



ELSEVIER

International Journal of Forecasting 18 (2002) 117–123

international journal
of forecasting

www.elsevier.com/locate/ijforecast

Seasonal adjustment of inventory demand series: a case study

Everette S. Gardner Jr.^{*}, Joaquin Diaz-Saiz

Center for Global Manufacturing, College of Business Administration, University of Houston, Houston, TX 77204-6282, USA

Abstract

This paper analyzes procedures for seasonal adjustment of inventory demand series at a large US auto parts distributor, BPX Holding Corporation of Houston, TX. The company's forecasting system made no attempt to classify demand series as seasonal or nonseasonal. All demand series were assumed to be seasonal. They were seasonally-adjusted using a multiplicative decomposition procedure, then forecasted with exponential smoothing. We show that simple methods of identifying seasonal series, coupled with an additive decomposition procedure, can make significant reductions in forecast errors and safety stock investment. We also discuss forecasting implementation problems in inventory control systems. © 2002 International Institute of Forecasters. Published by Elsevier Science B.V.

Keywords: Forecasting; Time series; Seasonal adjustment; Inventory; Distribution

1. Introduction

Our experience is that seasonality is treated rather casually in forecasting practice for inventory control. In many companies, demand series are classified as seasonal or nonseasonal based on management judgment, usually with no information from statistical testing. Other companies do not bother with classifying demand series and simply assume that all are seasonal. Regardless of how demand series are classified, it is common to assume that every seasonal pattern is multiplicative. If forecasting practice follows research, this assumption is reasonable.

A literature search found no evidence that additive or mixed additive/multiplicative models have ever been tested in empirical time series research.

This paper presents a case study in which improved methods of identifying seasonal time series, coupled with additive seasonal decomposition, made significant reductions in forecast errors and safety stock investment. The study was performed at BPX Holding Corporation of Houston, TX, a large auto-parts distributor. Section 2 describes seasonal adjustment and forecasting procedures at BPX. In Section 3, we discuss problems in designing alternative procedures, implementation, and measuring improvements. Comparisons of the recommended system to current practice are given in Section 4. Conclusions are offered in Section 5.

^{*}Corresponding author. Tel.: +1-713-743-4744; fax: +1-713-743-4940.

E-mail address: egardner@uh.edu (E.S. Gardner Jr.).

2. BPX seasonal adjustment and forecasting procedures

At the time of the study, the company operated 24 distribution centers, which supplied 350 company-owned stores and about 1600 affiliated stores. Stock levels at each distribution center were based on forecasts of seasonally-adjusted aggregate demand for the stores supplied by that center. Forecasts were provided to a variation of the economic order quantity (EOQ) model to determine order quantities, while safety stocks were set at a multiple of the mean absolute forecast error, usually referred to as the MAD (mean absolute deviation). Details of the EOQ and relative safety stock allocation varied considerably depending on the importance of the item.

Seasonal identification was not an issue at BPX. Every demand series was treated as multiplicative seasonal, with indices obtained through one of two procedures. Firstly, some series were decomposed using what appeared to be a standard ratio-to-moving-average procedure. This procedure was undocumented and had been in use for many years. We compared results for the BPX procedure and the ratio-to-moving average procedure in Makridakis, Wheelwright, and Hyndman (1998) and found that they were virtually identical. One complication in the BPX procedure was that demands were intermittent for some series, especially during seasonal trough periods, a problem which made it impossible to compute multiplicative seasonal indices. For intermittent series, BPX added a constant before decomposition and removed it afterward.

The second procedure for obtaining seasonal indices was to select them from a library of predetermined indices. Some years ago, BPX decomposed demand data for groups of similar items in order to develop predetermined indices. Over the years, they were subjectively modified many times by inventory managers; at the time

of our study, it was not possible to determine the origin or rationale for any particular set of predetermined indices.

