

Plotting basic control charts: tutorial notes for healthcare practitioners

M A Mohammed,¹ P Worthington,² W H Woodall³

¹ Department of Public Health and Epidemiology, University of Birmingham, Birmingham, UK; ² Prism Europe Consultancy Ltd, Isle of Barra, Outer Hebrides, UK; ³ Department of Statistics, Virginia Polytechnic Institute and State University, USA

Correspondence to:
M A Mohammed, Department of Public Health and Epidemiology, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK; m.a.mohammed@bham.ac.uk

Accepted 18 September 2007

ABSTRACT

There is considerable interest in the use of statistical process control (SPC) in healthcare. Although SPC is part of an overall philosophy of continual improvement, the implementation of SPC usually requires the production of control charts. However, as SPC is relatively new to healthcare practitioners and is not routinely featured in medical statistics texts/courses, there is a need to explain the issues involved in the selection and construction of control charts in practice. Following a brief overview of SPC in healthcare and preliminary issues, we use a tutorial-based approach to illustrate the selection and construction of four commonly used control charts (xmr-chart, p-chart, u-chart, c-chart) using examples from healthcare. For each control chart, the raw data, the relevant formulae and their use and interpretation of the final SPC chart are provided together with a notes section highlighting important issues for the SPC practitioner. Some more advanced topics are also mentioned with suggestions for further reading.

There is considerable interest in the use of statistical process control (SPC) in healthcare.¹⁻⁵ While SPC is part of an overall philosophy aimed at delivering continual improvement, the implementation of SPC usually requires the production of control charts.¹ Since SPC is relatively new to healthcare and does not routinely feature in medical statistics texts/courses, there is a need to explain some of the basic steps and issues involved in selecting and producing control charts.¹ The case for SPC has been made previously^{1 6}; here we show how to plot control charts. With the help of illustrative examples, this paper aims to provide guidance on the selection and construction of control charts that are relevant to healthcare. Some more advanced topics will also be mentioned with suggestions for further reading.

The primary objective of using control charts¹ is to distinguish between common (chance) and special (assignable) causes of variation. The former is generic to any (stable) process and its reduction requires action on the constraints of the process,⁷ whereas special cause variation requires investigation to find the cause and, where appropriate, action to eliminate it. The control chart is one of an array of quality improvement techniques that can be used to deliver continual improvement. Other quality improvement tools have been described elsewhere.⁸⁻¹¹

PRELIMINARY ISSUES

Before constructing control charts, it is essential to have a clear aim and clear action plans on how special cause data points will be investigated, or

how a process exhibiting only common cause variation might be improved. Ensure that individuals involved in the project are aware of the SPC approach to understanding variation.¹ Furthermore, the data collection method should be designed to provide a dataset that adequately reflects the underlying process. Issues such as measurement system analysis, sampling, rational subgrouping and operational definitions are relevant to this and have been discussed elsewhere.⁸⁻¹¹

In industrial practice, it has been found useful to distinguish between two phases in the application of control charts.¹² In phase I, historical data are used to provide a baseline, assess stability, detect special causes and estimate the parameters that describe the behaviour of the process. Phase II consists of ongoing monitoring with data samples taken successively over time and an assumed underlying probability distribution which is appropriate to the process. Shewhart control charts are highly recommended for phase I whereas cumulative sum (CUSUM) and exponentially weighted moving average (EWMA) charts have been shown to detect smaller process shifts in phase II, although it is generally acknowledged that a combination of charts is likely to be advantageous.¹²

INTERPRETING A CONTROL CHART

Control charts include a plot of the data over time with three additional lines—the centre line (usually based on the mean) and an upper and lower control limit, typically set at ± 3 standard deviations (SDs) from the mean, respectively. When data points appear, without any unusual patterns, within the control limits the process is said to be exhibiting common cause variation and is therefore considered to be in statistical control (or stable). However, control charts can also be used to identify special causes of variation. There are several guidelines that indicate when a signal of special cause variation has occurred on a control chart. The first and foremost is that a data point appears outside the control limits. Since the control limits are usually set at 3 SD from the mean, we can state that for a process which is producing normally distributed data (only really demonstrable in artificial simulations) the probability of a point appearing outside the precisely determined phase II control limits (upper or lower limit) is about 1 in 370. Although such a calculation gives some guidance about the “statistical performance” of this particular rule, it is important not to deem normality as a precondition on the use of a control chart¹¹⁻¹³—we return to this issue later.

Developing research and practice

Several other tests can also detect signals of special cause variation based on patterns of data points occurring within the control limits.⁸⁻¹¹ Although there is disagreement about some of the guidelines, three rules are widely recommended:

- ▶ A run of eight (some prefer seven) or more points on one side of the centre line.
- ▶ Two out of three consecutive points appearing beyond 2 SD on the same side of the centre line (ie, two-thirds of the way towards the control limits).
- ▶ A run of eight (some prefer seven) or more points all trending up or down.

Lee and McGreevey¹⁴ recommended the first rule and the trend rule with six consecutive points either all increasing or all decreasing. Any test for runs must be used with care. For example, Davis and Woodall¹⁵ showed that the trend rule does not detect trends in the underlying parameter, and the CUSUM chart can work better than the runs rules in phase II, as shown by Champ and Woodall.¹⁶ Practitioners should note that, as the number of supplementary rules used increases, the number of false alarms will also tend to increase.

Another rule is that patterns on the control chart should be read and interpreted with insight and knowledge about the process. This will help the user to identify those unusual patterns that also indicate special cause variation but may not be covered in the three rules above. For example, imagine a process that performs suboptimally every Monday (and in so doing, produces a data point near, but not beyond, the lower control limit). If a day-by-day control chart is plotted, the pattern for Monday would be repeated every seven data points. Although the three rules above do not capture this scenario,

clearly, consideration of the process would raise the question, why Mondays? On the other hand, there is a natural human tendency to see patterns in purely random data.

SELECTING THE RIGHT CONTROL CHART

There are many different types of control chart,^{5 10 11} and the chart to be used is determined largely by the type of data to be plotted. Two important types of data are: continuous (measurement) data and discrete (or count or attribute) data. Continuous data involve measurement—for example weight, height, blood pressure, length of stay and time from referral to surgery. Discrete data involve counts (integers)—for example, number of admissions, number of prescriptions, number of errors and number of patients waiting.

For continuous data that are available a point at a time (ie, not in subgroups) the *xmr*-chart (also known as the individuals chart) is often appropriate. For discrete data, the *p*-chart, *u*-chart and the *c*-chart are relevant. We will show how to construct the *xmr*-chart, *p*-chart, *u*-chart and *c*-chart using worked examples. Several software packages that produce control charts are now available. We used WinChart Professional, developed by Prism Europe Consultancy (<http://www.winchart.net/index.htm>) to produce our charts, but all the charts can be easily prepared using popular spreadsheet packages.

