Public Health &
Intelligence

GUIDANCE ON STATISTICAL PROCESS CONTROL CHART RULES AND INTERPRETATIONS

Document Control	
Version	Version 1.1
Date Issued	December 2016
Author	David Caldwell
Comments to	david.caldwell1@nhs.net

Version	Date	Comment	Author
Version 1.0	December 2016	Initial Draft for consideration	David Caldwell
Version 1.1	February 2017	Updated Draft following SAG comments	David Caldwell

CONTENTS

1 INTRODUCTION

Statistical Process Control (SPC) charts can be a very useful tool in identifying variations in a process over time. There are two types of variation that we are interested in:

- **Special cause variation** is variation in a process which has an underlying cause.
  - It is due to irregular or unnatural causes that are not inherent in the design of the process
  - It reflects a 'different mode' of the process
  - It affects some, but not necessarily all aspects of the process
  - It results in an 'unstable' process that is not predictable
- **Common cause variation** is variation in a process which is down to random noise.
  - It is inherent in the design of the process
  - It reflects the 'business as usual' state of the process
  - It is due to regular, natural or ordinary causes
  - Affects all the outcomes of a process
  - Results in a 'stable' distribution that is predictable[1]

Variation in any process is normal and expected, and the main goal of SPC charts is to highlight this unnatural variation. Areas of special cause variation may be of interest to the user for quality control/improvement purposes and can be signals for underlying issues. A signal on an SPC chart does not necessarily mean that something is wrong, nor does it tell the user what may be wrong should a problem actually be present. SPC charts should be used as a springboard for further investigation only.

There are a wide range of different types of SPC charts, each with different underlying methodologies and calculations, but the interpretations should be consistent across these different types. The purpose of this paper is to give a brief overview of the different types of SPC chart, some guidance on the rules that should be applied to the charts, and how to interpret their output.

An interactive tool has been produced to automate the production of SPC charts for the user. This is an Excel-based tool which is hosted on the Quality Indicators pages of the ISD website.

2 TYPES OF STATISTICAL PROCESS CONTROL CHART

The type of SPC chart that should be used depends on the type of data, whether the process is static or dynamic (point-in-time comparisons or trends over time), and the underlying probability distribution the data is expected to follow. Choosing the correct chart type is an important step, as using the incorrect chart type may lead to incorrect conclusions, causing one of two types of statistical error:

- Type I Error: Incorrectly flagging a signal when there is no signal actually present (false positive).
- Type II Error: Failing to flag a signal when a signal is actually present (false negative).

Two types of SPC chart will be covered in this paper: run charts and Shewhart control charts.

[1] :Lloyd, R. (2004). *Quality Health Care: A Guide to Developing and Using Indicators*.

2.1 RUN CHART

The most basic type of chart that a user can plot is a run chart. A run chart is a simple graph where a process is plotted against time, with a measurement of central tendency, usually the median. Run charts have no upper or lower control limits, which can make them easier to produce than typical Shewhart charts. Their simplicity also means they lack the same level of statistical sensitivity when compared to a typical Shewhart chart.

The purpose of a run chart is to try and identify areas of non-random variation within a process. This can be done by highlighting any 'long' runs of consecutive data points that lie above or below the median line. This is not the only rule, however. A list of run chart rules agreed for adoption in NHS Scotland can be found in Table 1.

Run charts should only be plotted if 10 – 15 data points are available, as they may not be meaningful with fewer data points. This is mostly due to the instability of the centreline, which affects the decisions that are made, due to the low number of data points. To help ensure maximum effectiveness, data points should be collected as frequently as possible, but not to the detriment of the sample size. For graphs of percentages or rates each data point should be based on a denominator size of at least 10, if possible, to avoid high volatility (random variation) in the data caused by small numbers, as this could disguise any special cause variation.

Table 1: Run Chart Rules agreed for NHS Scotland

Rule	Description
Shift	A run of 6 or more consecutive data points above or below the baseline. Points on the centreline neither break nor contribute to a shift
Trend	A run of 5 or more consecutive data points which are increasing or decreasing. Where two or more data points are the same value ignore duplicates when counting
Too many/too few runs	**Too many or too few runs**: a run is a consecutive series of data points above or below the median. As for shifts, do not count points on the median: a shift is a sort of run. If there are too many or too few runs (i.e. the median is crossed too many or too few times) that's a sign of non-random variation. You need to look up a statistical table (see Perla et al, 2011) to see what an appropriate number of runs to expect would be. An easy way to count the number of runs is to count the number of times the line connecting all the data points crosses the median and add one.
Astronomical Data Point	A data point which is distinctly different from the rest of the process. As there are no upper or lower control limits, a judgement should be made by the user on what would constitute an astronomical data point

Table 1 is not an exhaustive list of run chart rules, as other rules can be applied that are appropriate to the data. Other rules should **only** be applied if they are agreed to be appropriate for the data prior to the chart's production. The rules in table 1 should be sufficient for the vast majority of charts and should only be expanded upon for special cases.

