International Journal of Forecasting 19 (2003) 669–684
international journal of forecasting
www.elsevier.com/locate/ijforecast

Shrinkage estimators of time series seasonal factors and their effect on forecasting accuracy

Don M. Miller[a,*], Dan Williams[b]

[a]Virginia Commonwealth University, Box 844000, Richmond, VA 23284, USA
[b]Baruch College, 17 Lexington Ave., C-301, New York, NY 10010, USA

Abstract

This paper shows that forecasting accuracy can be improved through better estimation of seasonal factors under conditions for which relatively simple methods are preferred, such as relatively few historical data, noisy data, and/or a large number of series to be forecasted. In such situations, the preferred method of seasonal adjustment is often ratio-to-moving-averages (classical) decomposition. This paper proposes two shrinkage estimators to improve the accuracy of classical decomposition seasonal factors. In a simulation study, both of the proposed estimators provided consistently greater accuracy than classical decomposition, with the improvement sometimes being dramatic. The performances of the two estimators depended on characteristics of the series, and guidelines were developed for choosing one of them under a given set of conditions. For a set of monthly, M-competition series, greater forecasting accuracy was achieved when either of the proposed methods was used for seasonal adjustment rather than classical decomposition, and the greatest accuracy was achieved by following the guidelines for choosing a method.

Keywords: Seasonality; Time series; Shrinkage estimators; Empirical Bayes; Classical decomposition

1. Introduction

Many forecasting environments exhibit one or more of the following characteristics: few historical data, large amounts of random variation, and a large number of series to be forecasted. Important business examples include demand forecasting for manufacturing and inventory control, sales forecasting to support budget planning, and forecasting of service usage in new public programs. The presence of any one of these characteristics makes the use of simple methods desirable. They are easier to use, and research indicates that they are as accurate as more sophisticated methods (Makridakis & Hibon, 2000), especially when there is considerable random variation in the data (Makridakis et al., 1982).

A common approach to forecasting seasonal series with simple methods such as exponential smoothing is to deseasonalize the data, forecast the deseasonalized data, then reseasonalize the forecast. This approach was used for executing many of the relatively simple forecasting methods tested in the M-competition (Makridakis et al., 1982) and the M3-competition (Makridakis & Hibon, 2000). When simple forecasting methods are used, the preferred method for seasonal adjustment is often ratio-to-moving-averages (classical) decomposition

*Corresponding author. Tel.: +1-804-828-8410; fax: +1-804-828-8884.
E-mail addresses: dmmiller@vcu.edu (D.M. Miller), daniel_williams@baruch.cuny.edu (D. Williams).

0169-2070/02/$ – see front matter
doi:10.1016/S0169-2070(02)00077-8

(Macauley, 1930; also described in Makridakis, Wheelwright & Hyndman, 1998). In the M-competitions, for example, classical decomposition was used for seasonal adjustment in conjunction with simple and Holt exponential smoothing (Holt, 1957), damped-trend exponential smoothing, the robust-trend model (Meade, 2000), and the Theta model (Assimakopoulos & Nikolopoulos, 2000). Indeed, the use of classical decomposition for seasonal adjustment in the M-competition led to forecasts that were just as accurate as those that were produced using the Census II, X-11 model (Shiskin, Young & Musgrave, 1967).

The purpose of this paper is to examine the possibility of improving forecasting accuracy through better estimation of seasonal factors in situations where relatively simple methods are preferred. We propose two methods, each of which adjusts the classical decomposition seasonal factors toward 1.0, thereby shrinking the estimated seasonal variation. (Additive seasonal estimates are adjusted toward 0.) These methods are described in Section 2. Our evaluation consists of two stages. In Section 3, we present the results of a simulation in which the estimation accuracy of the proposed methods is compared to that of classical decomposition for a variety of time series under controlled conditions. We find that both shrinkage methods are generally more accurate than classical decomposition and their performances depend on characteristics of the time series. We develop guidelines on which to use for a given series. In Section 4, we examine the impact on forecasting accuracy when the proposed methods are used for seasonal adjustment of a set of real time series taken from the M-competition. In Section 5, we use two real data series to illustrate the effect of seasonal shrinkage. We summarize our findings and offer some concluding remarks in Section 6.

2. Methods for estimating seasonal factors

The classical decomposition model can be expressed in either additive or multiplicative form. The multiplicative model is $Y_t = T_t^* S_t^* E_t$, where T_t is the trend-cycle component, S_t is the seasonal component, and E_t is the irregular (random) component. The trend-cycle estimate T_t is a $2 \times J$ centered moving average at period t, where J = number of periods in one year. The ratio Y_t / T_t provides a preliminary seasonal factor at t. (It also is influenced by the random component, since $Y_t / T_t = S_t^* E_t$). For a given season, there will be several such preliminary seasonal factors. For example, with 4 years of monthly data (48 months), there are 36 preliminary seasonal factor estimates, because the centered moving average produces only 36 trend-cycle estimates. Consequently, there are three preliminary seasonal factor estimates for each month. The seasonal factor estimate for each season is the average of the set of preliminary seasonal factors for that season. Thus, the seasonal factor estimates are the means of J separate samples, each sample consisting of the preliminary seasonal factors for a given season. Finally, the factors are adjusted to force them to average 1.0.

Armstrong (1985, p. 163) asserts that seasonal factors increase error if there is a great deal of uncertainty. He proposes using *shrinkage*, that is, damping the estimated seasonal factors toward 1.0 (the mean of the J seasonal estimates), thereby shrinking the variance among them, as follows:

$$S_j^A = W^A + (1 - W^A) S_j, \quad j = 1, 2, \ldots, J \tag{1}$$

where S_j is the classical decomposition estimator of the seasonal factor for period j, and J = the number of seasons in a year. The weight W^A, which determines the degree of shrinkage, depends on the amount of uncertainty about seasonality. Armstrong suggests using $W^A = 1/(d^{1/2})$, where d = the number of years of data.

The methods we propose also apply shrinkage to existing estimates. Stein (1955) introduced the idea of shrinkage estimators, showing that it is possible to uniformly improve upon the usual maximum likelihood estimator (the vector of independent sample means) when estimating the mean of a multivariate normal distribution. Later, James and Stein (1961) presented a simple shrinkage estimator for which the improvement is substantial in certain situations. Shrinkage estimators may be viewed as forms of empirical Bayes estimation, an approach introduced by Robbins (1955) and further developed by many authors, notably Rutherford and Krutchkoff (1969), Efron and Morris (1973, 1975), and Morris (1983). Applications of empirical Bayes shrinkage methods to forecasting problems have been developed by Greis and Gilstein (1991) and Bunn and Vassilopoulos (1999).

Empirical Bayes methods derive from Bayesian statistical theory (Robbins, 1955). The set of parameters to be estimated are viewed as realizations of a prior distribution. Unlike Bayesian estimators, however, the prior distribution is estimated from the data rather than subjectively. Empirical Bayes methods are especially effective when the parameter values are similar, the information on each is weak (large within-sample variances, small sample sizes), and the number of populations is large. In this application, the multiple parameters are the true seasonal factors, and the information on each parameter consists of the classical decomposition estimate for that seasonal factor.

The estimators proposed here are based on assumptions of independence and equal variances among the classical decomposition seasonal factors. These are approximating assumptions. The estimated seasonal factors are not independent because (1) autocorrelation exists among trend-cycle values and possibly in the random component; and (2) forcing the seasonal estimates to average 1.0 (or to sum to 0.0 for additive factors) means there can be no more than $J-1$ independent factors. The variances of the estimated seasonal factors are likely to differ somewhat because (1) they depend on the underlying seasonal factors, which differ; and (2) there can be different numbers of preliminary seasonal factors for the J seasonal factor estimates. (This will be the case unless we have d complete years of data.) Nevertheless, the proposed shrinkage estimators have intuitive appeal and, as reported in Sections 3 and 4, they produce good results. In this regard, Morris (1983) states that the benefits of shrinkage estimators are mainly due to the richer model and less critically to how one estimates the parameters.

