Paper 0290-2017

Learning from Quality Improvement Data: Introduction to Statistical Process Control Charts

Ruth Croxford, Institute for Clinical Evaluative Sciences, Toronto, Ontario, Canada

ABSTRACT

Data is your friend. This paper discusses the use of data for quality improvement (QI). Measurement over time is integral to quality improvement, and statistical process control charts (also known as Shewhart or SPC charts) are a good way to learn from the way measures change over time in response to quality improvement efforts. The paper explains what an SPC chart is, how to choose the correct type of chart, how to create and update a chart using SAS®, and how to learn from the chart. The examples come from QI projects in health care, and the material is based on the Institute for Healthcare Improvement's Model for Improvement. However, the material is applicable to other fields, including manufacturing and business.

This paper is aimed at people newly considering a QI project, people who want to graph their data and need help with getting started, and at anyone interested in interpreting SPC charts created by someone else.

INTRODUCTION

Whether you are conducting a Quality Improvement (QI) project, or simply monitoring a process, it is important to track and learn from the behavior of measures over time. Statistical Process Control Charts (also referred to as Shewhart charts, or SPC charts) are a simple-to-use visual presentation of performance over time. While the interpretation of Shewhart charts is based on statistical theory, you don't need to know any statistics to use them. A small set of rules allow you to differentiate between random fluctuations in the measures and true "signals", helping you to react when something has changed, but also helping you refrain from reacting to random "noise".

Quality Improvement projects pass through a series of stages. After you have identified a set of measures that you hope to improve, baseline data is collected, either retrospectively or prospectively, in order to determine the current state of affairs – the starting point for your improvement efforts. Once the baseline level of performance has been established, you're ready to introduce a change that you hope will result in an improvement. Over the life of the QI project, you'll probably introduce several changes, collecting data after each change to see if the change has resulted in an improvement. Once the project is over, you may continue to collect data in order to insure that that the improvements are sustained. Shewhart charts are used throughout the life of the project – first to establish the baseline values, then to determine whether each change really has resulted in an improvement, and finally, to guard against backsliding.

Even in the absence of a QI project, Shewhart charts allow you to monitor processes so that you can react as quickly as possible if performance deteriorates, allowing you to identify improvements in performance so that they can be sustained and also protecting you from wasting time and energy overreacting to what may simply be random fluctuations in performance.

In this brief introduction to Shewhart charts, I'll describe what they are, explain the simple rules for interpreting them, and introduce SAS's® SHEWHART procedure. Much more information is available from the references listed at the end of the paper.

Please note that all of the data shown in this paper are fictitious, created with the help of a random number generator strictly for the purpose of illustration

INTRODUCING THE SHEWHART CHART

Figure 1 illustrates the basic parts of a Shewhart chart. This sample chart presents the proportion of patients discharged from a hospital who were subsequently readmitted within 30 days. The horizontal, or X-axis, of the graph is time. This particular graph happens to depict data collected monthly, but data can be collected weekly or daily or data can be collected on individual consecutive patients. The X-axis will almost always show the passage of time because the goal is to understand how a measure is changing over time. The vertical, or Y-axis, shows the measure (in this case, proportion of patients readmitted).

Each Shewhart chart contains a center line, drawn at the average value of the values on the graph. In this graph, the average value is 0.143, meaning that on average, 14.3% of the patients discharged each month are readmitted.

Two additional lines appear on the Shewhart chart: the lower control limit (LCL) and the upper control limit (UCL). By default, SAS® has shaded the area between the lower and upper control limits. This area can be thought of as the range within which the measure will fluctuate due to random variation.

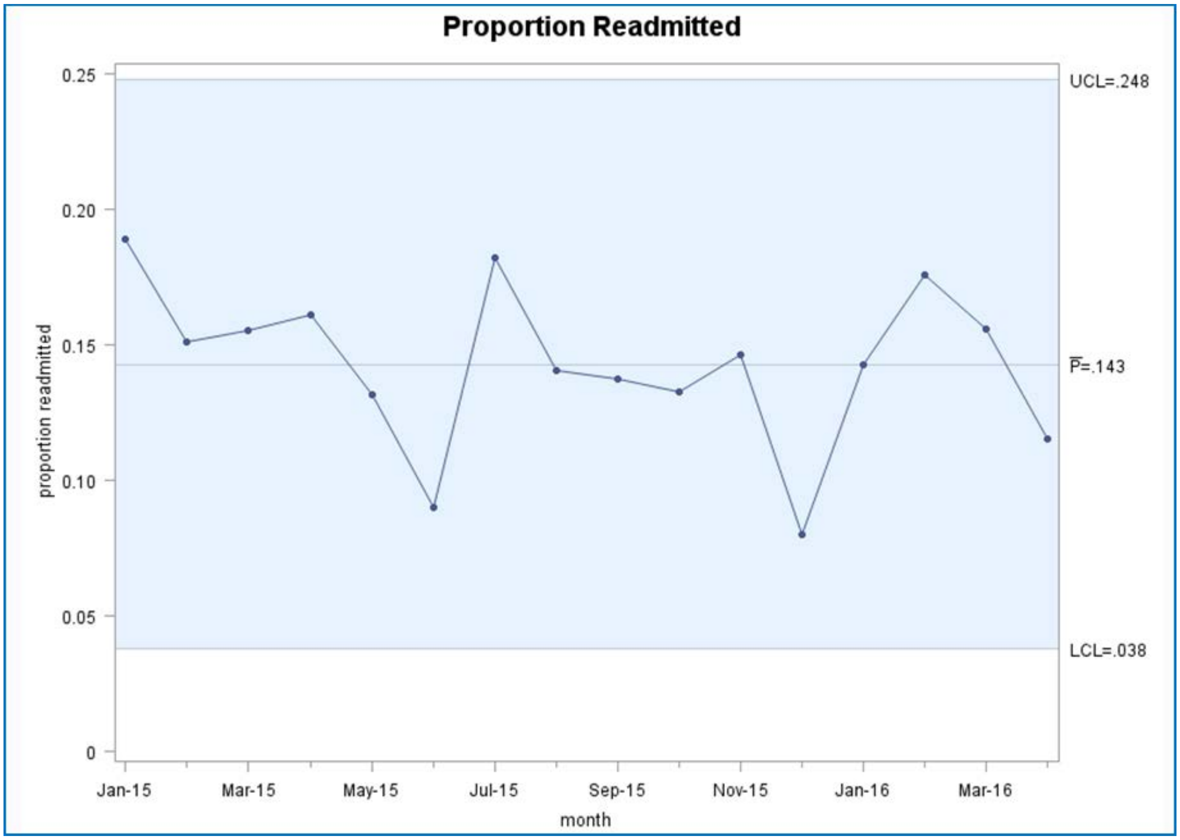

Figure 1. Sample Shewhart Chart

A SHORT DETOUR: TYPES OF DATA

There are many types of data, and many types of Shewhart charts. Only three types of data will be discussed here:

- Continuous: a continuous measure is one that can take on any value within a range of plausible values. Examples of continuous measures are hospital length of stay, waiting time to see a doctor, and blood pressure and other biological measurements. If you're trying to decide whether a particular measure is continuous, don't obsess too much. Yes, length of stay is typically measured in days or hours, and does not really take on "any value within a range of plausible values".

