Stable Seasonal Pattern Models for Forecast Revision: A Comparative Study

Phillip M. Yelland
Sun Microsystems Inc.

MPK 16-158
16 Network Circle,
Menlo Park, CA 94025
United States.

e-mail: Phillip.Yelland@Sun.Com
tel: (650) 786-6864
fax: (650) 786-7671

Biography

Phillip M. Yelland is a senior staff engineer at Sun Microsystems Laboratories. As a researcher, his work places particular emphasis on the application of theoretical results to the practical conduct of business, and currently centers on the use of statistical and microeconomic techniques for operations management. He has an M.A. and Ph.D from the University of Cambridge in England, and an M.B.A. from the University of California at Berkeley.

Stable Seasonal Pattern Models for Forecast Revision: A Comparative Study

Abstract

Many firms prepare forecasts at the beginning of each financial quarter that predict total sales over the upcoming quarter. Such forecasts may be used to make financial projections, or to plan manufacturing capacity and materials purchases. As weekly sales are recorded during the quarter, these quarterly forecasts are often revised, allowing plans and projections to be adjusted appropriately. A formal basis for these forecast revisions may be found in so-called *stable seasonal pattern models*, which are based on the observation that in many instances, the sales that accrue during a given period of a quarter follow a regular pattern. This paper discusses a number of stable seasonal pattern models — several from the literature, two that are novel — which have been evaluated for making forecast revisions at Sun Microsystems, Inc. Commonalities between the models are elucidated using a general theoretical framework, and a straightforward sample-based mechanism is described that affords great flexibility in the design and use of stable seasonal pattern models. The paper culminates in a detailed comparison of the performance of new and existing stable seasonal pattern models with respect to Sun's sales data.

Keywords: Marketing/Production Forecasting, Seasonal Variation

1 Introduction

Many businesses begin each financial quarter with a forecast of total sales expected over the quarter. Such a forecast might be used in making financial projections, or to plan manufacturing capacity and make materials purchases to satisfy anticipated sales. As the quarter progresses, the original forecast is often revised in the light of the actual sales that accrue week-by-week, so that financial projections or manufacturing and supply arrangements can be adjusted accordingly.

A number of *stable seasonal pattern models* (*SSPM*s) have been advanced in the research literature that provide a formal basis for such interim forecast revisions. Though they differ in detail, all SSPMs are based on the observation that the sales that occur during a given period of a quarter typically conform to a regular pattern involving the end-of-quarter total. Having discerned this pattern, therefore, it should be possible to use the actual sales made during part of the quarter to adjust an initial estimate of the end-of-quarter total.

This paper discusses a selection of stable seasonal pattern models that have been evaluated by Sun Microsystems Inc. (a major vendor of network computer products) for making week-by-week revisions of quarterly sales forecasts. We begin by illustrating product sales that are typical of those encountered at Sun, and patterns that show up in their cumulative weekly totals. Next, a general framework is used to describe existing stable seasonal pattern models from the literature, and two new SSPMs are introduced. We also show how the application of a simple form of *importance sampling* (Tanner, 1996, p. 54) affords great flexibility in the design and use of stable seasonal pattern models. Using this sampling mechanism, several existing SSPMs, along with the two novel ones, are applied to Sun's forecast revision problem. With simulations based on actual sales series of selected Sun products, the forecast performance of each model and its fit to the data are determined and compared with those of the others. The paper concludes with general comments about the SSPMs examined, and possible directions for future development.

2 Sales and bookings at Sun

Much of the sales forecasting activity within Sun is focused on *bookings* placed during a particular time period, such as a week or quarter. A period's bookings comprise tentative orders for products placed during that period. Such bookings are subject to possible cancellation at a later date, and no revenue is recorded until items are actually manufactured and shipped. In spite of their tentative nature, bookings are used widely within the company as an indication of customer demand, since the number of actual sales realized during a given period are often significantly affected by deliberate actions on the part of the company, such as the allocation of manufacturing capacity, procurement decisions, inventory/backlog policy and the like.

Figure 1 displays the weekly bookings for a product that is representative of many of those encountered at Sun. The series runs from the product's introduction to its eventual obsolescence. Two features of the series are notable, and also typical for Sun's products: First, cancelled orders and certain forms of stock rotation on the part of distributors are represented by *negative* bookings, so that it is possible to record a negative weekly bookings total. Second, weekly bookings are difficult to forecast directly with accuracy, since they are volatile and non-stationary. Thus while it might be possible to derive quarterly bookings totals by aggregating forecasts of weekly bookings (from a seasonal ARIMA model, for instance), our experience with such an approach has been disappointing.

Forecasting with a stable seasonal pattern model, on the other hand, begins by with a direct forecast of quarterly bookings. Figure 2 depicts the quarterly bookings totals of the product in Figure 1. Though such quarterly bookings series are by nature much shorter than the corresponding weekly series, they are normally more amenable to forecasting, in part because they are much smoother, but in the main because a substantial amount of prior information that can be brought to bear on them. Sales force, financial and marketing estimates of product sales, for example, are normally calculated on a quarterly basis, and managers have a much clearer picture of product lifecycles at the quarterly level. Such prior information makes it possible to use Bayesian methods to produce initial quarterly bookings forecasts with reasonable accuracy, as demonstrated in (Yelland, 2003). Once an initial forecast for a given quarter's total has been prepared, the stable seasonal pattern model may be used to update it as the quarter progresses.

As an example of patterns that emerge in weekly bookings — the pattern, in fact, that forms the basis of the new SSPMs introduced in this paper — the box-and-whisker plot in Figure 3 displays the distributions of cumulative weekly bookings for the product in Figure 1 as fractions of the total sales in the corresponding quarters. So over the 17 quarters of the product's lifecycle, it appears that about 50% of the total quarterly bookings have accrued by week 7 of the quarter, 80% by week 11, and so on.

3 Stable seasonal pattern models

Work on stable seasonal pattern models dates back over four decades, and during that period a wide variety of such models have been introduced. The framework set out below, which builds on that given by Chen and Fomby (1999), attempts to draw together most of the salient common characteristics of SSPMs from the literature for the purposes of this paper. Though sufficient for its application here, the framework is not intended to be definitive or wholly general; it is unabashedly Bayesian in nature, and proceeds from less fundamental assumptions than the work of Chen and Fomby, for example.

Consider a series of weekly bookings, b_{qw}, indexed by quarter q, and week $w = 1, 2, \ldots, 13$ within the quarter. At the end of week w of quarter q, we have observed a sequence of weekly bookings $\mathbf{b}_{qw} \triangleq (b_{q1}, ..., b_{qw})$ in the current quarter. We also have an information set, $\mathcal{H}_{qw}$, containing the history of the bookings process. The precise contents of $\mathcal{H}_{qw}$ are elaborated below, but we can assume that it contains at least records of weekly bookings for the quarters preceding q:

$$\{b_{q',1}, \ldots, b_{q',13}\} \in \mathcal{H}_{qw}, \text{ for all } q' < q \tag{1}$$

Let $B_q \triangleq \sum_{i=1}^{13} b_{qi}$ denote the sum of all bookings in the quarter q — that is, the total of the bookings observed in weeks $1, \ldots, w$ and those still to be recorded in weeks $w+1, \ldots, 13$. The forecast to be revised specifies a prior distribution for the quarterly total, $p(B_q \mid \mathcal{H}_{qw})$, and updating the forecast in light of the current week's bookings amounts to the calculation of the posterior distribution $p(B_q \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})$.

A stable seasonal pattern model rests on a summary quantity, $S(\mathbf{b}_{qw})$, derived from $\mathbf{b}_{qw}$, whose distribution is determined by the end-of-quarter total B_q and a collection of parameters that are independent of B_q. This distribution determines a likelihood for B_q, so that having observed $\mathbf{b}_{qw}$, an application of Bayes' rule furnishes the required posterior.

In formal terms:

Definition A *stable seasonal pattern model* for the bookings process described above consists of:

- A function $S(\cdot)$ of the observed bookings in a quarter, $\mathbf{b}_{qw}$.
- A probability distribution $\mathrm{F}(\cdot,\cdot)$, indexed by the quarterly bookings total B_q and a parameter vector, $\boldsymbol{\theta}_w$, associated with the same week in every quarter.

Such that:

- The value $S(\mathbf{b}_{qw})$ is distributed according to $\mathrm{F}(B_q, \boldsymbol{\theta}_w)$ independently of the process history:

$$S(\mathbf{b}_{qw}) \,|\, B_q, \boldsymbol{\theta}_w, \mathcal{H}_{qw} \;\sim\; \mathrm{F}(B_q, \boldsymbol{\theta}_w) \tag{2}$$

- B_q and $\boldsymbol{\theta}_w$ are a priori independent:

$$p(B_q, \boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw}) = p(B_q \,|\, \mathcal{H}_{qw})\; p(\boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw}) \tag{3}$$

- $S(\mathbf{b}_{qw})$ is a *sufficient statistic* for B_q (given $\mathcal{H}_{qw}$), implying that:

$$p(B_q \,|\, \mathbf{b}_{qw}, \mathcal{H}_{qw}) = p(B_q \,|\, S(\mathbf{b}_{qw}), \mathcal{H}_{qw}) \tag{4}$$

■

In principle, $S(\mathbf{b}_{qw})$ may be an arbitrary vector, including $\mathbf{b}_{qw}$ itself (c.f. (Chen & Fomby, 1999), for example). However, in the great majority of existing SSPMs — and all of those examined in detail in this paper — it is simply the sum of the components of $\mathbf{b}_{qw}$. In the following, the distribution $\mathrm{F}(\cdot,\cdot)$ is often referred to as the model's *seasonal distribution*.

To demonstrate how an SSPM facilitates the calculation of the revised forecast, $p(B_q \,|\, \mathbf{b}_{qw}, \mathcal{H}_{qw})$, note that by applying Bayes' rule to assumption (4), we have that:

$$\begin{aligned} p(B_q \,|\, \mathbf{b}_{qw}, \mathcal{H}_{qw}) &= p(B_q \,|\, S(\mathbf{b}_{qw}), \mathcal{H}_{qw}) \\ &\propto p(S(\mathbf{b}_{qw}), B_q \,|\, \mathcal{H}_{qw}) \end{aligned} \tag{5}$$

The latter expression may be obtained by marginalizing the seasonal parameters $\boldsymbol{\theta}_w$ out of the expression $p(S(\mathbf{b}_{qw}), B_q, \boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw})$, which may in turn be simplified derived using assumptions (2) and (3):

$$\begin{aligned} & p(S(\mathbf{b}_{qw}), B_q \,|\, \mathcal{H}_{qw}) \\ &\quad = \int p(S(\mathbf{b}_{qw}), B_q, \boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw})\, d\boldsymbol{\theta}_w \\ &\quad = \int p(S(\mathbf{b}_{qw}) \,|\, B_q, \boldsymbol{\theta}_w, \mathcal{H}_{qw})\; p(B_q, \boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw})\, d\boldsymbol{\theta}_w \\ &\quad = \int p(S(\mathbf{b}_{qw}) \,|\, B_q, \boldsymbol{\theta}_w)\; p(B_q \,|\, \mathcal{H}_{qw})\; p(\boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw})\, d\boldsymbol{\theta}_w \\ &\quad = p(B_q \,|\, \mathcal{H}_{qw}) \int p(S(\mathbf{b}_{qw}) \,|\, B_q, \boldsymbol{\theta}_w)\; p(\boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw})\, d\boldsymbol{\theta}_w \end{aligned} \tag{6}$$

To summarize:

$$p(B_q \,|\, \mathbf{b}_{qw}, \mathcal{H}_{qw}) \propto p(B_q \,|\, \mathcal{H}_{qw}) \int p(S(\mathbf{b}_{qw}) \,|\, B_q, \boldsymbol{\theta}_w)\; p(\boldsymbol{\theta}_w \,|\, \mathcal{H}_{qw})\, d\boldsymbol{\theta}_w \tag{7}$$

Observe that the integral in the final term is essentially the marginal likelihood for B_q according to the seasonal distribution in (2).

