

Which Chart to Use?

Individual Values or Averages?

Donald J. Wheeler

Process behavior charts are the interface between your data and your brain. But you have to begin by making a choice about which type of chart to use. You can either plot the individual values themselves, or you can organize your data into rational subgroups and plot the subgroup averages. This paper will discuss the issues involved and provide guidelines for when to use each chart.

THE ROLE OF PROCESS BEHAVIOR CHARTS

Your data almost always possess some sort of time-order sequence. In most cases this order will be linked to the operation of some underlying process. Yet most statistical techniques ignore this time-order sequence. Process behavior charts use this temporal order to characterize the behavior of these underlying processes. In this regard they are fundamentally different from virtually all other statistical procedures. Rather than trying to fit some type of mathematical model to the data, they use the time order in the data to characterize the underlying process as being either predictable or unpredictable.

When the process is being operated predictably we can extrapolate from the product measured to the product not measured, and even to the product not yet produced with some reasonable degree of confidence that the extrapolation will prove to be correct. Since these extrapolations are essential for continuing operations, this question of predictability is of primary importance.

When the process is being operated unpredictably we can be confident that something or someone is making changes to the process. When these changes are unplanned and unintentional, their detection will be the first step in removing their effects from the process. As we find the "assignable causes" behind these unplanned process changes we can make them part of the set of controlled inputs for the process. And as a result the process will display less variation and will also behave more predictably.

Thus, the objective of a process behavior chart is characterization and discovery. It tells us when we need to look for assignable causes of exceptional variation and when we do not need to

do so. And the choice of which chart to use is only a means to this end, not the end itself.

CHARTS FOR INDIVIDUAL VALUES

The simplest type of process behavior chart is the chart for individual values and moving ranges (the XmR chart). For this chart we plot the original data in their time-order sequence and use the differences between successive values as our two-point moving ranges. We use the average value as the central line for the individual values, while the average moving range is the central line for the moving ranges. These two averages are then used to compute the limits for both portions of the chart according to standard formulas.

As long as the successive values are logically similar (i.e. we are comparing apples to apples), and as long as the time period between successive values is long enough to capture the routine, short-term variation for the underlying process, the XmR chart will display the process behavior for all to see. The limits on the XmR chart filter out virtually all of the routine process variation so that we may find any evidence of exceptional variation that may be present in the data.

The strength of the XmR chart is that every time we get a value we plot a point, and each of these points has to sink or swim on its own. This makes the connection between the XmR chart and the context for the data extremely simple and direct. And this direct connection facilitates the use and interpretation of the XmR chart.

AVERAGE AND RANGE CHARTS

The average and range chart divides the original time series into subgroups containing two or more values each and then computes the subgroup averages and the subgroup ranges. According to Dr. Walter Shewhart, when we place two values together in the same subgroup we are making a "human judgment" that these two values were collected under "essentially the same conditions." The name given to this judgment process is rational subgrouping, and its objective is to end up with subgroups that contain the routine process variation and nothing but the routine process variation.

In practice we often create subgroups by collecting a few values at one specific time and placing these in one subgroup. Then, at a later time, another set of values is collected, etc. As long as these subgroups reasonably represent the routine process variation the subgrouping may be said to be rational.

So, when using an average and range chart we need to choose both a subgroup size and a subgroup frequency. The ability to choose the subgroup frequency will allow us to obtain subgroups at a rate that is consistent with the time period over which the process might change. The ability to independently select the subgroup size will allow us to take advantage of the increased sensitivity that comes with larger subgroup sizes. However, there is a diminishing

return here: The costs will increase with the subgroup size while the sensitivity will only increase with the square root of the subgroup size.

