

USE OF ERRORS OF PREDICTION IN IMPROVING FORECAST ACCURACY: AN APPLICATION TO WOOL IN AUSTRALIA

B. M. S. LEE and A. BUI-LAN*

Bureau of Agricultural Economics, Canberra, A.C.T. 2601

Over the past 20 years, the BAE has produced commodity forecasts in *BAE Trends in Australian Agricultural Commodities, Farm Costs and Farm Incomes* and other commodity outlook publications. There have been several studies on the accuracy of the forecasts, either comparing BAE forecasts with naive forecasts or making simple calculations of biases and standard deviations. Such approaches are not sufficiently informative in that they do not give any indication of what has caused the inaccuracies, nor of the timing of significant changes in prediction accuracy. In this analysis, an attempt is made to improve upon this situation using some tools of time series analysis.

Introduction

Predicting the short-term outlook for Australian agriculture is one of the objectives of the BAE. From 1952 through 1975, first estimates of gross revenues for each financial year for the major agricultural commodities appeared in the September issues of the quarterly publication *BAE Trends in Australian Agricultural Commodities, Farm Costs and Farm Incomes (Trends)*. Since June 1975, these forecasts have been reported in the June issue of *Trends* in order to provide earlier information for the users. Although the BAE has used econometric models and time-series models (such as simple trend extrapolation) to generate commodity forecasts, the forecasts are adjusted subjectively by commodity analysts before release.

Freebairn (1975) was the first to analyse systematically the accuracy of BAE forecasts. Using a nonparametric sign test, Freebairn concluded that BAE forecasts of production and price for several commodities, including wool, were significantly more accurate than the mechanistic forecasts based on a second-order autoregressive model. However, the 'price' series considered by Freebairn was simply a unit-value series derived as the quotient of the actual (or forecast) revenues and quantities during the period 1966 through 1973. Also, by the time the first estimates appeared in September, almost one-quarter of the financial year to be forecast had been observed.

The objectives in this study are to ascertain the level of accuracy of past BAE forecasts, to identify factors that might have led to inaccuracies in the past and to suggest ways of improving the current forecast. One of the common approaches to this task has been to compare various measures of forecasting accuracy for the series with other mechanical forecasting procedures. However, we believe that such an approach falls short of pin-pointing the possible weaknesses of a forecasting system and may not lead to improved forecasting procedures.

* We would like to thank Jim Longmire, Henry Haszler, Doug Williams and other colleagues at the BAE for helpful comments. All remaining errors are ours.

Since the predictive errors of a particular forecast normally contain information about the adequacy of the forecasting method, an analysis of the predictive errors of past BAE forecasts is useful and cumulative sum (cusum) techniques (Harrison and Davies 1964) will be used for this purpose. If the predictive errors do not appear to be purely random (i.e. zero-mean white noise), then an Autoregressive and Moving Average (ARMA) model will be identified and estimated to predict the error of the current forecast. A revised forecast will then be made.

Properties of an Optimal One-Step-Ahead Linear Prediction

Let $(X_t, t = 1, 2, \dots, T)$ be a univariate, discrete time, measurable time series which is to be forecast on the basis of a specific set of information, A_n , available at time n . It would be impossible to find an optimal (minimum variance unbiased) forecasting procedure without further restrictions on the information set or on the loss function (criterion of optimality). However, if we consider only quadratic loss functions (which may not be strictly appropriate but are mathematically convenient) and restrict the information set to consist only of past values of the series being forecast (i.e. $A_n = (X_{n-j}, j \geq 0)$), then it can be shown that the expected value of the target variable, conditional on past values of that variable, is the optimal one-step ahead linear predictor (Granger and Newbold 1977, p. 121). That is:

$$(1) \quad P_n = E(X_{n+1} | A_n),$$

where P_n = predictor at time n for the value of X at time $n + 1$;

A_n = set of past history of X up to and including time n ; and

$E(\)$ = conditional mathematical expectation.

If X_t is a zero-mean, stationary and purely non-deterministic series with finite variance, then it has an infinite moving average representation:

$$(2) \quad X_t = \sum_{j=0}^{\infty} C_j e_{t-j},$$

where e_t = a zero-mean white noise process with finite variance σ^2 ; and

C_j = any real number with $C_0 = 1$.

