



Contents lists available at ScienceDirect

Int. J. Production Economics

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

Bayesian forecasting for low-count time series using state-space models: An empirical evaluation for inventory management

Phillip M. Yelland*

Sun Microsystems Laboratories, 16 Network Circle, Menlo Park, CA 94025, USA

ARTICLE INFO

Available online 22 August 2008

Keywords:

Inventory management
Low-count time series
Bayesian statistics
State-space models

ABSTRACT

Inventories of optional components in discrete manufacturing are often subject to so-called *low-count* demand patterns. Quantities demanded from such inventories in any given period are sufficiently small that it may be unrealistic to forecast them with conventional models based on the normal distribution, and specialized models may be required. Fortunately, the statistical treatment of low-count time series has been the focus of much recent research. This paper recounts an attempt to apply some of this research to forecasting demands for optional parts at Sun Microsystems, a manufacturer and vendor of network computer products. Specifically, we compare the forecast performance of three simple state-space models using demand data obtained from Sun's inventory management records. The models are estimated using Bayesian methods, producing forecasts in the form of full predictive distributions. The accuracy of these probabilistic forecasts is compared using techniques borrowed from the field of meteorology, allowing us to assess the suitability of the candidate models for this type of application.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Fig. 1 displays a number of time series comprising the units of a selection of manufacturing parts used over a 78-week period in the operations of Sun Microsystems Inc., a manufacturer and vendor of network computing products. The parts in the figure are a subset of a larger sample of some 100 optional parts (that is, parts whose inclusion in a product depends upon configuration choices), drawn at random from Sun's inventory records.

Managing inventories for optional parts can be troublesome. In principle, with firm orders in hand for finished goods, materials requirements planning can be used to calculate parts demand over the short term through a straightforward bill-of-materials (BOM) "explosion" (Clement et al., 1995). In practice there are normally very many different types of such thinly demanded compo-

nents, so that the administrative overhead of entering them into the BOM and maintaining the correct BOM entries in the face of product changes, changes in parts specification or supplier, etc., is often prohibitive. Planning at longer horizons could also be achieved by BOM explosion of finished good demand forecasts, but even if all the parts are actually in the BOM, their presence or absence in the final product depends on particular configuration choices, which must themselves be forecast. In many instances, therefore, it is often expedient to forecast demands for optional parts directly.

As Fig. 1 illustrates, the time series in the sample are of a fairly idiosyncratic nature—an impression corroborated by Fig. 2—which displays the marginal distribution of weekly demand across the entire sample. From the latter, it is clear that the bulk of the values in the series are positive integers between 0 and 4, with zero occurring quite frequently (in fact, weeks with 0 units of demand constitute approximately 40% of all the weeks in the sample). None of the parts experienced weekly demands in excess of nine units during the period of observation.

* Tel.: +1 650 786 6864; fax: +1 650 786 7671.
E-mail address: phillip.yelland@sun.com

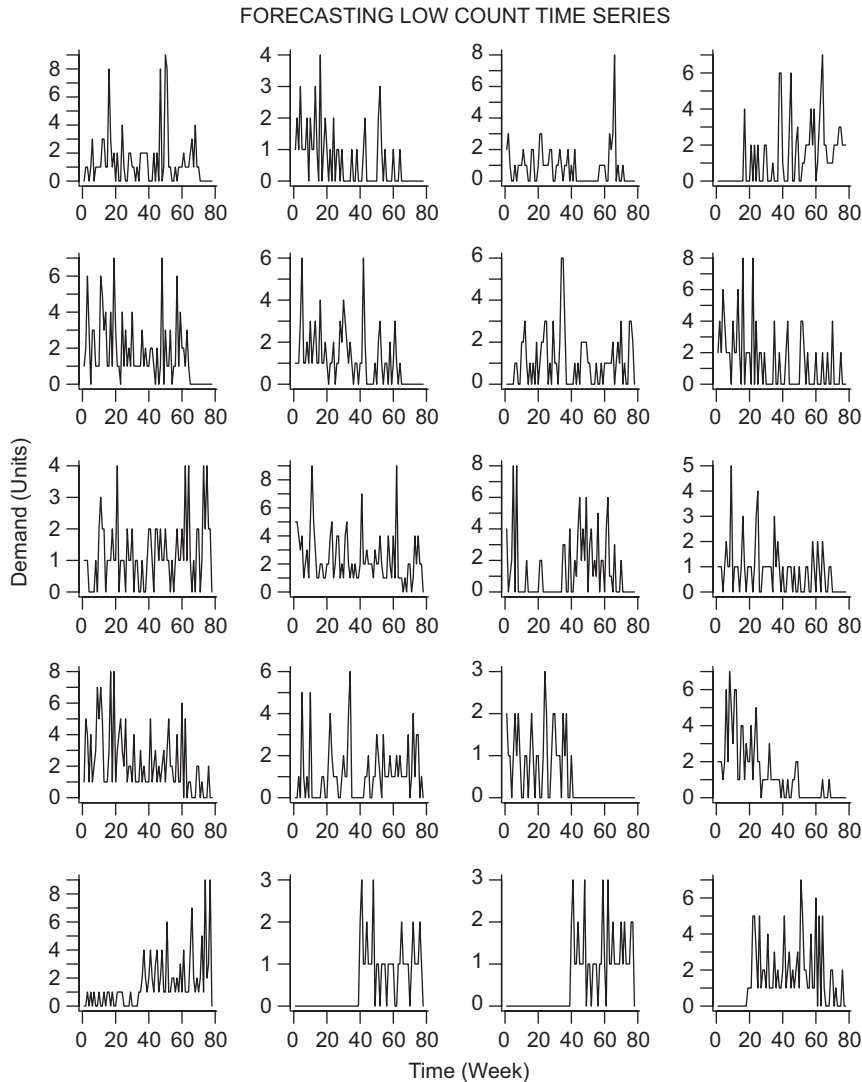


Fig. 1. Sample inventory demands.

McCabe and Martin (2005) refer to time series of this type as *low-count series*, distinguished in that low counts are poorly approximated by that staple of mainstream forecasting models, the normal distribution. (In contrast, series comprising larger count values are approximated much more felicitously.)

2. Forecasting low-count time series

Statistical modeling and prediction of low-count time series has become the focus of much attention in recent years; surveys of this work may be found in Cameron and Trivedi (1998), McKenzie (2003) and Winkelmann (1997). Many authors in the field use Cox's (1981) taxonomy to distinguish two different types of low-count model: in *observation-driven* models, dependence between values in a time-series is represented directly, usually by some form

of autoregressive or moving average mechanism. A *parameter-driven* model, in contrast, uses an underlying latent process to induce dependence between observations. No model of either type has achieved the dominance that was long enjoyed by Box–Jenkins/ARIMA models in forecasting continuously valued series. However, of the observation-driven models, the so-called *integer autoregressive* or INAR models described in Al-Osh and Alzaid (1990) and McKenzie (2003) have become increasingly popular. Correspondingly, most of the parameter-driven models use some form of *state-space* formulation, wherein an unobserved state vector evolves according to a Markovian process, with observations conditional upon some function of the state vector. Thorough discussion of state-space models may be found in West and Harrison (1997) and Durbin and Koopman (2001). In this paper, we concentrate on state-space models. While the potential of observation-driven models

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات