University of Groningen

A general method for addressing forecasting uncertainty in inventory models

Prak, Dennis; Teunter, Ruud

Published in:
International Journal of Forecasting

DOI:
[10.1016/j.ijforecast.2017.11.004](https://doi.org/10.1016/j.ijforecast.2017.11.004)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2019

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Prak, D., & Teunter, R. (2019). A general method for addressing forecasting uncertainty in inventory models. *International Journal of Forecasting*, 35(1), 224-238.
<https://doi.org/10.1016/j.ijforecast.2017.11.004>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

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.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

A general method for addressing forecasting uncertainty in inventory models

Dennis Prak*, Ruud Teunter

Department of Operations, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands

ARTICLE INFO

Keywords:

Demand forecasting
Inventory control
Safety stock
Parameter uncertainty
Bayesian methods

ABSTRACT

In practice, inventory decisions depend heavily on demand forecasts, but the literature typically assumes that demand distributions are known. This means that estimates are substituted directly for the unknown parameters, leading to insufficient safety stocks, stock-outs, low service, and high costs. We propose a framework for addressing this estimation uncertainty that is applicable to any inventory model, demand distribution, and parameter estimator. The estimation errors are modeled and a predictive lead time demand distribution obtained, which is then substituted into the inventory model. We illustrate this framework for several different demand models. When the estimates are based on ten observations, the relative savings are typically between 10% and 30% for mean-stationary demand. However, the savings are larger when the estimates are based on fewer observations, when backorders are costlier, or when the lead time is longer. In the presence of a trend, the savings are between 50% and 80% for several scenarios.

© 2017 International Institute of Forecasters. Published by Elsevier B.V. All rights reserved.

1. Introduction

Inventory control depends heavily on forecasts of the future demand, and yet the inventory control literature exhibits a separation between demand forecasting and inventory decision making. Since Harris (1913) established the economic order quantity model, a wide range of different models have been developed, with varying review structures, cost frameworks, and demand characteristics. Most (medium-sized and large) companies nowadays use such inventory models through either specific inventory control software or more general ERP software. However, inventory models generally rely on a complete certainty of the future demand distribution, which never applies in practice. Although a considerable amount of research has been devoted to the optimal forecasting of various types of demands, the interface between forecasting and decision making remains ill-studied. This paper presents a general

framework for estimating unknown demand parameters and including the estimation uncertainty in the inventory decision. The framework can be applied more generally to any optimization model that depends on some unknown variable that has to be forecasted, but we focus here on inventory control models.

Inventory control textbooks also generally leave the relationship between forecasting and inventory decision making unaddressed, see e.g. Hillier and Lieberman (2014), Waters (2012) and Zipkin (2000). Hax and Candea (1984) discuss how the distribution of demand forecast errors can be derived empirically, by updating the probabilities with incoming demand observations. However, although this yields a consistent estimate of the forecast error distribution, it ignores all uncertainty around that estimate at any point in time. For example, the estimated forecast error distribution after one observation would have all of its mass at zero (unless some prior distribution is used), since the estimated mean of the demand distribution will exactly equal that one observation. Treating point estimates as true parameters means that one ignores part of the uncertainty about demand, whereas inventories are

* Corresponding author.

E-mail addresses: d.r.j.prak@rug.nl (D. Prak), r.h.teunter@rug.nl (R. Teunter).

kept to account for this uncertainty. Inventory calculations that ignore this forecasting uncertainty are flawed, even if the estimators used are unbiased and have the minimum variance. This leads to insufficient safety stocks, resulting in frequent stock-outs, high costs (e.g., backorder costs, lost sales and emergency shipment costs), and not achieving the target service level. Moreover, in practice, the number of demand observations that are available for estimating the demand parameters is often limited, because, firstly, most companies do not store long histories of demand observations, and secondly – and more importantly – the underlying demand process is subject to frequent changes. The fewer historical demand observations that are used, the more volatile the parameter estimates, and thus the greater the negative effect of ignoring their uncertainty.

Bayesian inventory modeling provides an exact framework for incorporating unknown demand parameters into inventory decision models. A class of distributions and a prior parameter distribution are specified, and a predictive parameter distribution is obtained via Bayes' rule. The Bayesian forecasting literature is rich and deals with many different types of demand models, such as normally distributed dynamic models and autoregressive models, but it also deals with more general models with demand distributions that belong to the exponential family. Notably, though, applications of Bayesian theory in inventory modeling are very rare. Exact treatments exist only for simple single-period and/or single-parameter demand models. If the demand distribution contains several parameters, then typically some of them (e.g., the variance for normally distributed demand) will be considered to be known. This restricts the applicability of Bayesian inventory theory, while the negative effects of ignoring parameter uncertainty are typically largest in models with several parameters, multiple periods and/or positive lead times (since the forecast errors over the forecast horizon are correlated). Furthermore, in practice, it is difficult to specify the demand distribution exactly, making methods such as exponential smoothing popular. Standard Bayesian methods do not exhibit freedom in the choice of parameter estimator, but rather fix the demand distribution and choose a prior parameter distribution. Finally, the monetary consequences of ignoring parameter uncertainty in inventory decisions have not been studied much in the literature, as the focus in the forecasting literature has been on the accuracy of demand predictions, whereas the inventory control literature has typically ignored forecasting altogether.

This paper has two aims. First, we present a framework for addressing estimation uncertainty in inventory models that starts from the chosen parameter estimators, rather than from a prior distribution. The distributions of the unknown parameters are modeled using Bayes' rule on the sample distributions of the parameter estimators. By taking the expectation of the lead time demand distribution function with respect to the (now stochastic) unknown parameters, we obtain a predictive lead time demand distribution which can be applied in the inventory decision model. We assume that no information about the unknown parameters is available other than the available data. Second, we perform an extensive numerical study to

show that the cost of ignoring the parameter uncertainty is substantial for many widely-used inventory models and demand distributions. Furthermore, we demonstrate how this cost compares to that resulting from a misspecification of the demand model or the selection of suboptimal parameter estimators.

We demonstrate this method for a discrete-time, continuous-review inventory model with a fixed lead time and linear holding and backorder costs. The distribution of the estimation errors can be found exactly in the case of a normally distributed demand, but we will also discuss an approximative method, which is useful for two reasons. Firstly, the method is robust to demand distribution misspecification. Secondly, an exact derivation of the parameter estimation errors may not always be available, whereas the approximative method follows straightforwardly for every maximum likelihood estimator.

The remainder of this paper is structured as follows. Section 2 provides a literature review. Section 3 presents the modeling framework, the inventory model and the demand models, and derives the order levels according to both the classical and corrected (both exact and approximative) approaches. Section 4 performs a numerical study comparing the performances of the different approaches. Section 5 discusses the robustness of the approaches to misspecification of the demand model. Section 6 concludes.

2. Related literature

The mismatch between forecasting and inventory control has been pointed out occasionally by previous researchers. Fildes and Beard (1992) discuss the correlation between future forecast errors explicitly, while Silver, Pyke, and Peterson (1998) remark that the variability of the forecast error depends in a complicated fashion on both the demand model and the forecast procedure. Srijbosch, Heuts, and van der Schoot (2000) discuss the importance of deriving the lead time forecast error rather than only the forecast error per time instant. Toktay and Wein (2001) describe a production model where demand forecast updates are incorporated in the production decision, and also discuss the effect of forecast errors. All of these authors focus on the mean of demand. Beutel and Minner (2012) discuss an integrated framework of least squares demand forecasting and inventory decision making, and find that if the demand parameters are replaced by estimates that are made based on only a few observations, the actual service levels will undershoot their targets significantly. Prak, Syntetos, and Teunter (2017) provide adjustments to the standard calculations in order to incorporate the uncertainty of both parameters in the case of a mean-stationary normally distributed demand, for a service level model. They show that the service levels achieved by classical approaches undershoot their targets significantly, especially in situations where few previous demand observations are available. However, their method holds solely for mean-stationary, normally distributed demand, under a service level constraint.

The forecasting literature has paid a considerable amount of attention to Bayesian approaches to demand

forecasting. When applying Bayesian forecasting, one is seeking to derive a predictive demand distribution, rather than just a point forecast. This is key to addressing the estimation uncertainty in the eventual inventory decision properly. A broad discussion of Bayesian forecasting is provided by [West and Harrison \(1997\)](#), who present predictive distributions for dynamic linear models, autoregressive models, and non-linear models, among others. Key articles in this stream are those by [Harrison and Stevens \(1971\)](#) and [West, Harrison, and Migon \(1985\)](#). Bayesian forecasting is especially convenient when dealing with varying model parameters, whereas traditional methods typically assume that these remain constant for at least some time. Specific applications to demand forecasting are provided by [Spedding and Chan \(2000\)](#) and [Yelland \(2010\)](#).

Although applications of Bayesian inventory control were pioneered by [Dvoretzky, Kiefer, and Wolfowitz \(1952\)](#) and [Scarf \(1959\)](#), they still remain scarce. Numerical Bayesian applications to inventory control involve mainly one-parameter demand distributions, or, if the distribution has more than one parameter, impose the restriction that at least one parameters is known. [Azoury and Miyaoka \(2009\)](#) and [Chen \(2010\)](#) use a normal distribution of which the standard deviation is assumed to be known, while [Rajashree Kamath and Pakkala \(2002\)](#) use a lognormal distribution with a known standard deviation. This restriction is convenient because it ensures that the posterior demand distribution is again (log)normal, so that the analysis is similar to that without the Bayesian treatment of the unknown mean. [Azoury and Miyaoka \(2009\)](#) state that uncertainty in the demand standard deviation is of minor importance, citing [Gelman et al. \(1995\)](#). However, [Gelman et al. \(1995\)](#) discuss mainly regression models from economics, where the standard deviation only affects the standard errors of the estimated parameters. In inventory control, on the other hand, the standard deviation of demand relates directly to the safety stock (a key decision variable), and therefore plays a vital role. Another complication of applying Bayesian methods to inventory control is that one has to specify a prior distribution, which is a subjective element. [Hill \(1997, 1999\)](#) use a non-informative prior.