The predictive error ($u_n = P_n - X_n$) of the optimal predictor can be shown to follow a zero-mean white noise process and one can, therefore, expose the sub-optimality of a forecasting system by showing that the sequence of predictive errors fails to be a zero-mean white noise process. In this paper, cusum techniques are employed to reveal the bias of BAE forecasts and to correct for it in preparing subsequent forecasts.

Cusum Analysis of Predictive Errors

The departure of the mean of predictive errors from zero can be detected by means of the Shewhart Control Chart. This chart plots the predictive errors over time with control limits imposed at two (or three) standard deviations above and below the target horizontal line of zero. Any point that lies outside the limits signals an out-of-control situation.

However, in this Shewhart Control Chart approach, information that there may be runs of forecast errors which travel on only one side of the mean level, suggesting bias of the forecasting system, is neglected. In

cusum analysis, this phenomenon is taken into account and thus, the analyst is able to detect any inherent bias much faster than if the control chart method is used. Theoretically, cusum analysis is the sequential test applied in reverse.

Let $S_i = \sum_{r=1}^i u_r$, $i = 1, \dots, T$ be the series of cumulative sums of predictive errors which is plotted, instead of the individual predictive errors, against the sequence number $i = 1, \dots, T$. Note that each sum is obtained from its predecessor by the simple operation of adding on the new predictive error (i.e. $S_i = S_{i-1} + u_i$). If the mean value of the predictive errors remains close to zero, some of them will be positive and some negative so that the cusum chart will be essentially horizontal and the forecasting system is said to be unbiased. However, if the average value of the process rises to a new constant level, more of the predictive errors will become positive and the cusum chart will be a straight line sloping upwards. Similarly, if the average value of the process falls to a constant level below zero, the general slope of the chart will be downwards. Cusum charts are thus interpreted by the average slopes of the line which is graphed. The further the current mean process level is from zero, the steeper will be the slope of the cusum chart. In fact, the slope of the line joining the a^{th} and b^{th} points ($b > a$) on the chart is given by:

$$(3) \quad (S_b - S_a)/(b - a) = \sum_{r=a+1}^b u_r / (b - a),$$

which is the average bias between the time interval (a, b) . Thus the chart indicates quickly bias over any number of sequential forecasts.

Since the visual picture depends to some extent on the scales chosen for the axes of the chart, a convention is adopted so that one unit on the horizontal distance between two plotted points equals 2σ on the vertical scale. With this system of scaling, the mean path of the chart will make an angle of $\pi/4$ radians (or 45 degrees) with the horizontal when the mean value of the process is 2σ away from zero. It is generally desirable that no slope should exceed $\pi/3$ radians (or 60 degrees) since angles greater than this are relatively insensitive to changes in the mean values of the series.

Corresponding to the critical lines of the sequential analysis, a similar rule is needed to decide when a major change of slope has occurred on the cusum chart. One method for making these decisions was proposed by Barnard (1959) and involved superimposing a V-shaped mask over the chart. The axis of the V is set in a horizontal position with the vertex of the V at a distance d ahead of the most recent point on the chart. If any of the two limbs of the V are crossed, the slope of the cusum measured from the most recent value must exceed $\tan \theta$, where θ is the angle between each of the limbs of the V and the horizontal. The parameter $\tan \theta = 0.37$ with $d = 2$ corresponds to the 2-standard deviation control chart (Woodward and Goldsmith 1964, p. 15).

It should be noted that, while the cusum chart reveals quickly any significant runs of bias in the forecasting mechanism, statistical tests based on cusums are not the most powerful. This is due to the intractability of the distribution of cusums and the difficulty in assessing the significance of the departure of the cusum path away from the target of

zero. A complementary and more powerful test statistic (*BDE* test) proposed by Brown et al. (1975) was based on the cusum of squares of the predictive errors. That is, the quantities:

$$(4) \quad S_r = \frac{\sum_{j=1}^r u_j^2}{\sum_{j=1}^T u_j^2}, \quad r = 1, \dots, T,$$

are plotted and tested against a Beta distribution with mean r/T under the null hypothesis that the u_j s are independent zero-mean white-noise of the Gaussian type. The power of this *BDE* test against a number of alternatives has been demonstrated in McCabe and Harrison (1980).

