

Cogent Economics: Control Charts - A Primer

Once upon a time, there was a savings and loan company that specialized in residential mortgages. The company had 9 branches, each of which handled loan applications. The motto in this company was volume, volume, volume. The branches sought to process as many loans as possible, with little concern for the number of defective loans they approved. After some time, the management noticed that although a lot of loans were getting approved, repurchase and servicing issues were mounting. So they started paying attention to quality. If you're reading this, then you've probably reached this enlightened state already. You know that quality is important. The next step is determining where to make process improvements in order to improve quality. Control charts are a tool that we at Cogent have integrated into our software to help you accomplish that task.

Control Charts Defined

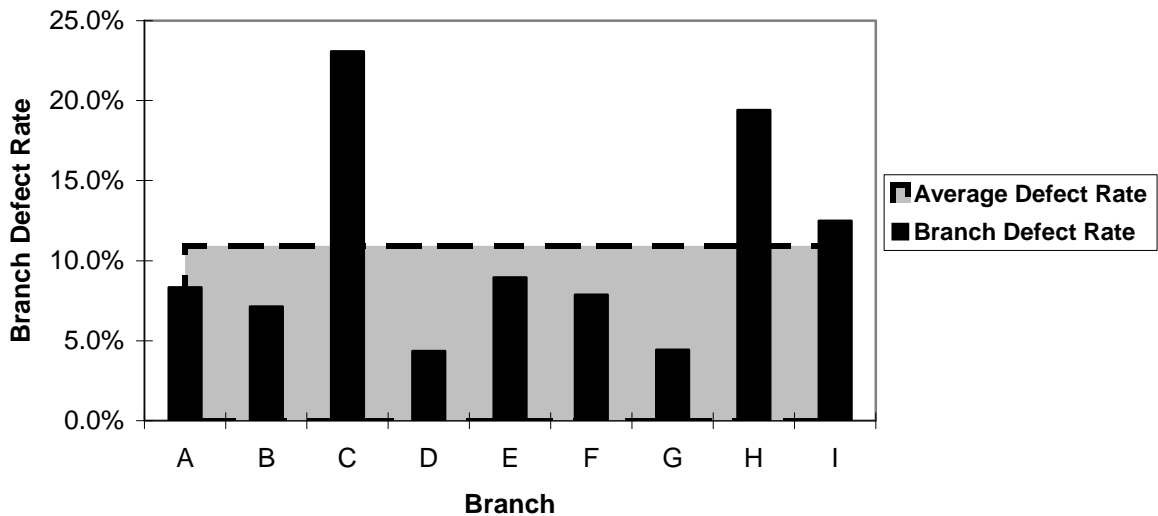
What is a control chart? A control chart is a graphical approach to quality control. It displays the defect rate for a given segment of the origination process against an upper limit on what that defect rate ought to be. If the defect rate is above the upper control limit (UCL), then we say that the process is out of statistical control, which means that something other than chance errors is probably causing the defects; that is, that they are caused by some structural flaw in the origination process. If the defect rate is below the UCL, then it's possible that chance errors account for the variance. This document has two goals. The first is to elucidate the importance and effectiveness of control charts as a statistically valid method for identifying process flaws. The second is to explain how the UCL is calculated so that the analysis behind this process is clear. Control charts are an extremely powerful method for analyzing data quickly, and their utility is enhanced by our software's ability to target samples for easy comparison.

It is also important to understand what a control chart does not do. A control chart does not tell you whether your company is doing better than the company next door. That comparison is not possible at this time, because there is no industry standard for what constitutes a defective loan. The upper control limit should also not be confused with your threshold for what constitutes acceptable quality. In other words, it is theoretically possible, though unlikely, that your branches will all have an 80% defect rate, but they will still be in statistical control. This is because the upper control limit is a value that reflects actual product counts, sample sizes, and defect rates, rather than your optimum defect rate (which is probably closer to zero). What the control chart *does* allow you to do is discover whether one particular branch or facet of the origination process is operating out of sync with the rest. In other words, you can highlight anomalous behaviors so that you can examine them further and determine if changes need to be made.

