European Journal of Operational Research 234 (2014) 163–173

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

A semi-parametric approach for estimating critical fractiles under autocorrelated demand

Yun Shin Lee *

KAIST College of Business, Korea Advanced Institute of Science and Technology, 85 Hoegi-Ro, Dongdaemoon-Gu, Seoul 130-722, Republic of Korea

ARTICLE INFO

Article history:
Received 27 March 2013
Accepted 21 October 2013
Available online 1 November 2013

Keywords:
Forecasting
Newsvendor model
Autocorrelated demand
Model misspecification
Forecast bias
Retail operations

ABSTRACT

Forecasting critical fractiles of the lead time demand distribution is an important problem for operations managers making newsvendor-type inventory decisions. In this paper, we propose a semi-parametric approach to forecasting the critical fractile when demand is serially correlated. Starting from a user-defined but potentially misspecified forecasting model, we use historical demand data to generate empirical forecast errors of this model. These errors are then used to (1) parametrically correct for any bias in the point forecast conditional on the recent demand history and (2) non-parametrically estimate the critical fractile of the demand distribution without imposing distributional assumptions. We present conditions under which this semi-parametric approach provides a consistent estimate of the critical fractile and evaluate its finite sample properties using simulation and real data for retail inventory planning.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Forecasting critical fractiles of the lead time demand (LTD) distribution is an important problem for operations managers making newsvendor-type inventory decisions. It corresponds to finding a cost-minimizing order quantity that optimally balances the costs of understocking and the costs of overstocking in a single period. The critical fractile solution has found numerous applications in inventory theory (see review by Khouja (1999)), particularly in the management of highly perishable or seasonable goods. It can also be repeatedly applied when one takes a multi-period inventory model with a myopic order up-to-policy (e.g., Graves, 1999; Lee, So, & Tang, 2000), which is noted as an appropriate policy for major retail chains (Smith & Agrawal, 2000). In this paper, we consider the problem of forecasting critical fractiles when demand is serially correlated.

The standard parametric approach to this problem starts with a demand forecasting model and estimates the critical fractile as if the chosen model were correctly specified and estimated (Silver, Pyke, & Peterson, 1998). Specifically, this approach assumes that (1) the fitted forecasting model produces unbiased LTD forecasts and (2) the forecast error follows a Gaussian distribution. However, in most practical situations, it is not possible to know the correct demand generating process and satisfy these assumptions (Chatfield, 1995). There are several reasons why this is the case. First, managers will often have to produce simultaneous forecasts for a very large number of items. For example, a typical US grocery store carries anywhere between 15,000 and 60,000 SKUs (FMI, 2012). It is impractical to expect that one can specify the correct model for each individual SKU. Instead, simple forecasting models, such as exponential smoothing, are used in real-life scenarios, without testing for their validity (Taylor, 2007). Second, the choice of the demand forecasting model used for inventory management might not be under the control of the operations/inventory planning department. Instead, it is often dictated by the sales and marketing group, which may well provide biased (usually overoptimistic) forecasts due to incentive misalignment (Goodwin, 1996; Oliva & Watson, 2009). For instance, Manary and Willems (2008) report that at Intel, the sales and marketing group, which is responsible for feeding demand forecasts directly into the manufacturing resource planning system, consistently overestimates demand across virtually all SKUs. Third, forecast errors may not follow a Gaussian distribution (Bookbinder & Lordahl, 1989; Fricker & Goodhart, 2000).

Under model uncertainty, the parametric approach inevitably produces biased fractile estimates. It is important to understand and reduce the bias as the biased fractile estimate leads to increased inventory costs, as well as failure to meet required service levels (Badinelli, 1990; Kim & Ryan, 2003). Moreover, it can amplify the increase in demand variability as one moves up a supply chain, a phenomenon known as the bullwhip effect (Chen, Drezner, Ryan, & Simchi-Levi, 2000).

When we need to be concerned by the possibility that the forecasting model has been misspecified and/or the demand series is non-Gaussian, empirical approaches are widely used in practice

* Tel.: +82 2958 3339; fax: +82 2958 3359.
E-mail address: yunshin@business.kaist.ac.kr

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.10.055

as a useful alternative. The non-parametric approaches (e.g., Bookbinder & Lordahl, 1989; Levi, Roundy, & Shmoys, 2007) are established to negate the need to make any assumptions about the demand model and to use solely the observed samples of demand to forecast critical fractiles. However, they are limited to independent demand processes and cannot be applied to serially correlated demand processes. Alternatively, one can begin with a user-specified forecasting model that captures autocorrelation in the process and collect past LTD forecast errors of the chosen model. One then uses an empirical distribution of these forecast errors, together with a point forecast for an upcoming period, to construct a LTD distribution and thereby calculate a critical fractile non-parametrically. This traditional empirical approach, which is often used in practice (Manary & Willems, 2008) and cited in textbooks (Cachon & Terwiesch, 2006), avoids the assumption that the forecasts are unbiased by checking whether the mean of past forecast errors is zero. If the mean of errors is not equal to zero, then it is used to correct for the bias in the most recent point forecast when calculating a critical fractile. However, this empirical approach can only deal with the unconditional forecast bias (such as the forecast bias in unconditional mean) and is not valid under autocorrelated demand. When the underlying demand is serially correlated, the point forecasts we make are conditional on the recent demand history (e.g., when a first-order autoregressive model (AR(1)) is the chosen forecasting model, the point forecasts are calculated conditional on the most recent demand) and the bias in the point forecasts, if there is any, would also be conditional on the recent demand history. Note that demand autocorrelation is indeed found to be present in numerous practical settings, such as consumer goods, fuel, food and machine tools (Chopra & Meindl, 2001; Erkip, Hausman, & Nahmias, 1990; Nahimas, 1993) and failure to model demand autocorrelation can have negative effect on the stockout rate experienced by a firm (Charnes, Marmorstein, & Zinn, 1995).