Monitoring the BAE Forecasts on Prices and Production of Wool

In this section, we illustrate the use of cusum techniques to monitor predictive errors. The BAE forecasts of wool prices and production are examined. Because of the relative importance of wool in the Australian agricultural sector, the accuracy of the BAE's forecasts for wool are likely to be of interest to analysts.¹

Wool prices (1957-58 through 1979-80)

Actual and predicted greasy wool prices (Table 1) are plotted in Figure 1 and a conventional Shewhart chart for monitoring the predictive errors




FIGURE 1—Actual and predicted values—wool prices

¹ While other agricultural commodities, such as beef, should also be monitored we have not included them here because to do so would have lengthened the paper considerably and served no further illustrative purpose.

TABLE 1
Average Auction Prices for Greasy Wool^a

t	Financial year	Actual (1)	Forecast (2)	Predictive error (3) = (2) - (1)	Cusum (4)	Relative cusum of squared errors (5)
		c/kg	c/kg	c/kg		
1	1957-58	114.7	147	32.3	32.3	0.091
2	1958-59	89.2	96	6.8	39.1	0.096
3	1959-60	106.2	107	0.8	39.9	0.096
4	1960-61	95.6	97	1.4	41.3	0.096
5	1961-62	99.5	101	1.5	42.8	0.096
6	1962-63	108.3	99	-9.3	33.5	0.104
7	1963-64	128.0	107	-21.0	12.5	0.142
8	1964-65	105.5	119	13.5	26.0	0.158
9	1965-66	110.4	103	-7.4	18.6	0.163
10	1966-67	104.5	101	-3.5	15.1	0.164
11	1967-68	92.0	85	-7.0	8.1	0.168
12	1968-69	98.5	81	-17.5	-9.4	0.195
13	1969-70	82.8	79	-3.8	-13.2	0.196
14	1970-71	64.7	64	-0.7	-13.9	0.197
15	1971-72	75.3	66	-9.3	-23.2	0.204
16	1972-73	183.8	105	-78.8	-102.0	0.749
17	1973-74	181.2	185	3.8	-98.2	0.750
18	1974-75	127.0	126	-1.0	-99.2	0.750
19	1975-76	143.3	130	-13.3	-112.5	0.765
20	1976-77	182.7	150	-32.7	-145.2	0.859
21	1977-78	187.2	185	-2.2	-147.4	0.860
22	1978-79	205.2	195	-10.2	-157.6	0.869
23	1979-80	243.7	205	-38.7	-196.3	1.000
24	1980-81		245			

^a The data in columns (1) and (2) are from BAE (1981).

is provided in Figure 2. Although there is a strong indication of under-prediction in the financial year 1972-73, the predictive errors in other years appear to be quite close to the target of zero. The source of the error in 1972-73 was, according to research officers in the BAE, the oil price boom associated with a period of strong demand for wool, an unexpected fall in wool production due to drought and market intervention by the Australian Wool Corporation (AWC).

A *V*-mask has been superimposed onto the cusum chart in Figure 3. The cusum path cuts the upper arm of the *V*-mask from above at three points, namely, 1971-72, 1975-76 and 1977-78, indicating significant deviation from a zero-mean white-noise process. This is also confirmed by using a cusum of squares test. The hypothesis that the predictive errors are independent normal variates is rejected at the one per cent level, and the cusum squares chart (Figure 4) indicates that significant changes occurred after 1971-72.

From the cusum chart in Figure 3, three changes in slope can be identified. Between 1957-58 and 1961-62, there was an average over-prediction of about 8.6c/kg (since the cusum increases from zero to 42.8c/kg over a five-period interval). From 1962-63 to 1971-72, there was an average under-prediction of about 6.6c/kg (it decreases from




FIGURE 2—Shewhart chart for predictive errors—wool prices




FIGURE 3—Cusum chart (with V-mask) for predictive errors—wool prices




FIGURE 4—Cusum of squares of predictive errors—wool prices

42.8c/kg to -23.2c/kg over a ten-period interval). The accuracy of the BAE forecasts apparently deteriorated after 1971-72 as the cusum decreases from -23.2c/kg to -196.3c/kg over an eight-period interval, giving an average under-prediction of about 21.6c/kg. The explanations offered by BAE officers for such a deterioration include the higher inflation in the 1970s, market intervention by the AWC which began in the early 1970s, and the unexpected oil and commodity boom in the late 1970s.