The first of the estimators we propose is an adaptation of the James–Stein estimator. When viewed from an empirical Bayes perspective, the prior distribution is assumed to be $\theta_j \sim Normal\ (1, A)$, $j=1,\ldots,J$, where θ_j is the underlying seasonal factor for season j (e.g. see Efron & Morris, 1973). A represents the variation among the underlying (true) seasonal factors, which must be estimated from the data. The James–Stein estimator is:

$$S_j^{\mathrm{J-S}} = W^{\mathrm{J-S}} + (1 - W^{\mathrm{J-S}})\, S_j, \quad j = 1, 2, \ldots, J \tag{2}$$

where S_j and J are as defined previously. Although the form of the James–Stein estimator is the same as that of Armstrong's estimator, they differ in the determination of the shrinkage parameters W^{A} and $W^{\mathrm{J-S}}$. Following Morris (1983), the James–Stein shrinkage parameter is:

$$W^{\mathrm{J-S}} = \left(\frac{J-3}{J-1}\right)\left(\frac{V}{V+A}\right) \tag{3}$$

where V = the variance (due to sampling error) of each estimated seasonal factor S_j, A = the variance among the true seasonal factors, and J = the number of seasons (i.e. $J=12$ for monthly series).

The variance among the J classical decomposition seasonal estimators is approximately $V+A$, which exceeds the variance among the true seasonal factors by V. (Following Morris, 1983, this result would be exact if the seasonal estimators were independent and unbiased.) Fig. 1 indicates that the approximation is good for the simulation results in Section 3. Thus, the degree of shrinkage is determined primarily by the relationship between V and A. If the sampling error variance for individual seasonal estimators (V) is large compared to the variation among the true seasonal factors (A), then the value of

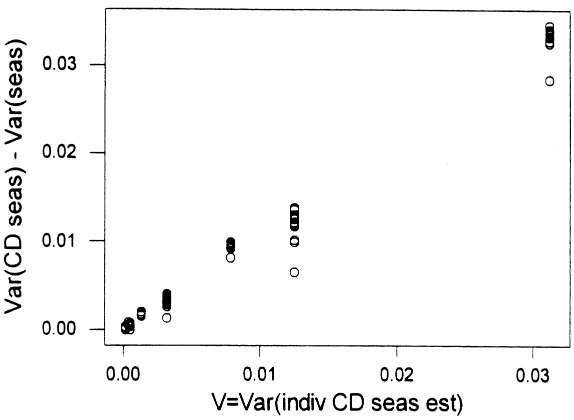

Fig. 1. Difference in variation among seasonal factors (classical decomposition − actual) versus variance of individual seasonal factor estimates (classical decomposition).

W^{J-S} is near 1, resulting in a large amount of shrinkage. But if V is small compared to A, the value of W^{J-S} is near 0, resulting in little shrinkage. The factor $(J-3)/(J-1)$ is a small-sample correction. For monthly series, for example, this factor is $(J-3)/(J-1)=9/11=0.818$.

V and A are unknown but may be estimated from the data as follows. To estimate the variance of S_j, $i=1, 2, \ldots, J$, which we assume to be the same for all seasons, we use the average of the estimated variances for $V_1, V_2, \ldots, V_J$:

$$\hat{V} = \frac{\sum_{j=1}^{J} \hat{V}j}{J} \tag{4}$$

where

$$\hat{V}_j = \frac{\sum_{k=1}^{K_j} (S_{jk} - S_j)^2}{K_j(K_j - 1)} \tag{5}$$

$S_{j,k}$ represents the preliminary seasonal factor estimate for season j in year k, and K_j = the number of preliminary seasonal factors for season j.

A is estimated by

$$\hat{A} = \frac{\sum_{j=1}^{J} (S_j - 1.0)^2}{J-1} - \hat{V} \tag{6}$$

$\hat{A}$ is set to zero if (6) becomes negative.

Since the James–Stein model is based on an assumed normal prior distribution, it seems reasonable to expect the James–Stein estimator to perform best when the empirical distribution of the J estimated seasonal factors is approximately symmetrical and single-peaked, that is, similar to a normal distribution.

The second shrinkage estimator that we propose was developed by Lemon and Krutchkoff (1969). The Lemon–Krutchkoff (L–K) empirical Bayes estimator estimator is nonparametric in regard to the prior distribution. For the application at hand, this means that no assumption is made regarding the prior distribution of the seasonal factors. This is appealing since there is no particular reason to expect, in general, any particular form of distribution among the underlying seasonal factors. We may anticipate that the L–K estimator is likely to perform best relative to the James–Stein estimator when the empirical distribution of the estimated seasonal factors is decidedly non-normal, e.g. asymmetric.

The L–K estimator for the seasonal factor for season $j*$ is:

$$S_{j^*}^{L-K} = \sum_{j=1}^{J} W_{j^*,j} S_j, \quad j = 1, 2, \ldots, J \tag{7}$$

That is, the L–K estimator for a given season is a weighted average of the J classical decomposition estimators. A different set of weights is used in estimating each seasonal factor. The weights indicate the relative likelihoods of observing the estimate S_{j*} if the true seasonal factor is S_j, that is:

$$W_{j^*,j} = \frac{L_{j^*,j}}{\sum_{j=1}^{J} L_{j^*,j}} \tag{8}$$

where $L_{j*,j}$ is the likelihood associated with S_{j*}, given the true seasonal factor is S_j. Since each classical decomposition seasonal estimator is a mean of the preliminary estimates for that season, it seems reasonable to assume that its sampling distribution (in Bayesian terms the conditional distribution) is approximately the normal distribution. Thus:

$$L_{ij} = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{S_i - S_j}{\sigma}\right)^2 \right] \tag{9}$$

where σ is estimated by $\hat{V}^{1/2}$ and $\hat{V}$ was defined previously.

As an illustration, suppose we have quarterly data for which the classical decomposition seasonality estimates are $S_1=1.12$, $S_2=1.15$, $S_3=1.13$, and $S_4=0.60$. In developing the L–K estimator for the first quarter seasonal factor S_1^{L-K}, the likelihoods are the following:

- L_{11} is the statistical likelihood of observing a classical estimate of $S_1=1.12$ if the true 1st quarter seasonality factor is 1.12 (the value of S_1)
- L_{12} is the statistical likelihood of observing a classical estimate of $S_1=1.12$ if the true 1st quarter seasonality factor is 1.15 (the value of S_2)
- L_{13} is the statistical likelihood of observing a

classical estimate of $S_1 = 1.12$ if the true 1st quarter seasonality factor is 1.13 (the value of S_3)
- L_{14} is the statistical likelihood of observing a classical estimate of $S_1 = 1.12$ if the true 1st quarter seasonality factor is 0.6 (the value of S_4)

In determining S_1^{L-K}, L_{11} will always be the greatest of these likelihoods. The magnitudes of L_{12}, L_{13}, and L_{14} depend on the distances of S_2, S_3, and S_4 from S_1, respectively. Therefore, L_{12} and L_{13} would be the next largest likelihoods and nearly as large as L_{11}, since S_2 and S_3 are nearly equal to S_1. L_{14}, however, would be much smaller—perhaps zero—because S_4 differs greatly from S_1. For this example, the L–K estimator S_1^{L-K} would effectively pool S_1, S_2, and S_3, while disregarding S_4.

