ARTICLE IN PRESS

International Journal of Forecasting ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

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.

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).

https://doi.org/10.1016/j.ijforecast.2017.11.004

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 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. Strijbosch, 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