- Classification: if there are only two possible outcomes, the measure is a classification variable. Examples of classification variables are hospital readmissions (a patient is either readmitted or not), and purchase returns (merchandise is either returned or it is not). If you have measured a percentage, you have a classification measure.
- Count: Count data arise when you have counted the number of discrete events (termed "nonconformities") per "opportunity". What differentiates count data from classification data is that it must be possible for the event being counted to occur more than once per "opportunity". Medication errors is a count variable because each time a medication is administered (the "opportunity"), it is possible to make more than one error (e.g., both the wrong drug and the wrong dose).

The purpose of a Shewhart chart is to detect unexpected changes in a measure. Generally, continuous variables are more sensitive than count variables, and count variables are more sensitive than classification variables. Therefore, if there is more than one way to capture the same information, you will generally be able to identify improvements faster if you use counts rather than classifications. Some examples of count vs. classification measures are shown in Table 1.

Count	Rate	Classification	Percentage
Number of aggressive behaviors exhibited by residents in a long-term care residence	Behaviors per resident over the course of the time period	Did a given resident exhibit aggression during the time period (yes/no)	Percentage of residents exhibiting aggressive behavior at least once during the time period
Number of visits to the emergency department made by long-term care residents	ED visit per 100 residents during the time period	Did a given patient visit the ED at least once during the time period (yes/no)	Percentage of residents with at least one ED visit during the time period
Number of medication errors (each time a medication is dispensed more than one error can be made)	Error rate per 1,000 doses	Was an error made (yes/no)	Percent of doses with at least one error
Number of infections	Infections per patient-day	Was there at least one new infection on a given day (yes/no)	Percent of days with at least one new infection.

Table 1. Some examples of measures which can be expressed either as a count or a classification variable. When there is a choice, using a count variable will generally detect changes faster.

BACK TO THE SHEWHART CHART

We are now ready to take a second look at the Shewhart chart shown in Figure 1. I've already mentioned that the center line is drawn at the mean, or average value ($CL = \mu$, where the Greek letter μ is used to signify the mean value). The lower control limit is drawn 3 sigma (σ) below the center line ($LCL = \mu - 3\sigma$); the upper control limit is drawn 3 sigma above the center line ($UCL = \mu + 3\sigma$). The value of sigma depends on both the type of measure being graphed (continuous, classification, count) and the number of patients or medication doses, etc. that were used to calculate each point on the graph. If you've taken a Statistics course, you've probably already guessed that the calculation of sigma assumes that count data follows a Poisson distribution, classification data follows a binomial distribution, and continuous data follows a normal distribution. And if you haven't taken a Statistics course, don't worry. You don't need this information to use the charts.

Despite the notation, sigma is *not* a standard deviation. The definition is similar to that of a standard deviation, but not exactly the same.

Because the lower and upper control limits are placed three sigma away from the center line, the area between the two limits can be divided into 6 bands, as shown in Figure 2.

Outt

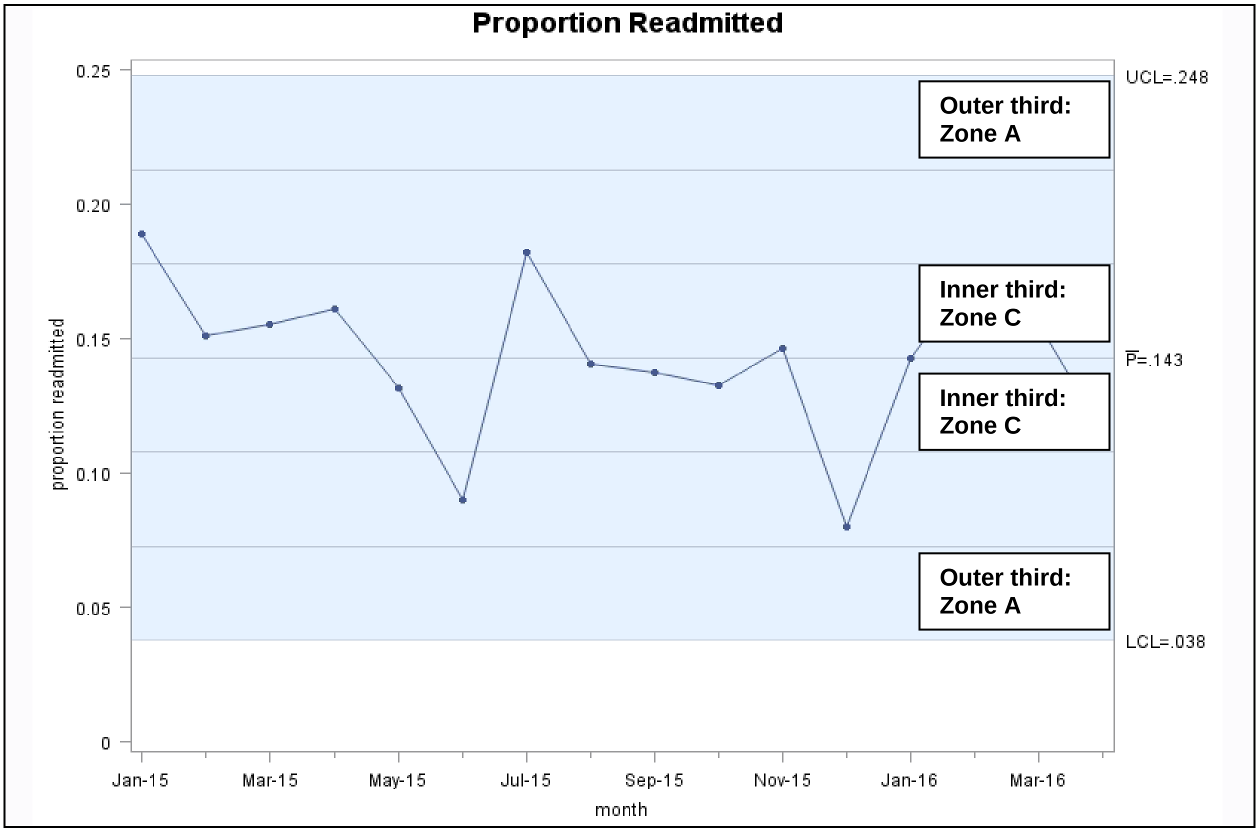

Figure 2. Shewhart chart, showing the location of the 1- and 2-sigma lines. The lines located 3 sigma away from the center line are referred to as the lower and upper control limits. The outer third of the graph refers to the two sections furthest away from the center line; the inner third of the graph refers to the two sections adjacent to the center line.

I mentioned that the value of sigma depends on the number of patients observed during each time period. In Figure 1 the lower and upper control limits are flat. However, if the number of patients discharged varies from month to month, the control limits will be jagged, as shown in Figure 3. In a month such as August 2015, when the number of discharges is high, the lower and upper control limits are closer together, indicating that the data point for that month has more certainty (it is more precisely measured). In a month such as February 2016, when the number of discharges is low, the lower and upper control limits are further apart, indicating that the data point for that month has less certainty (it is less precisely measured).