In some SSPMs, the assertion in (2) is expressed indirectly using a function, $T(\cdot,\cdot)$, of $S(\mathbf{b}_{qw})$ and B_q, together with a distribution $\mathrm{G}(\cdot)$ indexed by $\boldsymbol{\theta}_w$ alone, such that:

$$T(S(\mathbf{b}_{qw}), B_q) \,|\, \boldsymbol{\theta}_w \;\sim\; \mathrm{G}(\boldsymbol{\theta}_w) \tag{8}$$

A transformation of variables derives the distribution of $S(\mathbf{b}_{qw})$ from (8); in particular, when both $T(\cdot,\cdot)$ and $S(\mathbf{b}_{qw})$ are scalar (the case for all the models examined here), using the expression $T^{(1,0)}$ to denote the partial derivative of *T* with respect to its first argument:

$$p(S(\mathbf{b}_{qw}) \,|\, B_q, \boldsymbol{\theta}_w) \;=\; \left| T^{(1,0)}(S(\mathbf{b}_{qw}), B_q) \right| \; p(T(S(\mathbf{b}_{qw}), B_q) \,|\, \boldsymbol{\theta}_w) \tag{9}$$

It's possible to distinguish between what might be termed *sequential* and *non-sequential* updating methods for use with SSPMs, according to the way in which the information set, $\mathcal{H}_{qw}$, changes as a quarter progresses. In a non-sequential method, the information set used in each week of the quarter is the same, so that, $\mathcal{H}_{qw} = \mathcal{H}_{q1}$ for $w = 1, \ldots, 13$, and all information pertaining to bookings observed in the current quarter is conveyed by $S(\mathbf{b}_{qw})$. In contrast, sequential methods update the information set as each week's bookings are observed. They derive $\mathcal{H}_{qw+1}$ by adding an updated estimate of B_q to $\mathcal{H}_{qw}$, so that the forecast distribution $p(B_q \,|\, \mathcal{H}_{qw+1})$ reflects the most recent set of bookings.

Table 1 reviews the more prominent stable seasonal pattern models from the literature in the context of the general framework above. In each case, the seasonal distribution assumed in each model is stated in accordance with the form of equation (2) or (8), and the updating method intended for the model is noted.

As we can see from Table 1, the binomial and normal distributions are staple features of SSPMs. The former represents a quarter's bookings as a collection of independent events of two classes, namely those appearing on or before the week when the forecast revision is to be made, and those appearing afterwards. The normal distribution is used as an approximation to the binomial for large enough bookings totals in the second model of Oliver (1987), and for the sake of expedience in (Chang & Fyffe, 1971) and (Guerrero & Elizondo, 1997). The choice of the inverse Pareto distribution in (Mendoza & de Alba, 2002) is also one of convenience. It should be noted that neither the binomial nor the inverse Pareto distribution provides support for negative values like those found occasionally in Sun's bookings totals. While Chen and Fomby (1999) use a model very similar to Oliver's binomial model in their "discrete" case (where bookings totals are small), the normal distribution in their second model represents a generalization of the binomial — rather than an approximation, as with Oliver — to accommodate variables with measurement error or random perturbation. Chen and Fomby actually propose a richer model that also accounts for the incremental (as distinct from cumulative) weekly bookings within a quarter, and which is discussed later on in this paper.

Table 1 includes SSPMs with both sequential and non-sequential updating. All of the models updated non-sequentially base their inferences on the cumulative bookings in the current quarter up to and including the current week. The sequentially-updated models of Chang and Fyffe and Oliver use the Kalman filter to update the forecast of end-of-quarter bookings in each week. Guerrero and Elizondo (1997) describe non-sequential and sequential updating methods for their model, though the integrated moving average technique they describe for sequential updating places restrictions on the kinds of quarterly totals series that can be accommodated.

To simplify the Bayesian updating operation formalized in (7), the majority of the SSPMs in the table restrict the form of prior forecast for quarterly totals, $p(B_q \mid \mathcal{H}_{qw})$ so that it is conjugate with the likelihood induced by the seasonal distribution of the SSPM. (For full statistical inference, the Kalman filter used in the Chang and Fyffe and Oliver models relies implicitly on the conjugacy of a normally-distributed prior forecast.) The restriction to conjugate priors makes for convenient and efficient closed-form updating, but can be burdensome in practice. Guerrero and Elizondo do not require conjugate priors, since they use frequentist updating procedures, and bypass the operation in (7). Instead, as was indicated above, they impose fairly severe structural restrictions on the series used to represent quarterly bookings totals.

Of those that do use Bayesian principles for forecast updating, only the SSPM of Mendoza and de Alba uses a fully Bayesian approach to estimate the parameter vector, $\boldsymbol{\theta}_w$, associated with the seasonal distribution. The others might be termed “empirical Bayes” models (Carlin & Louis, 2000), in that they use either least squares or maximum-likelihood methods to produce estimates for the seasonal parameters.

4 Two simple stable seasonal pattern models

This section introduces two new simple SSPMs developed at Sun to update bookings forecasts. We make no claim that these new SSPMs are uniformly superior to the existing ones, but they do appear to be well suited to the bookings processes encountered at Sun, as we demonstrate in the following sections.

Both of the models are based on the observation (illustrated informally in Figure 3) that cumulative weekly bookings for Sun's products are typically a constant proportion of the quarterly total. One obvious way of formalizing this is to assert that the ratio of cumulative bookings in week w to the total quarterly bookings in each quarter is normally distributed around some mean μ_w with variance ν_w. Of course, using a normal distribution for such ratios gives positive support to cumulative bookings that are less than zero or greater than the quarterly total in some weeks, but such an occurrence is not wholly out of the question in the presence of negative bookings like those we saw in Section 2. An SSPM based on normally-distributed cumulative ratios may be summarized in terms of the general framework in the previous section:

$$\begin{aligned} &S(\mathbf{b}_{qw}) = \textstyle\sum_{i=1}^{w} b_{qi} \\ &\boldsymbol{\theta}_w = (\mu_w, \nu_w) \\ &\left.\begin{aligned} &T(s, B_q) = B_q^{-1} s \\ &\mathrm{G}((\mu,\nu)) = \mathrm{N}(\mu,\nu) \end{aligned}\right\} \quad \Rightarrow\ B_q^{-1} \textstyle\sum_{i=1}^{w} b_{qi} \mid \mu_w, \nu_w \sim \mathrm{N}(\mu_w, \nu_w) \end{aligned} \tag{10}$$

Then by (9) and (7), forecast revision takes the following form:

$$\begin{aligned} p(B_q \mid \mathbf{b}_{qw}, \mathcal{H}_{qw}) &\propto p(B_q \mid \mathcal{H}_{qw}) \int \left|B_q^{-1}\right| p(B_q^{-1} \textstyle\sum_{i=1}^{w} b_{qi} \mid \mu_w, \nu_w)\, p(\mu_w, \nu_w \mid \mathcal{H}_{qw})\, d(\mu_w, \nu_w) \\ &= p(B_q \mid \mathcal{H}_{qw}) \left|B_q^{-1}\right| \int p(B_q^{-1} \textstyle\sum_{i=1}^{w} b_{qi} \mid \mu_w, \nu_w)\, p(\mu_w, \nu_w \mid \mathcal{H}_{qw})\, d(\mu_w, \nu_w) \end{aligned} \tag{11}$$

As we demonstrate next, standard results for Bayesian inference in normal models allow the integral in the above to be solved exactly.

Consider a forecast revision that takes place in week w of quarter $T+1$, with knowledge of the bookings in the corresponding week of quarters $1, \ldots, T$. We use r_{qw} to denote the ratio of the cumulative bookings by week w of quarter q to total quarterly bookings:

$$r_{qw} \triangleq B_q^{-1} \textstyle\sum_{i=1}^{w} b_{qi} \tag{12}$$

The independence assumption in (2) implies that:

$$p(r_{1w}, \ldots, r_{Tw} \mid \mu_w, \nu_w) = \prod_{q=1}^{T} p(r_{qw} \mid \mu_w, \nu_w) \tag{13}$$

Therefore, with $p(\mu_w, \nu_w)$ as the initial prior for the model parameters, we have:

$$p(\mu_w, \nu_w \mid \mathcal{H}_{T+1,w}) \propto \left[\prod_{q=1}^{T} p(r_{qw} \mid \mu_w, \nu_w)\right] p(\mu_w, \nu_w) \tag{14}$$

Substituting into (11) and expanding out the joint density of μ_w and ν_w, the updated forecast in week w of quarter $T+1$ is given by the rather formidable expression:

$$p(B_{T+1}\,|\,\mathbf{b}_{T+1,w},\mathcal{H}_{T+1,w}) \propto$$

$$p(B_{T+1}\,|\,\mathcal{H}_{T+1,w})\left|B_{T+1}^{-1}\right|\iint p(r_{T+1,w}\,|\,\mu_w,\nu_w)\left[\prod_{q=1}^{T} p(r_{q\,w}\,|\,\mu_w,\nu_w)\right]p(\mu_w,\nu_w)\,d\mu_w\,d\nu_w \qquad (15)$$

where

$$r_{q\,w} \sim \mathrm{N}(\mu_w,\nu_w), \text{ for } q=1,\ldots,T+1$$

Fortunately, it is possible to simplify this expression markedly by selecting the reference prior $p(\mu_w,\nu_w)\propto 1/\nu_w$ to convey a (not unreasonable) lack of prior information about the model parameters (the precise choice of non-informative prior is largely for convenience — see (Bernardo and Smith, 2000, section 5.4), for example, for a discussion). Then standard results for Bayesian inference (Gelman et al., 1995, section 3.2) imply that the double integral in the above (which is proportional to the *posterior predictive distribution* for r_{T+1}) takes the form of a scaled Student's t distribution, yielding a formula for the revised quarterly forecast:

$$p(B_{T+1}\,|\,\mathbf{b}_{T+1,w},\mathcal{H}_{T+1,w}) \propto p(B_{T+1}\,|\,\mathcal{H}_{T+1,w})\left|B_{T+1}^{-1}\right|p(r_{T+1,w}\,|\,\overline{r},s^2)$$

where

$$r_{T+1,w}\,|\,\overline{r},s^2 \sim \mathrm{t}_{T-1}(\overline{r},s^2) \qquad (16)$$

$$\overline{r}=\frac{1}{T}\sum\nolimits_{t=1}^{T} r_t,$$

$$s^2=\frac{T+1}{T(T-1)}\sum\nolimits_{t=1}^{T}\left(r_t-\overline{r}\right)^2$$

The mechanics of the next step in the forecast revision process — the combination of the prior forecast with the likelihood from the stable seasonal pattern model specified in the above — are taken up in the next section.