The inventory control literature also contains a few non-Bayesian studies about decision making under parameter uncertainty. [Ritchken and Sankar \(1984\)](#) estimate the quantiles of the normal distribution with an unknown mean and variance consistently in order to optimize the reorder level in a single-period inventory model. [Janssen, Strijbosch, and Brekelmans \(2009\)](#) perform a simulation study involving adding a mark-up to the traditional reorder level. [Strijbosch and Moors \(2005\)](#) search numerically for the optimal batch size in an (r, Q) policy with a normally distributed demand and unknown parameters. [Strijbosch, Syntetos, Boylan, and Janssen \(2011\)](#) find that, for non-stationary demand series, the estimated demand variance again needs to be adjusted to account for the uncertainty in the parameter estimates. They use a simple moving average (SMA) and single exponential smoothing (SES) on a large historical dataset of demands for several products, and find that, for the majority of the items, time horizons of two to eight periods are optimal for SMA, and smoothing

parameters between 0.21 and 0.4 are optimal for SES (note that the search range was restricted to values between 0 and 0.4). This suggests that, in practice, estimates can often be based on only a few demand observations, which is in accordance with our own experience. In all of these studies, the conclusion is that ignoring the forecasting uncertainty leads to inventories that are too low, and to undershooting target service levels.

[Fildes \(1983\)](#) found that a multivariate Bayesian forecasting system typically does not outperform classical methods such as exponential smoothing. It has also been observed more recently that exponential smoothing performs well in forecasting competitions ([Ali & Boylan, 2011, 2012](#); [Gardner, 2006](#)), which may be one of the reasons why practitioners typically resort to classical methods. [Bermúdez, Segura, and Vercher \(2010\)](#) discuss the fact that the prediction intervals generated by maximum likelihood in the Holt-Winters model are too narrow, and show that a Bayesian approach performs better. This points to the key problem with using classical estimation methods for inventory decisions: whereas Bayesian forecasting leads to predictive distributions that naturally include estimation uncertainty, classical methods do not. [Ord, Koehler, and Snyder \(1997\)](#) discuss several ways of overcoming this issue, and find that a simulation method – which draws possible ‘true’ parameters based on observed estimates – provides the most reliable results in terms prediction intervals for future demand. However, the effects on the eventual inventory decision and the resulting costs are not addressed.

In conclusion, although Bayesian methods of deriving predictive forecast distributions exist in the forecasting literature, the inventory control literature applies these only in limited contexts. When making inventory decisions, practitioners typically apply classical methods and ignore estimation uncertainty, and there exists little insight into the severity of this problem. Whichever forecasting method is used, the estimation uncertainty should be taken into account in the inventory decision. We present a general approach for facilitating this in any combination of demand model and inventory system. Starting from a chosen parameter estimator, we perform an analysis along classical lines based on the sample distributions of the estimators, but use Bayes’ rule to find predictive parameter distributions and eventually a predictive lead time demand distribution. We assume no prior information on the parameters. Furthermore, in contrast to [Ord et al. \(1997\)](#), we perform an extensive numerical study of the resulting inventory costs. We benchmark the performance of the corrected approach against that of the classical approach that ignores estimation uncertainty, in an attempt to obtain insights into the consequences of ignoring estimation uncertainty, and their relationship to other costs, especially those of misspecifying the demand model or choosing a suboptimal parameter estimator.

3. Model

This section begins by describing the general framework that we apply throughout this paper, then discusses the specific inventory model and the several demand models and parameter estimators which we use to demonstrate

the framework. We derive the optimal policy under uncertainty of the demand parameters, and also discuss a robust approximative method based on the asymptotic sample distributions of the parameter estimators.

3.1. General framework

The approach to handling unknown demand parameters that is taken in this paper starts from a cost function that is to be minimized, and that depends in some way on a (lead time) demand distribution with unknown parameters which is transformed into a predictive demand distribution as follows. Denote the vector of parameters by θ and the required distribution function by $F(x; \theta)$. We estimate θ by $\hat{\theta}$, and rewrite θ as a function of $\hat{\theta}$ and the estimation error ϵ_θ . The estimation error ϵ_θ is a random vector of which we either derive or approximate the probability distribution. This is achieved via the sample distribution of the estimator, or its asymptotic distribution (which is available for every maximum likelihood estimator, for example), and then Bayes' rule is used with an uninformative prior to transform that distribution into the distribution of the unknown parameter. The predictive demand distribution function is now $E_\theta(F)$, where, for the sake of brevity, our notation omits the fact that the distribution of θ is conditional on the observed data. Numerically, this boils down to an integration with an order that is equal to the number of unknown parameters. The inventory decision model should use this predictive distribution to arrive at the inventory decision. If the demand model contains many unknown parameters (such as in a multiple regression model), then the lead time demand distribution may be reparameterized so as to reduce the dimension of integration. We will discuss this further in Section 3.3.3. Finally, we remark that classical approaches arrive at a predictive distribution by simply substituting the point estimates thus obtained for the true parameters.

3.2. Inventory model

We will illustrate our approach by focusing on a discrete-time, single-item, continuous-review inventory model with fixed linear holding and shortage costs. That is, the time horizon is $t = 1, 2, \dots$, the holding costs per unit per time unit are $h > 0$, the shortage costs per unit per time unit are $b \geq h$, and full back-ordering is assumed. Orders can be placed free of charge at every time instant and arrive after a fixed lead time $L \geq 0$. Because there are no fixed order costs, there is no need for batch orders, and so-called base-stock or order-up-to level policies can be applied that keep the inventory position (the current inventory level plus outstanding orders) constant. An optimal order-up-to level is derived at every review, and the order of events at any review time instant is as follows: first outstanding orders arrive, then demand occurs and costs are incurred, and finally new orders can be placed. In principle, it is possible that the order-up-to level could be smaller than the current inventory position, but it is unlikely, since it can only occur if the difference between the optimal order-up-to levels from one period to the next is larger than

the demand realization. In these rare cases, we allow for negative orders.

The costs arising from the inventory model at any time instant are easy to characterize. Denote the inventory level (the inventory on hand) at time t , after the arrival of the demand at time t , by $IL(t)$. Then, the costs at time t are given by

$$C(t) = h(IL(t))^+ + b(IL(t))^-,$$

where $(x)^+ = \max\{x, 0\}$ and $(x)^- = \max\{-x, 0\}$.

We refer to the current time instant as n , and derive an optimal order-up-to policy that is characterized by an order-up-to level S_n for the current inventory position, which includes outstanding orders. Furthermore, we assume that we are currently at the end of time n , and that we have observed demands at times $1, \dots, n$. Orders that are placed at time n will arrive at the beginning of time $n + L$. The next possible order could be placed at time $n + 1$, and will arrive at the beginning of time $n + L + 1$. Thus, decisions made at time n affect the costs at time $n + L$, which are driven by the inventory level directly after the arrival of the demand at time $n + L$, but before the next ordering, which we denote by $IL(n + L)$, as above. Therefore, the inventory level at time $n + L$ is given by the inventory position (the inventory level plus outstanding orders) after the ordering at the end of time n , minus the demands at times $n + 1, \dots, n + L$. Denote the inventory position at the end of time n by $IP(n)$ and denote the sum of the demands at times $k, \dots, m$ by $D_{[k,m]}$. Then, we can write

$$IL(n + L) = IP(n) - D_{[n+1, n+L]}.$$

We find that at time n , we face the problem of minimizing

$$TC(S_n) = hE(S_n - D_{[n+1, n+L]})^+ + bE(S_n - D_{[n+1, n+L]})^-.$$

Observe that this model can be decomposed, in the sense that the optimal order-up-to level at time n can be found independently of both former and future order-up-to levels. Therefore, in what follows we omit the subscript n and refer to the optimal order-up-to level at time n as S . We are now back in the framework of Section 3.1. The model discussed can be extended easily to a periodic review framework or a setting where the time horizon is continuous rather than discrete, for example.

3.2.1. Decision making

Denote the (predictive) distribution function of lead time demand $D_{[n+1, n+L]}$ by $F_{D_{[n+1, n+L]}}$. Since

$$\begin{aligned} E(S - D_{[n+1, n+L]})^+ &= \int_{-\infty}^S (S - s) dF_{D_{[n+1, n+L]}}(s) \\ &= \int_{-\infty}^S \int_s^S dx dF_{D_{[n+1, n+L]}}(s) \\ &= \int_{-\infty}^S \int_{-\infty}^x dF_{D_{[n+1, n+L]}}(s) dx \\ &= \int_{-\infty}^S F_{D_{[n+1, n+L]}}(x) dx, \end{aligned}$$

and similarly

$$E(S - D_{[n+1, n+L]})^- = \int_S^\infty [1 - F_{D_{[n+1, n+L]}}(x)] dx,$$

we can write

$$TC(S; \mu, \sigma) = h \int_{-\infty}^S F_{D_{[n+1, n+L]}}(x) dx + b \int_S^{\infty} [1 - F_{D_{[n+1, n+L]}}(x)] dx.$$

Using

$$\frac{d}{dS} \int_{-\infty}^S F_{D_{[n+1, n+L]}}(x) dx = F_{D_{[n+1, n+L]}}(S),$$

and

$$\frac{d}{dS} \int_S^{\infty} [1 - F_{D_{[n+1, n+L]}}(x)] dx = -[1 - F_{D_{[n+1, n+L]}}(S)],$$

we find the first order condition

$$F_{D_{[n+1, n+L]}}(S) = \frac{b}{b+h}, \quad (1)$$

which is a newsvendor equation and can be solved for S . Since

$$\frac{d^2 TC}{dS^2} = (b+h)f_{D_{[n+1, n+L]}}(S) > 0,$$

it follows that TC is convex and the first order condition leads to the global minimum.

Under the assumption of i.i.d. normally distributed demand with a known mean μ and standard deviation σ per time unit, the equation defining the optimal order-up-to level has a closed-form solution which is given by

$$S = L\mu + \sqrt{L}\Phi^{-1}\left(\frac{b}{b+h}\right)\sigma,$$

where Φ denotes the distribution function of the standard normal distribution. Classical inventory control methods substitute point estimates directly for the unknown demand parameters in Eq. (1), then calculate the ‘optimal’ order-up-to level, which we denote by $\hat{S}$. However, this order-up-to level is not optimal under parameter uncertainty. The optimal inventory decision that is derived using the predictive distribution is denoted by S^* . The remainder of this section is concerned with deriving such predictive distributions that can be used in Eq. (1) in order to find the optimal order-up-to level under parameter uncertainty.

3.3. Demand models and derivation of the predictive lead time demand distribution

The inventory model described in the previous section depends on the distribution of demand in the time interval $[n+1, n+L]$. This section presents four different demand models, derives the corresponding lead time demand distributions, discusses how the parameters can be estimated and their estimation errors modeled, and finally considers how a predictive lead time demand distribution can be obtained. The choice of a certain demand model is typically made either manually by the decision maker, or automatically using forecasting software by comparing goodness-of-fit measures. We first discuss the commonly used model of mean-stationary demand, where we compare the sample mean and the SES estimator. Whereas the former is the minimum variance unbiased estimator

for this model, the latter is popular because in practice it is often not known whether the demand is stationary, and if so, for how long it has been so. Next, we consider the trended demand model, which is selected when the demand appears to be stationary around a linear trend. This is also a common choice in textbooks and forecasting software, for example for fashion-sensitive products that show an upward trend at the beginning of their sales cycle and a downward trend at the end. Finally, we consider the random walk model, where the next demand observation is equal to the previous observation plus a random noise term. This model is realistic for products with demand patterns that are non-stationary and do not follow specific trends.