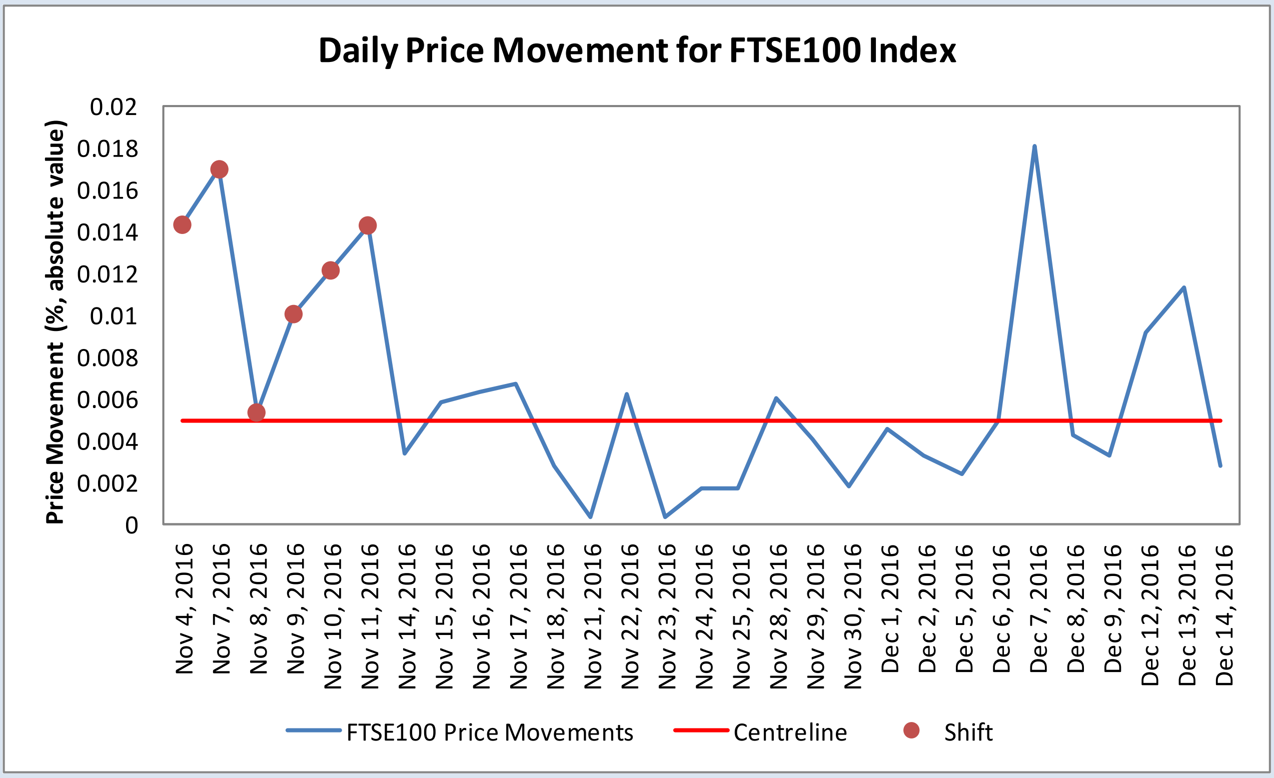

Figure 1: Example run chart of daily price movements for the FTSE100 Index

Figure 1 shows an example run chart: the absolute value of the daily price movement (%) of the FTSE100 Index from 4 November 2016 to 14 December 2016. The data features an upward shift in the first 6 days of the process. No trends or astronomical data points are present. Dec 7 could possibly have been considered astronomical, had some of the early November data points been closer to the median.

Run charts are a good chart to choose if the user is looking for something quick, simple and basic, but their lack of statistical power makes them less able to reliably identify areas of special cause variation (i.e. there will be more type II errors with run charts) than Shewhart charts, so if the user is looking for something that is more robust, a run chart is unlikely to be sufficient. However, as a starting point it is always better to plot a run chart first before any other type of SPC chart.

2.2 SHEWHART CONTROL CHARTS

Shewhart Control Charts are better at reliably identifying areas of special cause variation within a process. The inclusion of upper and lower control limits allows a degree of measurement of distance from the centreline. Unlike run chart rules, these limits also remove the element of subjectivity in judging whether a point is an astronomical data point or not as points are either outliers or they are not. As mentioned before, there are many different types of chart available, and the correct chart to use depends on the type of data being collected. A list of chart types and a brief summary of the types of data they serve is found in Table 2 (see section 2.4 and section 2.5 for examples of their application).

Table 2: Types of SPC chart

Chart Name	Appropriate For
XmR Chart	Continuous data (see below for definition), where each data point is an individual measurement and not an aggregate or average of multiple data points
Xbar & S Charts	Continuous data, where each data point is an average of multiple data points
P Chart	Discrete data. Data should be a proportion/percentage of successes/failures
C Chart	Count data, where the “area of opportunity” is equal for each data point
U Chart	Count data, where the “area of opportunity” is unequal for each data point
T Chart	For rare events data, recorded as time between unlikely events (continuous).
G Chart	For rare events data, recorded as the number of events between unlikely events (discrete).

NOTE: Table 2 is not an exhaustive list and is not intended to be prescriptive. A judgement on the type of chart to use should be made depending on what is appropriate for the data and what measures stakeholders are tracking. This should be used as a guide only.

Continuous data can occupy any value within a range. Typically a measurement of some kind, they are unrestrained by “categories”, and can be infinite in range. Examples of continuous data would be measurements of height, weight or length; monetary values and by convention, anything which can be considered workload or throughput is typically handled as continuous data.