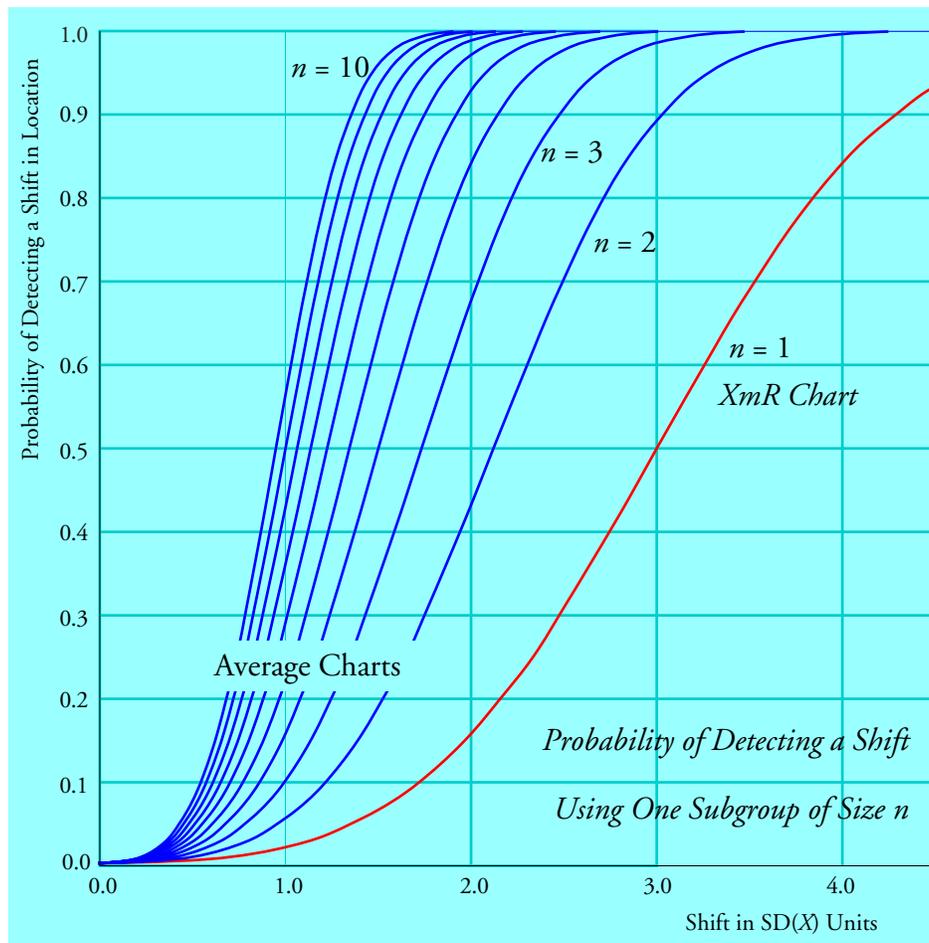


Figure 1: The Effect of Subgroup Size Upon Power for Average Charts

This diminishing return may be seen in the traditional power function curves. Power functions plot the probability of detecting a shift on the vertical scale against the size of the shift in location on the horizontal scale. In Figure 1 we can see how the increase in subgroup size results in charts with greater sensitivity to process shifts, yet these increases become progressively smaller with increasing subgroup size, n .

WHICH CHART TO CHOOSE

When deciding on which chart to use you have a choice between using the original data on a chart for individual values or arranging the data into subgroups and using an average and range chart. In the light of Figure 1 it would seem that an XmR chart suffers an extreme disadvantage in sensitivity when compared with any average chart. But the comparison made in Figure 1 is

not the whole story. In making our choice between the two types of charts it is helpful to make a distinction between “periodically collected data” and what I shall call, for lack of a better term, “widget type data.”

Widget type data are data where the subgroup size is independent of the subgroup frequency. These data tend to occur when people can collect any number of items to measure at any point in time. Here we will have to choose both the subgroup size and the subgroup frequency. For such data, as seen in Figure 1, an average and range chart will generally be the best chart to use.

Periodically Collected Data:	Widget Type Data:
The frequency for individual values is fixed by time or event, so that subgroup size determines subgroup frequency.	We may choose the subgroup size independently of the subgroup frequency.
Organize data so that successive differences capture routine process variation:	Organize data into Rational Subgroups that are internally homogeneous:
Use XmR Chart	Use Average and Range Chart

Figure 2: Which Chart to Use?

Periodically collected data are data for which the subgroup frequency depends upon the choice of subgroup size. Here each value is uniquely identified with a specific period of time, a specific batch, or a specific event. As a result of the inherent frequency attached to each value, any change in subgroup size will also require a change in the time period covered by the subgroup. Since the process might change between one value and the next, periodically collected data will have a logical subgroup size of $n = 1$, and as we will see below, the XmR chart will provide the most sensitive analysis for these data.