The normal seasonal distribution assumed in this model is somewhat heterodox, since common practice would constrain the fractions $r_{q\,w}$ to lie in the interval $[0,1]$. For the purposes of comparison with existing SSPMs, therefore, we consider a variant of the foregoing model that describes the weekly cumulative fractions by using a logistic normal, rather than a normal, seasonal distribution. In calibrating this model, the occasional negative ratio is discarded as an outlier. Within the general framework, this second model is specified as follows:

$$S(\mathbf{b}_{q\,w})=\sum\nolimits_{i=1}^{w} b_{q\,i}$$

$$\boldsymbol{\theta}_w=(\mu_w,\nu_w) \qquad (17)$$

$$\left.\begin{array}{l} T(s,B_q)=\log\left(\dfrac{B_q^{-1}s}{1-B_q^{-1}s}\right) \\ \mathrm{G}((\mu,\nu))=\mathrm{N}(\mu,\nu) \end{array}\right\} \Rightarrow \log\left(\frac{r_{q\,w}}{1-r_{q\,w}}\right)\Big|\,\mu_w,\nu_w \sim \mathrm{N}(\mu_w,\nu_w)$$

In this case, the forecast revision formula analogous to (16) is:

$$p(B_{T+1} \mid \mathbf{b}_{T+1,w}, \mathcal{H}_{T+1,w}) \propto p(B_{T+1} \mid \mathcal{H}_{T+1,w}) \left| \frac{1}{(1-r_{T+1,w})\sum_{i=1}^{w} b_{T+1,i}} \right| p\left(\log\left(\frac{r_{T+1,w}}{1-r_{T+1,w}}\right) \mid \overline{r}, s^2\right)$$

where

$$\log\left(\frac{r_{T+1,w}}{1-r_{T+1,w}}\right) \mid \overline{r}, s^2 \sim \mathrm{t}_{T-1}(\overline{r}, s^2) \tag{18}$$

$$\overline{r} = \frac{1}{T}\sum_{t=1}^{T} r_t,$$

$$s^2 = \frac{T+1}{T(T-1)}\sum_{t=1}^{T}\left(r_t - \overline{r}\right)^2$$

5 Importance sampling for forecast updating

To make practical use of the two simple SSPMs described in the previous section, we need a suitable representation of the prior forecast $p(B_q \mid \mathcal{H}_{qw})$ in the revision formulae (16) and (18). As we observed in Section 3, existing SSPMs rely on conjugate prior forecast distributions to allow closed-form updates. Unfortunately, as we also pointed out earlier, such conjugacy requirements substantially constrain the form of the quarterly forecast and the seasonal distribution that can be used in an SSPM. Furthermore, when quarterly forecasts take the form of samples produced using Monte Carlo techniques (as they are at Sun — see (Yelland 2003)), expressing them in a closed form conjugate with any seasonal distribution is inconvenient, if not impossible. Since Monte Carlo techniques are becoming increasingly widespread as a means of estimating more detailed and realistic statistical models, an inability to deal with them effectively would constitute a significant impediment to the use of SSPMs. Fortunately, a very simple updating technique accommodates both sample-based quarterly forecasts and a wide range of seasonal distributions in SSPMs.

First, note that the density of the revised forecast distribution in its general form in (7) — and in its specific forms in (16) and (18) — may be written:

$$p(B_q \mid \mathbf{b}_{qw}, \mathcal{H}_{qw}) = c\, L(B_q \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})\, p(B_q \mid \mathcal{H}_{qw}) \tag{19}$$

Here, the term $L(B_q \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})$ is the likelihood of B_q determined by the SSPM, and c is a normalizing constant.

Many characteristics of the revised forecast, including the mean, standard deviation and quantiles of the forecast distribution, may be derived by calculating — for some function $\psi(\cdot)$ of B_{T+1} — expectations of the form:

$$\int \psi(B_q)\, p(B_q \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})\, dB_q \tag{20}$$

Importance sampling (Tanner 1996, p. 54) approximates the value in (20) with the sum:

$$\sum_{i=1}^{N}\left(\frac{\omega_i}{\sum_{j=1}^{N}\omega_j}\psi(\hat{B}_i)\right) \tag{21}$$

Here, $\hat{B}_1,\ldots,\hat{B}_N$ are samples drawn from an *importance distribution* with density $h(\cdot)$ that is intended roughly to resemble $p(B_q \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})$, and the *importance weights* $\omega_1,\ldots,\omega_N$ are defined (up to the same constant of proportionality):

$$\omega_i \propto \frac{p(\hat{B}_i \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})}{h(\hat{B}_i)} \tag{22}$$

In the case of forecast revision using an SSPM, a very expedient choice for the approximating distribution is the prior quarterly forecast, which as we indicated above, is computed by Monte Carlo sampling in our work at Sun. Thus the samples $\hat{B}_1,\ldots,\hat{B}_N$ are given directly by the sample-based quarterly forecast, and in light of (19), the importance weights are simply:

$$\begin{aligned}\omega_i &\propto \frac{p(\hat{B}_i \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})}{p(\hat{B}_i \mid \mathcal{H}_{qw})}\\ &\propto \frac{L(\hat{B}_i \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})\, p(\hat{B}_i \mid \mathcal{H}_{qw})}{p(\hat{B}_i \mid \mathcal{H}_{qw})}\\ &= L(\hat{B}_i \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})\end{aligned} \tag{23}$$

The values in (21) become:

$$\sum_{i=1}^{N}\left(\frac{L(\hat{B}_i \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})}{\sum_{j=1}^{N} L(\hat{B}_j \mid \mathbf{b}_{qw}, \mathcal{H}_{qw})}\psi(\hat{B}_i)\right) \tag{24}$$

In summary, therefore, pertinent characteristics of the revised forecast can be derived simply by weighting each sample value in the prior forecast with the (normalized) likelihood it is given by the SSPM. As Tanner (1996) points out, provided the empirical distribution of prior forecast sample is not too far from the revised one (i.e. provided the prior forecast gives sufficient support to values around the actual end-of-quarter total), expressions like (24) are reasonably close approximations of the corresponding true values.

6 Comparing forecast performance

In the process of selecting a suitable stable seasonal pattern model at Sun, we set out to compare the performance of seven candidate SSPMs on a representative sample of Sun's bookings data. The candidate models are listed below, together with the abbreviations used to refer to them in the following discussion:

DU — The "dummy" SSPM that merely repeats the prior quarterly forecast in each week.

BN — A model based on binomial seasonal distribution for cumulative weekly bookings, like the first models in (Oliver, 1987) and (Chen & Fomby, 1999).

LR — A model based on the linear regression of quarterly totals against cumulative weekly bookings. It is essentially a Bayesian version of the frequentist models in (Guerrero & Elizondo, 1997).

MA — The inverse Pareto model of (Mendoza & de Alba, 2002).

CF — The second model in (Chen & Fomby, 1999).

SN — The simple normal SSPM described in Section 4.

LN — The logistic normal SSPM in Section 4.

The importance sampling technique of the previous section allowed us to use the same prior quarterly forecasts with all the candidate models. All of the candidate models (including the novel SN and LN models) were updated non-sequentially. The sequential updating methods used with the models of Chang and Fyffe (1971) and Oliver (1987) rely on the use of the Kalman filter, and so they are awkward to adapt to sample-based priors (see (West & Harrison, 1989, chapter 15) for details of the techniques involved).

As test data, we used the weekly bookings histories of 25 products from Sun's range of computer servers. The first of these series is displayed in Figure 1, and the remaining 24 are shown in Figure 4. The products represented range from budget-priced single-processor servers with sales in the tens of thousands of units per quarter to expensive multi-processor models selling in the hundreds per quarter. End-of-quarter totals for any one product also vary, reflecting the product's progression through its lifecycle. Since the products were of differing longevities and differing maturities at the time of selection, and because the initial quarters of some products were lost to bookkeeping procedures, the corresponding bookings series vary in length from 5 to 17 quarters, with an average length of around 11 quarters (all series were truncated if necessary to an integral number of quarters).

To produce test calibration data, successively longer initial segments of each bookings series were taken, beginning with a segment of 4 quarters in length (the minimum length required reliably to calibrate all the models tested), and repeatedly adding another quarter's bookings. The quarter of bookings following the calibration quarters was used as a holdout sample. Thus a complete bookings series $\left(w_1,\ldots,w_{13},\ldots,w_{13n+1}\ldots,w_{13n+13}\right)$ would yield the calibration-holdout sets:

$$\begin{gathered}
\begin{Bmatrix}\left(w_1,\ldots,w_{13},\ldots,w_{13\times3+1}\ldots,w_{13\times3+13}\right)\\ \left(w_{13\times4+1}\ldots,w_{13\times4+13}\right)\end{Bmatrix},\\
\begin{Bmatrix}\left(w_1,\ldots,w_{13},\ldots,w_{13\times3+1}\ldots,w_{13\times3+13},w_{13\times4+1}\ldots,w_{13\times4+13}\right)\\ \left(w_{13\times5+1}\ldots,w_{13\times5+13}\right)\end{Bmatrix},\\
\vdots\\
\begin{Bmatrix}\left(w_1,\ldots,w_{13},\ldots,w_{13\times3+1}\ldots,w_{13\times3+13},w_{13\times4+1}\ldots,w_{13\times4+13},\ldots,w_{13(n-1)+1}\ldots,w_{13(n-1)+13}\right)\\ \left(w_{13n+1}\ldots,w_{13n+13}\right)\end{Bmatrix}
\end{gathered} \tag{25}$$

To produce a test case for the candidate SSPMs, a calibration-holdout set was matched with a prior quarterly forecast. In keeping with our practice at Sun, where SSPMs are used to revise quarterly forecasts produced independently by judgment or by statistical tools such as that described in (Yelland, 2003), we used informative priors in the test, introducing errors intended to mirror the errors we customarily encounter in actual quarterly forecasts. Accordingly, each prior quarterly forecast took the form of 1,000 samples from a normal distribution, the mean of which was itself normally distributed around the true quarterly total, B_q, with a standard deviation of 20% of B_q. (Within reasonable limits, the actual distribution chosen for the quarterly priors affects the relative performance of the SSPMs only mildly; furthermore heuristic forecasts are commonly described in terms of normal distributions, and they provide a reasonable approximation to the output of our statistical tools.) The standard deviation of the forecast was calculated so as to give approximately correct confidence intervals, facilitating the comparison of forecast uncertainty detailed in the next section. In precise terms, the prior quarterly forecast distribution took the form:

$$
\begin{aligned}
\mathrm{N}(\mu, \sigma^2), \text{ where } \mu &\sim \mathrm{N}(B_q, (0.2B_q)^2), \\
\sigma &\sim \mathrm{Uniform}(0.5d, 3d), \\
d &= \left|\mu - B_q\right| / \Phi_{90}, \\
\Phi_{90} &\text{ is the } 90^{\text{th}} \text{ percentile of the standard normal distribution}
\end{aligned}
\tag{26}
$$

Each calibration-holdout set from the bookings series was paired with 32 different generated priors, yielding approximately 5,000 test cases; in the following, the number of test cases is denoted *N*. Each of the candidate models was run on the same *N* test cases. For each model and each test case, the model's parameters were estimated from the calibration data in the test case and then weeks 1 to 12 of the holdout quarter's bookings were passed successively to the calibrated model, along with the test case's prior quarterly forecast, thus simulating a quarter's weekly updates. (Of course, by the end of week 13 of the holdout quarter, the quarterly total is known.) In each week, the updated forecast was summarized by calculating the median and the 10^{th} and 90^{th} percentiles of the posterior distribution produced by the model. When testing models such as LN or BN, with seasonal distributions that give no support to weekly bookings totals that are negative or that exceed the quarterly total, the posterior distribution for any week in which such totals occurred (which in fact amounted to less than 0.1% of the test cases) simply replicated the posterior distribution computed in the previous week.

