Invited Comments on "Exponential Smoothing:

The State of the Art - Part II" by E.S. Gardner, Jr.

James W. Taylor
Saïd Business School, University of Oxford

International Journal of Forecasting, 2006, Vol. 22, pp. 671-672.

Address for Correspondence:

James W. Taylor
Saïd Business School
University of Oxford
Park End Street
Oxford OX1 1HP, UK

Tel: +44 (0)1865 288927
Fax: +44 (0)1865 288805
Email: james.taylor@sbs.ox.ac.uk

1. Introduction

Gardner's 1985 review has been a tremendous help to my research. For many years, I wouldn't leave home without it. I have no doubt that the sequel will also be of lasting use to me, as well as to numerous others, especially those starting out on research in this area.

The updated review provides comprehensive coverage of a large and rapidly growing literature. A fair number of these papers are rather mathematical, and Gardner has done a sterling job of summarising their essential message and contribution. Comment and opinion is present throughout, which is surely beneficial as it provides us with the perspective of someone who seems to have read pretty much every paper that has ever been written on the subject. I could merrily comment on numerous sections of the paper that I find interesting, but, instead, I have opted to say a few words about the two topics highlighted at the very end of the paper, method selection and empirical validation, and then to describe a couple of finance applications that were touched on only briefly: volatility and quantile forecasting.

2. Method Selection and Empirical Studies

Given that in 1985 Gardner emphasised the need for guidance in selecting among methods, it is surprising that so little has been done in this area. My own experiments on method selection have left me frustrated, as I find aggregate selection of the damped additive trend method hard to beat. Empirical studies, dealing with trending data, really should include this method as a benchmark. I don't suppose the damped multiplicative trend method will prove to be so broadly useful, but it would certainly be interesting to see some more empirical results, as I have only applied it to the M3-Competition data.

Selection between methods, for individual series, in an automated fashion over time, could be termed 'switching', and a natural alternative to this is smooth transition between methods. The switching and smooth transition combining methods of Deutsch et al. (1994)

would seem to be relevant here. Indeed, Gardner's call for more attention to be given to method selection seems to me to prompt research into combining methods, specifically designed for exponential smoothing. A further thought regarding selection, switching and smooth transition is that it may result in a rather complex predictive distribution. Implementation of a simpler method might then be justified on the basis that its predictive distribution is more easily estimated, even if its point forecast accuracy is not the best.

3. Volatility Forecasting

The empirical finding that series of returns often exhibit volatility clustering has led to the development of a variety of univariate methods for forecasting conditional volatility. Along with GARCH models, exponential smoothing is commonly used. In addition to the appeal of simplicity, exponential smoothing has been shown to compete well with alternatives in terms of accuracy (see Poon and Granger, 2003). It is presumably these features that led to it being recommended in the RiskMetrics technical document (RiskMetrics, 1996).

The use of exponential smoothing for volatility forecasting is rather different to the exponential smoothing applications reviewed by Gardner. With daily log returns, r_t, the mean is very often assumed to be zero or a small constant value, and attention turns to predicting the variance, σ_t^2. As variance is unobservable, exponential smoothing is applied to the squared returns or to the squared residuals, where the residual is defined as the return minus the mean. With no apparent trend, and little seasonality, the simple exponential smoothing method has been used:

$$\hat{\sigma}_{t+1}^2 = \alpha\, r_t^2 + (1-\alpha)\,\hat{\sigma}_t^2$$

The variance forecast for the return over a holding period of h days is given as the one-day forecast multiplied by h. This can be quite different to the corresponding multiperiod forecast produced by GARCH models because the latter tend to be mean-reverting. The simple exponential smoothing method can be viewed as a special case of the *integrated* GARCH model (IGARCH), which is a non-stationary version of GARCH. Smooth transition exponential smoothing for volatility forecasting is a simple development of the smooth transition GARCH models (see Taylor 2004a). There is surely scope for transferring other ideas from the voluminous GARCH literature to exponential smoothing. The increasing use of intraday data also presents opportunities for the use of more ambitious versions of exponential smoothing.

The conditional covariance, $\sigma_{12,t}$, between two series of returns, r_{1t} and r_{2t}, can also be estimated using exponential smoothing:

$$\hat{\sigma}_{12,t+1} = \alpha\, r_{1t} r_{2t} + (1-\alpha)\,\hat{\sigma}_{12,t}$$

This approach has been popular, as it is somewhat simpler to implement than GARCH-based covariance estimation. To ensure that the covariance matrix is positive semi-definite, RiskMetrics proposes the use of a common value of 0.06 for α in the estimation of all elements in the matrix. This arbitrarily chosen value has often been used regardless of whether or not a full covariance matrix is being estimated.

4. Quantile Estimation

Assessing value at risk (VaR) amounts to estimating tail quantiles of the conditional distribution of a series of financial returns. As with volatility, the unobservable nature of quantiles means that their prediction is not straightforward. Boudoukh et al. (1998) propose a form of exponentially weighted average for quantile estimation. It involves allocating to the most recent year of daily returns, exponentially decreasing weights, which sum to one. The returns are then placed in ascending order and, starting at the lowest return, the weights are summed until a value of θ is reached. The θth quantile estimate is set as the return that corresponds to the final weight used in this summation. Linear interpolation is used if the estimate falls between two returns. Dunsmuir et al. (1996) essentially use the same method to estimate the median in a robust approach to point forecasting for non-financial data. The method of Boudoukh et al. is equivalent to the simplest case of exponentially weighted quantile regression (EWQR), which amounts to exponential smoothing of the cumulative distribution function (see Taylor, 2004c). There would seem to be the potential for other considerations of exponential smoothing for forecasting conditional quantiles, and, indeed, conditional probability densities.

Additional references to those in Gardner (2006)

Boudoukh, J., M. Richardson, R.F. Whitelaw. 1998. The best of both worlds. *Risk* **11** May 64-67.

Deutsch M., Granger C.W.J., Teräsvirta, T. (1994). The combination of forecasts using changing weights. *International Journal of Forecasting*, 10, 47-57.

Dunsmuir, W.T.M., Scott D.J., Qiu, W. (1996). The distribution of the weighted moving median of a sequence of iid observations, *Communications in Statistics: Simulation and Computation*, 25, 1015-1029.

Poon, S., C.J.W. Granger. (2003). Forecasting volatility in financial markets: A review, *Journal of Economic Literature*, 41, 478-639.

RiskMetrics. (1996). *RiskMetrics*™ *Technical Document*, fourth edition, J.P. Morgan/Reuters.