SELECTING THE TYPE OF SHEWHART CHART

You don't need to know the equations for sigma. In order to obtain the correct lower and upper control limits, you just need to tell the graphing software what type of data you are graphing. Unfortunately, it is not quite as simple as saying, "I have classification data". Rather, each type of data is associated with one or two chart types. For example, if your measure is a classification variable, request a P chart in order to obtain the correct graph.

Figure 4 shows the association between the data type and the choice of Shewhart chart for five scenarios:

- Count data, when the number of "opportunities" is the same for each time period: C chart
- Count data, when the number of "opportunities" differs from time to time: U chart. A U chart can also be used when the number of opportunities is the same for each time period. The only reason for using a C chart is that you don't know or don't want to specify the number of opportunities.
- Classification data: P chart
- Continuous data, when each data point represents a single value (when all of the data for the time period has been aggregated into a single average value or when you are graphing data separately for each consecutive patient): I chart
- Continuous data, when the data for each time period has not been aggregated: X-bar & S chart

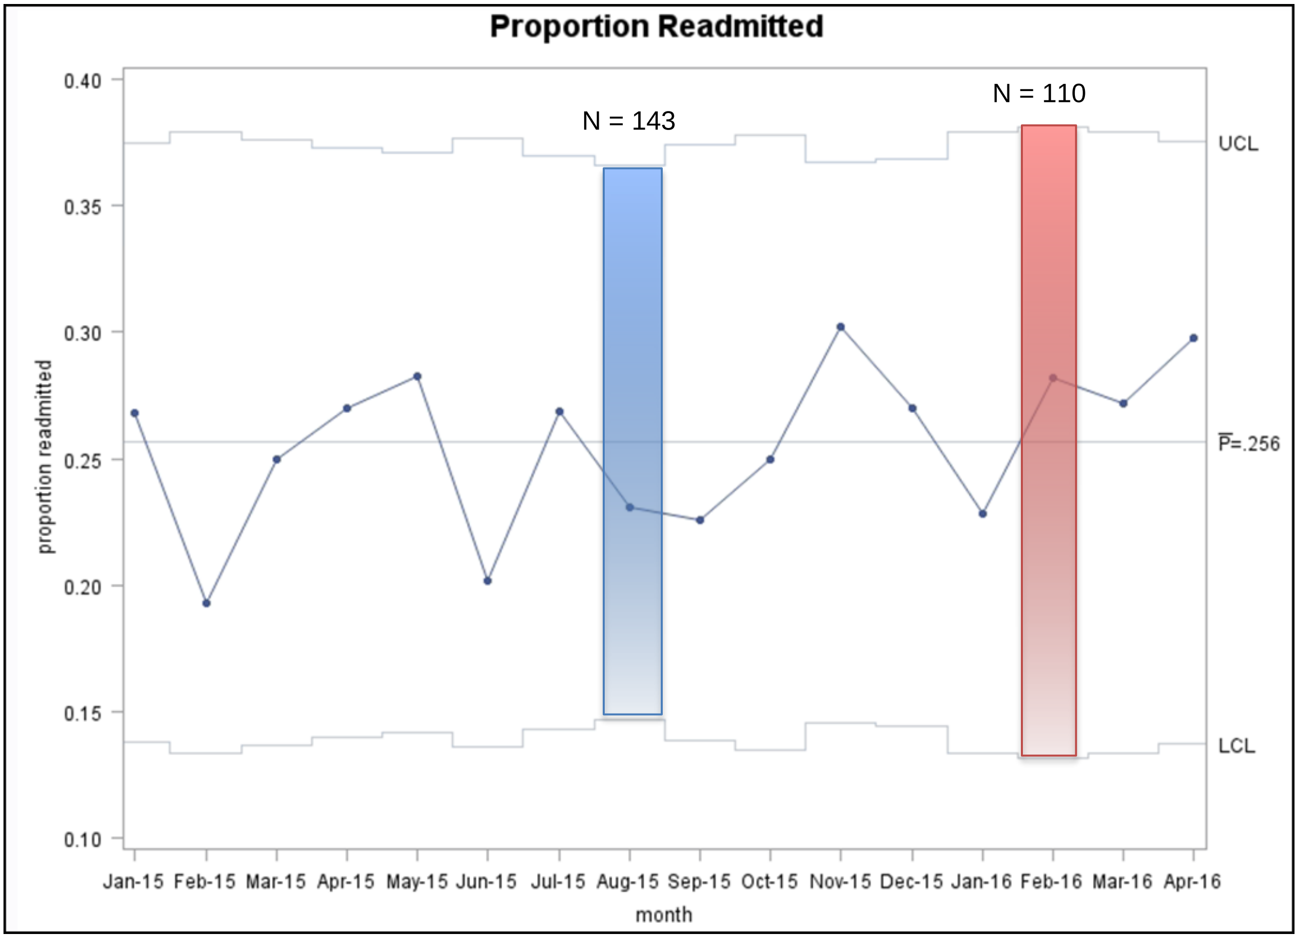

Figure 3. When the number of units examined is different for each period the control limits will be jagged. When the number of units is larger, the control limits will be closer together, reflecting the greater precision in the measure; when the number is smaller, the control limits will be further apart.

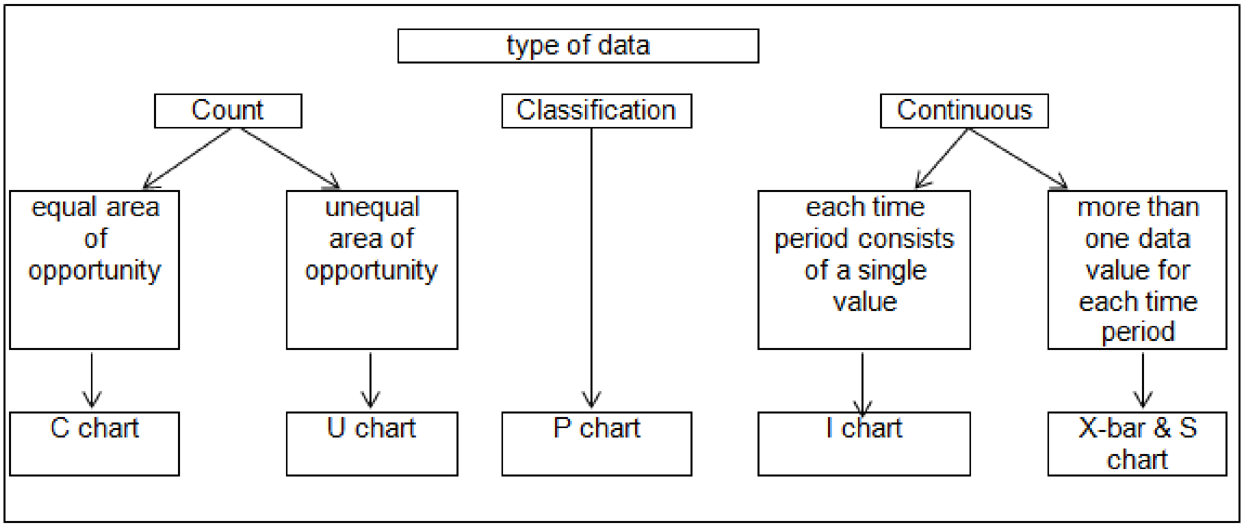

Figure 4. Relationship between type of data and type of Shewhart chart.