7 Test results

The results of the test described above comprised the summaries of $7 \times N \times 13$ distributions, namely, for each of the 7 candidate models in each of the *N* test cases, the beginning prior for the quarterly total together with the revised distributions produced in the light of bookings in weeks 1 to 12.

To compare the performance of the candidate models in providing point forecasts, we selected the median of the updated forecast distributions to represent point forecasts of the end-of-quarter total. We felt that the median would provide a point estimate that was more robust to possible sampling errors than the mean of the distribution, though as it transpires, the results are not altered substantively by selecting the mean instead.

Unfortunately, summarizing and comparing the performance of point forecasts across series of widely differing scales like those in the test is a complex and frequently vexatious problem, as (Armstrong, 2001) indicates. In our study, we followed the precept set out in (Armstrong & Collopy, 1992), using a yardstick based on the *relative absolute error* (*RAE*).

Let m index the candidate models, reserving the index 0 for the "dummy" model (DU) that simply repeats the prior quarterly forecast. Let q index the N simulated quarters in the test, and w index weeks in which forecasts were made. Let $f_{m,q,w}$ denote the point forecast produced by model m in week w for simulated quarter q, and let B_q be the actual end-of-quarter bookings total. Then the relative absolute error of forecast $f_{m,q,w}$ is defined:

$$\mathrm{RAE}_{m,q,w} \triangleq \frac{| f_{m,q,w} - B_q |}{| f_{0,q,w} - B_q |} \tag{27}$$

In other words, the RAE is simply the ratio of the absolute error in the forecast produced by model m to the absolute error of the dummy model's forecast for the same week. Taking the logarithm of the RAE makes for a more symmetrical distribution of measurements (supporting the use of convenient hypothesis tests such as the t test), and negating the result makes the resulting comparisons more intuitive (a higher score reflects a better forecast). Thus our performance measure, here denoted Γ, is defined:

$$\Gamma_{m,q,w} \triangleq -\log\left(\mathrm{RAE}_{m,q,w}\right) \tag{28}$$

A summary the model m's performance across all test cases for a given week w can be derived by calculating the arithmetic mean of the corresponding Γs. (In practice, outlying values of $\Gamma_{m,q,w}$ which may result from aberrantly small forecast errors are discarded before taking the mean.) The result is equivalent to the negated logarithm of the geometric mean of the RAE — commonly abbreviated as the *GMRAE* — of the model:

$$\begin{aligned}\overline{\Gamma}_{m,w} &\triangleq \frac{1}{N}\sum_{q=1}^{N}\left[-\log\left(\mathrm{RAE}_{m,q,w}\right)\right] \\ &= -\log\left[\left(\prod_{q=1}^{N}\mathrm{RAE}_{m,q,w}\right)^{\frac{1}{N}}\right] \\ &= -\log\left(\mathrm{GMRAE}_{m,w}\right)\end{aligned} \tag{29}$$

A value $\overline{\Gamma}_{m,w} = 0$ indicates that on average the model did as well as the dummy model in the week in question; the more positive (negative) the value of $\overline{\Gamma}_{m,w}$, the better (worse) the model did.

The forecast performance of the candidate models (relative to that of the dummy model) is displayed in Figure 5, where the bars signify the values of $\overline{\Gamma}_{m,w}$ for each of the models in weeks 1 to 12. The names of the corresponding model is displayed against each bar, with the symbol "+" whenever a t test on the values $\Gamma_{m,q,w}$ indicates that the difference between the mean value $\overline{\Gamma}_{m,w}$

for the model in question and that of the model with the highest value of $\overline{\Gamma}_{m,w}$ is not statistically significant.

Recall that the test described in the previous section also yielded confidence bounds associated with the forecasts in the form of 10th and 90th percentiles of the revised distributions. We would expect such an 80% confidence interval, if accurate, to contain the actual quarterly total 80% of the time, and to the extent that coverage deviates from 80%, the interval can be regarded as in error. To this end, Figure 6 displays the actual coverage of the forecast confidence interval for each model in each week.

(In view of the fact — noted by one of the referees to this paper — that models such as the binomial and Mendoza-de Alba were devised for forecasting applications without negative bookings, we repeated the tests in this section for a modified data-set in which negative bookings were eliminated. The results — which are available from the author on request — differ only insignificantly from those presented here.)

8 Discussion

On the basis of the summaries in Figure 5, it appears that the overall point forecast performance of the two simple models of Section 4 is certainly respectable, even if it is not uniformly superior to that of the other candidate models. The Guerrero-Elizondo LR model performs better than the SN and LN models in the very early weeks of the quarter (reflecting those test cases when the LR model's constant term accurately reflects an approximately constant quarterly bookings total for a particular product). In weeks 4 through 12, however, either one or both of the simple models is amongst the best performers, with the SN model excelling in the middle weeks of the quarter, and the LN model towards the end. The accuracy of the Chen-Fomby (CF) model is only once amongst the best, though it does overtake the LR model in the latter half of the quarter, and is superior to the SN model in week 12. The Mendoza-de Alba (MA) model appears ill-suited to this particular data set, and its performance is largely mediocre, though it improves markedly in the latter weeks of the quarter. With the exception of weeks 5, 6 and 7 (where it ranks only second-to-last), the binomial model (BN) is the poorest forecaster, though it too improves as the quarter progresses.

Turning to the coverage of the confidence intervals displayed in Figure 6: As was indicated above, the simulated prior forecasts were calculated to give correct confidence intervals in aggregate and this is made evident by the highlighted markers for the DU model, which simply reiterate the coverage of the initial forecast in each week of the figure. Relative to the initial forecast, only the Mendoza-de Alba (MA) model produces adequate coverage, though the confidence intervals it produces in the later weeks of the quarter are actually slightly too wide. It is clear that the confidence intervals produced by the binomial model (BN) are far too narrow, and those of the related CF (Chen-Fomby) model are consistently worse than the other models barring BN. In all by the final two weeks, the LR (linear regression), SN (simple normal) and LN (logistic normal) models tend to produce intervals of approximately the same width, with marginally less coverage than required.

Though the focus of this appraisal is on the suitability of the candidate models for forecasting, an examination of their respective in-sample fit might be illuminating. Unfortunately, with the longest of the test series providing only some 17 data points, and with many series in the test set considerably shorter, it is very difficult to draw statistically significant conclusions about goodness-of-fit. It is possible to produce a rough picture, however, by pooling summaries of fit to each of the series in the following manner: Each model is calibrated using all the data points from each series in turn, and for each week of the quarter, a bootstrap sample of the projected ratio of cumulative bookings in that week to the end-of-quarter total is generated from the model (see (Davison & Hinkley, 1997), for example, or (Gelman et al., 1996) for procedural details). The bootstrap samples, together with the corresponding actual ratios for each series are collected together. Figure 7 presents a comparison of projected to actual ratios in the form of quantile-quantile plots for each candidate model in each week of the quarter. Each plots the quantiles of the actual ratios (on the x-axis) against the corresponding quantiles of projected ratios calculated from the model. Each plot also displays the (dotted) line $x = y$, to which a plot would conform if the populations of projected and actual ratios were to match completely. Plots are annotated with the p-values of the corresponding 2-sample Kolmogorov-Smirnov test, with an "R" appended if the test rejects at the 5% level the hypothesis that the projected and actual populations are identical. Due to the possible confounding of points from different series, the results in Figure 7 should be treated with caution. Nonetheless, they do provide at least an indication of the models' in-sample fit, and a bad fit in the aggregate is unlikely to result from a uniformly good fit to the individual series.

Caveats notwithstanding, the poor fits of both the BN and MA models — both of which fail to fit in any week of the quarter — may help explain their erratic forecast performance. It is interesting to note, however, that while the CF model has demonstrably the best in-sample fit, its forecast performance is almost invariably inferior to those of the LN and SN models, even though the latter fit less well. Given the overly-narrow forecast intervals derived using the CF model, the fact that the CF model is calibrated using maximum-likelihood parameter estimates, rather than using the full Bayesian approach employed with the LN and SN models, it is possible that the CF model is over-fitting the data. A similar contention applies to the richly-parameterized LR model, which fits better in the latter weeks of the quarter (when its forecast performance is relatively poorer) than it does in the first weeks (when, as we pointed out earlier, the constant in the model delivers a better forecast).

9 Conclusion

Our tests indicate that for this application at least, the simple stable seasonal pattern models described in Section 4 perform well in comparison with those in the literature. This is surprising, perhaps, given that they are essentially "theory-free" descriptions of the bookings process, incorporating little in the way of insightful explanation of the process itself. Such findings resonate with a theme in the forecasting literature (Allen & Fildes, 2001); namely, that simple, empirically-based models frequently do better than more complex ones that seek to embody more accurate accounts of the phenomena they represent.

In this connection, it is interesting to contrast the simple SN and LN models with the CF candidate model. In point of fact, the latter derives its characterization of cumulative weekly totals from an account of incremental weekly bookings given by the *Gaussian-multinomial distribution* defined by Chen and Fomby in their paper. The Gaussian-multinomial is a multivariate distribution that generalizes the multinomial distribution, providing a specification of the joint distribution of bookings in the weeks of a quarter. One might expect that such a detailed description might add to the forecast accuracy of the CF model, but as we have seen, in this particular setting its forecast performance is middling. Furthermore, Figure 8 summarizes a bootstrap goodness-of-fit test for the CF model of the same nature as that illustrated in Figure 7, this time involving incremental rather than cumulative weekly bookings. As the figure shows, despite its superior fit to cumulative bookings, there is evidence that the CF model fails consistently to fit even the marginal distributions of the incremental bookings — still less their joint distribution. (It should be noted that Chen and Fomby mention a generalized form of the Gaussian-multinomial in their paper, which may yield better results. However, they provide no details of its calibration, which depends in any case on the selection of suitable observable covariates in each application.)

However, it would certainly be possible to improve upon the performance of the simple SSPMs. For example, as Figure 5 illustrates, neither the SN nor the LN model provides the best forecasts for every week of the quarter, and in fact the LR (regression-based) model out-performs them both in the first three weeks. One remedy might be to switch to the best-performing SSPMs as the quarter progresses, from LR to SN and thence to LN, as the results in Figure 5 imply. However, our test results suggest that none of the candidate SSPMs contribute greatly to forecast accuracy with only one or two week's bookings data — consider that in weeks one and two, Figure 5 indicated that even the best-performing LR model yields errors on average 90% those of the initial forecast. At the opposite end of the quarter, it's normally too late in weeks 11 and 12 to use a revised forecast effectively. Such considerations might make it hard in practice to justify the added complication introduced by switching models.