Basic Statistics: Mean and Standard Deviation

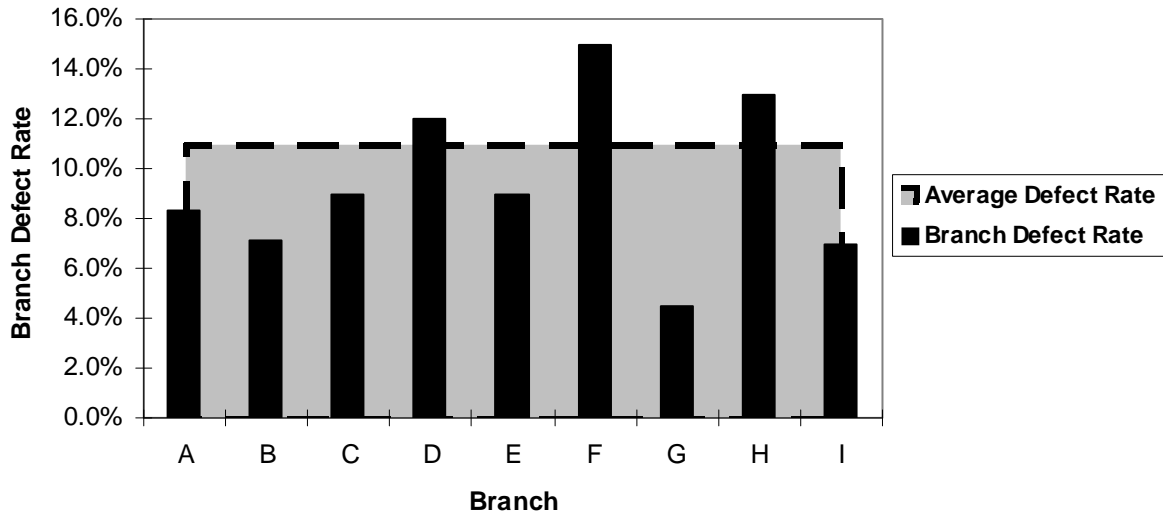
Let's return to our mythical company at the moment of their epiphany. The managers at our firm realized that the quality of loan applications deserved consideration, so they instituted a penalty for defective loans. The process was simple enough: every month, each branch's approximate defect rate was determined by reviewing a sample of their approved applications. The defect rates were then added together and divided by 9 (the number of branches) to produce the average, or **mean**. Every branch that had a defect rate below the average was rewarded, and every branch with a defect rate above the average was penalized. They devised a chart for the month of January, shown below.

Defect Rates Against the Average in January



Branches C, H, and I were all penalized for their poor performance in January. When February's chart was released, the average defect rate was about the same, but generally, the culprits were a different set of branches:

Defect Rates Against the Average in February



The problem with this process became readily apparent 8 or 9 months into it. None of the branches seemed to be exceptionally bad, but they all had good or bad months. The same branch that was rewarded one month was penalized the next month, only to be rewarded again the month after that. There was no noticeable trend in quality improvement, and nobody could tell whether the results obtained from the company's review had any statistical validity. As a matter of fact, they didn't, because the important question is not whether the defect rate for a branch is above or below average, but rather whether it is *so far* above the average that the defect rate cannot be explained away by chance. This required a separate calculation, involving the **standard deviation**.

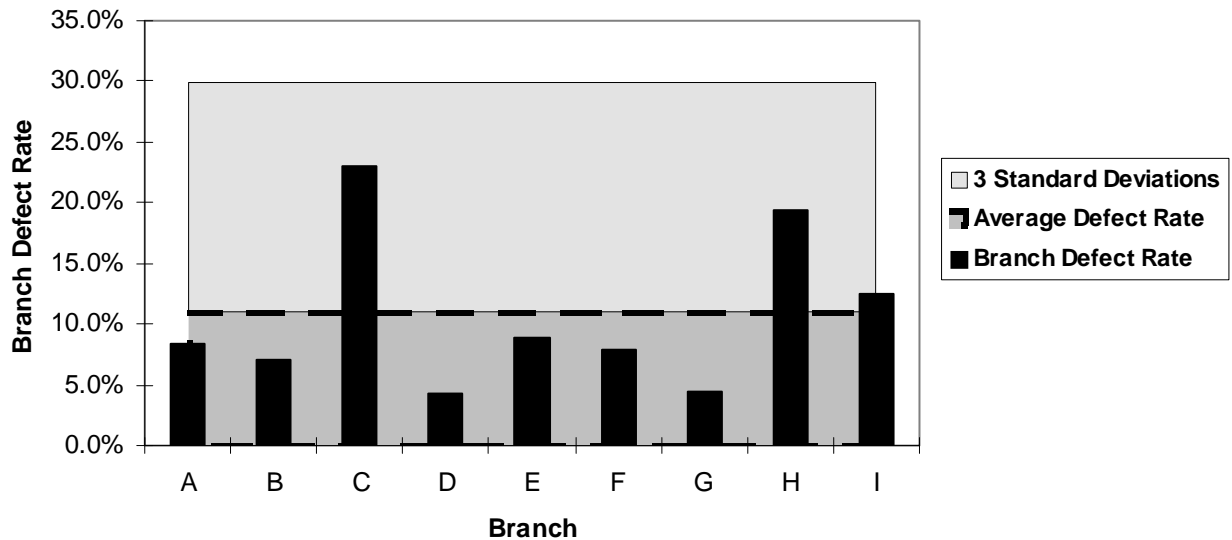
The standard deviation is generally a less familiar concept to those without a statistical background. The standard deviation is a calculation made by analyzing how far away from the mean a typical observation in a sample lies. In the case of our example, the standard deviation would measure approximately how far each branch's defect rate ought to be from the mean.¹