More generally, seasonal factors that are similar to the seasonal factor being estimated—say, S_{j*}—receive relatively large weights, and seasonal factors that differ from S_{j*} beyond a certain amount receive virtually no weight. When the statistical distribution of seasonal factors is asymmetric, this is a particularly appealing characteristic. Suppose, for example, that sales outcomes are the same on average for 11 months but increase sharply each December for the holiday season. In this case, the estimates for the first 11 months would essentially be pooled, while the estimate for December would be virtually unchanged. For another example, suppose the estimated seasonal factors consisted of two sets of similar factors—say, around 0.9 for the first 6 months and about 1.1 for the last 6 months. Then the L–K estimators would effectively pool each set separately.

3. Examination of proposed methods: a simulation study

In order to discover the conditions under which the proposed methods are more accurate than classical decomposition, if any, we performed a simulation. This allowed us to control the underlying conditions and to compare seasonal estimates to known seasonal factors. Monthly time series data were generated from a multiplicative model $X_t = T_t^* S_t^* E_t$, where T_t is a trend component, S_t is a seasonal component, and E_t a random component. The controlled conditions, shown in Table 1, were chosen to be representative of the monthly series in the M-competition (Makridakis et al., 1982). They are (1) the number of years of monthly data: either 3 or 6, (2) the trend T: either none or a constant 1 unit change in level from the previous period, (3) the actual seasonal factors (13 sets, discussed below), and (4) the variation of the random component (four levels, discussed below). The initial level of each series was arbitrarily set to 100. The seasonal factors were chosen to reflect two features: the variation among seasonal factors (four levels) and the asymmetry (if any) of the distribution of seasonal factors (four levels). Asymmetry is measured by the coefficient of skewness and serves as a simple proxy for non-normality in the distribution of seasonal factors, a condition for which the L–K estimator is expected to outperform the J–S estimator. The random component was generated from a lognormal distribution by exponentiating values generated randomly from a normal distribution using the Data Analysis Tool Pak feature of Microsoft Excel. This produced a mild, positive skewness in the distribution of random factors that is typical for multiplicative models. (Note that taking the logarithm of the multiplicative model produces an additive model with a normally distributed random term.)

Table 1
Simulation design factor levels

No. years	SD (E)	SD (S)	Skewness (S)	Trend
3	0.025	0.00	0	0
6	0.05	0.05	0.6154	1% per month
	0.125	0.15	1.4035	
	0.25	0.35	2.8868	

A complete $2 \times 4 \times 4 \times 4 \times 2$ factorial design requires 256 combinations. However, when the variation among seasonal factors is zero, there can be no skewness in their distribution. Consequently, only 13 combinations of levels for these two factors were required, rather than 16. Thus, only 208 treatment combinations were required, rather than 256. For each treatment combination, we generated 500 time series randomly and independently. The resulting coefficients of variation were approximately 2.0–2.5% for mean square error estimates and 1.0–1.25% for MAPE estimates. The simulated factor levels are shown in Table 1.

Table 2
Seasonal factors used in the simulation

Period	1	2	3	4	5	6	7	8	9	10	11	12	13
1	1.000	0.986	0.973	0.958	0.965	0.957	0.919	0.874	0.896	0.899	0.812	0.707	0.757
2	1.000	0.986	0.946	0.975	0.965	0.957	0.839	0.925	0.896	0.899	0.623	0.824	0.757
3	1.000	0.986	0.919	1.067	0.977	0.957	0.758	1.201	0.931	0.899	0.435	1.469	0.838
4	1.000	0.986	0.946	0.958	0.977	0.957	0.839	0.874	0.931	0.899	0.623	0.707	0.838
5	1.000	0.986	0.973	0.975	0.977	0.957	0.919	0.925	0.931	0.899	0.812	0.824	0.838
6	1.000	0.986	1	1.067	0.977	0.957	1	1.201	0.931	0.899	1	1.469	0.838
7	1.000	0.986	1.027	0.958	0.977	0.957	1.081	0.874	0.931	0.899	1.188	0.707	0.838
8	1.000	0.986	1.054	0.975	0.977	0.957	1.161	0.925	0.931	0.899	1.377	0.824	0.838
9	1.000	0.986	1.081	1.067	0.977	0.957	1.242	1.201	0.931	0.899	1.565	1.469	0.838
10	1.000	0.986	1.054	0.958	1.029	0.957	1.161	0.874	1.087	0.899	1.377	0.707	1.203
11	1.000	0.986	1.027	0.975	1.101	0.957	1.081	0.925	1.304	0.899	1.188	0.824	1.709
12	1.000	1.159	1	1.067	1.101	1.476	1	1.201	1.304	2.111	1	1.469	1.709
Total	12	12	12	12	12	12	12	12	12	12	12	12	12
SD	0	0.05	0.05	0.05	0.05	0.15	0.15	0.15	0.15	0.35	0.35	0.35	0.35
Skew	0	2.8868	0	0.6154	1.4035	2.8868	0	0.6154	1.4035	2.8868	0	0.6154	1.4035

The seasonal factors that were used in the simulation are provided in Table 2, along with their corresponding standard deviation and skewness values. The specific seasonal factor values in each set were chosen arbitrarily in order to (1) produce the desired standard deviations and coefficients of skewness and (2) reflect realistic seasonal patterns based on the experiences of the authors in forecasting a great variety of real data series. These seasonal patterns include (1) a seasonal increase during the same month each year, (2) a seasonal increase during the same month of each quarter of the year, (3) increasing values throughout the year (lowest early in the year, highest at end of year), and (4) a sinusoidal, wavelike pattern.

To insure that the conditions of the simulation were realistic, the factor levels were chosen to approximate the range of conditions for the monthly series in the M-competition. To this end, we selected all 68 monthly series in the subset of 111 series used in that study. The distributions of the estimated design factors over these series are summarized in Table 3, where SD(E), SD(S), Skew(S), and Trend were estimated for each series by applying classical decomposition.

3.1. Simulation results

Performance comparisons were based on (1) the ratio of mean squared errors (MSE) in estimating the 12 seasonal factors, that is, the MSE for a proposed method divided by the MSE for classical decomposition and (2) the corresponding ratio of mean absolute percentage errors (MAPE). MSE and MAPE are defined as follows:

$$\text{MSE} = \frac{\sum_{j=1}^{12} (\hat{S}_j - S_j)^2}{12} \tag{10}$$

Table 3
Distribution of estimated design factors for 68 monthly series from M-competition

	No. years	SD(E)	SD(S)	Skew(S)	Trend (% per month)
Minimum	2.5	0.0027	0.0008	0.0108	0.01
1st Quartile	3.5	0.0299	0.0460	0.3130	0.19
2nd Quartile	5.4	0.0568	0.1159	0.6155	0.45
3rd Quartile	7.1	0.0824	0.1867	1.4112	1.07
Maximum	10.5	0.2610	0.4873	2.5914	67.84

Table 4
Summary of results over all simulated conditions

	N	Mean	Minimum	Q1	Median	Q3	Maximum
Seasonal series							
MSE: A/CD	192	17.523	0.202	0.707	2.404	18.526	150.881
MSE: J–S/CD	192	0.826	0.219	0.775	0.951	0.992	1.001
MSE: L–K/CD	192	0.677	0.407	0.486	0.660	0.784	1.214
MSE: L–K/J–S	192	0.962	0.422	0.669	0.815	1.190	2.395
MAPE: A/CD	192	2.717	0.444	0.787	1.427	3.097	13.786
MAPE: J–S	192	0.881	0.417	0.850	0.962	0.993	1.000
MAPE: L–K	192	0.776	0.587	0.634	0.782	0.853	1.110
MAPE: L–K/J–S	192	0.923	0.587	0.800	0.875	1.078	1.549
Non-seasonal series							
MSE: A/CD	16	0.264	0.179	0.179	0.264	0.350	0.350
MSE: J–S/CD	16	0.144	0.125	0.126	0.143	0.161	0.163
MSE: L–K/CD	16	0.345	0.330	0.330	0.344	0.357	0.367
MSE: L–K/J–S	16	4.231	3.965	3.978	4.266	4.444	4.503
MAPE: A/CD	16	0.507	0.423	0.423	0.507	0.591	0.591
MAPE: J–S/CD	16	0.332	0.313	0.314	0.332	0.350	0.352
MAPE: L–K/CD	16	0.551	0.543	0.543	0.551	0.558	0.563
MAPE: L–K/J–S	16	1.915	1.847	1.849	1.920	1.976	1.987

$$\text{MAPE} = 100\,\frac{\sum_{j=1}^{12}\left|\dfrac{\hat{S}_j - S_j}{S_j}\right|}{12} \tag{11}$$