An alternate approach would be to use a more flexible seasonal distribution, such as a mixture (West and Harrison, 1997, Chp. 12) or the *skew-normal* distribution of (Azzalini, 1985). Or a more general form of transformation such as a Box-Cox power transformation (Box & Cox, 1964) might improve upon the simple logit used in the LN model. Again, the calibration and extrapolation of such models would be more involved than is the case with the SN and LN models.

We've also investigated the incorporation of covariates into the simple seasonal pattern models. One might expect, say, that a given class of products experiences a unique bookings pattern throughout a quarter, or that the position of a quarter in the financial year affects the pattern of weekly bookings. Technically, such additions are not difficult to accommodate. Thus far, however, we have found no covariates for use in such models (including lagged values of the ratios themselves computed from earlier quarters) that yield forecast revisions significantly better than those produced by the simple models without covariates.

One factor we have found to be critical to the effective use of SSPMs is an accurate statement of uncertainty in the prior quarterly forecast used as the basis for revision during the quarter. With a prominent judgmental component in our forecasting process (see (Yelland, 2003) for details), the initial quarterly forecasts we encounter in practice err not infrequently on the side of over-confidence. An over-confident quarterly forecast retards the adjustments made by an SSPM, resulting in any forecast inaccuracies being carried late into the quarter. We have found that procedural safeguards against overconfidence in the quarterly forecasts — such as those described in (Arkes, 2001) — can help significantly improve the performance of the SSPM.

While we hope that we have demonstrated the suitability of the simple SSPMs for the purposes of forecast revision at Sun, since to date, our experience with these models has been restricted largely to Sun's bookings data. However, the simple normal (SN) model in particular has been in use at the company for some time, and it has been used to update quarterly forecasts produced by the model in (Yelland, 2003) for a range of products in a variety of market circumstances. The reception afforded the resulting forecast improvements by operations planners and financial analysts at the company has been consistently enthusiastic. In light of the apparent effectiveness, and the demonstrable technical simplicity and straightforward implementation offered by the simple normal SSPM, we would certainly recommend it for consideration in similar forecasting situations.

10 Acknowledgement

The author is very grateful to the editorial staff of the *IJF*, and to two anonymous referees, whose comments and suggestions contributed greatly to the improvement of this paper.

11 References

Allen, G. & Fildes, R. (2001). Econometric forecasting, in (Armstrong, 2001a), pp. 303 – 362.

Arkes, H., 2001. Overconfidence in judgmental forecasting, in (Armstrong, 2001a), pp. 495 – 515.

Armstrong, J. (ed) (2001a). *Principles of Forecasting*, Kluwer.

Armstrong, J. (2001b). Evaluating forecasting methods, in (Armstrong, 2001a), pp. 443 - 472.

Armstrong, J. & Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons, *International Journal of Forecastin*g, 8, pp. 69 – 80.

Azzalini A. (1985). A class of distributions which includes the normal ones. *Scand. J. Statist.*, 12, pp. 171 – 178.

Box, G. E. P. & Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society (Series B)*, 26, pp. 211-252.

Carlin, B. & Louis, T. (2000). *Bayes and Empirical Bayes Methods for Data Analysis* (2nd ed.), Chapman & Hall/CRC.

Chang, S. H. & Fyffe, E. E. (1971). Estimation of forecast error for seasonal style-goods sales, *Management Science*, 16, pp. 89 – 96.

Chen, R. & Fomby, T. (1999). Forecasting with stable seasonal pattern models with an application to Hawaiian tourism data, *Journal of Business and Economic Statistics*, 17(4), pp. 497 – 504.

Davison, A. & Hinkley, D. (1997). *Bootstrap Methods and their Application*. Cambridge University Press.

Gelman, A., Meng, X.-L. & Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies, *Statistica Sinica* 6, pp. 733 – 807.

Gelman, A., Carlin, J., Stern, H., & Rubin, D. (1995). *Bayesian Data Analysis*. Chapman and Hall/CRC.

Guerrero, V. M. & Elizondo, J. A. (1997). Forecasting a cumulative variable using its partially accumulated data. *Management Science*, 43, pp. 879 – 889.

Mendoza, M. & de Alba, E. (2001). Forecasting an accumulated series based on partial accumulation: a Bayesian method for short series with seasonal patterns, *Journal of Business and Economic Statistics*, 19(1), pp. 95 – 102.

Oliver, R. M. (1987). Bayesian forecasting with stable seasonal patterns. *Journal of Business and Economic Statistics*, 5(1), pp. 77 – 85.

Tanner, M. (1996). *Tools for Statistical Inference* (3rd ed.), Springer.

West, M. & Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models* (2nd ed.), Springer.

Yelland, P. (2003). *A Model of the Product Lifecycle for Sales Forecasting*, Sun Microsystems Laboratories Technical Report 2003-0715.

Reference	Seasonal Distribution	Model Features	Forecast Revision	Model Estimation
Chang and Fyffe (1971)	$b_{qw} \sim \mathrm{N}(m_w B_q, \sigma_w^2)$	$S(b_{q1},\dots,b_{qw}) = b_{qw}$ $\boldsymbol{\theta}_w = (m_w, \sigma_w)$	Sequential (Kalman filter)	OLS, MLE
Oliver (1987) *(two models)*	$\sum_{i=1}^{w} b_{qi} \sim \mathrm{Bin}(B_q, \lambda_w)$	$S(b_{q1},\dots,b_{qw}) = \sum_{i=1}^{w} b_{qi}$ $\boldsymbol{\theta}_w = (\lambda_w)$	Non-sequential (conjugate prior)	Unspecified
	$\sum_{i=1}^{w} b_{qi} \sim \mathrm{N}\left(B_q\, p_w, B_q\, p_w(1-p_w)\right)$	$S(b_{q1},\dots,b_{qw}) = \sum_{i=1}^{w} b_{qi}$ $\boldsymbol{\theta}_w = (p_w)$	Sequential (Kalman filter)	Unspecified
Guerrero and Elizondo (1997)	$B_q - \left(\beta_{w0} + \beta_{w1} \sum_{i=1}^{w} b_{qi}\right) \sim \mathrm{N}(0, \sigma_w^2)$	$S(b_{q1},\dots,b_{qw}) = \sum_{i=1}^{w} b_{qi}$ $\boldsymbol{\theta}_w = (\beta_{w0}, \beta_{w1}, \sigma_w)$ $T(s, B_q) = B_q - \left(\beta_{w0} + \beta_{w1}\, s\right)$	Non-sequential (non-informative prior), Sequential (IMA)	OLS
Chen and Fomby (1999) *(two models)*	$\sum_{i=1}^{w} b_{qi} \sim \mathrm{Bin}(B_q, \gamma_w)$	$S(b_{q1},\dots,b_{qw}) = \sum_{i=1}^{w} b_{qi}$ $\boldsymbol{\theta}_w = (\gamma_w)$	Non-sequential (conjugate prior)	MLE
	$\sum_{i=1}^{w} b_{qi} \sim \mathrm{N}\left(B_q\, \gamma_w, \sigma^2\, \gamma_w(1-\gamma_w)\right)$	$S(b_{q1},\dots,b_{qw}) = \sum_{i=1}^{w} b_{qi}$ $\boldsymbol{\theta}_w = (\gamma_w, \sigma)$	Non-sequential (conjugate prior)	MLE
Mendoza and de Alba (2002)	$B_q \left(\sum_{i=1}^{w} b_{qi}\right)^{-1} \sim \mathrm{InvPareto}(\theta_w)$	$S(b_{q1},\dots,b_{qw}) = \sum_{i=1}^{w} b_{qi}$ $\boldsymbol{\theta}_w = (\theta_w)$ $T(s, B_q) = B_q\, s^{-1}$	Non-sequential (conjugate prior)	Bayesian (non-informative prior)

Table 1: Existing stable seasonal pattern models

Key:

$\mathrm{N}(\cdot,\cdot)$	—	Normal distribution	OLS	—	Ordinary least squares
$\mathrm{Bin}(\cdot,\cdot)$	—	Binomial distribution	MLE	—	Maximum likelihood estimation
$\mathrm{InvPareto}(\cdot,\cdot)$	—	Inverse Pareto distribution; $p(r) = \theta_w\, r^{1-\theta_w}$	IMA	—	Integrated moving average

