

Control charts in practice

In the first issue of *Significance*, **Roland Caulcutt** suggested that a common feature of many successful companies is their development of a “management-by-fact” culture. He further suggested that such a culture requires that managers focus on processes, and that control charts be used at all levels throughout the organisation. In this follow-up article he explains how the various types of control chart are set up and used in practice.

What is a control chart?

A control chart is a graph that can help a “process manager” to make better decisions about the actions that are necessary to achieve the best performance from the process. The simplest, but perhaps the most useful, of the many types of control chart is the individuals chart. Figure 1 is an example of this genre. The five horizontal lines on the graph are intended to aid in its interpretation and are used, together with various rules, to make decisions about the process performance. The centre line is drawn at the mean of the data. The lower and upper warning lines, labelled LWL and UWL, are placed at a distance of two standard deviations from the centre line. The lower and upper control lines, labelled LCL and UCL, are placed at a distance of three standard deviations from the centre line.

Whereas the positioning of the lines is universally agreed, the number and precise details of the rules used in the interpretation of the chart

are certainly not. For simplicity, I shall use the following three rules throughout this article.

Rule 1. Conclude that the process changed if we find *one* or more points above the upper control line, or *one* or more points below the lower control line.


Rule 2a. Conclude that the process changed if we find *two* or more consecutive points between the upper warning line and the upper control line.

Rule 2b. Conclude that the process changed if we find *two* or more consecutive points between the lower warning line and the lower control line.

Rule 3. Conclude that the process changed if we find *eight* or more consecutive points on the same side of the centre line.

Using these three rules to aid our interpretation of Figure 1, we would conclude that the process did not change. Of course, you could argue that the process changed many times, as each point differs from the previous point, but

Figure 1. Individuals chart: delivery times



“A control chart shows us recent performance of the process and predicts, within limits, the performance we can expect in the future”





Figure 2. Standard deviations from (a) equation (3) and (b) equation (1)

I would suggest that these differences are due to the natural random variation of the process. We are looking for real changes. Because we find no such changes, we declare the process to have been stable during the period when these 50 transactions took place.

The distinction between random variation (due to so-called common causes) and real changes (due to so-called assignable or special causes) is very important if we are to manage a process effectively. (For a fuller discussion of this point see "Management by fact" in the first issue of *Significance*.)

Note that the position of the warning lines and the control lines tells us nothing about what is acceptable to the customer or the manager.

The position of each line is based on data from the process. Thus the lines indicate the performance we might reasonably expect from the process, not what we would hope to get. As the

Table 1. Two processes

	Process A	Process B
Value of each item	£0.10	£50 000.00
Production rate	5000 per hour	1 per shift
Cost of inspecting 1 item	£3.00	£400.00
Inspect	5 items every 30 min	Every item
Type of data	Grouped	One at a time
Control chart	Mean chart	Individuals chart
Additional control chart	Range chart	Moving range chart

process is stable, we predict that future performance will give data that lie between the control lines.

A control chart shows us recent performance of the process and predicts, within limits, the performance we can expect in the future. The prediction will be useful only if the future random variation bears some resemblance to that in the chart. But the extent of the variation that we observe in a control chart such as Figure 1 may well depend on the time period during which the data were gathered, for it is well known that many processes perform very consistently for short periods but less consistently over longer periods. For this reason we must give careful consideration to how we calculate the standard deviation of the data.

The process standard deviation

Standard deviation is often defined as the root-mean-squared deviation about the mean, which gives rise to the equations

$$\sigma = \sqrt{\{\sum(x - m)^2/n\}} \quad (1)$$

and

$$\sigma = \sqrt{\{\sum(x - m)^2/(n - 1)\}}. \quad (2)$$

These equations are widely used and are built into electronic calculators and computer software such as Microsoft Excel, MINITAB, SAS and other popular packages. However, these equations are *not* normally used when we assess the process standard deviation in order to set up a control chart. For control charts the universally accepted equation is

$$\sigma = (\text{mean range})/(\text{Hartley's constant}). \quad (3)$$

(Many readers will realise that " σ " is the internationally agreed symbol for population standard deviation, and " s " the agreed symbol for sample standard deviation. Thus we would expect s to be used in the three formulae above. However, the "quality community" defies the convention, using σ rather than s .)