The *xmr*-chart

Consider the data in the top panel of fig 1, which shows the systolic blood pressure (mmHg) readings for a patient (A Ibrahim, personal communication, 2005) in the morning over 26 consecutive days. Since these are measurement data, we shall

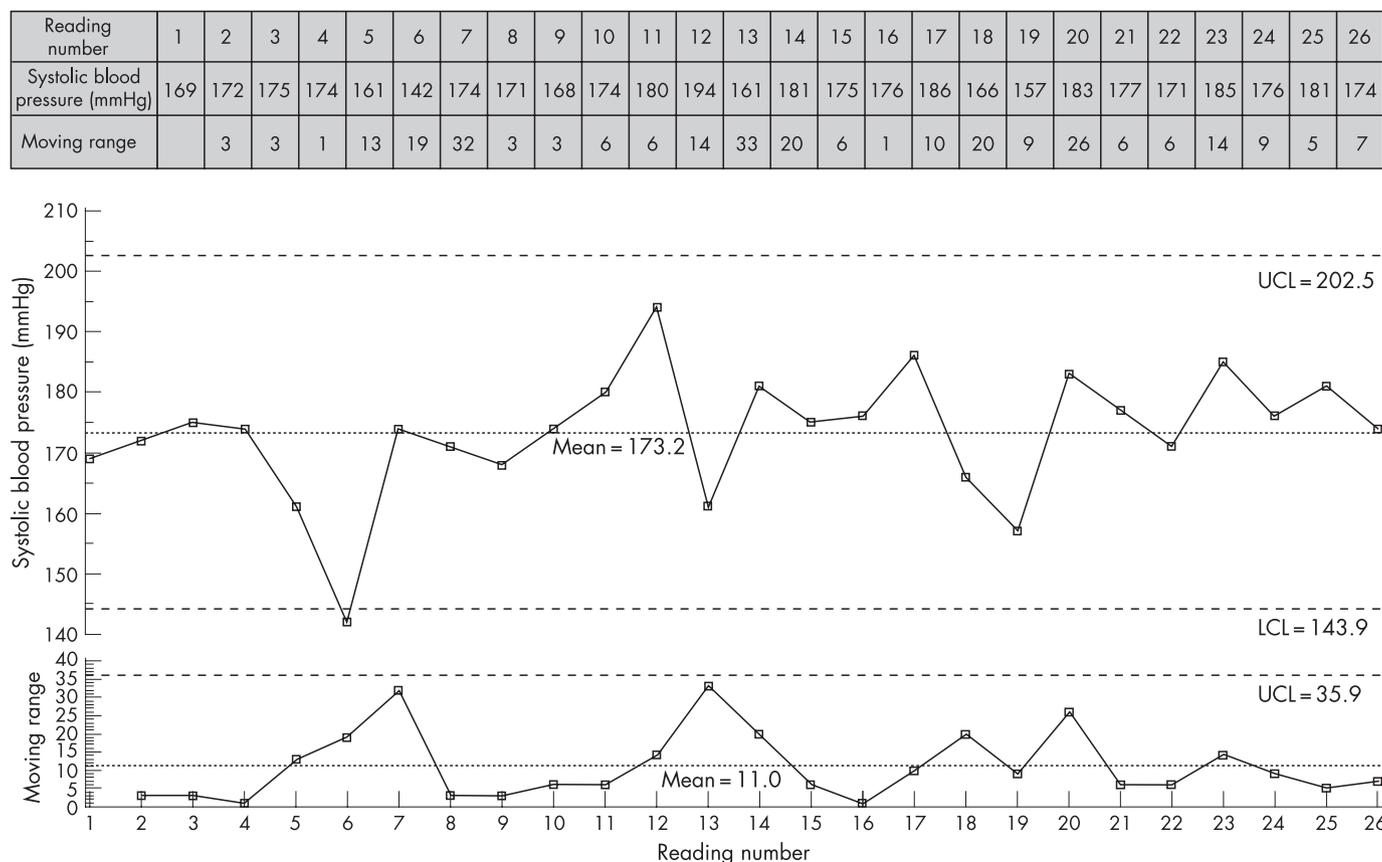


Figure 1 Example *xmr*-chart constructed by using systolic blood pressure readings of a patient. The top panel shows the data and the moving ranges, the middle panel is the *x*-chart and the lower panel is the *mr*-chart.

plot them using an *xmr*-chart. The *xmr*-chart consists of two charts—the *x*-chart and the *mr*-chart. The *x*-chart is a control chart of the n observed values, x_1, x_2, \dots, x_n , and the *mr*-chart is a control chart of the moving ranges of the data (explained below).

The first step in producing an *xmr*-chart is to calculate the magnitudes of the differences between successive values of the data (ie, the moving ranges). In our case, the difference between the second and first reading is $172-169=3$, ..., and the difference between the twenty-sixth reading and the twenty-fifth reading is $174-181=-7$, but we drop the minus sign as we are only interested in the magnitude of the difference, not the direction. The moving ranges are also shown in the top panel of fig 1.

In the second step, to produce the *x*-chart, plot the blood pressure readings (x) on a chart against time. The central line for this chart is given by the mean of x , which is denoted by \bar{x} and is calculated thus:

$$\text{Mean, } \bar{x} = \frac{\sum_{i=1}^n x_i}{n} = \frac{4503}{26} = 173.2$$

So, \bar{x} indicates where to place the central line (ie, at 173.2 on the y-axis). The control limits are based on Shewhart three-sigma limits with sigma (the process standard deviation) being estimated by the mean moving range (\overline{mr}) divided by the empirical constant 1.128. Thus, since $3/1.128 = 2.66$, the control limits for the *x*-chart are given by the formula:

$$\bar{x} \pm 2.66(\overline{mr})$$

The value of \overline{mr} is obtained as follows:

Mean moving range,

$$\overline{mr} = \frac{\sum_{i=2}^n mr_i}{n-1} = \frac{275}{26-1} = 11$$

Hence the control limits are:

$$\begin{aligned} \bar{x} \pm 2.66(\overline{mr}) &= 173.2 \pm 2.66(11) \\ &= 143.9 \text{ for the lower control limit and} \\ &= 202.5 \text{ for the upper control limit} \end{aligned}$$

The middle panel of fig 1 shows the *x*-chart for these data. The reading at day 6 is indicated as being unusually low.

