

# Forecasting the failure of component parts in computer systems: A case study

Everette S. Gardner, Jr\*

*Center for Global Manufacturing, University of Houston, Houston, TX 77204-6282, USA*

## Abstract

This paper is an applied study in forecasting the failure of component parts in computer systems to aid in production planning and inventory control. The aim is to develop a reasonably simple forecasting system that can be operated by managers rather than statisticians. Monthly failures of components are shown to be related to cumulative shipments from the factory after a time lag of several months. This relationship is modelled using discounted-least-squares regression, a methodology that appears to be rare in practice. Simulated forecast accuracy is then compared with two alternatives, exponential smoothing and a combination of regression and smoothing forecasts. Discounted regression proves to be the most accurate and is therefore implemented.

**Keywords:** Regression analysis; Combining forecasts; Monitoring forecasts; Exponential smoothing

## 1. Introduction

The setting for this study is a small manufacturing company in the computer industry. The company supplies a variety of component parts, such as disk drives and controllers to customers who assemble or enhance computer systems for end users. These customers include computer dealers and manufacturers, corporations, and some large government agencies. Routine sales forecasting is a minor problem not considered here because most production is based on contracts placed well in advance. However, forecasting the number of components that will fail once they are installed and put into service and the timing of these failures are important problems. To maintain its competitive position, the company must stand ready to replace all failed components, usually on very short notice.

The forecasting model proposed for this application is a discounted-least-squares (DLS) regression to predict monthly component failures as a function of cumulative shipments over time. To benchmark forecast accuracy, comparisons are made between two alternatives: (1) exponential smoothing, perhaps the most reasonable alternative model of the data; and (2) a combination of forecasts from regression and exponential smoothing.

In this paper, Section 2 examines the company's production and component failure data. Section 3 discusses the considerations involved in model selection. In Section 4, the results of forecasting tests are presented. Forecast monitoring problems are discussed in Section 5. Finally, conclusions and implementation of the forecasting system are discussed in Section 6. This research can be replicated. Copies of all data in disguised form are available from the author.

\* Tel: (713) 743-4720; fax: (713) 743-4807.

## 2. Data analysis

The company provided monthly time series production and failure data for nine components. Five components were selected for analysis, with time series of the following lengths: 24, 27, 29, 37, and 48 months, respectively. The components are identified simply by the letters A–E. An agreement between the author and the company prohibits disclosure of information on the nature of the components. One component was eliminated from the study because the data were absurdly ill-behaved, containing numerous jump-shifts up and down in the level of failures due to manufacturing problems. Management agreed to treat this component as a special case, for which judgmental forecasting was necessary. Three other components were not analyzed because they had been in production for periods ranging from 6 to 13 months.

At the beginning of the study, management hoped that monthly shipments of components would be of some assistance in predicting monthly failures several months later. Unfortunately, correlation analysis showed that this idea was not feasible. To illustrate, Fig. 1 plots monthly failures in period  $t$  versus monthly shipments in period  $t - 3$  for component C. There is no significant relationship when using this lag structure, and the same was true for all other components and lag structures.

However, a strong relationship was found between monthly failures and *cumulative shipments*

of components. This relationship makes some sense because samples showed that the failed components in a given month could be traced to a wide range of previous production dates. In Fig. 2, monthly failures in period  $t$  are plotted versus cumulative shipments through period  $t - 3$  for component C. The correlation coefficient between these variables is 0.92. For each of the other components, a coefficient of at least 0.90 was found using the same variables.

Figure 2 suggests the use of some form of regression model to predict monthly failures, using cumulative shipments as the independent variable. A time series model is also a possibility for this data. Figure 3 shows a time series plot of monthly failures for component C. The pattern in this graph is typical of components A–D. Each time series in this group contains a definite trend with a large random component and no evidence of any change in structure. However, as shown in Fig. 4, the data for component E (48 observations) does display a marked change in structure. Production slowed during the third year, and finally ceased in month 42. Growth in monthly failures disappeared during the third year, while later data fluctuate around a constant level. Based on experience of the industry, management believed that the constant-level structure would continue for some time, followed by a gradual decline in failures as the computers in which the component was installed became obsolete. Management also expected the other components to eventually behave like this.




Fig. 1. Component C. Monthly failures in period  $t$  versus monthly shipments in period  $t - 3$ .