So we have the general guidelines shown in Figure 2. While issues of rational sampling and rational subgrouping may occasionally result in different choices than those shown there, these guidelines are sufficient as a starting point.

SENSITIVITY OF THE CHARTS FOR PERIODICALLY COLLECTED DATA

If we are working with widget type data, increasing the subgroup size will, up to a point, result in an increase in the sensitivity of the average chart. This point is reached when we begin to subgroup unlike things together and our subgroups become non-homogeneous. As the subgroup size increases beyond this point the average chart will have *less* sensitivity rather than

more. The homogeneity within the subgroups is much more important to the sensitivity of the charts than is the size of the subgroups. This is why small subgroup sizes are always preferred with an average and range chart.

This also explains why periodically collected data tend to work best on an XmR chart. When each value represents a different point in time, placing several values together in a single subgroup is an implicit assumption that those time periods are all alike. Keeping those values separate on an XmR chart allows us to see if those time periods differed.

“So which type of chart is the more sensitive for periodically collected data?” With periodically collected data the subgroup size is not independent of the subgroup frequency. Thus, we cannot simply compare different subgroup sizes without also taking into account the different amounts of data they require. This means that Figure 1 no longer tells the right story. To make realistic comparisons we have to use a fixed amount of data.

Say that we obtain one value per time period and that the process average shifted by some amount at some time in the past as shown in Figure 3. We shall look at the sensitivities of an XmR chart and various average charts at three points in time following the shift. We shall compare the power function curves when we have one point available following the shift, when we have two points available following the shift, and when we have four points available post shift.

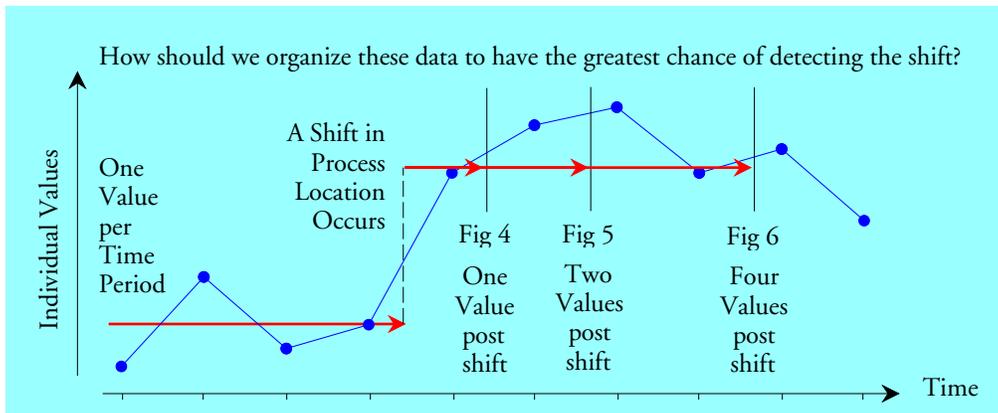


Figure 3: The Three Points in Time for the Following Comparisons

With only one point available following the shift we can plot this one point on an XmR chart, or we might assume that this one value is the last value needed to complete a subgroup on an average chart. Under this assumption Figure 4 shows the power functions for the XmR chart and for nine average charts with subgroup sizes ranging from $n = 2$ to $n = 10$. With only one point following the shift on each chart we can only use detection rule one—a point beyond one of the three-sigma limits.