In the last step, for producing the moving range (*mr*) chart, plot the moving ranges against time. The centre line is simply the mean moving range ($\overline{mr}=11$) and the upper control limit is given by $3.267(\overline{mr})$ which gives 35.9. The lower control limit for a *mr*-chart is taken to be 0. The *mr*-chart is shown in the lower panel of fig 1.

Notes

Although the estimation of the process standard deviation (sigma) is based on an underlying normal distribution, it is important to appreciate that the *xmr*-chart is robust (like other exploratory tools) to departures from the assumption of normality. Thus it is not necessary to ensure that the data behave according to the normal distribution before plotting a control chart.^{5 10-12} When the data are clearly not normally

distributed the statistical performance measures of the charts evaluated under the assumption of normality cannot be trusted to be accurate.⁵

One characteristic of time-ordered data that should be assessed is independence over time. Independence means that there is no relationship or autocorrelation between successive data points. With positively autocorrelated data, there will be a tendency for high values to follow high values and low values to follow low values, resulting in cyclic behaviour. With negatively autocorrelated data (less common in applications) high values tend to follow low values and vice versa. Wheeler¹⁷ shows that even with moderate autocorrelations, the *xmr*-chart behaves well (ie, the control limits remain similar), but when the correlation is large (eg, with a lag 1 correlation coefficient $r>0.6$) then the control limits need to be widened by a correction factor. However, Maragah and Woodall¹⁸ found that even moderate autocorrelation can adversely influence the statistical performance of a chart to the extent that time-series based methods may be more appropriate for autocorrelated data. A common approach is to model the time series and use control charts on the residuals to detect changes in the process.

It is possible that sometimes the lower control limit will be below 0. If the data under question come from a process in which negative measurements are not feasible then the lower control limit is customarily reset to 0 to reflect this. It is often stated that 20–25 data points are needed on an *xmr*-chart before the limits are sufficiently “firmed-up” so that if a process is showing common cause variation over these 20–25 data points then we can confidently conclude that it is stable.⁹ This is reasonable advice, but it does not imply that a control chart with fewer data points is not useful.¹⁷ With fewer data points, the resulting control limits are regarded as being “soft”¹⁷ or “provisional” control limits. In practice, where there is a stream of data, “provisional” control limits can be computed and refined (firmed up) with additional data in due course. Another important reason for re-computing the control limits is when knowledge about the underlying process deems it necessary (eg, a revised definition of the data, material changes to the process, etc.). Jensen *et al*¹⁹ reviewed the literature on the effect of estimation on control chart performance. They showed that sample sizes need to be fairly large to achieve the statistical performance expected under the known parameters case.

Limitations

The *xmr*-chart is a robust, versatile chart that has been used with a variety of processes. However, when the underlying data exhibit seasonality or the data are for rare events then preliminary work with the data is required before it can be placed on an *xmr*-chart.²⁰ There is disagreement in the literature on the use of the *xmr*-chart.⁵ Some have argued that the EWMA chart is superior to the *xmr*-chart in phase II applications. Others have shown that for time-between-events data, or successes-between-failures data, other charts^{21 22} are superior. However, the *xmr*-chart continues to be a useful, if not always optimal, tool for identifying special causes of variation in many practical applications.¹⁷ Even when the data are from a highly skewed distribution (eg, exponential waiting times, see Normality section below) a straightforward double square root transformation ($y = \sqrt{\sqrt{x}}$) will often render the data suitable for an *xmr*-chart.²³ According to Wheeler, the *xmr*-chart is even useful when considering count (attribute) data (see Wheeler^{17 20} for further details).

It is customary to produce the *xmr*-chart using the mean (of the data and the moving ranges) statistic for the centre line.

However, the median statistic may also be used, but the constants used in the formulas to produce the three-sigma limits will be different¹¹ (3.145 instead of 2.66 for the \bar{x} -chart and 3.865 instead of 3.267 for the mr -chart). Wheeler and Chambers¹¹ suggest that when the mean moving range appears to be inflated by just a few data points, xmr -charts based on the median may be a more suitable choice, because the median statistic, although less efficient in its use of the data compared with the mean, affords greater resistance (to extreme data points) than the mean.

Normality

It is sometimes argued that a precondition to the use of a control chart is for assumptions of normality to be met.^{13 17} This is not correct. (This issue has been discussed in depth by Woodall.¹²) Control charts have been shown to be robust to the assumptions of normality. Furthermore during phase I applications the distributional assumption cannot even be checked because the underlying process may not be stable. As special causes of variation are removed the process becomes more stable. Also the form of the hypothesised distribution becomes more apparent and useful in defining the statistical performance characteristics of the control chart. According to Woodall,⁵ during phase I applications, practitioners need to be aware that the probability of signals of special cause variation depend primarily on the shape of the underlying distribution, the degree of autocorrelation and the number of samples.

The preceding arguments do not imply that there is nothing to gain in checking assumptions of normality. A useful exploratory graphical method for assessing normality is the normal probability plot^{10 23} (which is available in most statistical software packages). A normal probability plot is a graphical method for determining if a given set of data are consistent with an underlying normal distribution. We will illustrate its use with data ($n = 100$) generated from an exponential distribution with a rate parameter of 0.2, producing a highly skewed simulated time-to-event data set. The left panel of fig 2 shows a normal probability plot using the raw data. If the data were plausibly normal then the points would be expected lie close to

the diagonal line of a normal probability plot. This is clearly not the case. However, if we transform the data using the double square root transformation $y = \sqrt{\sqrt{x}}$ (see fig 2 right panel), the transformed data are now much more plausibly normal and therefore more appropriate for plotting on an xmr -chart. So, with the caveat that normality is not a precondition on the use of a control chart, we suggest that the production of a normal probability plot is potentially useful, especially in identifying highly skewed data. But ensure that concepts of “statistical outliers” on the normal probability plot are not confused with “signals” of special cause variation on a control chart. The former may be used as evidence against normality (in a formal statistical test for normality) but the latter may guide a practitioner towards finding and removing special causes of variation.

The p -chart

Consider the data in the top panel of fig 3. The data show the number of patients who were admitted with a fractured neck of femur and the number who died over 24 consecutive quarters (M Narayan-Lee, personal communication, 2006). To find out whether variation in mortality over time is consistent with common cause variation, we shall use a p -chart (where p stands for proportion).