3.3.1. Mean-stationary, normally distributed demand

Denoting the demand at time t by D_t , we assume that $D_t \stackrel{\text{i.i.d.}}{\sim} N(\mu, \sigma^2)$, where μ and σ^2 are unknown. Models with normally distributed demands have the drawback that the demand can theoretically be negative (although the probability of that occurring can be made arbitrarily small); nevertheless, they are applied widely in the literature, and especially in textbooks, such as those by [Axsäter \(2006\)](#) and [Silver et al. \(1998\)](#). It now follows easily that

$$D_{[n+1, n+L]} \stackrel{\text{i.i.d.}}{\sim} N(L\mu, L\sigma^2).$$

The unknown parameters μ and σ^2 are estimated efficiently by the sample mean and sample variance; that is,

$$\hat{\mu} = \frac{1}{n} \sum_{t=1}^n D_t$$

and

$$\hat{\sigma}^2 = \frac{1}{n-1} \sum_{t=1}^n (D_t - \hat{\mu})^2.$$

Observe that $\hat{\sigma}^2$ is not equal to the maximum likelihood estimator for σ^2 , which divides by n rather than by $n-1$. As was discussed by [Davidson and MacKinnon \(2003\)](#), for example, $\hat{\sigma}^2$ is unbiased, whereas the maximum likelihood estimator is not. However, $\hat{\sigma}^2$ still has the same asymptotic distribution as the maximum likelihood estimator, which we use later for our approximate modeling approach.

All that remains is to find the exact relationship between the true parameters μ and σ^2 , and their estimates $\hat{\mu}$ and $\hat{\sigma}^2$. We start from the following results, which are proven by [Miller and Miller \(2004\)](#), for example:

$$\frac{\mu - \hat{\mu}}{\sigma/\sqrt{n}} \sim N(0, 1),$$

and, independently,

$$\frac{(n-1)\hat{\sigma}^2}{\sigma^2} \sim \chi_{n-1}^2,$$

where χ_{n-1}^2 denotes a χ^2 distribution with $n-1$ degrees of freedom. Given σ^2 , we can now write

$$\mu = \hat{\mu} + \frac{\sigma Z}{\sqrt{n}},$$

where Z is standard normally distributed, and

$$\sigma^2 = \frac{(n-1)\hat{\sigma}^2}{X},$$

where X is χ^2_{n-1} -distributed. This defines the predictive parameter distribution as necessary in the framework described in Section 3.1. The integral is taken first with respect to σ^2 and then with respect to μ . We denote the resulting order-up-to level by S^* .

This modeling approach fits in the Bayesian framework when an uninformative prior is used. If the normal distribution is assumed for demand and an uninformative prior distribution is used, then the posterior distribution of μ , given σ^2 , is normal with mean $\hat{\mu}$ and variance σ^2/n . Also, the marginal posterior distribution of σ^2 is scaled inverse χ^2 with scale parameter $\hat{\sigma}^2$ and $n-1$ degrees of freedom, see e.g. Gelman et al. (1995). With Bayesian analysis, one can show that using a non-central Student's- t distribution (the posterior distribution that results when a non-informative prior is used on a normal distribution) as the predictive lead time demand distribution leads to the same results. However, this one-to-one correspondence with a standard Bayesian analysis does not always exist.

If it is not possible or desirable to derive the distribution of the estimation error exactly, then one can resort to an asymptotic approximation of the error distribution. Every maximum likelihood estimator is asymptotically normal with a variance that follows directly from the Fisher information matrix. Since the estimator for μ in the previous section is the maximum likelihood estimator and the estimator for σ^2 is asymptotically equivalent to the maximum likelihood estimator, we can use maximum likelihood theory to derive the asymptotic distributions of the estimators. These asymptotic distributions hold even if the demand distribution is not normal, making this approach robust and practically appealing. However, note that this alternative approach is sub-optimal if the demand distribution is specified correctly and an exact derivation is available.

As was shown by Greene (2011), in our model,

$$\hat{\mu} - \mu \stackrel{a}{\sim} N(0, \sigma^2/n)$$

and

$$\hat{\sigma}^2 - \sigma^2 \stackrel{a}{\sim} N(0, 2\sigma^4/n).$$

Replacing σ^2 by $\hat{\sigma}^2$ results in a consistent estimator of these asymptotic distributions. Hence, we find $\mu \approx \hat{\mu} + \sqrt{\frac{\hat{\sigma}^2}{n}}Z_1$ and $\sigma^2 \approx \hat{\sigma}^2 + \sqrt{\frac{2\hat{\sigma}^4}{n}}Z_2$, where Z_1 and Z_2 are independent and standard normally distributed. We are now back in the framework of the previous section.

Note that although σ^2 is restricted to be positive, its normal distribution model allows the integral to be evaluated at negative values of the demand variance. As the distribution function is not defined for negative values of the variance, this will lead to an overestimation of the true demand variance, and therefore to safety stocks that are too high. We correct for this by using a truncated normal distribution with appropriately bounded support to model Z_2 , and denote the resulting order-up-to level by S' .

3.3.2. Single exponential smoothing for mean-stationary, normally distributed demand

Consider the same demand model as above, $D_t \stackrel{i.i.d.}{\sim} N(\mu, \sigma^2)$, but now estimate μ using single exponential smoothing (SES) with a smoothing parameter of α . This method is applied frequently in practice, as it is often not known with certainty whether demand has been stationary, and if it is, for how long it has been. The variance σ^2 is again estimated by the sample variance. As above, we are interested in the lead time demand distribution, which is characterized by

$$D_{[n+1, n+L]} \stackrel{i.i.d.}{\sim} N(L\mu, L\sigma^2).$$

The SES estimator at time n can be written in expanded form as

$$\hat{\mu} = \alpha(D_n + (1-\alpha)D_{n-1} + (1-\alpha)^2D_{n-2} + \dots + (1-\alpha)^{n-2}D_2) + (1-\alpha)^{n-1}D_1,$$

and its sample distribution is normal with mean μ and variance

$$\begin{aligned} \text{Var}(\hat{\mu}) &= \sigma^2 \alpha^2 \sum_{i=0}^{n-2} (1-\alpha)^{2i} + \sigma^2 (1-\alpha)^{2(n-1)} \\ &= \sigma^2 \frac{\alpha - \alpha^2 + 2(1-\alpha)^{2n}}{(\alpha-2)(\alpha-1)}. \end{aligned}$$

We can now write, conditional on the true variance,

$$\mu = \hat{\mu} + \sqrt{\text{Var}(\hat{\mu})}Z,$$

where Z is standard normally distributed. Since we estimate σ^2 by the sample variance, we still have

$$\sigma^2 = \frac{(n-1)\hat{\sigma}^2}{X},$$

where X is χ^2_{n-1} -distributed. Observe that, for this sample distribution to hold true, the deviations of the demand observations when estimating σ^2 should be taken with respect to the sample mean, not with respect to the SES estimator for μ . We have now completely specified the predictive parameter distribution.

The approximate modeling method in the case of the SES estimator uses its asymptotic distribution. By letting $n \rightarrow \infty$ in the expression for the sample variance of the SES estimator, we find the familiar result that its asymptotic variance is $\frac{\alpha}{2-\alpha}\sigma^2$. Replacing σ^2 by its estimator, we find that the asymptotic variance can be estimated by $\frac{\alpha}{2-\alpha}\hat{\sigma}^2$.

Thus, we find that $\mu \approx \hat{\mu} + \sqrt{\frac{\hat{\sigma}^2 \alpha}{2-\alpha}}Z_1$ and $\sigma^2 \approx \hat{\sigma}^2 + \sqrt{\frac{2\hat{\sigma}^4}{n}}Z_2$, where Z_1 and Z_2 are independent and standard normally distributed. As before, we use a truncated normal distribution for σ^2 .

3.3.3. Normally distributed, trended demand

As a natural extension of the mean-stationary demand model, let us now consider a trended demand model. Denoting the demand at time t by D_t , we assume that the demand is distributed normally and may have a linear trend, $D_t = \alpha + t\beta + v_t$, $v_t \stackrel{i.i.d.}{\sim} N(0, \sigma^2)$. The parameters α , β , and σ^2 are unknown and have to be estimated.

In this model, we find that $D_{[n+1, n+L]}$ is normal with mean

$$\begin{aligned}\mu_{\text{lead}} &= L\alpha + (n+1 + \dots + n+L)\beta \\ &= L\alpha + \frac{1}{2}(L^2 + 2nL + L)\beta\end{aligned}$$

and variance $L\sigma^2$. The variance of lead-time demand is again easy to deal with, but the mean requires some further analysis.

It follows from OLS theory (Davidson & MacKinnon, 2003) that, for a general regression model $y_t = x_t'\beta + \epsilon_t$, the OLS estimator for the parameter vector β is given by $\hat{\beta} = (X'X)^{-1}X'y$, where the vector y contains all elements y_t and the matrix X contains all row vectors x_t . Furthermore, the error variance σ^2 is estimated by

$$\hat{\sigma}^2 = \frac{(y - X\hat{\beta})'(y - X\hat{\beta})}{n - k},$$

where k denotes the number of regressors in the model. It can be shown that, conditional on the true variance, we can write $\beta = \hat{\beta} + \epsilon_\beta$, for

$$\epsilon_\beta \sim \text{MVN}(0, \sigma^2(X'X)^{-1}),$$

where MVN denotes a multivariate normal distribution. We also have that $\sigma^2 = \frac{\hat{\sigma}^2(n-k)}{\epsilon_{\sigma^2}}$, where $\epsilon_{\sigma^2} \sim \chi_{n-k}^2$.

Applying the OLS theory discussed above to our trend model, we find that at time n the parameters α and β are estimated by

$$\begin{pmatrix} \hat{\alpha} \\ \hat{\beta} \end{pmatrix} = \begin{pmatrix} n & \sum_{t=1}^n t \\ \sum_{t=1}^n t & \sum_{t=1}^n t^2 \end{pmatrix}^{-1} \begin{pmatrix} \sum_{t=1}^n D_t \\ \sum_{t=1}^n tD_t \end{pmatrix}. \quad (2)$$

Then, the OLS predictor of demand at time $t > n$ is $\hat{\mu}_t = \hat{\alpha} + t\hat{\beta}$, and that of the lead time demand is $\hat{\mu}_{\text{lead}} = L\hat{\alpha} + \frac{1}{2}(L^2 + 2nL + L)\hat{\beta}$.