Table 4 provides a summary of results for the MSE ratios and the MAPE ratios over all conditions simulated. Results are shown for a multiplicative model. We also performed the simulation using an additive model and got similar results (results not shown). The results for the non-seasonal series are shown separately, since it is obvious that any method that shrinks the classical decomposition seasonals toward 1.0 has to do better in this case. With either model, the MSE ratios and MAPE ratios were quite similar to each other, with the caveat that the MSE ratios are somewhat more extreme in both directions because they deal with squared quantities. Trend was never a factor; results for series with trend were almost indistinguishable from those for series with no trend. Since the patterns of results are approximately the same for either MSE ratios or MAPE ratios, we discuss results only for MSE ratios for the multiplicative model in the remainder of the paper.

3.2. Overview of results

Although Armstrong's ad hoc method sometimes provided more accurate seasonal estimates than classical decomposition, its overall performance is disappointing. For 25 percent of the simulated conditions, the Armstrong MSE is less than half that of classical decomposition (first quartile ratio = 0.48), but for half the conditions, its MSE is more than 80 percent greater (median ratio = 1.82). Its overall performance is far inferior to that of the empirical Bayes methods, probably because it accounts for only one of the conditions believed to affect the expansion of variance among CD seasonal factors (number of years).

The empirical Bayes methods are almost always more accurate that classical decomposition, and often by a substantial amount. For the nonseasonal series, the mean MSE ratios for the James–Stein (J–S) method ranged from 0.125 to 0.163, while the MSE ratios for the Lemon–Krutchkoff (L–K) method ranged from 0.330 to 0.367. Clearly, the J–S estimator is superior when the data are nonseasonal. Indeed the mean MSE for the J–S estimator was less than that of the L–K estimator for all 16 of the nonseasonal conditions simulated. For seasonal series, there was greater variation in results. The mean MSE ratios for the James–Stein (J–S) and Lemon–Krutchkoff (L–K) methods were 0.83 and 0.68, respectively. The corresponding medians were

0.95 and 0.66, respectively. The J–S estimator is never less accurate that classical decomposition (CD) under any conditions tested (3rd quartile ratio = 0.99; maximum ratio = 1.00). The L–K estimator is substantially more accurate that classical decomposition under most conditions (3rd quartile ratio = 0.78) but can be less accurate (maximum value = 1.21). When the empirical Bayes methods are good, they are very good. For 25% of the conditions simulated, J–S provided MSE ratios of less than 0.56, and L–K provided ratios of less than 0.45. In comparing L–K to J–S, the median MSE ratio was 0.815 (favoring L–K). The mean MSE for L–K was less than that for J–S for 116 of the 192 conditions simulated (60.4%).

Since overall results for the empirical Bayes methods were outstanding, and those for the Armstrong method were disappointing, we focus exclusively on the former in the remainder of the analysis.

3.3. Classical decomposition

We now consider the effects of the simulation factors. First, under what conditions, if any, does classical decomposition inflate the variation among seasonal factors? Recall that, in theory, the variance among the CD factors exceeds the variance among the true factors by approximately V, the variance of the CD seasonal estimator for a given season. Fig. 1 is a plot of the difference between the variances among CD seasonal estimates and among the true seasonal factors versus $\hat{V}$. All differences are positive (or zero), indicating that the variance among CD seasonal factors is virtually always inflated to some degree. The result supports theory: The degree of variance inflation is mostly explained by the variance of the individual CD seasonal factor estimates. (Regression slope approximately = 1.0; correlation coefficient $\rho = 0.995$. We note that Fig. 1 represents 208 conditions, and some of the symbols represent multiple points.) Therefore, the inflation of the variance among CD seasonal estimates increases as random variation increases and as the number of observations decreases, since these factors largely determine the variance of the individual CD seasonal factor estimates. We note that these are the very conditions for which classical decomposition is usually the method chosen for seasonal adjustment prior to forecasting.

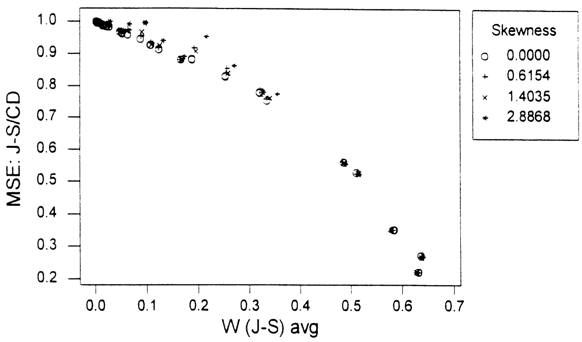

Fig. 2. Relative accuracy of the James–Stein method vs. classical decomposition: Ratio of MSEs vs. W^{J-S}.

3.4. James–Stein vs. classical decomposition

Fig. 2 is a plot of the ratios of the mean squared errors of J–S and CD for different values of W^{J-S} (nonseasonal series excluded). Smaller ratios indicate greater accuracy for J–S. The plotting symbols indicate the skewness of the distribution of seasonal factors in each case. J–S is never less accurate than classical decomposition and, as W^{J-S} increases, its relative advantage increases. The improvement in MSE is greater than 10% if W^{J-S} exceeds 0.2. For values of W^{J-S} above 0.5, the relative improvement with J–S becomes astonishing. [Values of W^{J-S} exceed 0.5 when the variance V of the individual CD seasonal factor is substantially greater than the variance A among the seasonal factors; see Eq. (3)]. The effect of skewness among seasonal factors is almost negligible.

3.5. Lemon–Krutchkoff vs. classical decomposition

Fig. 3 is a plot of the ratios of the mean squared errors of L–K and CD for different values of W^{J-S} (nonseasonal series excluded). Smaller ratios indicate greater accuracy for L–K. The plotting symbols indicate the skewness of the distribution of seasonal factors in each case. L–K is usually more accurate than CD. Like J–S, it provides striking reductions in MSE for values of W^{J-S} above 0.5. However, when the distribution of seasonal factors is symmetric and W^{J-S} is less than about 0.2, the MSE for L–K can exceed that of CD by up to 20%. On the other hand, when the distribution of seasonal factors is skewed, L–K provides substantial improvement (20–60%

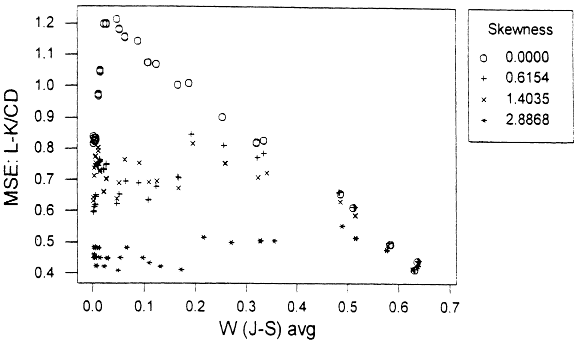

Fig. 3. Relative accuracy of the Lemon–Krutchkoff method vs. classical decomposition: Ratio of MSEs vs. W^{J-S}.

reductions in MSE) even for very small values of W^{J-S}. The amount of MSE reduction increases as seasonal skewness increases.