To plot a p -chart for the above data, use the notation given in the data rows of fig 3. Plot the proportion of deaths (p) on the y-axis and the time (quarters) on the x-axis and then add the central line and the control limits as they are calculated. To determine the central line, compute the overall mean mortality:

$$\bar{p} = \frac{\sum_{i=1}^n x_i}{\sum_{i=1}^n n_i} = \frac{345}{1406} = 0.25$$

Thus, \bar{p} indicates where to place the central line (ie, at 0.25 on the y-axis). To derive the control limits we apply this overall \bar{p} to each of the quarters cases (n_i) in turn using the formula:

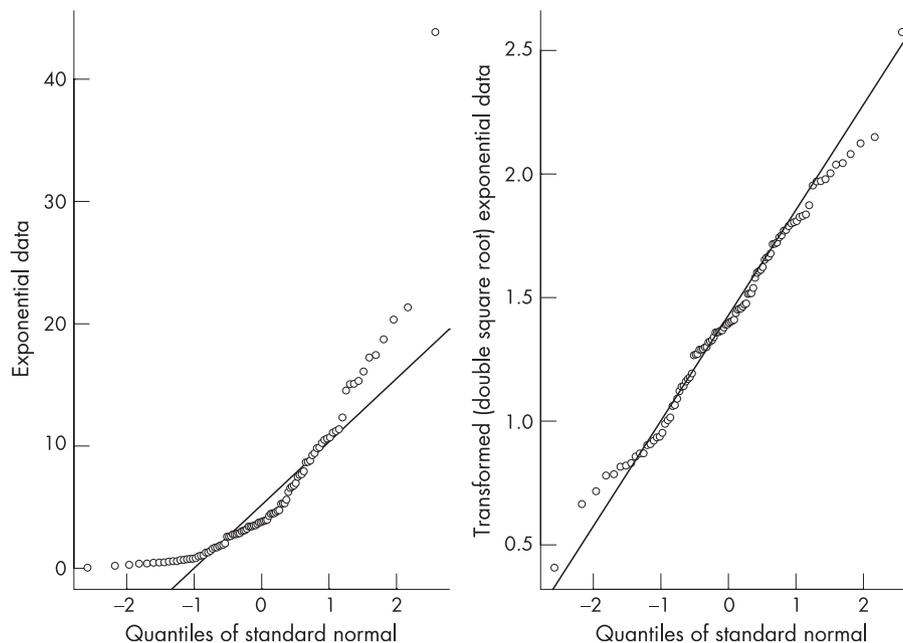


Figure 2 Normal probability plots. Left panel: the raw exponential data ($n = 100$, rate parameter = 0.2); right panel: the transformed data.

$$\bar{p} \pm 3 \sqrt{\bar{p}(1-\bar{p})/n_i}$$

So, for example, for quarter 1 ($n_1 = 56$) we get the following control limits:

$$0.25 + 3\sqrt{0.25(1-0.25)/56} = 0.42 \text{ (upper control limit) and}$$

$$0.25 - 3\sqrt{0.25(1-0.25)/56} = 0.07 \text{ (lower control limit)}$$

This calculation is repeated for each quarter to finally produce the control chart shown in fig 3. The control limits thus produced are stepped (see fig 3) because they reflect the changes in the sample sizes between quarters. There is no evidence of special cause variation. Improvement in this case requires a closer look at the constraints of the process,⁷ perhaps with the aid of a Pareto chart¹⁰ showing the most frequent reasons for death.

Notes

The assumptions of the p -chart¹⁷ are that the events (deaths in our case) are: (a) binary (can only have two states, eg, alive/dead, infected/not infected, admitted/not admitted, etc); (b) have a constant underlying probability of occurring; and (c) that they are independent of each other. One simple process that can meet all these requirements is the repeated tossing of a coin, but it is indeed impossible for a healthcare process to meet all these assumptions exactly. Nevertheless experience indicates that the

p -chart is useful in practical circumstances, even where there is gross departure from the ideal conditions. For instance, Deming²⁴ shows an example of inspection data that appeared to have much less variation than expected. The control limits were very wide and the data appeared to “hug” the central line. This “under-dispersion” was subsequently traced to an inspector who was too frightened to record the proper failure rates, choosing instead to make up the figures to be just below the minimum target. Clearly, assumptions (b) and (c) were being violated.

Conversely, when the observed variation is much greater than expected a substantial number of the data points fall outside the control limits. This is termed “over-dispersion” and often indicates that assumption (b) and/or (c) have been violated. While there are statistical techniques for allowing for “over-dispersion”, this technical fix²⁵⁻²⁷ does not itself address the fundamental question of why “over-dispersion” exists. This requires detective work to understand the underlying process and learn why it is behaving in this way. (Wheeler recommends using the xmr -chart in this situation.¹⁷)

When using the p -chart to plot percentages, all the proportions, the centre line and the control limits are multiplied by 100. In the case where the upper control limit is above 1, or the lower control limit is below 0, such limits are customarily reset to 1 and 0, respectively, because proportions cannot be negative or larger than 1. We do not have a signal if the observed value was to fall exactly on the reset limit.

The standard deviation of a binomial variable is usually derived using $\sqrt{\bar{p}(1-\bar{p})/n_i}$, hence giving the control limits as

Quarter	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
Number of patients admitted with fracture neck of femur (n)	56	53	45	44	46	50	48	48	42	46	41	55	48	58	62	72	57	60	65	100	77	76	84	73
Number died (x)	11	12	11	12	4	12	15	13	9	14	10	12	14	19	17	24	13	13	15	25	24	16	15	15
Proportion died (x/n)	0.20	0.23	0.24	0.27	0.09	0.24	0.31	0.27	0.21	0.30	0.24	0.22	0.29	0.33	0.27	0.33	0.23	0.22	0.23	0.25	0.31	0.21	0.18	0.21
Centre line	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25
Upper control limit	0.42	0.42	0.44	0.44	0.44	0.43	0.43	0.43	0.44	0.44	0.45	0.42	0.43	0.41	0.41	0.40	0.42	0.41	0.41	0.37	0.39	0.39	0.39	0.40
Lower control limit	0.07	0.07	0.05	0.05	0.06	0.06	0.06	0.06	0.05	0.06	0.04	0.07	0.06	0.08	0.08	0.09	0.07	0.08	0.09	0.12	0.10	0.10	0.10	0.09

