Journal of Business Research 68 (2015) 1739–1741

Contents lists available at ScienceDirect

Journal of Business Research

Conservative forecasting with the damped trend

Everette S. Gardner Jr. *

Bauer College of Business, University of Houston, Houston, TX 77204-6021, USA

ARTICLE INFO

Article history:
Received February 2014
Received in revised form January 2015
Accepted January 2015
Available online 3 April 2015

Keywords:
Forecasting
Time series
Extrapolation
Exponential smoothing
Damped trend

ABSTRACT

The "Golden Rule" checklist by Armstrong, Green, and Graefe, in this issue (referred to as AGG below), is a systematic procedure for implementing conservative forecasting principles, and it should help close the long-standing gap between theory and practice. The checklist is both a practical tool and an empirical research agenda. Trend damping is an important part of the checklist, and AGG rely on subjective judgment about when and how damping should be done. I recommend a more objective approach based on the damped trend method of exponential smoothing, which has a long record of success in empirical research.

1. Introduction

By now the empirical research is overwhelming that conservatism improves forecast accuracy, but this research has not carried over to forecasting practice. One reason may be that researchers have never presented a systematic way to implement conservatism. The Golden Rule checklist by Armstrong, Green and Graefe, 2015–in this issue (referred to as AGG below), admirably fills this need.

AGG's conservatism theme reminds me of Stephen Schnaars' book *Megamistakes* (1989), which reviews published forecasts over a 30-year period for new technologies, products, and markets. Schnaars (1989, p. 55) offers the following advice, "Cut or damp *any* trend estimates with which you are provided. Do not be swayed by the sophistication of the forecasting method or the forecaster. Be suspicious. Be especially suspicious of forecasts that are based on accelerating trends in growth. In the past they have led to the largest errors".

According to Schnaars, the history of business forecasting is mostly one of over-forecasting, and thus damping the trend is a simple way to prevent unrealistic forecasts. AGG primarily rely on subjective judgment about when and how damping should be done, but I recommend supplementing their guidance with a model-based approach to damping as discussed below.

2. History of the damped trend

The history of the damped trend method of exponential smoothing is important here because it reinforces AGG's conclusion that forecasting practice has not improved with time. Two early ideas on trend extrapolation influenced development of the damped trend. First, Gilchrist's (1976) *Statistical Forecasting* discusses different means of fitting growth curves, including the simple modified exponential, a model that produces gradually decreasing growth (or decline) toward an asymptote. Gilchrist (1976) suggests the modified exponential might be fitted using discounted least squares, which is R. G. Brown's (1963) classic rationale for exponential smoothing. Second, Scott Armstrong's *Long Range Forecasting* (1978) argues for judgmentally reducing the weight on uncertain trends, which is much the same thing as trend damping. Armstrong and I had many discussions in the early days about trend extrapolation, but our paths eventually diverged. Armstrong took a judgmental approach to trend extrapolation set forth in his many papers on rule-based forecasting, while I took a statistical approach.

During the last few years of my Navy career in the 1980s, I served as Director of Operations Analysis, where I was responsible for numerous projects in forecasting for inventory control. I was concerned about over-forecasting of inventory demands for repair and maintenance parts, a problem that usually leads to excess stocks, and it was clear that a new method was needed to extrapolate trends. At that time, Eddie McKenzie was a Visiting Professor at the Naval Postgraduate School, and we collaborated on the damped trend method of exponential smoothing (Gardner & McKenzie, 1985, 1988, 1989). I also developed the Autocast software (Gardner, 1986) to implement the damped trend. Links to free software for the damped trend are now available at http://www.forlab.eu/forecasting-software.

In Gardner (1990), I published a report on damped trend performance in a real inventory system. Compared to simple exponential smoothing, the method used prior to my study, the damped trend reduced inventory investment by 7% (worth $30 million). An alternative strategy was to keep the same level of inventory investment and use

* Tel.: +1 713 410 1244; fax: +1 713 743 4940.
E-mail address: egardner@uh.edu.

http://dx.doi.org/10.1016/j.jbusres.2015.03.033
0148-2963/

the damped trend to improve customer service, which is measured by the delay time to fill customer orders; the damped trend reduced delay time by 19%, from 32 days to 26 days. I also tested exponential smoothing with a linear trend, which consistently overestimated demand and produced ridiculous levels of inventory investment.

Since 1990, damped trend exponential smoothing has performed well in empirical research (see Gardner, 2006, for a review). The M3 competition (Makridakis & Hibon, 2000) is the only empirical study in which another method, the Theta method of forecasting (Assimakopoulos & Nikolopoulos, 2000), was more accurate than the damped trend. However, in the telecommunications data (Gardner & Diaz-Saiz, 2008), the Theta method performed much worse than the damped trend, and no other empirical comparisons to the damped trend appear to be available. The M3 results are peculiar because Hyndman and Billah (2003) show that the Theta method is equivalent to simple exponential smoothing with a fixed drift (linear trend term) equal to half the slope of a straight line fitted to the data.