The covariance matrix of the OLS estimators for α and β is given by

$$\begin{aligned}V &= \sigma^2 \begin{pmatrix} n & \sum_{t=1}^n t \\ \sum_{t=1}^n t & \sum_{t=1}^n t^2 \end{pmatrix}^{-1} \\ &= \frac{\sigma^2}{n \sum_{t=1}^n t^2 - (\sum_{t=1}^n t)^2} \begin{pmatrix} \sum_{t=1}^n t^2 & -\sum_{t=1}^n t \\ -\sum_{t=1}^n t & n \end{pmatrix} \\ &= \sigma^2 \begin{pmatrix} \frac{4n+2}{n(n-1)} & \frac{6}{n(n-1)} \\ \frac{6}{n(n-1)} & \frac{12}{n(n-1)} \end{pmatrix},\end{aligned}$$

where we use that

$$\sum_{t=1}^n t = \frac{n(n+1)}{2}$$

and

$$\sum_{t=1}^n t^2 = \frac{n(n+1)(2n+1)}{6}.$$

Since the predictor of the mean of the lead time demand is a linear transformation of $\hat{\alpha}$ and $\hat{\beta}$, conditional on the true demand variance, it is normally distributed with mean μ_{lead} and variance

$$\begin{aligned}\text{Var}(\hat{\mu}_{\text{lead}}) &= \left(L \quad \frac{1}{2}(L^2 + 2nL + L) \right) \\ &\quad \times V \begin{pmatrix} L \\ \frac{1}{2}(L^2 + 2nL + L) \end{pmatrix}.\end{aligned}$$

Again conditional on the true variance, we can now write

$$\mu_{\text{lead}} = \hat{\mu}_{\text{lead}} + \sqrt{\text{Var}(\hat{\mu}_{\text{lead}})}Z,$$

where Z is standard normally distributed. For σ^2 , we find that

$$\sigma^2 = \frac{(n-2)\hat{\sigma}^2}{X},$$

where X is χ_{n-2}^2 -distributed. This specifies the predictive parameter distribution. Observe that, by modeling μ_{lead} directly, rather than α and β , we have reduced the dimension of the integral leading to the predictive lead time demand distribution from three to two, thus simplifying its use in practice.

Replacing σ^2 with its estimator in $\text{Var}(\hat{\mu}_{\text{lead}})$ leads straightforwardly to an estimator $\widehat{\text{Var}}(\hat{\mu}_{\text{lead}})$ of the asymptotic variance of $\hat{\mu}_{\text{lead}}$. Again, the asymptotic variance of $\hat{\sigma}^2$ is estimated by $2\hat{\sigma}^4/n$. Hence, we find $\mu_{\text{lead}} \approx \hat{\mu}_{\text{lead}} + \sqrt{\widehat{\text{Var}}(\hat{\mu}_{\text{lead}})}Z_1$ and $\sigma^2 \approx \hat{\sigma}^2 + \sqrt{\frac{2\hat{\sigma}^4}{n}}Z_2$, where Z_1 and Z_2 are independent and standard normally distributed. For σ^2 , we again use a truncated normal distribution.

3.3.4. Demand is a normal random walk

The trend model described in the previous subsection still assumes that the demands at different time instants are independent of each other. Here, we demonstrate the procedure for a typical model of autocorrelated demand. Denoting the demand at time t by D_t , we assume that

$$D_t = D_{t-1} + v_t,$$

where $v_t \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma^2)$. That is, demand follows a random walk and the increments are distributed normally. The parameter σ^2 is unknown and has to be estimated based on n previous observations. In this model we find that, conditional on D_n ,

$$D_{[n+1, n+L]} = LD_n + Lv_{n+1} + (L-1)v_{n+2} + \dots + v_{n+L},$$

which is normally distributed with mean LD_n and variance

$$\begin{aligned}\sigma_{\text{lead}}^2 &= \sigma^2(L^2 + (L-1)^2 + (L-2)^2 + \dots + 1) \\ &= \sigma^2 L(L+1)(2L+1)/6.\end{aligned}$$

The unknown parameter σ^2 can be estimated by the sample variance of the difference series of the previous n

Table 1

Numerical results: mean-stationary, normally distributed demand.

Parameters						Order-up-to levels			Safety stock diff.		Total expected holding and shortage costs				
$\hat{\mu}$	$\hat{\sigma}^2$	n	b	h	L	Class.	Approx.	Exact	Approx.	Exact	Class.	Approximate		Exact	
10	4	5	20	1	5	57.5	61.0	63.8	+47.9%	+84.4%	26.1	21.3	(−18.2%)	20.5	(−21.2%)
10	4	10	20	1	5	57.5	59.2	60.2	+23.9%	+36.8%	15.0	13.7	(−8.4%)	13.6	(−9.4%)
10	4	20	20	1	5	57.5	58.4	58.8	+12.0%	+17.7%	11.6	11.3	(−2.9%)	11.3	(−3.1%)
10	4	100	20	1	5	57.5	57.6	57.7	+2.4%	+3.5%	9.7	9.7	(−0.1%)	9.7	(−0.2%)
10	4	5	50	1	5	59.2	64.2	69.1	+54.2%	+107.1%	42.9	29.7	(−30.7%)	27.1	(−36.6%)
10	4	10	50	1	5	59.2	61.8	63.2	+27.6%	+43.2%	20.2	17.0	(−16.1%)	16.6	(−17.8%)
10	4	20	50	1	5	59.2	60.5	61.1	+13.9%	+20.1%	14.3	13.5	(−5.6%)	13.4	(−6.1%)
10	4	100	50	1	5	59.2	59.5	59.6	+2.8%	+3.9%	11.4	11.3	(−0.3%)	11.3	(−0.3%)
10	4	5	100	1	5	60.4	66.5	73.8	+58.3%	+128.3%	65.4	38.9	(−40.5%)	33.1	(−49.4%)
10	4	10	100	1	5	60.4	63.6	65.5	+30.1%	+48.7%	25.8	19.7	(−23.6%)	19.0	(−26.3%)
10	4	20	100	1	5	60.4	62.0	62.7	+15.3%	+22.1%	16.5	15.1	(−8.6%)	15.0	(−9.3%)
10	4	100	100	1	5	60.4	60.7	60.9	+3.2%	+4.2%	12.5	12.5	(−0.4%)	12.5	(−0.4%)
10	4	5	20	1	10	110.6	119.1	123.8	+81.2%	+125.9%	51.4	37.0	(−28.1%)	35.6	(−30.7%)
10	4	10	20	1	10	110.6	115.1	116.7	+43.1%	+58.0%	26.9	22.5	(−16.6%)	22.2	(−17.4%)
10	4	20	20	1	10	110.6	112.9	113.6	+22.7%	+28.9%	18.8	17.5	(−7.0%)	17.5	(−7.3%)
10	4	100	20	1	10	110.6	111.1	111.2	+4.9%	+5.9%	14.1	14.0	(−0.4%)	14.0	(−0.4%)
10	1	5	20	1	5	53.7	55.5	56.9	+47.9%	+84.4%	13.0	10.7	(−18.1%)	10.3	(−21.1%)
10	1	10	20	1	5	53.7	54.6	55.1	+23.9%	+36.8%	7.5	6.9	(−8.5%)	6.8	(−9.5%)
10	1	20	20	1	5	53.7	54.2	54.4	+12.0%	+17.7%	5.8	5.7	(−2.9%)	5.6	(−3.2%)
10	1	100	20	1	5	53.7	53.8	53.9	+2.4%	+3.5%	4.9	4.8	(−0.1%)	4.8	(−0.2%)
20	4	5	20	1	5	107.5	111.0	113.8	+47.9%	+84.4%	26.1	21.4	(−18.2%)	20.6	(−21.2%)
20	4	10	20	1	5	107.5	109.2	110.2	+23.9%	+36.8%	15.0	13.8	(−8.4%)	13.6	(−9.4%)
20	4	20	20	1	5	107.5	108.4	108.8	+12.0%	+17.7%	11.6	11.3	(−2.9%)	11.3	(−3.1%)
20	4	100	20	1	5	107.5	107.6	107.7	+2.4%	+3.5%	9.7	9.7	(−0.1%)	9.7	(−0.2%)

observations. The $n - 1$ differences are each i.i.d. normally distributed with mean 0 and variance σ^2 , so that the theory for σ^2 , as derived for the mean-stationary model, still holds. For the exact approach we can write

$$\sigma^2 = \frac{(n-2)\hat{\sigma}^2}{X},$$

where X is χ_{n-2}^2 -distributed, and for the approximate approach we have $\sigma^2 \approx \hat{\sigma}^2 + \sqrt{\frac{2\hat{\sigma}^4}{n}}Z$, where Z is standard normally distributed. Replacing that standard normal distribution by a truncated normal distribution, we are back in the familiar framework, and the predictive lead time demand distribution follows directly.

4. Numerical study

This section numerically evaluates the order-up-to levels and costs that result from the corrected and classical approaches for the inventory model under study and the four demand models. For each set-up, we start from the observed parameter estimates and derive the order-up-to levels S^* (corresponding to the exact, corrected approach, based on the exact sample distribution of the estimator), S' (corresponding to the approximate, corrected approach, based on the asymptotic sample distribution of the estimator), and $\bar{S}$ (corresponding to the classical approach, which ignores parameter uncertainty). We then perform 1,000,000 replications in which true parameters are drawn according to the earlier derived probability distributions, and the costs are evaluated based on the full information cost function using these true parameters. We report the differences between the three methods in terms of order-up-to levels and resulting average costs, as well as the

differences in terms of safety stock, by subtracting the estimated mean of the lead time demand from the order-up-to levels.

4.1. Mean-stationary, normally distributed demand

We start with the mean-stationary demand model, define a base case of parameter settings, and then vary one parameter at a time. The base case of the comparison has $\hat{\mu} = 10$ and $\hat{\sigma}^2 = 4$. The holding cost per unit per time unit is 1, while the shortage cost is 20. The lead time is five time units. The other scenarios are constructed by varying the shortage costs from 50 to 100, doubling the lead time to 10 time units, decreasing $\hat{\sigma}^2$ to 1, and doubling $\hat{\mu}$ to 20. Within every scenario, the number of historical observations corresponding to the current demand distribution may be 5, 10, 20, or 100. As was discussed in Section 2, the usable history of demand observations is typically very small in practice, which makes 5 or 10 observations realistic choices.