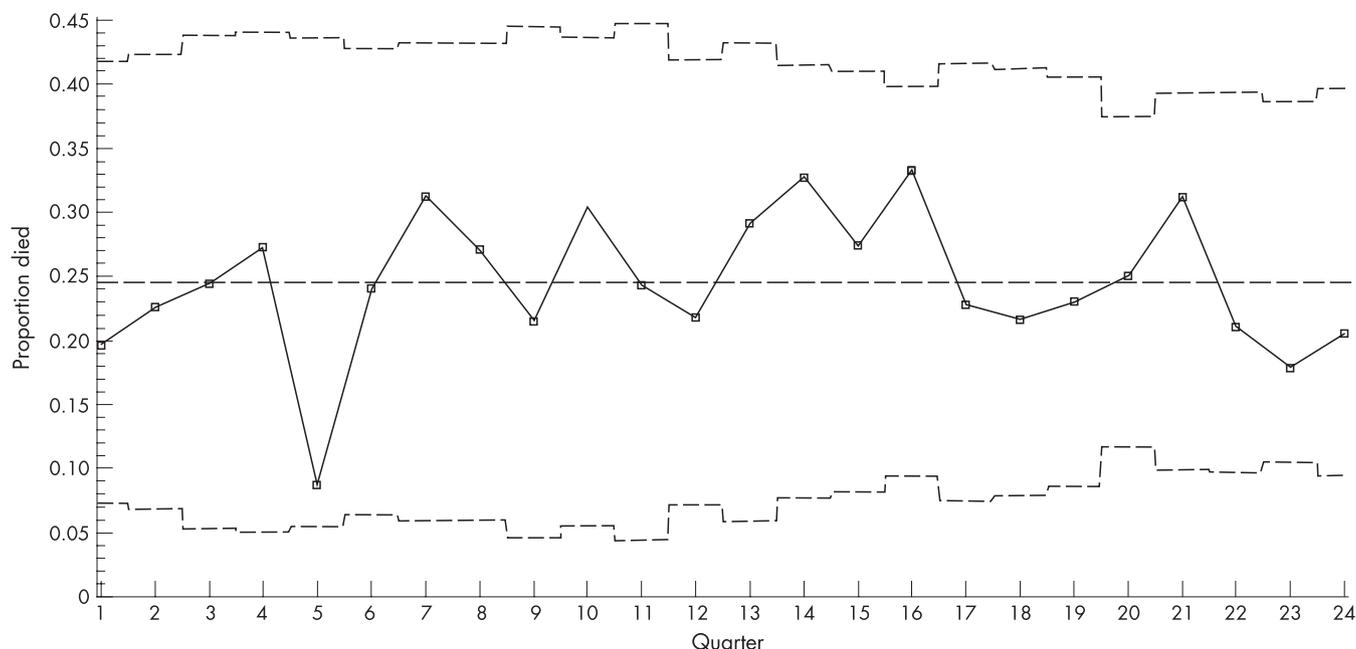


Figure 3 Example p -chart based on the proportion of deaths following admission to hospital for patients with a fractured neck of femur.

Developing research and practice

$\bar{p} \pm 3\sqrt{\bar{p}(1-\bar{p})/n_i}$. According to Xie *et al*²⁸ this approximation is good as long as $n_i\bar{p}(1-\bar{p}) > 5$ and $0.1 \leq \bar{p} \leq 0.9$ as is the case in the worked example above, because under these conditions the binomial distribution is fairly symmetrical and so the Shewhart three-sigma concept works well. However, when these conditions are not satisfied, then a different approach is required. One method is to calculate the exact limits using the probability distribution function of the binomial distribution, which, with modern software such as popular spreadsheet packages and statistical packages, is relatively straightforward. Other approaches involving transformations and regression-based limits have also been documented.^{10 29}

A common application of a graph similar to the p -chart is a comparison of the performance of healthcare providers^{5 30} over a fixed period. In this instance there is no time order to the data. The proportions, such as the proportions of readmissions, are plotted as a function of the subgroup sizes (n). The resultant chart is a "funnel" shape which is visually attractive and statistically intuitive because the funnel shows how the variation due to "chance" reduces with increasing sample sizes. This avoids the misleading ranking of the providers in league tables and its attendant negative consequences.^{3 30}

The u -chart

Consider the data in the top panel of fig 4 which shows the number of falls in a hospital department over a 13-month period.³¹ To investigate the type of variation in these data we shall use a u -chart.

To plot a u -chart for the above data, plot the fall rate per patient-day (u) on the y-axis and the time (months) on the x-axis. Then add the central line and the control limits as they are calculated. To determine the central line, compute the overall fall rate (\bar{u}) which is given by the total number of falls divided by the total number of patient-days:

$$\bar{u} = \frac{\sum_{i=1}^n u_i}{\sum_{i=1}^n n_i} = \frac{39}{10501} = 0.00371 \text{ (fall rate on "average" per patient-day)}$$

Thus, \bar{u} indicates where to place the central line (ie, at 0.00371 on the y-axis). To derive the control limits we apply this overall average \bar{u} and the number of days in each time period (n_i) in turn using the formula:

$$\bar{u} \pm 3\sqrt{\bar{u}/n_i}$$

So for example, for October 2004, ($n_i = 1048$) we get the following control limits:

$$0.00371 + 3\sqrt{0.00371/1048} = 0.00936 \text{ (upper control limit) and}$$

$$0.00371 - 3\sqrt{0.00371/1048} = -0.00193 \text{ (lower control limit)}$$

Month	Oct 04	Nov 04	Dec 04	Jan 05	Feb 05	Mar 05	Apr 05	May 05	Jun 05	Jul 05	Aug 05	Sep 05	Oct 05
Number of falls	1	4	3	4	2	3	5	2	0	2	6	2	5
Number of patient days	1048	896	918	995	866	896	864	930	732	630	492	622	612
Falls per patient day	0.000954	0.004464	0.003268	0.004020	0.002309	0.003348	0.005787	0.002151	0.000000	0.003175	0.012195	0.003215	0.008170
Centre line	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714	0.003714
Upper control limit	0.009361	0.009822	0.009748	0.009510	0.009927	0.009822	0.009934	0.009709	0.010471	0.010998	0.011956	0.011045	0.011104
Lower control limit	-0.001934	-0.002394	-0.002320	-0.002082	-0.002499	-0.002394	-0.002506	-0.002281	-0.003044	-0.003570	-0.004529	-0.003617	-0.003676