Further evidence for refuting the hypothesis that the predictive error of BAE forecasts for wool prices follows a zero-mean, white-noise process is based on an examination of the correlograms of the autocorrelation function (ACF) and the partial autocorrelation function (PACF). It was found that both the ACF and PACF have peaks at lags 4 and 7. After experimenting with several ARMA models, the following model was found to be the most satisfactory (standard errors in parentheses):

$$u_t = -6.01 + e_t - 0.33e_{t-4} - 0.47e_{t-7},$$

$$\begin{matrix} (6.10) & (0.21) & (0.21) \end{matrix}$$

$$Q(8) = 1.09$$

where u_t = prediction error at time t ;

e_t = white-noise term at time t ; and

$Q(8)$ = the portmanteau statistic which equals the sum of squares of the first eight autocorrelations of the residuals and should be compared with a chi-square distribution with three degrees of freedom.

Based on this model, the BAE forecast for 1980-81 is estimated to under-predict the actual wool prices by about 5.6c/kg. Thus, it is suggested that the forecast be revised upward from 245c/kg to 250c/kg. The probability is about 70 per cent that the actual wool prices will fall within the range 232c/kg–269c/kg. This revised forecast is preferable since the track record of BAE wool-price forecasts is such that they tend to be under-predicted.² Note that such adjustment is relevant at this particular time point only and is not intended for all future BAE forecasts of wool prices because the structure of the predictive error process may change in subsequent years.

Wool production (1957-58 through 1979-80)

It is difficult to forecast production by wool types and the BAE, in conjunction with other representatives on the Australian Wool Production Forecasting Committee, has confined its attention to total wool production in greasy equivalent. In Figure 5, actual and predicted wool production based on Table 2 for the years 1957-58 through 1979-80 are shown.

The BAE forecast appears to be quite accurate in predicting production and the Shewhart control chart (Figure 6) fails to detect any significant bias. Even the more powerful cusum and cusum of squares graphs (Figures 7 and 8) do not indicate any systematic deviation of the predictive errors from a zero-mean, white-noise process.




FIGURE 5—Actual and predicted values—wool production

² Theil's (1971) U_2 -coefficient for the price series is about 0.70, indicating that BAE forecasts are better than a naive no-change forecast in predicting actual changes (ΔA_t). Note that this is more informative than simply counting turning-points errors which has loss 1 if (ΔP_t) (ΔA_t) is negative and zero otherwise, whereas for the U_2 statistic the loss is proportional to $(\Delta P_t - \Delta A_t)^2$. Thus, a large turning-point error suffers more loss than a small turning-point error in the U_2 statistic.

TABLE 2
Annual Production of Wool (greasy equivalent)^a

t	Financial year	Actual (1)	Forecasts (2)	Predictive errors (3) = (2) - (1)	Cusum (4)	Relative cusum of squared errors (5)
		kt	kt	kt		
1	1957-58	650	727	77	77	0.166
2	1958-59	722	629	-93	-16	0.408
3	1959-60	762	758	-4	-20	0.408
4	1960-61	737	726	-11	-31	0.412
5	1961-62	771	767	-4	-35	0.412
6	1962-63	759	774	15	-20	0.419
7	1963-64	810	788	-22	-42	0.432
8	1964-65	809	822	13	-29	0.437
9	1965-66	754	768	14	-15	0.442
10	1966-67	799	740	-59	-74	0.540
11	1967-68	803	824	21	-53	0.552
12	1968-69	884	831	-53	-106	0.631
13	1969-70	923	906	-17	-123	0.639
14	1970-71	891	947	56	-67	0.727
15	1971-72	882	878	-4	-71	0.727
16	1972-73	735	808	73	2	0.876
17	1973-74	701	737	36	38	0.912
18	1974-75	794	772	-22	16	0.926
19	1975-76	754	788	34	50	0.958
20	1976-77	703	725	22	72	0.972
21	1977-78	677	683	6	78	0.973
22	1978-79	706	685	-21	57	0.985
23	1979-80	706	729	23	80	1.000
24	1980-81		686			

^a The data in column (1) are from ABS (1980) and in column (2) from Australian Wool Production Forecasting Committee (1980).

Nevertheless, three changes in slope can be identified here. A major under-prediction was registered in 1958-59. From 1957-58 through 1969-70, the magnitude of under-prediction was an average of 9 kt over a 13-year period, possibly due to better forecasting techniques. After 1970-71, the accuracy in forecasting production reversed, with an average over-prediction of 16 kt over a nine-year period up to 1979-80.