To use equation (3) we must put the data into subgroups, calculate the range of each subgroup and then calculate the mean of the

ranges. With the data in Figure 1 there is no natural grouping, so the choice of subgroup size (n) is quite arbitrary. The default option, $n = 2$, was used in Figure 1. Hartley's constant, for a group size of 2, is 1.128. (Hartley's constant and other constants used in setting up control charts can be obtained from Caulcutt or Oakland in the bibliography and many other texts on statistical process control.)

The effect of using equation (3) rather than equation (1) is illustrated by the two individuals charts in Figure 2. The standard deviation in Figure 2(a), which came from equation (3), is 5.46, but the standard deviation in Figure 2(b), which resulted from the use of equation (1), is 12.45. Clearly the chart in Figure 2(a) detects many changes that the chart in Figure 2(b) fails to indicate.

Subgrouping of data

The standard deviations in Figure 1 and Figure 2(a) were calculated by equation (3) after putting the data into subgroups of size 2. Clearly $n = 2$ is the smallest possible subgroup size, and is preferred when there is no natural grouping of the data. However, when a natural grouping exists it should not be ignored, as it may well offer the best possible subgrouping.

Consider for example the two processes summarised in Table 1. Process B produces high value items. The cost of testing or inspecting each item is much less than the value of the item. Thus, we test every item and we get one result every 8-hour shift. This is often described as one at a time data. Process A produces low value items. The cost of testing or inspecting an item is greater than the value of the item. Thus, 100% inspection would be foolish, so we take a sample of five items every 30 minutes. Clearly our data fall into natural groups with $n = 5$. We preserve this grouping as we plot a mean chart in which each point is the mean of five results.

Figure 3 is a mean chart. Each point plotted in the chart shows the mean of five response times, recorded by an operator who repeatedly logged on to a particular web site. He made five attempts to log on, every half-hour, throughout one particular day. Such a chart could help us to decide whether the accessibility of the web site changed during the day. It appears to have been stable, with an average response time of approximately 15 seconds.

The standard deviation of the data in Figure 3 is calculated from the ranges of the subgroups by using equation (3). Because each subgroup contains five results, Hartley's constant is 2.326. Furthermore, because we are plotting means in Figure 3, the control lines are placed three standard errors from the centre line (i.e. $3\sigma/\sqrt{n}$, where n is the subgroup size.)

Both the mean chart and the individuals chart have the same purpose, namely to detect changes in the process mean. But the mean chart or the individuals chart could be mislead-





Figure 3. Mean chart of response time

ing if the process became more variable or less variable. For this and other reasons, these charts are often used in conjunction with other charts designed to check or to monitor the variability of the process. Figure 4 is such a chart.

Each point in Figure 4 represents the range of one subgroup of data. The stability we see suggests that the variability of the process did not change during the day when this investigation was carried out. The word "variability" in the previous sentence should be taken to mean "short-term variability", as five attempts to log on to the web site would take only 2 minutes, say. (For details of the positioning of the decision lines in a range chart see one of the texts in the bibliography.)

Figure 4. Range chart for response time



Other control charts

We have examined the individuals chart, the mean chart and the range chart. There are many other charts in widespread use. Unfortunately, it is not possible to include a detailed description of every chart in this short article, but the following notes will give some indication of the various charts and their use.

- The moving mean chart and the moving range chart are used by people who have one at a time data but wish to harness the extra power of grouping and averaging.
- The exponentially weighted moving average chart is a type of moving mean chart




Figure 5. CUSUM chart for sales volume

which has found use in the chemical industry, where many processes are prone to drifting.