Fig. 2. Component C. Monthly failures in period  $t$  versus cumulative shipments as of period  $t-3$ .




Fig. 3. Component C. Time series plot of monthly failures.




Fig. 4. Component E. Time series plot of monthly failures. Production ceased in month 42.

Although the time series are too short for a conclusive statistical analysis, no evidence of a seasonal pattern was found in any of them. Management also did not believe that there was any reason to expect a seasonal pattern.

### 3. Model development

One approach to modelling component failures relies on an engineering estimate of the mean-time-between-failures (MTBF) for a component. Given the probability law followed by the MTBF or some function of the MTBF, predictions of failure rates can be made for any time period. However, management ruled out any attempt to develop a model using MTBF estimates. A previous consultant had attempted such a model, but management found the predictions to be useless. The problem was that the true MTBF changes continuously, as design, manufacturing, and quality control improve with experience. Thus replacement parts often have different failure characteristics than the originals. A related problem was that company engineers found it difficult to quantify failure characteristics, given the extremely complex interactions between the components and the variety of hardware and software configurations in which they are employed.

The correlation results suggest a regression approach, at least while a component is in production, with cumulative shipments as an explanatory variable for monthly failures several months later. A time lag of 3 months between cumulative shipments and failure was selected in the regression. There was no need to consider shorter time lags, because historical data showed that at least 3 months elapse between shipments and initial failure reports. Time lags longer than 3 months served only to reduce the excellent correlation coefficients discussed above. It is important to emphasize that there is no need to forecast more than one step ahead in this application. The regression makes each prediction using cumulative shipments data recorded 3 months earlier. Thus, a one-step-ahead forecast gives the company a total of 4 months of lead-time to adjust production plans.

There is justification for a DLS regression rather than the usual ordinary-least-squares

model. It is reasonable to believe that time to failure of components is exponential, so the number of failures depends directly on the number of components in use (cumulative shipments) and the total failure rate. The latter is unlikely to be a constant, the assumption made in an ordinary-least-squares model. Instead, the total failure rate should decrease over time, making DLS regression more appropriate for tracking changes.

By following the recursive updating scheme in Gilchrist (1976), the DLS regression was easy to program in a spreadsheet. All personnel involved in production planning were familiar with spreadsheets, which simplified implementation of the model. Managers and production planners were able to answer many of their own questions about DLS by tracing the spreadsheet formulas and performing sensitivity analysis on model parameters.

An alternative forecasting model is needed to benchmark accuracy. Given the short time series available, some form of exponential smoothing appears to be the only practical alternative. For reasons explained below, the damped-trend smoothing model of Gardner and McKenzie (1985) was selected. This model is written:

$$S_t = S_{t-1} + \phi T_{t-1} + h_1 e_t \quad (1)$$

$$T_t = \phi T_{t-1} + h_2 e_t \quad (2)$$

$$\hat{X}_t(m) = S_t + \sum_{i=1}^m \phi^i T_t \quad (3)$$

$S$  is the local level of the series,  $T$  is the local trend estimate, and  $X$  is the forecast, made at origin  $t$  for  $m$  steps ahead. The smoothing parameters are  $h_1$  and  $h_2$ , chosen independently. The trend estimate is modified with an autoregressive-damping parameter,  $\phi$ . Although the model is stable over a wide range of parameters, all three parameters are usually constrained to the range 0–1.

The reason for selecting this model is that it includes several special cases which often lead to greater forecast accuracy than when the usual assumption of a linear trend in every time series is made. For example, when the trend is erratic, model-fitting yields  $0 < \phi < 1$ , so that growth in the forecasts is damped. In series with stronger trends, model-fitting yields  $\phi = 1$ , and the fore-

casts are identical to those from Holt's linear trend model. For examples of these results with empirical data, see Gardner and McKenzie (1985).

Another alternative is to combine the forecasts from the DLS and smoothing models. A simple average of forecasts was used, because Clemen's (1989) review concludes that there is little evidence in favor of more sophisticated methods of combining forecasts. See also the arguments by Armstrong (1989) and Makridakis (1989) on this point. Another reason for using a simple average is that the time series are probably too short to support the analysis required by more sophisticated methods.

