L. W. G. and Moors, J. J. A. (1999), Simple expressions for safety factors in inventory control, CentER Discussion Paper 112, Tilburg University, The Netherlands.
- Strijbosch, L. W. G. and Moors, J. J. A. (2005), The impact of unknown demand parameters on (R, S) -inventory control performance, *European Journal of Operational Research* **162**, 805–815.
- Strijbosch, L. W. G. and Moors, J. J. A. (2006), Modified normal demand distributions in (R, S) -inventory control, *European Journal of Operational Research* **172**, 201–212.
- Strijbosch, L. W. G., Moors, J. J. A. and de Kok, A. G. (1997), On the interaction between forecasting and inventory control, FEW Research Memorandum 742, Tilburg University, The Netherlands.
- Syntetos, A. A. and Boylan, J. E. (2006), On the stock control performance of intermittent demand estimators, *International Journal of Production Economics* **103**, 36–47.
- Tyworth, J. E. and O’Neill, L. (1997), Robustness of the normal approximation of lead-time demand in a distribution setting, *Naval Research Logistics* **44**, 165–186.
- Zeng, A. Z. and Hayya, J. C. (1999), The performance of two popular service measures on management effectiveness in inventory control, *International Journal of Production Economics* **58**, 147–158.
- Zipkin, P. H. (2000), *Foundations of Inventory Management*, McGraw-Hill Higher Education, Boston.