In this paper, we present an approach that corrects for the conditional forecast bias prior to using an empirical distribution of past forecast errors to forecast critical fractiles. The proposed approach can be viewed as an extension of the traditional empirical approach, which is suited to an autocorrelated demand process. Using a sample of forecast errors of the chosen forecasting model, we first estimate a parametric (linear) bias-adjusting regression model, conditional on the recent demand history. We then use the empirical residuals of this bias-adjusting regression model to estimate the critical fractile non-parametrically. The framework is semi-parametric as we specify the parametric bias-adjusting model but use the empirical distribution of the model errors to calculate the critical fractile.

Using asymptotic theory, we show that if the demand process follows a stationary autoregressive (AR(p)) demand process, the semi-parametric approach provides consistent estimates of the critical fractile, independently of the user-specified forecasting model, as long as it is within the autoregressive integrated (ARI($\tilde{p}, \tilde{d}$)) class. This is because if the user-specified model is misspecified, we can correct for the conditional forecast bias in constructing critical fractiles using the parametric bias-adjusting regression. No specific parametric assumptions about the forecast error distribution are necessary. The proposed approach therefore has practical relevance as it is applicable to forecasting models that are commonly used in practice (e.g., the autoregressive, random walk and independent demand models) and is robust against incorrectly specified error distributions. We also demonstrate that when the chosen forecasting model is correctly specified (i.e., the conditional bias is equal to zero), the semi-parametric approach collapses to the traditional empirical approach and the asymptotic variance of the critical fractile estimation is reduced. A smaller asymptotic variance is of practical importance as it implies lower inventory costs for a given sample size. The benefit of robustness against the unbiasedness of the forecasting model is traded off against the loss in efficiency resulting from a higher asymptotic variance.

To investigate the finite sample properties of the proposed framework we report on a simulation study. We demonstrate that, in comparison to the traditional parametric and empirical approaches, the proposed semi-parametric approach can significantly reduce inventory costs by correcting for a conditional forecast bias before estimating critical fractiles non-parametrically with empirical forecast errors. The cost reduction is increasing in sample size, lead time, and service level. Perhaps more importantly, our simulation results indicate that the proposed approach is a good heuristic for estimating critical fractiles under more general conditions than those specified by our asymptotic analysis. In particular, we show that the semi-parametric approach works well for the more general autoregressive moving average (ARMA) demand processes and under a broader class of user-specified forecasting models, which includes autoregressive integrated moving average (ARIMA) models. Finally, we illustrate the practical applicability of the proposed approach using real data in retailing. We find that the semi-parametric approach can significantly reduce inventory costs – by up to 57% – and also meet the required service levels better, compared to best-practice alternatives that fail to account for model uncertainty.

This paper is organized as follows: Section 2 introduces the critical fractile estimation problem when an arbitrary demand forecasting model is used; In Section 3, we describe the semi-parametric approach to estimating critical fractiles; Section 4 provides conditions for the asymptotic validity of the proposed approach; Section 5 provides numerical experiments to investigate the small sample properties of our asymptotic results; Section 6 shows the practical applicability of the proposed approach in retail inventory planning; and Section 7 concludes. All proofs are provided in the Supplementary material.

2. Estimating critical fractiles with an arbitrary demand forecasting model

We study the critical fractile estimation problem with a fixed replenishment lead time L where demand Y_t is serially correlated. Our main variable of interest is the lead time demand (LTD) in period $t, Y_t^L = \sum_{\tau=1}^{L} Y_{t+\tau}$. We allow the LTD to be dependent on the demand history $X_t = (Y_t, Y_{t-1}, \ldots, Y_{t-\dot{w}+1})'$, where $\dot{w}$ denotes the time-invariate length of an observation window. The distribution of the LTD in period t, conditional on the demand history $X_t = x$, is then defined as $\Psi_t(y|x) = \Pr(Y_t^L \leqslant y|X_t = x)$.

We assume that in period t, we first observe Y_t and determine the inventory target Q_t. Note that the inventory target Q_t needs to be calculated conditional on $X_t = x$ because the underlying demand process is serially correlated (i.e., Q_t is state-dependent). We also adopt the standard assumption of linear overage and underage costs, which we denote by h and s, respectively. The inventory target Q_t is then given by the following well-known Newsboy-type result (Cachon & Terwiesch, 2006):

$$Q_t = \min\{y : \Psi_t(y|x) \geqslant K\}, \tag{1}$$

where $K = \frac{s}{s+h}$ is the targeted service level (i.e., the likelihood of being able to meet demand with inventory) that minimizes the expected cost by balancing the costs of understocking and the costs of overstocking. (1) implies that the optimal inventory target Q_t is given by the Kth quantile of the conditional LTD distribution and this is known as the *critical fractile* solution.

We note that the critical fractile solution can be repeatedly applied when one takes a multi-period inventory model with a myopic order-up-to policy, in which the inventory position is

Download English Version:

https://daneshyari.com/en/article/479874

Download Persian Version:

https://daneshyari.com/article/479874

Daneshyari.com