Observe that, as an optimal result, the fill rate in a model with holding and shortage costs is equal to $\frac{b}{b+h}$. It follows that the base case set-up corresponds to a fill rate of 95%, and the set-ups with $b = 50$ and $b = 100$ correspond to fill rates of 98% and 99%, respectively, which are typical settings in practice. Table 1 shows the resulting order-up-to levels of the corrected approaches and the classical approach, as well as the percentage differences between the average safety stocks and costs of each. The safety stock is calculated by subtracting the expected lead time demand from the order-up-to level.

It is clear from Table 1 that the calculated order-up-to level that results from the classical method (which ignores parameter uncertainty and is therefore insensitive to

variations in n) is substantially too low, especially when n is small, resulting in a significant cost increase. Based on five observations in the base case, the cost benefit of treating the estimation uncertainty in μ and σ^2 correctly is 21%, with a difference in safety stock level of 84%. As n increases, both the safety stock mark-up and the cost benefit decrease, as extra information results in more accurate estimates, meaning that the effects of parameter uncertainty are less severe. With 10 observations, the cost difference is still 9%, with 20 observations it is reduced to 3%, and with 100 observations the difference is only 0.2%.

The shortage costs have a substantial impact on the actual cost benefit. As the corrected order-up-to levels are larger than the classical order-up-to levels, it is expected and confirmed by the results that the costlier backorders are, the larger the cost benefit of applying the correction is. If $b = 100$, then even with $n = 20$ historical observations, the exact approach leads to a cost advantage of 9% relative to the classical approach, which increases to 49% for $n = 5$. The lead time also has a big effect: if the lead time is doubled to 10, the cost benefits increase, especially in the cases where the estimates are not based on only a few observations. This can be explained by the fact that the autocorrelative effect of the estimation errors is magnified if the lead time increases. For example, if μ is underestimated, then the demand in all time instants during the lead time is underestimated.

If the demand variance is decreased, then this decreases the fluctuation in lead time demand, and hence the safety stock mark-up. However, the percentage safety stock difference and cost benefit of the exact approach remain the same. Thus, the demand variance can be seen as a scaling parameter for the uncertainty in demand. A similar observation holds for the mean.

Finally, we compare the approximate approach to the exact approach. As one would expect from asymptotic theory, the exact approach outperforms the approximate approach, but the difference decreases as n increases. For $n = 10$ and larger, the exact and asymptotic approaches typically yield very similar results. The shape of the distribution that the asymptotic approach assumes for the demand variance error term only holds asymptotically, and thus, two effects come into play that dictate the benefit of the asymptotic approach relative to the classical approach. Firstly, the smaller n is, the larger the effect of the estimation uncertainty, and hence, the larger the possible gain. However, at the same time, the smaller n is, the more severe is the misfit of the shape of the asymptotic normal distribution for the demand variance error term, which has a negative effect on the gain. In conclusion, it is worthwhile to derive the exact error term distribution (if possible) if very few data points are available, but the approximate method (which is easy to derive and always available) suffices for 10 or more data points.

We now inspect the cost differences between the different approaches in more detail for the base case scenario. We do this by means of the violin plots in Fig. 1, which show a combination of a box plot and a kernel density plot. The exact approach typically sets the largest order-up-to level,

Fig. 1. Violin plots of the cost differences between the classical approach and the improved approaches, for $\hat{\mu} = 10$, $\hat{\sigma}^2 = 4$, $b = 20$, $h = 1$, $L = 5$, and $n = 10$.

approach. This implies that if μ and σ are underestimated, then the exact approach yields the largest benefits, followed by the asymptotic approach, as backorder costs are much larger than holding costs. This is reflected in the tails of the plots. Observe that there are also cases where the exact and asymptotic approaches lead to slightly larger costs than the classical approach. This typically occurs when μ and/or σ are severely overestimated. In such situations, the safety stocks of the classical approach were already too high, and adding an extra mark-up leads to larger holding costs. Overall, though, the cost consequences of underestimating the demand parameters are more severe, as is confirmed by the fact that the corrected approaches lead to substantial cost benefits on average relative to the classical approach. The average cost of the classical approach is 15.0, whereas that of the approximate approach is 13.7 and that of the exact approach is 13.6. Thus, setting an extra safety-stock mark-up incurs slightly larger costs in situations where the demand parameters are overestimated; however, at the same time the decision-maker is protected against very large costs when the demand parameters are underestimated, which has a larger effect on average. These effects behave similarly for the demand models that we discuss next.

4.2. Single exponential smoothing for mean-stationary, normally distributed demand

Table 2 shows the results of the numerical comparisons for the mean-stationary normal demand model, where SES

Table 2

Numerical results: single exponential smoothing for mean-stationary, normally distributed demand.

Parameters							Order-up-to levels			Safety stock diff.		Total expected holding and shortage costs				
$\hat{\mu}$	$\hat{\sigma}^2$	α	n	b	h	L	Class.	Approx.	Exact	Approx.	Exact	Class.	Approximate		Exact	
10	4	0.2	5	20	1	5	57.5	59.7	64.8	+30.4%	+97.8%	29.3	25.0	(−14.7%)	22.1	(−24.7%)
10	4	0.2	10	20	1	5	57.5	59.4	60.7	+26.2%	+42.9%	16.1	14.5	(−10.3%)	14.2	(−11.8%)
10	4	0.2	20	20	1	5	57.5	59.3	59.8	+24.9%	+31.3%	13.7	12.6	(−7.9%)	12.6	(−8.2%)
10	4	0.2	100	20	1	5	57.5	59.3	59.4	+24.7%	+25.9%	12.6	11.8	(−6.5%)	11.8	(−6.5%)
10	4	0.2	5	100	1	5	60.4	64.6	75.5	+39.6%	+144.9%	76.6	50.9	(−33.6%)	35.5	(−53.6%)
10	4	0.2	10	100	1	5	60.4	63.8	66.2	+32.5%	+55.3%	28.6	20.8	(−27.3%)	19.8	(−30.8%)
10	4	0.2	20	100	1	5	60.4	63.4	64.2	+28.7%	+36.3%	21.1	16.9	(−20.1%)	16.7	(−20.7%)
10	4	0.2	100	100	1	5	60.4	63.1	63.2	+25.5%	+26.8%	17.8	15.2	(−14.8%)	15.2	(−14.8%)
10	4	0.2	5	20	1	10	110.6	116.0	126.1	+52.0%	+147.5%	59.7	46.0	(−23.0%)	39.1	(−34.6%)
10	4	0.2	10	20	1	10	110.6	115.5	117.8	+47.0%	+68.4%	30.0	24.1	(−19.7%)	23.6	(−21.3%)
10	4	0.2	20	20	1	10	110.6	115.4	116.1	+45.5%	+53.0%	25.0	20.8	(−16.9%)	20.7	(−17.2%)
10	4	0.2	100	20	1	10	110.6	115.3	115.5	+45.3%	+46.7%	23.0	19.4	(−15.7%)	19.4	(−15.7%)
10	1	0.2	5	20	1	5	53.7	54.9	57.4	+30.4%	+97.8%	14.6	12.5	(−14.7%)	11.0	(−24.6%)
10	1	0.2	10	20	1	5	53.7	54.7	55.3	+26.2%	+42.9%	8.0	7.2	(−10.3%)	7.1	(−11.7%)
10	1	0.2	20	20	1	5	53.7	54.7	54.9	+24.9%	+31.3%	6.8	6.3	(−7.9%)	6.3	(−8.2%)
10	1	0.2	100	20	1	5	53.7	54.7	54.7	+24.7%	+25.9%	6.3	5.9	(−6.5%)	5.9	(−6.6%)
20	4	0.2	5	20	1	5	107.5	109.7	114.8	+30.4%	+97.8%	29.2	24.9	(−14.7%)	22.1	(−24.6%)
20	4	0.2	10	20	1	5	107.5	109.4	110.7	+26.2%	+42.9%	16.1	14.5	(−10.3%)	14.2	(−11.8%)
20	4	0.2	20	20	1	5	107.5	109.3	109.8	+24.9%	+31.3%	13.7	12.6	(−7.9%)	12.6	(−8.2%)
20	4	0.2	100	20	1	5	107.5	109.3	109.4	+24.7%	+25.9%	12.6	11.8	(−6.5%)	11.8	(−6.5%)
10	4	0.5	5	20	1	5	57.5	62.7	65.9	+70.8%	+113.5%	33.1	24.7	(−25.4%)	23.8	(−28.2%)
10	4	0.5	10	20	1	5	57.5	62.3	63.6	+65.2%	+82.5%	24.5	18.4	(−25.1%)	18.2	(−25.9%)
10	4	0.5	20	20	1	5	57.5	62.2	62.8	+63.5%	+71.9%	21.8	16.5	(−24.2%)	16.5	(−24.5%)
10	4	0.5	100	20	1	5	57.5	62.2	62.3	+63.3%	+64.9%	20.1	15.4	(−23.1%)	15.4	(−23.1%)
10	4	0.5	5	100	1	5	60.4	69.0	77.5	+82.8%	+164.2%	90.3	45.0	(−50.2%)	38.2	(−57.6%)
10	4	0.5	10	100	1	5	60.4	68.1	70.7	+73.4%	+98.3%	54.3	26.2	(−51.7%)	25.3	(−53.4%)
10	4	0.5	20	100	1	5	60.4	67.6	68.6	+68.5%	+78.4%	43.6	22.1	(−49.4%)	21.9	(−49.8%)
10	4	0.5	100	100	1	5	60.4	67.1	67.3	+64.4%	+66.1%	37.2	19.9	(−46.6%)	19.9	(−46.6%)
10	4	0.8	5	20	1	5	57.5	66.2	70.3	+117.7%	+171.5%	49.1	31.6	(−35.7%)	30.4	(−38.0%)
10	4	0.8	10	20	1	5	57.5	65.7	67.4	+110.6%	+132.6%	37.5	23.4	(−37.6%)	23.1	(−38.2%)
10	4	0.8	20	20	1	5	57.5	65.6	66.3	+108.5%	+119.1%	33.8	21.1	(−37.8%)	21.0	(−38.0%)
10	4	0.8	100	20	1	5	57.5	65.5	65.7	+108.1%	+110.2%	31.4	19.7	(−37.3%)	19.7	(−37.3%)
10	4	0.8	5	100	1	5	60.4	74.3	85.0	+133.0%	+236.0%	151.7	57.2	(−62.3%)	48.7	(−67.9%)
10	4	0.8	10	100	1	5	60.4	73.0	76.3	+121.1%	+152.7%	100.6	33.5	(−66.7%)	32.3	(−67.9%)
10	4	0.8	20	100	1	5	60.4	72.4	73.7	+114.7%	+127.4%	84.7	28.1	(−66.8%)	27.9	(−67.1%)
10	4	0.8	100	100	1	5	60.4	71.8	72.1	+109.5%	+111.7%	74.8	25.4	(−66.1%)	25.4	(−66.1%)

is used to estimate the mean. We focus on the results that provide new insights relative to the previous section. In particular, it is interesting to observe that all costs in Table 2 are higher than those in Table 1. This reflects the fact that the SES estimator has a larger sample variance than the sample mean, which is the minimum variance unbiased estimator for the mean-stationary normal demand model. For given estimators $\hat{\mu}$ and $\hat{\sigma}^2$, the classical order-up-to levels in this section are the same as those in the previous section. However, given these estimates, the true values for μ and σ^2 can vary more under the SES estimator than under the sample mean, and therefore, even when applying the exact correction that leads to the optimal decision under the mean-stationary normal demand model, the costs using the SES estimator are larger than those using the sample mean.