3.6. Lemon–Krutchkoff vs. James–Stein

We have seen that each of the empirical Bayes methods provides greater accuracy than CD under all or most circumstances. How do they compare to each other? Fig. 4 is a plot of the ratio of the mean squared error of L–K to that of J–S for different values of W^{J-S} (nonseasonal series excluded). The figure shows results only for $W^{J-S} < 0.6$. If $W^{J-S} > 0.6$, the MSE ratio dramatically favors J–S. (W^{J-S} was less than 0.6 for 85% of the M-competition series.) For $W^{J-S} < 0.6$, if the distribution of seasonal factors is symmetric (skewness = 0), then the MSE of the L–K estimator is generally about 20% greater than that of the J–S estimator. This, apparently, is the price paid for using a nonparametric estimator

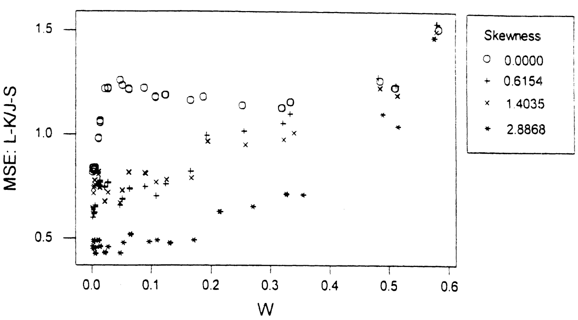

Fig. 4. Relative accuracy of the Lemon–Krutchkoff method vs. the James–Stein method: Ratio of MSEs vs. W^{J-S} (for $W^{J-S} <$ 0.6).

rather than a parametric estimator whose distributional assumptions are essentially correct. If the distribution of seasonal factors is skewed, the relative performances depend on W^{J-S} and the degree of skewness. The ratio of mean squares favors L–K as W^{J-S} decreases and skewness increases and it favors J–S as W^{J-S} increases and skewness decreases.

3.7. Recommendations suggested by the simulation

Based upon the relative performances in the simulation of CD and the two empirical Bayes methods, we suggest the following guidelines for choosing a method of estimating seasonal factors:

1. *Choose the James–Stein method if* (1) W^{J-S} *exceeds* 0.5 *or* (2) $W^{J-S} > 0.2$ *and the distribution of seasonal factors is approximately symmetric*. (The analysis is not extensive enough to specify a threshold value of skewness with confidence. For the analysis of real data in Section 4, we used 'coefficient of skewness < 0.5' as the operational definition of a symmetric distribution.) The J–S method's MSE was uniformly as small or smaller than that of CD for all values of W^{J-S}. It was smaller than the Lemon–Krutchkoff MSE for $W^{J-S} > 0.2$ and skewness = 0.0. Additionally, it was as small or smaller than the Lemon–Krutchkoff MSE for $W^{J-S} > 0.5$, regardless of the skewness of the distribution of seasonal factors.
2. *Choose the Lemon–Krutchkoff method if the distribution of seasonal factors is skewed and* W^{J-S} *is less than* 0.5. The L–K method's MSE was uniformly as small or smaller than that of CD if the distribution of season factors is skewed. It was as small or smaller than that of J–S if the distribution of season factors is skewed and $W^{J-S} < 0.3$. Between $W^{J-S} = 0.3$ and $W^{J-S} = 0.5$, the MSEs are relatively close, depending on the degree of skewness.
3. *Choose either classical decomposition or the James–Stein method if the distribution of season factors is symmetric and* W^{J-S} *is less than* 0.2. Since J–S's MSE was usually smaller by less than 10% in this case, using CD can be justified based upon its greater simplicity. (But note that W^{J-S} must be determined anyway.)

These guidelines are summarized in Table 5.

Table 5
Guidelines for choosing a method

W^{J-S}	Recommended method	
	Symmetric seasonals	Skewed seasonals
<0.2	Either classical decomposition or James–Stein	Lemon–Krutchkoff
0.2 to 0.5	James–Stein	Lemon–Krutchkoff
>0.5	James–Stein	James–Stein

3.8. Examples of the effects of the James–Stein and Lemon–Krutchkoff estimators

Fig. 5 illustrates the typical effects of the James–Stein and Lemon–Krutchkoff methods in adjusting classical decomposition seasonal estimates. Fig. 5a shows the results of a simulation for conditions that are advantageous to the J–S method. For three years of monthly data, the symmetrically distributed seasonal factors are 0.85 each month for January through March; 0.95 for April–June, 1.05 for July–September, and 1.15 for October–December. The standard deviation of the random component was 0.125; the standard deviation among these seasonal factors is 0.117; and there is no trend. When this pattern was replicated 500 times, the average value of W^{J-S} was 0.545, indicating substantial shrinkage. The average mean squared errors were 0.0173 for J–S and 0.0194 for L–K vs. 0.0329 for classical decomposition. Seasonals are plotted only for January, April, July, and October, to simplify the figure. The connecting lines track the relationship of each true seasonal factor to the corresponding CD and J–S estimates. We observe that (1) classical decomposition exaggerated the variation among seasonal factors; (2) the J–S method reduced the variation so that the estimates more closely resemble the original seasonals; and (3) the L–K estimates are more accurate than the CD estimates but less accurate than those of J–S.

Fig. 5b illustrates the effect of the L–K method for a simulated series in which it is expected to provide more accurate estimates. For three years of monthly data, the seasonal factors are 0.78 for January and 1.02 for all other months. The coefficient of skewness for these factors is −2.88. The standard deviation of the random component was 0.125; the standard deviation among these seasonal factors is 0.069; and there is no trend. When this pattern was replicated 500 times, the average value of W^{J-S} was 0.196, indicating modest J–S shrinkage. The average mean squared errors were 0.00069 for L–K vs. 0.00138 for J–S and 0.00148 for classical decomposition. The seasonal factors are plotted only for January, April, July, and October, for simplicity. Classical decomposition exaggerated the variation among seasonal factors. The J–S method shrinks the estimates for April, July, and October (the months with similar factors) nicely toward the correct values, but it over-adjusts the estimate for January (the

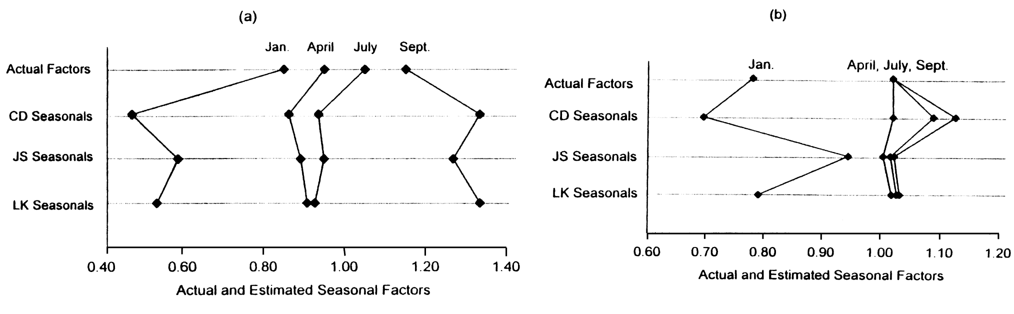

Fig. 5. Two examples of the effects of James–Stein and Lemon–Krutchkoff methods in adjusting classical decomposition estimators.

extreme factor). The L–K method applies appropriate shrinkage for the months with similar values but adjusts the extreme factor more appropriately.

4. Impact on forecasting accuracy

The simulation described in Section 3 clearly suggests that the empirical Bayes methods provide more accurate seasonal estimates than their classical counterpart. Does this apparent superiority lead to more accurate forecasts? We took the following general approach. The analysis is based on the 68 monthly time series that are included in the subset of 111 series used in the M-competition. For each series, we developed three sets of seasonal factors by applying classical decomposition and the two empirical Bayes methods. Each set of seasonal factors was used to deseasonalize the original data. We developed a forecast based on each of the three sets of deseasonalized data, then reseasonalized each forecast. Then we assessed the accuracy of each forecast.