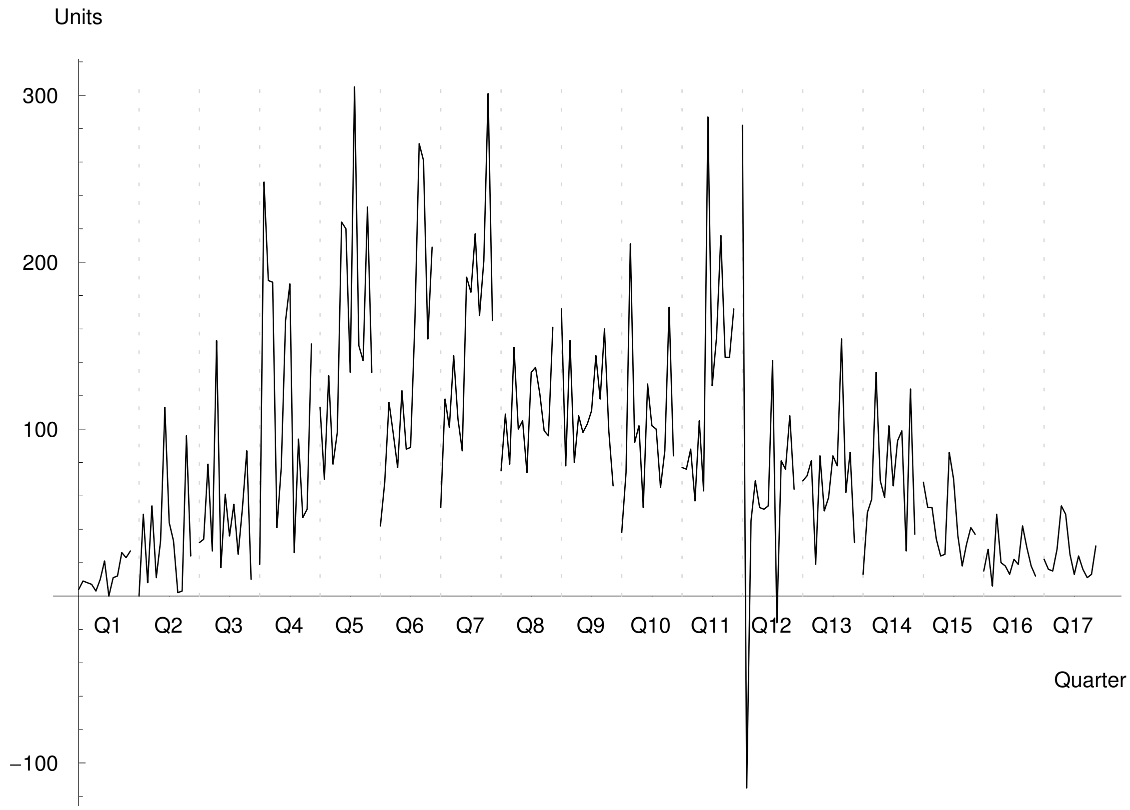

Figure 1: Weekly sales series for a sample product

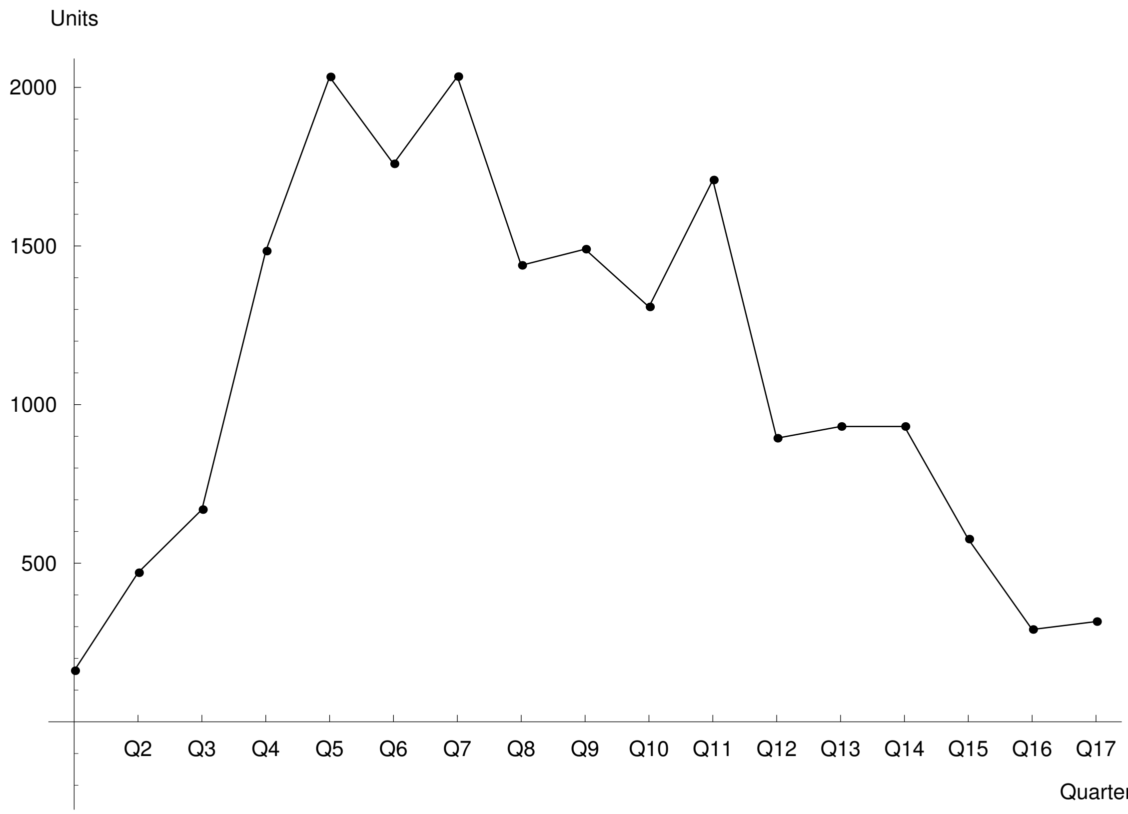

Figure 2: Quarterly bookings totals for the sample product

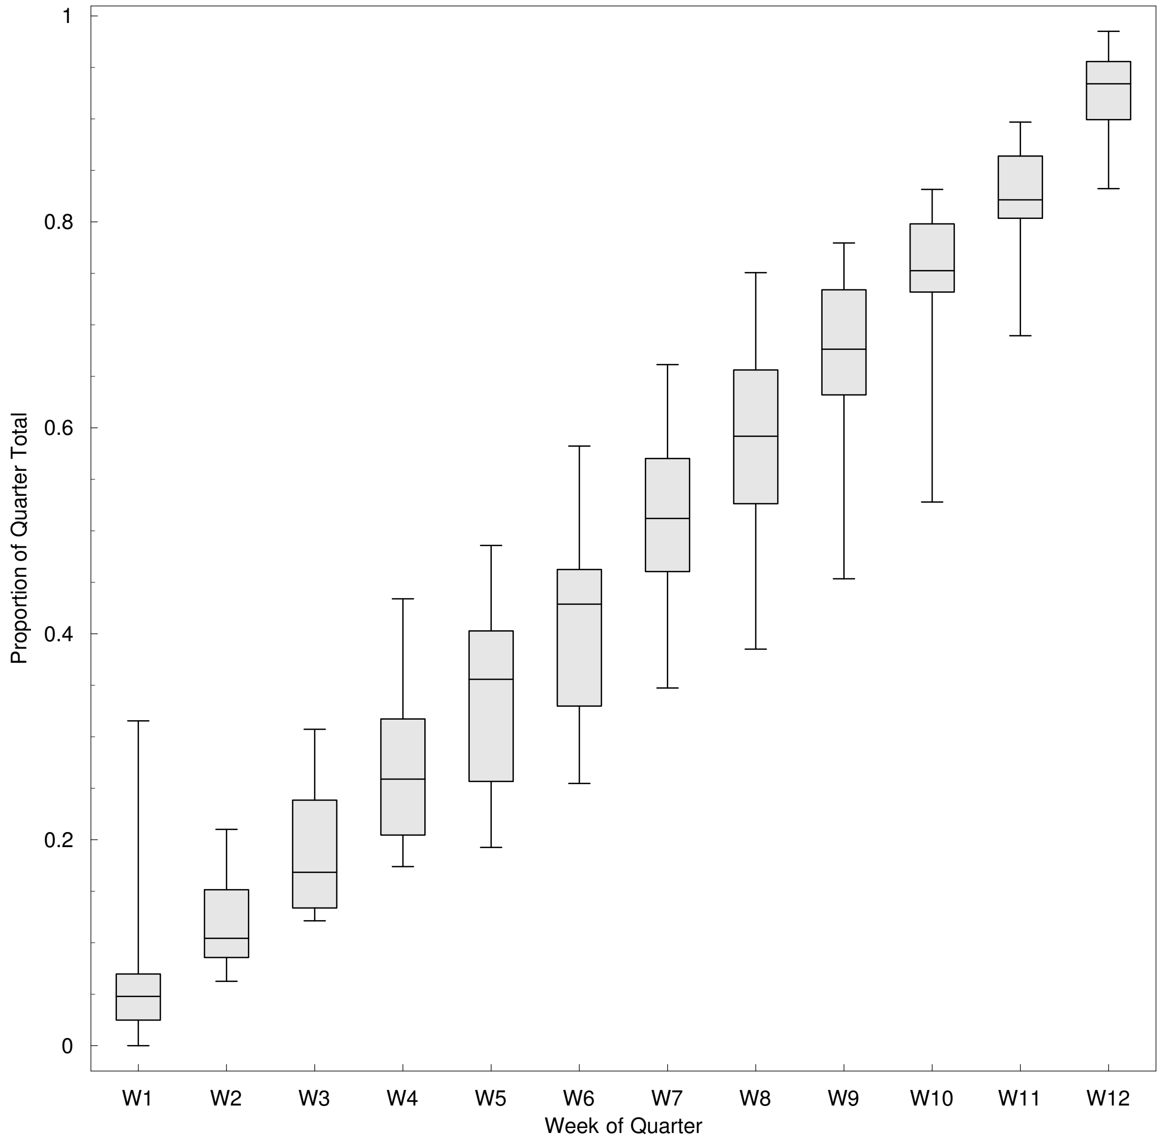