BPX forecasted seasonally-adjusted demand with the Trigg and Leach (1967) adaptive method of simple exponential smoothing. As each error is observed, this method sets the value of the smoothing parameter equal to the absolute value of a ratio, the smoothed forecast error divided by the smoothed MAD. The aim is to apply moderate smoothing parameters during periods of stability in the time series; when the structure of the series changes, parameters should automatically increase to shorten the response lag in the forecasts. As discussed in Gardner (1985), there is no evidence that adaptive smoothing parameters offer any significant advantage in forecast accuracy. However, adaptive parameters require little maintenance, a considerable advantage in this application which involved forecasting more than 100,000 time series every 4 weeks.

3. Research design

A new manager at BPX noticed some demand series in which seasonal adjustment produced curious results. He asked us to review seasonal adjustment procedures and suggest alternatives. We argued for expansion of the study to include exponential smoothing methods, including those designed specifically for intermittent series (Croston, 1972; Johnston & Boylan, 1996). We also suggested testing model identification procedures for exponential smoothing, such as those proposed in Gardner and McKenzie (1988). However, we were overruled. The manager was convinced that simple exponential smoothing was the only reasonable forecasting method and asked us to concentrate on the seasonal adjustment problem.

The manager pre-selected four large distribution centers for the study, in Florida, Minnesota,

Missouri, and California. At each center, the study was limited to ‘fast-movers’, those inventory items with sufficient demand to require regular forecast updates and stock replenishment decisions. ‘Slow-movers’ or items with infrequent demand were excluded from the study. We agreed with the justification, that slow movers would likely require different procedures and should be treated in a separate study.

We proposed the following plan: draw random samples of demand series from each distribution center. Test each series for seasonality and if seasonal, perform an additive rather than multiplicative decomposition because of the intermittent nature of many series. We could see no evidence of consistent trends, so additive decomposition should also give reasonable results in continuous series. Additive decomposition was performed using a Delphus (1997) product, Peer Planner (2000), which gives results identical to the procedure in Makridakis et al. (1998).

This was a simple plan but it was not clear how the results should be measured. Management was not interested in the typical summary forecast error measures found in empirical research. Instead, they wanted to know how seasonal adjustment affected inventory performance. It may not be obvious that alternative seasonal adjustment procedures can result in very different stock levels and replenishment workloads. This is because the timing and quantities of stock replenishment decisions are functions of the forecasts. Another consideration is that seasonal adjustment affects customer service because safety stocks are a multiple of the MAD of the forecast errors.

To measure inventory effects, we considered developing tradeoff curves between inventory investment and customer service as in Silver and Peterson (1985) or Gardner (1990). We also considered response surfaces showing the tradeoffs among replenishment workload, investment, and service (see Gardner and Dan-

nenbring (1979)). Unfortunately, time pressures ruled out these types of tradeoff analysis. BPX had recently acquired several smaller companies and had difficulty merging inventories of the acquisitions with BPX stocks. Seasonality appeared to be an issue in managing the new inventories, so we were under pressure to produce quick results.

Therefore, we decided to keep the study as simple as possible and estimate the effects of seasonal adjustment on only one performance measure, the aggregate MAD of forecast errors at each distribution center. Any reduction in aggregate MAD would reduce safety stock investment with no impact on customer service. Although the MAD is widely used as a measure of the dispersion of forecast errors in inventory control, Brown (1982) showed that it can be seriously misleading. We agree with Brown but BPX programming constraints were such that the MAD would have to be retained, at least in the short term.

Another problem in research design was in the identification of seasonal series. BPX had 3 years of inventory demand history available. The company operated with 13 4-week accounting periods per year, so we had time series of 39 observations. We decided not to attempt to implement autocorrelation analysis because of the short series as well as the antiquated BPX computer system, which dated to the 1950s. Processing time to update forecasts was already a serious issue and it was not feasible to add significant additional computations. We also anticipated problems in understanding autocorrelation analysis within BPX. Therefore, we chose a simpler method of seasonal identification, a comparison of the variance of the original demand series with the seasonally-adjusted series. If seasonal adjustment made a significant reduction in variance at the 95% significance level, the series was declared seasonal. Certainly this method of seasonal identification is limited and can be misleading.