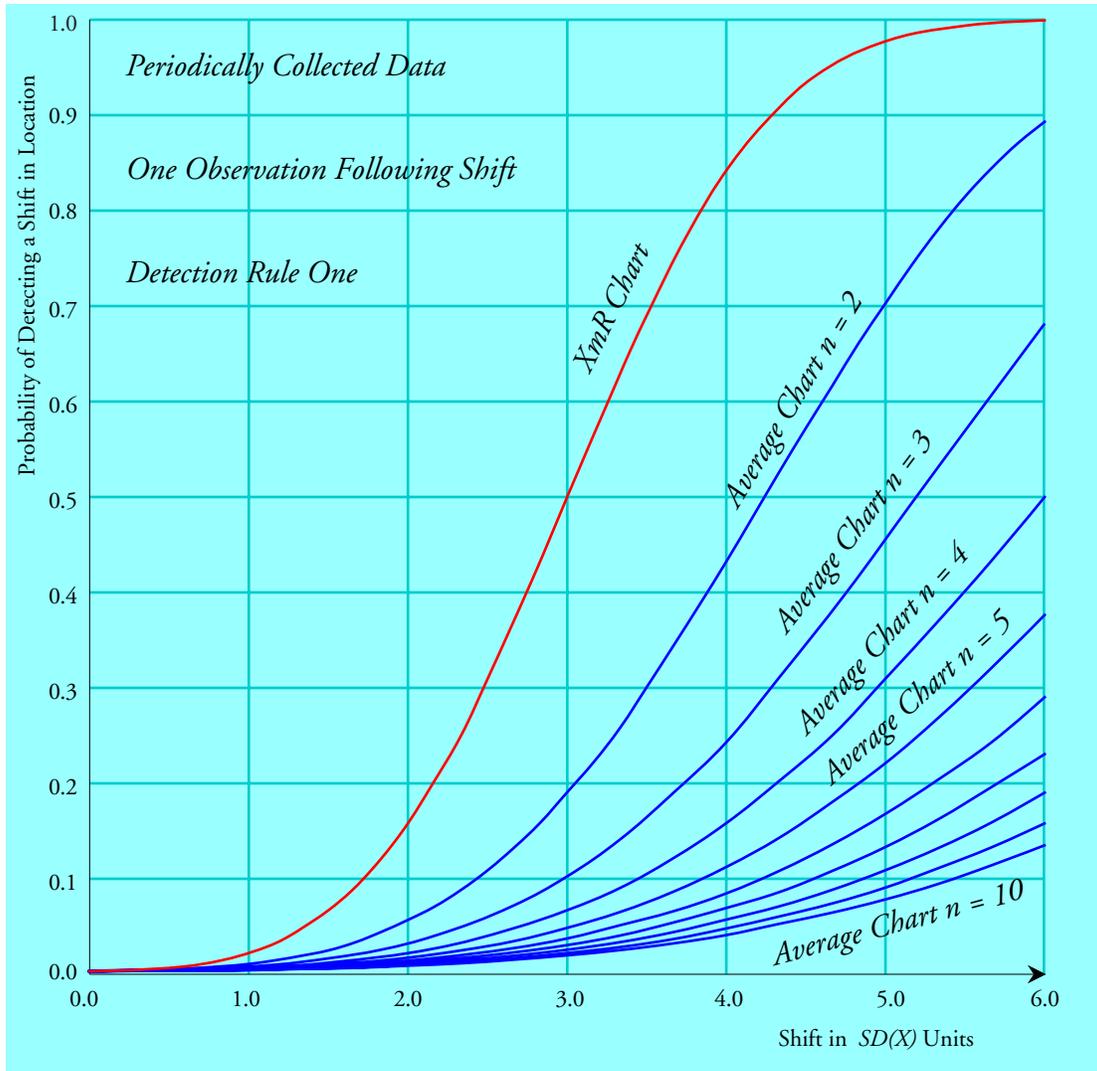


Figure 4: Power Functions for Periodically Collected Data for One Point Post Shift

The hierarchy in Figure 4 is exactly the opposite of the hierarchy shown in Figure 1. This happens because the average charts dilute the effect of the shift by combining data from before the shift with the one value following the shift. Because of this dilution their power functions are all dominated by the power function for the XmR chart, and the sensitivity drops with increasing subgroup size. *Regardless of the size of the shift, with periodically collected data, the XmR chart will be more likely to detect that shift on the first value following the shift than will any possible average chart.*

Because the XmR chart considers each value on its own, it will always have the greatest initial sensitivity of any charting technique when used with periodically collected data. When your data occur one-at-a-time you will want to interpret them as they occur, and the XmR chart lets you do this.

Next we consider what happens when we have two values to use following the shift. While

an average chart with $n = 2$ will no longer suffer the effects of a dilution of the effect of the shift, average charts with larger subgroup sizes will still be combining pre-shift data with post-shift data. At the same time, with two points on the XmR chart, we can consider using both detection rule one and detection rule two—a run of two out of three points beyond two-sigma. Figure 5 shows the relevant power functions.