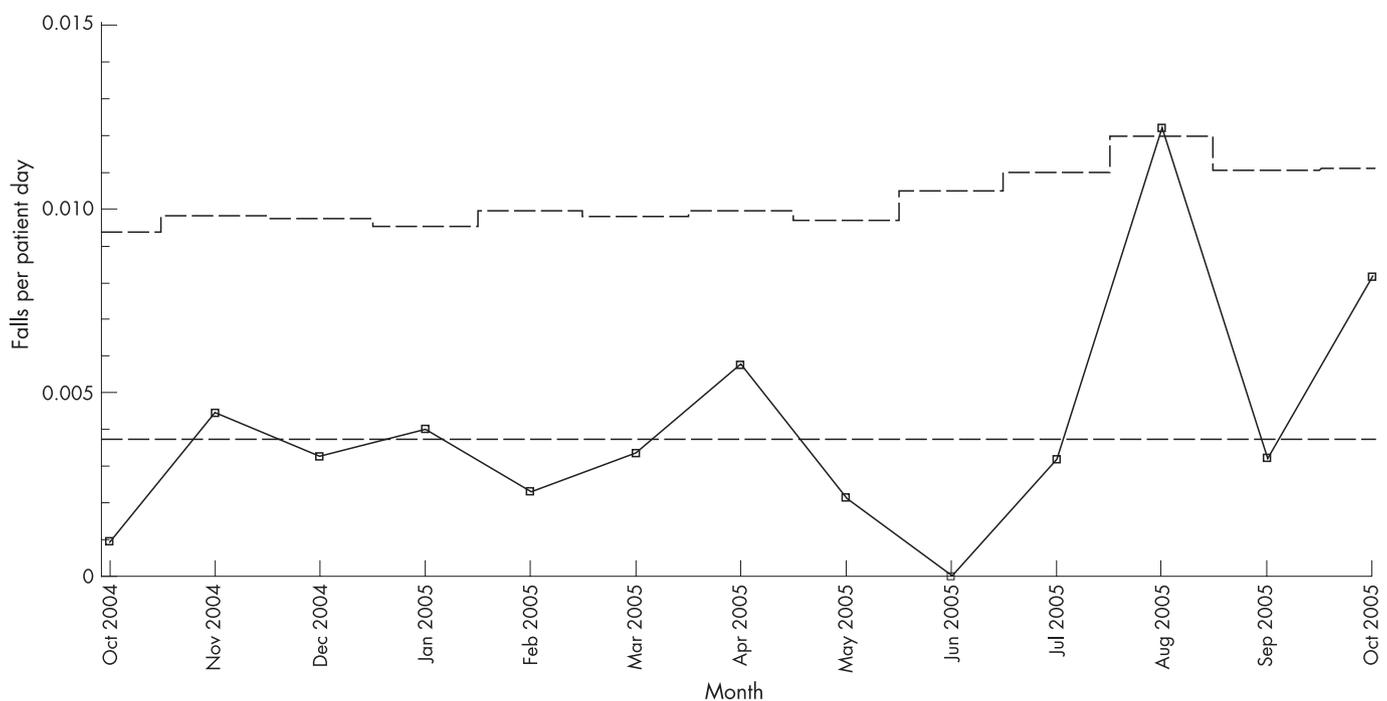


Figure 4 Example u -chart based on number of falls per patient-day over a 13-month period.

This calculation is repeated for each time period to finally produce the control chart shown in fig 4. When the calculated lower control limit falls below 0 it is customarily reset to 0 because count/attribute data cannot fall below 0. As can be seen from fig 4, the control limits thus produced are stepped because they reflect the changes in the area of opportunity (number of patient-days) between sampling periods. There is evidence of special cause variation in August 2005, which merits further investigation.

Notes

The assumptions of the u -chart are that the events: (a) occur one at a time with no multiple events occurring simultaneously or in the same location; and (b) are independent in that the occurrence of an event in one time period or region does not affect the probability of the occurrence of any other event. Although it is rare for all these assumptions to be met exactly in practice, experience indicates that the u -chart is useful in practical circumstances, even where there is a marked departure from the ideal conditions.

Typically low frequency events (eg, number of major complications following surgery) are plotted using u -charts. This is appropriate because the calculations for the control limits are based on an underlying Poisson distribution which is reasonable for infrequent events. To aid communication though, the low frequency event is often re-expressed—for example, using our falls data, in October 2004 there was 1 fall in 1048 patient-days, which is 0.000954 falls per patient-day, but this can be re-expressed as a rate per 100 patient-days by multiplying by 100 ($1/1048 \times 100 = 0.0954$). Such an adjustment will make no difference to the analysis since a u -chart with the y-axis as falls per 100 patient-days will give the same messages as the one with falls per patient-day, albeit in a unit of analysis which may be more intuitive to healthcare practitioners.

The c -chart

Consider the data in the top panel of fig 5 which shows the number of emergency admissions over 23 consecutive Mondays (1 December 2003 to 3 May 2004) to one large acute hospital in England (personal communication). If we regard the people in this hospital's catchment area as being the underlying population from which these admissions occur, we can see that (a) the events are relatively low frequency (in comparison with the size of the underlying population) but (b) the size of that underlying population is unknown. Assuming that the underlying population is large and fairly constant, we can analyse these data using a c -chart, which is essentially a u -chart with $n_i = 1, i = 1, 2, \dots, n$.

For a c -chart, plot the number of admissions (c) on the y-axis and the days (time) on the x-axis, and then add the central line and the control limits as they are calculated. To determine the central line, compute the overall mean number of admissions (\bar{c}) which is given by the total number of admissions divided by the total number of days:

$$\bar{c} = \frac{\sum_{i=1}^n c_i}{n} = \frac{1847}{23} = 80.3 \text{ (admissions on "average" per Monday)}$$

Thus, \bar{c} indicates where to place the central line (ie, at 80.3 on the y-axis). The control limits are set at $\bar{c} \pm 3$ times the square root of \bar{c} since, as can be shown, the standard deviation of a Poisson distribution is the square root of the mean, viz, $\bar{c} \pm 3\sqrt{\bar{c}}$, this gives:

$$80.3 + 3\sqrt{80.3} = 107.2 \text{ (upper control limit) and}$$

$$80.3 - 3\sqrt{80.3} = 53.4 \text{ (lower control limit)}$$

The final c -chart is shown in fig 5.

Monday number	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
Number of emergency admissions (c)	82	63	86	72	82	88	95	68	81	77	92	69	83	86	86	86	68	87	74	62	100	85