If a larger value of the smoothing parameter α is chosen, then relatively more weight is given to recent observations. The SES estimator gives exponentially decreasing weights to older observations, meaning that increasing α increases

the sample variance of the estimator, and thus also the necessary safety stock mark-up and the cost difference between the corrected approaches and the classical approach. Furthermore, the cost benefit of applying the corrected approaches remains very large (up to 66% even for $n = 100$), especially for large values of α . The discussion in Section 2 suggests that values between 0.21 and 0.4, and possibly even larger, are realistic. As in the previous section, we again find that the approximate approach performs very similarly to the exact approach if there are 10 or more observations. Furthermore, a longer lead time increases the possible cost differences, and the percentage cost differences are invariant to the demand parameterization. In conclusion, although using the SES estimator in the mean-stationary normal demand model leads to a cost disadvantage compared to using the sample mean, the cost benefit that can be realized by correcting for the parameter uncertainty is significantly larger for SES, and is also substantial in proportion to the extra cost associated with using SES instead of the sample average.

Table 3

Numerical results: normally distributed, trended demand.

Parameters							Order-up-to levels			Safety stock diff.		Total expected holding and shortage cost				
$\hat{\alpha}$	$\hat{\beta}$	$\hat{\sigma}^2$	n	b	h	L	Class.	Approx.	Exact	Approx.	Exact	Class.	Approximate		Exact	
10	1	4	5	20	1	5	97.5	119.7	131.0	+298.3%	+449.6%	138.7	71.7	(−48.3%)	67.4	(−51.4%)
10	1	4	10	20	1	5	122.5	131.7	133.7	+124.2%	+151.2%	42.8	25.6	(−40.3%)	25.2	(−41.1%)
10	1	4	20	20	1	5	172.5	176.6	177.3	+55.9%	+64.4%	20.2	15.8	(−21.4%)	15.8	(−21.7%)
10	1	4	100	20	1	5	572.5	573.2	573.3	+10.2%	+11.3%	10.6	10.4	(−1.5%)	10.4	(−1.5%)
10	1	4	5	50	1	5	99.2	128.3	149.8	+315.1%	+548.3%	298.2	109.0	(−63.4%)	93.8	(−68.6%)
10	1	4	10	50	1	5	124.2	136.3	139.4	+130.8%	+164.6%	75.7	31.8	(−58.0%)	30.9	(−59.1%)
10	1	4	20	50	1	5	174.2	179.6	180.5	+58.7%	+68.0%	28.9	18.9	(−34.5%)	18.8	(−35.0%)
10	1	4	100	50	1	5	574.2	575.2	575.3	+10.6%	+11.8%	12.5	12.2	(−2.6%)	12.2	(−2.6%)
10	1	4	5	100	1	5	100.4	134.4	167.7	+326.3%	+645.3%	547.9	156.1	(−71.5%)	119.8	(−78.1%)
10	1	4	10	100	1	5	125.4	139.5	143.8	+135.3%	+176.2%	121.6	37.6	(−69.1%)	35.8	(−70.6%)
10	1	4	20	100	1	5	175.4	181.7	182.8	+60.6%	+71.0%	38.6	21.2	(−45.1%)	21.0	(−45.6%)
10	1	4	100	100	1	5	575.4	576.6	576.7	+11.0%	+12.1%	13.9	13.4	(−3.5%)	13.4	(−3.6%)
10	1	4	5	20	1	10	215.6	289.9	322.2	+705.1%	+1010.9%	472.1	204.8	(−56.6%)	192.2	(−59.3%)
10	1	4	10	20	1	10	265.6	295.1	300.0	+280.3%	+326.1%	139.0	61.3	(−55.9%)	60.5	(−56.5%)
10	1	4	20	20	1	10	365.6	378.4	379.6	+121.3%	+133.3%	53.4	31.7	(−40.5%)	31.6	(−40.7%)
10	1	4	100	20	1	10	1165.6	1167.8	1167.9	+20.9%	+22.2%	17.0	16.2	(−5.0%)	16.2	(−5.0%)
10	1	1	5	20	1	5	93.7	104.9	110.5	+298.3%	+449.6%	69.2	35.7	(−48.4%)	33.5	(−51.5%)
10	1	1	10	20	1	5	118.7	123.4	124.4	+124.2%	+151.2%	21.5	12.8	(−40.2%)	12.6	(−41.0%)
10	1	1	20	20	1	5	168.7	170.8	171.1	+56.0%	+64.4%	10.1	7.9	(−21.3%)	7.9	(−21.6%)
10	1	1	100	20	1	5	568.7	569.1	569.2	+10.2%	+11.3%	5.3	5.2	(−1.5%)	5.2	(−1.5%)
10	0.5	4	5	20	1	5	77.5	99.7	111.0	+298.3%	+449.6%	138.7	71.6	(−48.3%)	67.3	(−51.5%)
10	0.5	4	10	20	1	5	90.0	99.2	101.2	+124.2%	+151.2%	42.8	25.6	(−40.1%)	25.3	(−40.9%)
10	0.5	4	20	20	1	5	115.0	119.1	119.8	+55.9%	+64.4%	20.1	15.8	(−21.2%)	15.8	(−21.5%)
10	0.5	4	100	20	1	5	315.0	315.7	315.8	+10.2%	+11.3%	10.6	10.4	(−1.5%)	10.4	(−1.5%)
20	1	4	5	20	1	5	147.5	169.7	181.0	+298.3%	+449.6%	139.0	71.9	(−48.3%)	67.4	(−51.5%)
20	1	4	10	20	1	5	172.5	181.7	183.7	+124.2%	+151.2%	42.7	25.5	(−40.2%)	25.2	(−41.0%)
20	1	4	20	20	1	5	222.5	226.6	227.3	+55.9%	+64.4%	20.1	15.8	(−21.4%)	15.8	(−21.6%)
20	1	4	100	20	1	5	622.5	623.2	623.3	+10.2%	+11.3%	10.6	10.4	(−1.5%)	10.4	(−1.5%)

4.3. Normally distributed, trended demand

Table 3 shows the comparative results for the trended demand model. As many effects have already been already interpreted for the previous two demand models, we focus again on the insights that are specific to the trended demand model. If we compare Tables 1 and 3, we can see that both the percentage safety stock mark-ups and the percentage cost savings are significantly larger for the trended demand than for the mean-stationary demand. This reflects the autocorrelative effect of future forecast errors. In the trend model, it is not only the forecast errors of the mean parameter that accumulate over the lead time, but also those of the trend parameter. If a trend is underestimated, this has a substantially larger effect on the error in the lead time demand forecast than if only a mean is underestimated. For example, if the trend parameter is underestimated by 0.1, then the (one-period) demand forecast for five periods later is off by 0.5. On the other hand, if the mean parameter is underestimated by 0.1, then the demand forecast for five periods later is still only off by 0.1. Thus, the consequences of ignoring parameter uncertainty are substantially larger in the trended model than in the mean-stationary model.

In the base case, the safety stocks that result from the classical approach have to be adjusted by 11% for 100 observations, to 450% when only five observations are available. This mark-up increases to more than 1000% (so with a multiplication factor of 11) if the lead time is doubled to

10. For this longer lead time, the necessary mark-up even with 100 observations is still 22%, and the resulting cost saving is 5%, which increases to 59% for $n = 5$. When the backorder costs are set at 100 (and the lead time is five), savings of between 4% for $n = 100$ and 69% for $n = 5$ can be achieved. The percentage cost differences are invariant to the demand parameterization, to both $\hat{\alpha}$ and $\hat{\beta}$, and we again observe that the approximate and exact approaches perform nearly equally well for 10 observations or more.

4.4. Demand is a normal random walk

When the demand follows a random walk with normally distributed increments and the model is specified correctly, there is only one parameter left to estimate: the variance of the increments. Thus, the safety stock mark-ups and cost savings shown in Table 4 are not as dramatic as those in previous sections. However, if the backorder cost is set to 100, savings of 9% for $n = 10$ or 31% for $n = 5$ are still possible. Typically, though, the possible savings become very small if there are more than 10 observations available, and they are negligible in many settings for $n > 100$. Interestingly, for this model a longer lead time does not lead to larger percentage cost savings of the corrected approaches relative to the classical approach. This is easy to explain, as the underestimation of the variance of the increments does not have a multiplicative effect when future demands are aggregated, other than with a mean or trend parameter, while the proportional error remains constant over the

Table 4

Numerical results: normal random walk.