Although the official forecasts seemed to predict actual change fairly accurately (Theil's (1971) U_2 -coefficient = 0.69) the official forecasts managed to predict actual change fairly accurately, they tended to over-predict the level of production in recent years. As a matter of fact, the latter half of the cusum path slopes upward over time. BAE officers attributed this inaccuracy, in part, to the fact that there was a run of drought years in the 1970s. It may be that the production forecasts were adversely affected by drought years simply because they are not predictable. In view of this inherent bias, one should be able to improve the BAE forecast, *if* one could predict droughts accurately.

Alternatively, an ARMA model was estimated for the series of predictive errors. The ACF and PACF suggested that the series is stationary,




FIGURE 6—Shewhart chart for predictive errors—wool production




FIGURE 7—Cusum chart (with V-mask) for predictive errors—wool production




FIGURE 8—Cusum of squares of predictive errors—wool production

with a possible four-year cycle. After several trials, a relatively satisfactory model was found to be:

$$u_t + 0.256u_{t-4} = 4.424 + e_t - 0.385e_{t-3},$$

(0.17) (7.35) (0.25)

$$Q(8) = 0.47.$$

This model predicts that the BAE forecast for 1980-81 may be biased upward by around 6.7 kt. It is, therefore, suggested that the forecast be revised downward from 686 kt to 680 kt. The 70 per cent confidence interval would be about 650 kt to 710 kt.

The BAE forecasts for the year 1980-81 can be revised as in Table 3. It should be noted that the 70 per cent confidence interval is reported instead of the more common 95 per cent interval because the variability of forecasts is such that meaningful intervals are necessarily narrower than those with more than 95 per cent confidence. A 70 per cent interval is adequate to provide some indication of the variability of forecasts.

TABLE 3
Revisions of the BAE Forecasts

Item	BAE forecast	Suggested revisions	
		Forecast	70 per cent interval
Average wool price (greasy)	245c/kg	250c/kg	(232c/kg to 269c/kg)
Wool production (greasy equivalent)	686 kt	680 kt	(650 kt to 710 kt)

Benefits from the Analysis

Owing to data limitations, it is difficult to assess the improvements that might be made as a result of such an analysis. Also, it is expensive and unproductive to identify a new ARMA model on receipt of an additional observation to the data set. Nevertheless, it can be demonstrated that even a simple adjustment of the BAE forecasts after an 'out-of-control' situation is detected would have led to a substantial reduction in the sum of squared errors.

Suppose that the analyst has used the cusum chart and detected that the BAE forecasts for wool prices have been consistently biased downward since 1973-74. He then estimates the expected bias in 1974-75 by dividing the cusum into the number of records which form it (i.e. average bias = $-98.5/17$ or -5.8). This implies that he will revise the BAE forecast in 1974-75 upwards from 126c/kg to 131.8c/kg. Similar adjustments are made for the subsequent years (Table 4). The sum of squared errors from 1974-75 through 1979-80 of these revised forecasts is 1830.6, a reduction of 36 per cent compared with the sum of squared errors associated with the original BAE forecasts.

TABLE 4
Errors of Adjusted Price Forecasts

Year	1974-75	1975-76	1976-77	1977-78	1978-79	1979-80
	c/kg	c/kg	c/kg	c/kg	c/kg	c/kg
Adjusted price forecasts	131.8	135.5	155.9	192.3	202.0	212.2
Error	4.8	-7.8	-26.8	5.1	-3.2	-31.5

The predictive errors for wool production have a cusum graph converging to zero and no out-of-control signal is found (Figure 7). This implies that the present BAE wool production forecasting system is fairly adequate and is unlikely to be improved upon. Possibly, one could have improved the BAE forecasts of wool production in the 1960s but not in the 1970s. However, if the user mechanically adjusted the BAE forecasts after 1974-75 by the above procedure, he would have obtained the results shown in Table 5. The sum of squared errors is 3090, a reduction of 1 per cent on the sum of squared errors associated with the original forecasts. This indicates a negligible gain, as is expected from the cusum analysis.

TABLE 5
Errors of Adjusted Production Forecasts

Year	1974-75	1975-76	1976-77	1977-78	1978-79	1979-80
	kt	kt	kt	kt	kt	kt
Adjusted production forecasts	796.8	787.1	722.4	679.4	681.3	726.4
Error	-24.2	33.1	19.4	2.4	-24.7	20.4

Discussion and Conclusion

The forecast of production is more accurate than the forecast of prices. This is understandable, given the continual effort in obtaining economic survey information on supply conditions. While supply is more stable due to biological constraints, prices tend to be more unstable because they are affected by both supply and demand factors. More unstable general economic conditions during the 1970s have compounded the difficulty in forecasting wool prices.