Discrete data are observations where each possible observation is distinct from one another. It can be categorical e.g. number of red balloons handed out a fair; or numeric, typically count data e.g. the number of falls in a care home each month; or binomial data e.g. how many red balloons are handed out each hour. Typically, when putting a binomial process into an SPC chart, the data should be presented as a proportion of successes, or a proportion of failures.

Area of Opportunity is the number of possible “defects” per sample. E.g. with data which are the number of care home residents who have had a fall each month, the area of opportunity would be the total number of residents

There is an alternative school of thought regarding Shewhart Control Charts, which proposes that the XmR chart is sufficient for all data and the wider array of charts do not provide enough of a benefit to justify the risk of misunderstanding or misapplying a chart due to their complexity. The pros and cons of this approach can be found in Section 3.

Figure 2 is a flow chart which illustrates the decision making process for choosing the correct chart type.

Formulae for calculating control limits and centrelines for all charts found in table 2 can be found in Appendix A.

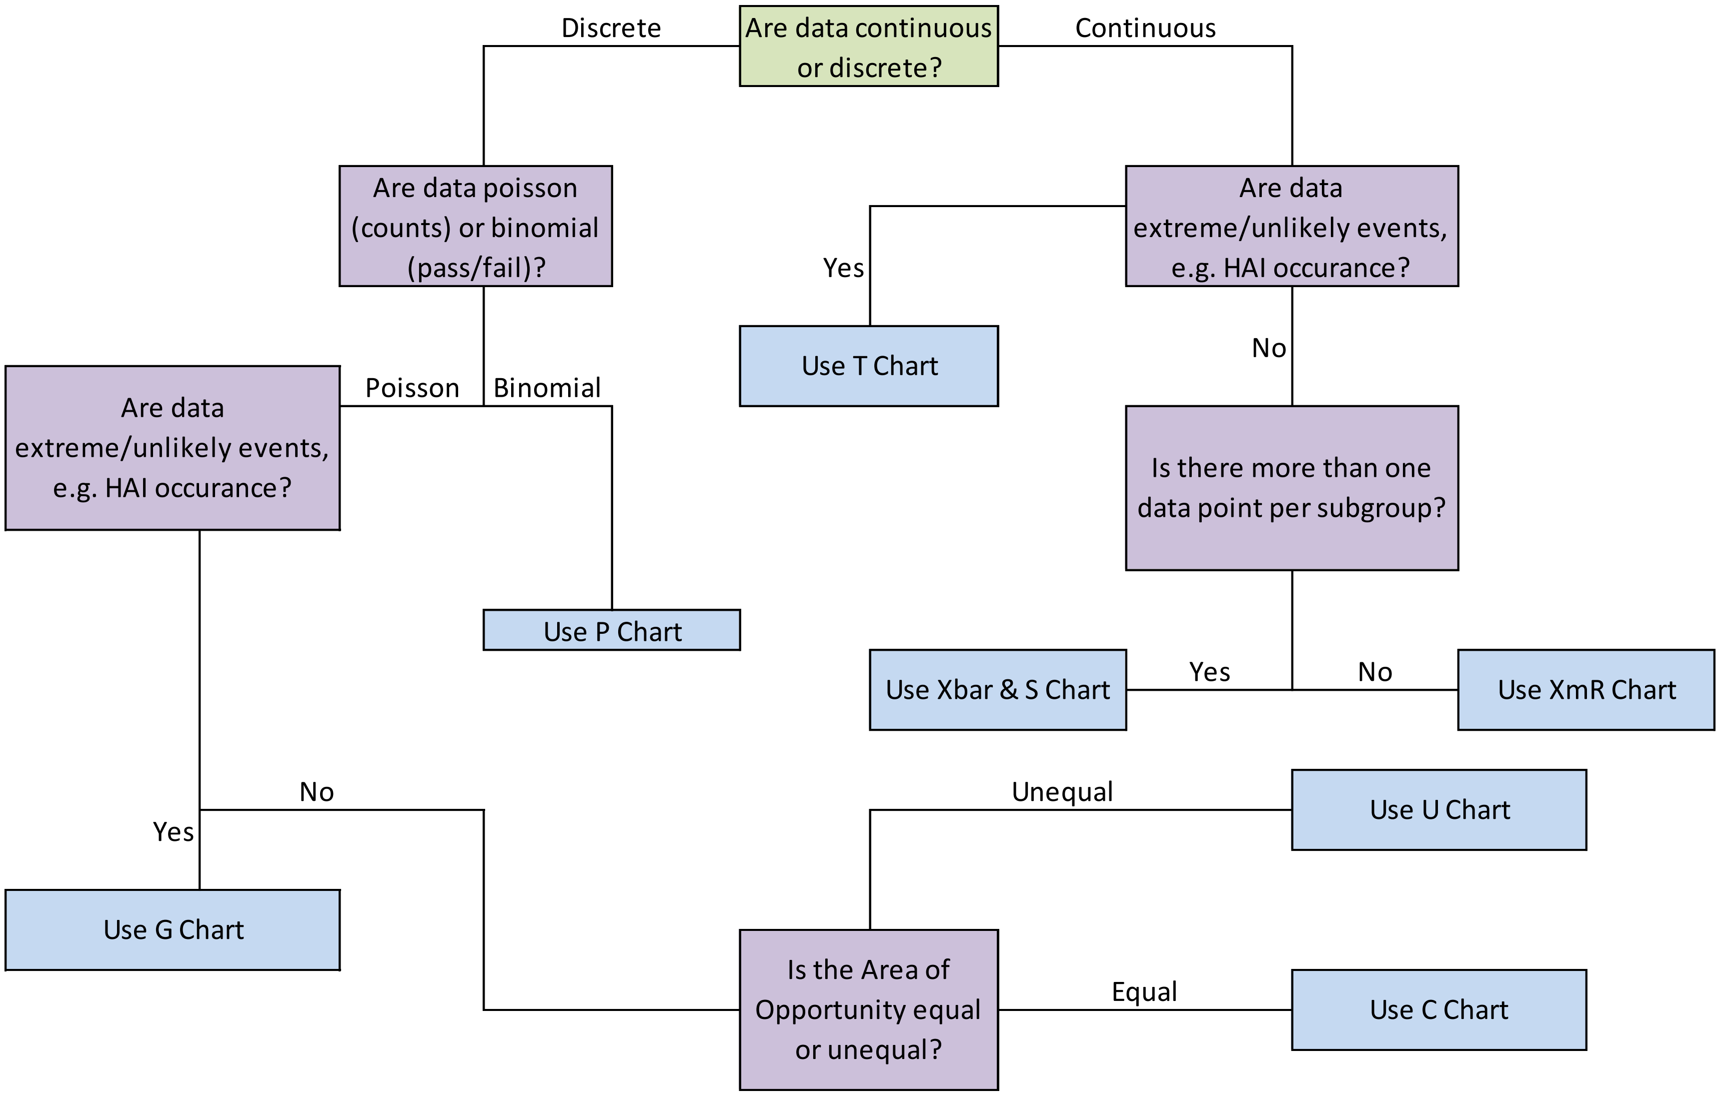

Figure 2: Flow chart illustrating decision-making process for choosing the most appropriate type of Shewart Control Chart

2.3 SHEWHART RULES

In order to interpret data from a control chart correctly, a set of consistent rules must be applied to highlight any special cause variation. These rules highlight data points which may be of interest in order to prompt the user of the data to investigate further. These rules should be agreed upon before creating the control chart based on the user's knowledge of the data, as creating the chart first could lead to a bias from the user who may feel compelled to choose a set of rules which show something more in line with what they want to see.