The rule of thumb is that if you go one standard deviation in either direction from the mean, you should encompass just a little more than two thirds of the data, or about 68 percent. If you go two standard deviations away, you enclose a little more than 95 percent, and within three standard deviations, you should have accounted for 98 percent of your data. Any data points that are more than 3 standard deviations from the mean are considered statistical anomalies that probably did not occur by chance. The managers at

¹ The standard deviation is also called the root mean squared error, or RMS error. To calculate it, subtract the mean from each branch's defect rate, square that number, and *then* take the mean of all the squared errors. *Then*, take the square root of that average, and that number is the standard deviation.

our firm bought a statistics textbook and looked at their January defect rate chart again, this time plotted against the benchmark of three standard deviations.

Defect Rates Against 3 Standard Deviations

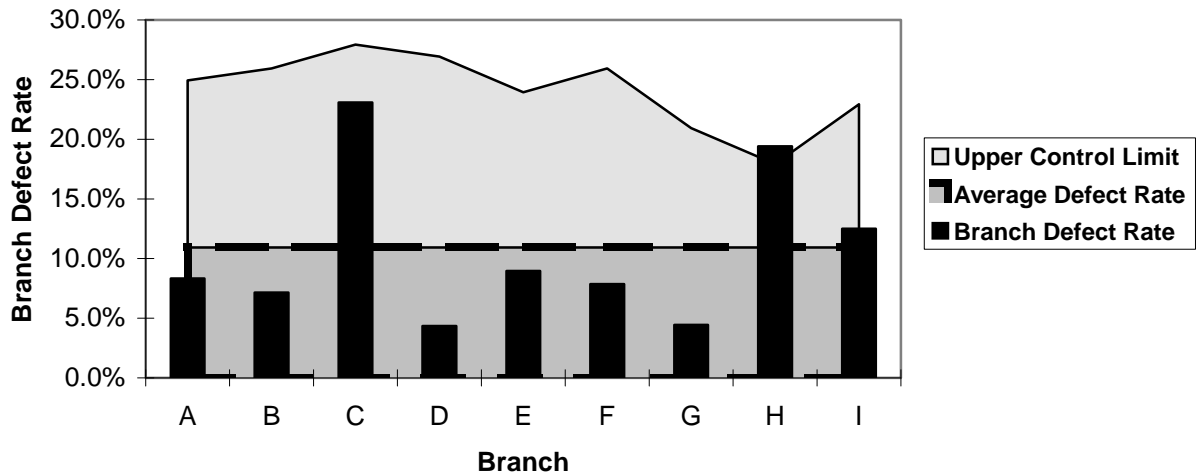


What a relief! None of the branches had defect rates above three standard deviations from the mean. But as time wore on, the bottom line got worse. Something needed to change. The company decided to invest its resources in correcting the problems at Branch C, because it had the highest absolute defect rate. Was this the right decision?

The Art of Control Charts: Calculating the True Upper Control Limit

You probably guessed that, in fact, the company was wrong. The reason why the company was wrong is because the measurement of three standard deviations isn't an appropriate standard to apply to all the branches. That number needs to be adjusted to account for differences between the branches (for example, each branch probably sampled a different number of loans). Finally, the managers at our company acquired the Cogent ProductionQC system, and used it to calculate a statistically appropriate upper control limit. The upper control limit is three standard deviations multiplied by an adjustment variable that is unique to each branch. The resulting chart displayed the branches' defect rates against the upper control limit, and it looked like this:

Statistical Control Chart: By Branch (3 Std. Dev.)



The chart indicates that Branch C, although it has the highest absolute defect rate, is in statistical control. The number of defects generated at Branch C can be anticipated because of the particularities of that branch (the size of the sample of loans they reviewed, for example). Branch H, on the other hand, although it has a lower absolute defect rate, is out of statistical control. This means that there are errors occurring at Branch H that most likely result from process flaws rather than sampling error. This chart tells the management exactly where they should target their resources and attention in order to have the most significant impact on the bottom line.

The Power of Control Charts in the Cogent System:

With the Cogent software, you can construct control charts that allow you to compare branches with each other, but also evaluate appraisers, underwriters, etc. The purpose of control charts, as articulated earlier, is to ensure that all cylinders are firing and no one component of your origination process contributes disproportionately to your overall defect rate. By learning to read control charts and understanding the logic behind them, you can quickly and easily sift through your entire set of practices and isolate statistically significant causes of error. By enabling you to establish specific targets for reform, we hope to help you efficiently improve your procedures to the point where they meet your satisfaction.