#### 4. Forecasting tests

Forecast accuracy results for both mean-absolute-percentage-error (MAPE) and root-mean-

squared-error (RMSE) criteria are summarized in Table 1. RMSE rather than MSE was computed because the RMSE approximates the standard deviation of the forecast errors, an important measure for inventory control of the components.

Series E will be used to explain how Table 1 was developed. First, the DLS and smoothing models were fitted to months 1–12. The smoothing model was initialized by fitting an ordinary-least-squares trend to the data for months 1–12, using time as the independent variable. The initial trend was set equal to the slope, with the initial level set equal to the intercept. Following the fit, model parameters were retained, and forecasts and errors were computed during months 13–48. The forecast periods refer to the months during which a one-step-ahead forecasting simulation was conducted. At the end of each month during the forecast periods, the DLS and smoothing forecasts were averaged to obtain

Table 1  
Forecast accuracy comparisons

Series	Forecast periods	MAPE			RMSE		
		Regression	Smoothing	Combined	Regression	Smoothing	Combined
A	13–24	<b>7.6</b>	11.9	9.0	<b>26.0</b>	43.4	32.8
	17–24	<b>6.1</b>	11.0	7.6	<b>22.7</b>	45.3	32.2
	21–24	<b>5.3</b>	6.4	5.8	23.0	23.3	<b>22.7</b>
B	13–27	<b>9.8</b>	10.4	10.1	64.5	65.0	<b>64.1</b>
	17–27	<b>9.0</b>	10.2	9.6	66.3	<b>64.2</b>	64.7
	21–27	<b>9.3</b>	12.0	10.6	<b>67.4</b>	72.5	68.9
	25–27	<b>7.7</b>	14.7	11.2	<b>55.6</b>	84.8	68.2
C	13–29	6.1	8.3	<b>5.8</b>	23.4	29.5	<b>22.9</b>
	17–29	<b>6.0</b>	7.1	6.4	<b>24.3</b>	26.5	24.4
	21–29	<b>6.6</b>	7.6	7.0	<b>27.5</b>	30.2	27.7
	25–29	<b>6.7</b>	9.1	7.9	<b>27.2</b>	34.7	30.0
	29	<b>5.3</b>	12.1	8.7	<b>18.5</b>	43.2	30.8
D	13–37	9.9	10.0	<b>9.8</b>	17.5	<b>16.2</b>	16.6
	17–37	<b>8.5</b>	10.9	9.3	<b>17.1</b>	19.7	17.5
	21–37	<b>7.6</b>	11.0	8.8	<b>15.9</b>	18.9	16.0
	25–37	<b>6.2</b>	10.7	7.8	<b>13.0</b>	18.9	14.2
	29–37	5.1	5.5	<b>5.0</b>	11.8	11.5	<b>11.1</b>
	33–37	<b>6.5</b>	6.7	6.6	<b>14.1</b>	14.6	14.2
	37	3.8	<b>2.7</b>	3.3	8.0	<b>5.7</b>	6.9
E	13–48	<b>6.4</b>	7.9	6.8	<b>49.7</b>	64.8	54.2
	17–48	<b>5.3</b>	8.3	6.0	<b>48.3</b>	70.6	54.4
	21–48	<b>5.4</b>	6.0	5.6	<b>50.9</b>	59.1	53.0
	25–48	<b>5.2</b>	6.6	5.7	<b>53.0</b>	64.7	56.4
	29–49	<b>5.2</b>	5.3	5.1	54.1	57.9	<b>54.0</b>
	33–48	<b>4.6</b>	6.1	5.3	<b>50.5</b>	65.9	57.3
	37–48	<b>4.2</b>	5.7	4.9	<b>47.4</b>	65.7	55.8
	41–48	<b>3.1</b>	3.6	3.4	<b>30.8</b>	35.7	33.2
	45–48	<b>3.1</b>	4.3	3.7	<b>26.3</b>	37.3	31.9

a combined forecast. In order to evaluate the stability of forecast performance over time, the number of fit periods was incremented by four throughout each series. Again, in series E, each model was refitted at the end of months 12, 16, 20, 24, and so on. There were a total of 28 refittings for the five time series. Alternative increments in the number of fit periods of three and six were also tested, but made no significant difference to the results. To simplify the DLS model, the search for the discount factor was constrained to four possibilities: 0.7, 0.8, 0.9, and 1.0. In most cases, 0.7 was the best choice, so this value was selected for all components. As one of the referees pointed out, it is certainly possible that lower discount factor values would yield better results. However, this possibility was not explored in any detail, because management was uncomfortable with the rapid rate of change in the regression slope caused by lower discount factors.