Figure 3: Cumulative proportions by week for the sample product

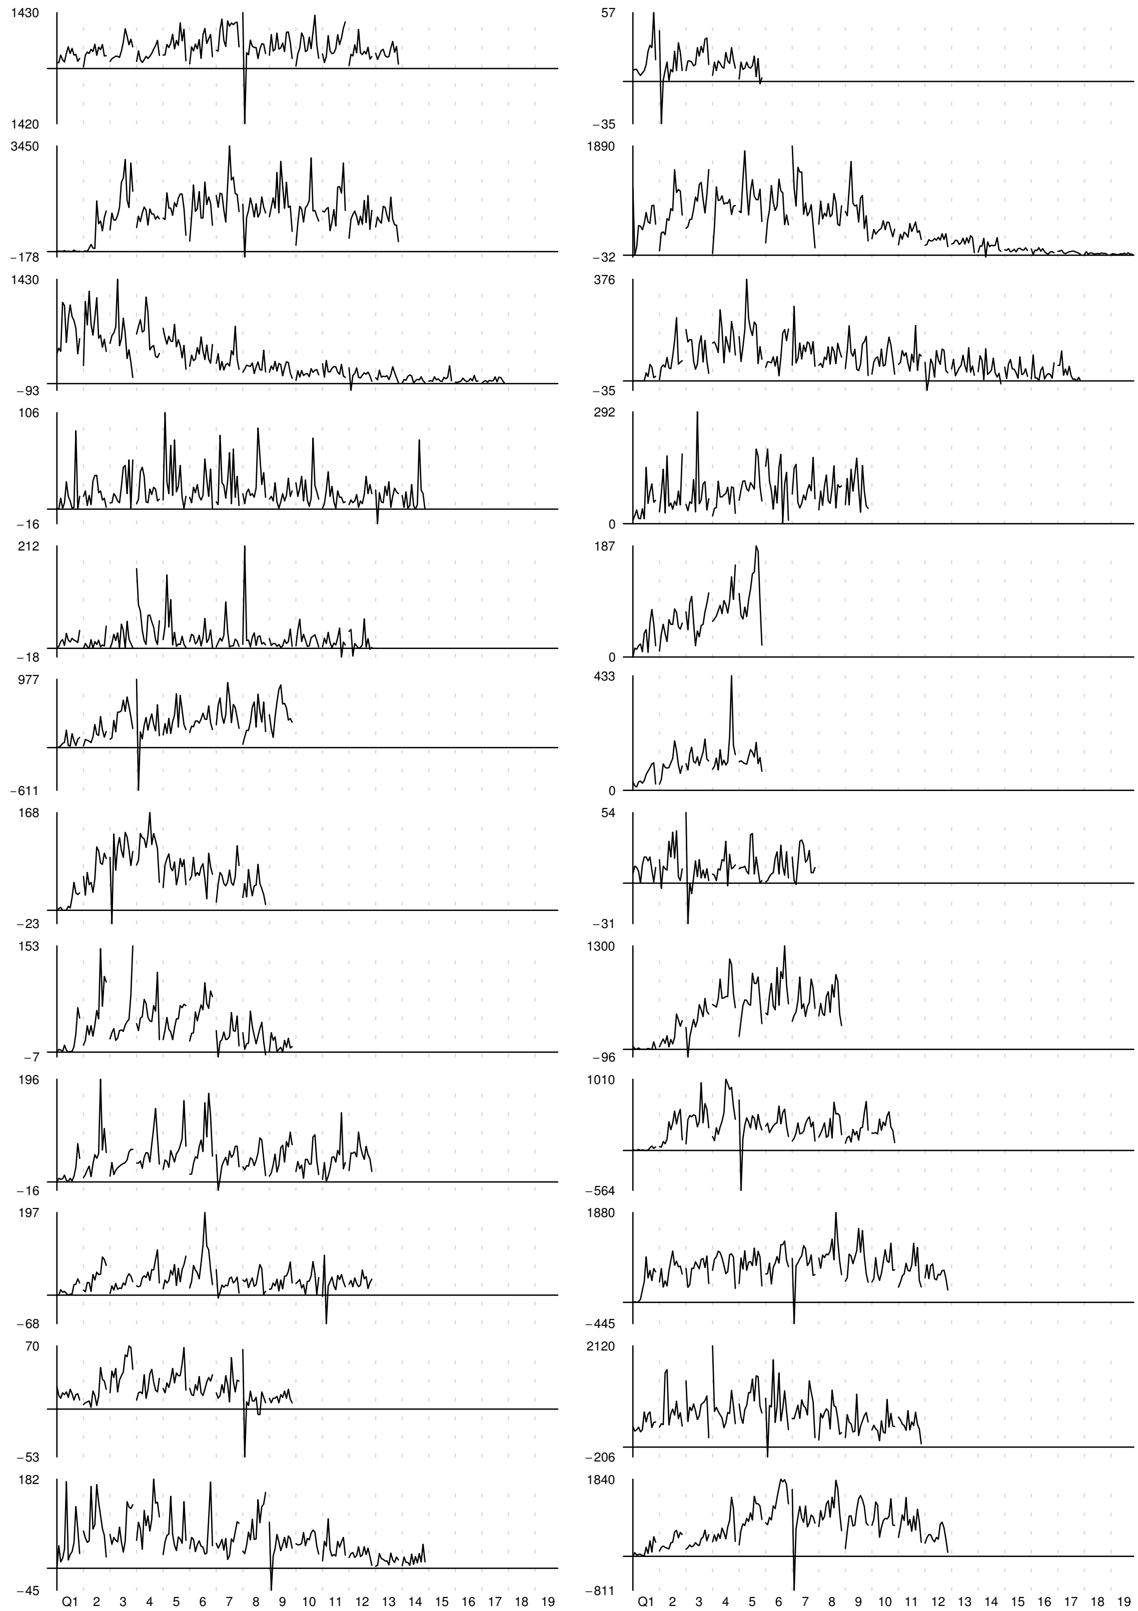

Figure 4: Test bookings series

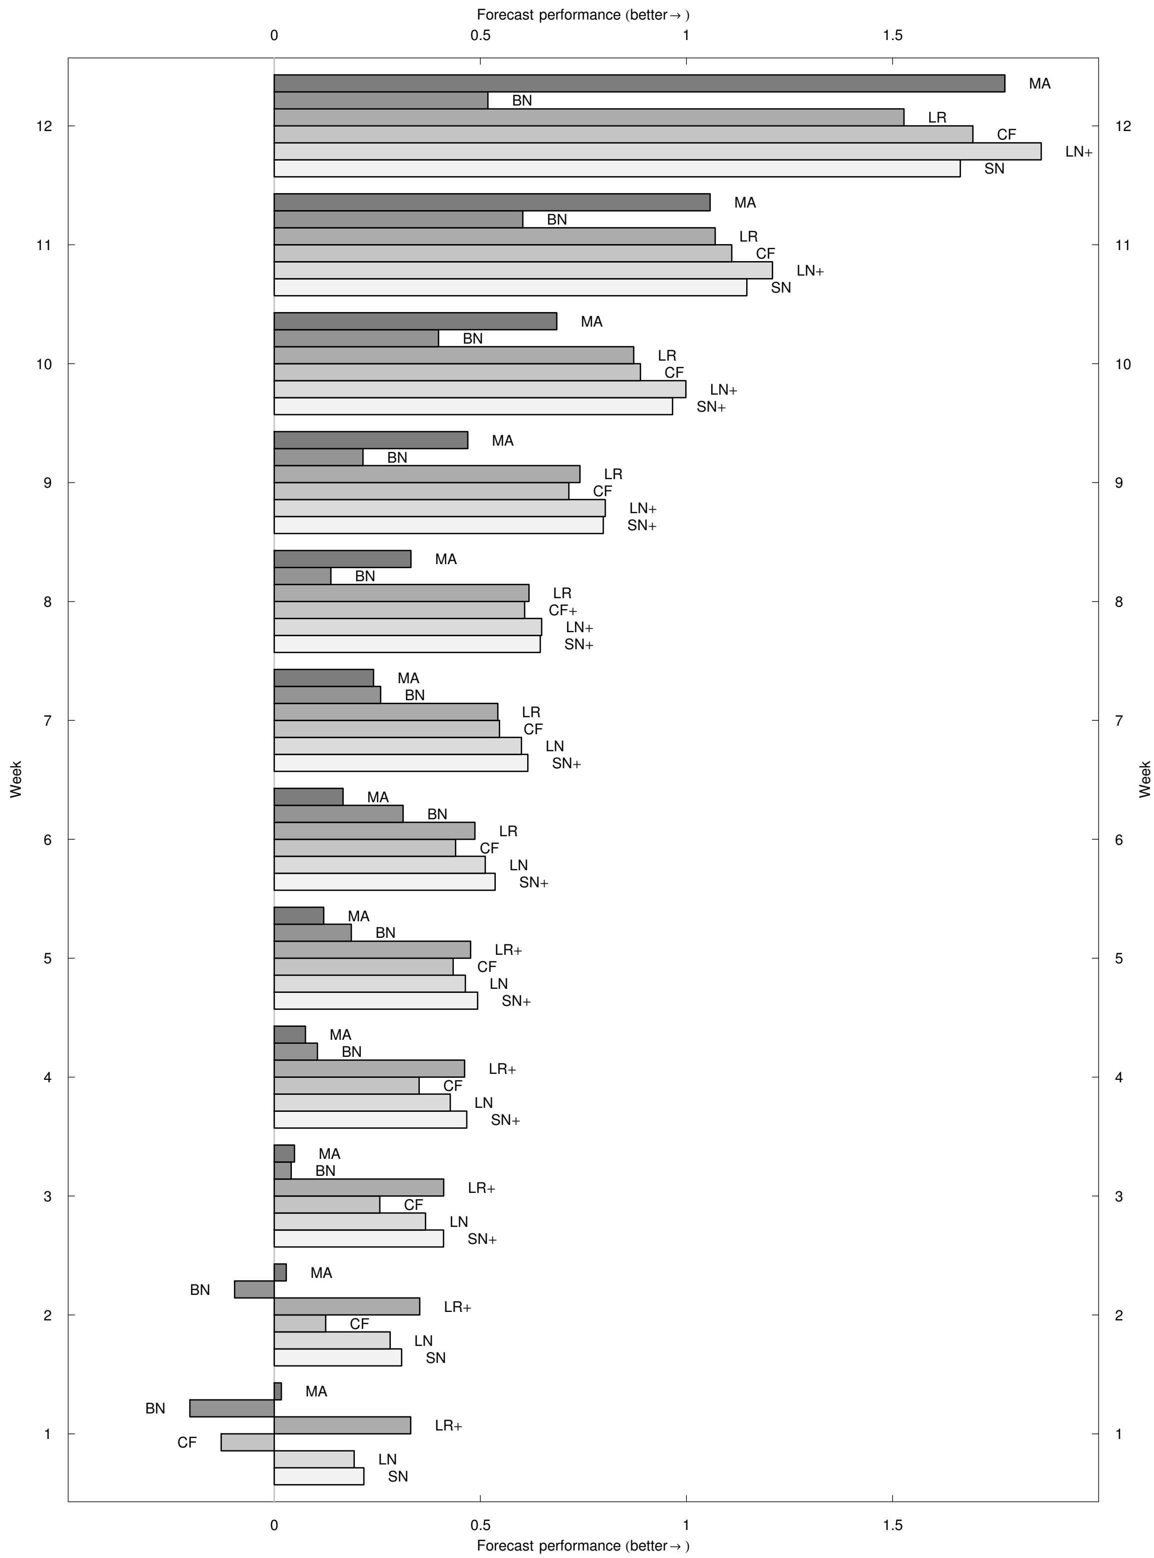

Figure 5: Relative point forecast performance