To make the types of chart more concrete, Figure 5 shows what the data might look like for each type of chart. Parts A, B, and C show (fictitious) data for a long-term care residence. The measure concerns transfers of residents to the Emergency Department (ED). A given resident may be transferred more than once during a period. In part A, only the number of transfers has been recorded. The assumption is that the total number of residents is approximately the same from week to week. Since our long-term care home is generally operating at full capacity, this assumption is good enough for QI purposes. These data can be graphed using a C chart. In part B, we have also recorded the number of residents for each week (i.e. the number of "opportunities" for a transfer to occur). These data can be graphed using a U chart. In part C, we did not count the number of transfers. Instead, each resident was classified as having had at least one transfer, yes or no. Since this is a classification measure, a P-chart is used.

Parts D and E show (fictitious) data for a physician's office which has had complaints about the waiting time to see the doctor. The receptionist recorded the wait time in minutes for every 10th patient. In part D, the receptionist averaged the times for each day. Wait times are continuous data. Because there is only one data value for each day – that being the average wait – an I chart will be used to graph the data. Part E shows the same data, but this time the receptionist has retained the individual wait times. Notice that there do not need to be the same number of data points for each time period. On day 1, 6 values (patients 1 – 6) were recorded. On day 4, only 3 values were recorded. Because there are multiple measurements for each time period, an X-bar & S chart can be used. This chart makes the best use of the data, because we can examine not only the daily averages, but also the variation in the values from day to day.

Week	Number of transfers to the ED
1	13
2	15
3	17
4	15
5	13
6	18
7	10
8	15
9	13
10	17
11	19
12	13
13	11
14	13
15	16
16	18
17	15
18	19
19	10

A

Week	Number of transfers to the ED	Number of residents
1	13	143
2	15	140
3	17	145
4	15	139
5	13	142
6	18	135
7	10	137
8	15	142
9	13	135
10	17	133
11	19	137
12	13	141
13	11	139
14	13	140
15	16	145
16	18	139
17	15	142
18	19	147
19	10	143

B

Week	Number of residents transferred to the ED at least once	Number of residents
1	11	143
2	14	140
3	12	145
4	15	139
5	9	142
6	15	135
7	10	137
8	14	142
9	11	135
10	16	133
11	19	137
12	13	141
13	9	139
14	10	140
15	14	145
16	17	139
17	11	142
18	17	147
19	10	143

C

Day	p1	p2	p3	p4	p5	p6	p7
1	31	25	27	15	12	21	
2	10	26	31	25	19		
3	9	45	32	18	23	27	25
4	13	16	35				

D

Day	average wait
1	21.8
2	22.2
3	35.6
4	21.3

E

Figure 5 . Examples of the data that might be graphed using each type of Shewhart chart.

COMMON CAUSE AND SPECIAL CAUSE VARIATION

So far, we've looked at the anatomy of Shewhart charts (center line, lower and upper control limits) and at the type of measures we might graph (classification, count, and continuous). You've learned that the location of the lower and upper control limits depends on the data type of the measure to be graphed, and that the way to inform the graphing software of the type of data is by specifying the type of chart. But what is the use of the control limits?

In the 1920's, Walter Shewhart developed the theory of common cause and special cause variation. The theory holds that there are two origins of variation in a process. Common-cause variation is the natural or expected variation, or random noise. Common cause variation results in a random variation of a measure

around its average. Special cause variation is variation attributed to a specific cause. Any system with special-cause variation is a learning opportunity. If the variation is in an undesirable direction, we will try to determine the cause and change things to make it more difficult for the problem to reoccur. If the variation is in a desirable direction, we'll again try to determine the cause, but we'll try to make it a permanent part of our procedures.

A system which has only common cause variation is called a stable system, or a system that is in "statistical control". "Stable" is not necessarily the same as good. For example, if a patent's systolic blood pressure predictably averages 165, varying between 160 and 170, it is stable, but not good. The performance of a stable system is predictable and therefore provides a basis for planning. The other important thing to understand about a stable system is that its performance will not change without a redesign of the process. If the level of performance is not acceptable, effort must be expended to change it.

Nor is special cause variation necessarily bad. An unusual value or unexpected trend may represent a good result that you'd like to understand and sustain. In fact, the goal of QI is to identify and implement changes which result in special cause variation in the desired direction.

Humans are very good at discerning patterns in data – unfortunately, we are very good at discerning patterns even in randomly fluctuating data. The usefulness of Shewhart charts lies in their ability to distinguish between common and special cause variation, using a small set of rules which are based on statistical theory combined with experience.

Actually, because the rules arise from a combination of statistics and heuristics, there are several systems of rules. The five rules specified by the Institute for Healthcare Improvement (IHI) for healthcare data are illustrated in Figure 6 the rule set used in other fields, such as manufacturing, are similar. These rules have been shown in practice to do a good job of distinguishing between special and common causes of variation, and thus to minimize the total cost of overreaction and underreaction to variation.

The IHI rules identifying special cause variation are: (A) A point outside the control limits (either above the upper control limit or below the lower control limit); (B) At least 6 consecutive points all increasing or all decreasing; (C) Eight or more consecutive points all on the same side of the center line (all above the line or all below the line); (D) Two out of three points in the outer third of the chart (the two points can either be on the same side of the center line, or they can be on opposite sides of the center line; the points can be consecutive, or there may be an intervening point); (E) Hugging (at least 15 consecutive points in the inner third of the chart.

Some of the rules will seem obvious, others may be puzzling. The "hugging" rule (at least 15 consecutive points in the inner third of the chart) and the "two out of three" rule (two out of three consecutive points in the outer third of the chart) represent an unexpected decrease or increase, respectively, in the amount of variation in the data rather than a change to the long-term average behavior.

The rules are not perfect. If you generate enough random numbers, each of these patterns will eventually show up by chance. However, the rules have been shown in practice to do a good job of distinguishing between special and common causes of variation, and thus to minimize the total cost of overreaction (treating variation as special cause when it is actually common cause) and underreaction (treating variation as common cause when it is due to a special cause) to variation.