Table 3 contains standard **Shewhart Control Chart Rules** for identifying special cause variation agreed for Improvement Programmes in NHS Scotland.

Table 3: NHS Scotland Shewhart Control Chart Rules

Rule	Description
Outlier	Data point(s) more than three sigma from the centreline
Shift	A run of 8 or more consecutive data points above or below the centreline
Trend	A run of 6 or more consecutive data points. Where two or more data points are the same value ignore duplicates when counting
Outer One – Third	Two out of three consecutive data points which sit close to one of the control limits (within 2 and 3 sigma)
Inner One - Third	15 or more consecutive data points that lie close to the centreline (within 1 sigma).

Again, these are standard rules which should be sufficient for the vast majority of charts but they can be adapted accordingly **only if** it is appropriate for the data and any changes have been agreed beforehand.

2.4 CONTROL CHART EXAMPLES

This section will include 3 examples of SPC charts. Each example will present a dataset, and follow the flow chart for deciding which chart to use step-by-step.

Example 1
Hospital Standardised Mortality Ratios are calculated for all SPSP-participating hospitals in Scotland. SPC charts can be very useful for this measure to identify areas of unnatural variation.

Is the data discrete? – No.
HSMR is a calculated measure which can take any value greater than 0.

Is there more than one observation per subgroup? – No.
There is only one measurement each quarter calculated for each hospital.

As the data are continuous and consist of individual measurements, the preferred chart for this would be an XmR chart. If HSMR was being calculated and then averaged for multiple hospitals each quarter, then an Xbar chart would be suitable, with the process being the average of each measurement across all sites (note: the number of sites should remain constant).