Parameters						Order-up-to levels			Safety stock diff.		Total expected holding and shortage costs				
D_n	$\hat{\sigma}^2$	n	b	h	L	Class.	Approx.	Exact	Approx.	Exact	Class.	Approximate	Exact		
10	4	5	20	1	5	74.7	76.3	85.7	+6.4%	+44.3%	63.3	61.8	(−2.3%)	58.4	(−7.7%)
10	4	10	20	1	5	74.7	75.1	78.1	+1.5%	+13.4%	38.4	38.2	(−0.4%)	37.8	(−1.6%)
10	4	20	20	1	5	74.7	74.8	76.1	+0.2%	+5.6%	33.7	33.7	(−0.0%)	33.6	(−0.4%)
10	4	100	20	1	5	74.7	74.7	75.0	+0.0%	+1.0%	31.4	31.4	(−0.0%)	31.4	(−0.0%)
10	4	5	50	1	5	80.6	84.0	102.1	+11.3%	+70.2%	100.8	93.9	(−6.9%)	81.7	(−19.0%)
10	4	10	50	1	5	80.6	82.0	86.5	+4.7%	+19.4%	48.7	47.7	(−2.2%)	46.4	(−4.7%)
10	4	20	50	1	5	80.6	81.2	83.0	+2.0%	+7.9%	40.5	40.3	(−0.5%)	40.0	(−1.1%)
10	4	100	50	1	5	80.6	80.7	81.0	+0.3%	+1.4%	36.7	36.7	(−0.0%)	36.7	(−0.0%)
10	4	5	100	1	5	84.6	89.6	117.6	+14.5%	+95.6%	150.8	133.5	(−11.5%)	104.8	(−30.5%)
10	4	10	100	1	5	84.6	86.9	93.1	+6.9%	+24.6%	58.5	55.8	(−4.7%)	53.3	(−8.8%)
10	4	20	100	1	5	84.6	85.7	87.9	+3.3%	+9.8%	45.7	45.1	(−1.2%)	44.8	(−2.0%)
10	4	100	100	1	5	84.6	84.8	85.1	+0.6%	+1.7%	40.4	40.4	(−0.0%)	40.4	(−0.1%)
10	4	5	20	1	10	165.5	169.7	194.5	+6.4%	+44.3%	168.2	164.3	(−2.3%)	155.2	(−7.7%)
10	4	10	20	1	10	165.5	166.5	174.2	+1.5%	+13.4%	101.5	101.2	(−0.4%)	99.9	(−1.6%)
10	4	20	20	1	10	165.5	165.6	169.1	+0.2%	+5.6%	89.3	89.2	(−0.0%)	88.9	(−0.4%)
10	4	100	20	1	10	165.5	165.4	166.1	+0.0%	+1.0%	83.0	83.0	(−0.0%)	83.0	(−0.0%)
10	1	5	20	1	5	62.4	63.2	67.9	+6.4%	+44.3%	31.8	31.0	(−2.3%)	29.3	(−7.7%)
10	1	10	20	1	5	62.4	62.6	64.0	+1.5%	+13.4%	19.2	19.1	(−0.4%)	18.9	(−1.6%)
10	1	20	20	1	5	62.4	62.4	63.1	+0.2%	+5.6%	16.9	16.9	(−0.0%)	16.8	(−0.4%)
10	1	100	20	1	5	62.4	62.4	62.5	+0.0%	+1.0%	15.7	15.7	(−0.0%)	15.7	(−0.0%)
20	4	5	20	1	5	124.7	126.3	135.7	+6.4%	+44.3%	63.6	62.1	(−2.3%)	58.6	(−7.7%)
20	4	10	20	1	5	124.7	125.1	128.1	+1.5%	+13.4%	38.4	38.2	(−0.4%)	37.7	(−1.6%)
20	4	20	20	1	5	124.7	124.8	126.1	+0.2%	+5.6%	33.7	33.7	(−0.0%)	33.6	(−0.4%)
20	4	100	20	1	5	124.7	124.7	125.0	+0.0%	+1.0%	31.4	31.4	(−0.0%)	31.4	(−0.0%)

lead time. Furthermore, the random walk model has no estimation error around the level of future demands. We conclude that the cost consequences of ignoring parameter uncertainty are less severe for the random walk demand model with normal increments than for the other demand models.

5. Demand misspecification

The numerical results in the previous section hold under the assumption that the demand model was specified correctly. However, in practice it may be difficult to specify whether demand is stationary, has a trend, or is autocorrelated, and if so, for how long such has been the case. We have already considered using the SES estimator instead of the sample mean in the mean-stationary demand model in Section 4.2. The underlying demand model – that of mean-stationary, normally distributed demand – was specified correctly and used in the analysis. In this section, on the other hand, we study a scenario where the demand model is not specified correctly and the analysis is based on that incorrect specification, examining the negative cost consequences and the performances of the classical and corrected approaches in the presence of such a misspecification. Specifically, we consider the scenario where a trend is fitted to a data series that is actually mean-stationary. We focus on the base case scenarios of the previous section and compare the order-up-to levels and costs of the classical and corrected approaches, in order to analyze how the cost consequences of demand misspecification compare to the cost consequences of ignoring parameter uncertainty.

Assume that the decision maker suspects that the demand is normally distributed and stationary around a

linear trend, and acts accordingly, whereas actually the demand is mean-stationary. Table 5 shows the numerical results of this scenario. We consider the case where the point estimate for the trend parameter is zero, which is the most likely scenario in the absence of a trend. Note that the uncertainty around the trend parameter is still taken into account by the proposed methods, so the misspecification affects the results even if the point estimate for the trend parameter is zero. The costs are uniformly higher than those in Table 1. For five observations, the classical approach overall results in costs that are approximately three times as large as those without misspecification. The exact approach reduces this to approximately a factor of two for cases with a lead time of five time instants.

It is noteworthy that the asymptotic approach outperforms the exact approach for this case of demand misspecification. The difference is especially large if few observations are used, and the two approaches yield similar results for $n = 10$ or more. Thus, the asymptotic approach proves to be valuable in this scenario, suggesting that the asymptotic approach can yield better results than the exact approach if the decision maker is unsure of the demand specification.

Comparing Tables 1 and 5, we find that, for the base case with $n = 5$, the total cost (in absolute units) without misspecification is 26 for the classical approach, 21 for the approximate approach, and 21 for the exact approach. Under demand misspecification, these figures are 67, 38, and 45, respectively. This implies that the cost effect of misspecifying the demand is 41 for the classical approach, 17 for the approximate approach, and 23 for the exact approach. The gain from applying the approximate correction factor compared to the classical approach with misspecification is

Table 5

Numerical results: misspecification, trend fitted to the mean-stationary demand.

Parameters								Order-up-to levels			Safety stock diff.		Total expected holding and shortage costs				
$\hat{\alpha}$	$\hat{\beta}$	$\hat{\sigma}^2$	n	b	h	L		Class.	Approx.	Exact	Approx.	Exact	Class.	Approximate		Exact	
10	0	4	5	20	1	5		57.5	79.7	91.0	+38.7%	+58.4%	67.1	38.0	(−43.4%)	44.8	(−33.2%)
10	0	4	10	20	1	5		57.5	66.7	68.7	+16.1%	+19.6%	29.8	20.5	(−31.2%)	21.3	(−28.6%)
10	0	4	20	20	1	5		57.5	61.6	62.3	+7.3%	+8.4%	17.4	14.5	(−16.6%)	14.6	(−16.0%)
10	0	4	100	20	1	5		57.5	58.2	58.3	+1.3%	+1.5%	10.5	10.4	(−1.4%)	10.4	(−1.4%)
10	0	4	5	50	1	5		59.2	88.3	109.8	+49.1%	+85.4%	131.5	49.4	(−62.4%)	63.2	(−51.9%)
10	0	4	10	50	1	5		59.2	71.3	74.4	+20.4%	+25.6%	48.6	25.0	(−48.5%)	26.5	(−45.6%)
10	0	4	20	50	1	5		59.2	64.6	65.5	+9.1%	+10.6%	24.0	17.3	(−27.7%)	17.5	(−26.9%)
10	0	4	100	50	1	5		59.2	60.2	60.3	+1.7%	+1.8%	12.4	12.1	(−2.3%)	12.1	(−2.3%)
10	0	4	5	100	1	5		60.4	94.4	127.7	+56.3%	+111.3%	226.1	59.5	(−73.7%)	80.8	(−64.3%)
10	0	4	10	100	1	5		60.4	74.5	78.8	+23.3%	+30.4%	72.6	28.6	(−60.6%)	30.6	(−57.8%)
10	0	4	20	100	1	5		60.4	66.7	67.8	+10.5%	+12.2%	30.8	19.3	(−37.2%)	19.6	(−36.4%)
10	0	4	100	100	1	5		60.4	61.6	61.7	+1.9%	+2.1%	13.8	13.4	(−3.3%)	13.4	(−3.3%)
10	0	4	5	20	1	10		110.6	184.9	217.2	+67.3%	+96.5%	149.3	91.2	(−38.9%)	119.9	(−19.7%)
10	0	4	10	20	1	10		110.6	140.1	145.0	+26.7%	+31.1%	66.7	42.7	(−35.9%)	46.5	(−30.2%)
10	0	4	20	20	1	10		110.6	123.4	124.6	+11.6%	+12.7%	36.0	26.2	(−27.0%)	26.9	(−25.2%)
10	0	4	100	20	1	10		110.6	112.8	112.9	+2.0%	+2.1%	16.5	15.9	(−4.1%)	15.9	(−4.1%)
10	0	1	5	20	1	5		53.7	64.9	70.5	+20.7%	+31.2%	33.6	19.0	(−43.3%)	22.4	(−33.2%)
10	0	1	10	20	1	5		53.7	58.4	59.4	+8.6%	+10.5%	14.9	10.2	(−31.3%)	10.6	(−28.6%)
10	0	1	20	20	1	5		53.7	55.8	56.1	+3.9%	+4.5%	8.7	7.3	(−16.7%)	7.3	(−16.1%)
10	0	1	100	20	1	5		53.7	54.1	54.2	+0.7%	+0.8%	5.3	5.2	(−1.4%)	5.2	(−1.4%)
20	0	4	5	20	1	5		107.5	129.7	141.0	+20.7%	+31.2%	67.1	38.1	(−43.3%)	44.9	(−33.1%)
20	0	4	10	20	1	5		107.5	116.7	118.7	+8.6%	+10.5%	29.9	20.5	(−31.4%)	21.3	(−28.8%)
20	0	4	20	20	1	5		107.5	111.6	112.3	+3.9%	+4.5%	17.5	14.6	(−16.6%)	14.7	(−16.0%)
20	0	4	100	20	1	5		107.5	108.2	108.3	+0.7%	+0.8%	10.5	10.4	(−1.4%)	10.4	(−1.4%)

29. For the exact correction factor, the gain is 22. If the demand model is misspecified, then the approximate method (which performs best in that scenario) does not yield lower costs than the classical method under a correctly specified demand model. Therefore, we conclude that the cost effect of misspecifying the demand model is larger than the cost reduction that can be achieved by applying the correction. However, the figures above show that the latter is still substantial, as the gain of 29 units corrects 71% of the loss due to demand misspecification. Thus, it is still worthwhile to apply a corrected method (especially the approximate method) when the demand is misspecified, though the largest gain can be achieved by specifying the demand model correctly. As another example, in the base case with $n = 20$, the cost of the classical method under demand misspecification is six units higher than that under a correct specification of the demand model, but applying the corrected methods under misspecification leads to a gain of three units relative to using the classical method. Thus, we find that 50% of the loss due to demand misspecification can be corrected for by applying the corrected methods.

6. Conclusion

An adaptation of inventory models in order to incorporate parameter estimation uncertainty for future demand has been proposed. The approach involves four steps:

1. Formulate the decision model in terms of the lead time demand distribution function.
2. Estimate the parameters and the distributions of their estimation errors efficiently.
3. Replace the true parameters with an appropriate function of the point estimates and estimation errors in the lead time demand distribution function.
4. Use the expectation of the lead time demand distribution function with respect to the estimation errors as the new predictive lead time demand distribution in the decision model.