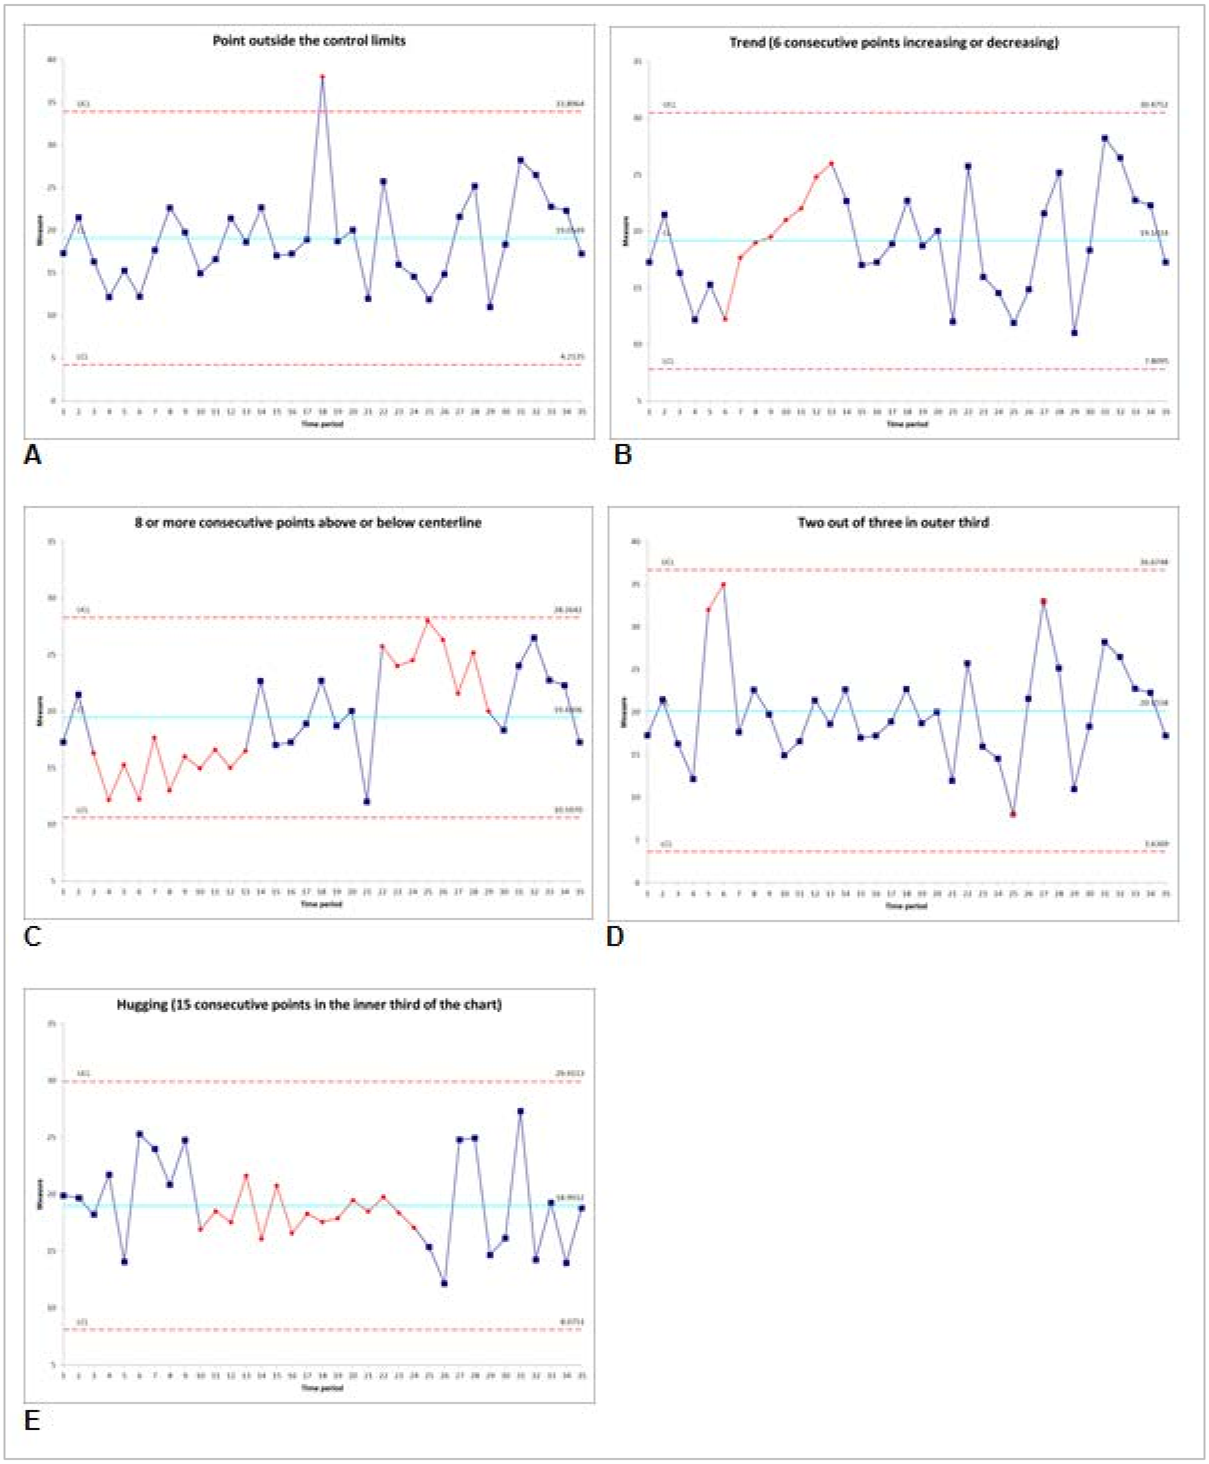

Figure 6: The 5 patterns indicating special cause variation. (A) a point outside the control limits; (B) at least 6 consecutive points all increasing or decreasing (trend); (C) 8 consecutive points above or below the center line; (D) two out of three points in the outer third of the chart; (E) 15 consecutive points in the inner third of the chart (hugging)

THE LIFECYCLE OF A QI PROJECT

Quality Improvement projects have a lifecycle, and Shewhart charts are used throughout most of the project's life. Once the goals of the project have been specified, and the measures that will be used to assess the project decided on, baseline data for all of the measures is collected and graphed. The ideal starting place is a stable baseline for each measure, since it is hard to identify progress in a measure which is already unpredictable. Thus, the initial graphs are checked for stability and before work on improvement starts, measures with unstable baseline patterns should be stabilized.

Once a stable baseline has been obtained, the center line, and lower and upper control limits of the graph of the measure will be frozen prior to implementing the first improvement idea. Data collected after the change is made will be added to the graph, but won't affect the frozen baseline limits. This allows the new post-change data to be compared to the baseline. The hope is that the newly added data will show special cause variation (in the desired direction).

Once it has been established that the change resulted in an improvement, a new (and improved) baseline is obtained and frozen. Subsequent improvements will be compared to the new baseline, not to the original (and no longer applicable) baseline. Now, the cycle starts again, with a new change idea, continued data collection and graphing, and the hope for another successful display of special cause variation. Each successful change, as evidenced by special cause variation, restarts the cycle of establishing a baseline, freezing the center line and control limits, introducing the next change and assessing the results against the most recent frozen baseline.

Once the project has been completed, data monitoring persists, in order to tell if the improvements are sustained even after attention turns to new projects.