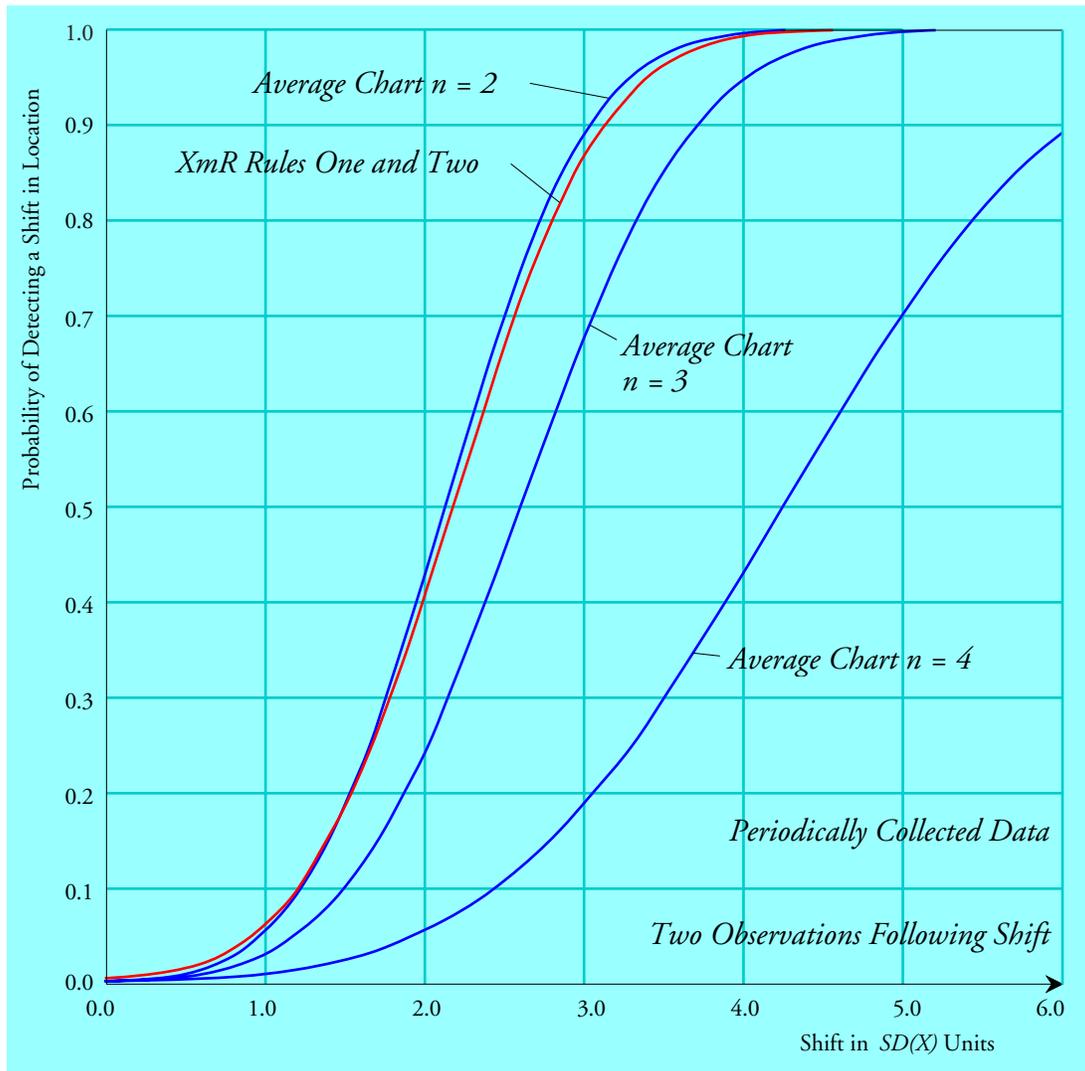


Figure 5: Power Functions for Periodically Collected Data for Two Points Post Shift

With two observations following a shift, an average chart with $n = 2$ effectively becomes as sensitive as an XmR chart that uses detection rules one and two. While the power function for the average chart with $n = 2$ is very slightly higher than the power function for the XmR chart, these differences in theoretical power are too small to make any difference in practice.

And when you only have two points following the shift, all average charts with subgroup

sizes greater than two will be less sensitive than the XmR chart.