We have demonstrated the approach for a model with holding and shortage costs, using various demand models. The method can be applied to any inventory model using any demand process and parameter estimator, as long as the distribution of its estimation error can be derived or approximated. The typical results of this approach are safety stock mark-ups that increase with the degree of uncertainty in the model. The degree of uncertainty increases if more parameters have to be estimated, if the estimates are based on small numbers of observations, if the lead time is large, or if the variance of the demand is large with respect to the mean. The method allows for freedom in the choice of the approximation method of the estimators' error distributions. We discussed the exact error distribution and the asymptotic approach, which is typically easier to derive and robust to demand model misspecification.

The cost benefits of dealing with estimation uncertainty properly in inventory models are twofold: the expected costs are lower and the cost variance is reduced. The point estimate of a parameter in the model is sometimes close to the true parameter, leading to lower costs, but it can also be relatively far away from the true parameter, leading to very high costs if no correction is applied. By incorporating a distribution of the estimation error, one acknowledges that the actual parameter value may be either smaller or larger than the point estimate. A safety stock mark-up is established, which reduces the chance of extraordinarily high future costs, since shortage costs tend to be higher than holding costs. This two-sided benefit is even more beneficial for risk-averse decision makers. The asymptotic approach that is generally applicable performed almost as well as the exact approach when the estimates were based on 10 observations or more. This implies that the correction framework discussed in this paper can also be applied to demand distributions for which the exact error distribution cannot be derived.

One particular case where very large cost savings can be achieved is the trend model, because mis-estimating a trend and aggregating the lead time demands has a much larger effect than when only a mean is mis-estimated, for example. On the other hand, the random walk model showed relatively low possible cost savings. In that model, parameter uncertainty (in only the variance of the increments) plays only a minor role, as long as the model is specified correctly and the variance of the lead time demand is derived correctly.

We also studied the severity of ignoring parameter uncertainty relative to choosing a suboptimal estimator or mis-specifying the demand model. Although the latter has a larger effect on costs, the savings that can be achieved by applying the correction are still substantial. Also, the approximate method showed favorable results in this case due to its robustness. In conclusion, we find that the current standard of separating forecasting and decision making has severe cost consequences, which could be prevented easily by applying the simple derived four-step method to any inventory model in which forecasts are used.

In practice, the message that we aim to transmit to inventory decision makers is that inventories are held as a protection against the uncertainty around future demand, but that a major part of this uncertainty is ignored if parameter estimates are treated as the true parameters of the demand distribution. This uncertainty comprises a large part of the total uncertainty around future demand, especially since, in practice, estimates are typically based on only a few observations, and therefore have large errors. Ignoring this leads to safety stocks that are too low, and therefore to frequent stock-outs. This in turn implies that target service levels are not achieved, that customers are dissatisfied, and that high backorder costs are incurred. This paper proposes a framework that can correct this flaw for general combinations of demand models and parameter estimators.

There are ample opportunities for further research, along three main strands. The first strand involves studying the four-step method in different inventory models,

under different demand distributions, and for different processes. The second strand involves studying different approaches to modeling the estimation error distribution, such as bootstrapping. The final strand involves searching for efficient parameterizations and model formulations so that the four-step approach can be implemented and the optimal solution calculated efficiently even for models with many parameters. A related issue is that of the choice of estimator. The parameter estimators that are used generally are based on loss functions that may be inappropriate for some inventory models. For example, least squares considers an underestimation to be just as damaging as an overestimation, whereas if $b > h$, an underestimation of a mean, for example, has a larger cost effect than an overestimation. Thus, an interesting future research topic would be to make a connection between the actual loss in the inventory model and the loss function of the estimator.

Acknowledgments

Funding: This work was supported by The Netherlands Organisation for Scientific Research (NWO) [grant number 406-15-014].

References

- Ali, M. M., & Boylan, J. E. (2011). Feasibility principles for downstream demand inference in supply chains. *The Journal of the Operational Research Society*, 62(3), 474–482.
- Ali, M. M., & Boylan, J. E. (2012). On the effect of non-optimal forecasting methods on supply chain downstream demand inference. *IMA Journal of Management Mathematics*, 23(1), 81–98.
- Axsäter, S. (2006). *Inventory control*. (2nd ed.). New York: Springer-Verlag.
- Azoury, K. S., & Miyaoka, J. (2009). Optimal policies and approximations for a Bayesian linear regression inventory model. *Management Science*, 55(5), 813–826.
- Bermúdez, J. D., Segura, J. V., & Vercher, E. (2010). Bayesian forecasting with the Holt-Winters model. *The Journal of the Operational Research Society*, 61(1), 164–171.
- Beutel, A. L., & Minner, S. (2012). Safety stock planning under causal demand forecasting. *International Journal of Production Economics*, 140(2), 637–645.
- Chen, L. (2010). Bounds and heuristics for optimal Bayesian inventory control with unobserved lost sales. *Operations Research*, 58(2), 396–413.
- Davidson, R., & MacKinnon, J. G. (2003). *Econometric theory and methods*. (1st ed.). New York: Oxford University Press.
- Dvoretzky, A., Kiefer, J., & Wolfowitz, J. (1952). The inventory problem: II. Case of unknown distributions of demand. *Econometrica*, 20(3), 450–466.
- Fildes, R. (1983). An evaluation of Bayesian forecasting. *Journal of Forecasting*, 2(1), 137–150.
- Fildes, R., & Beard, C. (1992). Forecasting systems for production and inventory control. *International Journal of Operations and Production Management*, 12(5), 4–27.
- Gardner, E. S. (2006). Exponential smoothing: the state of the art — Part II. *International Journal of Forecasting*, 22(4), 637–666.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (1995). *Bayesian data analysis*. (2nd ed.). London: Chapman and Hall.
- Greene, W. H. (2011). *Econometric analysis*. (7th ed.). London: Pearson.
- Harris, F. W. (1913). How many parts to make at once. *Factory, the Magazine of Management*, 10(2), 135–136.
- Harrison, P. J., & Stevens, C. F. (1971). A Bayesian approach to short-term forecasting. *The Journal of the Operational Research Society*, 22(4), 341–362.
- Hax, A. C., & Candea, D. (1984). *Production and inventory management*. (1st ed.). Englewood Cliffs: Prentice-Hall.

- Hill, R. M. (1997). Applying Bayesian methodology with a uniform prior to the single period inventory model. *European Journal of Operational Research*, 98(3), 555–562.
- Hill, R. M. (1999). Bayesian decision-making in inventory modelling. *IMA Journal of Mathematics Applied in Business and Industry*, 10(2), 165–176.
- Hillier, F. S., & Lieberman, G. J. (2014). *Introduction to operations research*. (10th ed.). Boston: McGraw Hill.
- Janssen, E., Strijbosch, L., & Brekelmans, R. (2009). Assessing the effects of using demand parameters estimates in inventory control and improving the performance using a correction function. *International Journal of Production Economics*, 118(1), 34–42.
- Miller, I., & Miller, M. (2004). *John E. Freund's mathematical statistics with applications*. (7th ed.). London: Pearson.
- Ord, J. K., Koehler, A. B., & Snyder, R. D. (1997). Estimation and prediction for a class of dynamic nonlinear statistical models. *Journal of the American Statistical Association*, 92(440), 1621–1629.
- Prak, D. R. J., Syntetos, A. A., & Teunter, R. H. (2017). On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research*, 256(2), 454–461.
- Rajashree Kamath, K., & Pakkala, T. P. M. (2002). A Bayesian approach to a dynamic inventory model under a unknown demand distribution. *Computers & Operations Research*, 29(4), 317–422.
- Ritchken, P. H., & Sankar, R. (1984). The effect of estimation risk in establishing safety stock levels in an inventory model. *The Journal of the Operational Research Society*, 35(12), 1091–1099.
- Scarf, H. (1959). Bayes solutions of the statistical inventory problem. *Annals of Mathematical Statistics*, 30(2), 490–508.
- Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling*. (3rd ed.). New York: John Wiley & Sons.
- Spedding, T. A., & Chan, K. K. (2000). Forecasting demand and inventory management using Bayesian time series. *Integrated Manufacturing Systems*, 11(5), 331–339.
- Strijbosch, L. W. G., Heuts, R. M. J., & van der Schoot, E. H. M. (2000). A combined forecast-inventory control procedure for spare parts. *The Journal of the Operational Research Society*, 51(10), 1184–1192.
- Strijbosch, L. W. G., & Moors, J. J. A. (2005). The impact of unknown demand parameters on (R, S)-inventory control performance. *European Journal of Operational Research*, 162(3), 805–815.
- Strijbosch, L. W., Syntetos, A. A., Boylan, J. E., & Janssen, E. (2011). On the interaction between forecasting and stock control: The case of non-stationary demand. *International Journal of Production Economics*, 133(1), 470–480.
- Toktay, L. B., & Wein, L. M. (2001). Analysis of a forecasting-production-inventory system with stationary demand. *Management Science*, 47(9), 1268–1281.
- Waters, C. D. J. (2012). *Inventory control and management*. (2nd ed.). Hoboken: John Wiley & Sons.
- West, M., & Harrison, J. (1997). *Bayesian forecasting and dynamic models*. (2nd ed.). New York: Springer.
- West, M., Harrison, P. J., & Migon, H. S. (1985). Dynamic generalized linear models and Bayesian forecasting. *Journal of the American Statistical Association*, 80(389), 73–83.
- Yelland, P. M. (2010). Bayesian forecasting of parts demand. *International Journal of Forecasting*, 26(2), 374–396.
- Zipkin, P. H. (2000). *Foundations of inventory management*. (1st ed.). Boston: McGraw-Hill.

Dennis Prak is a Ph.D. Candidate at the Department of Operations, University of Groningen, The Netherlands. His main research interest is in the interface of demand forecasting and inventory control, particularly in how forecasting uncertainty should result in inventory level mark-ups. More generally, he seeks solutions to combine demand forecasting and inventory decision making into hybrid models, whereas both fields are typically separated. His articles have appeared in e.g. the *European Journal of Operational Research*.

Ruud Teunter holds a chair of operational research and is Head of Department Operations at the University of Groningen. His research interests include (intermittent demand) forecasting, inventory control, spare parts management, maintenance and reverse logistics. He has published over 70 papers in international OR and OM journals such as the *European Journal of Operational Research*, *Production and Operations Management* and *Journal of Operations Management*. He is currently editor of the *European Journal of Operational Research*.