USING SAS® TO CREATE SHEWHART CHARTS

Shewhart charts are produced using the SHEWHART procedure, which is part of the SAS/QC module. The basic syntax is requires you to specify the type of chart, the measure to be graphed, and the name of the variable which specifies the time period or "subgroup":

```
proc shewhart data = dataset_name;
   chart_type measure * subgroup;
run;
```

It's impossible to show the full range of possibilities without reproducing the User's Guide. Below are a few examples, to get you started.

EXAMPLE 1: C CHART

For this example, monthly data has been collected on the number of visits to the Emergency Department by residents of a long-term care facility. Because the facility consistently operates are capacity, we are willing to assume that the denominator (number of residents each month) is relatively constant. Because a given resident can have more than one ED visit, the number of visits is an example of a count variable.

The "markers" option requests the data points on the graph. Without it, only the lines would be shown.

```
proc shewhart data = ED;
   cchart ED_visits * month / markers;
  title "C-chart example: number of ED visits, constant denominator";
run;
```

The C chart is shown in Figure 7.

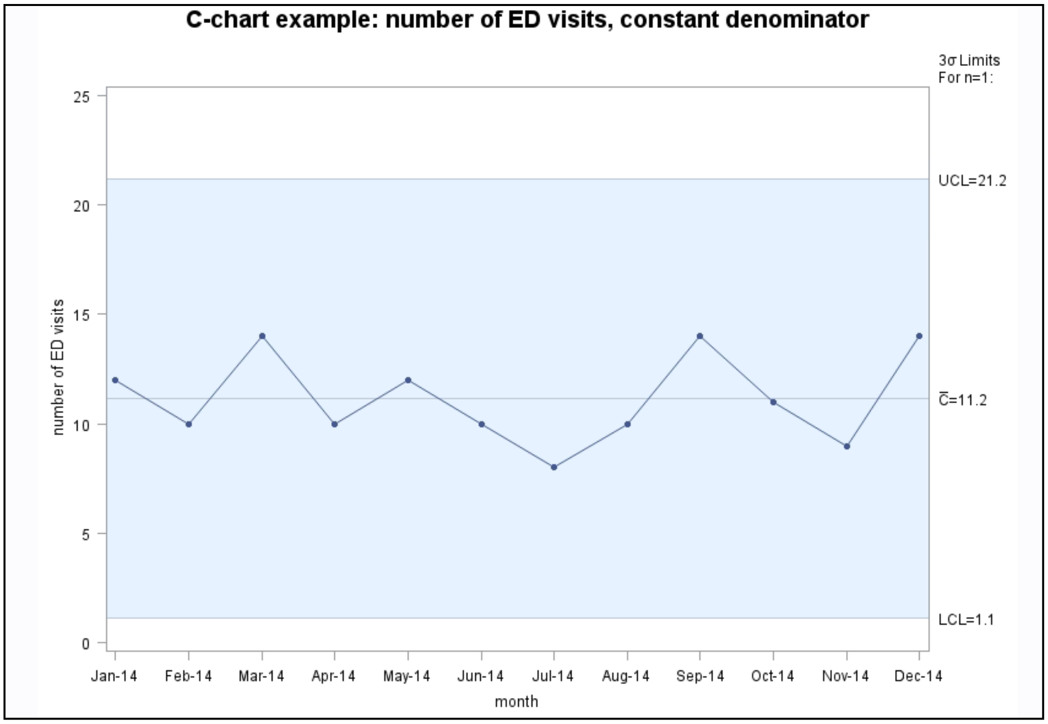

Figure 7. C chart created in Example 1.

EXAMPLE 2; P CHART

This example plots the proportion of patients who were readmitted within 30 days after discharge. Because a discharged patient is either readmitted or not, readmission is a classification variable, and a P chart will be used to plot the data. In this example, we don't want to assume that the number of patients discharged is the same each month. Instead, we'd like to provide the SHEWHART procedure with both the proportion of patients readmitted and also the denominator (number of discharges). The HISTORY option indicates that both variables will be provided. The names of the two variables must share a common root, with the name of the variable holding the proportion ending in 'u' and the variable holding the denominator ending in 'n'. The RENAME statement changes the names of the original variables, num_discharges and proportion_readmit, to conform to the required naming convention.

```
symbol v = dot height = 0.8;
proc shewhart history = readmissions
  (rename = (num_discharges = admitn proportion_readmit = admitp));
  pchart admit * month;
  title "P-chart: variable sized denominator";
run;
```

The resulting chart is shown in Figure 8..

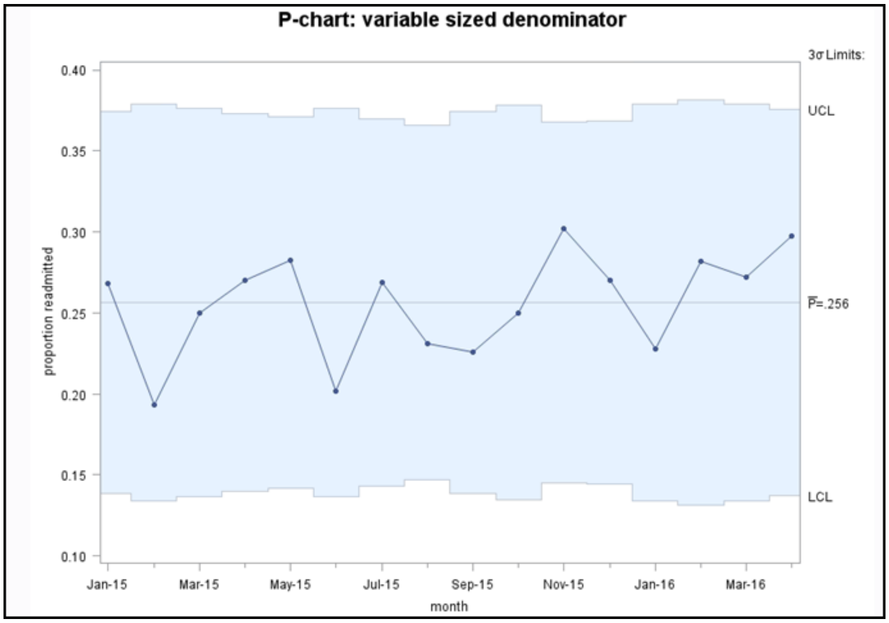

Figure 8. P chart created in example 2.

EXAMPLE 3: I CHART AND TESTING FOR SPECIAL CAUSE VARIATION

SAS provides 8 different tests for special cause variation. Desired tests are requested by number: the tests and the corresponding test number are shown in Table 2. You'll notice that while rules 1, 3, 5, and 7 correspond to IHI rules for special cause variation, there is no equivalent of the rule regarding 8 points all on one side of the center line.

In this example, average wait time has been calculated for each month. Because wait time is a continuous variable and because there is only one value for each time period, an I chart is called for.