Even if seasonality is not present, the variance of the seasonally adjusted series can be small because outliers and irregularities have been smoothed. In some cases, just the opposite can occur. However, as discussed below, the variance test worked well compared to both autocorrelation and graphical analysis of the test series.

4. Empirical results

We drew several small random samples of demand series to estimate variances and computed sample sizes that would give reasonable confidence intervals around our results. Our original sample totaled 290 series. Few were identified as seasonal by the variance-ratio test, a conclusion that management found difficult to accept. Therefore management asked for an additional sample of the same size. When this sample was completed, management asked for further small samples drawn only from temperature control (heating and air conditioning) parts because these parts were thought to be more seasonal than others. Tables 1 and 2 present results for all samples combined at each center,

stratified by temperature control and all other parts.

Table 1 shows the percentage of series in which variance was significantly reduced by seasonal adjustment. BPX seasonal adjustment procedures proved hopeless, failing to make significant variance reductions in any series. Additive adjustment was only slightly more successful, making significant variance reductions in only 11% of the 691 series. Notice that additive adjustment was somewhat more successful in temperature-control parts, especially in Minnesota, and less successful in the others.

To confirm the accuracy of the variance-ratio test for seasonality, we performed four additional tests: (1) autocorrelation analysis of the series both before and after adjustment, (2) autocorrelation analysis of the residuals, (3) comparison of the variances of original series with seasonally-differenced series, and (4) examination of plots of each year superposed to look for evidence of consistent seasonal peaks. Regardless of the test, there was no significant difference in the number of series identified as seasonal.

Because the data were given in 4-week periods, trading day variations were avoided.

Table 1
Identification of seasonal demand series

Inventory		Number of series	Percentage of series in which variance was significantly reduced	
			BPX adjustment	Additive adjustment
Florida	Fast-movers	189	0	5%
	Temp. control	27	0	11%
Minnesota	Fast-movers	139	0	13%
	Temp. control	26	0	65%
Missouri	Fast-movers	139	0	7%
	Temp. control	28	0	29%
California	Fast-movers	115	0	3%
	Temp. control	28	0	29%
Total		691	0	11%

Table 2
MAD reduction from additive seasonal adjustment

Inventory		Option A: adjust series when significant variance reduction occurs			Option B: adjust series when any variance reduction occurs		
		Percent MAD reduction	95% confidence limits		Percent MAD reduction	95% confidence limits	
			Lower	Upper		Lower	Upper
Florida	Fast-movers	18	15%	20%	16	14%	18%
	Temp. control	22	14%	30%	22	16%	28%
Minnesota	Fast-movers	21	17%	24%	18	15%	20%
	Temp. control	50	38%	61%	43	33%	52%
Missouri	Fast-movers	18	15%	20%	17	15%	19%
	Temp. control	31	21%	41%	19	11%	27%
California	Fast-movers	11	8%	14%	19	16%	21%
	Temp. control	32	22%	41%	20	13%	27%
Weighted averages	Fast-movers	17	16%	18%	17	16%	18%
	Temp. control	33	28%	38%	26	22%	29%
Total		20	19%	21%	19	17%	20%

However, it is possible that all tests could be confounded by the timing of holiday periods. In graphical analysis, we found no evidence of outliers due to the timing of holidays. This is perhaps not surprising in that the data represent highly aggregated wholesale sales, which should be less affected by the timing of holidays than retail sales.

Estimated percent MAD reductions are shown in Table 2. To compute current MAD values, each demand series was adjusted using BPX indices. Next, Trigg-and-Leach exponential smoothing was run on the seasonally-adjusted data and the MAD was computed using data from the last year. In option A, on the left side of the table, new MADs were computed by applying additive seasonal adjustment only when a significant reduction in variance occurred; otherwise, the original demand series was smoothed. Overall, we estimated a 20% reduction in MAD, with a 95% confidence interval from 19 to 21%.