In a review of evidence-based forecasting, Armstrong (2006) recommends the damped trend as a well established forecasting method that should improve accuracy in practical applications. In a review of forecasting in operational research, Fildes, Nikolopolus, Crone, and Syntetos (2008) concludes that the damped trend can “reasonably claim to be a benchmark forecasting method for all others to beat”. Additional empirical evidence using the M3 competition data is given by Hyndman, Koehler, Ord, and Snyder (2008), who found that the use of the damped trend method alone compares favorably to complex method selection procedures using statistical information criteria.

Despite these successes, the damped trend is virtually unknown outside a small group of academics. In Gardner (2006), I reviewed 65 empirical studies in forecasting with exponential smoothing that used real data (not counting forecasting competitions and the many associated papers). Only 7 papers even considered the damped trend as a forecasting alternative, and 5 of those were my own. Since 2006, I am aware of only one other paper (Acar & Gardner, 2012) that used the damped trend with real data.

Supply chain management is a likely area for application of the damped trend because production plans and inventory levels depend directly on forecasting. Most supply chain and operations management textbooks include a discussion of exponential smoothing, but I can find none that include the damped trend. Furthermore, I can find only one mention of the damped trend (Li, Disney, & Gaalman, 2014) in any of the supply chain or operations management journals. Few commercial software packages include the damped trend, and the method appears to be unknown in forecasting support systems and enterprise resource planning systems.

Resistance to the damped trend is difficult to understand, especially in the supply chain field. Perhaps it is not so much resistance to the damped trend as it is resistance to conservative forecasting methods, which tend to be simple. The damped trend is relatively simple, whereas many forecasters and their clients have long been seduced by complex methods, as AGG point out.

3. Subjective trend damping

AGG present four rules for conservatism in trend extrapolation, all implemented by trend damping, and all based on a subjective evaluation of the data. I agree that one should be cautious in all of the situations the rules address, but I have reservations about three of the rules. I do agree with Rule 3.3.1, which calls for damping when the series is variable or unstable. I believe a purely statistical approach makes sense here because it is hard to judgmentally determine how variable or unstable a series must be before the trend should be damped.

When the observed trend conflicts with causal forces, a condition called a “contrary series”, AGG recommend damping the trend heavily toward a no-change forecast or even ignoring such trends altogether (Rule 3.3.2). This recommendation could be difficult to implement because it assumes that one can detect changes in causal forces in real time or else predict when future changes will occur. A related problem is deciding how long one should wait before deciding a contrary time series has become the new standard.

When the forecast horizon is longer than the historical series, Rule 3.3.3 recommends damping the trend toward zero or averaging the trend with analogous series. I think this rule is too cautious. In practice, forecast horizons are often longer than the *relevant* historical series. For example, I regularly make long-range forecasts of the numbers of frequent flyer miles that will be redeemed at several airlines. Lengthy time series are available, but events such as mergers and drastic changes in redemption policies have often left me with forecast horizons that are longer than relevant history and with strong reasons to expect growth throughout those horizons.

When short- and long-term trend directions are inconsistent, the authors recommend damping the short-term trend toward the long-term trend (Rule 3.3.4). I understand AGG's reasoning, but I believe more research is needed on this problem. I know of no evidence to support damping toward the long-term trend, and I can think of many examples in which the forecasts could go wrong. Long-term trends do not keep on forever, and a change will start with what appears to be a short-term trend that departs from the long-term. The product life cycle curve is a case in point. The trend changes dramatically at the end of the growth phase of the cycle, and damping back to the long-term trend will lead to over-forecasting.

4. A model-based approach to trend damping

Use a more objective, model-based approach to trend damping. Simply fit the damped trend to the data and let the model determine how the trend should be extrapolated. The damped trend is constructed by adding a damping parameter ranging from 0.0 to 1.0 to modify the trend component in Holt's (2004) linear trend method of exponential smoothing. You can think of the damping parameter as a rough measure of the strength of the trend. If the data contain a strong trend with little noise, the parameter will be estimated at a value near 1.0 during the model-fitting process, and the forecasts will be about the same as a linear trend. If the data are noisy or if the trend is erratic, the parameter will be less than 1.0, and the trend will decline in both relative and absolute terms each period in the future. Eventually the forecasts will bend toward a horizontal straight line or a saturation point. In messy data, the parameter will be much less than 1.0, and any trend in the forecasts will be negligible. As a rule of thumb, whenever the parameter is fitted at less than about 0.50, one might as well forecast with simple exponential smoothing.

Another interpretation of how the damped trend works is based on Brown's (1963) original thinking about how time series are generated. Brown looked at a time series as a sequence of local segments of data. Within each segment, the trend is constant, but the trend can change from one segment to the next, and the change can be sudden. That is, the trend frequently restarts with a new slope (which could be zero). The mean run length before the trend changes is computed as follows: damping parameter / (1 − damping parameter). This ratio is an indicator of the risk in trend extrapolation. When the damping parameter is fitted at 0.5, the mean run length is only one period, and it would be foolish to extrapolate any trend. When the parameter is fitted at 0.9, the mean run length is 9 periods, so there is less risk in trend extrapolation.