In addition to requesting certain tests of special cause variation, the SAS code below requests that the area between the lower and upper control limits not be shaded (CINFILL = NONE), that the one- and two-sigma lines be shown (ZONES), that any instances of special cause variation be depicted using a red star, and that the rule underlying each instance of special cause variation be labeled on the graph. The 'r' in the name of the chart requests an 'Individuals-moving-range' graph. The moving range is the absolute value of the difference between the current month and the previous month. Small values on the moving range graph indicate periods of relatively low variation in the measure.

```
proc shewhart data = wait_times;
    irchart wait_time * month / cinfill = none   zones
                                tests = 1 3 5 7 8
                                testsymbol = star ctestsymbol = red
                                testlabel1 = 'beyond cntrl limits'
                                testlabel3 = 'trend'
                                testlabel5 = '2 out of 3'
                                testlabel7 = 'hugging'
                                testlabel8 = '8 one side';
  title "Wait time: Continuous variable with special cause tests";
run;
```

The chart produced by this code is found in Figure 9.

Test number	
1	One point outside the control limits
2	Nine points in a row in the inner third of the chart, all on one side of the center line
3	Six points steadily increasing or decreasing (a trend)
4	14 points alternating up and down
5	two out of three points in the outer third of the chart or beyond
6	Four out of five points in the outer two-thirds of the chart or beyond
7	15 points in a row in the inner third of the chart (hugging)
8	Eight points in a row on either side of the center line with no points in the inner thid.

Table 2. Special cause variation rules provided by the SHEWHART procedure

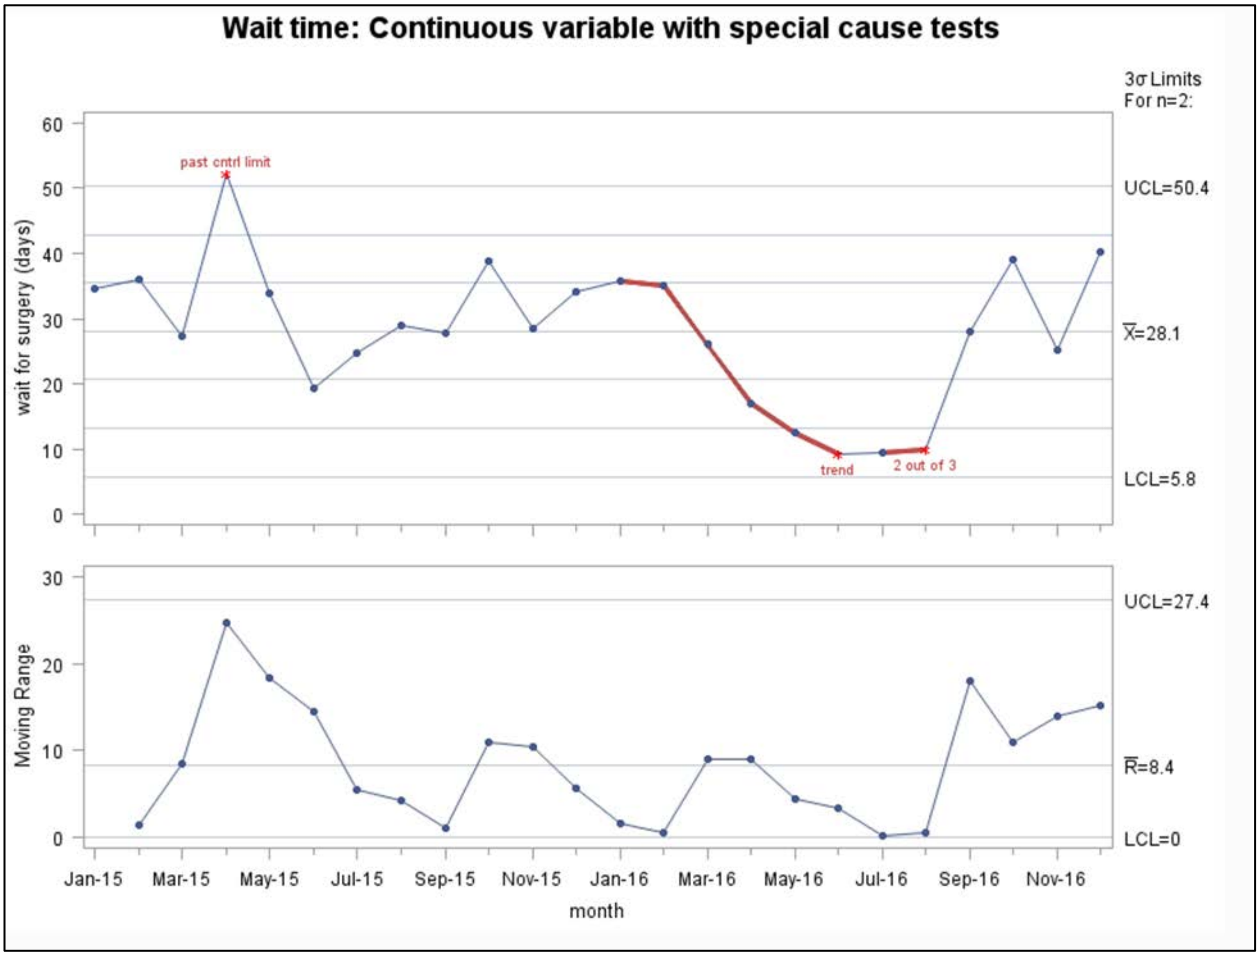

Figure 9. I chart created in example 3

EXAMPLE 4: COLLECTING BASELINE DATA, FREEZING THE BASELINE, AND ADDING NEW DATA AFTER IMPLEMENTING A CHANGE

This example looks at my efforts to decrease my commute time. The first step was to collect and graph baseline data, saving the results to a file using the OUTLIMITS option. I calculated the average commute time each week. Because I had only a single measure for each time period, an I chart was used (the NOCHART2 option suppresses the moving range chart). The graph in Figure 10 shows that I had a stable process, with an average commute time of just over 60 minutes.

```
proc shewhart data = baseline;
  irchart commute_time * week /
       cinfill = none nochart2
       outlimits = commute_limits;
  title "baseline data";
run;
```