The tendency to under-predict prices seems to indicate that BAE forecasters are conservative in making their forecasts, possibly due to a different loss function being used. When econometric models produce forecasts which represent major departures from the latest observed value, judgment is likely to be used to moderate the forecasts. Although there may be abundant prior information as to what is likely to happen in the following year, it is still quite a difficult task to incorporate properly such information into the forecasts.

The cusum technique has been used to help detect the bias in the forecasts in the last 20 years. Changes in variability of the forecast errors can also be detected if more than one forecast is made for a particular item at the same point of time. Such techniques are not infallible but are useful in providing an indication of the way in which the forecasting scheme needs to be adjusted. The probability of missing a warning of bias by this method can be shown to be much smaller than that by the conventional control charts for a given sample size. Some comparisons of the power of the cusum technique are given by Johnson (1961).

Box and Jenkins techniques, which have obvious problems when applied to small samples of observations, have been used to fit various ARMA models to the predictive error series in an attempt to refine the residuals to be as close as possible to a zero-mean white-noise process. An adequately identified and estimated model could then be used to predict the 1980-81 forecast errors which are used to revise the current BAE forecast. The revised forecast should, theoretically, be closer to that obtained from an optimal linear predictor. This adjustment would apply only to the first estimates and not to the revised estimates in later quarters of the year. Users of the wool forecasts should be warned against adjusting automatically future forecasts using the approaches given in this paper. The use of these techniques is dependent on the system currently in use remaining intact, that is, on the forecasters not having benefited from the knowledge of the 'bias' in their previous forecasts.

One can only test whether our revision would lead to any improvement of the current BAE forecasts by comparing them to the actual values next year. Nevertheless, some simple adjustments, aided by the cusum graphs, lead to an improvement of more than 35 per cent over the BAE forecasts for wool prices.

Finally, as a first step in an attempt to improve the forecasting system, all the forecast errors should at least be subject to as stringent a test as was proposed in this paper. Considerable insight could be gained from the graphs of cusum and cusum of squares of the predictive errors. In particular, the techniques have helped identify the fact that predictions of wool prices have been more difficult in the 1970s than in the 1960s.

The cusum chart reveals inherent bias in the original forecasts, which would otherwise pass unnoticed by other conventional analyses. Removal of such bias would very likely lead to improved forecasts in future movements of the wool market.

References

- ABS (1980), *Wool Statistics*, Canberra (and previous issues).
- Australian Wool Forecasting Committee (1980), Press Releases, July (and previous issues).
- BAE (1981), *Trends in Australian Agricultural Commodities, Farm Costs and Farm Incomes*, Canberra, December (and previous issues).
- Barnard, G. A. (1959), 'Control charts and stochastic processes', *Journal of the Royal Statistical Society Series B* 21(2), 239-71.
- Brown, R. L., Durbin, J. and Evans, J. M. (1975), 'Techniques for testing the constancy of regression relationships over time', *Journal of the Royal Statistical Society Series B* 37(2), 149-92.
- Freebairn, J. W. (1975), 'Forecasting for Australian agriculture', *Australian Journal of Agricultural Economics* 19(3), 154-74.
- Granger, C. W. J. and Newbold, P. (1977), *Forecasting Economic Time Series*, Academic Press, New York.
- Harrison, P. J. and Davies, O. L. (1964), 'The use of cumulative sum (CUSUM) techniques for the control of routine forecasts of product demand', *Operations Research* 12(2), 325-33.
- Johnson, N. S. (1961), 'A simple theoretical approach to cumulative sum charts', *Journal of the American Statistical Association* 56(296), 835-40.
- McCabe, B. P. M. and Harrison, M. J. (1980), 'Testing the constancy of regression relationships over time using least squares residuals', *Applied Statistics* 29(2), 142-8.
- Theil, H. (1971), *Applied Economic Forecasting*, North Holland, Amsterdam.
- Woodward, R. H. and Goldsmith, P. L. (1964), *Cumulative Sum Techniques*, ICI Monograph No. 3, Oliver and Boyd, Edinburgh.