What happens when we have four values to use following the shift? Average charts with $n = 2, 3,$ or 4 will no longer suffer the effects of a dilution of the effect of the shift. Their subgroups will once again be homogeneous. However, average charts with subgroup sizes larger than 4 will still be combining pre-shift data with post-shift data which will dilute the effects of the shift and deflate the power functions. At the same time, with four points on the XmR chart, we can consider using detection rules one, two, and detection rule three—a run of four out of five beyond one-sigma. Figure 6 shows the relevant power functions.

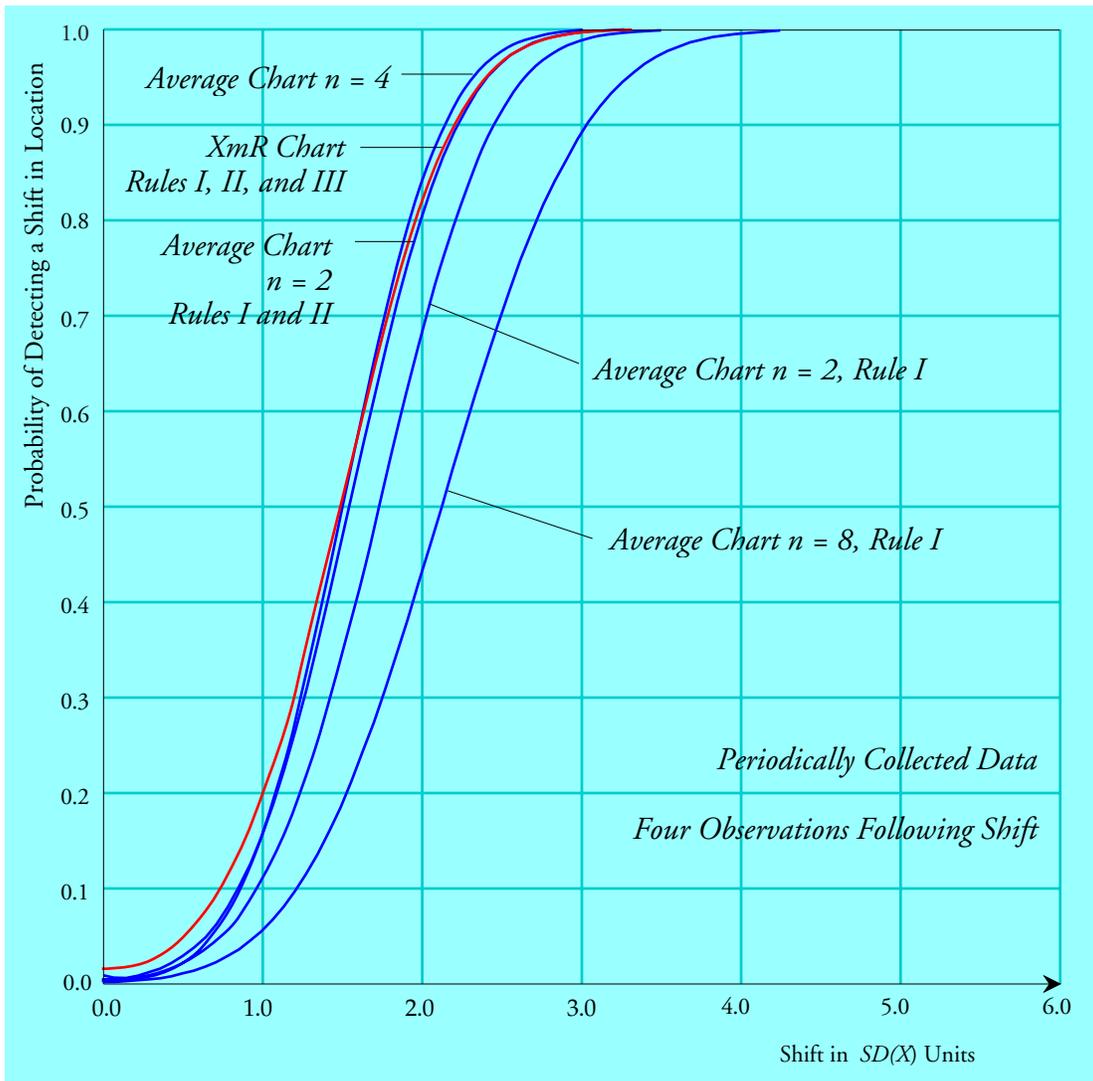


Figure 6: Power Functions for Periodically Collected Data for Four Points Post Shift

Once again, the best of the average charts is on a par with the XmR chart using the Western Electric zone tests.