For the exponential smoothing model, a grid search was made over the range 0–1 for all parameters. This search was repeated for each subset of fit periods. Owing to the strong growth patterns in these series, model-fitting produced a near-linear trend in most cases, with a  $\phi$  value near 1.0. The smoothing results can be replicated using the Autocast II software (Gardner, 1992). See the user's guide for more details on initial values of model components and the exact procedures used in the parameter search.

The MAPE results suggest that both the re-

gression and smoothing models are adequate representations of every time series. This impression was confirmed by testing for autocorrelation in the fitted residuals of the two models. No significant autocorrelation was found in any case. Judged by the MAPE, the regression model was consistently more accurate than smoothing. Comparisons are made for 30 sets of forecast periods in Table 1. Only one comparison, the last month of data in series D, favors the smoothing model. The RMSE results also favor the regression model. Regression was more accurate in 22 comparisons, while smoothing was better in only three.

Combining the forecasts was not an effective strategy in these time series. There are only a few cases in which combining improves the MAPE or RMSE, and in many cases combining incurs a significant penalty. This is not surprising, because the regression model thoroughly dominates the smoothing model.

The one-step-ahead forecast simulation for series E is plotted in fig. 5. A matter of concern with this series is that the DLS regression is not appropriate after the structural change to a constant-level series. The problem is that the regression slope coefficient becomes insignificant around period 40, as shown in Fig. 6. Although the regression forecasts behave much like simple exponential smoothing when the slope coefficient approaches zero, simpler models are always desirable in practice. Thus, management was advised to switch the forecasting model to simple




Fig. 5. Component E. One-step-ahead forecast simulation using DLS model.




Fig. 6. Component E. Monthly values of regression slope coefficient.

smoothing when the DLS slope coefficient is not significantly different from zero.

$$MSE_t = e_t^2 + \beta MSE_{t-1} \tag{5}$$

$$r_t = COV_t / MSE_t \tag{6}$$

**5. Forecast monitoring**

In any forecasting system for operational decisions, it is highly desirable to monitor the forecast errors in order to ensure that the system remains in control. One obvious indicator of an out-of-control condition is the first-order autocorrelation in the forecast errors. A DLS model to track the autocorrelation is (Gardner, 1983):

$$COV_t = e_t e_{t-1} + \beta COV_{t-1} \tag{4}$$

The tracking signal  $r$  is an estimate of the autoregressive parameter on successive errors,  $COV$  is the smoothed covariance,  $MSE$  is the smoothed mean-squared-error, and  $\beta$  is the discount factor.

Using  $\beta = 0.9$ , Gardner's simulation results show that a control limit of 0.48 gives a probability of about 1% that a false alarm will occur when the forecasting model is actually in control. The autocorrelation signal was implemented using this combination of discount factor and control limit for all time series, and no alarms of

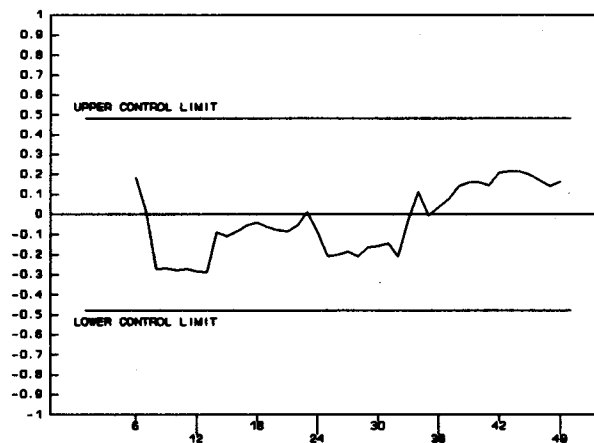


Fig. 7. Component E. Behavior of the autocorrelation tracking signal.

any kind were encountered. This performance was confirmed by a review of the detailed results for each series. There was no reason to believe that alarms should have occurred. An example of the behavior of the tracking signal for component E is shown in Fig. 7. Some positive autocorrelation develops as a result of the change in structure of this series, but is not significant. It should be noted that there is disagreement with the control limits for tracking signals developed in Gardner (1983). See the discussion in McClain (1988) and the comments by Sweet and Wilson (1988). However, it is not clear that this disagreement is relevant here, because the autocorrelation signal performed well with a variety of actual data using the control limit given above. Certainly I can see no reason to experiment with alternative control limits.