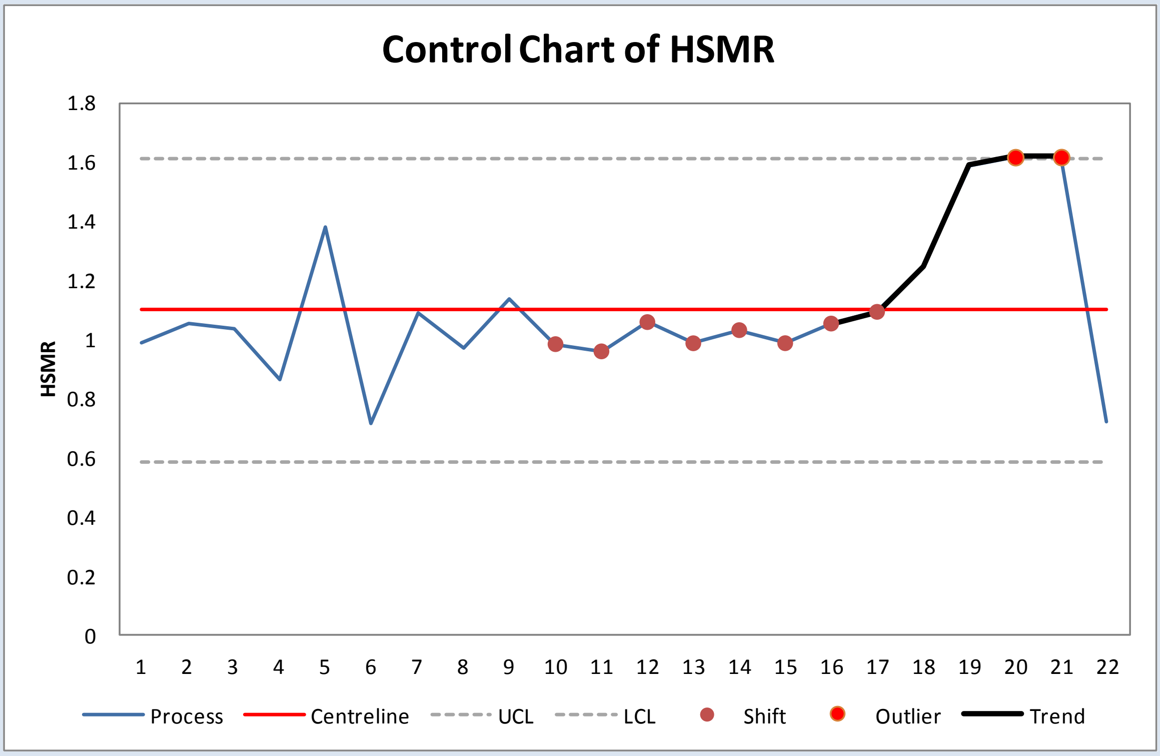

The chart above shows that the process looked relatively stable until quarter 19, where there was a large spike, followed by two outlying quarters. This likely pulled the centreline up slightly, leading to the downward shift in the centre. So, additional investigation of the outlying quarters would be necessary.

Example 2

You have a large jar of jelly beans, and from the jar you take 25 scoops of beans. You measure the number of green beans in each scoop and wish to monitor this process with an SPC chart.

Is the data discrete? – Yes.
Are there only two possible outcomes? – Yes (green or not green).

As the data are binomial (acceptable/unacceptable), then the best chart for this type of data is a P chart. The best way to plot this process is with the proportion of successes to the number of overall beans. In this case, the process will be the proportion (percentage) of green beans in each scoop.

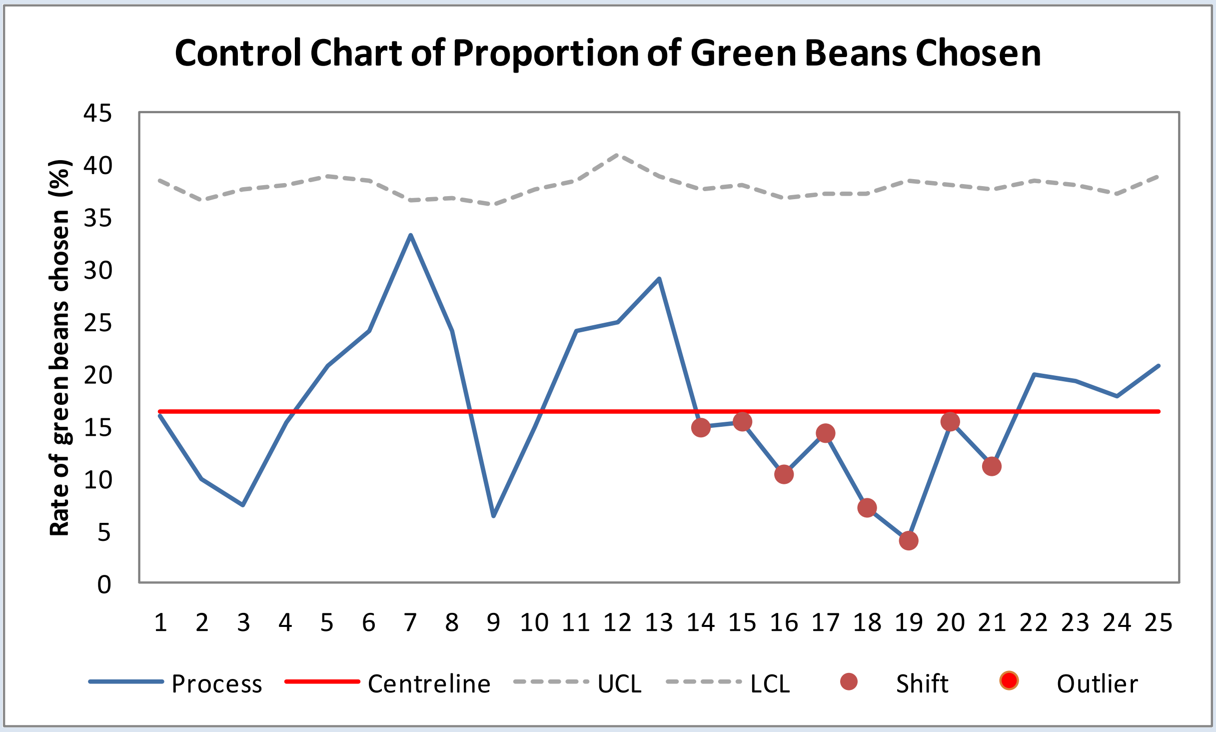

The variable control limits come from the variable size of each sample. Scoops with more beans will have tighter limits as the predicted proportion of green beans obtained by random chance should tend towards the true mean (which the centreline is an estimate of) when the sample size increases.

The above P chart shows no outliers, but does show a downward shift. There is also no lower control limit as the calculated limit was less than 0, and a negative percentage is impossible to achieve in this context.