More specifically, for each series we (1) withheld 18 months of data (the same data that were withheld in the M-competition); (2) used the remaining data to estimate the seasonal factors for each series using CD, J–S, and L–K; (3) used each set of seasonal estimates to deseasonalize the data; (4) forecasted each deseasonalized series using conventional exponential smoothing methods; (5) reseasonalized the forecasts; (6) compared the reseasonalized forecasts to the withheld actuals and recorded the MAPE over several forecasting horizons. We categorized each series in regard to the recommended method of seasonal adjustment, based upon Table 5. To do this, we recorded the value of W^{J-S} that was used for J–S and the coefficient of skewness of the CD seasonal factors. We found that several series were unsuitable for this analysis. Six series had too few data to support CD, which requires 36 observations in order to provide two preliminary seasonal estimates for each season. One series was deemed inappropriate for the analysis due to its bizarre behavior. (It drops from 5000 to 200 over 10 months, with a precipitous decline near the end of the series.) Six series were deemed non-seasonal, because the variance of the data was increased by seasonal adjustment. These exclusions left 55 seasonal series available for analysis. Of these, 31 satisfied the Table 5 criteria for a Lemon–Krutchkoff recommendation. (We used 'coefficient of skewness >0.5' as the operational definition of an asymmetrical distribution.) Ten met the criteria for choosing the James–Stein estimator, and 14 met the criteria for choosing either CD or the James–Stein estimator. We compared the performances of the methods for each group of series and for the entire set of series as a group.

In developing forecasts of deseasonalized data, we used three types of exponential smoothing: simple (SES), Holt's method (Holt, 1957), and damped-trend (Gardner & McKenzie, 1985). For each series and set of seasonal factors, we chose the exponential smoothing method that produced the smallest MAPE. Parameter values for the exponential smoothing methods were those that minimized one-step-ahead mean squared error in model fitting. The parameter values were chosen from the following ranges:

- SES: α between 0.01 and 0.9
- Holt: α between 0.01 and 0.9; β between 0 and 0.15
- Damped-trend: α between 0.01 and 0.9, β between 0 and 0.15; ϕ between 0.9 and 1.0

Each forecast was evaluated over five horizons: 1, 3, 6, 12, and 18 months. Since the scales of the series differed substantially, MAPE was used to measure forecast error over a given horizon. For example, over a 3-month horizon:

$$\text{MAPE}_3 = 100\frac{\left[\left|\frac{e_1}{X_1}\right| + \left|\frac{e_2}{X_2}\right| + \left|\frac{e_3}{X_3}\right|\right]}{3} \qquad (12)$$

where: $e_t = X_t - F_t$; X_t = actual for period t; F_t = forecasted value for period $t = \hat{S}_t * F\text{-deseas}_t$; $\hat{S}_t$ = estimated seasonal factor for period t; $F\text{-deseas}_t$ = forecasted deseasonalized value for period t.

Note that, since each error measure (e.g. MAPE_3 above) reflects forecasting errors month-by-month rather than cumulatively, it is especially sensitive to the accuracy of seasonal factor estimates. To summarize the performance of a given method of seasonal adjustment using a given forecasting method, we used the average MAPE over the group of series. Thus, for a 3-month horizon:

$$\text{average MAPE} = \frac{\sum_{i=1}^{n} \text{MAPE}_{3,i}}{n} \quad (13)$$

where n = number of series.

4.1. Results

Table 6 provides the average MAPEs resulting from the use of each method for several groups of series: (1) the 31 series for which the recommended method of seasonal adjustment is Lemon–Krutchkoff; (2) the 10 series for which the recommended method is James–Stein; (3) the 14 series for which the recommended method is James–Stein but CD is considered a suitable alternative; and (4) all 55 series. The table also provides ratios of the average MAPEs for each pair of methods.

For the 31 series for which the recommended method of seasonal adjustment is Lemon–Krutchkoff, use of L–K indeed leads to the smallest average MAPE. For the 10 series for which the James–Stein method is recommended, the results are somewhat spotty, perhaps due to the small number of series. Use of J–S generally produced more accurate forecasts than the other two methods. For the 14 series for which the recommended series is J–S, but CD is also considered suitable, the choice of method for seasonal adjustment did not make a substantial difference in forecasting accuracy. Both empirical Bayes methods produced slightly more accurate forecasts than CD over all horizons.

Thus, the recommendations that resulted from the simulation in Section 2 held up rather well when applied to the M-competition series. When a given method was recommended, it indeed led to the most

Table 6
Forecasting accuracy of three methods of seasonal adjustment for 55 M-competition series

Horizon	Average MAPE				Ratios of average MAPEs[b]					
	CD	J–S	L–K	Rec[a]	J–S/CD	L–K/CD	L–K/J–S	Rec/CD	Rec/J–S	Rec/L–K
L–K recommended: 31 *series*										
1	6.56	6.58	5.93		1.002	0.904**	0.902*			
3	7.72	7.66	7.05		0.992	0.914***	0.921**			
6	8.86	8.94	8.52		1.009	0.962***	0.953***			
12	10.43	10.32	10.17		0.989	0.975**	0.986**			
18	10.90	10.73	10.61		0.984**	0.973**	0.989			
J–S recommended: 10 *series*										
1	7.07	7.47	7.75		1.058	1.097	1.037			
3	10.44	9.06	9.72		0.868***	0.932**	1.073***			
6	10.93	10.53	10.71		0.964	0.980	1.017			
12	13.06	12.28	12.92		0.941*	0.990	1.052*			
18	14.35	13.62	14.32		0.949*	0.998	1.052			
CD or J–S recommended: 14 *series*										
1	5.00	4.85	4.93		0.972	0.987	1.016			
3	5.94	5.72	5.90		0.962**	0.994	1.033			
6	5.86	5.71	5.78		0.976	0.986	1.011			
12	7.07	7.00	7.02		0.991	0.993	1.002			
18	8.40	8.38	8.38		0.997	0.998	1.001			
All 55 *series*										
1	6.256	6.301	6.006	5.828	1.007	0.960	0.953	0.932	0.925	0.970
3	7.761	7.421	7.243	6.903	0.956**	0.933***	0.976	0.889***	0.930**	0.953**
6	8.473	8.407	8.221	7.980	0.992	0.970***	0.978*	0.942***	0.949**	0.971
12	10.053	9.831	9.868	9.507	0.978**	0.982**	1.004	0.946***	0.967*	0.963*
18	10.891	10.657	10.717	10.217	0.979**	0.984**	1.006	0.938***	0.959	0.953

[a] Rec = 'use recommended method.' For the third group of series (Recommended = 'CD or J–S'), the J–S estimates were always chosen.
[b] Asterisks indicate P-values for paired t-tests: * $P<0.10$; ** $P<0.05$; *** $P<0.01$. All tests are one-tailed, except for the comparisons of L–K vs. J–S over all 55 series, which are two-tailed.

accurate forecasts on average. However, a forecaster might prefer to choose one method of seasonal adjustment for all series. When the same method is used for all 55 series, both empirical Bayes methods generally produced more accurate forecasts than CD. Compared to CD, J–S provided reductions in average MAPE ranging from 0 to 2.2%. L–K provided reductions in average MAPE ranging from 1.6 to 6.7% when compared to CD. When applied to all 55 series, L–K was the most accurate method on average.

How much advantage is derived from always using the recommended method, rather than choosing one method for all series? Table 6 shows that this approach indeed leads to more accurate forecasts.

On average, using the empirical Bayes methods indeed leads to reduced forecasting error. But how consistently does each method provided greater accuracy than classical decomposition? These results are shown in Table 7, using average MAPE over the forecasting horizon as the accuracy criterion. Essentially identical results were achieved using average RMSE; these are not shown.