- The standard deviation chart (s-chart) is an alternative to the range chart.
- The np-chart and the p-chart are based on the binomial distribution and are used to plot attribute data such as the "number defective" and "proportion defective" where there is a known sample size.
- The c-chart and the u-chart are based on the Poisson distribution and are used with attribute data such as the "number of defects" and the "number of defects per unit".
- The cumulative sum (CUSUM) chart could be described as an "integral of error chart" as the plotted function is the cumulative sum of the deviations of the data points from a specified target; see Figure 5.

The CUSUM chart is particularly useful for detecting smaller changes, but it is easily misinterpreted if we fail to apply the following rules.

- The CUSUM chart will have an upward slope if the data are above target.

- The CUSUM chart will have a downward slope if the data are below target.
- Changes in slope in the CUSUM chart indicate changes in mean within the data.

The CUSUM in Figure 5 would suggest changes in sales volume around weeks 7, 17, 23, 34, 42 and 46. These indicated changes agree with those we found in the individuals chart, Figure 2(a). Of course, many of these changes are so large that discovering them does not require the use of a CUSUM chart.

Comparing alternative charts

Several of the charts listed above share the same purpose. The individuals chart, the mean chart, the moving mean chart and the CUSUM chart are all used to detect changes in the mean level of the measured variable. When so many alternatives are available, how can you choose the best chart for your particular application? If you were considering buying a car you would compare the price and the performance of the available models. When choosing from two or more control charts, it

would be wise to consider their relative performance and their ease of use.

As the purpose of a control chart is to detect changes, the performance of a chart can be quantified by the speed with which it indicates change. Of course, any chart is likely to detect a large change very quickly, but a chart may require many points to be plotted before offering convincing proof of a small change. It is common practice, therefore, to compare the average run lengths of alternative charts over a range of change sizes.

Figure 6 offers average run length curves for three mean charts, with sample sizes of 1, 4 and 9. Clearly the mean chart with the largest sample size will detect changes more quickly. But of course the larger sample size demands more time and effort in the sampling and testing or inspecting of the selected items. Similar sets of curves can be produced to illustrate the benefits of using additional decision rules, such as rules 2 and 3 mentioned earlier. Further average run length curves could be used to compare the relative power of mean charts and CUSUM charts, for example.

In "Management by fact" in the first issue of *Significance* I emphasised that managers in organisations need to focus on processes and then to use control charts to obtain the best performance from these processes. This article has described many types of control charts and attempted to illustrate in general terms how control charts are set up and used. Much greater detail is offered in the texts listed in the bibliography.

Bibliography

- Caulcutt, R. (1995) *Statistics in Research and Development*, 2nd edn. London: Chapman and Hall.
- Champ, C. W. and Woodall, W. H. (1987) Exact results for Shewhart control charts. *Technometrics*, 28, 393–399.
- Deming, W. E. (1993) *The New Economics*. Cambridge: MIT Press.
- Grant, E. L. and Leavenworth, R. W. (1996) *Statistical Quality Control*. New York: McGraw-Hill.
- Neave, H. R. (1990) *The Deming Dimension*. Knoxville: SPC.
- Oakland, J. S. (2003) *Statistical Process Control*, 5th edn. Amsterdam: Butterworth-Heinemann.
- Shewhart, W. A. (1939) *Statistical Method from the Viewpoint of Quality Control*. New York: Dover Publications.
- Wheeler, D. J. (1993) *Understanding Variation*. Knoxville: SPC.

Roland Caulcutt is a Principal Consultant with Caulcutt Associates in Sherborne, Dorset. Caulcutt Associates offers training and consultancy in applied statistics, process improvement and human dynamics. Six sigma blackbelt training is currently a major focus, but "Consultancy skills for statisticians" may be better known to readers of *Significance*.

Figure 6. Mean charts (using rule 1 only)