Example 3

A factory that produces computer accessories takes a sample of 10 keyboards every hour and counts the number of loose keys per keyboard, but they take a random sample across all models of keyboard they produce.

Is data discrete? – Yes.
Are there only two possible outcomes? – No. The data is a count of loose keys, not a classification of whether or not the keyboard had a loose key or not.
Is the area of opportunity variable? – Yes. Different keyboards have different numbers of keys which could be loose, and each sample has a different combination of potentially loose keys.

As the data in this example are poisson counts, a C or U chart would be suitable, and as there is room for one keyboard to have more loose keys than another, the chart used should be a U chart.

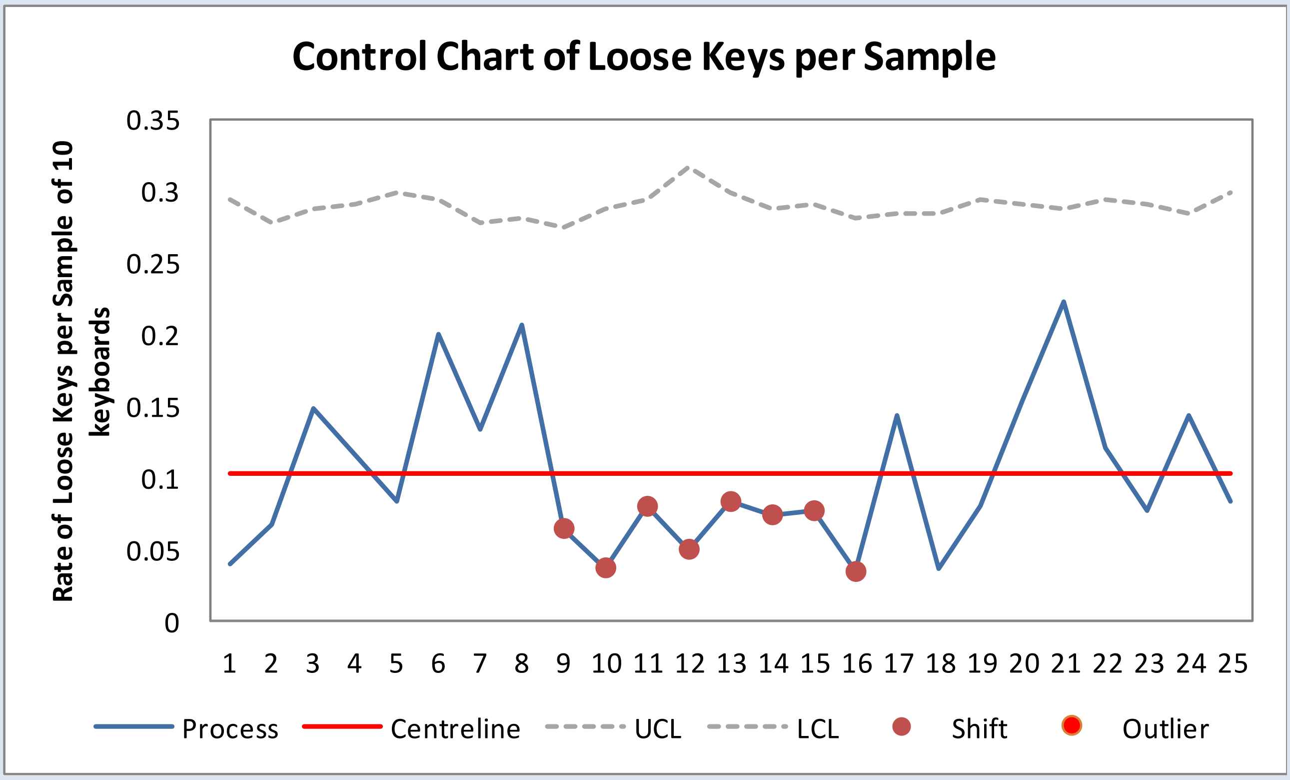

The mean number of defects (unadjusted for sample size) found per sample is 2.7 (~0.1 per 100 keys per sample). Most samples fluctuate closely around this number and no samples are out of control, however there is a downward shift from sample 9 to sample 16, indicating that further investigation is required to determine what was happening there.

Example 4

A care home monitors how many of its residents have had a fall each month, and wants to produce an SPC chart to ensure the issue is kept controlled.

Is data discrete? – Yes.
Are there only two possible outcomes? – No. The data is a count of residents who have had a fall.
Is the area of opportunity variable? – No. The number of residents in this example is constant each month.

As the data in this example are Poisson counts, a C or U chart would be suitable, and as the number of residents in the home is constant month-to-month, a C chart is appropriate here.