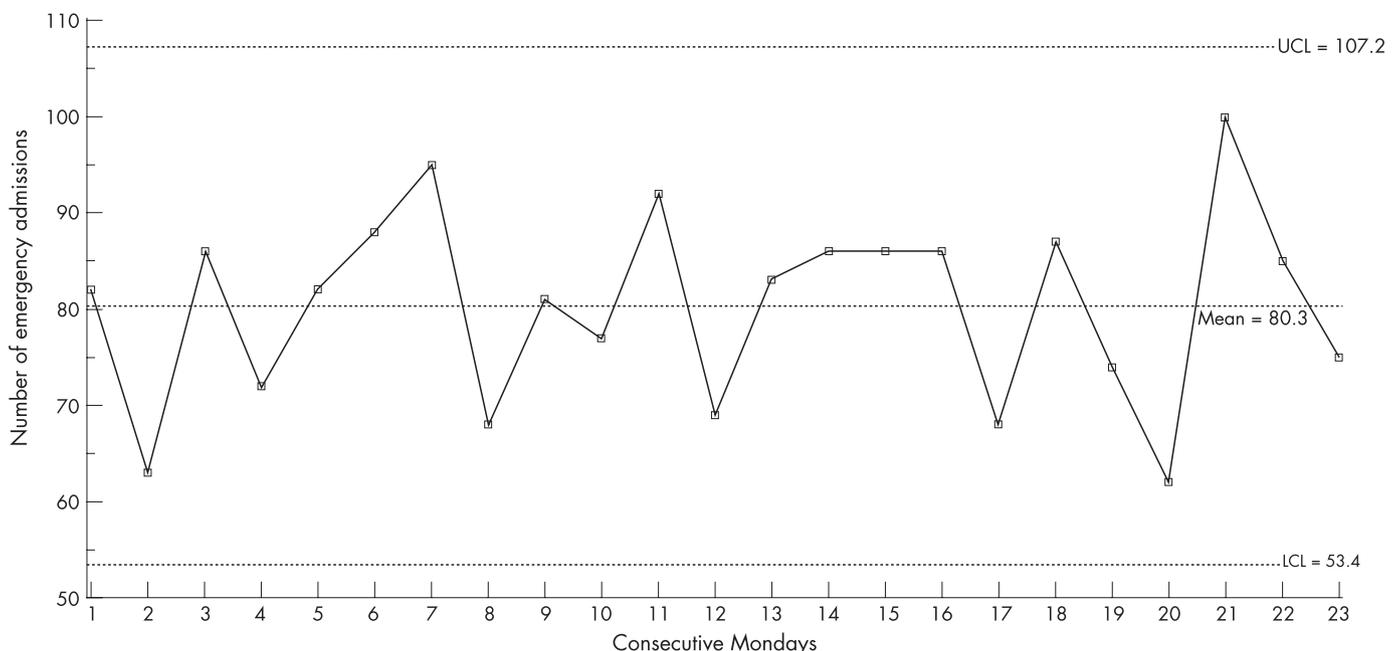


Figure 5 Example c -chart using number of emergency admissions on consecutive Mondays.

Notes

The assumptions of the c -chart are that the events: (a) occur one at a time with no multiple events occurring simultaneously or in the same location; and (b) are independent in that the occurrence of an event in one time period or region does not affect the probability of the occurrence of any other event. Although it is impossible for all these assumptions to be met exactly in practice, experience indicates that the c -chart is useful in practical circumstances, even where there is significant departure from the ideal conditions. It is important to note that the u -chart is preferred to the c -chart when sample sizes or areas of opportunity are available and vary considerably. When such an area of opportunity information is not available then the c -chart is the only real option, provided that it is reasonable to assume that the underlying areas of opportunity are relatively large and fairly constant—as is the case in our emergency admissions example here.

The control limits for a c -chart are derived from $\bar{c} \pm 3\sqrt{\bar{c}}$. When $\bar{c} > 10^{10}$ as is the case here, the Poisson distribution is roughly symmetrical and so the Shewhart three-sigma concept works well. However, when $\bar{c} > 10$ is not satisfied, then a different approach is required. One method is to calculate the exact limits using the probability distribution function of the Poisson distribution with the help of computer software or statistical tables.¹⁰ Other approaches, using transformations as well as regression-based limits, have also been suggested.²⁹ When $\bar{c} < 9$ the computation of the lower control limit produces a value below 0,¹⁰ which is customarily reset to 0 because the process is not capable of producing a negative count. However, the negative lower control limit is also an indication that the Poisson distribution is not symmetrical and the computation of exact limits is then recommended. Furthermore, according to Wheeler,¹⁷ the c -chart will break down when the mean number of events falls below 1, at which point plotting the time between events can be helpful.²⁰

In this worked example we plotted the number of emergency admissions on consecutive Mondays instead of consecutive days. The reason for this is related to rational subgrouping. Some processes are made up of different underlying cause–effect subprocesses and an emergency admission to hospitals is one example. So for instance we do not expect the probability of an emergency admission to hospital to be the same for weekends (Saturday/Sunday) and weekdays, nor do we expect Fridays to be the same as other weekdays. In other words, the process exhibits seasonality. In order not to violate the requirements of rational subgrouping, we split the emergency admissions data by day of week and used Mondays as an illustrative example. If a control chart for emergency admissions for consecutive days is required then methods which first deseasonalise the data must be used.²⁰

Other types of control charts

We have considered the four types of control chart that will have frequent applicability to healthcare. Many more types of control chart and related plots are documented in the literature.^{10–17} These include EWMA control charts, CUSUM plots, Q-charts²² and number-between-events charts,²¹ to mention but a few. While several of these charts have apparently superior performance compared with the basic charts described here, it is important to note that generally speaking, they are often not meant to completely replace Shewhart control charts; rather they are considered to be useful adjuncts.¹² An interesting feature of the application of control charts to healthcare is the incorporation of patient case-mix adjustment, which has been described

elsewhere.^{5 25 32–34} Case-mix adjustment is often required to make meaningful comparisons between healthcare providers.³⁵

CHOICE OF THREE-SIGMA LIMITS

Shewhart argued—on the basis of mathematical theory (Tchebysheff theorem)³⁶ and also on the basis of experience and pragmatism—that stable processes produce variation within limits and that, in the search for economic control of variation, these limits can usefully be set at three-sigma. Several decades of experience in a whole range of application domains, including healthcare, has shown that three-sigma limits are indeed useful. Nevertheless this does not imply that the limits cannot be changed. According to Nelson,¹⁵ the rationale for widening or narrowing the limits is a judgement call in which the costs of looking for special cause variation, when it does not exist, need to be balanced against the costs of overlooking such a signal, when it does exist. However we must emphasise that wherever the control limits are set at a finite positive distance from the centre line we risk one of two types of mistake²⁴:

- ▶ type I: to treat an outcome resulting from a common cause as if it were a special cause;
- ▶ type II: to treat an outcome resulting from a special cause as if it were a common cause.

It is impossible to reduce the frequency of both errors to 0, but what we can do is minimise the economic losses due to either kind of mistake. Nevertheless, despite some of the controversies regarding the theoretical basis of control charts,¹² there is ample evidence, increasingly from healthcare, to show that in the context of continual improvement, control charts offer major benefits which are well worth the effort.

Acknowledgements: We thank Mr A M Ibrahim for giving us permission to use his blood pressure measurements data and Mrs M Narayan-Lee for the fractured neck of femur and emergency admissions data sets used in this article. Thanks are also due to StatIt.com (www.statit.com) for kind permission to use the falls data.

Competing interests: None declared.