The average chart with $n = 4$ will only have one subgroup following the shift. The average chart with $n = 2$ will have two subgroups following the shift, and so here we could use detection rules one and two. Both of these average charts have a power function that is indistinguishable from that of the XmR chart using detection rules one, two, and three. All three of the charts listed on the left in Figure 6 are essentially operating with maximum power for the amount of data available.

The average chart with $n = 2$ used with only detection rule one is slightly less sensitive to intermediate sized shifts than the three charts above. Finally, the average chart for $n = 8$ still lags behind the other charts because the mixture of the pre-shift and post-shift data dilutes the effect of the shift and results in a subgroup that is not homogeneous. Average and range charts do not work as well as they might when we bury signals within the subgroups.

These examples serve to make the point. Average charts simply cannot become fully sensitive to a shift until all of the values within a subgroup are obtained following that shift. Therefore, when you use an average chart with periodically collected data you are making a commitment to waiting at least n time periods before the chart will become fully responsive to process changes. While this can make sense when the process changes slowly, it risks missing some other types of process changes.

SUMMARY

When tracking periodically collected data, the XmR chart will have the greatest initial sensitivity of any technique.

When we use the XmR chart with the Western Electric zone tests to track periodically collected data it will have essentially the maximum possible sensitivity that can be achieved with any statistical technique.

These two mathematical facts provide the rigorous foundation for the guidelines in Figure 2.

The choice of a chart requires judgment. Rational sampling and rational subgrouping are essential to the effective use of process behavior charts. Rational subgrouping requires judgment based on an understanding of the way the data were obtained, what the data represent, and the various ways in which the context for each value might differ. Rational sampling requires judgment about the way the process might behave and the time scale for possible process changes.

The guidelines in Figure 2 suggest certain categorical choices, but in the end the best chart will depend upon the context for the data and the questions to be answered. This is why robotic approaches to the creation of process behavior charts are doomed. While computers are good at computing, *the software never knows the context for the data*. It never knows how the data were collected. It never knows how the process is likely to change over time. It doesn't even know

what the data represent. And it certainly does not know which questions are of interest to the user. About all that robotic subgrouping can do is to blindly try various subgroupings and see which one satisfies some algorithm.

In my career I have corresponded and consulted with many people who were perplexed by the results of their robotic subgroupings. The results they obtained simply did not make sense, and when this happened it was invariably due to the fact that they were using irrational subgroups or irrational sampling schemes. So let the computers create the charts, but do not surrender the choices regarding subgrouping and sampling to the computer.

When choosing an interface between your data and your brain it is essential to organize things so that they make sense. When you use an appropriate process behavior chart it will allow you to not only identify opportunities for process improvement but to also communicate those opportunities to others. This is why rational sampling and rational subgrouping are the heart of the process behavior chart technique. They are the essence of the way of thinking that is the foundation for using process behavior charts successfully.

So follow the guidelines, and if you are ever uncertain about how to subgroup your data, remember that it is hard to screw up the subgrouping on an *XmR* chart.

FOR FURTHER READING

This article draws upon a large body of practical and technical material. For more on rational subgrouping and rational sampling see my articles of the same names in *Quality Digest Daily*, June 1, 2015; and *Quality Digest Daily*, July 1, 2015.

For those who are concerned with how to chart counts or rates based on counts see my article "What about *p*-charts?" *Quality Digest Daily*, Oct. 3, 2011.

For more on the *XmR* chart see my articles: "What Makes the *XmR* Chart Work," *Quality Digest Daily*, Dec.3, 2012; "A History of the Chart for Individual Values" *Quality Digest*, July 10, 2017; "Individual Charts Done Right and Wrong," *Quality Digest Daily*, Feb. 2, 2010; and "But the Limits Are Too Wide!" *Quality Digest Daily*, Jan. 7 2013.

For those who are interested in learning more about power functions, both the theory and the formulas may be found in a white paper: "Power Functions for Process Behavior Charts" by Donald J. Wheeler and Rip Stauffer, at <spcpress.com/pdf/DJW321.pdf>. This white paper represents an extension of material originally published by Donald J. Wheeler in the *Journal of Quality Technology* v. 15, no. 4, October 1983, pp.155-169.