For the 31 series for which the Lemon–Krutchkoff method is recommended, L–K resulted in more accurate forecasts than CD for 66.45% of all forecasts (31 series, five horizons per series). The corresponding result for J–S vs. CD is 52.26%. Thus, a much greater advantage was achieved from using L–K than from using J–S, which exhibited almost no advantage over CD.

For the 10 series for which the James–Stein

Table 7
Relative frequency with which each proposed method provided a smaller MAPE than classical decomposition[a]

Horizon	% LK<CD	% J–S<CD	% L–K<J–S	% Rec[b]<CD
L–K recommended: 31 *series*				
1	61.29	54.84	51.61	
3	61.29	48.39	58.06	
6	70.97**	41.94	70.97**	
12	67.74*	51.61	58.06	
18	70.97**	64.52*	48.39	
All horizons	66.45***	52.26	57.42*	
J–S recommended: 10 *series*				
1	50.00	70.00	60.00	
3	90.00	90.00	0.00***	
6	80.00*	70.00	10.00***	
12	60.00	70.00	40.00	
18	60.00	60.00	30.00	
All horizons	68.00***	72.00***	28.00***	
CD or J–S recommended: 10 *series*				
1	57.14	57.14	42.86	
3	42.86	57.14	28.57***	
6	42.86	71.43*	42.86	
12	42.86	57.14	42.86	
18	50.00	57.14	50.00	
All horizons	47.14	60.00*	41.43*	
All 55 *series*				
1	58.18	58.18	50.91	61.82*
3	56.36	63.64	40.00	65.45**
6	63.64**	60.00*	52.73	70.91***
12	56.36	60.00*	50.91	65.45**
18	63.64**	61.82*	52.73	65.45**
All horizons	59.64***	60.73***	49.45	65.82***

[a] P-values for sign tests: * $P<0.10$; ** $P<0.05$; *** $P<0.01$. All tests are one-tailed, except for the comparisons of L–K vs. J–S over all 55 series, which are two-tailed.
[b] Rec = 'use recommended method.' For the third group of series (Recommended = 'CD or J–S'), the J–S estimates were always chosen.

method is recommended, J–S resulted in more accurate forecasts than CD for 72% of all forecasts. The corresponding result for L–K vs. CD is 68%. Here, J–S was only marginally superior to L–K.

For the 14 series for which either CD or J–S is recommended, J–S resulted in more accurate forecasts than CD for 60% of all forecasts. The corresponding result for L–K vs. CD is 47.14%. Here, L–K exhibited no advantage over CD, while J–S was superior to CD.

If the same method were used for all 55 series, L–K resulted in more accurate forecasts than CD for 59.64% of all forecasts, while J–S resulted in more accurate forecasts than achieved CD for 60.73% of all forecasts. Overall, then, L–K and J–S were approximately equally effective compared to CD with respect to forecasting accuracy. However, use of the recommended method in all cases led to more accurate forecasts than were obtained using CD for 65.82% of all forecasts.

5. Examples with real data

We use two real series to illustrate the effects of shrinkage. The first series consists of transformed sales data for 'Company *X*' (Chatfield & Prothero, 1973) over a period of 77 months. The second is series 54 from the M-competition, consisting of 'private company' data values over a period of 74 months. In order to observe the impact of seasonal shrinkage on forecasting accuracy, we withheld the last 18 periods of data for each series. Thus, in each case there are slightly fewer than 5 years of monthly data available for forecasting.

Fig. 6a,b shows plots of these series. Both exhibit trend. The Company *X* series is characterized by an obvious pattern of seasonal variation and comparatively little random variation. In contrast, the seasonal pattern for the private company series is not apparent, and random variation appears to be much greater than any seasonal variation. Thus, it can be anticipated that little shrinkage of classical decomposition seasonals is required for the Company *X* data, whereas considerable shrinkage might be needed for the private company data.

The analysis bears out these expectations. For Company *X* data, the James–Stein shrinkage weight was $W^{J-S} = 0.015$; thus virtually no shrinkage was suggested. Forecasts of this series are essentially unaffected by the application of seasonal shrinkage. As this series illustrates, using shrinkage estimates neither helps nor hurts forecasting accuracy when seasonal variation dominates random variation.

For the private company data, classical decomposition suggests the presence of considerable seasonal variation. Fig. 6c shows that the classical decomposition factors range from 0.75 to 1.23, and they are distributed symmetrically (coefficient of skewness = 0.038). The range of the estimated seasonal factors is surprising since seasonal variation of this magnitude was not apparent in the plot of the data; it underscores the possibility that classical decomposition has exaggerated seasonal variation. Indeed, the shrinkage weight for this series is $W^{J-S} = 0.439$; thus, considerable shrinkage is suggested. The factors that result from using the James–Stein estimator are shown in Fig. 6c. We forecasted this series using damped-trend Holt exponential smoothing (which fit the data better than Holt or simple exponential smoothing). We deseasonalized the data using (1) the classical decomposition factors and (2) the James–Stein seasonals. In each case, we forecasted the deseasonalized data over an 18-month horizon, then reseasonalized the forecast. Using CD seasonals, the MAPEs were 29.3 (over 6 months), 23.03 (over 12 months), and 22.94 (over 18 months). Using J–S seasonals, the corresponding MAPEs were 28.10, 18.79, and 16.26. Such a reduction of forecasting error for this single series exemplifies the forecasting results in Section 4.

6. Conclusions

What have we learned? The simulation results were approximately the same for both multiplicative and additive models, and they were not sensitive to the choice of MSE or MAPE as an error statistic. The simulation confirms that classical decomposition tends to exaggerate seasonal variation, especially under conditions for which simple methods are often preferred (relatively short series, noisy data). The proposed empirical Bayes methods were superior to classical decomposition for almost all conditions, often by a substantial amount. The James–Stein