In Gardner and McKenzie (2011), we tested automatic model-fitting using the 3003 time series in the M3 competition and obtained very conservative results. Some degree of trend damping occurred in 84% of the time series; in about half of those series, damping was so powerful that the trend in the forecasts was negligible. The model allowed strong trends in the forecasts (with the damping parameter = 1.0) only about 12% of the time. Another finding in this study is that the parameters found during model-fitting often define a special case of the damped trend rather than the damped trend itself. For example, one important special case is simple exponential smoothing with a damped drift term rather than a smoothed trend. Simple smoothing with damped drift occurred in 25% of the series; this method describes a fixed early trend that gradually dies out, behavior that may seem odd but is quite common in the M3 series. We also found that a random walk model with damped drift occurred in 8% of the series.

5. The damped trend as a benchmark

The damped trend will not always beat other forecasting methods, and it may under-forecast in some time series, especially strongly trending series. However, on average the damped trend is the single most accurate extrapolation method. The damped trend should do well with high levels of noise and with inconsistent trends, conditions that probably describe most business data. The damped trend is conservative, and I recommend it as a benchmark for evaluating other methods and for making subjective decisions about the effects of causal forces, amount of variability, length of the forecast horizon, and directions of short and long-term trends.

AGG's conservatism checklist should be a valuable tool in practice. I hope that the checklist becomes a tool for researchers as well because it is actually a research agenda, with many opportunities for empirical testing of the guidelines for which evidence is limited. I would add one more rule to the extrapolation guidelines: Run the damped trend before doing anything else.

References

Acar, Y., & Gardner, E. S., Jr. (2012). Forecasting method selection in a global supply chain. *International Journal of Forecasting*, *28*(4), 842–848.
Armstrong, J. S. (1978). *Long-range forecasting.* New York: Wiley-Interscience.
Armstrong, J. S. (2006). Findings from evidence-based forecasting: Methods for reducing forecast error. *International Journal of Forecasting*, *22*(3), 583–598.
Armstrong, J. S., Green, K, & Graefe, A (2015). Golden rule of forecasting: Be conservative. *Journal of Business Research*, *68*(8), 1717–1731 (in this issue).
Assimakopoulos, V., & Nikolopoulos, K. (2000). The theta model: A decomposition approach to forecasting. *International Journal of Forecasting*, *16*(4), 521–530.
Brown, R. G. (1963). *Smoothing, forecasting and prediction of discrete time series.* Englewood Cliffs, NJ: Prentice-Hall.
Fildes, R., Nikolopolus, K., Crone, S., & Syntetos, A. (2008). Forecasting and operational research: A review. *Journal of the Operational Research Society*, *59*(9), 1150–1172.
Gardner, E. S., Jr. (1986). *Autocast exponential smoothing system.* Bridgewater, NJ: Core Analytic.
Gardner, E. S., Jr. (1990). Evaluating forecast performance in an inventory control system. *Management Science*, *36*(4), 490–499.
Gardner, E. S., Jr. (2006). Exponential smoothing: The state of the art—Part II. *International Journal of Forecasting*, *22*(4), 637–666.
Gardner, E. S., Jr., & Diaz-Saiz, J. (2008). Exponential smoothing in the telecommunications data. *International Journal of Forecasting*, *24*(1), 170–174.
Gardner, E. S., Jr., & McKenzie, E. (1985). Forecasting trends in time series. *Management Science*, *31*(10), 1237–1246.
Gardner, E. S., Jr., & McKenzie, E. (1988). Model identification in exponential smoothing. *Journal of the Operational Research Society*, *39*(9), 863–867.
Gardner, E. S., Jr., & McKenzie, E. (1989). Seasonal exponential smoothing with damped trends. *Management Science*, *35*(3), 372–376.
Gardner, E. S., Jr., & McKenzie, E. (2011). Why the damped trend works. *Journal of the Operational Research Society*, *62*(6), 1177–1180.
Gilchrist, W. (1976). *Statistical forecasting.* London: Wiley.
Holt, C. C. (2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, *20*(1), 5–10.
Hyndman, R. J., & Billah, B. (2003). Unmasking the theta method. *International Journal of Forecasting*, *19*(2), 287–290.
Hyndman, R., Koehler, A., Ord, J., & Snyder, R. (2008). *Forecasting with exponential smoothing.* Berlin: Springer.
Li, Q., Disney, S., & Gaalman, G. (2014). Avoiding the bullwhip effect using damped trend forecasting and the order-up-to replenishment policy. *International Journal of Production Economics*, *149*, 3–16.
Makridakis, S., & Hibon, M. (2000). The M3-competition: Results, conclusions and implications. *International Journal of Forecasting*, *16*(4), 451–476.
Schnaars, S. P. (1989). *Megamistakes.* New York: Free Press.