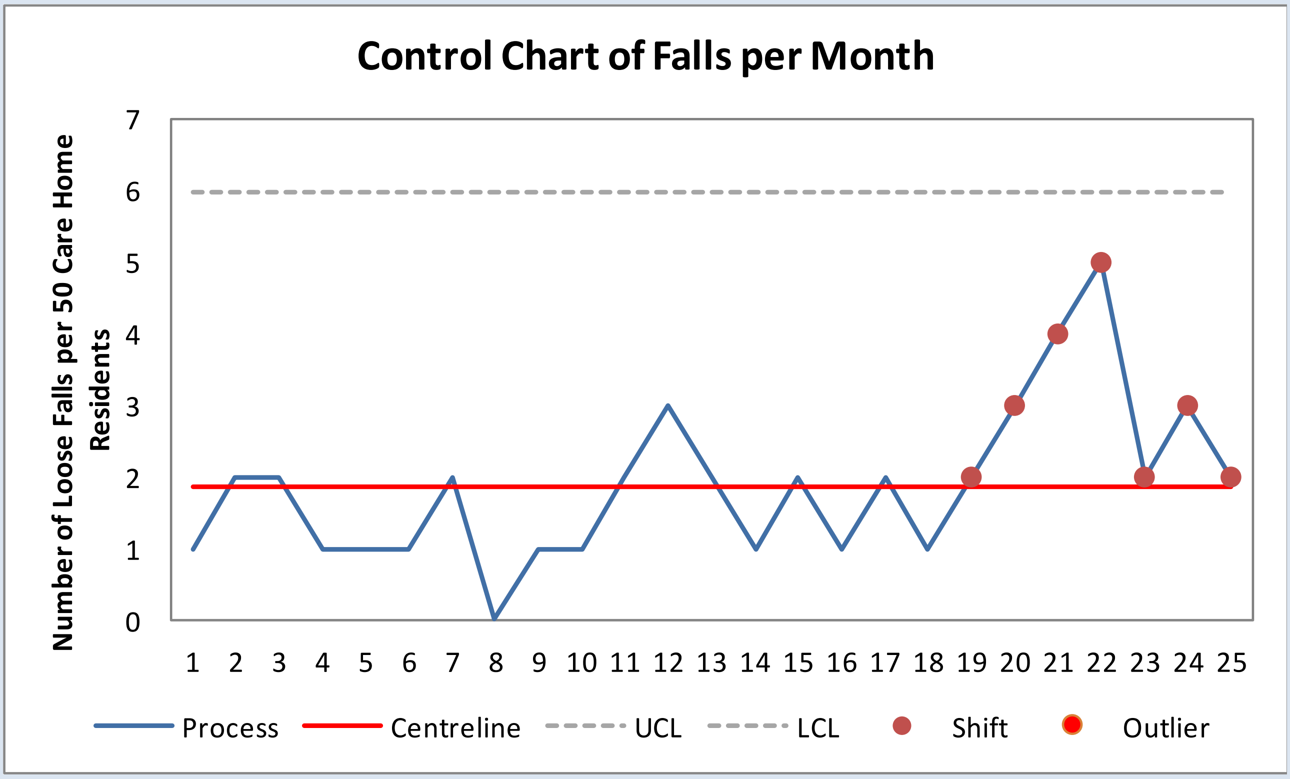

The mean number of defect falls found per sample is around 2. Most samples fluctuate closely around this number or below, however an upward shift is present in the final 6 data points.

NOTE: these data could also be presented as a percentage (number of residents who fell/total number of residents).

2.5 FURTHER EXAMPLES

Data	Flow Chart Process	Type of SPC to use
Number of patients that go through a GP surgery each week	• Is data discrete? No. Data are throughput which, by convention, are considered continuous. • Is there more than one data point per subgroup? No. Data are single measurements.	XmR chart.
Percentage of patients that died within 30 days of admission to hospital	• Is data discrete? Yes. • Is data poisson or binomial? Binomial (patients can only be dead or alive after 30 days)	P chart.
Number of morbidities patients report when admitted to A&E departments each day	• Is data discrete? Yes. • Is data poisson or binomial? Poisson (data are counts of morbidities) • Is the area of opportunity equal? No. The number of patients may vary day-to-day and the number of morbidities each patient can report is theoretically infinite.	U chart.
Average daily spend on antibiotics prescribed by GP practices across a health board	• Is data discrete? No. Monetary values are continuous. • Is there more than one datapoint per subgroup? Yes. Each GP practice contributes towards the health board average.	Xbar & S chart.
Percent of hospitals across the UK which make the target of 95% of patients seen within 4 hours in A&E	• Is data discrete? Yes. • Is data poisson or binomial? Binomial.	P chart, but data must first be expressed as a proportion.
The number of days between hospital acquired infections (HAI) for a single hospital site	• Is data discrete? No. Data are time measurements. • Are data unlikely events? Yes.	T chart.
The number of patients prescribed a certain drug between two adverse drug events	• Is data discrete? Yes. • Are data unlikely events? Yes.	G chart.

3 THE XMR CHART

As mentioned previously in Section 2.2, there is an alternative school of thought around Shewhart chart selection which proposes only an XmR chart is necessary. This section lays out the arguments for and against this approach. The decision over which approach to take is entirely down to the user and the decision should be made before producing any SPC chart.

	Wider range of Shewhart Charts	XmR Charts
Advantages	• When applied properly, these charts are correct and provide more statistically appropriate results • The control limits vary according to the size of each individual sample	• They are simple to understand • In most improvement situations, the results can be deemed 'appropriate enough' • Does not require any assumptions about underlying data distributions • Quick and easy to apply to practical situations • They are recommended by influential members of the quality improvement community[2]
Disadvantages	• Assumptions about underlying data distributions may not be met[3] • Lack of understanding of these probability distributions can lead to errors or lack of user confidence in choice of chart • The greater complexity of these charts makes understanding and automation more difficult	• They are not correct for all types of data • A large portion of the quality improvement community does not recommend them over the wider range of Shewhart Charts.