After the results in option A were presented, management requested that we discard the

significance test and perform additive seasonal adjustment when any reduction in variance occurred. The results, shown as option B on the right side of Table 2, are similar. Why? Detailed examination showed that most additional series adjusted in option B contained very weak seasonal patterns. Thus seasonal indices were quite small and made little difference in the variance of the series.

The estimates in Table 2 are not sensitive to computational procedures. Even though management was not interested in smoothing methods for intermittent series, we tested Croston's (1972) method for both BPX and additive-adjusted data and found no significant difference in estimated MAD reductions. Next, we computed a fitted MAD over all 39 observations of each series, again with no significant difference. Finally, we used only 32 observations to develop additive seasonal indices and fit the exponential smoothing model; a one-step-ahead forecast simulation over the last seven observations produced no significant difference in results.

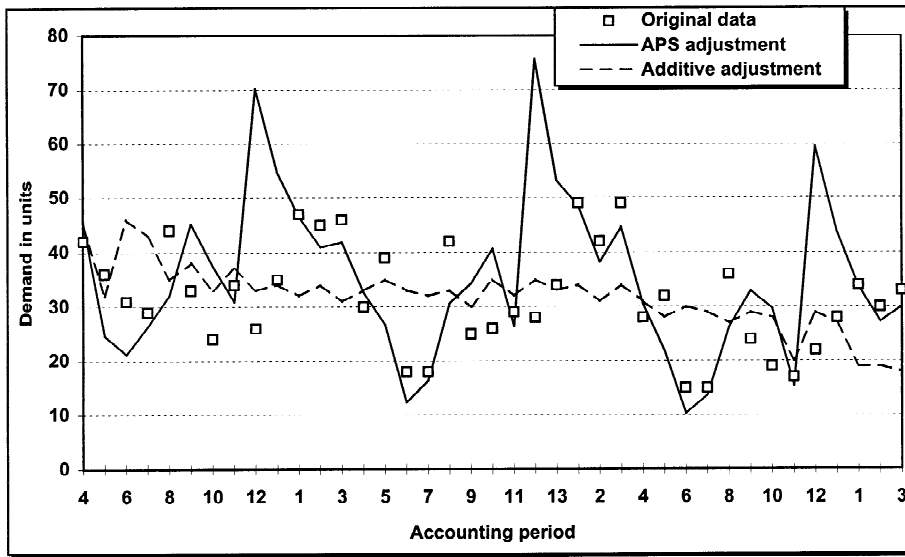


Fig. 1. Effects of seasonal adjustment with continuous demand.

Examples of individual time series are given in Figs. 1 and 2. A continuous time series is shown in Fig. 1, with both BPX and additive-adjusted data, while an intermittent series is shown in Fig. 2. In both cases, BPX-adjusted

data varies over a range considerably larger than that of original data while the additive-adjusted data are well-behaved. Figs. 1 and 2 are typical. In most sample series, BPX seasonal adjustment served only to make forecasting more difficult.

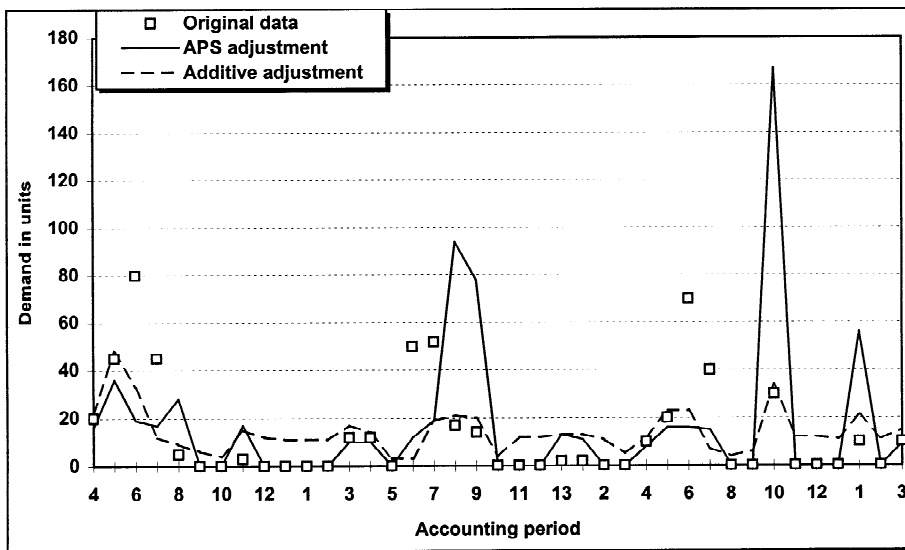


Fig. 2. Effects of seasonal adjustment with intermittent demand.