## 6. Conclusions

The DLS regression model, supported by the tracking signal, was implemented shortly after the study was completed. The model has been in operation for more than a year at the time of writing. Production plans are adjusted on a monthly schedule, so that stocks of spare components are maintained at the level of forecast demand plus three standard deviations of the forecast error. This strategy is simple, easy to understand, and gives a small probability that a shortage of replacement components will occur. The forecasting spreadsheets have been modified to make the safety stock computations automatically and to plot on-hand stocks versus target stock levels, much like quality control charts. The first set of charts revealed that more stock was on hand than was necessary to provide shortage protection at the three-standard-deviation level. Since then, stocks have gradually been reduced to target levels. In the future, stocks will be reduced further, because experience has shown that the three-standard-deviation level provides more shortage protection than is actually required.

The results of this study must be interpreted with caution because the time series analyzed are short. However, product life cycles in the com-

puter industry are also short, and forecasting is not an elective activity. Without forecasting, there is no sensible basis for inventory control of the component parts.

Judging from the forecasting literature, DLS estimation of an explanatory regression model is rare in practical applications. In Fildes' review (1985) of econometric forecasting, DLS is not even mentioned. An independent search of the literature revealed only one discussion of an explanatory DLS regression model, that of Agnew (1982), although his application was hypothetical. The performance of DLS in this case study suggests that the methodology should at least be considered in practical applications.

Future research will explore other approaches to modelling structural change in this and similar data on component failures. As one of the referees pointed out, DLS has rather shaky theoretical foundations compared with its alternatives. For example, a dynamic regression [see Harvey (1991); Snyder (1984)] may yield more reliable estimates of mean-squared-errors for safety stock determination. It will be worthwhile to know whether more complex approaches actually produce better results from a practical point of view.

## References

- Agnew, R., 1982, "Econometric forecasting via discounted least squares", *Naval Research Logistics Quarterly*, 2, 291–302.
- Armstrong, J.S., 1989, "Combining forecasts: The end of the beginning or the beginning of the end", *International Journal of Forecasting*, 5, 585–588.
- Clemen, R.T., 1989, "Combining forecasts: A review and annotated bibliography", *International Journal of Forecasting*, 5, 559–584.
- Fildes, R., 1985, "Quantitative forecasting—the state of the art: econometric models", *Journal of the Operational Research Society*, 36, 549–580.
- Gardner, E.S., 1983, "Automatic monitoring of forecast errors", *Journal of Forecasting*, 2, 1–21.
- Gardner, E.S., 1992, *Autocast II: Business Forecasting System* (Delphus, Inc., Morristown, PA).
- Gardner, E.S. and E. McKenzie, 1985, "Forecasting trends in time series", *Management Science* 31, 1237–1246.
- Gilchrist, W., 1976, *Statistical Forecasting* (Wiley, New York).
- Harvey, A.C., 1991, *Forecasting with Structural Time Series Models and the Kalman Filter* (Cambridge University Press, Cambridge).



- Makridakis, S., 1989, "Why combining works?", *International Journal of Forecasting*, 5, 601–604.
- McClain, J., 1988, "Dominant tracking signals", *International Journal of Forecasting*, 4, 563–572.
- Snyder, R.D., 1984, "Inventory control with the gamma probability distribution", *European Journal of Operational Research*, 17, 373–381.
- Sweet, A. and J. Wilson, 1988, "Pitfalls in simulation-based evaluation of forecast monitoring schemes", *International Journal of Forecasting*, 4, 573–580.

**Biography:** Everette S. GARDNER, Jr. is Professor of Operations Management and Director of the Center for Global Manufacturing at the University of Houston. He received the PhD in Production and Operations Management from the University of North Carolina at Chapel Hill. He has written numerous research articles in forecasting, inventory control, and production planning. He is co-author of *Quantitative Approaches to Management* (eighth edition) and author of *The Spreadsheet Operations Manager* and *The Spreadsheet Quality Manager*, all published by McGraw-Hill.