These advantages and disadvantages are for information purposes only and are not intended to try and sway a user towards one approach or the other. The decision to use just an XmR chart is down to the user(s) and should be agreed upon before production of any chart begins.

4 SUMMARY

Choosing the correct chart is an important step in monitoring a process, and ensures that incorrect conclusions around process stability are not drawn. However, chart selection is not an exact science and there is room for manoeuvrability. It is also important that any decisions regarding chart selection and rules are made prior to producing the chart, as making these decisions retrospectively can lead to a bias towards a chart which is more agreeable with what the user would like to see.

Run charts and Shewhart control charts are not the only type of SPC chart available, and a process over time is not the only type of data which can be put into an SPC chart. Data which are point-in-

[2] Stauffer (2010), Balestracci (2014)
[3] Wheeler (1996)

time comparisons, e.g. crude mortality rates by hospital for one quarter in Scotland could be put into an appropriate funnel plot to be compared against the national average. Like Shewhart control charts, there are a number of approaches to funnel plots (based on Xbar+S, P or U charts) depending on the data.

5 FURTHER INFORMATION

There are a wide range of online resources on the use and calculation of SPC charts. As mentioned in the introduction, the Quality Indicators team have produced an interactive tool in producing SPC charts. The tool takes the user through the selection process step-by-step and decides what the best chart is to use.

The Quality Improvement hub has pages dedicated to Shewhart control charts and run charts. The Healthcare Data Guide[4] is an excellent reference guide for this type of analysis.

6 ACKNOWLEDGEMENTS

This document was prepared by the Quality Indicators Team from both knowledge of the topic within PHI and reference material. This document was circulated to individuals with Quality Indicators and wider PHI for comments and feedback. Their input was much appreciated and we are grateful for their advice. Any error or lack of clarity that remains is, however, our responsibility.

7 CONTACT INFORMATION

If after reading this document you have any unanswered queries on SPC charts, or would like to discuss any aspects further, please contact the Quality Indicators Team, via NSS.isdQualityIndicators@nhs.net.

We would also welcome any feedback from you on this document.

[4] See http://eu.wiley.com/WileyCDA/WileyTitle/productCd-0470902582.html

8 APPENDIX A: HOW TO CALCULATE CENTRELINE AND CONTROL LIMITS

8.1 XmR Chart

$$\bar{x} = \frac{\sum_{i=1}^{m} x_i}{m}$$

$$MMR = \frac{\sum_{i=2}^{m} |x_i - x_{i-1}|}{m-1}$$

$$Control\ Limits = \bar{x} \pm 2.66 * MMR$$

Where x_i is the ith datapoint and m is the number of datapoints.

8.2 Xbar and S Chart

$$\bar{x} = \frac{\sum_{i=1}^{m} \sum_{j=1}^{n} x_{ij}}{mn}$$

$$Control\ Limits = \bar{x} \pm 3 \frac{\sum_{i=1}^{m} \sqrt{\frac{\sum_{j=1}^{n} \left(x_{ij} - \bar{x}\right)^2}{n-1}}}{m}$$

Where x_{ij} is the ith datapoint from group *j* and m is the number of groups and *n* is the number of data points per group.

8.3 P Chart

$$\bar{p} = \frac{\sum_{i=1}^{m} \sum_{j=1}^{n} \begin{cases} 1 & if\ x_{ij} is\ defective \\ 0 & otherwise \end{cases}}{\sum_{i=1}^{m} n_i}$$

$$Control\ Limits = \bar{p} \pm 3 \sqrt{\frac{\bar{p}(1-\bar{p})}{n_i}}$$

Where x_{ij} is the ith datapoint from group *j* and m is the number of groups and *n* is the number of data points per group.

8.4 C Chart

$$\bar{c} = \frac{\sum_{i=1}^{m} x_i}{m}$$

$$Control\ Limits = \bar{c} \pm 3\sqrt{\bar{c}}$$

Where x_i is the total number of "defects" from group *j* and m is the number of groups.

8.5 U Chart

$$\bar{u} = \frac{\sum_{i=1}^{m} x_i}{mn}$$

$$Control\ Limits = \bar{u} \pm 3\sqrt{\frac{\bar{u}}{n}}$$

Where x_i is the total number of "defects" from group *j*, m is the number of groups and *n* is the sample size for each group.

8.6 T Chart

$$T_i = t_i^{\left(\frac{1}{3.6}\right)}$$

$$MMR = \frac{\sum_{i=2}^{n} |T_i - T_{i-1}|}{n}$$

$$\bar{T} = \frac{\sum_{i=1}^{n} T_i}{n}$$

$$\bar{t} = \bar{T}^{3.6}$$

$$Upper\ Control\ Limit = \left(\bar{T} + (2.66MMR)\right)^{3.6}$$
$$Lower\ Control\ Limit = \left(\bar{T} - (2.66MMR)\right)^{3.6}$$

Where t_i is the time between the *i*th event and *n* is the number of events.

8.7 G Chart

$$\bar{g} = \frac{\sum_{i=1}^{k} g_i}{k}$$

$$Upper\ Control\ Limit = \bar{g} + 3\sqrt{\bar{g}(1+\bar{g})}$$
$$Lower\ Control\ Limit = \bar{g} - 3\sqrt{\bar{g}(1+\bar{g})}$$

Where g_i is the number of occurrences between the *i*th event and *k* is the number of events.