REFERENCES

1. **Mohammed MA.** Using statistical process control to improve the quality of health care. *Qual Saf Health Care* 2004;**13**:243–5.
2. **Benneyan JC, Lloyd RC, Plsek PE.** Statistical process control as a tool for research and healthcare improvement. *Qual Saf Health Care* 2003;**12**:458–64.
3. **Adab P, Rouse A, Mohammed MA, et al.** Performance league tables: the NHS deserves better. *BMJ* 2002;**324**:95–8.
4. **Diaz M, Neuhauser D.** Pasteur and parachutes: when statistical process control is better than a randomized controlled trial. *Qual Saf Health Care* 2005;**14**:140–3.
5. **Woodall WH.** The use of control charts in health-care and public-health surveillance (with discussion). *Journal of Quality Technology* 2006;**38**:104–5.
6. **Mohammed MA, Cheng KK, Rouse A, et al.** Bristol, Shipman and clinical governance: Shewhart's forgotten lessons. *Lancet* 2001;**357**:463–7.
7. **Dettmer HW.** *Goldratt's theory of constraints: a systems approach to continuous improvement.* Wisconsin, USA: ASQ Press Milwaukee, 1997.
8. **Kelly DL.** *How to use control charts for healthcare.* Wisconsin, USA: Quality Press, 1999.
9. **Carey RG.** *Improving healthcare with control charts: basic and advanced SPC methods and case studies.* Wisconsin, USA: Quality Press, 2003.
10. **Ryan TP.** *Statistical methods for quality improvement,* 2nd edn. New York: John Wiley & Sons, 2000.
11. **Wheeler DJ, Chambers DS.** *Understanding statistical process control,* 2nd edn. Knoxville, TN: SPC Press, 1992.
12. **Woodall WH.** Controversies and contradictions in statistical process control. *Journal of Quality Technology* 2000;**32**:341–50.
13. **Nelson LS.** Technical aids: notes on the Shewhart control chart. *Journal of Quality Technology* 1999;**31**:124–6.
14. **Lee K, McGreevey C.** Using control charts to assess performance measurement data. *Jt Comm J Qual Improv* 2002;**28**:90–101.
15. **Davis RB, Woodall WH.** Performance of the control chart trend rule under linear shift. *Journal of Quality Technology* 1988;**20**:260–2.
16. **Champ CW, Woodall WH.** Exact results for Shewhart control charts with supplementary runs rules. *Technometrics* 1987;**29**:393–9.
17. **Wheeler D.** *Advanced topics in statistical process control. The power of Shewhart's charts.* Knoxville, TN: SPC Press, 1995.

18. **Maragah HD**, Woodall WH. The effect of autocorrelation on the retrospective \bar{x} -chart. *Journal of Statistical Computation and Simulation* 1992;**40**:29–42.
19. **Jensen WA**, Jones-Farmer LA, Champ CV, *et al*. Effects of parameter estimation on control chart properties: a literature review. *Journal of Quality Technology* 2006;**38**:349–64.
20. **Wheeler D**. *Making sense of data*. Knoxville, TN: SPC Press, 2003.
21. **Benneyan JC**. Number-between g -type statistical quality control charts for monitoring adverse events. *Health Care Manag Sci* 2001;**4**:305–418.
22. **Quesenberry C**. Statistical process control geometric Q -chart for nosocomial infection surveillance. *Am J Infect Control* 2000;**28**:314–20.
23. **Montgomery DC**. *Introduction to statistical quality control*, 5th edn. New York, NY: John Wiley & Sons, 2005.
24. **Deming WE**. *Out of the crisis*. Cambridge, MA: Massachusetts Institute of Technology, 1986.
25. **Aylin P**, Best N, Bottle A, *et al*. Following Shipman: a pilot system for monitoring mortality rates in primary care. *Lancet* 2003;**362**:485–91.
26. **Spiegelhalter D**. Handling over-dispersion of performance indicators. *Qual Saf Health Care* 2005;**14**:347–51.
27. **Mohammed MA**, Laney D. Over-dispersion in healthcare performance data: Laney's approach. *Qual Saf Health Care* 2006;**15**:383–4.
28. **Xie M**, Goh TN, Kuralmani V. *Statistical models and control charts for high quality processes*. Kluwer Academic Publishers, 2002.
29. **Ryan TP**, Schwertman NC. Optimal limits for attributes control charts. *Journal of Quality Technology* 1997;**29**:86–98.
30. **Spiegelhalter D**. Funnel plots for comparing institutional performance. *Stat Med* 2005;**24**:1185–202.
31. **Statit Software**. Data source available at http://www.statit.com/support/quality_practice_tips/ucharts_pcharts_icharts.htm (accessed October 2006).
32. **Tekkis PP**, McCulloch P, Steger AC, *et al*. Mortality control charts for comparing performance of surgical units: validation study using hospital mortality data. *BMJ* 2003;**326**:786–8.
33. **Hart MK**, Lee KY, Hart RF, *et al*. Application of attribute control charts to risk-adjusted data for monitoring and improving health care performance. *Qual Manage Health Care* 2003;**12**:15–19.
34. **Grigg O**, Farewell V. An overview of risk-adjusted charts. *J R Stat Soc A* 2004;**167**:523–39.
35. **Lilford RJ**, Mohammed MA, Spiegelhalter D, *et al*. Use and misuse of process and outcome data in managing performance of acute medical care: avoiding institutional stigma. *Lancet* 2004;**363**:1147–54.
36. **Shewhart WA**. *Economic control of quality of manufactured product*. New York, NY: D Van Nostrand Company, 1931 (reprinted by ASQC Quality Press, 1980)

APPENDIX

Summary of formulae for plotting four basic control charts

Type of data	Type of chart	Centre line	Three-sigma control limits
Continuous data (x) of which there are n items	\bar{x} -chart	$\bar{\bar{x}} = \frac{\sum_{i=1}^n X_i}{n}$	$\bar{\bar{x}} \pm 2.26 (\overline{mr})$
	mr -chart	$\overline{mr} = \frac{\sum_{i=2}^n mr_i}{n-1}$	$3.267 (\overline{mr})$ (upper control limit only)
Proportion data (p) given by x/n where x is the number of occurrences and n is its number of opportunities.	p -chart	$\bar{p} = \frac{\sum_{i=1}^n X_i}{\sum_{i=1}^n n_i}$	$\bar{p} \pm 3 \sqrt{\bar{p}(1-\bar{p})/n_i}$
Count data (u) with n observations each in a known "area of opportunity" n_i	u -chart	$\bar{u} = \frac{\sum_{i=1}^n u_i}{\sum_{i=1}^n n_i}$	$\bar{u} \pm 3\sqrt{\bar{u}/n_i}$
Count data (c) with n observations each with a constant "area of opportunity"	c -chart	$\bar{c} = \frac{\sum_{i=1}^n c_i}{n}$	$\bar{c} \pm 3\sqrt{\bar{c}}$