5. Conclusions

Management extrapolated the results in Table 2 (option A) to the entire inventory and estimated that the MAD reduction would in turn reduce safety stocks by about 20% or \$5 million. Another benefit would be improved accuracy in purchase quantities although we made no attempt to estimate this. We did not have the data to verify the dollar reduction or to estimate the percentage reduction in total inventory value. The savings were never realized because the company entered bankruptcy proceedings just as the study was completed.

Several conclusions follow from this research. Firstly, it may be possible to reduce inventory investment in other companies which ignore seasonality testing of demand series. The research design outlined above is a simple way to estimate changes in investment.

Secondly, it is dangerous to assume that every seasonal demand series is multiplicative. Inventory managers should consider additive seasonal models because they work well with intermittent data and are more robust to outliers. A related point was noted by one of the referees. Even when a series is known to be influenced by seasonality, this does not mean that the type of seasonality can be well identified and estimated by a seasonal adjustment method. When this is the case, it is better not to seasonally adjust.

Finally, we believe that researchers should evaluate additive and mixed additive/multiplicative seasonal models in empirical studies of forecast accuracy. In a literature search, we could find no evidence that such models have ever been used in empirical research. We were surprised by the lack of evidence but it is confirmed in a recent survey paper on univariate forecasting (Fildes et al., 1998), which makes no mention of alternative methods of modeling seasonality.

References

- Brown, R. G. (1982). In: 2nd ed, *Advanced service parts inventory control*, Materials Management Systems, Norwich, Vermont.
- Croston, J. D. (1972). Forecasting for intermittent demand. *Operational Research Quarterly* 3, 289–303.
- Delphus, Inc (1997). *Peer planner*, Morristown, New Jersey.
- Fildes, R. et al. (1998). Generalizing about univariate forecasting methods: further empirical evidence. *International Journal of Forecasting* 14, 339–358.
- Gardner, E. S. (1985). Exponential smoothing: the state of the art. *Journal of Forecasting* 4, 1–38.
- Gardner, E. S. (1990). Evaluating forecast performance in an inventory control system. *Management Science* 36, 490–499.
- Gardner, E. S., & Dannenbring, D. G. (1979). Using optimal policy surfaces to analyze aggregate inventory tradeoffs. *Management Science* 25, 709–720.
- Gardner, E. S., & McKenzie, E. (1988). Model identification in exponential smoothing. *Journal of Operational Research Society* 39, 863–867.
- Johnston, F. R., & Boylan, Y. E. (1996). Forecasting for items with intermittent demand. *Journal of Operational Research Society* 47, 113–121.
- Makridakis, S. et al. (1998). In: 3rd ed, *Forecasting: methods and applications*, John Wiley, New York.
- Silver, E. A., & Peterson, R. (1985). In: 2nd ed, *Decision systems for inventory management and production planning*, Wiley, New York.
- Trigg, D. W., & Leach, A. G. (1967). Exponential smoothing with an adaptive response rate. *Operational Research Quarterly* 18, 53–59.

Biographies: Everette S. GARDNER, Jr. is Professor of Decision and Information Sciences and Director of the Center for Global Manufacturing at the University of Houston, where he teaches courses in operations management. His research interests are in forecasting and management of distribution inventories.

Joaquin DIAZ-SAIZ is Associate Professor of Decision and Information Sciences at the University of Houston, where he teaches courses in statistics. His research interests are in time series analysis and Bayesian methods.