Damped Seasonality Factors: Introduction

Forthcoming, *International Journal of Forecasting*

J. Scott Armstrong
The Wharton School
University of Pennsylvania, Philadelphia, PA

armstrong@wharton.upenn.edu

Previous research has shown that seasonal factors provide one of the most important ways to improve forecast accuracy. For example, in forecasts over an 18-month horizon for 68 monthly economic series from the M-Competition, Makridakis et al. (1984, Table 14) found that seasonal adjustments reduced the MAPE from 23.0 to 17.7 percent, an error reduction of 23%. On the other hand, research has also shown that seasonal factors sometimes increase forecast errors (e.g., Nelson, 1972).

So, when forecasting with a data series measured in intervals that represent part of a year, should one use seasonal factors or not? Statistical tests have been devised to answer this question, and they have been quite useful. However, some people might say that the question is not fair. Why does it have to be either/or? Shouldn't the question be "to what extent should seasonal factors be used for a given series?"

Damping as a Basic Strategy for Forecasting

One solution to "to what extent" relies on damping. Basically, damping says that the forecaster is more conservative as uncertainty increases. In 1978, in summarizing research by others, I concluded that trends should be damped (Armstrong 1978, p.153). As nearly as I can tell, only two people took action: Gardner and Mackenzie (1985) provided convincing evidence that damping improved accuracy. Just as important, they provided an operational procedure. Their effort led to one of the more important advances in extrapolation.

Again drawing on the research of others, I concluded in Armstrong (1978, 148-150) that seasonal factors should be damped. Here also only two people listened. Miller and Williams (2003, 2004) have obtained convincing evidence that damped seasonal factors improve forecast accuracy. They have also developed an operational procedure for doing this.

Is the Miller-Williams' procedure optimal? Along with the panelists (including Miller and Williams), I am skeptical that it is. However, I expect that their *statistical* procedures for damping seasonality will prove to be nearly optimal. Their two papers provide evidence from simulations and from analyses of monthly series. For the real data, damping improved accuracy in about 60% of the series, with about a 4% error reduction.

Process for the Special Section on Damped Seasonal Factors

I reviewed early versions of the current Miller-Williams paper in 2001, then in 2002. Following that, I sent their version to five reviewers widely recognized as leaders in the use of seasonal factors. All who were asked agreed to participate. Three others joined as coauthors on the commentaries. In doing their reviews, all checked the Miller-Williams procedures and conducted additional analyses.

Miller and Williams made revisions to address reviewers' concerns. Authors of two commentaries conducted additional analyses. A revised paper and written commentaries were prepared for presentation at the International Symposium on Forecasting in Merida, Mexico, June 2003. The session drew much interest. It also generated a heated discussion by some of those in the audience.

Miller and Williams conducted extensive analyses to test possible threats to validity that were raised at the symposium. This revised version was then sent to the commentators so they could revise their commenty. Additional suggestions were made by the commentators and by another independent reviewer, and Miller and Williams again revised their paper in light of the comments. In all, the process extended over several years and involved ten reviewers. While there have been many additional analyses that have led to substantial improvements in the paper, the original findings have held up well over time.

Implementation

Major software providers such as Forecast Pro and SAS have gone on record as being willing to provide new features upon requests from clients. If you are using seasonal factors to make forecasts, you should ask them and other software developers to implement the Miller-Williams Seasonality procedures.

However, you do not need to wait for the software providers. Miller and Williams have provided full disclosure of their procedures. In addition, they provide spreadsheet software for their procedures as part of the public domain. (See under software at forecastingprinciples.com.) You could direct your current software provider to this procedure.

Beyond Miller-Williams

The major potential for gain beyond the Miller-Williams procedures will come from using additional information. Two sources are promising: domain knowledge and analogous series.

In a pilot study involving a small number of time series, Fred Collopy and I (1998) found that damping based on the amount of data *and* on domain experts' expectations about trend reduced forecast error by more than six percent (see these results at the Researchers page of forecastingprinciples.com). Miller and Williams (2003) showed that if our adjustment were based only on the number of years of data (fewer years implying more uncertainty and thus more damping) there was no improvement over the traditional approach. In effect it likely led to too much damping. Thus, the overall gain we found was apparently due to domain knowledge.

Another approach to damping is to draw the seasonal factors towards the estimates of the seasonal factors for one or more analogous series, For example, you could average seasonal factors for a brand with those for that brand's product class. This is expected to be especially useful when the analogous series involve considerably more data. Bunn and Vassilopoulos (1999) provide procedures for combining seasonal factors across analogous data, along with supporting evidence for such a procedure.

Conclusion

Historically, seasonal factors have been concerned with estimates of historical data, not with forecasting. The Miller-Williams procedures enable organizations to adjust seasonal factors so as to make more accurate forecasts. This has an immediate payoff to forecasters – they can reduce forecast errors by about four percent. In addition, their papers have paved the way for improvements in the estimation of seasonal factors for forecasting, such as through the use of domain knowledge and analogous series.

References

Armstrong, J. S. (1978, 1985, 2nd ed.). *Long-Range Forecasting: From Crystal Ball to Computer*. New York: John Wiley. 1985 edition available in full-text at hops.wharton.upenn.edu/forecast.

Armstrong, J. S. & F. Collopy (1998), Speculations about seasonal factors. In full text on Researchers' page at forecastingprinciples.com.

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

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

Makridakis, S., A. Andersen, R. Carbone, R. Fildes, M. Hibon, R. Lewandowski, J. Newton, E. Parzen & R. Winkler (1984). *The Forecasting Accuracy of Major Time-series Methods.* Chichester: John Wiley.

Miller, D. M. & D. Williams (2004). Shrinkage estimators for damping X-12-ARIMA seasonals. *International Journal of Forecasting*, forthcoming.

Miller, D. M. & D. Williams (2003). Shrinkage estimators of time series seasonal factors and their effect on forecasting accuracy. *International Journal of Forecasting*, 19, 669-684.

Nelson, C.R. (1972). The prediction performance of the FRB-MIT-Penn model of the U.S. economy. *American Economic Review*, 5, 902-917.