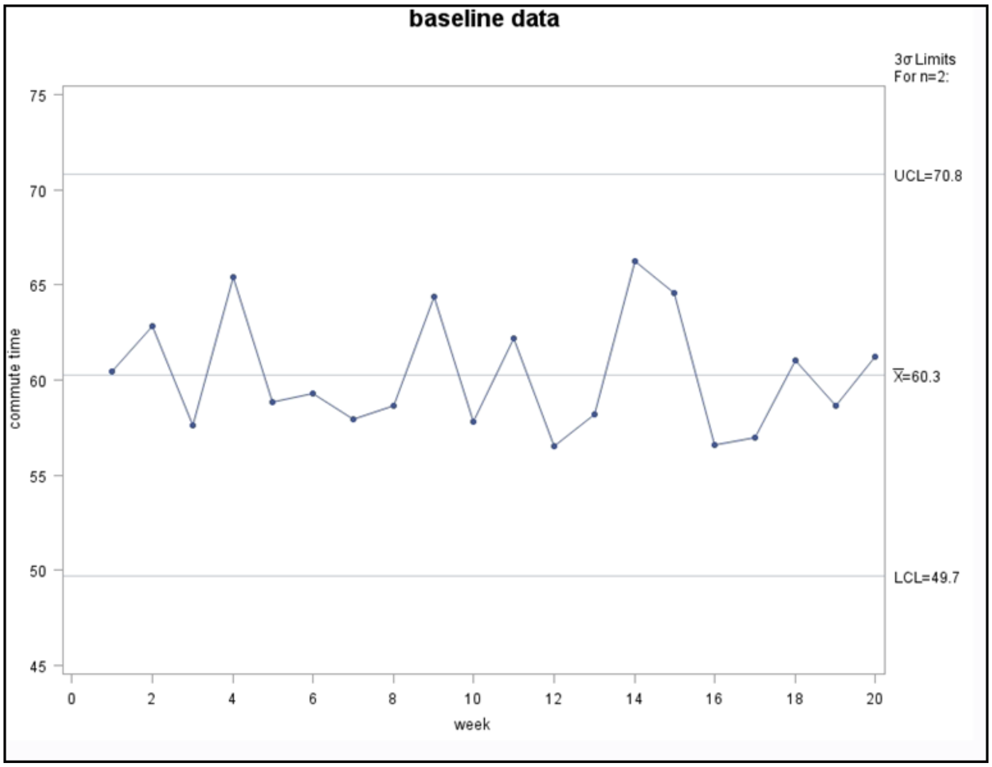

Figure 10. Graph of baseline data for example 4.

After making a change to my commute, I continued to collect data on a weekly basis. The new data was plotted, using the frozen baseline limits from the file created above.

In the resulting graph, shown in Figure 11, you can see that the new data represents special cause variation when compared to the frozen baseline control limits. I conclude that I have successfully reduced my commute time – at least on paper!

```
proc shewhart data = all limits = commute_limits;
   irchart commute_time * week /
           cinfill = none nochart2
           tests = 1 3 5 7 8
           testsymbol = star
           ctestsymbol = red;
    title "baseline limits frozen, extended";
run;
```

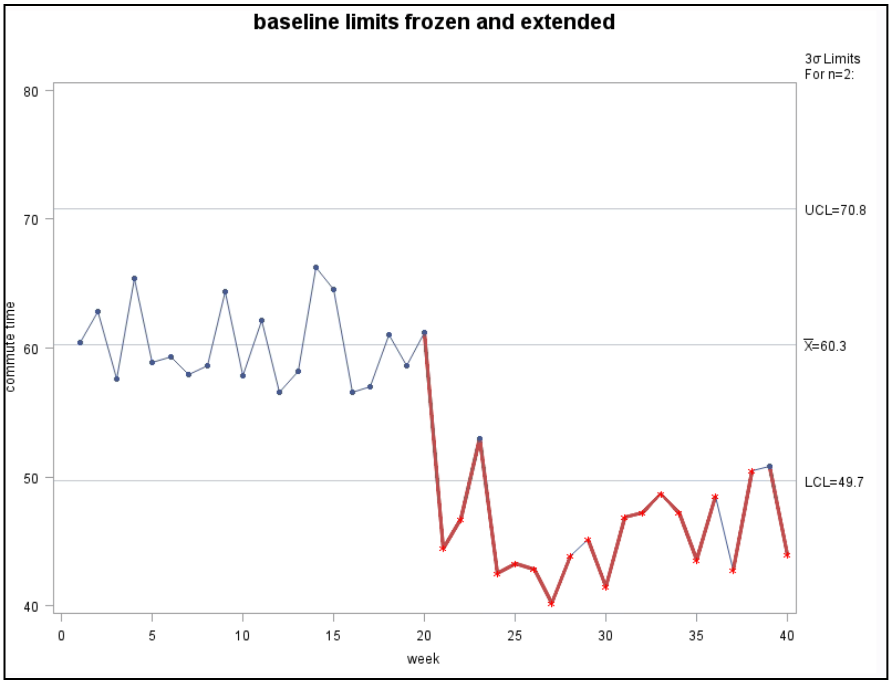

Figure 11. Graph showing new data points compared with the frozen baseline control limits.

CONCLUSION

Shewhart charts are a tool for learning about the behaviour of a measure over time. They allow us to differentiate between common cause and special cause variation, striking a balance between reacting too early to random fluctuations and reacting too late to real shifts in performance. In the context of Quality Improvement work, they allow us to tell if the changes we're testing produce variation in the desired direction – variation which cannot be explained by chance fluctuations in the data.

Shewhart charts make use of both the average level of performance and the amount of natural variation expected in the data. Because the amount of variation expected to arise naturally depends on the type of measure being graphed, we need to select the appropriate type of chart for the type of data being graphed.

Statistical theory underlies the rules used to differentiate common cause variation from special cause variation. Special cause variation is always a learning opportunity. In the context of a quality improvement project, the goal is to identify changes to a process which produce special cause variation in the desired direction.

RECOMMENDED READING

- *A great source of information on run charts and Shewhart charts in the context of quality improvement in healthcare is Provost LP, Murray SK. 2011. The Health Care Data Guide: Learning from Data for Improvement. San Francisco, C: Jossey-Bass.*
- *For SAS documentation, see: SAS Institute Inc. 2011. SAS/QC® 9.3 User's Guide. Cary, NC: SAS Institute Inc. The SHEWHART procedure is covered in chapter 13.*
- *The Institute for Healthcare Improvement provides a number of great videos on Quality Improvement. A seven-minute video on run charts can be accessed here: http://www.ihi.org/education/IHIOpenSchool/resources/Pages/AudioandVideo/Whiteboard7.aspx*

Two short videos (5 minutes and 9 minutes long) on Control Charts can be found here: http://www.ihi.org/education/IHIOpenSchool/resources/Pages/AudioandVideo/Whiteboard13.aspx

- *An easy to read article on run charts is Perla RJ, Provost LP, Murray SK. The run chart: a simple analytical tool for learning from variation in healthcare processes. BMJ Qual Saf. 2011; 20: 46-51. The article is available at http://www.med.unc.edu/cce/files/education-training/The%20run%20chart%20a%20simple%20analytical%20tool.pdf*
- *If you are interested in graphing rare events, two useful references are*
  - *Barbara A. Clear. "Charting Rare Events Effectively" Accessed 26 March, 2017. http://www.qualitydigest.com/inside/six-sigma-article/charting-rare-events-effectively.html*
  - James Benneyan. "Measuring Rare Events and Time-Between Measures" Accessed 26 March, 2017. *www.ihi.org/resources/Pages/Presentations/MeasuringRareEventsandTimeBetweenMeasures.aspx*
- Another video from the Institute for Healthcare Improvement ties a lot of ideas together: Frank Federico, "How Can Data Drive Reliability". Accessed 26 March, 2017. *http://www.ihi.org/education/IHIOpenSchool/resources/Pages/Activities/FrankReliability5.aspx*

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Ruth Croxford
Institute for Clinical Evaluative Sciences, Toronto, Ontario, Canada
ruth.croxford@ices.on.ca
www.ices.on.ca