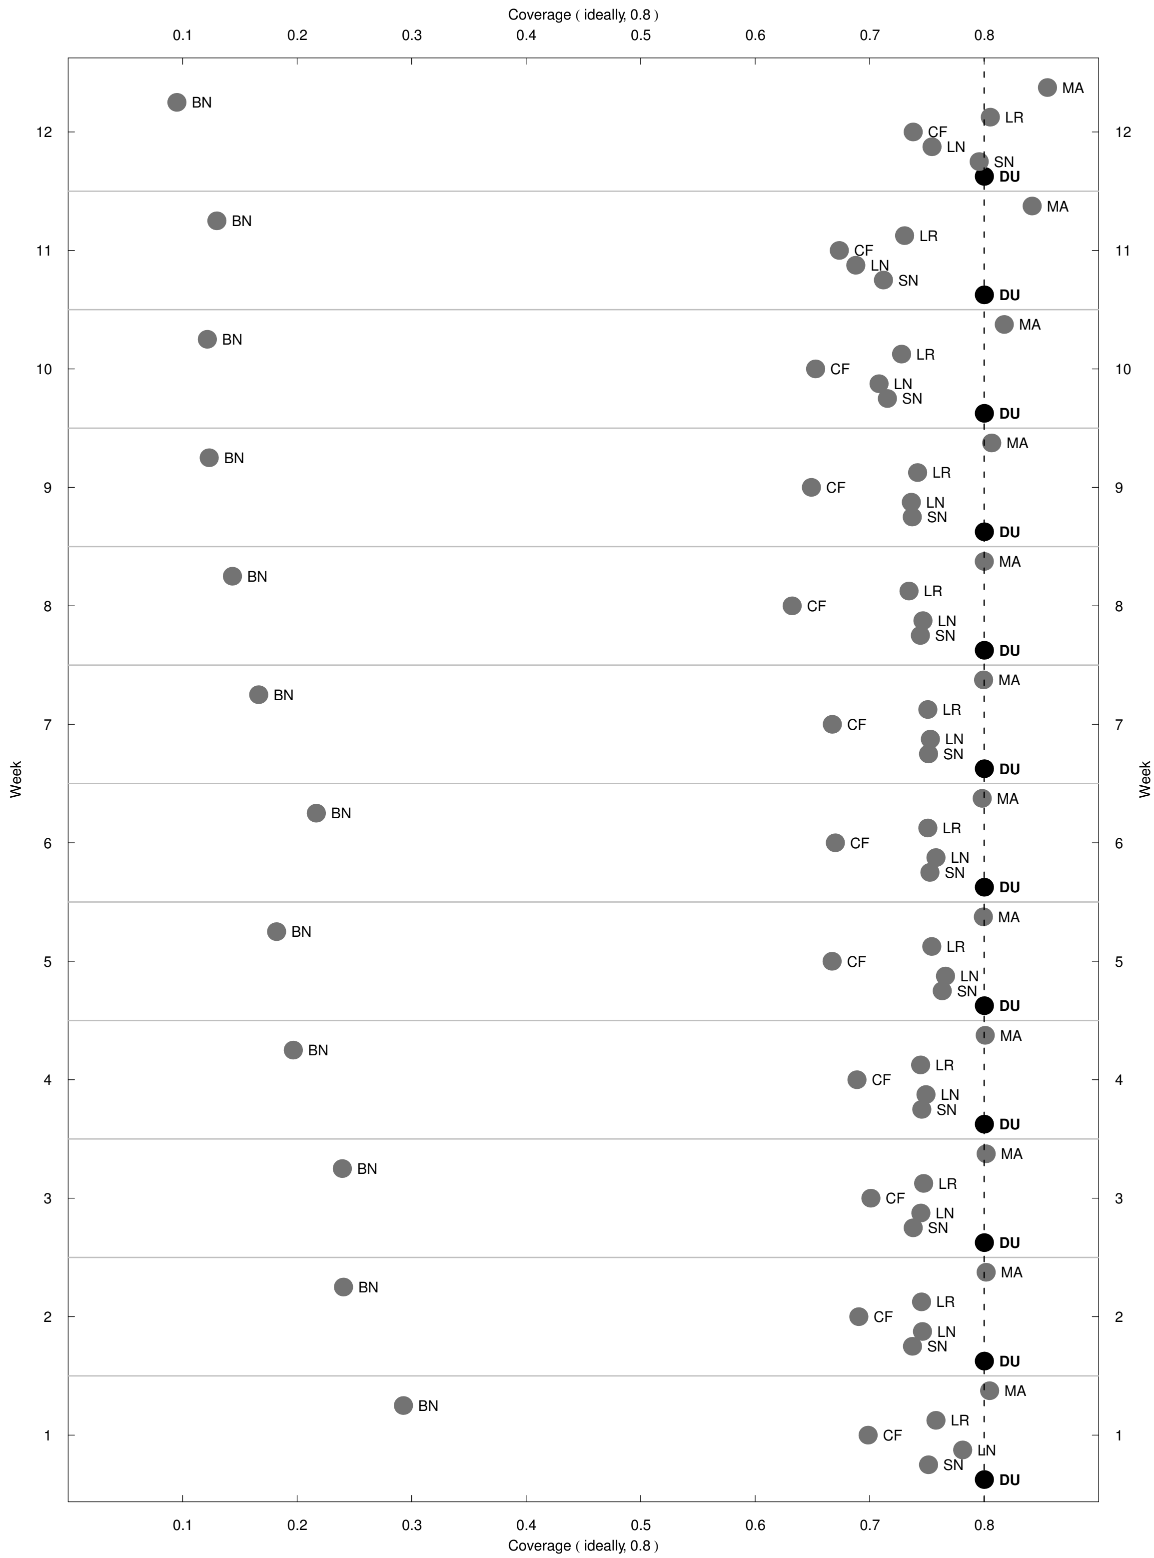

Figure 6: Forecast confidence accuracy

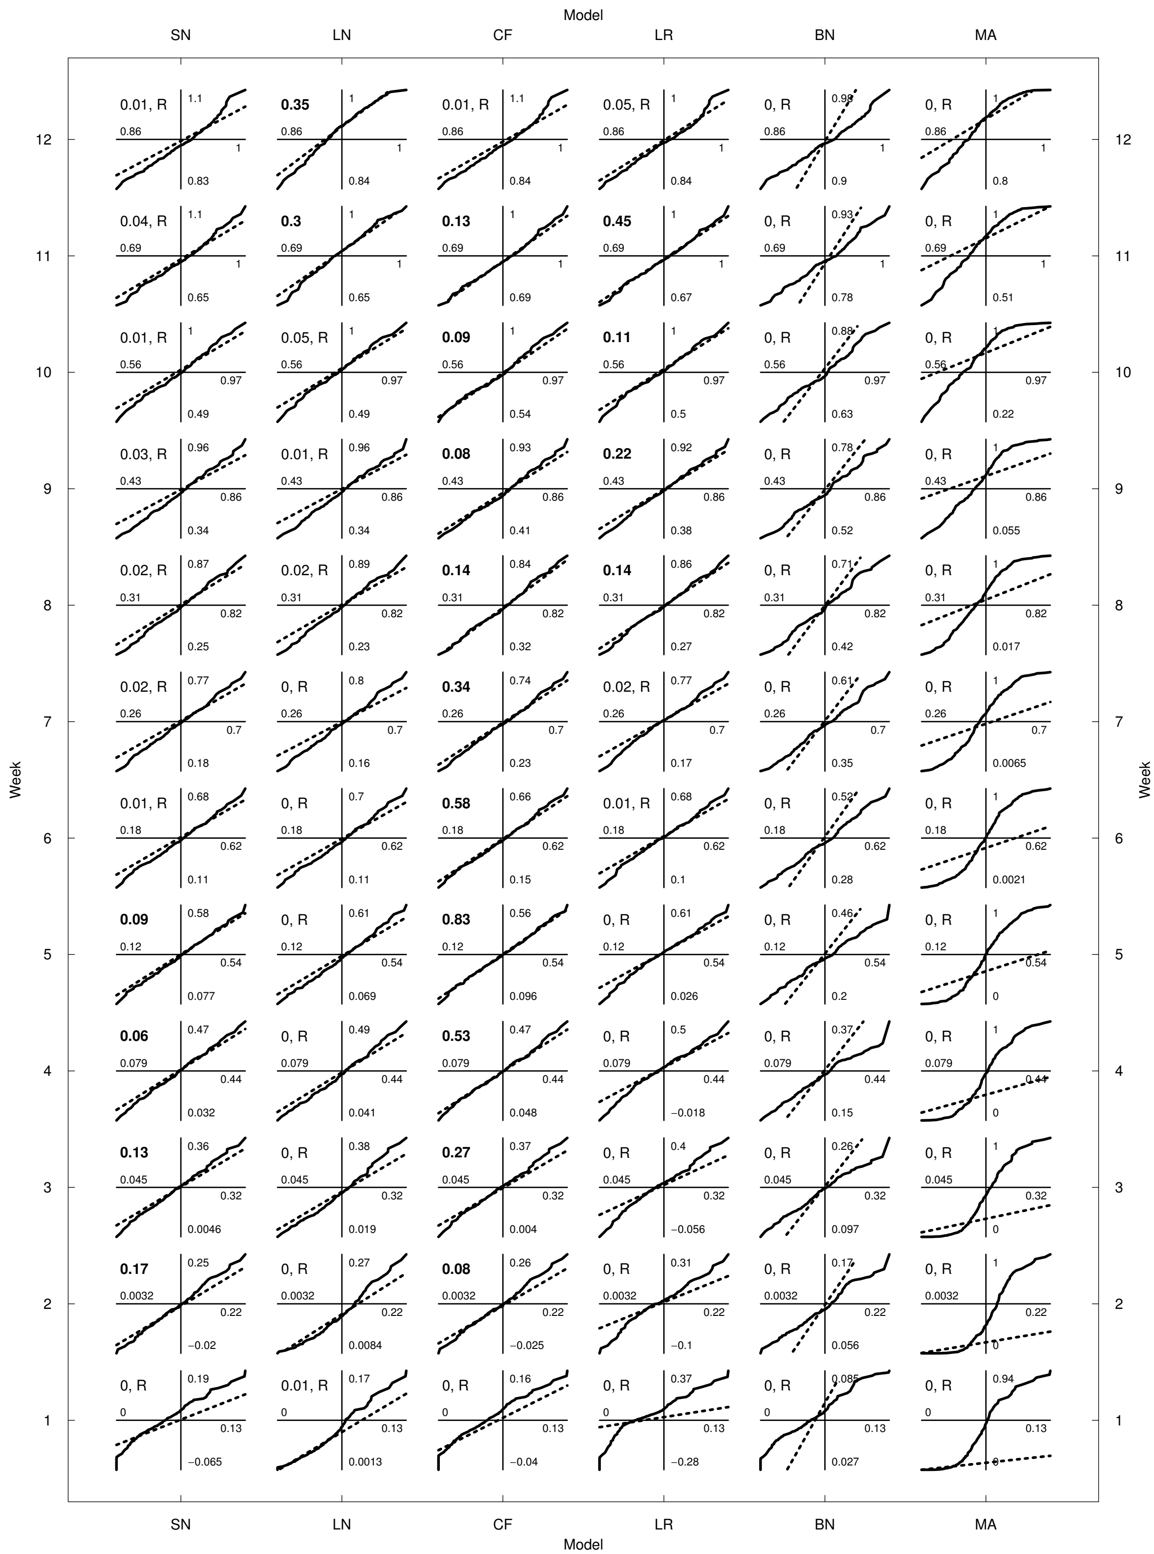

Figure 7: Quantile-quantile plots

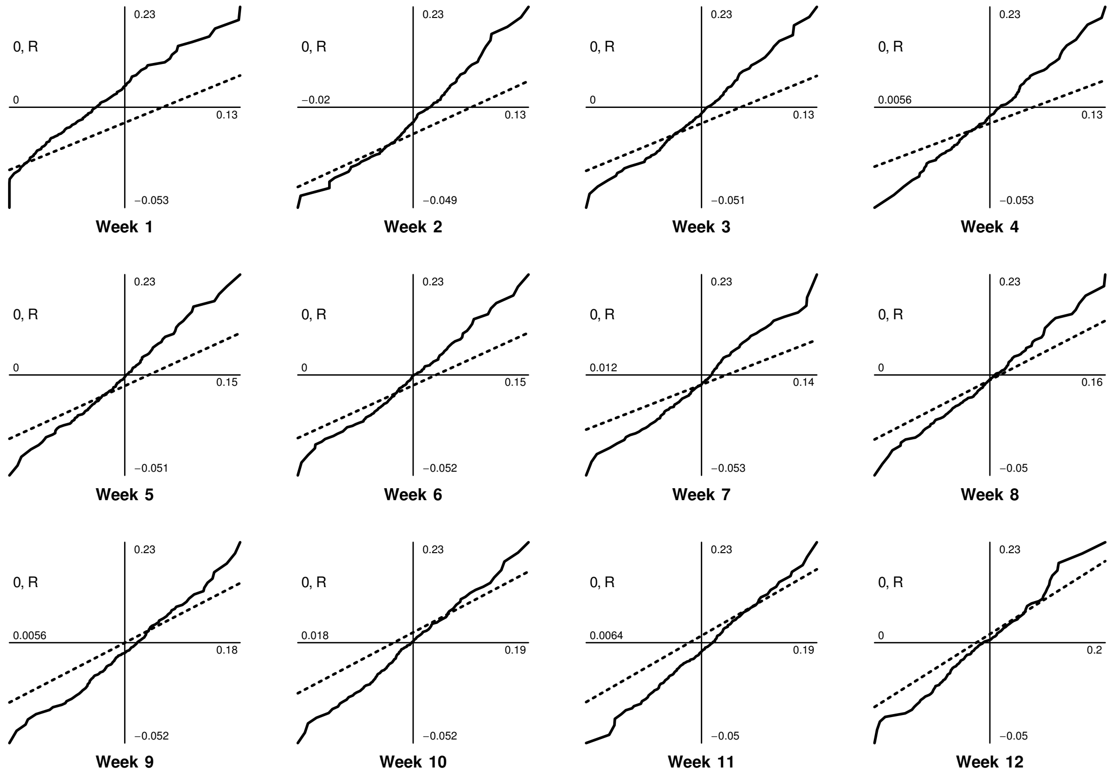

Figure 8: Model CF in-sample fit for incremental bookings