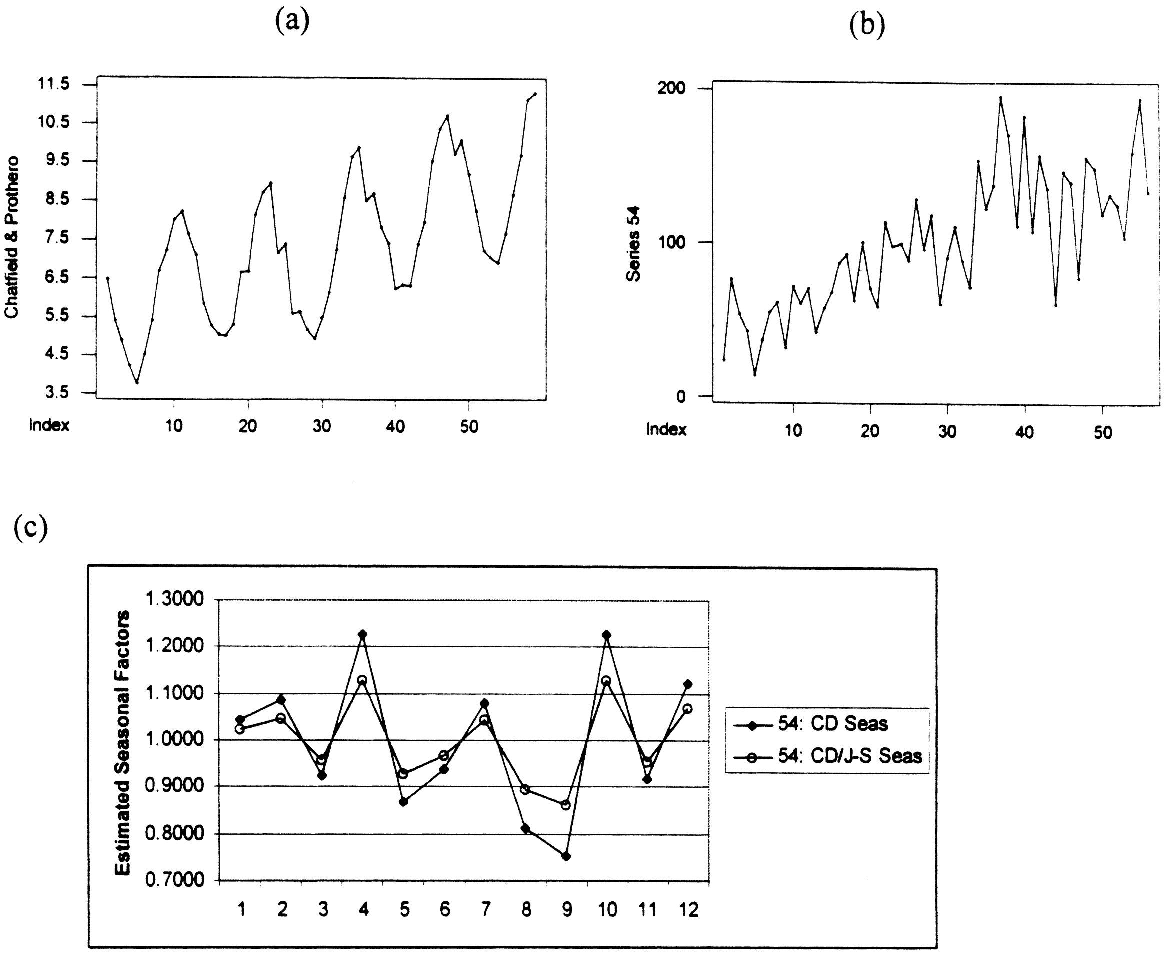

Fig. 6. Plots of (a) Chatfield and Prothero Company *X* data; (b) series 54: private company data; (c) CD and CD/J–S seasonal factor estimates for series 54: private company data.

method was never less accurate than classical decomposition and sometimes achieved reductions of nearly 80% in MSE. The Lemon–Krutchkoff method was usually, but not always, more accurate than classical decomposition. When the distribution of seasonal factors is asymmetric, the Lemon–Krutchkoff method was often superior to the James–Stein method. Specific guidelines were developed for choosing between the James–Stein method and the Lemon–Krutchkoff method. Generally, as W^{J-S} becomes smaller and skewness among seasonals becomes greater, conditions favor the Lemon–Krutchkoff method. Larger values of W^{J-S} and less skewness among seasonals favor the James–Stein method.

The examination of forecasting results based on the M-competition data supported the simulation findings and revealed a decided forecasting advantage when empirical Bayes methods were used for seasonal adjustment rather than classical decomposition. Using either empirical Bayes method rather than classical decomposition led to more accurate forecasts for 60% of all series, and using the guidelines for choosing an empirical Bayes method led to more accurate forecasts for 66% of all series. Since the estimation of seasonal factors is only one among many sources of forecasting error with real data, these are impressive results. Of the two empirical Bayes methods, the guidelines recommended the

Lemon–Krutchkoff method more often than the James–Stein (56% Lemon–Krutchkoff vs. 44% James–Stein).

The simulation and the forecasting analysis of M-competition data strongly indicate that the empirical Bayes methods are superior to classical decomposition as a method for seasonal adjustment in forecasting. Since this paper specifically addresses forecasting when simple methods are preferred, it is fair ask whether they can be considered simple methods. We believe the answer is yes. Both methods are easily implemented within a spreadsheet. Executing classical decomposition requires about five calculations (five columns in a spreadsheet). The James–Stein method requires four additional calculations. The Lemon–Krutchkoff method requires considerably more (about 27 additional calculations, all very similar), but they are not difficult to implement and, once in place, they are performed routinely.

Acknowledgements

This research was supported by a grant from VCU School of Business Faculty Excellence Fund.

References

Armstrong, J. S. (1985). *Long-range forecasting*. New York: Wiley.

Assimakopoulos, V., & Nikolopoulos, K. (2000). The theta model: a decomposition approach to forecasting. *International Journal of Forecasting*, *16*, 521–530.

Bunn, D. W., & Vassilopoulos, A. I. (1999). Comparison of seasonal estimation methods in multi-item short-term forecasting. *International Journal of Forecasting*, *15*, 431–443.

Chatfield, C., & Prothero, D. L. (1973). Box–Jenkins seasonal forecasting: problems in a case study. *Journal of the Royal Statistical Society A*, *36*, 295–336.

Efron, B., & Morris, C. (1973). Stein's estimation rule and its competitors—an empirical Bayes approach. *Journal of the American Statistical Association*, *68*, 117–130.

Efron, B., & Morris, C. (1975). Data analysis using Stein's estimator and its generalizations. *Journal of the American Statistical Association*, *70*, 311–319.

Gardner, E. S., & McKenzie, E. (1985). Forecasting trends in time series. *Management Science*, *31*, 1237–1246.

Greis, N. P., & Gilstein, C. Z. (1991). Empirical Bayes methods for telecommunications forecasting. *International Journal of Forecasting*, *7*, 183–197.

Holt, C. C. (1957). *Forecasting seasonals and trends by exponentially weighted moving averages*, Research memorandum no. 52. Office of Naval Research.

James, W., & Stein, C. (1961). Estimation with quadratic loss. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, vol. 1. Berkeley, CA: University of California Press, pp. 361–379.

Lemon, G. H., & Krutchkoff, R. G. (1969). Smooth empirical Bayes estimators: with results for the binomial and normal situations. In *Proceedings of the symposium on empirical Bayes estimation and computing in statistics*, *Texas Tech University mathematics series no. 6*, pp. 110–140.

Macauley, F. R. (1930). *The smoothing of time series*. National Bureau of Economic Research.

Makridakis, S., Anderson, A., Carbone, R., Fildes, R., Hibon, M., Lewandowski, R., Newton, J., Parzen, E., & Winkler, R. (1982). The accuracy of extrapolation (time series) methods: results of a forecasting competition. *Journal of Forecasting*, *1*, 111–153.

Makridakis, S., & Hibon, M. (2000). The M3-Competition: results, conclusions and implications. *International Journal of Forecasting*, *16*, 451–476.

Makridakis, S., Wheelwright, S. C., & Hyndman, R. J. (1998). *Forecasting methods and applications*, 3rd ed. New York: Wiley.

Meade, N. (2000). A note on the robust trend and ARARMA methodologies used in the M3 competition. *International Journal of Forecasting*, *16*, 517–519.

Morris, C. N. (1983). Parametric empirical Bayes inference: theory and applications. *Journal of the American Statistical Association*, *78*, 47–55.

Robbins, H. (1955). An empirical Bayes approach to statistics. *Proceedings of the Third Berkeley Symposium on Probability and Statistics*, *1*, 157–163.

Rutherford, J. R., & Krutchkoff, R. G. (1969). An empirical Bayes smoothing technique. *Biometrika*, *56*(1), 133–137.

Shiskin, J., Young, A. H., & Musgrave, J. C. (1967). *The X-11 variant of the Census II method seasonal adjustment program*, Technical paper no. 15. Bureau of the Census.

Stein, C. (1955). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *Proceedings of the Third Berkeley Symposium on Probability and Statistics*, *1*, 197–206.

Biographies: Don MILLER teaches courses in business forecasting and statistics at Virginia Commonwealth University. He has worked as a forecaster at Xerox Corporation and as a forecasting consultant to industry and state government for over 20 years.

Dan WILLIAMS teaches courses in forecasting, research methods, and government at Baruch College. He forecast Virginia Medicaid expenditures from 1987 through 1995. Over this time, Virginia was recognized for its financial management with specific mention of Medicaid forecasting.